

Dagobert – A New Coin Recognition and Sorting System

Michael Nölle¹, Harald Penz², Michael Rubik²,
Konrad Mayer², Igor Holländer², Reinhard Granec²
ARC Seibersdorf research GmbH

¹Video- and Safety Technology , ²High Performance Image Processing
A-2444 Seibersdorf
michael.noelle@arcs.ac.at

Abstract. In this paper we describe the pattern recognition methods used for the identification of coins in the new coin recognition and sorting system Dagobert developed at the ARC Seibersdorf research centre. The purpose of Dagobert is to sort high volumes of coins. In particular the coin detection task, the features used to pre-select meaningful potential master coins for the verification process and the verification/rejection method itself will be detailed. Some recognition results are given at the end of the paper, indicating that the approach fulfils the given task effectively.

Keywords: Coin recognition, rotational invariant features, density dissimilarity measures, polar transformation, feature correlation.

1 Background and Introduction

The changeover from 12 European currencies to the Euro created a unique situation. Great volumes of money had to be physically returned to the national banks of the member states. Charity organisations took the opportunity to appeal for funds.

In Austria alone, by European standards a medium sized country with a population of 8 million, the charitable donations amounted to several hundred tons of cash. Unfortunately, the coins could only be collected as a potpourri of currencies. Quite often donors would generously contribute their left over holiday money accumulated over several years or their entire small cash collections. The sheer volume of material rules out any attempt to separate the money manually and calls for an automatic processing device.

Existing coin sorting machines usually work with only one currency at the time and therefore, cannot be used. Statistics evaluated on a small fraction of the material showed that more than 100 currencies are to be expected. This means that potentially two to three thousand different coin faces have to be recognised. Throughout this paper, by coin face we refer to the obverse and reverse side of the coin.

In general the material is not in mint condition. Circulation over decades has caused abrasion and dirtiness. Surprisingly, we found in addition that the mint

process itself is less accurate in many cases than one might expect. Out of centre imprints occur frequently as well as noticeable changes of the coin die over the years.

In this paper we will describe the pattern recognition approach used for the coin recognition and sorting system *Dagobert* designed and built at ARC Seibersdorf research GmbH. Section 3 will introduce the problem in more detail as well as the approach taken. Previous work is reviewed in Section 2. In Section 4 we will discuss some results based on data gathered in the laboratory as well as data from the actual production. Concluding remarks and future work are presented in Section 5.

2 Previous work

Several coin recognition approaches are mentioned in the literature.

Fukumi et al describe a system based on a rotation-invariant neural network that is capable of distinguishing Japanese coins ([4]), a 500 yen and a 500 won piece. Rotational invariance is achieved by explicitly generating the rotational group for a coarse model of the coin in a preprocessing step and feeding the results into a neural network. One drawback of the neural network approach is that it is not apparently clear how rejection of coins should be expressed. It is essential to be able to reject coins as it is impossible to know in advance which types of coins will be presented to the system.

Davidsson ([3]) compares several strategies, namely induction of decision trees ([10]), neural networks and Bayesian classifiers. He derives a variation of the decision tree algorithm that will reject coins if their defining attributes are outside an acceptance region. However, it is difficult to extend the approach to images.

Finally, Adameck et al presented an interesting method for a coin recognition system based on colour images ([1]). Similar to our approach translational invariance is achieved through segmentation, whereas rotational invariance is a result of a polar coordinate representation and correlation. Their system uses a special hardware to assure that no fraud coins are accepted by the system. In our case there is no risk that fraud images of coins will be presented to the system. Therefore, the use of colour seems to increase the computational costs unnecessarily.

3 Dagobert approach

For the sake of brevity we will restrict ourselves to the pattern recognition problem resulting from the overall aim of separating high volumes of diverse coins. For the recognition task, we assume that from a heap of money, coins are already singled out and put on a conveyor belt where a camera observes one coin at a time. The grey level images define the input to the recognition system. The major goals for the recognition task are:

- Separate coins from at least 30 different currencies. The desired currencies or more concrete the corresponding coin faces define the *recognition pattern set (RPS)*. False classifications should to be less than 0.01%.
- Reject coins that are not in the RPS.
- Real time performance. Real time in our case is defined by the performance of the mechanical parts of the Dagobert system. At present five to six coins per second can be loaded onto and taken from the conveyor belt.

As we could not find or access any substantial ground data, i.e. master coin collections or catalogues covering all required currencies, a first and quite laborious step was to gather the master data to build up the RPS.

The recognition process itself is roughly divided into three parts which are described below:

1. Coin detection: detect the coin in the current image, separate it from the background and determine the convex hull, centre of gravity and perimeter.
2. Coin pre-selection: for the current coin determine a short list SL of l objects from the RPS that minimise a dissimilarity measure.
3. Coin verification: for all objects in SL calculate the distance (error) to the current coin. Reject the current coin if the distance is too big, otherwise classify the current coin as the object with the smallest distance.

Step 2 purely serves for computational efficiency. As will be seen below the verification involves some computationally expensive algorithms that have to be performed for all master coin patterns under consideration. Pre-selecting promising candidates into the short list considerably reduces the costs for verification.

3.1 Coin detection

As long as the background conditions, i.e. the conveyor belt, can be controlled sufficiently, the coin detection task becomes almost trivial. If the conveyor belt is homogenous and always darker (brighter) than the coins, a simple (automatic) threshold operation suffices for the separation. This will be the assumption throughout this paper. Once the current coin is separated from the background we can calculate the centre of gravity, perimeter and convex hull.

But it should be noted that the coin detection may become the toughest part of the recognition and will be subject to future work.

3.2 Coin verification

In general the appearance of one coin pattern varies considerably with respect to its grey values. These variations frequently are inhomogeneous. This suggests that for recognition purposes grey values by themselves will not give us appropriate results.

On the other hand edge information remains more or less stable or at least degrades gracefully. Therefore, we based the coin recognition algorithms of Dagobert on edges. In principle any edge detector may be used for this purpose. But

from our experience the Canny edge operator and the Laplacian of Gaussian method ([2], [8], [5]) work satisfactorily. As a result of the edge operator we either get a binary (edge) image or a list of coordinates at edge pixel locations.

Let $I : \mathbb{M} \times \mathbb{N} \rightarrow \mathbb{R}_{[0,1]}$ be an intensity image. $\mathbb{M} \times \mathbb{N}$ gives the index space and $\mathbb{R}_{[0,1]}$ the intensity values taken from the closed interval $[0, 1]$.

$$E(x, y) = \begin{cases} 1, & \text{if } I(x, y) \text{ is an edge point;} \\ 0, & \text{else,} \end{cases} \quad (3.1)$$

is the binary edge image and $E^c = \{(x - x_m, y - y_m) | E(x, y) = 1\}$ the list of cartesian edge point coordinates with the centre of gravity (x_m, y_m) as origin. The polar coordinate representation of E^c is given by:

$$E^p = \{(\theta, \rho) | (x, y) \in E^c\}, \quad (3.2)$$

where

$$\begin{aligned} \theta &= \arctan y/x, \quad x = \rho \cos \theta, \\ \rho &= \sqrt{x^2 + y^2}, \quad y = \rho \sin \theta. \end{aligned}$$

Assume E_m is a master edge image that corresponds to E_a , the current coin

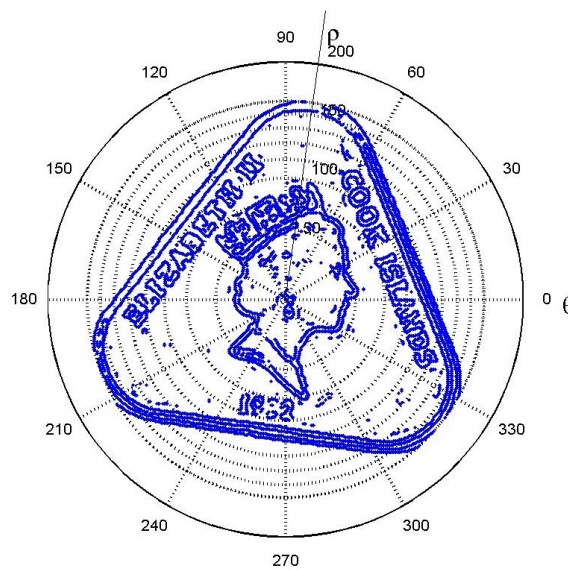


Fig. 1. Polar coordinate representation of an edge image.

edge image. Then in general there is an unknown rotation ϕ around the centre of gravity that aligns both edge images. In polar coordinates this rotation transforms into a cyclic translation in the angular direction. To determine ϕ we may

deploy a fast correlation method ([8], [5]) based on the edge images. Although correlation methods based on the fast Fourier transformation perform efficiently, there are some drawbacks using the edge images directly. First, to preserve the visual information the resolution of the edge image cannot be too small. Depending on the diameter of the coin we typically get coin image resolutions from 100×100 to 300×300 pixel's and the correlation would add significantly to the overall computational costs. Secondly, the outer border, which in most cases contains a substantial part of the edge points, usually does not help to find ϕ as it comprises too many symmetries. To avoid both drawbacks we suggest calculating the correlation on a two dimensional edge density function restricted to the inner part of the coin. This is given by:

$$H_{i,j}^d = |\{(\theta, \rho) \in E^p | \theta_{i-1} \leq \theta < \theta_i, \rho_{j-1} \leq \rho < \rho_j\}|,$$

$$E_{i,j}^d = H_{i,j}^d / N, i = 1, \dots, n; j = 1, \dots, l; N = \sum_{i,j}^{n,l} H_{i,j}^d. \quad (3.3)$$

$\{\theta_0, \dots, \theta_n\}$ and $\{\rho_0, \dots, \rho_l\}$ are the discrete resolutions in angular and distance directions, respectively. Now, we may estimate ϕ by correlating E_m^d and E_a^d . By choosing a high resolution in the angular direction (i.e. $n \geq 512$) and a coarse resolution (i.e. $l \leq 16$) in the distance direction, omitting the coin borders ($\rho_l < c \max_{(\theta, \rho) \in E^p}(\rho)$, $c \approx 0.9$), we found that ϕ usually may be determined up to $\pm 0.5^\circ$. Once ϕ is known, we may align the current coin image with the master. This is done efficiently by only calculating the rotated coordinates for the edge points in E_a^c , resulting in the rotated current coin edge image E_{a_ϕ} . From here we compute two distance measures:

$$e_{abrasion} = \frac{1}{|E_m^c|} \sum_{(x,y) \in E_m^c} 1 - E_{a_\phi}^{d_s}(x, y), \quad (3.4)$$

$$e_{dirt} = \frac{1}{|E_{a_\phi}^c|} \sum_{(x,y) \in E_{a_\phi}^c} 1 - E_m^{d_s}(x, y), \quad (3.5)$$

where E^{d_s} is the result of applying s morphological dilation operations to the binary edge image E in order to counteract the remaining uncertainty of the angular position. $e_{abrasion}$ tells us how many expected (master) edge points are missing, whereas e_{dirt} sums the additional edge points in the current edge image. If these errors are higher than given thresholds we have to dismiss the match.

In general we cannot know which master coin corresponds to the current coin image. Therefore, we have to calculate Equations 3.4 and 3.5 for all possible master coin candidates. In the next subsection we will derive features that help us to keep this set small.

3.3 Coin pre-selection features

The Dagobert system is equipped with two additional sensors measuring the thickness as well as the rough diameter of the current coin. At the same time

they are used to trigger the imaging process. For the main results in this paper we did not use these measurements, but they deliver valuable information for the pre-selection process as well. The production system uses these measurements for a coarse pre-selection of potential master coins into the short list *SL*. In this section we derive features that help us to refine this first pre-selection. Solely based on the edge information we derive three additional types of features that in turn are invariant against rotations of the coin. The first two use an edge density function similar to the one given by Equation 3.3 in the polar coordinate system whereas the third will be computed in cartesian coordinate space.

Let $\rho_0 = 0$, $\rho_l = \max_{(\theta, \rho) \in E^p}(\rho)$ and choose $\rho_1, \dots, \rho_{l-1}$ such that

$$\rho_i^2 - \rho_{i-1}^2 = \rho_{i+1}^2 - \rho_i^2. \tag{3.6}$$

By Equation 3.6 all circlets with inner and outer radius ρ_{i-1}, ρ_i , respectively cover the same amount of area, $a = 2\pi(\rho_i^2 - \rho_{i-1}^2)$ and thereby, have an equal probability for containing edge points. Let $N = |E^p|$. The distribution of edge points F^ρ over the distance may be expressed as

$$F^\rho(i) = |\{(\theta, \rho) \in E^p | \rho_{i-1} \leq \rho < \rho_i\}|/N, \tag{3.7}$$

for $i = 1, \dots, l$. Likewise the distribution of edge points over the angles is given by

$$F^\theta(i) = |\{(\theta, \rho) \in E^p | \theta_{i-1} \leq \theta < \theta_i\}|/N, \tag{3.8}$$

for $i = 1, \dots, n$. F^θ is not invariant against rotations of the coin, but the modulus of the discrete Fourier transform, $F^\theta = |DFT(F^\theta)|$, is. F^ρ naturally is invariant against rotations.

For the third feature consider a function f^c that is invariant against cyclic rotations of binary strings of length d , i.e. for $u, v \in \{0, 1\}^d$

$$f^c(u) = f^c(v) \Leftrightarrow \exists 0 \leq t < d \forall 0 \leq i < d \ u_{i \oplus t} = v_i, \tag{3.9}$$

where \oplus is the modulus d addition and $\{0, 1\}^d$ describes words of length d over a two symbol alphabet. As an example consider two binary strings '00101' and '10100', $d = 5$, that are aligned by $t = 2$. f^c defines an equivalence relation on $\{0, 1\}^d$ as it is reflexive, symmetric and transitive. Therefore, it is complete, i.e. able to separate all patterns of binary strings of length d . If d is small (i.e. $d \leq 16$) the function can be realised as a function table and calculated explicitly offline. See Figure 2 as an example. Let k be the number of equivalence classes produced by f^c and $f^c : \{0, 1\}^d \rightarrow \{1, \dots, k\}$. Then

$$F^c(i) = |\{(x, y) \in E^c | f^c(C_r(x, y)) = i\}|/N, \tag{3.10}$$

for $i = 1, \dots, k$, gives the distribution of binary strings on circles centred at edge points of the edge image. $C_r(x, y)$ determines the contents of the edge image on a circle with radius r at centre (x, y) (cf. Figure 2). In order to represent the given information r should be chosen in relation to k such that $k \approx 2\pi r$. By construction F^c is invariant against rotations of the edge image.




Fig. 2. Left: a binary string of pixels centred at each edge point is mapped to an invariant class. Middle: number of equivalence classes (#) in relation to the length of the binary string (d). Right: an example of a function table for $d = 4$.

F^ρ , F^θ and F^c define density functions on binary images. In previous work ([7]) we showed that for given density functions p_1, p_2 of a probability variable X with discrete events x_1, \dots, x_n , we may define a metric function F_{sin} by

$$F(p_1, p_2) = \sum_{x \in X} \sqrt{p_1(x)} \sqrt{p_2(x)},$$

$$F_{\text{sin}}(p_1, p_2) = \sqrt{1 - F^2(p_1, p_2)}.$$

It is easy to verify that $0 \leq F_{\text{sin}}(p_1, p_2) \leq 1$, as p_1, p_2 are density functions, i.e. $\sum_{x \in X} p_{1,2} = 1$ (see [7] for further details). Moreover, $F_{\text{sin}}(p_1, p_2) = 0 \Leftrightarrow F(p_1, p_2) = 1 \Leftrightarrow p_1 = p_2$. F_{sin} can be used to compare discrete density functions.¹ Let M be a set of master edge images that belong to the same coin face. If F_a^ρ , $F_{m_i}^\rho$ and F_a^θ , $F_{m_i}^\theta$ and F_a^c , $F_{m_i}^c$ are the density functions of the current coin E_a and a master coin edge image $m_i \in M$, respectively, we may define the distance of the current coin and the master by

$$d(E_a, M) = \min_{m_i \in M} (w_1 F_{\text{sin}}(F_a^\rho, F_{m_i}^\rho) + w_2 F_{\text{sin}}(F_a^\theta, F_{m_i}^\theta) + w_3 F_{\text{sin}}(F_a^c, F_{m_i}^c)),$$

$$1 = \sum_{i=1}^3 w_i. \tag{3.11}$$

¹ Other functions may be used as well for this comparison, like the χ^2 ([9]) or Jensen Shannon divergence ([6]). However, from our experience F_{sin} tends to give sharper results if the density functions to be compared are quite close to each other.

In the experiments we could not verify that different weights are advantageous. Therefore, we used $w_i = 1/3, i = 1, 2, 3$. By computing the distance from the current coin to all master coins and selecting those l master coins with smallest distance for the short list SL of hypothesis, we can effectively keep the list of master candidates for the verification process small.

4 Results

As mentioned before additional thickness and area sensors are used for the Dagobert system. Based on their measurements a first rough pre-selection of potential master coins is determined. This provides us with a set of master coins that have almost the same diameter and that have to be distinguished by the methods given in the previous sections.

Test set one	coins	%	coins	%	coins	%	coins	%
Coin in RPS	214	$l=1$	214	$l=2$	214	$l=3$	214	$l=6$
recognised	196	91.59	202	94.39	205	95.79	205	95.79
rejected (false negative)	18	8.41	12	5.61	9	4.21	9	4.21
not recognised	0	0	0	0	0	0	0	0
	coins				%			
Coin not in RPS	334				$l=1, \dots, 6$			
rejected	333				99.7			
false positive	1				0.3			

Table 1. Recognition (upper) and rejection (lower) results. l indicates the length of the short list SL .

A closer analysis of the material showed that the biggest subclass of coins we found so far is given by those with a diameter of 24mm. For this reason some of the results presented in this section are based on images of 24mm coins as collected in the laboratory. In particular we chose 39 master coin classes to form the RPS. The master coin patterns are based upon five coin pictures (if available) to keep accidental fluctuations small. Five or more additional pictures per master coin class were taken for testing purposes. In total the experiments were carried out on 214 coins that belonged to one of the master classes but were not used for training. These together with 334 coins from 60 (unknown) classes that had to be rejected form the *test set one*. All master coin patterns are defined on (almost) the same number of coins and the distribution of test coins is uniform with respect to the master coin classes. This insures that the results have no accidental bias.

Table 1 summarises the results of the recognition experiments. The overall result is quite encouraging as almost 96% of the known coins are classified correctly ($l \geq 3$) and more than 99% of the unknown coins are correctly rejected.

Test set two	coins	%
Coin in RPS	12192	1=6
recognised	10292 (10337)	84.42 (84.79)
rejected (false negative)	1842	15.11
not recognised	58 (13)	0.48 (0.11)
	coins	%
Coin not in RPS	757	1 = 6
rejected	736 (745)	97.23 (98.41)
false positive	21 (12)	2.77 (1.58)
Test set two	sorting result [%]	
Coins correctly sorted	99.24 (99.76)	
Coins incorrectly sorted	0.76 (0.24)	

(a)

(b)




Fig. 3. (a) Recognition (upper) and rejection (middle) results, sorting results (lower). The numbers in brackets indicate the results if coins that are not distinguishable by one face are grouped together in one class. (b) Example of a Spanish coin that exists in an old (bottom) and new (middle) version. A current coin is given at the top.

The length of the short list has some influence on the recognition result. $d(E_a, M)$ necessarily is not as exact as $e_{abrasion}$ and e_{dirt} , as it does not take the entire geometric relationship of the edge points into account. Therefore, sometimes coins are accidentally ranked higher in the short list than they should be. On the rejection results the length of the short list had no influence.

Most importantly the number of coins that were incorrectly sorted into the valid set is very small. Additionally, coins were either recognised correctly or rejected, which insures that sorted coin piles are 'clean', i.e. almost all coins belong to the same class.

To generate a bigger unbiased benchmark test set for this application has proven almost impossible. Naturally, there are more local coins in the collections, in our case from Austria and the surrounding countries, than from further abroad. The same is true for defining the master coin patterns unless data from the coin producers is available, which usually is not the case. Therefore, the available training sets for some currencies contain a very limited number of coins.

As a second test set we took 12949 coin images and validated them manually. They were tested against 913 master coin patterns of all diameters in the RPS. Figure 3(a) shows the results. Again the set of incorrectly sorted coins is quite small. Some of the coins can only be distinguished knowing both faces (see Figure 3(b)). The figures in brackets show the results when we combine the master coin patterns of those coins into one. The number of incorrectly rejected coins seems to be too high. This has to be analysed further.

So far the system has sorted several tons of coins and is able to meet the realtime conditions, i.e. to process 5 to 6 coins per second. Using the obverse

and reverse face for the recognition task, between 85% and 90% of the material is sorted into classes defined in the RPS, which contains at present around 1500 patterns of coin faces. The rest is rejected. Random tests performed on classified sets of coins indicate that we seem to meet the goal of having less than 0.01% incorrect positive classifications.

5 Conclusion and future work

In this paper we described a new coin recognition and sorting system Dagobert that is presently being built at the ARC Seibersdorf research GmbH. The system is capable of sorting heaps of mixed coins. Thereby, coins from more than 30 countries can be recognised and separated. Unknown coins are rejected. Further research will be carried out to improve the recognition result and speed.

Acknowledgement: We would like to thank Janice Knight and Ian Glendingin for editing the manuscript and the helpful discussions.

References

1. M. Adameck, M. Hossfeld, and M. Eich, *Three color selective stereo gradient method for fast topography recognition of metallic surfaces*, Proceedings of Electronic Imaging, Science and Technology (Martin A. Hunt and Jeffery R. Price, eds.), Machine Vision Applications in Industrial Inspection XI, vol. SPIE 5011, January 2003, pp. 128–139.
2. John Canny, *A computational approach to edge detection*, IEEE Transactions on Pattern Analysis and Machine Intelligence **PAMI-8** (1986), no. 6, 679–698.
3. Paul Davidsson, *Coin classification using a novel technique for learning characteristic decision trees by controlling the degree of generalization*, Ninth International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems (IEA/AIE-96), Gordon and Breach Science Publishers, 1996, pp. 403–412.
4. Minoru Fukumi, Sigeru Omatu, Fumiaki Takeda, and Toshihisa Kosaka, *Rotation-invariant neural pattern recognition system with application to coin recognition*, IEEE Transactions on Neural Networks **3** (1992), no. 2, 272–279.
5. Rafeal C. Gonzalez and Richard E. Woods, *Digital image processing*, Addison-Wesley Publishing Company, 1993.
6. J. Lin, *Divergence measures based on the Shannon entropy*, IEEE Transactions on Information Theory **37** (1991), no. 1, 145 – 151.
7. Michael Nölle, *Distribution distance measures applied to 3-d object recognition – a case study*, Proceedings of the 25th Pattern Recognition Symposium of the German Association for Pattern Recognition, Magdeburg, Germany, September 10-12, 2003.
8. James R. Parker, *Algorithms for image processing and computer vision*, New York: John Wiley & Sons, Inc., 1997.
9. Puzicha, J., Rubner, Y., Tomasi, C., and Buhmann, J. M., *Empirical Evaluation of Dissimilarity Measures for Color and Texture*, Proc. of the International Conference on Computer Vision (ICCV'99), 1999, pp. 1165–1173.
10. J. R. Quinlan, *Induction of decision trees*, Machine Learning **1** (1986), no. 1, 81–106.