Computer Vision Winter Workshop 2008, Janez Perš (ed.)
Moravske Toplice, Slovenia, February 4–6
Slovenian Pattern Recognition Society, Ljubljana, Slovenia

Segmentation of Ancient Coins Based on Local Entropy and Gray Value Range

Sebastian Zambanini and Martin Kampel

Pattern Recognition and Image Processing Group
Institute of Computer Aided Automation
Vienna University of Technology
Favoritenstr. 9/1832, A-1040 Vienna, Austria
{zamba,kampel}@prip.tuwien.ac.at

Abstract *Numismatics deals with various historical aspects of the phenomenon Money. Fundamental part of a numismatists work is the identification and classification of coins according to standard reference books. Prior to any analysis a coin image has to be segmented into two areas: the area depicting the coin and the area belonging to the background. In this paper, we focus on the segmentation task as a preprocessing step for any automated coin identification system. We present a simple and fast method for coin segmentation, based on local entropy and gray value range. Results of the algorithms developed are shown for an image database of ancient coins.*

1 Introduction

Nowadays, ancient coins are becoming subject to a very large illicit trade. Thus, the interest in reliable automatic coin recognition systems within cultural heritage and law enforcement institutions raises rapidly. Traditional methods to fight the illicit traffic of ancient coins comprise manual, periodical search in auctions catalogues, field search by authority forces, periodical controls at specialist dealers, and a cumbersome and unrewarding internet search, followed by human investigation. However, these methods only prevent the illicit trade of ancient coins to a minor extent. To date, no automatic coin recognition system for ancient coins has been researched and thus applied successfully.

For the image-based recognition of ancient coins, initially a segmentation of the coin region has to be done in the image. Especially for the identification of stolen coins, a correct segmentation is a crucial step since the shape of the coin provides a substantial feature for its identification. An automatic segmentation method is also of great benefit for the indexing of new coins, since up to now numismatists have to perform this time-consuming task manually.

Major challenges that have to be faced in the segmentation of coins are caused by an improper image acquisition procedure, as can be seen in the two coin images of Fig. 1. In both images a hard shadow is cast at the coin border because of the insufficient illumination setup. Furthermore, the background of Fig. 1(b) was not well chosen because it has no consistent color.

In this paper a simple and fast method for coin segmentation, based on local entropy and gray value range, is presented. The underlying assumption of the proposed method is that the local entropy (i.e. the information content) and range of gray values is higher inside the coin and at the coin border, since the background shows a higher homogeneity than the coin. For that reason the local entropy and range are summed up and the final segmentation region is obtained by global thresholding.

The paper is organized as follows: Section 2 gives a short overview of the COINS project, which represents the framework where the segmentation task belongs to. An image acquisition system is described in Section 3. The coin segmentation itself is addressed in Section 4, consisting of an overview of the state of the art and a detailed description of the proposed method. Experiments on a set of 92 images are reported in Section 5. A conclusion is finally given in Section 6.

2 The COINS project

The COINS project focuses on technologies aimed at permanent identification and traceability of ancient coins. Furthermore, it devises strategies to facilitate the prevention and repression of illicit trade of stolen coins. To achieve these goals, three main activities have been identified: (1) standardization of numismatic data structures, (2) numismatic web search tool and (3) image based recognition tool for ancient coins (see Figure 2).

Documentation and inventory methodologies and tools based on international standards facilitate the interoperability and cross-border traceability. Starting from analysis of current forms for inventorying coins, a standardized domain ontology and multilingual thesaurus will be produced.

An automatic, unsupervised web search tool is greatly needed. Thus, core activity of the project addresses the development of a web tool that is integrating text search and image analysis and tailored to numismatic needs.

The goal of the image based coin recognition tool is twofold. On the one hand, it addresses the classification of an ancient coin based on its visual representation. On the other hand – the identification of individual coin based on peculiar features, as minting signs or use-wear traces. However, their procedures differ. Classification must ignore the individual features and emphasize the general

(a)

(b)

Figure 1: Examples of coin images acquired under improper conditions.

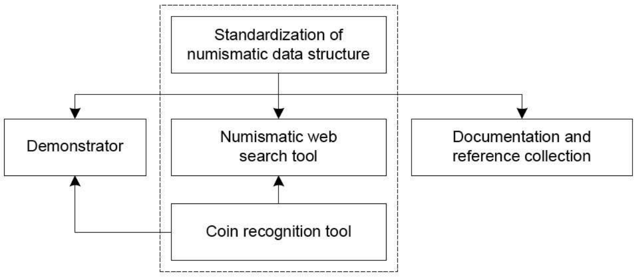

Figure 2: Core activities within the COINS project

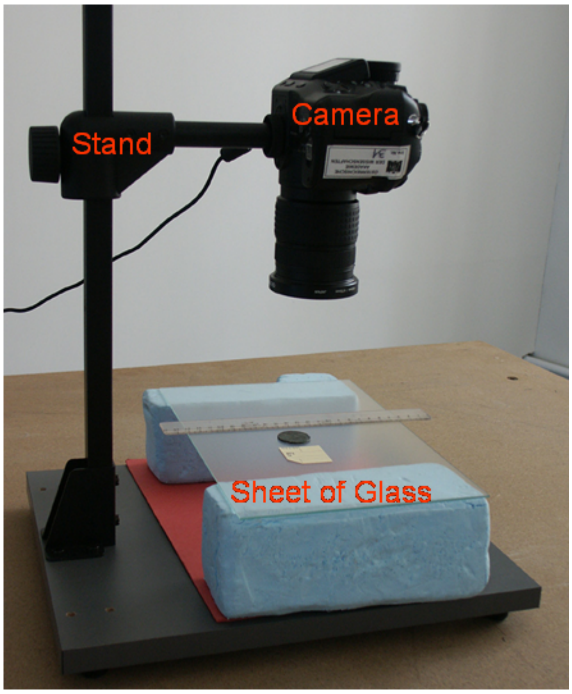

Figure 3: Acquisition system used at the *Kunsthistorisches Museum Vienna*.

ones, to assign an individual coin to a general category, whereas identification relies on individual unique features, which make that particular instance different from all other individuals in the same class.
To achieve these goals, the joined efforts of experts in different fields – technology, cultural heritage and law enforcement – are required.

3 Coin Image Acquisition

The performance of image recognition methods is highly related to the image quality. For the segmentation of coins in images, there are two main issues that lead to a higher degree of robustness in coin segmentation: the avoidance of shadow casts at the coin border and the use of a constant background with high contrast to the coin border. Fig. 3 exemplarily shows a proper coin image acquisition system concerning these goals that is currently used at *Kunsthistorisches Museum Vienna*. A digital single-lens reflex camera with a 60mm macro lens is mounted on a camera stand providing a constant distance and parallelism between the coin and the camera's image plane. For a constant background with high contrast to the coin a red sheet of paper is used. A ruler and a label is put next to the coin for a later determination of the image scale and a unique identification of the coin. For the avoidance of shadow casts at the coin border, the coin is laid on raised sheet of glass. The effect of this approach is shown in Fig. 4: if the coin is directly placed on the background paper a shadow is cast (Fig. 4(a)), whereas by a placement on a raised sheet of glass no shadows are visible in the image (Fig. 4(b)).
It is important to state here that this setup describes an easy way for numismatists to acquire images which allow an accurate segmentation due to the absence of shadows and high contrast to the background. Nevertheless, for the identification and classification of coins, a further goal has to be an optimal visualization of details with the avoidance of highlights on the coin. Therefore, more research has to be done in that area to define optimal illumination conditions and set of camera parameters.

4 Coin Segmentation

This section addresses the segmentation of the coins in the images. Coin segmentation deals with the division of the image into two regions: the region depicting the coin and the region belonging to the background.

(a)

(b)

Figure 4: (a) image of a coin directly placed on the background producing a shadow, (b) the same coin placed on a raised sheet of glass producing no shadows. In contrast to Fig. 3, a blue background was used here.

4.1 Image Segmentation

Image segmentation refers to the process of dividing the image into regions that correspond to structural units in the scene or distinguishing objects of interest. For our purpose, a segmentation method has to detect the coin region, regardless of other objects (e.g. a ruler or label identifying the coin) contained in the image. In general, image segmentation algorithms may follow three approaches [16]: Thresholding, edge-based segmentation and region-based segmentation.

4.1.1 Thresholding Thresholding methods define a range of brightness values (the *thresholds*) in the original image and select the pixels within this range as belonging to the foreground, whereas the remaining pixels are rejected to the background. The basic assumption of thresholding methods is that the gray levels of the object are significantly different from the gray values of the background. Thresholding techniques can either work globally, where a single threshold is applied to the whole image, or locally, where the image is divided into regions and each region has its own threshold. Besides that thresholding techniques differ in the way of finding optimal threshold values for a given image, e.g. the use of histogram information [4] or entropy of gray level distribution [15]. A survey is given in [14].

4.1.2 Edge-Based Segmentation This category of segmentation methods partitions an image based on abrupt changes in the intensity, i.e. edges found in an image by edge detectors [6]. Apart from that, edge-based segmentation differ in the further methodology. In edge relaxation, a global relaxation (optimization) process based on edge properties is used to form continuous boundaries of objects [2]. Border tracing methods are used to follow the objects borders from a known start point [1]. In the case of additional knowledge about the objects to be segmented, globally optimal borders can be found using graph searching [10] or dynamic programming [3]. For the segmentation of certain shapes like circles the Hough transform [8] can be applied.

4.1.3 Region-Based Segmentation Region-based segmentation methods try to partition or group regions according to common image properties, like color or texture. Split-and-Merge [7] combines two operations to segment an image: splitting, where the image is divided into a set of regions which are coherent within themselves, and merging, where adjacent splitted regions are merged together based on a similarity criterion. In the watershed segmentation [17] the image is considered as a topographic surface. According to that analogy, the watershed transform finds “catchment basins” and “watershed ridge lines” where the catchment basins theoretically correspond to the homogeneous gray level regions of this image.

4.2 Proposed Method

For typical images of coins (like the two shown in Fig. 4), the coin itself possesses more local information content and details than the rest of the image, i.e. the background. For that reason, our method is based on two filters providing a local measurement of information content in the image: the *local entropy* and the *local range of gray values*.

Local entropy of an image: Entropy is the measure of the information content in a probability distribution. For digital images the probability distribution is represented by the histogram of gray values [9]. If an image consists of N possible gray values whose actual frequencies of occurrence (i.e. the normalized image histogram values) are $p_1, p_2, ..., p_N$ the entropy of the image is defined as

$$H = -\sum_{k=1}^{N} p_k \cdot \log_2(p_k) \tag{1}$$

With local entropy the entropy of each pixel is computed individually by means of the gray values of the local neighborhood.

Local range of gray values: The local range of gray values is defined as the difference of the maximum and minimum gray value of a local neighborhood.

The outputs of these two filters are summed-up to build the final intensity image where the thresholding is applied on. For both filters a circular neighborhood with a radius of 3 pixels is used and both filter outputs are normalized to the range 0 to 1. In Fig. 5 the particular results of the entropy filter, the range filter and their summation, applied to a coin image, are shown. Note that the output of both filters is

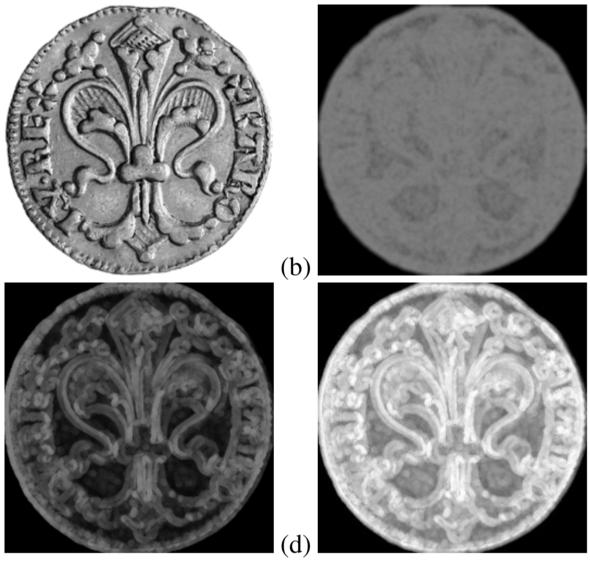

Figure 5: (a) original image, (b) output of local entropy filter, (c) output of local range filter, (d) sum of local entropy and local range.

higher for the region of the coin than for the region of the background, especially at the coin border.

To obtain the final coin segmentation from the intensity image shown in Fig. 5(d), a simple way would be to apply a global threshold and close all holes in the binary mask caused by homogeneous regions inside the coin. However, tests have shown that such a manually defined threshold does not perform well on the overall given test set. Therefore, a more sophisticated approach is used: we apply seven thresholds T_i ($T_i = 0.3, 0.4, ..., 0.9$) to the intensity image and compute a score for each achieved segmentation that represents the confidence to the given segmentation. Afterwards the segmentation with highest confidence is chosen. Since the shape of a coin is close to a circle, we use the *formfactor* [12] of the binary segmentation mask as confidence measure. The formfactor of a binary mask is computed as follows:

$$formfactor = \frac{4\pi A}{P^2} \tag{2}$$

where A is the area and P the perimeter of the the binary mask. The form factor is sensitive to both the elongation of a region and the jaggedness of its border. The higher the jaggedness of the border, the less the formfactor. The formfactor is equal to 1 for a circle and is less for any other shape. Since the final shape of the segmentation should be close to circle with a regular border, the formfactor provides a convenient measure for the confidence of the segmentation.
Since low thresholds can produce a coin segmentation that is near the rectangular shape of the whole image (providing a comparatively high formfactor), a segmentation is furthermore only accepted if the area of the segmented region is lower than 90 % of the image area. An example for the seven segmentations obtained with different thresholds is shown in Fig. 6. The segmentation obtained with $T_7 = 0.9$ shows the highest formfactor and is therefore chosen as the final segmentation (note that the segmentations with $T_1 = 0.3$ and $T_2 = 0.4$ are rejected since their area is more than 90 % of the overall image area).

(a) Original image.

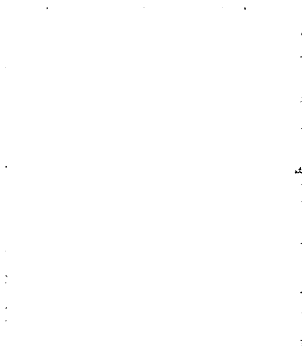

(b) $T_1 = 0.3$, formfactor = 0.646

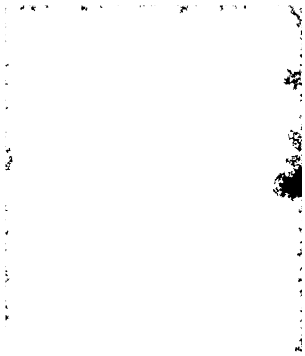

(c) $T_2 = 0.4$, formfactor = 0.165

(d) $T_3 = 0.5$, formfactor = 0.004

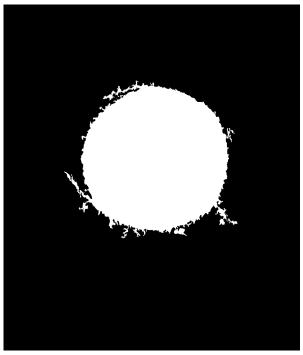

(e) $T_4 = 0.6$, formfactor = 0.093

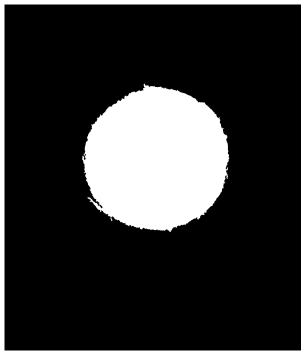

(f) $T_5 = 0.7$, formfactor = 0.428

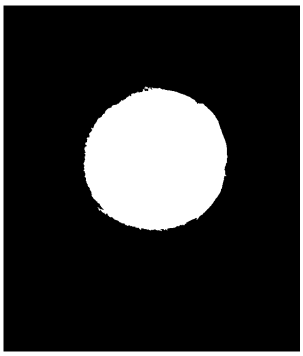

(g) $T_6 = 0.8$, formfactor = 0.562

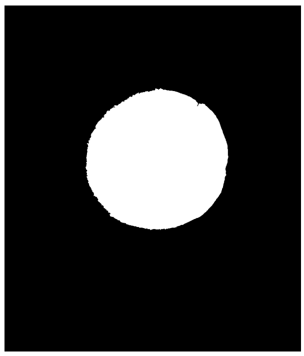

(h) $T_7 = 0.9$, formfactor = 0.691

Figure 6: Seven segmentation masks according to different thresholds T_i applied to to the intensity image.

For our method, the discretization of $T_i = 0.3, 0.4, ..., 0.9$ was chosen empirically. Tests have shown that a finer discretization does not improve the accuracy of the method.

5 Experiments

The proposed method was tested on a set of 92 images acquired at the *Kunsthistorisches Museum Vienna*, the *Fitzwilliam Museum* and the *Muzeul National de Istorie a României*, representing a wide range of different coin images. Six of the images used for evaluation are exemplarily shown in Fig. 7. The images of the evaluation set differ in various ways:

- Resolution: from 178×184 up to 1154×866.
- Color: color images (e.g. Fig. 7(a)) or gray-level images (Fig. 7(b)).
- Background: images with uniform white background (Fig. 7(b)), images with less uniform white background (Fig. 7(c) or colored background (Fig. 7(a)).
- Coin size relative to image size: images where the coin perfectly fits into the image frame (Fig. 7(b)) or images where the coin region makes only $\sim 15\%$ of the image (Fig. 7(c)).
- Illumination conditions: images with (Fig. 7(c)) and without shadow casts (Fig. 7(a)).

Image Fig. 7(a) represents a coin acquired with the acquisition system described in Section 3. In total 10 of such images are contained in the test set to reveal the acquisition system's suitability for coin segmentation. For the experiments presented here, all color images were converted to gray-level images. The segmentation of color images is topic for future research.

5.1 Setup

For each image a ground truth segmentation was manually obtained by means of a commercial image editing program. For the evaluation of the segmentation error a metric called *border error* is measured [5]:

$$\text{border error} = \frac{\text{number of misclassified pixels}}{\text{area of coin}} \quad (3)$$

A pixel is called *misclassified* if it is contained in the ground truth segmentation but not in the automatic or vice versa. Note that this error is independent of both the size of the coin and the size of the image.

5.2 Results

On the overall set of test images an average border error of 4.54 % is achieved, and the median of the results lies at 2.65%. Fig 8 shows results on images where the obtained coin border is outlined by a black or white line. Figure 8(a)-(c) belong to the best segmentation results with border errors of 0.29%, 0.37% and 0.50%, respectively. Figure 8(d)-(e) belong to the worst results with border errors of 15.81%, 11.71%, respectively. You see that shadows pose a problem to the method since they produce a strong edge not belonging to the actual coin border. However, on the image of Figure 8(f) the method correctly excludes the shadow from the segmentation, producing a border error of 1.93%. On ten images acquired with the proposed acquisition system, the algorithm shows an average border error of 0.67%. Although there are 11 particular images with border errors of above 10%, no image was segmented with a border of above 25%, a indication for high robustness under several types of coin images. In fact, the highest border error lies at 21.03% and 60 out of the 92 images could be segmented with a border error of less than 5%.

5.3 Comparison

In [19] segmentation of ancient coins was achieved by a adaptive thresholding method originally suggested by Yanowitz and Bruckstein [18]. The proposed method derives a threshold surface which is interpolated using tie points placed at positions obtained from thinned and thresholded gradient values. Applied to the same test set, this method achieves an average border of 65.54% and a median value of 77.71%. Fig. 9 shows the plot of border errors for all 92 images, achieved with the proposed algorithm based on local entropy and gray value range (solid line) and the algorithm proposed in [19] (dotted line). It can be seen that there are cases where the algorithm has a similar border error than our method (e.g. for image 9 and 11). However, the high number of images with border error greater or equal than 100% (38 in total) indicate its low robustness. Although the parameters of the method can be adjusted to perform well on a given type of coin image it is not able to handle the wide range of different images contained in the test set. Furthermore, the method is approximately four times slower than the proposed method.

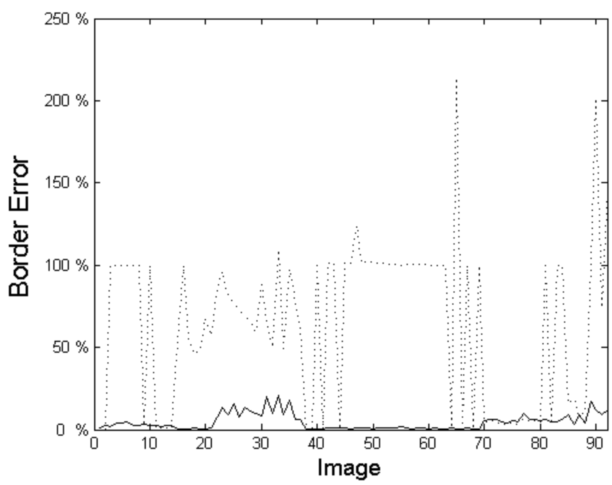

Figure 9: Border error plot of 92 segmentations.

Figure 7: 6 of the 92 coin images used for evaluation.

6 Conclusion

The paper proposes a method for a robust segmentation of ancient coins. It uses range and entropy filter under the assumption that the coin in the image provides more information and details than the background. The final segmentation mask is obtained by thresholding operation where the optimal threshold is found by means of the formfactor of the resultant binary mask.

The method shows promising results with a border error median of 2.65% and proves that local entropy and gray value range give a convenient estimate of the actual coin region. However, although the method's robustness is indicated by a maximum border error of 21.03% on a set of 92 test images, shadows still pose a problem. For a maximal segmentation performance a controlled acquisition system like the one presented in Section 3 is highly desired. This is also proofed by a comparison of the method's performance on the 10 "optimal" images acquired under the proposed system with the performance on the remaining images: on the optimally acquired images the average border error lies at 0.67% where on the remaining set it lies at 5.01%.

For the identification of stolen coins, where the coin shape serves as a highly discriminative feature, an exact segmentation is essential. Therefore, future research has the goal to improve the accuracy to satisfy the needs of automatic coin identification, especially in the occurrence of shadows.

Extensions of the presented method that has to be investigated for future research include:

- **Use of color:** the colors of metal limits the range of possible coin colors.
- **Shadow detection:** Shadows can be excluded from segmentation by shadow detection algorithms [13]. Shadow detection is thereby based on a analysis of edges with respect to the possibility that they are due to a material change as opposed to a shadow or other illumination effects.
- **Border tracing:** Border tracing methods [16] can be used to determine the exact border of the coin. Start points of border tracing could be found by a circle detection algorithm [11] followed by a maximum gradient search along the roughly estimated radius.

Acknowledgement

The authors would like to thank Dr. Klaus Vondrovec (*Kunsthistorisches Museum Vienna, Austria*), Dr. Mark Blackburn (*Fitzwilliam Museum, UK*) and Dr. Ernest Oberlaender Tarnoveanu (*Muzeul National de Istorie a României, Romania*) for providing the test images.

This work was partly supported by the European Union under grant FP6-SSP5-044450. However, this paper reflects only the authors' views and the European Community is not liable for any use that may be made of the information contained herein.

References

[1] S. A. Dudani. Region extraction using boundary following. In *Pattern Recognition and Artificial Intelligence*, pages 216–232, 1976.

(a) border error = 0.29%


(b) border error = 0.37%


(c) border error = 0.50%


(d) border error = 15.81%


(e) border error = 11.71%


(f) border error = 1.93%


(g) border error = 3.21%

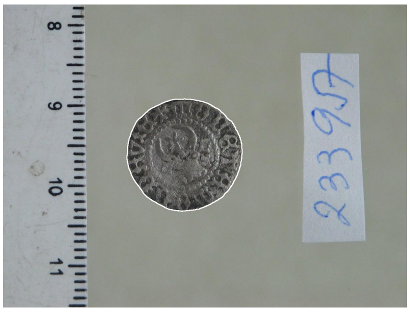
(h) border error = 0.54%


(i) border error = 5.46%

Figure 8: Results of the proposed segmentation method.

[2] J. S. Duncan and T. Birkhölzer. Reinforcement of linear structure using parametrized relaxation labeling. *IEEE Trans. Pattern Anal. Mach. Intell*, 14(5):502–515, 1992.

[3] M. A. Furst and P. E. Caines. Edge detection with image enhancement via dynamic programming. *Image Understanding*, 33(3):263–279, Mar. 1986.

[4] C. A. Glasbey. An analysis of histogram-based thresholding algorithms. *Graphical Model and Image Processing*, 55(6):532–537, 1993.

[5] G.A. Hance, S.E. Umbaugh, R.H. Moss, and W.V. Stoecker. Unsupervised color image segmentation. *IEEE Engineering in Medicine and Biology*, 15(1):104–111, January/February 1996.

[6] M. Heath, S. Sarkar, T. Sanocki, and K. Bowyer. Comparison of edge detectors. A methodology and initial study. *Computer Vision and Image Understanding: CVIU*, 69(1):38–54, 1998.

[7] S. L. Horowitz and T. Pavlidis. Picture segmentation by a directed split-and-merge procedure. In *Proceedings of the Second International Joint Conference on Pattern Recognition*, pages 424–433, 1974.

[8] P. V. C. Hough. Methods and means for recognizing complex patterns. US Patent 3,069,654, 1962.

[9] J. N. Kapur, P. K. Sahoo, and A. K. C. Wong. A new method for gray-level picture thresholding using the

entropy of the histogram. *Computer Vision, Graphics, and Image Processing*, 29:273–285, 1985.

[10] A. Martelli. An application of heuristic search methods to edge and contour detection. *Commun. ACM*, 19(2):73–83, 1976.

[11] A. A. Rad, K. Faez, and N. Qaragozlou. Fast circle detection using gradient pair vectors. In *Proceedings of the Seventh International Conference on Digital Image Computing: Techniques and Applications, DICTA 2003*, pages 879–888.

[12] J. C. Russ. *The Image Processing Handbook*. CRC Press, Boca Raton, FL, 2nd edition, 1995.

[13] E. Salvador. *Shadow Segmentation and Tracking in Real-Word Conditions*. PhD thesis, Ecole Polytechnique Fédérale de Lausanne (EPFL), 2004.

[14] M. Sezgin and B. Sankur. Survey over image thresholding techniques and quantitative performance evaluation. *J. Electron. Imaging*, 13(1):146–165, June 2004.

[15] A. G. Shanbhag. Utilization of information measure as a means of image thresholding. *Graphical Models and Image Processing*, 56(5):414–419, Sept. 1994.

[16] M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis, and Machine Vision*. Thomson-Engineering, 1999.

[17] L. Vincent and P. Soille. Watersheds in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):583–598, June 1991.

[18] S.D Yanowitz and A.M. Bruckstein. A new method for image segmentation. *Computer Vision, Graphics and Image Processing*, 46(1):82–95, 1989.

[19] Maia Zaharieva, Reinhold Huber-Mörk, Michael Nölle, and Martin Kampel. On ancient coin classification. In David Arnold, Alan Chalmers, and Franco Niccolucci, editors, *8th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST'07)*, pages 55–62. Eurographics, November 2007.