

Fast 3D-Vision System to Classify Metallic Coins by their Embossed Topography

Michael Hossfeld*, Weiyi Chu*, Markus Adameck⁺ Manfred, and Eich*

* *Hamburg University of Technology, Eissendorfer Strasse 38, D-21073 Hamburg, Germany*

⁺ *Hella KGaA Hueck & Co., Beckumer Strasse 130, D-59552 Lippstadt, Germany*

Received 13 September 2005; accepted 02 October 2006

Abstract

This paper presents a security-related machine-vision solution for real-time classification of moving objects with highly reflective metallic surfaces and complex 3D-structures. As an application example of our so called Three-Color Selective Stereo Gradient Method (Three-Color SSGM) a classification system for three main coin denominations of Euro coins is presented. Such coins are quickly moving in a coin validation system. The objective is to decide only from comparison of specially measured and processed 3D-surface information with characteristic topographical data stored in a database whether a coin belongs to one of the reference classes or has to be rejected as a foreign or counterfeit coin. Under illumination from a three-color light emitting diode equipped ring a single image of the moving coin is captured by a digital color camera. Exploiting the spectral properties of the illumination sources, which correspond to the special spectral characteristics of the camera, three independent subimages can be extracted. Comparison between these subimages leads to a discrimination between a coin with real 3D-surface and a counterfeit coin based on a photographic image of a coin of the same type. After the coin has been located and segmented, grey value based rotation and translation invariant features are extracted from a normalized image. In combination with template matching methods, a coin can be classified. Classification results will be reported for the three main coin denominations of Euro coins.

Key Words: Object Recognition, Structural Pattern Analysis, Specular Metallic Surfaces, Three-Color Selective Stereo Gradient Method (Three-Color SSGM)

1 Introduction

Modern coin validators measure precisely different coin features of which electromagnetic properties of the coin materials are the most important nowadays. On January 1st, 2002, Euro bank notes and coins began circulating within the European Monetary Union. At the same time the coin validator's task to verify genuine Euro coins and to reject coins of foreign currencies or counterfeit coins became more difficult for various reasons.

Each of the 12 members of the Monetary Union (before May 1st, 2004) embosses its own sets of eight different Euro coins with an individual minting pattern on the national sides. Several European manufacturers produce the metal alloys for the coin blanks which are embossed in 15 different European mints. Despite of

Correspondence to: <m.eich@tu-harburg.de>

Recommended for acceptance by <Walter F. Bischof>

ELCVIA ISSN:1577-5097

Published by Computer Vision Center / Universitat Autònoma de Barcelona, Barcelona, Spain

a quality control system implemented by the European Central Bank for all aspects of the Euro coin production, the large number of members and manufacturers increase the range of variation in the coin's properties, especially of the electric conductivity of the coin material.

To compensate this effect, the intervals for permitted measurements of coin properties in coin validators have to be widened compared to the previous single national coin detection scheme. A loss of identification quality and a slightly higher probability to falsely accept foreign or counterfeit coins would be the consequence. This scenario, therefore, has to be avoided especially because of the newly observed strong increase of counterfeit Euro coins in some countries [1].

Only coin properties not yet acquired in coin validators and thus used as additional information to the existing electromagnetic properties guarantees a more effective rejection of foreign and counterfeit coins. In addition, we expect the prospect of improved recognition quota of genuine Euros.

To identify such coin properties a wide range of possible sensor technology not yet used in modern coin validators had been intensively tested for applicability. It soon turned out that the obvious visual 3D-appearance of the coined pattern *as a whole* bore high discriminative potential that had not yet been exploited for classification purposes in coin validators today. Verifying as well a coin's visual structural information excludes effectively attempts of coin validator fraud with specially designed metallic tokens whose only purpose is to simulate the electromagnetic properties of a genuine coin in order to trigger an acceptance signal. Their optical appearance is insignificant.

Against the background of the presented facts we have developed the so called Three-Color Selective Stereo Gradient Method (Three-Color SSGM)[2]-[4]. It is in general a non-destructive and contact-free optical inspection system for real-time classification of moving objects with highly reflective metallic surfaces and complex structures.

A Three-Color SSGM-system which classifies minted patterns of moving Euro coins into classes corresponding to their country of origin serves as an application example. The rotation and translation invariant classification is based solely on 3D-topographical information of the embossed profile. 3D-information is not directly extracted. Instead image characteristics resulting from the fact that a coin has 3D-contour are analyzed to determine the coin class. Characteristic intervals for acceptable measurement values are derived from a training set. As a proof of the Three-Color SSGM classification principle, hardware and image processing algorithms of a PC-based laboratory demonstrator are described, for which the classification performance is presented and discussed. We concentrate our development primarily on three basic coin denominations 2 Euro, 1 Euro, and 50 Cent, but the method can easily be extended to other face values. It will be shown that the coin diameter can be determined with an accuracy of less than 0.1 mm (see chapter 4.2). As an important system property the average false rejection rate is determined to be less than 0.6% (see chapter 6). This "first coin insertion" result is much better than in many existing electromagnetic based coin validation systems. With a Three-Color SSGM-system installed in existing coin validators in addition to the electromagnetic sensors the robustness of such machines toward coin recognition and coin rejection can be improved.

An optically based visual pattern recognition system can, in principle, be deceived by photographic images of the objects to be classified. This aspect is normally neglected in industrial inspection tasks. But in the realm of security systems and anticounterfeiting tasks a classification system has to cope with this kind of trivial attempts of fraud. Therefore a 3D-verification algorithm has been developed which is as well described.

2 General Concept of the Three-Color Selective Stereo Gradient Method

The surface of coins exhibits a fine structure resulting from both embossing and wear and tear. Under illumination this structure reflects light characteristically towards an observer. Intensity values in a digital image of a coin carry information about the three-dimensional local shape of the embossing's structure. To exploit this information for classification purposes it is necessary to understand what governs the observed intensity. In this section we briefly address light reflection from diffuse and specular surfaces.

2.1 Determining Surface Orientation

If the elevation $h = h(x, y)$ of a surface above the $x - y$ -plane in a Cartesian coordinate system is given, the first partial derivatives of h with respect to x and y will be defined as

$$p(x, y) = \frac{\partial h(x, y)}{\partial x} \qquad q(x, y) = \frac{\partial h(x, y)}{\partial y} \qquad (1)$$

The vectors $(1, 0, p)$ and $(0, 1, q)$ are tangent to the surface in point (x, y) . The normal vector of the surface in this point is the cross-product of these two, resulting in $[-p, -q, 1]$. By convention the viewer looks into the $(0, 0, -1)$ -direction of the negative z -axis, so usually the opposite normal vector $[p, q, -1]$ is considered. The quantity (p, q) is called gradient of $h(x, y)$. The (p, q) -space is a nonlinear representation of the surface orientation.

A reflectance map $R(p, q)$ is a convenient model to describe the image intensities as a function of the surface orientation (p, q) . It takes into account a fixed object illumination, the surface reflectance of an object's material given a particular light source in form of an explicit mathematical description, and the imaging geometry. We assume, that the image projection is orthographic and that the incident illumination comes from a single distant point source. Even then the calculation of the surface orientation from the image intensity is difficult, because a non-linear first-order partial differential equation has to be solved for p and q . Since there is only one equation at hand, the system is underdetermined [5]-[6].

Until today, photometric stereo methods are well-known techniques for extracting surface normals [7]. The photometric stereo method according to Woodham [8] uses reflectance maps to solve the problem of finding the surface orientation. The idea is to vary the direction of incident illumination between successive views of the same object, while holding the viewing direction constant. If it is known how light is reflected at a surface, $R(p, q)$ can be calculated. Woodham showed for Lambertian surfaces that three views with their corresponding $R(p, q)$ are sufficient to uniquely determine both surface orientation and reflectance factor at each image point.

2.2 Reflectance Models

Besides rendering the correct 3D-shape of the object to create realistic images in computer graphics, a reflectance model is needed for each illuminated reflective surface in the scene. For modeling object shapes there are a number of known techniques such as range image merging, photometric stereo or shape from shading [9][10]. Some techniques are capable of modeling both shape and reflectance properties [11]-[13].

The most accurate way to model the reflectance mechanism alone is given by Maxwell's equations since they describe the physical nature of the interaction between electromagnetic waves and matter. Although there are reflectance models which are based on solving Maxwell's equations [14] such a direct approach is usually computationally very extensive.

The bi-directional reflectance distribution function (BRDF), defined by the National Bureau of Standards, formally describes the reflectance mechanism [15]. It is the angular dependent ratio of the reflected radiance dL_r (within solid angle $d\omega_r$) in direction towards the viewer to the incident irradiance dE_i (within solid angle $d\omega_i$) in direction of an infinitesimally small part of the surface's area:

$$f_r(\theta_r, \phi_r; \theta_i, \phi_i) = \frac{dL_r(\theta_r, \phi_r; \theta_i, \phi_i, E_i)}{dE_i(\theta_i, \phi_i)} \qquad (2)$$

where the polar angles θ and the azimuth angles ϕ together indicate a direction. The subscripts i and r denote incident and reflected radiant flux (see Fig. 1). Equation 2 assumes monochromatic light. If a BRDF term can be explicitly given for a surface, its brightness or shading can be accurately calculated in synthesized images.

BRDF data can be achieved in two ways: gonireflectometers can measure the BRDF for each object of a scene. Exhaustive measurements of the entire hemisphere is required here because image rendering algorithms must evaluate the BRDF in arbitrary directions [16] such that this method is not suitable for practical purposes.




Figure 1: Nomenclature for the definition of geometric quantities in the BRDF (adapted from [15]).

A second way is to assume a certain reflectance behaviour for the surfaces. A very simple and widely used model is the Lambertian reflection under which an object surface appears equally bright from all viewing directions [17]. The Lambertian BRDF can lead to satisfactory approximations of surface appearances but usually the surfaces appear somewhat unreal because the model neglects specular reflection components. Non-Lambertian parametric models for real-world objects can be divided into physically-based and empirically-based models [18]. Although they lead to highly accurate surface rendering, physically-based models like the Torrance and Sparrow reflection model [19] are still used only occasionally because of their complexity and because parameters are not readily available for real time computing [16].

The best known empirical model is the Phong reflectance model [20] because of its simplicity and its good surface rendering results. It describes the shading at a surface point as a superposition of Lambertian and specular behaviour.

$$L_{r,Phong} = \rho[\cos(\theta_i)(1 - d) + d] + W(\theta_i)[\cos(\gamma)]^n \quad (3)$$

where ρ is the reflection coefficient of the object's surface, d the environmental diffuse reflection coefficient, $W(\theta_i)$ the ratio of specular reflected light, γ the angle between direction of the reflected light and the direction of the viewer. The exponent n is a power which models the specularly reflected part of the incident light and depends on the material.

2.3 Reflection at specular metal surfaces

Reflection at metal surfaces differs significantly from diffuse reflection. Most notably it is dominantly specular, strongly forward oriented into a narrow solid angle, and therefore does not produce smooth shading on the surface. For our Three-Color Selective Stereo Gradient Method we choose the Phong reflectance model to describe this characteristic reflectance because it approximates the angle dependent reflection of metallic surfaces very well as we will show experimentally (see Fig. 2).

A 1-Euro-coin is illuminated under a fixed incidence angle $\theta_i = 45^\circ$ by a widened and collimated beam of a laser. The average surface's normal and the vectors of incident and reflected direction are coplanar. An image sensor is moved around the reflected maximum in discrete angle steps. At every position a grey-value image of the coin is taken. The camera's optical axis lies within the plane of the three vectors.

In all images the mean grey value of the same plane surface patch of $5 \times 5 \text{ mm}^2$ is measured. As expected, the maximum at $\gamma = 0$ of the measured reflection is found where the reflection angle θ_r equals θ_i . The angle

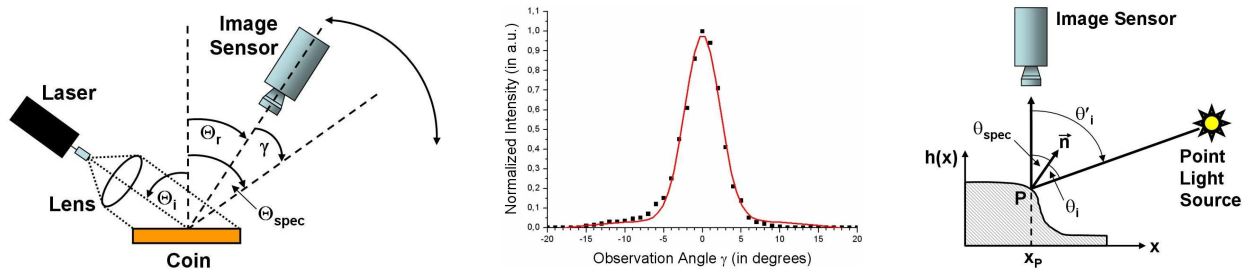


Figure 2: Reflection at specular metal surfaces. Left: Experimental setup for measuring the reflectivity property of a metallic coin. Center: The reflected specular intensity can be approximated with the Phong shading Model. Dark dots: Measurement; Line: Model fit. Right: Cut through a profile contour.

θ_r , which fulfills the reflection condition at the point of incidence on the surface, is named here as θ_{spec} . The measured data, indicated as dots in Fig. 2, middle, can be approximated by the predominant $\cos(\gamma)^n$ -term of the Phong shading model according to equ. 3, which yields $n > 600$. The extreme high value of n emphasizes the highly specular property of the metallic surface.

2.4 Modelling the embossed profile and selecting a gradient

In our model the three-dimensional nature of the embossed metal structure is described by an infinite number of infinitesimally small surface elements all tangential to the 3D-structure relief across the coin surface. By determining gradients we are able to both identify the embossed structure for classification and investigate the nature of the relief to sort out photos.

Determining surface orientations using Woodham's photometric stereo method [8] would be possible here, but requires demanding lighting with a great number of point light sources, illuminating the scene from different directions. Many images taken from one scene view by activating the light sources one at a time have to be analyzed to calculate all gradients of the objects. Due to this drawback, this method appears often to be very time consuming and not suitable for real-time applications*.

Still, Woodhams approach can be used here to classify coin embossings, if reasonable simplifying assumptions can be found: since coins are flat, thin, mostly circular objects, it is convenient to observe them by an image sensor from top view and to assume that image projection is orthographic. To get hold of the coined structure it is not necessary to determine *all* gradients, not even numerically. Under the assumption of a uniform albedo of the coin's surface, partial knowledge about one gradient value suffices for a safe classification.

This is explained in the schematic sketch of Fig. 2, right. A metallic ridge of the profile of a coin, illuminated by a distant point light source, is cut perpendicularly to the coin surface and drawn as seen from the side. It is observed by an image sensor from the top view. The azimuth angle ϕ_i of the incident illumination is fixed, as well as the camera, whose image sensor lies parallel to the coin surface.

For an arbitrary but then fixed polar angle θ'_i of incidence light measured now in a viewer-oriented coordinate system against the main surface normal of the coin, only one bright spot appears in the camera. The light which causes this spot, originates from that point P on the ridge, where the normal vector of its tangential small surface element bisects θ'_i . From all points on the profile this occurs for this and only for this surface element since only there the reflection condition is *locally* fulfilled for a given viewing and observation direction:

$$\theta_i = \theta_r = \theta_{spec} = \frac{\theta'_i}{2} \quad (4)$$

This results as a consequence of the strongly forward oriented specular metallic reflection.

*see for example [21], in which an inspection system for specular solder joints of surface mounted components is described. A hemisphere of 127 point light sources are used to extract all local surface orientations of each solder joint.

Simultaneously, the camera detects one gradient (p', q') with absolute value $|(p', q')|$ graphically. For every point on the ridge profile there is ideally only one θ'_i by which it appears as bright spot on the image sensor. Therefore, a minted edge appears as a single bright line as shown in Fig. 3, left. One predetermined absolute gradient value $|(p', q')|$ of interest can unambiguously be selected by choosing a special θ'_i using one single light source instead of three different light sources according to Woodham's photometric stereo if the albedo varies less over the observed area. This is the principle of the SSGM.




Figure 3: Left: A minted edge appears as a single bright line representing gradients with absolute value $|(p', q')|$. Center and right: Sometimes more than one point on the ridge can locally fulfill the reflection condition.

Depending on the individual form of a coin's profile and the angle θ'_i chosen two or even more points on the ridge can fulfil locally the reflection condition. This is shown schematically in Fig. 3, center. Since the gradient images are taken from a distant image sensor the separate two lines indicating the same absolute value $|(p', q')|$ are very close together and nearly merge into one single line. This effect has negligible influence on the recognition accuracy.

If *all* gradients with absolute value $|(p', q')|$ can be captured over the coin's surface the image sensor receives a representation of the embossed structure.

2.5 Development of a suitable illumination

A problem-adapted tailored illumination is half the solution of nearly every vision system. The demands on our illumination here are extensive. To detect all points with gradient value $|(p', q')|$, an azimuthally 360° -illumination ring with constant θ'_i for all illumination directions is necessary. Each source has to illuminate the object plane homogeneously to ensure rotation invariance of the features.

In chapter 4.5 we will show how the 3D-property of the embossing can be verified. The verification module is fundamentally based on multiple images showing the same coin profile from different illumination directions. To avoid stopping the coin in its motion for taking these images and to limit the number of sensors to one, we resorted to a scheme which encodes different illumination directions by colors.

We achieve good results with three different illumination directions, encoded by red, green, and blue light sources and arranged circularly in three 120° -ring-shaped sectors. Now, only one color image is necessary, which can be taken while the coin is in motion. The software module separates the color image into three independent subimages (see chapter 4.1), bearing disjunct information about the coin profile from different ranges of azimuth angles. Calculating difference images from pairs of these subimages allows conclusions to be drawn about the 3D-nature of the coin.

The difference image method requires that each detected gradient originates only from one single light source. This is important at the border of two adjacent colors. If a profile segment is simultaneously illuminated by two colors, the information about the profile contour in this segment will vanish in the difference image.

To ensure this feature, the illumination ring has to be composed of individual light sources. From the graph shown in Fig. 2 can be inferred that the reflected intensity of a metallic surface even from a point light source

deviates substantially (intensity $\geq 5\%$) from zero only between an angle of $\pm 8^\circ$. This indicates a separation of at least 16° azimuthally between each light source, making the 24° between each of 15 equally spaced sources an adequate azimuthal distance. Then the border condition is fulfilled.




Figure 4: Grey value and thresholded images of two top views on perfectly reflecting spheres with radius one. Its shadings are calculated according to the Phong reflection model.

This requirement seems to be inconsistent with the above mentioned requirement for a rotation invariant continuous illumination. But it can be shown that a continuous detection of a gradient can be built up with individual light sources as well. Fig. 4 shows 8-bit-grey value images of two top views on perfectly reflecting metallic spheres with radius one. Its shadings are calculated[†] according to the reflection model of Phong (equ. 3). The first sphere is illuminated by 5, the second by 15 point light sources, all emitting at the same polar angle. The geometric distances of the model correspond to the situation in the illumination ring of Fig. 5. The right part of Fig. 4 shows the same grey value images now thresholded by grey value 229. It can be seen that 15 individual light sources are enough to detect continuously one single gradient value.

3 Experimental Setup

In order to prove and to evaluate the working and recognition principle of the Three-Color SSGM, a laboratory demonstrator is set up, based on moving coins rolling down an inclined plane.

The illumination ring with inner diameter of 30 mm consists of 15 color LEDs[‡] separated azimuthally 24° and arranged in three 120° -sectors of 5 monochromatic LEDs circularly around the field of view (see Fig. 5, left). The polar angle of $\theta'_i = 63^\circ$ is the same for all LEDs. It is chosen as a compromise between achieving the brightest profile contours from the middle part of the edge curvature and the darkest background.




Figure 5: Left: Schematic top view on the illumination ring. 15 color LEDs are arranged in 3 sectors around the coin centered in the middle. Right: Side view on the experimental setup of the 3-Color-SSGM.

[†]with MATLAB[™]: The Mathworks, Inc., 3 Apple Hill Drive, Natick, MA, USA; Internet: www.mathworks.com

[‡]light-emitting diodes from Infineon Technologies AG; Red: LA E67B; Green: LT E67C; Blue: LB E67C

LED irradiance is always inhomogeneous (see Fig. 6, A). To compensate this effect, we use planoconvex plastic lenses in a special setup. By properly positioning the LED within the focal length of the lens and tilting the optical axis of lens and LED against each other, a very good illumination homogeneity of the object plane of 30 mm in diameter can be achieved even with LEDs (see Fig. 6, B and C). For image acquisition we use




Figure 6: Inhomogeneous illumination of a single LED: top view on a 90° -sideways illuminated surface of 30 mm in diameter without any lens. Detected radiant light along the arrow line as grey values plotted in the graph (A). Compensation of inhomogeneous illumination using a sideways positioned lens (B and C).

a digital, progressive, color CCD camera with 1024×768 pixels. The camera comes with an 1/3 inch sensor ICX204AL from Sony with a Bayer-Pattern and is connected to a digital frame grabber from EuresysTM.

The LEDs are chosen in such a way that their spectral emission characteristics meet the spectral selection function of the Bayer-Pattern mosaic filter in front of the image sensor (shown in Fig. 7).




Figure 7: Spectral transmission properties of the Bayer-Pattern mosaic filter elements (solid lines, abbreviated B.-P.) and the spectral emission characteristics of the three sorts of LEDs (dashed lines).

When a coin rolls down the guideway of the inclined plane into the recognition system, it passes a photoelectric IR-barrier in the middle of the ring. After 300 ms all 15 LEDs flash for 100 ms, the camera captures one single 8-bit grey value image, which is sent to the PC for further analysis.

4 Image Processing and Classification

We refer to the three main coin denominations 2 Euro, 1 Euro and 50 Cent simply as "coin sort" and to the countries of origin within each sort as "classes". The objective of the image processing algorithm is to decide from standardized information extracted from an image of the 3D-profile contour of the coin's minting with absolute gradient value $|(p', q')|$ to which coin sort and class a specimen belongs and whether the specimen is a photographic image or not. The image processing algorithm follows the schematic diagram shown in Fig. 8. Its individual moduls are described in the following sections.

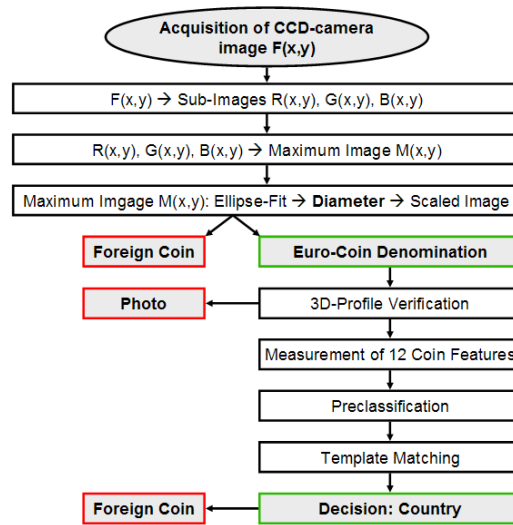


Figure 8: Schematic flow diagram illustrating the most important steps of the image processing algorithm, implemented to classify coins according to the Three-Color SSGM.

4.1 Data Preprocessing

After grabbing a coin image $F(x, y)$, three grey value subimages $R(x, y)$, $G(x, y)$, and $B(x, y)$ of the red, green, and blue illumination sector are extracted from $F(x, y)$. Because there are twice as many green filter elements as compared to blue or red ones, every second green Pixel is skipped in order to yield three subimages of same image dimension. The original image size of 1024×768 pixels is reduced to 511×383 pixels. A maximum image $M(x, y)$ is calculated according to

$$M(x, y) = \max[R(x, y); G(x, y); B(x, y)] \quad \forall (x, y) \in M \quad (5)$$

in which for each point (x, y) in the maximum image M the highest intensity in the set of the three pixels $R(x, y)$ or $G(x, y)$ or $B(x, y)$ is chosen. Since the background in each of the three images is mostly dark, this operation leads here to a contrast improvement. Original and maximum image are shown in Fig. 9. All coin images in this article are shown inverted because much more details can be visually perceived.

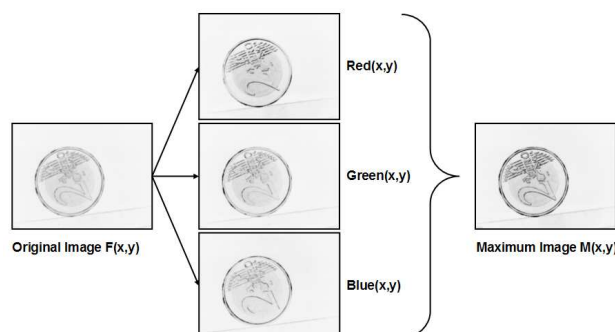


Figure 9: Inverted original image $F(x, y)$, three grey value subimages $R(x, y)$, $G(x, y)$, and $B(x, y)$ and maximum image $M(x, y)$ of the international side of a 2-Euro-coin.

4.2 Segmentation of the coin by an ellipse-specific point fit

In order to segment the coin, its edges have to be located. To achieve this objective a suitable number of straight lines parallel to the horizontal x-axis of the maximum image are chosen, such that they run through the middle part of the coin body and intersect its edge on each side. In analyzing the grey value profile taken along these straight lines from both ends toward its midpoints, edge coordinates are determined. This procedure leads to a set of points all lying on the coin's outer rim. In our case, a set of 32 parallel straight lines resulting in 64 edge points has proven to deliver satisfactory results.

The extracted edge points located on an arc are fit to the equation of an ellipse enclosing and defining the coin area in the image. Since the inclined plane is slightly tilted backwards to hold the coins on the narrow guideway, the coins appear as ellipses in the images[§]. Instead of using iterative least-square-methods, we have implemented a direct ellipse-specific algorithm introduced in 1998 by Halir and Flusser [22]. The main advantages of this method are that the geometric parameters for the ellipse are given in only *one* step and that the algorithm can be fed with an arbitrary number of edge points. This method turns out to be very robust and accurate. The diameter of coins belonging to one sort can be determined with an accuracy of less than 0.1 mm.

After the coin's region of interest (ROI) is determined as the result of the ellipse fitting, the coin area is copied into a new image. It always has the size of 256×256 pixels for *all* coins, independently of the coin's real object size. This is important for analyzing the grey value based features which we use for classification. Such a standardized image is called "scaled image".

4.3 Coin Features

In the scaled image characteristic circular regions are defined (see Fig. 10), for which the values for characteristic features are measured. The radii for the border circles of each region are fixed and used for each coin sort.




Figure 10: The inverted scaled image (left) is separated into three characteristic circular regions in order to extract classification relevant information.

From each of these regions two grey value based features are calculated: mean grey value and standard deviation of grey values. We use grey value based statistical features because they are rotation invariant by nature and can be determined very fast. Essential preconditions are homogeneous illumination and a calibration to achieve identical grey values from the coin profile no matter what the rotation position of the coin with respect to the individual illumination sector upon exposure.

A measure for the distribution of embossed structures in the binarized thin outer ring $F_{out,bin}$ is the distance D of the geometric image center (always point (127, 127)) to the center of gravity (X_{cog}, Y_{cog}) . This valuable additional feature is calculated according to

$$X_{cog} = \frac{\sum x \cdot F_{out,bin}(x, y)}{\sum F_{out,bin}(x, y)} \quad Y_{cog} = \frac{\sum y \cdot F_{out,bin}(x, y)}{\sum F_{out,bin}(x, y)} \quad (6)$$

$$D = \sqrt{(127 - X_{cog})^2 + (127 - Y_{cog})^2} \quad (7)$$

[§]The average numerical excentricity resulting from our setup is only 0,09.

This feature alone is capable of separating the international, Spanish, Luxembourgian, Greek, and Dutch class from all the other 2-Euro-classes.

Further five features are derived from one-dimensional grey value profiles, taken along five circles around the coin's center in the scaled image (see Fig. 11). The radii of the circles are fixed for all coins and chosen in such a way that they run through structurally characteristic regions of most of the 2-Euro-mintings. For each grey value profile, the dominant spatial frequency of its Fourier transform is determined denoting a characteristic periodicity.

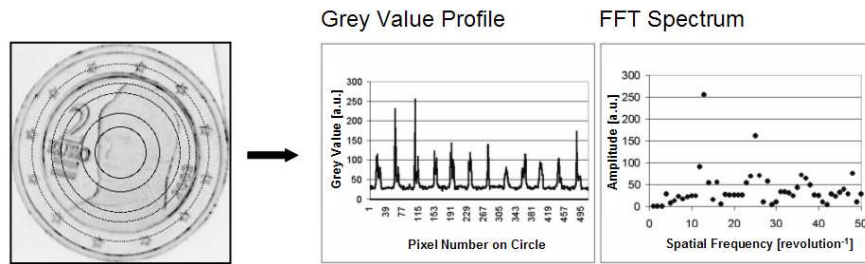


Figure 11: Left: Inverted scaled image of an Austrian 2-Euro. Grey value profiles are taken along 5 circles around the image center in point (127,127). Center: Plot of grey values taken along the outer circle which goes through the 12 stars. Right: Plot of spatial frequencies calculated according to a fast Fourier transform. The dominant frequency observed on the outer circle is 12 / revolution.

4.4 Template Matching

The last step in identifying a coin's structure consists of a template matching. For each class within each sort characteristic templates must be stored in a database (typically of size 300×100 pixels with a resolution of 0.2 mm per pixel). Template matching in scaled images is possible, but computational intensive because it requires to store each template in many possible rotational positions of the coin in front of the camera. Even if 360 different templates per class belonging to an angle-discretization of 1° were stored, the match value would be statistically rather unreliable. Since the match value strongly depends on the degree of structural congruence between the two pattern, the maximum achievable match value and therefore the best discrimination capability of the classifier could potentially be obtained only between two stored rotational positions thus leading to a false rejection.




Figure 12: Inverted scaled image of the international side of a 2-Euro-coin (left) and its unwarped image (right).

Instead of this, the grey values of the scaled image are transformed into a new image. To illustrate the transform process, a rotating "radius beam" of length $R = 127$ pixels originating in the center of the scaled image and ending at the coin edge can be envisaged. At each angular position Ψ of the beam relative to its starting orientation all grey values $I(x, y)$ are taken for (x, y) pairs given by the azimuthal orientation of the "radius beam" and transformed into a new $\Psi - I(r)$ -coordinate system. We call this image "unwarped image".

Fig. 12 shows an example of a 2-Euro-coin. In an unwarped image the structure is increasingly distorted from the rim to the coin center. However, it is rotationally invariant and thus much more efficient for a template matching routine, since a rotation of the coin results solely in a horizontal translation of the profile pattern. Therefore only one reference template taken from an unwarped image has to be stored in just one orientation.

The template searched for can be cut off at the left or right border of the unwarped image. To ensure that always the complete template lies within the unwarped image the radius beam rotates through an angle of 4π .

4.5 Verification of the 3D-profile

Each visual recognition system in which the classification is based on images of the objects to be recognized can be deceived by a photographic image of the objects (see Fig. 13 for an example).



Figure 13: Digital image of a Belgian 2-Euro-coin (left) and a photographic image of the same coin (right). Both can not easily visually be distinguished from each other.

The discrimination between a real object with 3D-topography and a photographic image of the object is achieved by difference images which are calculated from the three subimages R , G and B as:

$$D_{RG}(x, y) = |R(x, y) - G(x, y)| \quad (8)$$

$$D_{RB}(x, y) = |R(x, y) - B(x, y)| \quad (9)$$

$$D_{GB}(x, y) = |G(x, y) - B(x, y)| \quad (10)$$

The implemented verification strategy seeks to confirm the existence of 3D-edges. Here, all difference images are thresholded and added to a "sector-sum-image" $S(x, y)$ (see Fig. 14). The maximum image $M(x, y)$ is as well thresholded resulting in image $T(x, y)$. A threshold calculated from the sum of the mean grey value in the image and the standard deviation of grey values is a suitable automatically calculable threshold.

In the case of a 3D-profile as demonstrated in Fig. 14 the resulting image calculated as the difference between $S(x, y)$ and $T(x, y)$ is nearly empty. This is a direct consequence of the *disjunct* profile information in the three difference images, whose image sum $S(x, y)$ is identical to the binarized maximum image $T(x, y)$.

In the case of a photographic image, all subimages are identical, because photos are plane. The image content in the difference image between two subimages is zero, therefore $S(x, y)$ is empty. But neither the structure bearing the maximum image $M(x, y)$ nor its thresholded image $T(x, y)$ is empty, consequently the difference image between $S(x, y)$ and $T(x, y)$ now is *not* empty. Here again it becomes clear why homogeneous illumination is crucial. Otherwise, $S(x, y)$ would as well be not completely empty for a photo. In practice, due to slight production differences among the LEDs and tiny differences in positioning the LEDs into the ring results in small illumination inhomogeneities such that not all image contents vanishes in $S(x, y)$. As can be seen from the resulting images in Fig. 15, coins with 3D-profile and photographic images of such coins can be clearly distinguished by the mean grey value.

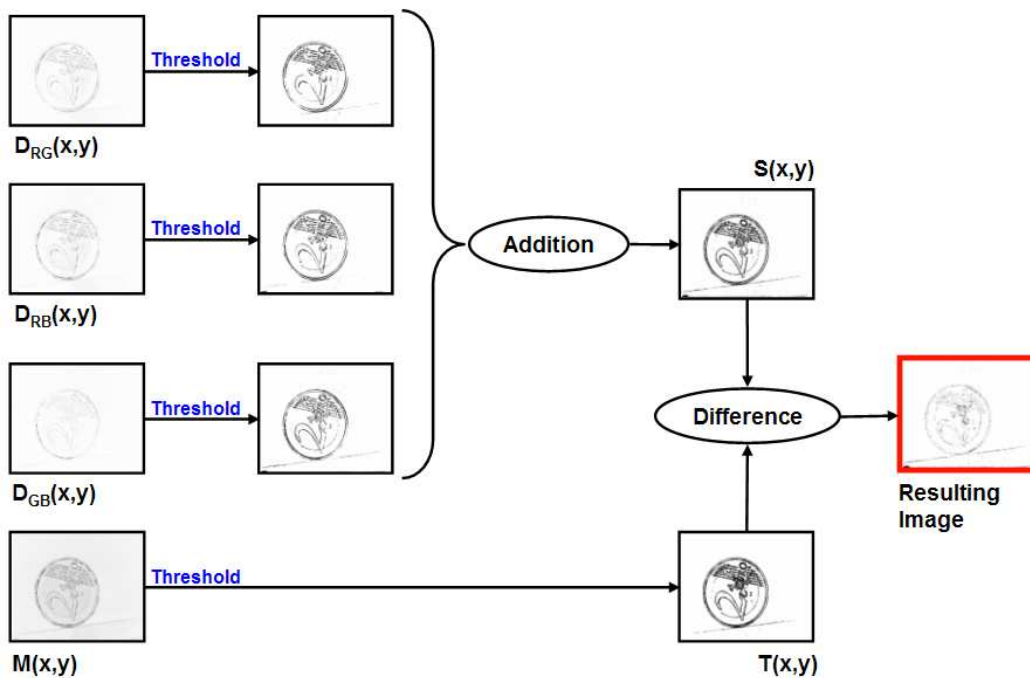


Figure 14: Schematic diagram of the 3D-verification module for a coin.




Figure 15: Inverted resulting images of the 3D-verification module for a coin (left) and a photographic image of a coin (right). They can be clearly distinguished in the algorithm by their mean grey values which yields here 3 for the coin and 22 for the photographic image.

5 Classification strategy

Following the flow diagram of Fig. 8 which illustrates the most important steps of the Three-Color SSGM coin classification algorithm, the coin diameter resulting from the ellipse fitting module determines to which sort a test coin belongs. If this value doesn't lie within one of the predetermined acceptance intervals for possible Euro coin denominations, a coin specimen will be treated as a false coin and rejected. In addition, only coin specimen with positively verified 3D-profiles enter the classification module.

For all 12 coin-describing features (seven features from the three ring regions and five dominant spatial frequencies, see section 4.3) for each class of each sort individual intervals of acceptance are defined. Measured values falling into these intervals are regarded as originating from "true" coins, otherwise from "false" coins. We call the set of all intervals of acceptance for all sorts and all classes "preclassification filter". To assign one of 13 possible class labels[¶] to a coin specimen the classification strategy follows a two-step cascading elimination strategy, by which in the first (the preclassification) step the number of 13 possible classes for a coin sort is reduced as far as possible.

[¶]Consisting of the names of the 12 countries of the members of the European Monetary Union plus one for the international side. Coins of the Vatican, the principality of Monaco and San Marino are neglected here.

For all countries of the beforehand determined coin sort, the 12 measured feature values are compared with the 12 intervals of acceptance. If for a current country only one feature value lies outside its allowed interval, this country will be not regarded a possible candidate in the decision process.

The final decision is achieved by a template matching as the second step. Due to the preclassification step, the computationally intensive template matching is carried out only for a small number of remaining classification candidates.

We have implemented the algorithm in a C++-based programm using the eVision Image Processing ToolboxTM^{||}. The time needed to classify a coin depends strongly on the hardware used. On a Pentium 4 - PC with 3 GHz and 1 GB RAM a complete coin classification process takes 0.21 s starting with loading the original coin image and ending with the classification decision. This is the reason why we regard our Three-Color SSGM a fast classification system.

6 Experimental Results and Discussion

In order to determine the intervals of permitted feature values for each sort and each class a training set of (arbitrary) 140 different coins per class is processed by the laboratory demonstrator as described. For each coin its original image $F(x,y)$ is stored for future verification purposes. This results in 5460 coin images. The distributions of coin diameters are determined to identify the diameter intervals for each sort. All coins of the training set have been completely and correctly separated into sorts. The standard deviation has been 0,08 mm for each coin sort.

For each class within each coin sort maximum and minimum values are determined from the training set for each feature and registered in the preclassification filter. It turns out that the preclassification is capable of reducing the initial number of 13 possible classes to at maximum only two with over 62% probability.

The human judgement is not perfect in finding the optimal template. Therefore, we designed an algorithm that allows to find class-specific templates automatically from visually well contrasted, unblurred and structurally most complete unwarped images among the set of all unwarped images of the training set. Given the minimum and maximum sizes of a template this programm tries out every possible cutout in the unwarped image and matches it with every available unwarped image of the own class *and all* other classes. The overlap interval between match values of the own class and match values of all other classes is determined. That template giving the highest match value with unwarped images from the own class and the smallest possible overlap interval is finally taken for the laboratory demonstrator. Simultaneously, the minimum match value to be achieved for recognizing a coin of the own class is determined for all classes and stored in the preclassification filter.

To test the discriminating power after the classifier training a set of 30 different foreign coins are processed from both coin sides resulting in 60 different images. The coin diameters are similar to the diameters of the three Euro coin sorts. A part of this set consists of coins with two-colored appearance which resemble the 2-Euro- and 1-Euro-coins. To evaluate the 3D-verification module high-quality photographic reproductions of all 2-Euro-classes are produced from one arbitrary chosen coin. From each reproduction 60 images are generated.

Table 1 summarizes the classification results for the test sets. An average false rejection rate (FRR) of 0,4% is achieved. This is an excellent result, taking into account that in the industrial field of coin validation systems, a FRR of 95% for "first coin insertion" in existing, only electromagnetic based, industrial coin validation systems is a common standard value. A coin has been classified as a Euro coin if no error occurred in the determination of its country of origin. For only ten Euro coins no name of a EU-country could be found and the coins were falsely rejected.

The classification of foreign coins and the performance of the 3D-verification module is summarized in table 2. It indicates that all coins of the test set are correctly classified as coins, all photos of 2-Euro-coins are correctly identified as photos and no foreign coin is classified as genuine Euro. This yields a false acceptance

^{||}Euresys s.a., Avenue du Pre-Aily, 14, B-4031 Angleur, Belgium; Internet: www.euresys.com

Euro coin sort	Classification-error	Errors of Preclassification-Module:		Errors of the T.M.-Module	No. of coins falsely rejected	FRR in %
		nothing preclassified	own country of origin not preclassified			
2 Euro	0	2	1	0	3	0,4
1 Euro	0	0	2	3	5	0,6
50 Cent	0	0	2	0	2	0,3

Table 1: Classification result for 2340 genuine Euro coins of the test set (60 different coins per class).

rate (FAR) of 0%. The embossing of one of the foreign coins was so flat that it has been classified as a photographic image and correctly rejected by the system.

Type of test set:	2 Euro	1 Euro	50 Cent	Photos	Foreign
No. of coins in the test set:	780	780	780	780	60
... classified as photo	0	0	0	780	1
... classified as "3D"	780	780	780	0	59
... classified as Euro-coin	780	780	780	0	0

Table 2: Performance of the 3D-verification module and determination of the classifier's discriminative power.

These results prove that grey value based features derived from Three-Color SSGM-images of the embossed topography are suitable to classify coins reliably. The most disturbing influence originates from fluctuations in the reflected light within each coin class (due to slightly different color of the minting background, degree of oxidation, dirt, or wear and tear) which leads to rather broad intervals for the features. An identification of the coin's profile only by statistical grey value based methods without template matching is still not realized. To achieve this, further research has to be done to improve the method. First approaches are to normalize the images with respect to the grey values of the background and to develop contour based features.

7 Conclusion

In this paper we have presented the Three-Color Selective Stereo Gradient Method (Three-Color SSGM) as a general security-related machine-vision solution for real-time classification of moving objects with highly reflective flat metallic surfaces and complex structures. In order to prove and to evaluate its working and recognition principle we developed a classification scheme for moving Euro coins. Analyzing only the measured characteristic 3D-surface topography has led to a distinction between 2-Euro-, 1-Euro- and 50-Cent-coins and, within each sort, to a classification in 13 reference classes belonging to a coin's country of origin. This classification cannot be deceived by a photographic image of the coin due to an algorithmic 3D-verification module in combination with a specially designed color illumination and a Bayer-pattern equipped camera sensor. The recognition method is based on preclassification by rotation and translation invariant statistical grey value features in combination with the determination of dominant frequencies of the grey value profile taken from five fixed circles. The preclassification reduces the original number of classes to a few possible classes. The final decision is achieved by a template matching.

The presented classification performance has proved that our Three-Color SSGM-approach is suitable to classify Euro coins with high accuracy. Due to a robust ellipse-specific point fit used to segment the coin in the grabbed image the diameter of coins belonging to one sort has been determined with an average accuracy of less than 0.1 mm. In existing industrial coin validation systems, a false rejection rate (FRR) for "first coin insertion" is a common standard value. Compared to this value the achieved average FRR of 0,4% of the

presented Three-Color SSGM-system is an excellent result. No errors have been observed in determining the coin sort and the country of origin within the test set. No coins have been classified as photos and all photos of 2-Euro-coins have been sorted out.

The Three-Color SSGM can easily be adopted to further Euro denominations, to other currencies and to related problems of 3D-surface analyses. With such a Three-Color SSGM-system installed in existing coin validators parallel to the electromagnetic sensors the robustness of such machines toward coin recognition and coin rejection can be improved.

References

- [1] "Euro coin counterfeiting", *Press Release of the European Anti-Fraud Offices (OLAF)*, 18th January 2005.
- [2] Markus Adameck, Michael Hossfeld, and Manfred Eich, "A Vision System for Classification of Metallic Tokens using the Selective Stereo Gradient Method", *Proceedings of SPIE* 4661:206-217, 2002.
- [3] Markus Adameck, Michael Hossfeld, and Manfred Eich, "Three Color Selective Stereo Gradient Method for fast Topography Recognition of Metallic Surfaces", *Proceedings of SPIE* 5011:128-139, 2003.
- [4] Michael Hossfeld, Markus Adameck, and Manfred Eich, "Machine Vision detects Counterfeit Coins", *Laser Focus World* 39(6):99-104, 2003.
- [5] Berthold K.P. Horn, "Obtaining Shape from Shading Information", *The Psychology of Computer Vision* 115-155, 1975.
- [6] Ariel Tankus, "Shape-from-Shading under perspective projection", *Int. J. of Comp. Vision* 63(1):21-43, 2005.
- [7] Ruo Zhang et al., "Shape from shading: A survey", *PAMI* 21(8):690-706, 1999.
- [8] Robert J. Woodham, "Photometric Method for Determining Surface Orientation from Multiple Images", *Optical Engineering* 19(1):139-144, 1980.
- [9] Katsushi Ikeuchi, "Determine surface orientation of specular surface by using the photometric stereo method", *PAMI* 13(11):1139-1153, 1991.
- [10] Holger Lange, "Advances in the Cooperation of Shape from Shading and Stereo Vision", *3DIMIEEE*, p. 46-58, 1999.
- [11] Yoichi Sato et al., "Object shape and reflectance modeling from observation", *Com. Graphics Proc., Annual Conference Series (SIGGRAPH)*, 1997.
- [12] Tianli Yu et al., "Recovering Shape and reflectance Model of Non-lambertian Objects from Multiple Views", *Proc. of the 2004 IEEE Computer Vision and Pattern Recognition*, Los Alamitos: II-226 - II-233, 2004.
- [13] Hideo Saito et al., "Recovery of shape and surface reflectance of specular object from rotation of light source", *Image and Vision Computing*, 21:777-787, 2003.
- [14] Petr Beckmann and Andre Spizzichino, "The scattering of electromagnetic waves from rough surfaces", Oxford (a.o.): Pergamon Press, 1963.
- [15] F. E. Nicodemus et al., "Geometrical Considerations and Nomenclature for Reflectance", *U.S. Dep. of Commerce / National Bureau of Standards, Washington. D.C.* 160:1-52, 1977.

- [16] Stephen R. Marschner, "Image-based bidirectional reflectance distribution function measurement", *Applied Optics*, p. 2592-2600, 2000.
- [17] J. H. Lambert, *Photometria sive de mensura et gradibus luminis colorum et umbrae*, Augsburg, Germany: Eberhard Klett 1975.
- [18] A. S. Georghiades, "Incorporating the Torrance and Sparrow Model of Reflectance in Uncalibrated Photometric Stereo", *ICCVIEEE 2003*, p. 816-823.
- [19] K.E. Torrance and E.M. Sparrow, "Theory for off-specular reflection from roughened surfaces", *Journal of the Optical Society of America (JOSA)*, 57(9):1105-1114, 1967.
- [20] Bui Tuong Phong, "Illumination for Computer Generated Pictures", *Communications of the ACM* 18(6):311-317, 1975.
- [21] Shree K. Nayar, "Specular Surface Inspection using Structured Highlight and Gaussian Images", *IEEE Transactions on Robotics and Automation* 6(2):208-218, 1990.
- [22] R. Halir, J. Flusser, "Numerically Stable Direct Least Squares Fitting of Ellipses", *Proc. 6th Int. Conf. on Computer Graphics and Visualization*, Plzen (CZ), 125-132, 1998.