Recognizing Ancient Coins Based on Local Features^{*}

Martin Kampel and Maia Zaharieva

Pattern Recognition and Image Processing Group - TU Vienna Favoritenstr 9 - 11, 1040 Vienna, Austria 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. The recognition of ancient coins is a highly complex task that requires years of experience in the entire field of numismatics. To date, no optical recognition system for ancient coins has been investigated successfully. In this paper, we present an extension and combination of local image descriptors relevant for ancient coin recognition. Interest points are detected and their appearance is described by local descriptors. Coin recognition is based on the selection of similar images based on feature matching. Experiments are presented for a database containing ancient coin images demonstrating the feasibility of our approach.

1 Introduction

Numismatics is at a point where it can benefit greatly from the application of computer vision methods, and in turn provides a large number of new, challenging and interesting conceptual problems and data for computer vision. For coin recognition we distinguish between two approaches: coin identification and coin classification. A coin classification process assigns a coin to a predefined category or type, whereas a coin identification process assigns a unique identifier to a specific coin. What makes this application special and challenging for object recognition, is that all the coins are very similar.

The first coins were struck in Asia Minor in the late 7th century BC. Since then coins are a mass product [1]. In the Antiquity coins were hammer-struck from manually engraved coin dies. Coins from the same production batch will have very much the same picture and also the same quality of its relief. Depending on the series of coins in question, the only varying details can be either part of the picture or legend or there can be a difference in a prominent detail such as the face of a figure. The scientific requirement is to assign a coin its correct number according to a reference book.

^{*} 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.

G. Bebis et al. (Eds.): ISVC 2008, Part I, LNCS 5358, pp. 11–22, 2008.



Fig. 1. Different coins of the same coin type



Fig. 2. Different image representations of the same coin

Ancient and modern coins bear fundamental differences that restrict the applicability of existing algorithms [14]. Due to their nature ancient coins provide a set of identifying features. The unique shape of each coin originates in the manufacturing process (hammering procedure, specific mint marks, coin breakages, die deterioration, etc.). Furthermore, the time leaves its individual mark on each coin (fractures, excessive abrasion, damage, corrosion, etc.). Eventually, identification of ancient coins turns out to be "easier" compared to classification. For example, Figure 1 shows ten different coins of the same coin type. A classification algorithm should ideally classify them all of the same class. However, they all provide complete different characteristics (see shape, die position, mint marks or level of details). At the same time, exactly those features enable the identification process.

In contrast, Figure 2 presents five pictures of one and the same coin. The pictures were taken using different acquisition setups, i.e. scan as well as fixed and free hand cameras with varying lighting conditions. The figure points out the challenges for an automated identification process as well as the importance of quality images for the process itself. Different lighting conditions can hide or show details on the coin that are significant for a successful identification process (e.g. compare the first and the third image in Figure 2).

The remainder of this paper is organised as follows: In Section 2 related work on recognizing coins is presented. Section 3 gives an overview of local features with respect to our needs. The coin recognition workflow is described in Section 4. The experiments performed and their results are presented in Section 5. We conclude the paper in Section 6 with discussion on the results achieved and an outlook for further research.

2 Related Work

Research on pattern recognition algorithms for the identification of coins started in 1991 when Fukumi et al. [2] published their work on rotation-invariant visual coin recognition using a neural networks approach. Also [3] is devoted to neural network design, but investigates the possibilities of simulated annealing and genetic algorithms. In 1996 Davidson [4] developed an instance-based learning algorithm based on an algorithm using of decision trees [5]. An industrial implementation of a neural networks approach is described in [6].

A more recent neural algorithm was published in [7]. This approach employs the output of a filter bank of Gabor filters fed into a back propagation network. The algorithm uses correlation in the polarspace and in combination with a neural networks. Khashman et al. implemented a neural coin recognition system for use in slot machines [8].

Huber et al. present in [9] a multistage classifier based on *eigenspaces* that is able to discriminate between hundreds of coin classes. The *Dagobert* coin recognition system presented by Nölle et al. [10] aims at the fast classification of a large number of modern coins from more than 30 different currencies. In their system coin classification is accomplished by correlating the edge image of the coin with a preselected subset of master coins and finding the master coin with lowest distance.

In [11] Maaten et al. present a coin classification system based on *edge-based* statistical features. It was developed for the MUSCLE CIS Coin Competition 2006 [12] focusing on reliability and speed. The coin classification method proposed by Reisert et al. [13] is based on gradient information. Similar to the work of Nölle et. al [10] coins are classified by registering and comparing the coin with a preselected subset of all reference coins.

Current research approaches for coin recognition algorithms possess mainly two limitations. On the one hand, the input digital image is well defined – there is always only one coin pictured and the image is taken under very controlled conditions (such as background, illumination, etc.). On the other hand, the algorithms focus mainly on the recognition of modern coins. Those assumptions facilitate the classification and identification process substantially. In the case of controlled conditions and the well known circular shape of modern coins, the process of coin detection and segmentation becomes an easier task. The almost arbitrary shape of an ancient coin narrows the amount of appropriate segmentation algorithms. Tests performed on image collections both of medieval and modern coins show that algorithms performing good on modern coins do not necessarily meet the requirements for classification of medieval ones [14]. The features that most influence the quality of recognition process are yet unexplored.

3 Local Image Features

Local features describe image regions around given interest points. Their application in the computer vision is manifold ranging from object and texture recognition [15] to robot localization [16], symmetry detection [17] and wide

Edited by Foxit Reader 14 M. Kampe**Copylight(ه) های** Foxit Software Company,2005-2008 For Evaluation Only.

baseline stereo matching [18]. Local features are already successfully used for object classification. Crucial influence on local feature based object recognition bear both the detection of interest points and their representation. Hence, in the following we give a short overview over top performing interest point detectors and local feature descriptors and discuss their applicability with respect to the identification of ancient coins.

3.1 Interest Point Detectors

In the literature exist a broad number of interest point detectors with varying level on invariance against rotation, scale or affine changes. Comparative studies on interest points and their performance evaluation can be found in [19,20].

The Harris corner detector [21] is based on local auto-correlation matrix of the image function. The squared first derivatives are averaged over a 41×41 Gaussian weighted window around an image point. If the auto-correlation matrix has two significant eigenvalues, an interest point is detected. However, the detected points are not invariant to scale and affine changes. To achieve scale invariance Mikolajczyk et al. [22] extend the Harris detector by selecting corners at location where a Laplacian attains an extrema in scale space (Harris-Laplace). The Harris-Affine detector [22,19] additionally uses second moment matrix to achieve affine invariance. Detected points are stable under varying lighting conditions since significant signal change in orthogonal directions is captured.

Hessian-Laplace localizes points at local maxima of the Hessian determinant in scale-space maxima of the Laplacian-of-Gaussian [15,19]. Detected keypoints are invariant to scale and rotation transformations. Similar to Harris-Affine, the Hessian-Affine detector provides in a next step affine invariance based on second moment matrix [19]. In contrary to the Harris-based detectors, Hessian interest points indicate the presence of blob like structures. Bay et al. [23] introduced recently a further detector based on the Hessian matrix – the Fast-Hessian detector. It approximates Gaussian second order derivative with box filter. To further reduce the computational time, image convolutions use integral images.

Tuytelaars et al. present in [18] further two methods to extract affine invariant regions. The *Geometry-based region* detector starts from Harris corners and uses the nearby edges identified by the Canny edge operator [24] to build a parallelogram. Keypoints are detected if the parallelogram goes through an extremum of intensity-based functions. The second method proposed – *Intensity-based region* detector – relies solely on the analysis of image intensity. It localizes interest points based on intensity function along rays originating from local extrema in intensity.

The *Maximally Stable Extremal Regions* (MSER) proposed by Matas et al. [25] are a watershed based algorithm. It detects intensity regions below and above a certain threshold and select those which remain stable over a set of thresholds.

The *Difference-of-Gaussian*(DoG) detector was introduced by Lowe as keypoint localization method for the Scale Invariant Feature Transform (SIFT) approach [26,15]. Interest points are identified at peaks (local maxima and minima) of Gaussian function applied in scale space. All keypoints with low contrast or keypoints that are localized at edges are eliminated using a Laplacian function.

Edited by Foxit Reader Copyright(@)rbyAhaxit SoftwateeCompanyF2QQ5e:2008 15 For Evaluation Only.

Detector	Interest points
Difference-of-Gaussian (DoG) [26]	968
Harris-Laplace [15]	204
Harris-Affine [19]	198
Hessian-Laplace [19]	1076
Hessian-Affine [19]	778
Fast-Hessian [23]	198
Geometry-based region (GBR) [18]	61
IBR [18]	184
Maximally Stable Extremal Regions (MSER) [25]	134

 Table 1. Average interest points detected

Common critic to edge-based methods is that it is more sensitive to noise and changes in neighboring texture. Interest point detectors which are less sensitive to changes changes in texture perform well in a classification scenario since they recognize and capture those features that are common for all instances in a given class. On the opposite, identification relies on those features that are unique for a given object. Due to their nature and manufacturing process, ancient coins are unique. Coins produced by the same die show the same picture. However, since they are hand-hammered, shape, texture and relief can vary to a large degree. In this particular scenario, texture-sensitive interest point detectors are expected to perform better. Table 1 shows average interest points extracted per detector for the dataset explained in Section 5.

As we will show in Section 5, the methods which detect most interest points do not necessarily perform the best. First, we are faced with the problem of overfitting (i.e. each coin is similar to all the other coins to some degree). Second, essential role play the information captured per interest point. Thus, in the next subsection we give a short overview of the local feature descriptors we used for the experiments.

3.2 Local Feature Descriptors

Given a set of interest points, the next step is to choose the most appropriate descriptor to capture the characteristics of a provided region. Different descriptors emphasize different image properties such as intensity, edges or texture. Please refer to [27] for a thorough survey on the performance of local feature descriptors. We focus our study on four descriptors which show outstanding performance with respect to changes in illumination, scale, rotation and blur.

(1) Lowe [15] introduced the *Scale Invariant Feature Transform (SIFT)* descriptor which is based on gradient distribution in salient regions – at each feature location, an orientation is selected by determining the peak of the histogram of local image gradient orientations. Subpixel image location, scale and orientation are associated with each SIFT feature vector.

Edited by Foxit Reader 16 M. Kampe**Copylight(ه) های** Foxit Software Company,2005-2008 For Evaluation Only.

(2) Mikolajczyk and Schmid [27] propose an extension of the SIFT descriptor – Gradient Location and Orientation Histogram (GLOH) – designed to increase the robustness and distinctiveness's of the SIFT descriptor. Instead of dividing the path around the interest points into a 4×4 grid, the authors divide it into radial and angular grid. A log-polar location grid with 3 bins in radial and 8 bins in angular directions is used. The gradient orientations are quantized into 16 bins which gives a 272 bin histogram further reduced in size using PCA to 128 feature vector dimension.

(3) Belongie et al. [28] introduce *Shape Context* as feature descriptor for shape matching and object recognition. The authors represent the shape of an object by a discrete set of points sampled from its internal or external boundaries. As starting points, edge pixels as found by an edge detector. Following, for each point the relative location of the remaining points is accumulated in a coarse log-polar histogram.

(4) *Speeded Up Robust Features* (SURF) [23] are fast scale- and rotation invariant features. The descriptor captures distributions of Haar-wavelet responses within the neighborhood of an interest point. Each feature descriptor has only 64 dimensions which results in fast computation and comparison.

In [27] complementary evaluation on the performance of local descriptors with respect to rotation, scale, illumination, and viewpoint change, image blur and JPEG compression, is presented. In most of the tests SIFT and GLOH clearly outperformed the remaining descriptors: shape context, steerable filters, PCA-SIFT, differential invariants, spin images, complex filters, and moment invariants. In [29] Stark and Schiele report that the combination of Hessian-Laplace detector with SIFT and GLOH descriptor outperforms local features such as Geometric Blur, k-Adjacent Segments and Shape Context in a object categorization scenario. For their evaluation the authors used three different datasets containing quite distinguishable objects such as cup, fork, hammer, knife, etc. By contrast, our two coin data sets possess very different characteristics in comparison with existing evaluation and application scenarios. Both data sets contain similar objects and both are targeted to evaluate identification performance.

4 Recognition Workflow

We define the workflow for the identification of ancient coins by five well-defined stages as shown in Figure 3. In the *preprocessing* step (1) coins contained in the image are detected and segmented. Essential influence on the process carries the image diversity, e.g. single or multiple objects pictured, varying lighting conditions, shadows, diverse background textures, etc. In the scenario of ancient coins identification the almost arbitrary shape of coin additionally impede the task of coin(s) detection and segmentation. Since our test database consists solely of images of single coin on an unitary background no preprocessing is required. Eventually, the applied local feature detectors locate interest points on the background (e.g. due to intensity change). However, their amount is minimal and has no influence on the identification process.



Fig. 3. The five stages of coins identification workflow

The goal of the *feature extraction* step (2) is twofold. First, local features algorithms are applied to extract local image descriptors for coins identification. Second, features that can be used to reduce the number of required feature comparisons by reducing the coins database can be extracted. Provided uncontrolled acquisition process, simple features such as area or perimeter of a coin are not eligible since the scaling factor is unknown. Other features such as shape descriptors can be used as basis for step (3) - *preselection* step [30].

Step (4) descriptor matching is performed by identifying the first two nearest neighbors in terms of Euclidean distances. A descriptor D_1 is accepted only if the distance ratio of the nearest (1.NN) to the second nearest (2.NN) neighbors is less then or equal to 0.5:

$$2d(D_1, D_{1.NN}) <= (D_1, D_{2.NN}). \tag{1}$$

In [15] Lowe suggests a distance ratio of 0.8. However, our experiments showed that for the case of lower inter-class differences (as all classes are coins), a lower distance ratio tends to keep more distinctive descriptors while eliminating a great part of the false matches. The value of 0.5 was determined experimentally and used throughout the tests. Furthermore, we apply a restriction rule to fasten the quality of the matches. Since each image in the database picture is a single ancient coin, a given keypoint can only be matched to a single point in a different feature set. Thus, all multiple matches are removed as they are considered to be unstable for the identification process.

Finally, an additional *verification* step (5) can assure the final decision. Provided images of both obverse and reverse side of a coin, each side is first identified separately. If both sides vote for the same coin identification, the coin is identified adequately. Otherwise, it is classified as unknown.

5 Experiments

For our experiments we used a dataset of images acquired at the Fitzwilliam Museum in Cambridge, UK. We used varying technical setups – scan as well as fixed and free hand cameras, and varying lighting conditions. The dataset consists of 350 images of three different coin types (10 to 16 coins à coin type, 3 to 5 pictures à coin side). Ground truth is encoded in the file names. For testing the recognition one image was selected as test images. Presented evaluations as [27,29] on the performance of local descriptors use different datasets containing quite distinguishable objects such as cup, fork, hammer, knife, etc. By contrast, our coin data set possesses very different characteristics in comparison with existing evaluation and application scenarios. The data set contains similar objects and is targeted to evaluate coin recognition performance.

In a first experiment we compare the performance of three descriptors on coin identification.

Figure 4 shows corresponding interest points detected by the different approaches. Despite the lower image quality of the input image, the rotation and scale change of the coin, the SIFT approach matches correctly against image of the same coin acquired by the scan device (see Figure 4(a)).

The Fast Approximated SIFT approach – Figure 4(b) – tends to detect keypoints mostly on the background of the image. The algorithm detects far more points than SIFT, e.g. for the example input image 8999 keypoints (by contrast keypoints detected by SIFT for the same image: 721). However, they lack of stability and distinctiveness. Eventually, each detected interest point is similar (i.e. being matched) to a large number of keypoints in the second image. The elimination of multiple matched points reduces the number of final matches by approximately 90%.

Performing manual pairwise comparison of the resulting matches, PCA-SIFT (see Figure 4(c)) seems to achieve almost the same amount on descriptors as SIFT for less computational time. However, the stability of the PCA-SIFT features is considerably lower since approximately 40% of the correct classified



(a) SIFT

(b) Fast Approximated SIFT

(c) PCA-SIFT

Fig. 4. Example matches for a given ancient coin acquired using a free hand camera (Input image on the left side and corresponding match on the right hand). Using the SIFT approach (a), the test coins was successfully matched against an image of the same coin acquired using scan device. The Fast Approximated SIFT fail to recognize the image (b). PCA-SIFT (c) matched against different coin of the same coin type.

Interest Point	(1) SIFT		(2) GLOH		(3) Shape		(4) SURF	
Detectors	CR	IR	CR	IR	CR	IR	CR	IR
DoG	90.57%	84.57%	60.00%	40.00%	61.14%	29.14%	82.57%	28.57%
Harris-Laplace	68.39%	50.86%	71.84%	53.45%	79.71%	61.45%	71.30%	28.12%
Harris-Affine	76.15%	55.46%	73.56%	54.31%	73.04%	53.04%	71.88%	27.83%
Hessian-Laplace	65.90%	47.28%	65.90%	47.28%	92.57%	82.00%	84.29%	32.29%
Hessian-Affine	71.63%	50.72%	68.48%	49.28%	88.00%	80.00%	79.43%	29.71%
Fast-Hessian	85.43%	79.43%	85.43%	78.29%	84.86%	72.29%	90.86%	78.29%
GBR	51.47%	27.36%	48.53%	24.76%	52.44%	29.64%	56.03%	15.31%
IBR	80.29%	60.57%	75.71%	50.57%	80.29%	55.14%	77.43%	25.14%
MSER	80.86%	68.29%	77.71%	64.29%	74.86%	58.00%	74.00%	28.29%

Table 2. Evaluation results on the recognition performance of the local image feature descriptors using the small database of ancient coins. CR shows the rate of correctly *classified* coins, and IR those of correctly *identified* ancient coins.



Fig. 5. Performance distribution of the interest point detectors

images are due matching of the obverse with the reverse side of a coin. The PCA reduction of feature vector size seems to lead to loss of valuable information for the identification process. In terms of identification rate, SIFT clearly outperforms both modifications by more than 10%.

The second experiment aims at evaluation of the performance of the presented interest point detectors and local descriptors with respect to recognition. We compare both classification (CR) and identification rate (IR) and show that a good classification rate is no guarantee for the distinctiveness and stability of the respective detectors or descriptors. Table 2 summarizes the results on the coin data set. The best classification rate of 92.57% was achieved with Shape Context combined with Hessian Laplace detector. The best identification rate of 84.57% was achieved with SIFT combined with DoG. The main reason for the significant difference between classification and identification rate is the nature of local descriptors. Local descriptors simply describe the close surroundings of given interest point. Dependent on the size of this box, matching, i.e. similar enough, descriptor can be found on multiple coins or even on the same coin or different sides of the same coin.

Figure 5 visualizes the performance distribution with respect to the interest point detectors. One can clearly identify four groups. The first one, low identification and low classification rate, is dominated by the GBR detector. Independent of the applied local feature descriptor the achieved performance is too low with a rate close or far bellow 50%. The second group, high classification and low identification rate, is defined by the use of the SURF descriptor. Independently of the applied interest point detector, the SURF descriptor shows high stability with respect to classification. The last conspicuous group, high classification and high identification rate, is dominated by the Fast Hessian detector.

6 Conclusion

In this paper, we described a strategy for the recognition of ancient coins based on local image features. The achieved recognition rates indicate the feasibility of the approach. SIFT features show outstanding performance in existing evaluations. However, the main drawback and critical point is their computational time. Benefits of the proposed system are in the field of coin recognition. Based on the promising results we plan to extend the evaluation on a recently recorded coin collection of 2400 images of 240 different coins. Future research will include methods in the field of optical character and symbol recognition. Furthermore, we will extend our work towards die and mint sign identification based on spatially constrained local features.

Acknowledgment. The authors want to thank Dr. Mark Blackburn and his team at the Coin Department, Fitzwilliam Museum, Cambridge, UK, for sharing their experience and providing the database of images of ancient coins.

References

- 1. Duncan-Jones, R.: Money and Government in the Roman Empire. Cambridge (1994)
- Fukumi, M., Omatu, S., Takeda, F., Kosaka, T.: Rotation-invariant neural pattern recognition system with application to coin recognition. IEEE Transactions on Neural Networks 3, 272–279 (1992)
- Mitsukura, Y., Fukumi, M., Akamatsu, N.: Design and evaluation of neural networks for coin recognition by using ga and sa. In: Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks, IJCNN 2000, vol. 5, pp. 178–183 (2000)

- 4. Davidsson, P.: Coin classification using a novel technique for learning characteristic decision trees by controlling the degree of generalization. In: Proc. of 9th Int. Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems (IEA/AIE 1996), pp. 403–412 (1996)
- Aha, D., Kibler, D., Albert, M.: Instance-based learning algorithms. Machine Learning 6, 37–66 (1991)
- Moreno, J.M., Madrenas, J., Cabestany, J., Laúna, J.R.: Using classical and evolutive neural models in industrial applications: A case study for an automatic coin classifier. In: Biological and Artificial Computation: From Neuroscience to Neurotechnology, pp. 922–931 (1997)
- Bremananth, R., Balaji, B., Sankari, M., Chitra, A.: A new approach to coin recognition using neural pattern analysis. In: Proceedings of IEEE INDICON 2005, pp. 366–370 (2005)
- Khashman, A., Sekeroglu, B., Dimililer, K.: Rotated Coin Recognition Using Neural Networks. Advances in Soft Computing, vol. 41, pp. 290–297. Springer, Berlin (2007)
- Huber, R., Ramoser, H., Mayer, K., Penz, H., Rubik, M.: Classification of coins using an eigenspace approach. Pattern Recognition Letters 26, 61–75 (2005)
- Nölle, M., Penz, H., Rubik, M., Mayer, K.J., Holländer, I., Granec, R.: Dagobert

 a new coin recognition and sorting system. In: Proc. of the 7th International Conference on Digital Image Computing - Techniques and Applications (DICTA 2003), Macquarie University, Sydney, Australia, pp. 329–338. CSIRO Publishing (2003)
- van der Maaten, L.J., Poon, P.: Coin-o-matic: A fast system for reliable coin classification. In: Proc. of the Muscle CIS Coin Competition Workshop, Berlin, Germany, pp. 7–18 (2006)
- Nölle, M., Rubik, M., Hanbury, A.: Results of the muscle cis coin competition 2006. In: Proceedings of the Muscle CIS Coin Competition Workshop, Berlin, Germany, pp. 1–5 (2006)
- Reisert, M., Ronneberger, O., Burkhardt, H.: An efficient gradient based registration technique for coin recognition. In: Proc. of the Muscle CIS Coin Competition Workshop, Berlin, Germany, pp. 19–31 (2006)
- Zaharieva, M., Kampel, M., Zambanini, S.: Image based recognition of ancient coins. In: Kropatsch, W.G., Kampel, M., Hanbury, A. (eds.) CAIP 2007. LNCS, vol. 4673, pp. 547–554. Springer, Heidelberg (2007)
- 15. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision 60, 91–110 (2004)
- Murillo, A.C., Guerrero, J.J., Sagüés, C.: Surf features for efficient robot localization with omnidirectional images. In: IEEE International Conference on Robotics and Automation, pp. 3901–3907 (2007)
- Loy, G., Eklundh, J.O.: Detecting symmetry and symmetric constellations of features. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3952, pp. 508–521. Springer, Heidelberg (2006)
- Tuytelaars, T., Van Gool, L.: Matching widely separated views based on affine invariant regions. International Journal of Computer Vision 59, 61–85 (2004)
- Mikolajczyk, K., Tuytelaars, T., Schmid, C., Zisserman, A., Matas, J., Schaffalitzky, F., Kadir, T., Van Gool, L.: A comparison of affine region detectors. International Journal of Computer Vision 65, 43–72 (2005)
- Schmid, C., Mohr, R., Baukhage, C.: Evaluation of interest point detectors. International Journal of Computer Vision 2, 151–172 (2000)

Edited by Foxit Reader

M. Kampel Copwight Greby Foxit Software Company, 2005-2008 For Evaluation Only.

- 21. Harris, C., Stephens, M.: A combined corner and edge detector. In: Alvey Conference, pp. 147–152 (1988)
- 22. Mikolajczyk, K., Schmid, C.: Scale & affine invariant interest point detectors. International Journal of Computer Vision 60, 63–86 (2004)
- 23. Bay, H., Tuytelaars, T., Van Gool, L.: SURF: Speeded up robust features. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3951, pp. 404-417. Springer, Heidelberg (2006)
- 24. Canny, J.: A computational approach to edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence 8, 679–698 (1986)
- 25. Matas, J., Chum, O., Urban, M., Pajdla, T.: Robust wide baseline stereo from maximally stable extremal regions. In: Proceedings of the Britisch Machine Vision Conference, London, vol. 1, pp. 384–393 (2002)
- 26. Lowe, D.G.: Object recognition from local schale-invariant features. In: International Conference on Computer Vision (ICCV 1999), Washington, DC, USA, vol. 2, pp. 1150–1157. IEEE Computer Society, Los Alamitos (1999)
- 27. Mikolajczyk, K., Schmid, C.: A performance evaluation of local descriptors. IEEE Transactions on Pattern Analysis and Machine Intelligence 27, 1615–1630 (2005)
- 28. Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. IEEE Transactions on Pattern Analysis and Machine Intelligence 24, 509–522 (2002)
- 29. Stark, M., Schiele, B.: How good are local features for classes of geometric objects. In: 11th International Conference on Computer Vision (ICCV 2007), Rio de Janeiro, Brazil (2007)
- 30. Zaharieva, M., Huber-Mörk, R., Nölle, M., Kampel, M.: On ancient coin classification. In: Arnold, D., Chalmers, A., Niccolucci, F. (eds.) 8th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST 2007), Eurographics, pp. 55–62 (2007)

22