Available online at www.sciencedirect.com
SCIENCE DIRECT®
Pattern Recognition Letters 26 (2005) 61–75
Pattern Recognition Letters
www.elsevier.com/locate/patrec

Classification of coins using an eigenspace approach

Reinhold Huber [a,*], Herbert Ramoser [a], Konrad Mayer [b], Harald Penz [b], Michael Rubik [b]

[a] *Advanced Computer Vision GmbH - ACV, Donau-City-Strasse 1, A-1220 Vienna, Austria*
[b] *ARC Seibersdorf research GmbH, High Performance Image Processing, A-2444 Seibersdorf, Austria*

Received 27 November 2003; received in revised form 3 June 2004
Available online 18 October 2004

Abstract

We present a vision-based approach to coin classification which is able to discriminate between hundreds of different coin classes. The approach described is a multistage procedure. In the first stage a translationally and rotationally invariant description is computed. In a second stage an illumination-invariant eigenspace is selected and probabilities for coin classes are derived for the obverse and reverse sides of each coin. In the final stage coin class probabilities for both coin sides are combined through Bayesian fusion including a rejection mechanism. Correct decision into one of the 932 different coin classes and the rejection class, i.e., correct classification or rejection, was achieved for 93.23% of coins in a test sample containing 11,949 coins. False decisions, i.e., either false classification, false rejection or false acceptance, were obtained for 6.77% of the test coins.

Keywords: Classification; Bayesian fusion; Multistage classifier; Multiple eigenspaces; Invariant eigenspace; Coin sorting

1. Introduction

Charity organizations such as the Austrian "Licht ins Dunkel" collect coins from all over the world as donations. The classification of these coins according to their denomination and country of origin is the field of application for the method presented in this paper. The "Licht ins Dunkel" coin collection took place in the course of the implementation of the Euro currency in twelve European countries at the turn of the year 2001/2002 when Austrian citizens were asked to donate any spare coins they might have at home. During this campaign 300 tons of coins coming from virtually all countries of the world but predominately from the twelve Euro member states were collected. Such a large variety of coins exceeds the capabilities of off-the-shelf electromechanical coin discriminating machines.

* Corresponding author. Tel.: +43 1 2696255 137; fax: +43 1 2696255 200.
E-mail address: reinhold.huber@acv.ac.at (R. Huber).

0167-8655/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.patrec.2004.09.006

Coin judging systems using electromechanical devices are commonly based on measuring weight, diameter, thickness, permeability and conductivity (Davidsson, 1996), oscillating electromagnetic field characteristics (Neubarth et al., 1998), and photo- and piezoelectric properties (Shah et al., 1986). Coin classification using image processing is described by Fukumi et al. (1992), and uses a neural network approach to discriminate between 500 Won and 500 Yen coins. Very recently, Nölle et al. (2003) and Fürst et al. (2003) described the Dagobert coin recognition and sorting system for high volumes of coins and a large number of currencies. Our approach uses the same mechanical setup as described in Fürst et al. (2003), but a method different from the matching approach described in Nölle et al. (2003).

A number of coin authentification methods employing optical means are to be found in patents, e.g., Hibari and Arikawa (2001) describe a system by which both sides of a coin are first imaged by cameras, then features are derived from binarized images, and finally results are combined for each side and with a magnetic sensor measurement. Another approach, also based on binarization followed by area measurement and comparison of coin centre and centre of gravity, is suggested by Onodera and Sugata (2002). Tsuji and Takahashi (1997) analyze one side of a coin by transformation of its image into polar coordinates and matching of profiles taken along angle direction.

The approach by Uenohara and Kanade (1997) for matching rotated patterns is similar to the approach presented. Uenohara and Kanade describe a method by which an eigenspace is constructed from uniformly rotated images. The application of their approach works through locating of a specific small pattern in a larger image, whereas our approach aims at classification for a large number of classes. In a later paper, Uenohara and Kanade (1998) describe an improvement of their template location method which is based on the discrete cosine transform (DCT). Parts of our approach are also based on utilization of the eigenspace approach. Our method differs from the one suggested by Uenohara and Kanade as the eigenspace is constructed from images rotated into reference orientation. Therefore, each image is contained only once in the set of images used to construct the eigenspace.

The eigenspace method is a subspace method used to construct the most expressive subspace. Linear discriminant analysis (LDA) (Belhumeur et al., 1997) and support vector machines (SVM) (Burges, 1998) construct most discriminative subspaces. For a given set of classes the subspace constructed by LDA maximizes the ratio of between-class scatter to within-class scatter. The disadvantage of LDA concerning our classification problem, by which varying subsets of classes are distinguished, is that for each combination of classes a specific subspace is required. Furthermore, for some classes only a single example coin exists, therefore it becomes impossible to estimate the within-class scatter. SVMs use the so called kernel mapping technique to map data which is not linearly separable into a hyperspace where it can be separated linearly. SVMs are able to discriminate between two classes, for $K > 2$ classes ensemble schemes such as one per class (OPC) and pairwise coupling (PWC) are employed (Goh et al., 2001). Unfortunately, those schemes are computationally very demanding, i.e., for a K class problem using OPC K binary classifications are necessary and for PWC $K(K-1)/2$ binary classifiers are needed.

We present a vision-based coin classification approach which is able to discriminate between a very large number of coins, i.e., coins from 30 countries and a total number of 932 different coin images. The presented approach is a multistage procedure (Pudil et al., 1992), which splits into the following stages:

1. Processing for translational and rotational invariance and hypotheses generation.
2. Eigenspace selection and classification.
3. Bayesian fusion and rejection.

Detection of the coin, i.e., removing translational dependency, processing for rotational invariance and generation of classification hypotheses is done in stage 1. A small number of classification hypotheses for obverse and reverse sides of the coins are retained for the subsequent steps.

Hypotheses are derived from the correlation between observed and reference images after mapping to polar coordinates. In stage 2 an appropriate eigenspace is selected depending on coin thickness and diameter estimation. Thickness and diameter measurements allow the construction of multiple eigenspaces, each of which spans only a portion of the thickness/diameter plane. Thus the number of coins represented in a single eigenspace is reduced. Depending on thickness and diameter up to 113 classes are represented in a single eigenspace; on the average 30.9 classes are comprised in a single eigenspace. Distances in eigenspace are evaluated for the hypotheses generated in stage 1 and are related to probabilities. Stage 3 delivers a final decision through Bayesian fusion applied to the coin sides. Bayesian fusion incorporates probabilities for both coin sides and knowledge about rotation angle coherence between obverse and reverse coin sides. A rejection mechanism based on thresholding the fused probability is used to sort out unknown coins and decisions of weak support. The rejection threshold is chosen so as to optimize a compound objective, i.e., maximization of correct classification (a known coin is correctly classified) and correct rejection (an unknown coin is not classified) and minimization of false classification (a known coin is classified into a wrong class), false rejection (a known coin is rejected), and false acceptance (an unknown coin is classified).

The mechanical setup, which is not discussed in this paper, comprises a number of devices (Fürst et al., 2003, Rupp et al., 2003). First, the coins are transported via a singularization device onto a conveyor belt. The conveyor belt moves the coins past two cameras to a sorting mechanism. Between the cameras there is a turning device, i.e., a mechanical apparatus for reversing the coin. The cameras use circular lighting and a protective device which blocks ambient light. Additionally, in front of each camera a light curtain allows thickness measurements for each coin. The conveyor belt finally transports the coins to a series of boxes into which the coins are pneumatically ejected according to the classification result.

In this paper we will describe image processing and coin classification. Section 2 summarizes the mathematical notation for the eigenspace approach and reviews the concept of multiple eigenspaces as applied to the coin classification problem. In Section 3 invariance with respect to translation, rotation and illumination is discussed for coin pattern processing. Section 4 introduces our approach to classification, fusion and rejection. Section 5 presents results and concluding remarks are found in Section 6.

2. Eigenspaces

Eigenspace decomposition for image analysis was introduced by Sirovich and Kirby (1987) and found numerous applications over the last years, most prominently in the field of face recognition. We start with the description of the mathematical procedure of eigenspace construction employing principal components analysis (PCA) in the next subsection. Subsequently, we discuss multiple eigenspaces in the context of coin recognition.

2.1. Principal components analysis for coins

We consider a set of M images B_1 to B_M. Each image B_i is of size $N * N$ pixels. The images are reformed into vectors Γ_1 to Γ_M, e.g., by scanning the image line by line. If all pixels of an image are used to produce a vector, each vector Γ_i has length $L = N^2$. An average vector Ψ and difference vectors ϕ_i are calculated by

$$\Psi = \frac{1}{M}\sum_{i=1}^{M}\Gamma_i, \quad \phi_i = \Gamma_i - \Psi, \quad i = 1,\ldots,M. \tag{1}$$

The covariance matrix for ϕ_i is defined as

$$C = \frac{1}{M}\sum_{i=1}^{M}\phi_i\phi_i^{\mathrm{T}} = AA^{\mathrm{T}}, \quad A = (\phi_1,\phi_2,\ldots,\phi_M). \tag{2}$$

Principal axes are obtained by the eigendecomposition of the covariance matrix C (Sirovich and Kirby, 1987; Turk and Pentland, 1991). The eigenvectors are sorted in non-increasing order depending on the corresponding eigenvalue. A small number M' of significant eigenvectors is retained from the ranked eigenvalues, a common practice

which leads to the most expressive features (Turk and Pentland, 1991).

A weighting factor ω_k corresponding to the kth eigenimage for a new reformed image Γ is obtained by projection onto the kth eigenspace component using

$$\omega_k = u_k(\Gamma - \Psi), k = 1, \ldots, M'. \quad (3)$$

The weight ω_k are arranged in an vector $\Omega = (\omega_k, \ldots, \omega_k)^{\mathrm{T}}$. For the coin recognition task, not the full images are reformed into a vector, only the interior pixels of the coin are rearranged into the vector Γ, see Fig. 1.

2.2. Multiple eigenspaces

Multiple, or modular, eigenspaces have been introduced for viewing orientation independent face recognition by Pentland et al. (1994), that is a set of view-based eigenspaces is constructed, each capturing the different classes under a common view. Recently, a self-organizing framework to construct multiple eigenspaces was suggested by Leonardis et al. (2002). An alternative approach is the parametric eigenspace, presented for recognition of 3D objects under varying pose and illumination by Murase and Nayar (1995). In the latter approach invariance to illumination and geometric transformations is achieved by varying pose and illumination for the images used to construct the eigenspace. In the coin recognition task, we regard coins as 2D objects taken under controlled illumination. Multiple eigenspaces are constructed to group coins depending on their diameter and thickness measurements. Thus each eigenspace is constructed from a subset of images and, consequently, only a moderate number of classes is represented in a single eigenspace. The advantages of multiple eigenspaces spread over the coin diameter and thickness plane are twofold. Firstly, construction of eigenspaces from a lower number of images is computationally less demanding, and secondly, classification can be performed faster and more reliably among the smaller number of classes present in a single eigenspace than between all considered cases.

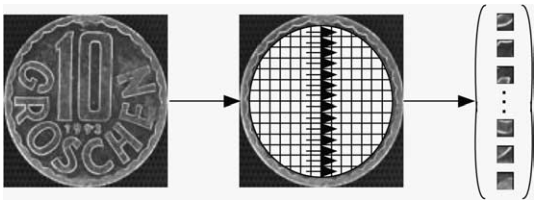

Fig. 1. Rearrangement of the inner pixels of a coin into a vector.

As diameter and thickness are provided with each coin image the diameter/thickness plane is partitioned into regularly spaced regions. For each such region an eigenspace is built, therefore region m is denoted by Es_m. A region Es_m is characterized by its minimum and maximum diameter dl_m, dh_m, and minimum and maximum thickness tl_m, th_m, respectively. The extensions of each region are D_{step} and T_{step}. The global minimum and maximum measures for diameter and thickness, $D_{\min}$, $D_{\max}$, $T_{\min}$, $T_{\max}$ are derived from coin statistics. In particular, for a total number of K coin sides, we find

$$\begin{aligned} D_{\min} &= \min_{i=1,\ldots,K}(dmn_i - \text{tol}), \\ D_{\max} &= \max_{i=1,\ldots,K}(dmx_i + \text{tol}), \\ T_{\min} &= \min_{i=1,\ldots,K}(tmn_i - \text{tol}), \\ T_{\max} &= \max_{i=1\cdots K}(tmx_i + \text{tol}), \end{aligned} \quad (4)$$

dmn_i, dmx_i, tmn_i, tmx_i are the limits calculated from the measurement statistics for coin i. A tolerance value tol accounts for variations of a specific coin and measurement errors. A coin side class B_i typically covers more than one eigenspace. The number of coin side classes per eigenspace is calculated as

$$\#BEs_m = \left| \bigcup_{i=1,\ldots,K} \begin{bmatrix} B_i : (dmn_i - \text{tol}) < dh_m \\ \cap\ (tmn_i - \text{tol}) < th_m \\ \cap\ (dmx_i + \text{tol}) > dl_m \\ \cap\ (tmx_i + \text{tol}) > tl_m \end{bmatrix} \right|. \quad (5)$$

Fig. 2 shows a partition of the diameter/thickness plane, where each cell ES_m denotes a single eigenspace. Each coin side B_i covers at least one eigenspace. Due to diameter and thickness tolerances and measurement errors, each B_i is generally represented in more than one eigenspace (see Fig. 2).

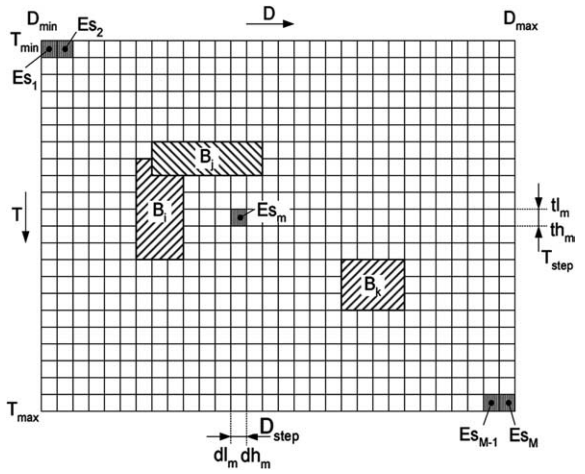

Fig. 2. Partitioning of the diameter/thickness plane into multiple eigenspaces.

Not each eigenspace is occupied by coin classes, e.g., combinations of large diameter with small thickness and vice versa do not exist. Some eigenspaces are occupied by a single coin side, which is unexpected as diameter and thickness should be the same regardless whether the obverse and reverse side of the coin have been looked at. This problem is caused by small measurement noise. Finally, some neighboring eigenspaces represent the same set of coin side classes. For the considered sample the most occupied eigenspaces are found around 2 mm in thickness and 24 mm in diameter, there a maximum #*BEs* of 113 was found. More specifically, D_{step}, T_{step} and tol were set to 0.3 mm, which results in 1000 potential eigenspaces, 494 of which were occupied by at least one coin side image. Out of those 494 coin side images 21 eigenspaces contained only one coin side image, which makes classification redundant, and 95 eigenspaces were replications of others. On the average the number of coin side classes per eigenspaces #*BEs* is 30.9.

3. Invariant preprocessing

Invariant preprocessing is a necessary step for template matching based pattern recognition methods such as the eigenspace approach. Rotation, scaling and translation are the three geometrical variants to be taken into account. Illumination is another varying quantity mainly resulting from the variance in reflection caused by abrasion and dirtiness of the coins illuminated with controlled lighting. Due to the static setup with known distance between camera and object, scale invariance is not addressed in our work. Translational invariance is achieved by coin detection through segmentation: the segmentation result holds the centered coin image. Section 3.1 discusses coin segmentation. For a more general discussion, please refer to the large number of works related to scale and translation invariant pattern recognition, e.g., the Fourier-Mellin transform (Altmann and Reitböck, 1984). Rotational invariance is achieved through coordinate mapping and correlation, described in detail in Section 3.2. Illumination invariance for the eigenspace approach is discussed in detail in Section 3.3.

3.1. Translational invariance

Detection of the coin employs a common segmentation approach and works reliably for controlled lighting conditions and clean background, i.e., a moderately dirty conveyor belt. Problems might be caused by very dark coins, i.e., coins which reflect only a small amount of light towards the camera. At that point, a multi-stage segmentation procedure is suggested. The outline of the segmentation method is:

1. Smoothing of the image to suppress the conveyor belt texture.
2. Edge filtering using a Laplacian of Gaussian approach followed by zero-crossing detection (Marr and Hildreth, 1980).
3. Labeling of the detected regions and selection of the region with largest bounding box as coin region candidate.
4. Form a blob by computing the convex hull of the coin region candidate.

Coin position and diameter are estimated from the detected blob, which directly delivers a translation invariant description.

3.2. Rotational invariance

Rotational invariance is sometimes approached via the use of geometrical moments (Hu, 1962), radial coding of features (Torres-Mendez et al., 2000), or using a mapping from Cartesian to polar coordinate representation, e.g., log-polar mapping (Kurita et al., 1998). We obtain rotational invariance by estimation of the rotational angle followed by a rotation into a reference pose. Angle estimation is performed for images transformed into polar coordinates. In the polar image shift invariance, corresponding to rotational invariance when mapped back to Cartesian coordinates, is achieved through cross-correlation. Cross-correlation is efficiently implemented using the fast Fourier transform (FFT) (e.g., Ballard and Brown (1982, p. 24)).

Rotational invariance for a coin image involves cross-correlation with reference images. The set of reference images for a specific coin image is, like in Section 2.2, determined from diameter and thickness measurements. To diminish illumination effects edge detection is performed for the coin image and the reference images. The used edge detector is described in Rothwell et al. (1995) and is an extension of the widely accepted edge detector suggested by Canny (1986). The edge image is mapped from Cartesian to polar coordinates (see Fig. 3 in the middle). The result of cross-correlation between the coin image to be classified and a set of reference images is used to derive class hypotheses, i.e., this corresponds to the outcome of the first stage in the proposed multi-stage scheme. In detail, for both sides of a coin under investigation rotational invariant processing and hypothesis generation proceeds as follows:

1. Estimation of coin thickness from thickness sensor measurement and estimation of coin diameter from coin detection as described in Section 3.1.
2. Selection of a set of reference images, each reference image is associated with a coin class, depending on thickness and diameter.
3. Cross-correlation of the coin side image under investigation with all reference coin images in the selected reference set, resulting in a cross-correlation value and associated rotation angle estimation for each reference class.

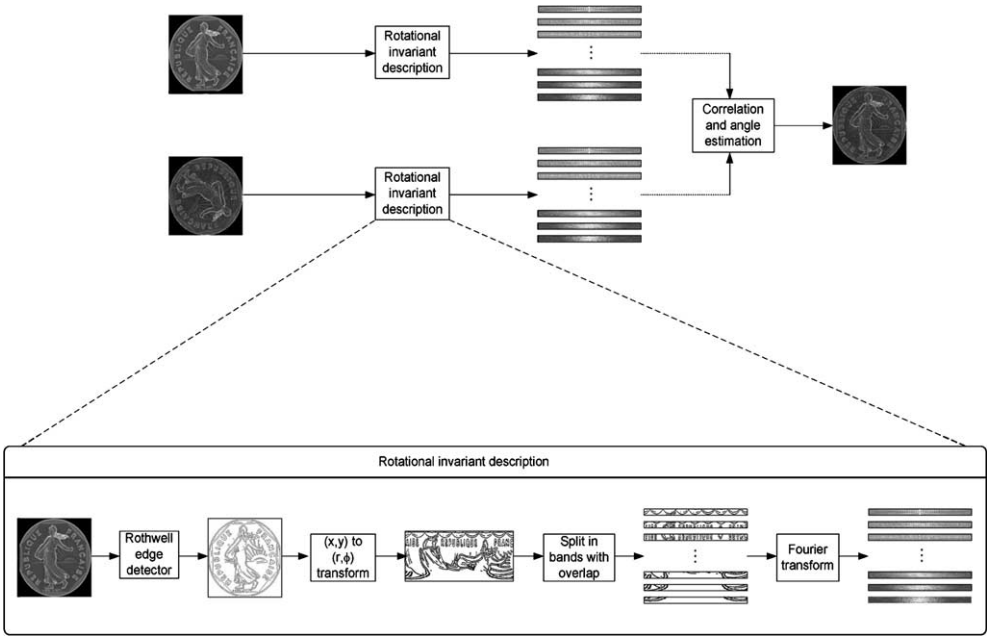

Fig. 3. Processing of coins to achieve rotational invariance.

4. Ranking of the reference set by the maximum correlation value and generation of a set of hypotheses for the highest-ranking classes.

To obtain reliable estimates for cross-correlation and rotation angle the polar image is split into n bands along the radius coordinate, corresponding to concentric rings in Cartesian coordinates. The peak of the correlation value K_i for band i is determined for each band and the position of the peak is taken as an estimate α_i for the rotation angle in band i. For angle data the von Mises distribution can be used as a model of normal distribution on a circular domain (Fisher, 1995). Under this assumption, the sample mean angle direction $\overline{\alpha}$ is estimated via

$$\overline{\alpha} = \begin{cases} \arctan(S/C) & \text{if } S \geqslant 0 \text{ and } C > 0, \\ \arctan(S/C) + \pi & \text{if } C < 0, \\ \arctan(S/C) + 2\pi & \text{if } S < 0 \text{ and } C > 0, \end{cases} \tag{6}$$

in which $C = \sum_{i=1}^{n} \delta_i \cos\alpha_i$, $S = \sum_{i=1}^{n} \delta_i \sin\alpha_i$, where $\delta_i = 1$ if band i contains a significant number of edge pixels in reference coin and coin under investigation, otherwise $\delta_i = 0$. A cross-correlation estimate K for the coin under investigation is calculated using $K = \frac{1}{n'}\sum_{i=1}^{n} \delta_i K_i$. The number of bands $n' \leqslant n$ used in cross-correlation and angle estimation varies between images and is simply obtained by $n' = \sum_{i=1}^{n} \delta_i$.

The computational complexity of the correlation procedure is $O(r n' M N \log_2 MN)$, where r is the number of reference images, M and N are the dimensions of an image band. Using cross-correlation based on FFT the time consumption was measured to be 0.05 ms for an image band of 8×256 pixels on a 3 GHz Pentium IV platform and optimized code. Using the decomposition of the diameter/thickness plane as shown in Fig. 2, the same example parameters as in Section 2.2 and $n' = n = 5$ bands, the worst case time consumption, i.e., $r = 113$, becomes $5 \times 113 \times 0.05 = 28.25$ ms, and on the average, i.e., $r = 30.9$, the time consumption is 7.725 ms.

3.3. Illumination invariance

To overcome limitations regarding illumination variation in the eigenspace approach a number of solutions were proposed, e.g., Murase and Nayar (1994) investigate the determination of the illumination which gives best discrimination. The PCA of edge images and smoothed edge images is

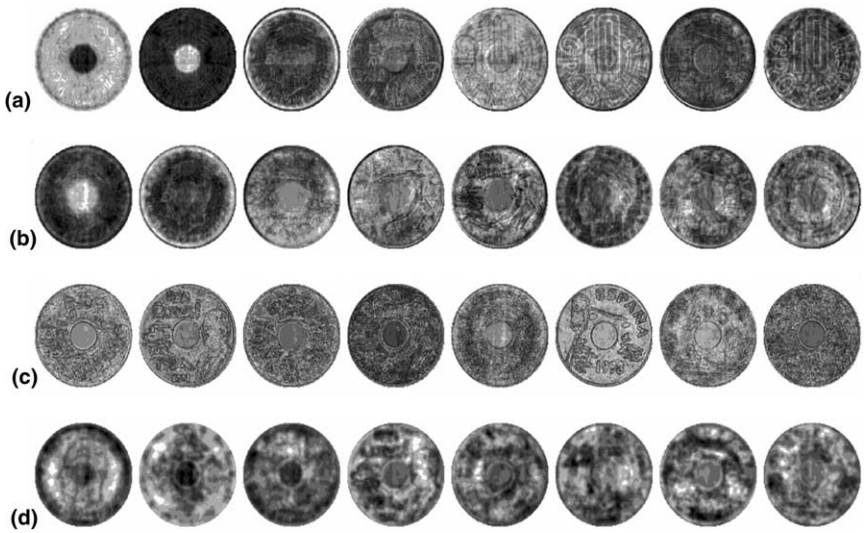

Fig. 4. First 8 eigenimages constructed from: (a) intensity images, (b) equalized intensity images, (c) edge images, (d) smoothed edge images.

suggested as an illumination invariant way of eigenspace construction in Yilmaz and Gökmen (2000), in Venkatesh et al., 2002 gradient images are used as input to PCA, and Bischof et al. (2001) use a set of gradient based filter banks applied to the eigenimage representation. We compare the eigenspace approach on original images with three approaches aimed at reduction of illumination variance in the following.

3.3.1. Intensity eigenspace

Fig. 4a shows the first eight eigenimages constructed from graylevel images, the top left image is the eigenimage corresponding to the largest eigenvalue. The eigenimages are regarded as the most expressive features (MEF) when sorted in non-increasing order of their corresponding eigenvalue (Swets and Weng, 1996). Fig. 5a gives the normalized cumulative sum of the sorted eigenvalues, i.e., the first 32 eigenimages retain approximately 78% of the variance present in the original set of intensity images.

3.3.2. Equalized intensity eigenspace

Histogram equalization results in an image with a uniform distribution of graylevels (e.g., Ballard and Brown (1982, p. 70)). Hence, histogram equalization is sometimes suggested as a way to achieve illumination invariance. Fig. 4b shows the most expressive eigenimages constructed from histogram equalized images. As shown in Fig. 5b approximately 60% of the variance present in the

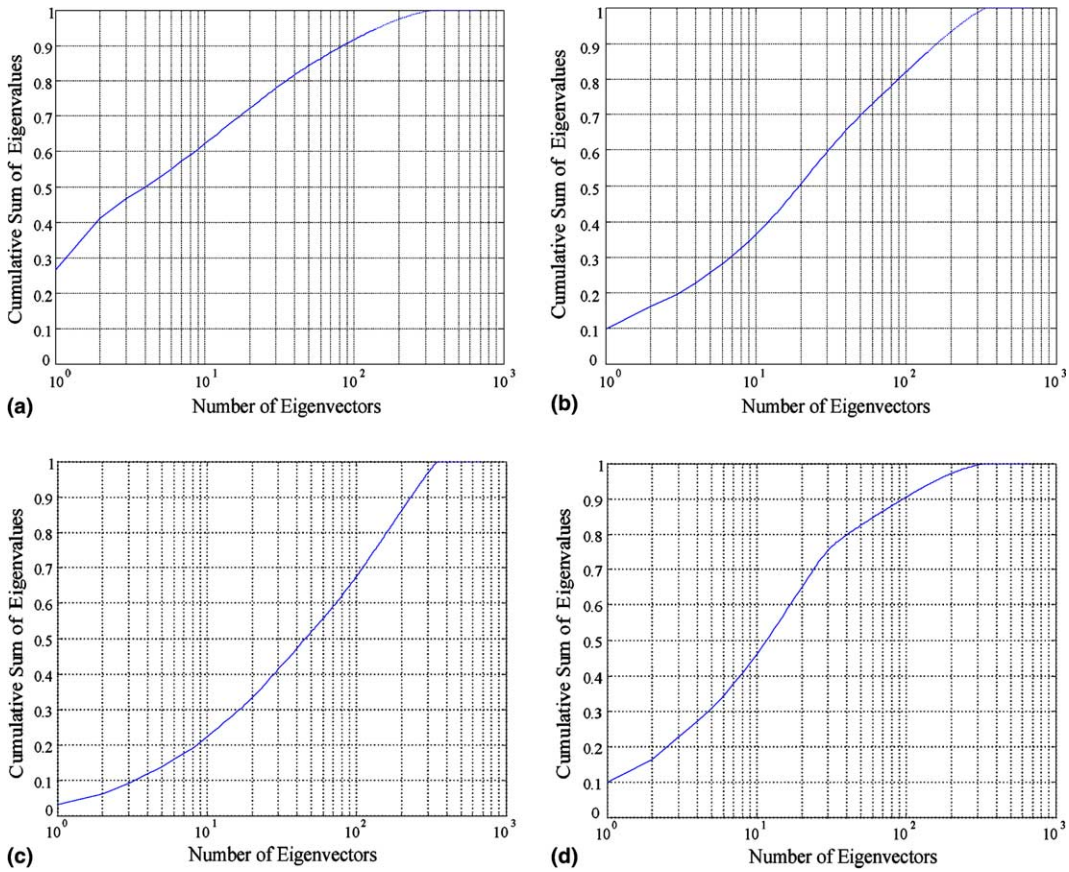

Fig. 5. Cumulative sum of eigenvalues depending on the number of eigenvectors for an eigenspace derived from (a) intensity images, (b) equalized intensity images, (c) edge images, (d) smoothed edge images.

original set of histogram equalized intensity images is contained in the first 32 sorted eigenimages. In terms of the MEF the equalized intensity eigenspace is significantly worse than the intensity eigenspace.

3.3.3. Edge eigenspace

Edges extracted by an edge detector, in our case the Rothwell edge detector (Rothwell et al., 1995), are robust features over a wide range of illumination conditions. Fig. 4c shows the most expressive eigenimages constructed from edge images. As shown in Fig. 5c approximately 42% of the variance present in the original set of edge images is contained in the first 32 sorted eigenimages. This low value is mainly caused by small errors in the preprocessing stages, i.e., imprecision in coin segmentation or determination of the rotation angle.

3.3.4. Smoothed edge eigenspace

To overcome the drawbacks of an eigenspace derived from edge images, eigenhills have been suggested by Yilmaz and Gökmen (2000). There eigenhills are derived from application of the PCA to edge images which are covered by a "membrane". We used a 2D Gaussian filter kernel with a σ_s of 1.5 to smooth the edge images which are of size 128×128 pixels. Fig. 4d shows the most expressive eigenimages constructed from smoothed edge images. As shown in Fig. 5d approximately 76% of the variance present in the original set of smoothed edge images is contained in the first 32 sorted eigenimages.

3.4. Eigenspace construction

The performance of the smoothed edge eigenspace with respect to MEF extraction is comparable to the intensity eigenspace with the advantage of providing an illumination invariant description, therefore this approach was selected. Smoothed edge eigenspaces were constructed for each subregion shown in Fig. 2 where $\#BEs \geqslant 2$. For each coin class at least one image for obverse and reverse side exists in the training set. The number of images for each coin class varies as training images were collected from the coin pool to be sorted. Therefore, for rare coins only a few examples are present and for prevalent coins a large number of examples could be gathered. We limited the maximum number of examples for a specific coin to 20. The variation of coins within a class, i.e., the within-class scatter, is mainly due to abrasion, dirtiness and differences in a very small region, namely the region containing the date.

4. Classification and rejection

For each coin a decision is made whether the coin is accepted or rejected. The classification system should have the option to reject coins which are invalid to achieve maximum reliability. Invalid coins are coins which are not known to the system, i.e., coins which have not been available at system training time. However, if the rejection mechanism is very severe, known coins will falsely be rejected. The trade-off between false acceptance rate (FAR) and false rejection rate (FRR) is an important performance measure in verification and recognition systems (Golfarelli et al., 1997). Table 1 shows the events false acceptance for invalid coins, measured by the FAR, and false rejection for valid coins, measured by FRR. A classification method should maximize correct classification for valid coins and correct rejection for invalid coins. Apart from FAR and FRR the case of wrong classification of a valid coin is also an undesired event. For that reason, we introduce the term false classification rate (FCR).

The final decision, i.e., classification or rejection, follows from single-side classifications of obverse and reverse side and fusion of both sides and with information from coin orientation coherence. As the coin passes a mechanical reversing unit (Fürst et al., 2003), which turns obverse and reverse side of the coin between the two image acquisitions, the difference in rotation angle between the

Table 1
Decision categories for valid and invalid coins

	Acceptance		Rejection
Valid coin	Correct classification	False classification	False rejection
Invalid coin	False acceptance		Correct rejection

two sides is an additional cue in the classification process.

4.1. Classification

It was shown by Murase and Nayar (1995) that the Euclidean distance between points in eigenspace is an approximation to the similarity between the corresponding images measured by the sum of squared differences. On the other hand, this distance measure does not take into account second-order statistics. Due to the small number of training examples per coin it is not possible to use a measure based on second order statistics, such as the Mahalanobis distance, as this would require a large number of examples per coin (Kittler et al., 1998). Therefore, we base our classification on a Euclidean distance measure.

For a specific eigenspace, we consider a set of reference coefficient vectors $\Omega_r = (\omega_{r1}, \ldots, \omega_{rD})^{\mathrm{T}}$, $r = 1, \ldots, R$, and an observed coefficient vector $\Omega^S = (\omega_1^s, \ldots, \omega_D^s)^{\mathrm{T}}$ corresponding to the coin side to be classified. Classification starts from two observation vectors together with a set of hypotheses, ranked by their corresponding correlation measure, was described in Section 3.2.

We introduce the following notation with typical values of parameters given in brackets:

- R, number of reference coefficient vectors (typically $R = K \cdot 20$);
- S, number of coin sides (usually $S = 2$);
- K, number of classes (typically $K = 2, \ldots, 113$);
- H, number of hypotheses (usually $H = 5$);
- D, dimension of coefficient vector (typically $D = 32$);
- G_h^s, hypothesis number h on side s;
- d_r^s, distance to rth coefficient vector on side s;
- l_r^s, label of rth coefficient vector on side s.

The selection of the parameter R is motivated by the balance between representing occurring variation within a coin and efficient construction of the eigenspace for a $R \times R$ covariance matrix for which the eigenproblem is solved. The eigenproblem for general $R \times R$ matrices require on the order of R^3 arithmetic operations (Pan and Chen, 1999), accordingly a small R is preferred. The number of coin sides S is obviously equal to 2. The maximum number of classes per eigenspace is determined by the partitioning of the diameter/thickness plane discussed in Section 2.2. The number of hypotheses H generated in stage 1 was limited to 5. This decision is motivated by observing the necessary number of hypotheses to ensure that the valid decision is included in the considered set of hypotheses. From a validated set of coins, it was observed, that the correct coin class is contained in 92.62%, 95.15%, 96.91%, 98.04% or 98.88% of all cases when retaining the first 1, 2, 3, 4 or 5 hypotheses, respectively. This means, a stage 1 classifier considering the highest ranking result would not do better than 92.62%. On the other hand, considering 5 hypotheses in the stage 3 classifier limits this classifier to 98.88%, which is a reasonable limit for practical application.

The distance to the rth coefficient vector on side s is calculated by the Euclidean distance

$$d_r^s = \sum_{i=1}^{D} (\omega_i^s - \omega_{ri})^2. \tag{7}$$

The class labels $l_r^s \in \{1, \ldots, K\}$ correspond to the distances d_r^s. The distance for hypothesis h on side s is derived as the average distance to coefficient vectors with class label $G_h^s \in \{1, \ldots, K\}$

$$D_h^s = \frac{1}{N_h} \sum_{r=1}^{R} d_r^s \delta_{rh}^s, \tag{8}$$

where N_h is the number of training samples for class G_h^s and the δ_{rh}^s are given by

$$\delta_{rh}^s = \begin{cases} 1 & \text{if } l_r^s = G_h^s, \\ 0 & \text{else.} \end{cases} \tag{9}$$

The conditional probability for observation Ω^s depending on hypothesis G_h^s on side s is estimated to be inversely proportional to the distance D_h^s

$$P^s(\Omega^s | G_h^s) = \frac{1}{D_h^s \sum_{i=1}^{H} 1/D_i^s}, \tag{10}$$

where the summation term in the denominator accounts for normalization.

4.1.1. A priori probabilities

A priori probabilities $P^s(G_h^s)$ are either set to equal probability, e.g., $P^1(G_h^1) = 1/H$ for side 1,

or are derived from the difference in rotation angle α_h for side 1 (e.g., obverse) and angle α_j for side 2 (e.g., reverse), e.g., $P^2(G_j^2) = P^2(G_h^1, G_j^2) = a + bP(\alpha_h^1, \alpha_j^2)$, $a + b = 1$. The weights a and b account for the fact that a number of coins exist which appear similar under rotation, e.g., some Danish coins have this property. The constant term is chosen relatively small, in our study $a = 0.08$ turned out to be a good choice. Fig. 6a shows the distribution of the absolute turning angle difference for an observed sample of coins.

The prior probability $P^2(\alpha_h^1, \alpha_j^2)$ is assumed normally distributed around zero angle difference for coins with same orientation on front and back side, denoted by $P^2(\alpha_h^1, \alpha_j^2) \equiv P(\alpha_0)$, and around 180° angle difference for coins turned upside down between sides, denoted by $P^2(\alpha_h^1, \alpha_j^2) \equiv P(\alpha_{180})$, respectively

$$P(\alpha_0) = \exp\left(-\frac{|\alpha_0|}{2\sigma^2}\right),$$
$$P(\alpha_{180}) = \exp\left(-\frac{|\alpha_{180}|}{2\sigma^2}\right), \tag{11}$$

where α_0 and $\alpha_{180} = |180 - \alpha_0|$ are defined by

$$\alpha_0 = \begin{cases} |\alpha_h^1 + \alpha_j^2| & \text{if } \alpha_h^1 + \alpha_j^2 \leqslant 180, \\ |\alpha_h^1 + \alpha_j^2 - 360| & \text{if } 180 < \alpha_h^1 + \alpha_j^2 \leqslant 540, \\ |\alpha_h^1 + \alpha_j^2 - 720| & \text{if } \alpha_h^1 + \alpha_j^2 > 540. \end{cases} \tag{12}$$

The parameters a and σ are chosen to approximate the observed distribution, e.g., $a = 0.08$ and $\sigma = 15$. Fig. 6b shows the prior probabilities depending on absolute turning angle difference.

4.2. Fusion

Fusion combines conditional and prior probabilities for each side using Bayes theorem

$$P^s(G_h^s|\Omega^s) = \frac{P^s(\Omega^s|G_h^s)P^s(G_h^s)}{\sum_{i=1}^{H} P^s(\Omega^s|G_i^s)P^s(G_i^s)}. \tag{13}$$

We concentrate on the nominator since the denominator is a constant term. Combination of both sides is done by the product rule (Kittler et al., 1998)

$$P(G_k|\Omega) = P(G_h^1 = G_j^2|\Omega^1, \Omega^2) = P^1(G_h^1|\Omega^1) \cdot P^2(G_j^2|\Omega^2), \tag{14}$$

where probabilities are only derived for hypotheses present for both sides. The product rule is a severe rule of combination, which is a desired behavior as we aim at a high rate of correct classification.

4.3. Rejection

Rejection is based on the a-posteriori probability $P(G_k|\Omega)$, where G_k is the class which maximizes $P(G_k|\Omega) = P^1(G_h^1|\Omega^1) \cdot P^2(G_j^2|\Omega^2)$, for all hypotheses G_h^1 and G_j^2 which are overlapping for obverse

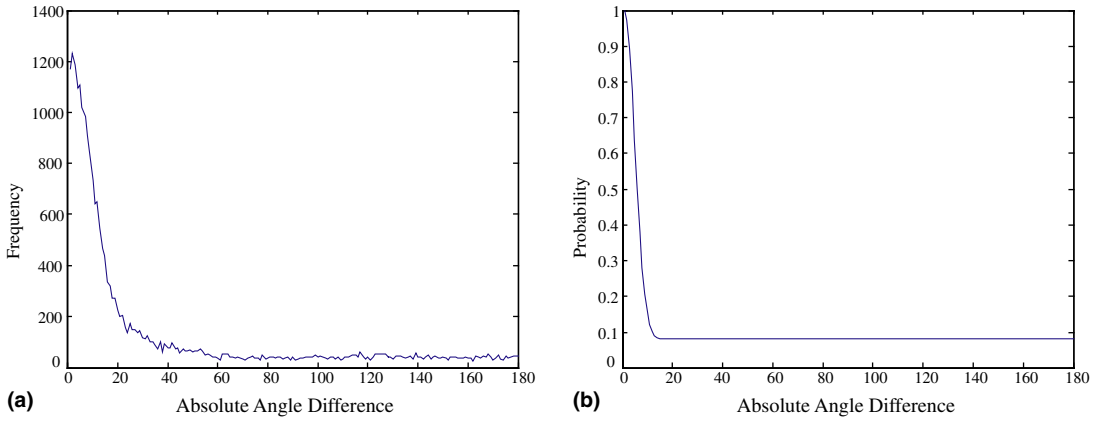

Fig. 6. Turning angle difference histograms: (a) distribution of the estimated absolute turning angle differences, (b) prior probability depending on absolute turning angle difference.

and reverse side. A coin pattern Ω is accepted to be of class G_k if $P(G_k|\Omega) \geqslant t$, and rejected if $P(G_k|\Omega) < t$, where t $(t \in [0,1])$ is the rejection threshold. The parameter t is used to tune the system towards the desired trade-off between false rejection and false acceptance. Section 5 will describe how to find an optimal value for t.

5. Results

Results are presented for a sample of 11,949 coins taken randomly from the collected money. Those coins have been manually labeled into valid and invalid coins. Valid coins are coins from 30 countries including most European countries, the USA, Canada and Japan. The portion of valid coins in the sample was 91.6%. The remaining 8.4% are dominated by coins from Asia, South-America, Africa and former socialist countries. Fig. 7 shows the distribution of fused probabilities $P(G_k|\Omega)$ for correctly classified valid coins as the solid line. Incorrectly classified valid coins are shown by the dashed line. The fused probability distribution for invalid coins is represented by the dotted line. Selection of threshold t on $P(G_k|\Omega)$ governs FAR, FCR and FRR, e.g., increasing t reduces FAR and FCR and increases FRR.

From a receiver operator characteristics (ROC) curve, as shown in Fig. 8, the tradeoff between FCR plus FAR and FRR can be identified. An operating point, corresponding to a specific t, is found on the ROC curve, e.g., for perfect classification with FAR + FCR ≈ 0, a very high FRR has to be taken into account (i.e., FRR > 0.5).

If the incorrect decisions FCR, FRR and FAR are equally weighted and we aim at minimization of the sum of false decisions FD = FCR + FAR + FAR. We find the optimum value for the rejection threshold t as the minimum of FD. This can be seen from Fig. 9, in which the minimum of FD is

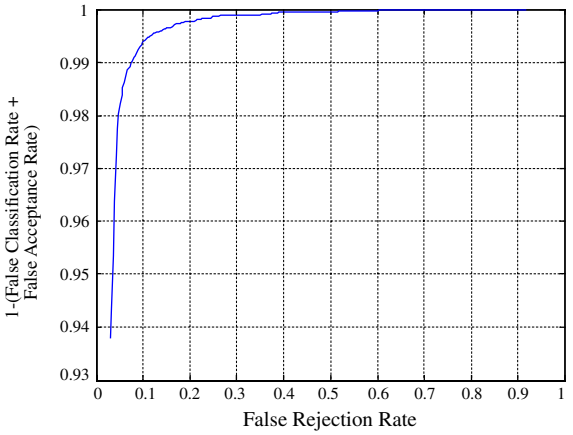

Fig. 8. Receiver operating characteristics.

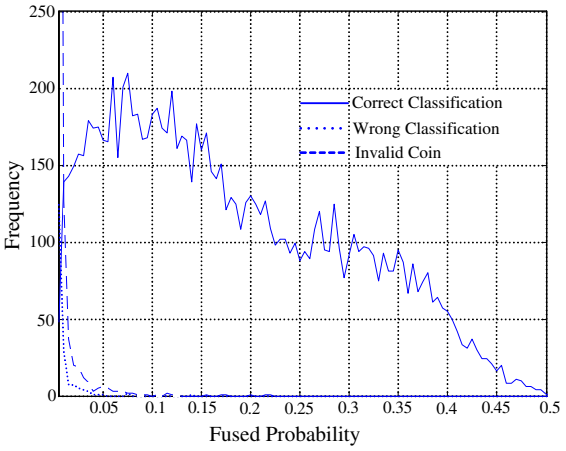

Fig. 7. Distribution of the fused probability with respect to classification result.

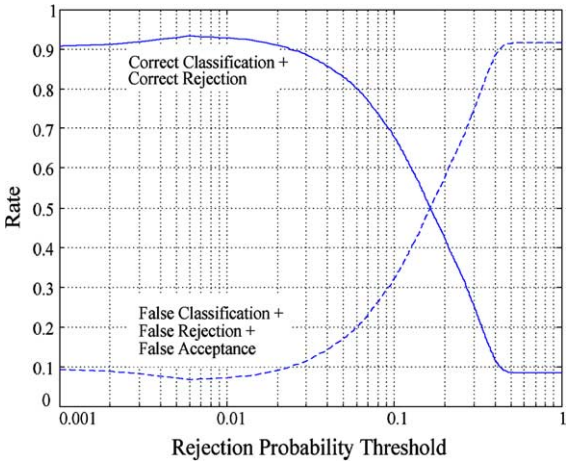

Fig. 9. Dependency of correct and false decision rates on rejection threshold.

Table 2
Summary of parameters

Total number of coins	11,949
Total number of coin side classes	932
Number of eigenspaces	494
Step sizes D_{step}, T_{step} in diameter/thickness plane in mm	0.3, 0.3
Tolerance parameter in diameter/thickness plane in mm	0.3
Minimum, average, maximum number of classes per eigenspace	2, 30.9, 113
Minimum, average, maximum number of coefficient vectors per eigenspace K	40, 618, 2260
Coin image size in pixels	128×128
Smoothed edge eigenspace parameter σ_s	1.5
Number of hypotheses H	5
Number of coefficients D	32
Parameter a in $P^2(\alpha_h^1, \alpha_j^2)$	0.08
Parameter σ in $P^2(\alpha_h^1, \alpha_j^2)$	15
Rejection threshold t	0.006

Table 3
Final results for a mixed sample of valid and invalid coins

	Acceptance		Rejection
Valid coin: 91.6%	Correct classification: 86.64%	False classification: 0.49%	False rejection: 4.52%
Invalid coin: 8.4%	False acceptance: 1.76%		Correct rejection: 6.59%

found for $t = 0.006$. At the same time correct decisions, i.e., correct classification and correct rejection rates, are maximized.

Considering only valid coins, i.e., the 91.6% coins included in the 30 countries mentioned above, and using no rejection mechanism, correct classification was made for 98.27% of valid coins, which is close to the practical optimum of 98.88% mentioned in Section 4.1. With rejection at the chosen level of $t = 0.006$, a percentage of correct classification of 94.54%, 0.53% false classification and 4.93% false rejection is achieved for valid coins.

Considering only invalid coins, i.e., the 8.4% coins not included in the 30 countries mentioned above, and rejection at the chosen level of $t = 0.006$ classification into any of the known coin classes happens for 20.47% of the unknown coins. Correct rejection of unknown coins is performed for 79.53% of invalid coins.

Examining at the mixed sample, a correct decision, i.e., correct classification or rejection, was made for 93.23% of all coins. False decisions, i.e., either false classification, false rejection or false acceptance, took place for 6.77% of all coins. Table 2 gives a summary of the parameters specific to the test sample and analysis method. Table 3 presents the final results.

The relatively high rate of false acceptance and false rejection could be reduced by considering coins from more than 30 countries. For example, for coins gathered in Austria adding of coins from the neighboring former socialist countries to the set of valid coins is advisable, as by visual analysis the majority of falsely accepted coins are of this origin.

6. Conclusion

We have presented a multistage approach to coin recognition applicable to coin collections comprising a large number of coin classes. Coins represent financial value only if the coins are sorted and returned to the respective national banks. A tunable system is required as national banks accept coins only if they are delivered with a high degree of purity. The rejection mechanism based on the probabilistic fusion result allows to adjust a tradeoff between rigorous classification (yielding high reliability against false acceptance but a higher rate of false rejections) versus tolerant

classification (yielding more false acceptances but fewer false rejections). Further improvement of results is expected through enlargement of the training set, which decreases the amount of falsely accepted coins.

Acknowledgment

The authors thank two anonymous reviewers for their valuable comments and suggestions which helped to improve the quality of the manuscript.

References

Altmann, J., Reitböck, H.J.B., 1984. A fast correlation method for scale- and translation-invariant pattern recognition. IEEE Trans. Pattern Anal. 6 (1), 46–57.

Ballard, D.H., Brown, C.M., 1982. Computer Vision. Prentice Hall, Englewood Cliffs, NJ.

Belhumeur, P.N., Hespanha, J.P., Kriegman, D.J., 1997. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. IEEE Trans. Pattern Anal. 19 (7), 711–720.

Bischof, H., Wildenauer, H., Leonardis, A., 2001. Illumination insensitive eigenspaces. In: Proceedings of International Conference on Computer Vision, pp. 233–238.

Burges, C.J.C., 1998. A tutorial on support vector machines for pattern recognition. Data Min. Knowl. Disc. 2 (2), 121–167.

Canny, J., 1986. A computational approach to edge detection. IEEE Trans. Pattern Anal. 8 (6), 679–698.

Davidsson, P., 1996. Coin classification using a novel technique for learning characteristic decision trees by controlling the degree of generalization. In: Proceedings of International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems.

Fisher, N.I., 1995. Statistical analysis of circular data. Cambridge University Press.

Fukumi, M., Omatu, S., Takeda, F., Kosaka, T., 1992. Rotation-invariant neural pattern recognition system with application to coin recognition. IEEE Trans. Neural Networ. 3 (2), 272–279.

Fürst, M., Kronreif, G., Wögerer, C., Rubik, M., Holländer, I., Penz, H., 2003. Development of a mechatronic device for high speed coin sorting. In: Proceedings of IEEE International Conference on Industrial Technology, Maribor, Slovenia.

Goh, K., Chang, E., Cheng, K.-T., 2001. SVM binary classifier ensembles for multi-class image classification. In: Proceedings of International Conference on Information and Knowledgement Management, Atlanta GA, USA, pp. 395–402.

Golfarelli, M., Maio, D., Maltoni, D., 1997. On the error-reject trade-off in biometric verification systems. IEEE Trans. Pattern Anal. 19 (7), 786–796.

Hibari, E., Arikawa, J., 2001. Coin discriminating apparatus. European Patent EP1077434.

Hu, M.-K., 1962. Visual pattern recognition by moment invariants. IRE Trans. Inform. Theor. 8, 179–187.

Kittler, J., Hatef, M., Duin, R.P.W., Matas, J., 1998. On combining classifiers. IEEE Trans. Pattern Anal. 20 (3), 226–239.

Kurita, T., Hotta, K., Mishima, T., 1998. Scale and rotation invariant recognition method using higher-order local autocorrelation features of log-polar images. In: Proceedings of Asian Conference on Computer Vision, volume II, pp. 89–96.

Leonardis, A., Bischof, H., Maver, J., 2002. Multiple eigenspaces. Pattern Recogn. 35 (11), 2613–2627.

Marr, D., Hildreth, E., 1980. Theory of edge detection. Proc. R. Soc. Lond. B 207, 187–217.

Murase, H., Nayar, S.K., 1994. Illumination planning for object recognition using parametric eigenspaces. IEEE Trans. Pattern Anal. 16 (12), 1219–1227.

Murase, H., Nayar, S.K., 1995. Visual learning and recognition. Int. J. Comput. Vision 14 (1), 5–24.

Neubarth, S.K., Gerrity, D., Waechter, M., Everhart, D., Phillips, A.C., 1998. Coin discrimination apparatus. Canadian Patent CA2426293.

Nölle, M., Penz, H., Rubik, M., Mayer, K., Holländer, I., Granec, R., 2003. Dagobert—a new coin recognition and sorting system. In: Proceedings of International Conference on Digital Image Computing—Techniques and Applications, Sydney, Australia.

Onodera, A., Sugata, M., 2002. Coin discrimination method and device. United States Patent US2002005329.

Pan, V.Y., Chen, Z.Q., 1999. The complexity of the matrix eigenproblem. In: Proceedings of the Annual ACM Symposium on Theory of Computing, pages 507–516, Atlanta, GA USA.

Pudil, P., Novovicova, J., Blaha, S., Kittler, J., 1992. Multistage pattern recognition with reject option. In: Proceedings of International Conference on Pattern Recognition, volume II, pp. 92–95.

Pentland, A., Moghaddam, B., Starner, T., 1994. View-based and modular eigenspaces for face recognition. In: Proceedings of Conference on Computer Vision and Pattern Recognition, Seattle WA.

Rothwell, C.A., Mundy, J.L., Hoffman, W., Nguyen, V.D., 1995. Driving vision by topology. In: Proceedings of International Symposium on Computer Vision, pp. 395–400.

Rupp, J., Prager, L., Rubik, M., 2003. Münzenwendestation. Austrian Patent Application A1514/2003.

Shah, G., Pester, A.L., Stern, C.M., 1986. Low power coin discrimination apparatus. Canadian patent CA1336782.

Sirovich, L., Kirby, M., 1987. Low-dimensional procedure for the characterization of human faces. J. Opt. Soc. Am. A 4, 519–524.

Swets, D.L., Weng, J., 1996. Using discriminant eigenfeatures for image retrieval. IEEE Trans. Pattern Anal. 18 (8), 831–836.

Torres-Mendez, L.A., Ruiz-Suarez, J.C., Sucar, L.E., Gomez, G., 2000. Translation, rotation, and scale-invariant object recognition. IEEE Trans. Syst. Man Cyb. 30 (1), 125–130.

Tsuji, K., Takahashi, M., 1997. Coin discriminating apparatus. European Patent EP0798669.

Turk, M., Pentland, A., 1991. Eigenfaces for recognition. J. Cogn. Neurosci. 3 (1), 71–86.

Uenohara, M., Kanade, T., 1997. Use of Fourier and Karhunen–Loeve decomposition for fast pattern matching with a large set of templates. IEEE Trans. Pattern Anal. 19 (8), 891–898.

Uenohara, M., Kanade, T., 1998. Optimal approximation of uniformly rotated images: Relationship between Karhunen–Loeve expansion and discrete cosine transform. IEEE Trans. Image Process. 7 (1), 116–119.

Venkatesh, B.S., Palanivel, S., Yegnanarayana, B., 2002. Face detection and recognition in an image sequence using eigenedginess. In: Proceedings of Indian Conference on Vision, Graphics and Image Processing, Ahmdebad, India.

Yilmaz, A., Gökmen, M., 2000. Eigenhills vs. eigenface and eigenedge. In: Proceedings of International Conference on Pattern Recognition, volume 2, pp. 827–830.