COIN-O-MATIC: A fast system for reliable coin classification

L.J.P. van der Maaten
P.J. Boon
MICC-IKAT, Universiteit Maastricht
P.O. Box 616, 6200 MD Maastricht, The Netherlands
telephone: (+31)43-3883901 fax: (+31)43-3884897
e-mail: {l.vandermaaten,p.boon}@micc.unimaas.nl

Abstract

Systems that perform fast and reliable classification of heterogeneous coin collections can be beneficial to charity organizations and financial institutions that collect unsorted coins. Existing coin classification systems cannot classify heterogeneous coin collections. We present a new coin classification system designed to perform reliable classification of heterogeneous coin collections. In this case, reliability means with a low number of incorrect classifications. COIN-O-MATIC uses a combination of coin photographs and sensor information in the coin classification. The system preprocesses the coin photographs, and classifies the coins using edge-based statistical features. The classification is verified using a mutual information measure of the coin image and an averaged coin image that corresponds to the classification. We measure the performance of the system on a test set supplied by the MUSCLE CIS benchmark. We show that our system classifies approximately 72% of the coins correctly, while misclassifying only 2% of the coins. Moreover, the presented system is computationally efficient.

1 Introduction

After the introduction of the Euro, charity organizations collected large numbers of pre-Euro coins in order to raise extra funds for their work. Current state-of-the-art coin sorting machines are not capable of sorting these coins. This is due to the large number of coin types and currencies that is present in the obtained coin collection. In general, coin sorting machines measure features such as area, thickness, and weight in order to classify coins, while ignoring visual features in the coins. Incorporation of visual features in the task of coin classification allows for classifying and sorting the heterogeneous coin collection of unsorted pre-Euro coins. Automatic sorting of heterogeneous coin collections is not only beneficial to charity organizations, but also for financial institutions that handle these collections, such as banks and change offices.

In addition, the development of a system for automatic classification of heterogeneous coin collections can be beneficial to the cultural heritage domain [13].
Because of the ignorance of visual features in coins, existing systems for automatic classification of coins can only be applied on coin collections with a limited number of coin classes [1, 3, 5, 8, 11]. Systems that classify heterogeneous coin collections by incorporating visual features are presented by Nölle *et al.* [9] and Huber *et al.* [6]. Nölle *et al.* [9] present a system that is based on measuring correlations between coin edge images under several rotations. Nölle *et al.* report a high sorting performance on a dataset of modern European coins. Huber *et al.* [6] present a feature-based system that classifies coins using eigenfeatures of a rotation-invariant representation of the coin images. Huber *et al.* present promising classification performances on a subset of the MUSCLE CIS benchmark dataset.
This paper presents COIN-O-MATIC: a new system for fast and reliable coin classification, which was developed to compete in the MUSCLE CIS benchmark competition. The system combines thickness sensor measurements with visual features in order to classify coins. In contrary to the system by Huber *et al.*, which uses eigenfeatures to classify coins, our system is based on edge-based statistical features. The focus in the presented work is on reliability and speed. Reliability is necessary, since misclassifications are very expensive in the MUSCLE CIS benchmark assessment scheme. Speed is an important issue in automatic coin classification, because systems performing classification of large coin collections should be able to process a large number of coins in a limited amount of time. Systems competing in the MUSCLE CIS benchmark should process 10,000 coin images within 8 hours.
The outline of the remainder of the paper is as follows. In section 2, we give a brief description of MUSCLE CIS benchmark specification. In section 3, we describe the workflow of COIN-O-MATIC. Section 4 describes our approach to the segmentation of the coin. Section 5 presents the edge-based statistical features that are used by the system. In section 6, we describe our approach to the classification of coins and to the verification of these classifications. In section 7, we present the results of the experiments with COIN-O-MATIC on a test set provided for the MUSCLE CIS benchmark. Conclusions and recommendations for further research are presented in section 8.

2 MUSCLE CIS benchmark

For the MUSCLE CIS benchmark competition, a new coin dataset was developed. This dataset contains approximately 100,000 coin images, corresponding to 50,000 coins. An example of a coin image from the dataset is shown in Figure 1(a). The dataset is divided into a fixed training set of 20,000 coins, and six fixed test sets of 5,000 coins. The training set contains 2,270 different coin faces, corresponding to 692 coin classes. In addition, the training set contains an averaged coin image for each coin face. An example of an averaged coin image is shown in Figure 1(b). In the test sets, approximately 400 of the coin classes appear. In addition, the test set contains 3-4% coins that are not in the training set, and that should be classified as unknown.
In order to assess the performance of the systems competing in the benchmark competition, an assessment scheme was developed. In the assessment scheme, 1 point is assigned for a correct

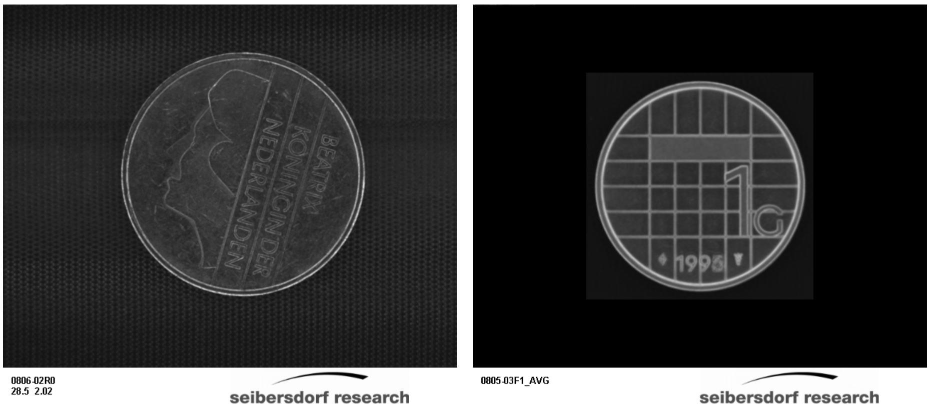

(a) Normal coin image. (b) Averaged coin image.

Figure 1: Examples of coin images.

coin classification, whereas 0 points are assigned to a classification as unknown (i.e. a rejection of the coin). An incorrect classification leads to a substraction of 100 points. In addition, 25 points are assigned for each coin class for which at least one coin was classified correctly. The aim of the systems in the benchmark competition is to maximize the number of assigned points. The high penalty for incorrect classifications in the assessment scheme indicates the importance of reliability of the classifications.

In addition to the assessment scheme, the MUSCLE CIS benchmark contains two exclusion criteria: (1) the program should classify at least 70% of the coins correctly, and (2) the program should process 5,000 coins within 8 hours on a 3 GHz computer with 1 GB of RAM. The latter criterium implies that a classification of a coin (i.e. the processing of two coin photographs) should be performed within 5.76 seconds.

3 COIN-O-MATIC

COIN-O-MATIC performs automatic classification of coins in five stages: (1) segmentation, (2) feature extraction, (3) preselection, (4) classification, and (5) verification. Segmentation is the separation of the coin from the background of the coin photograph. Feature extraction is the transformation of the segmented coin into an efficient and coin-specific representation. Preselection is the selection of possible coin classes based on area and thickness measurements. Area measurements are performed by counting the number of pixels in the segmented coins, whereas thickness measurements are obtained from a thickness sensor. The preselection stage is not addressed any further in this paper. Classification is the process of mapping the feature representation onto one of the selected coin labels, based on information gathered from the training process. Verification is checking whether two coin images have identical labels, based on visual comparison. Verification is necessary because the test sets contain unknown coins, that are not available in the training set (see section 2). The classification stage contains no explicit way to handle unknown coins.

Figure 2 gives an overview of the workflow of the system. Roughly, the system consists of four subsystems: (1) a segmentation subsystem, (2) a feature extraction subsystem, (3) a classification subsystem, and (4) a verification subsystem. These subsystems are indicated as rectangles in Figure 2. The functionality of the subsystems is discussed separately in sections 4, 5, and 6.

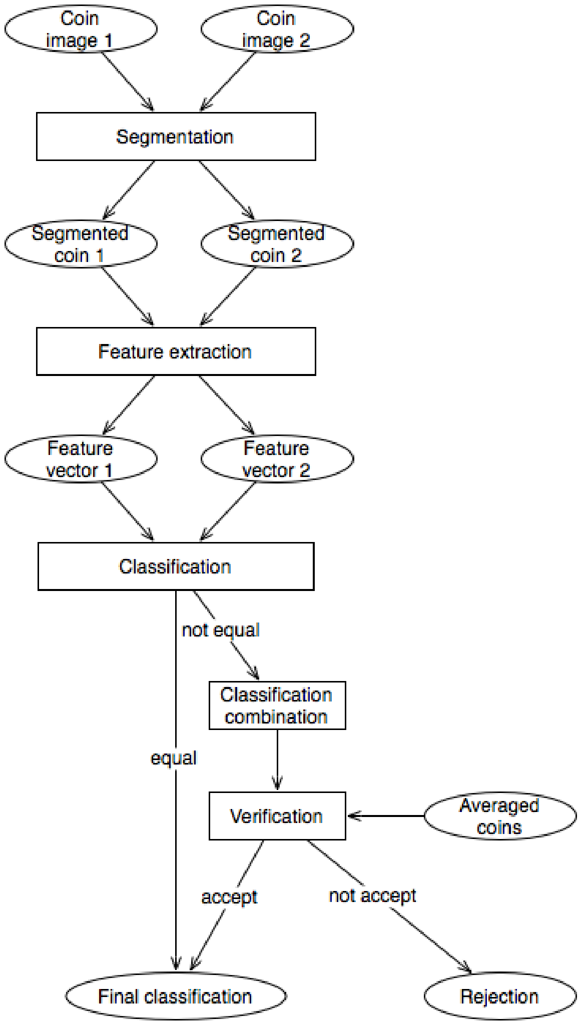

Figure 2: COIN-O-MATIC workflow.

4 Segmentation

Segmentation is the separation of a coin from the background of a coin photograph. Figure 1(a) shows an example of a coin photograph. We propose a two-stage approach for the segmentation of coins. First, we attempt a fast segmentation procedure, which is successful for approximately 85% of the coin photographs. We detect coin photographs for which the segmentation procedure fails, and apply a computationally more intensive segmentation procedure to these coin

photographs.
The fast segmentation procedure consists of three steps: (1) thresholding, (2) edge-detection, and (3) application of morphological operations. The thresholding step is necessary to remove the conveyor belt background from the coin image. We remove all pixels below a threshold of $t = 60$, a value that was determined experimentally. In the edge-detection step, we apply a Sobel edge-detection using a dynamic threshold. From the resulting edge image, we obtain a mask image by subsequently applying a dilation operation, a bucket fill operation from the upperleft corner, and a erosion operation. The resulting mask image is applied on the original coin image in order to segment the coin. An example of a segmented coin is shown in Figure 3.
Although the segmentation procedure described above is very efficient, it fails for very dark

Figure 3: Segmented coin.

coins, since it is not able to discriminate between the coin and the conveyor belt based on threshold t. Dark coins represent approximately 15% of the coin photographs in the dataset. We detect failed segmentations by checking whether the bounding box of the segmented coin is approximately square and reasonable large. If a failed segmentation is detected, we apply a computionally more expensive segmentation procedure. This procedure is roughly the same as the fast segmentation procedure, however, the thresholding is replaced by a convolution with a box filter. The convolution with the box filter removes the conveyor belt structure from the coin photographs, allowing for successful edge-detection on dark coins.

5 Feature extraction

Feature extraction is the extraction of efficient and coin-specific features from coin images. The resulting features can be used to train a classifier.
Important visual information in coins is contained in the stamp of the coin. The coin stamp information corresponds to edge information in the inner part of the coin. Therefore, our system uses features that measure edge-based statistical distributions. Edge-based statistical features yield strong results in e.g., writer identification [2].
This section describes three edge-based statistical features: (1) edge distance distributions (subsection 5.1), (2) edge angle distributions (subsection 5.2), and (3) edge angle-distance distributions (subsection 5.3). Only the latter distribution is used in COIN-O-MATIC, the first two distributions are presented to clarify the concept of edge angle-distance distributions.

The edge images used in the estimation of edge-based statistical distributions are obtained by median filtering the coin images and convolving them with two orthogonal Sobel kernels. A dynamic thresholding operation is applied on the resulting edge image to obtain edge pixels. Subsequently, the edge pixels corresponding to the coin outer border are removed, since they do not discriminate between coins.

5.1 Edge distance distributions

Edge distance distributions measure the distribution of the distances of edge pixels to the center of the coin [13]. The distribution is estimated by dividing the coins in circular concentric parts, as is illustrated in Figure 4(a). The number of edge pixels in each part is accumulated, and the resulting histogram is normalized in order to provide an estimation of the edge distance distribution. Edge distance distributions are rotation invariant by definition.
Edge distance distributions can be used in a multiscale approach, by measuring the histograms for various number of bins (e.g., for 2, 4, 8, 16, and 32 bins).

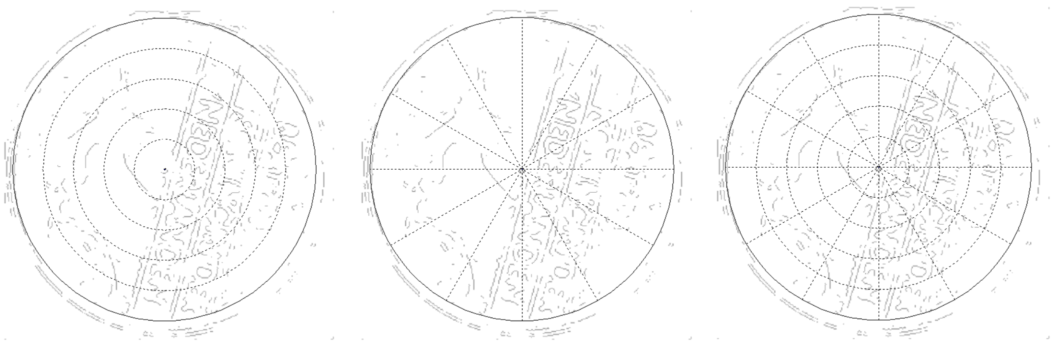

(a) Distance distribution. (b) Angle distribution. (c) Angle-distance distribution.

Figure 4: Edge-based statistical distributions.

5.2 Edge angle distributions

Although edge distance distributions were shown to be strong features in coin classification [13], they do not incorporate all information in coin edge images. In edge distance distributions, the relative angular distribution of the edge pixels is not represented. The relative angular distribution of edge pixels can be described using edge angle distributions. Edge angle distributions are measured by dividing the coin in pie-shaped parts, as is illustrated in Figure 4(b). The number of edge pixels in the parts is accumulated, and the resulting histogram is normalized in order to provide an estimation of the relative angular distribution of the coin edge pixels.
In contrary to edge distance distributions, edge angle distributions are not rotation invariant by definition. Rotation invariance of the edge angle feature can be obtained by computing the magnitude of the Fourier transform of the obtained histogram [12]. This step makes the histogram invariant under circular shifts (which correspond to rotations of the coin). In this respect, a large

number of bins in the histogram is required, since a rotation of the coin should imply a circular shift on the histogram, instead of a change in the histogram accumulators.

Edge angle distributions can be measured in a multiscale approach in the same way as edge distance distributions. However, one should note that the edge angle histogram should only be measured on fine scales (e.g., for 180 and 270 bins), because for coarser scales the rotation invariance of the feature is lost.

5.3 Edge angle-distance distributions

In order to give a good characterization of the distribution of edge pixels over a coin, angular and distance information should be combined. The combination can be constructed in two ways. First, the edge distance distributions and edge angle distributions can be measured seperately, and the resulting feature vectors can be combined into one feature vector in the classification stage. Second, a joint angle-distance distribution can be measured. We refer to the latter distributions as edge angle-distance distributions.

Edge angle-distance distributions measure the joint angle-distance distribution of edge pixels in the coin image. In practice, this implies dividing the coin image into parts as illustrated in Figure 4(c). The number of edge pixels is binned for each part. Normalization of the resulting histogram results in an estimation of the joint angle-distance distribution of edge pixels. The feature is made rotation invariant by computing the magnitudes of the Fourier transforms of all distance bands in the distribution.

In COIN-O-MATIC, we use edge angle-distance distributions that are measured using 2, 4, 8, and 16 distance bins, and 180 angle bins. This results in a feature vector with 5400 dimensions. To reduce the dimensionality of the feature vectors, PCA can be applied. We performed experiments, in which we reduced the dimensionality of the feature vectors to 200 using the SPCA procedure [10].

The computation of edge angle-distance features can be performed very fast. The number of distance and angle computations that has to be performed to estimate the distribution is equal to the number of edge pixels in the coin edge image (which is generally low). The computation of the rotation invariant feature can be done using the FFT algorithm, leading to a computational complexity of $O(km \log m)$, where k is the total number of distance bins (which is equal to 30 in our case), and m is the number of angle bins (which is equal to 180 in our case).

6 Classification and verification

In the previous section, we presented a number of coin-specific features, which can be used to assign a label to a coin image. This section describes how the features are used in order to obtain a reliable decision on the class of the coin. The decision process consists of two stages: (1) classification and (2) verification. In subsection 6.1, the classification stage is described. Subsection 6.2 describes the verification stage.

6.1 Classification

For the classification of a coin, we apply a nearest-neighbour approach in the constructed feature space. We selected a nearest-neighbour approach, because nearest-neighbour approaches usually yield good performances on problems with a high number of classes. In COIN-O-MATIC, we use a 3-nearest neighbour classifier.

Since a coin has two sides, we have to evaluate two coin images in order to classify a coin. The two coin images are first classified separately. If the classifications of both coin images are equal, we classify the coin accordingly. If the classifications of both coin images differ, we create a hitlist of the 15 nearest neighbours for the coin images, and compute a score for all classifications based on these two hitlists. In the computation of the scores, the score of a classification with position p on the hitlist is increased by $15 - p$. The classification that has received the maximum score in the procedure is selected as the final classification.

6.2 Verification

Once we have assigned a label to a test sample, the assigned label is verified in a verification process. Verification is necessary for two reasons: (1) the test set contains unknown coins that should be rejected by the system and (2) reliability of the classification is very important in the MUSCLE CIS benchmark, and cannot be guaranteed by the classification procedure.

We perform verification only for coins for which the two coin faces were classified differently (i.e. for classifications obtained using the hitlist procedure described in subsection 6.1). The acceptance of the other coin classifications is based on the assumption that it is unlikely for both classifiers to make exactly the same mistake.

The verification procedure is based on computing the mutual information of the test sample and the averaged coin image (see Figure 1(b)) that corresponds to the classification assigned to the test sample. In order to compute the mutual information measure, first the rotation of the test sample is normalized. This is done by maximizing the correlation of the averaged coin and the test sample under number of rotations. For computational reasons, these calculations are performed in polar space. After the normalization of the rotation of the test sample, the intensity gradients of the test sample and the averaged coin are computed. The two intensity gradients are blurred by a convolution with a Gaussian kernel. Subsequently, we compute the mutual information m of the resulting blurred intensity gradients. Mutual information is defined as

$$m = H(X) + H(Y) - H(X, Y) \tag{1}$$

where $H(X)$ is the Shannon entropy, defined by

$$H(X) = -\sum_i \log(p_i) * p_i \tag{2}$$

in which p_i is the relative frequency of a pixel value in the image. The base number of the logarithm is 256, since grayscale image pixels can have 256 different values. The value of m is used as a rejection value, on which a threshold t_m is applied. This leads to an acceptance or a rejection of the coin classification. Experimentally, we found $t_m = 0.15$ to be a good value.

Feature	*Verification*	*Correct*	*Incorrect*	*Unknown*	*Comp. time*
MSEADD	yes	72.1%	1.9%	27.0%	7,912 sec.
MSEADD	no	78.0%	16.8%	5.2%	4,723 sec.
MSEADD-PCA	yes	71.5%	3.5%	25.1%	4,187 sec.
MSEADD-PCA	no	74.5%	20.4%	5.2%	1,160 sec.

Table 1: Results on the MUSCLE CIS dataset.

7 Experiments

In order to evaluate the performance of the COIN-O-MATIC, we performed experiments on a test set supplied for the MUSCLE CIS benchmark. Table 1 presents the results for an experiment with a test set of 5,000 coins. We report the percentage of correct and incorrect classifications, as well as the percentage of coins that is rejected by the system (i.e. coins that were classified as unknown). In addition, Table 1 presents the the computation time consumed for classifying the entire test set, as measured on a 2.8 GHz PC with 1 GB RAM.

The results are reported for multi-scale edge angle-distance distributions (MSEADD) and for multi-scale edge angle-distance distributions on which PCA was applied (MSEADD-PCA). We present results for the case in which the verification is turned on (i.e. focus on low number of misclassifications) and for the case in which the verification is turned off (i.e. focus on high number of correct classifications). The small number of unknown classifications in the cases in which verification is turned off is caused by detected failed segmentations of coins.

Table 1 reveals that the verification stage is capable of rejecting a high number of misclassifications, while accepting a relatively high number of correct classifications. Furthermore, Table 1 reveals that the application of PCA on the multi-scale edge angle-distance distributions reduces the computational requirements of the system, with only a small loss in classification performance. This speed-up is caused by the reduced dimensionality of the feature space, which speeds up our (lazy) classifier. In the final COIN-O-MATIC system, we use multi-scale edge angle-distance distributions without PCA, and have turned the verification stage on.

Analysis of the misclassifications of the system leads to the observation that misclassifications are usually made for very dark coins, which lack the contrast necessary to construct a discriminating coin feature. However, most of the misclassification caused by a lack of contrast are eliminated in the verification stage.

8 Conclusions

We have presented a new system for fast and reliable coin classification, that uses edge-based statistical features in order to classify coin images. We have shown promising results for our system on a test set available for the MUSCLE CIS benchmark, achieving a good correct/incorrect classification ratio. In addition, the system is computationally efficient.

Future work should focus on improving the classification performance on dark coins, e.g., by the application of edge-enhancing filters [4]. Furthermore, the verification procedure should be improved in order to further reduce the percentage of incorrect classifications. We surmise the verification can be improved by a better preprocessing of the coin images, e.g. by applying contrast stretching or edge enhancing techniques, or by measuring the mutual information on a number of coarse-to-fine scales. A further speed-up of the system can be obtained by applying techniques such as LAESA [7] in the nearest-neighbour classifier.

Acknowledgements

This work was supported by NWO/CATCH under grant 640.002.401.

References

[1] M. Adameck, M. Hossfeld, and M. Eich. Three color selective stereo gradient method for fast topography recognition of metallic surfaces. *Proceedings of the SPIE*, 5011:128–139, 2003. 2

[2] M. Bulacu, L. Schomaker, and L. Vuurpijl. Writer identification using edge-based directional features. In *Proceedings of ICDAR 2003*, pages 937–941, 2003. 5

[3] P. Davidsson. Coin classification using a novel technique for learning characteristic decision trees by controlling the degree of generalization. In *Proceedings of the* 9^{th} *IEA/AIE*, pages 403–412, 1996. 2

[4] D.A. Forsyth and J. Ponce. *Computer vision: A modern approach*. Prentice Hall, 2003. 9

[5] M. Fukumi, S. Omatu, F. Takeda, and T. Kosaka. Rotation-invariant neural pattern recognition system with application to coin recognition. *IEEE Transactions on Neural Networks*, 3(2):272–279, 1992. 2

[6] R. Huber, H. Ramoser, K. Mayer, H. Penz, and M. Rubik. Classification of coins using an eigenspace approach. *Pattern Recognition Letters*, 26(1):61–75, 2005. 2

[7] M.L. Micó, J. Oncina, and E. Vidal. A new version of the nearest-neighbour approximating and eliminating search algorithm (AESA) with linear preprocessing time and memory requirements. *Pattern Recognition Letters*, 15(1):9–17, 1994. 10

[8] M. Nölle. Distribution distance measures applied to 3D object recognition: A case study. In *Proceedings of the* 25^{th} *Pattern Recognition Symposium of the German Association for Pattern Recognition*, pages 84–91, 2003. 2

[9] M. Nölle, H. Penz, M. Rubik, K. Mayer, I. Holländer, and R. Granec. Dagobert - a new coin recognition and sorting system. In *Proceedings of the DICTA 2003*, pages 329–338, 2003. 2

[10] M. Partridge and R. Calvo. Fast dimensionality reduction and Simple PCA. *Intelligent Data Analysis*, 2(3):292–298, 1997. 7

[11] P.A. Passeraub, P.-A. Besse, C. de Raad, O. Dezuari, F. Quinet, and R.S. Popovic. Coin recognition using an inductive proximity sensor microsystem. In *International Conference on Solid State Sensors and Actuators*, volume 1, pages 389–392, 1997. 2

[12] T. Szoplik and H.H. Arsenault. Rotation-variant optical data processing using the 2D nonsymmetric Fourier transform. *Applied Optics*, 24(2):168–172, 1985. 6

[13] L. van der Maaten and E. Postma. Towards automatic coin classification. In *Proceedings of the EVA-Vienna 2006*, page (in press), 2006. 2, 6