

Universidade do Minho

Escola de Engenharia

Departamento de Informática

Luís Salgado

**Medieval coin automatic recognition by
computer vision**

**Reconhecimento automático de moedas medievais
usando visão por computador**

July 2016



Universidade do Minho

Escola de Engenharia

Departamento de Informática

Luís Salgado

**Medieval coin automatic recognition by
computer vision**

**Reconhecimento automático de moedas medievais
usando visão por computador**

Master dissertation

Master Degree in Engenharia Informática

Dissertation supervised by
Professora Cristina Santos

July 2016

ACKNOWLEDGEMENTS

***To my parents,
who gave me the curiosity gene.***

***To my wife,
who patiently puts up with it.***

I would like to thank Cristina P. Silva for all the support, the precious guidance and the enthusiastic encouragement along this research, without her it would be so much harder to get this far. I am also very grateful to Pedro Henriques, who gave me the opportunity to accomplish this project and was always understanding towards my clumsiness in the university protocol.

Some numismatists promptly answered my appeal for categorized images of Portuguese medieval coins and respective permissions of use, which allowed the foundation of an image dataset on this subject. Among these are José Almeida, Paulo Silva, Mário Matos and Alberto Praça.

Finally, I have to both thank and apologize to my wife Clara and my kids for all the time I stole from family time (and, I confess, from household tasks too) to be able to conclude this research. To Clara I owe the revision of this text and its cleansing from the uncountable grammar mistakes.

To all, my enormous gratitude. This work wouldn't be possible without you.

ABSTRACT

Keywords: Computer vision, Image classification, Image segmentation, Machine Learning, Coins

The use of computer vision for identification and recognition of coins is well studied and of renowned interest. However the focus of research has consistently been on modern coins. The same algorithms that advertise a hit of 97% in recognition of modern coins are quite disappointing when applied to ancient coins.

This discrepancy has several reasons, starting with the nature of ancient coins, that are manually minted, having plenty variances and failures, frequent ripples and centuries of degradation which further deform the characteristic patterns, making their identification a hard task even for humans.

Another noteworthy factor in almost all similar studies is the controlled environments and uniform illumination of all images of the datasets. Though it makes sense to focus on the more problematic variables, this is an impossible premise to find outside the researchers' laboratory, therefore a problematic that must be approached.

This dissertation focuses on medieval and ancient coin recognition in uncontrolled "real world" images, thus trying to pave way to the use of vast repositories of coin images all over the internet that could be used to make our algorithms more robust.

In the first part of the dissertation I propose a fast and automatic method to segment ancient coins over complex backgrounds using a Histogram Backprojection approach. Results are compared against a proposed automation of GrabCut algorithm, and its benefits are demonstrated. Although the present report is oriented to ancient coin segmentation, the method can also be used in other contexts presenting thin objects with uniform colors.

In the second part, several state of the art machine learning algorithms are compared in the search for the most promising approach to classify these challenging coins.

RESUMO

Palavras-Chave: Visão por computador, Classificação de Imagens, Segmentação de imagens, *Machine Learning*, Moedas

O uso de visão por computador para identificação e reconhecimento de moedas é bastante estudado e de reconhecido interesse. No entanto o foco da investigação tem sido sistematicamente sobre as moedas modernas. Os mesmos algoritmos que anunciam um sucesso de 97% no reconhecimento de moedas modernas são bastante desapontantes quando aplicados a moedas antigas. Esta discrepância de resultados tem várias justificações, a começar pela natureza das moedas antigas que, sendo cunhadas à mão, apresentam bastantes variações e falhas, frequentes ondulações na sua superfície e séculos de degradação que deformam ainda mais os padrões característicos, tornando a sua identificação uma tarefa ingrata mesmo para o ser humano. Um outro fator digno de nota na quase totalidade dos estudos desta problemática é o uso de ambientes controlados e iluminação uniformizada entre todas as imagens dos *datasets*. Embora faça sentido focar-se nas variáveis mais problemáticas, esta é uma premissa impossível de encontrar fora do laboratório do investigador e portanto uma problemática que tem que ser estudada.

Esta dissertação foca-se no reconhecimento de moedas medievais e clássicas em imagens não controladas, tentando assim abrir caminho ao uso de vastos repositórios de imagens de moedas disponíveis na internet, que poderiam ser usados para tornar os nossos algoritmos mais robustos. Na primeira parte da dissertação, proponho um método rápido e automático para segmentar moedas antigas sobre fundos complexos, numa abordagem que envolve *Histogram Backprojection*. Os resultados são comparados com uma automação do algoritmo *GrabCut* e são demonstradas as mais-valias do método proposto. Embora este relatório seja orientado para a segmentação de moedas antigas, este método pode ser usado noutros contextos que incluam objetos planos de cor uniforme. Na segunda parte, o estado da arte de *Machine Learning* é testado e comparado em busca da abordagem mais promissora para classificar estas moedas complicadas.

CONTENTS

ACKNOWLEDGEMENTS.....	III
ABSTRACT.....	IV
RESUMO.....	V
CONTENTS.....	VI
LIST OF FIGURES.....	VIII
LIST OF TABLES.....	VIII
INTRODUCTION.....	1
STATE OF THE ART.....	3
2.1 Image recognition evolution	3
2.2 Coin recognition	4
2.2.1 Of modern coins	4
2.2.2 Of ancient coins.....	6
2.2.3 Other approaches	10
2.3 Other relevant studies.....	12
THE PROBLEM AND ITS CHALLENGES	13
3.1 The problem is that there are many problems	13
3.1.1 Classification	13
3.1.2 Identification.....	14
3.1.3 Style.....	16
3.2 And even more challenges.....	17
3.2.1 Medieval coins are evil.....	17
3.2.2 Ordinary images are not role models.....	17
3.2.3 Big data is hardly an option (but big contribution is not).....	18
3.3 Proposed approach – solution	19
3.3.1 The goal	19
3.3.2 The approach	19
3.3.3 The framework and test environment.....	20

THE METHOD: I – SEGMENTATION	21
4.1 Introduction.....	21
4.2 Related work	22
4.3 Proposed method	23
4.3.1 Step A.1 – Histogram Backprojection	25
4.3.2 Step A.2 – Border approximation by Convex Hull.....	27
4.3.3 Step B.1 – Second stage – Canny edge.....	28
4.3.4 Step B.2 – Refined border.....	28
4.3.5 Automatizing the method	29
4.4 Experiments and results	31
4.4.1 The datasets.....	31
4.4.2 Comparing method - GrabCut.....	32
4.4.3 Results	33
4.5 Discussion.....	34
4.6 Conclusions.....	35
THE METHOD: II – FEATURE EXTRACTION & RECOGNITION	36
5.1 Introduction.....	36
5.2 Proposed method	37
5.2.1 Overview.....	37
5.2.2 Feature Extraction.....	39
5.2.3 Feature organization	40
5.2.4 Machine Learning.....	42
5.3 Experiments and results	46
5.3.1 The dataset.....	46
5.3.2 The tests.....	47
5.3.3 Results and Discussion	48
5.3.4 Conclusion.....	50
FINAL CONCLUSIONS.....	51
6.1 Conclusions.....	51
6.2 Prospect for future work	51
BIBLIOGRAPHY	53
SUPPORT MATERIAL	60
A.1 Notes on examined methods on segmentation chapter.....	61

LIST OF FIGURES

Figure 1 - Proposed method for segmentation.	24
Figure 2 - Original image with rectangle model (left), result from HB (right).	26
Figure 3 - Histogram from HB result with a red line on the estimated ideal threshold value.	30
Figure 4 - Overview of the proposed classification method	38
Figure 5 - Simple example of discrimination on SVM.....	43
Figure 6 - A Decision Tree	44
Figure 7 - Sample images on the dataset.....	46
Figure 8 – Chart of the classification tests	48

LIST OF TABLES

Table 1 - Summary of empirical evaluation results	33
Table 2 - Computation times on empirical evaluation	34
Table 3 - Summary of Machine Learning results	49
Table 4 - Impact of preprocessing over de training methods with SIFT descriptors.....	49

INTRODUCTION

The use of computer vision for recognition of coins is well studied and of renowned interest in areas ranging from vending machines to fakes identification. In fact the problem is so relevant that it has been object of financing by the European Union (e.g. Project COINS [1], [2]) and international competitions for researchers (e.g. Muscle, 2006 and 2007). However the focus of research and testing has consistently been on modern coins, under controlled lighting conditions. The focus on ancient and medieval coins could provide valuable tools for archeologists' rapid finding identification; museum curators could be more easily guided on cataloging their collections, or identifying fake or stolen coins; and perhaps historians could find unsuspected relations between different coins from different areas, consequently establishing cultural relations.

The challenge is that the same algorithms that advertise a hit of 97% in recognition of modern coins are quite disappointing when applied to ancient coins.

The discrepancy in these results has several reasons, starting with the nature of ancient coins, that are hand minted, having plenty variances and failures. For example, if we take two pieces of 1 euro (from the same country) they are totally identical and practically indistinguishable from each other, in the case of ancient coins no two are alike copies (if they are equal, surely one of them is a fake). The well-defined standards and uniform thicknesses of modern coins do not exist in ancient coins, which present frequent ripples even within the same figure, not centered or incomplete designs, or irregular shapes. On top of all this are the softer metal alloys and centuries of degradation which further deform the characteristic patterns, making their identification a hard task sometimes even for humans. When we move from the classical coins (like Roman and Greek) to medieval coins, the task gets even harder, due to the severe declination in the artistic quality from artisans in the Middle Ages.

Another noteworthy factor in almost all similar studies is the controlled environments and uniform illumination of all images of the training and testing datasets. This choice makes sense, taking into account the amount of other more problematic variables that

must be managed in research algorithms to achieve minimally encouraging results. However this is an impossible premise to find in the real world, where the photographs that populate shops, auction houses, numismatic forums, archeology offices and sometimes even museums are the most disparate imaginable and often with awful lighting conditions. Thus, any methodology that aims to be applied to the world outside the researcher's laboratory will necessarily have to consider this disparity.

In an age when image recognition researchers seem to be battling over a few decimal points in success rates, working on medieval coins, especially in such hard conditions, does not foretell to be promising at all, thus, this is not a very appellative area to bet our time and careers on. Nevertheless, the technology is achieving a point that could already start to be used in the real world helping researchers in other areas (like archeologists, museum experts and numismatists) forwarding our global knowledge. This work aims to give a modest contribution to such advance.

STATE OF THE ART

2.1 IMAGE RECOGNITION EVOLUTION

From the primordial image matching to the hype on Convolutional Neural Networks, the evolution of image recognition was not always easy, but also not immune to fashion. In a much simplistic way, we could describe its progress more or less like this:

In 1999 Lowe [3] introduced Scale Invariant Feature Transform (SIFT) and researchers from all around the world praised the new messiah and changed the course of their works. SIFT allows a point inside an RGB image to be represented robustly by a low dimensional vector and find a way to be invariant to scaling and rotation, partially invariant to illumination changes and robust to local geometric distortion, a much welcome evolution from the previous ordinarily used Sum of Squared Distances (SSD).

The Local Descriptors were on the rise when it became clear the internet would be the future for everything digital. The images were going online and easily grown into large-scale datasets. No longer were the researchers confined to their laboratories with their own small datasets. Rapidly the big problem became “my dataset is bigger than yours”.

More ambitious object recognition problems were arising and raw SIFT was not enough to deal with deforming objects or occlusions. New methods had to be found. Visual Words (VW) were a clever way of applying the same principles from text matching to visual content. It can be represented by small parts of an image which carry some kind of information related to the features (such as the color, shape or texture), or changes occurring in the pixels, such as deformations or missing areas. Another popular idea arose at the time: the need to have some sort of binning structure for matching objects. Grids were initially placed around entire images, and later on they would be placed around object bounding boxes. Methods like Pyramid Match Kernel introduced powerful and hierarchical ways of integrating spatial information into the image matching.

A new feature descriptor arrived around 2005, and started a new reign: the Histogram of Oriented Gradients (HOG), based on the counting of occurrences of gradient orientation in localized portions of an image, came to simplify and greatly improve the processing speed. HOG and a new machine learning tool called Support Vector Machine (SVM) easily gained acceptance. A later technique called the Deformable Parts-based Model (DPM), helped reinforce even more the popularity and strength of the HOG technique.

As datasets became larger, an old and discredited method rose again (this time with many folds the computation power from the 80s): Deep Learning. It attempts to model high-level abstractions in data by using multiple processing layers with complex structures, or otherwise composed of multiple non-linear transformations. Suddenly everyone seems to be jumping into the Convolutional Neural Networks (CNN) bandwagon, Machine Learning is fusing with Artificial Intelligence, and the word of the day is BigData... even if the problem is small.

2.2 COIN RECOGNITION

2.2.1 Of modern coins

One of the first major advances in coin recognition was given by the Dagobert [4] project¹. Its purpose was to sort high volumes of modern coins. The coins were already singled out and put on a conveyor belt where a camera observed one coin at a time in ideal lighting conditions. The method relied on binarized edge information that was correlated with all possible master edge images stored in a database, finding the master coin with lowest distance. For edge information they used Canny edge operator and Laplacian of Gaussian (LoG), plus a polar coordinate representation, but also sensor information of coin diameter and thickness. The success rate was high but the use of sensors and such very controlled conditions make this method infeasible for the purpose of this work. The massive coin image database was later available to the public, and it still provides the best dataset existent for modern coin recognition.

¹ A presentation video of the Coin Classification Machine is available in www.youtube.com/watch?v=I6JiD2yEi8Q

Huber et al. [5] proposed the use of Eigenspaces in modern coin recognition. The method consisted in a preprocessing, performed to obtain a translationally and rotationally invariant description, followed by a second stage, in which an appropriate Eigenspace was selected.

The MUSCLE CIS Coin Competition² in 2006 and 2007 launched new ground in coin recognition investigation. The big winner in 2006 was Reisert et al. [6], who used gradient based orientations in order to achieve the most effective success rate. They segmented the coin from the background by applying the Hough transform, then normalized the region containing the coin and transformed it to polar coordinates. An angular image was then computed based on the image gradient orientations. The similarity between two different coins was computed by counting the number of pixels with which the two respective angles coincide. This similarity measure was fed to a Nearest Neighbor classifier that would find the best-matching coin within a given coin image database. Reisert would later improve the method [7]. The runner-up, called *COIN-O-MATIC* [8], also persists in paper citations. This one focused on reliability and speed, relying on the coin edge information, but also on sensor information. The Edge angle-distance distributions were calculated and classified using the Nearest Neighbor approach. In 2007 Maaten showed up again with a paper over partially occluded coins [9]. For that he used Texon-based texture classifiers and template matching based on gradient orientations. But Zaharieva paper [10] tested a promising new area for coin recognition: Scale-Invariant Feature Transform (SIFT) (originally presented by Lowe [3]), which would prove much more effective in ancient coins.

Other authors experienced new approaches around the same time. Neural Networks were tested by several researchers, like Khashman et al. [11]. Ghanem et al. [12] tried a Gabor wavelet approach for feature extraction followed by Nearest Neighbor classification.

Some later, if more modest, contributions in the modern coin recognition, are Märtens et al. [13], who used Cross-correlation (CC) matching (while admittedly CC is not invariant to imaging scale, rotation, illumination and perspective distortions, the authors claim that normalization of the CC can significantly improve the method); Vadivelan et al. [14], who

² Information and data from the competition in: muscle.caa.tuwien.ac.at/coin_past.php

experimented both Gabor wavelet and LBP operator features over several distance measurement methods and a Nearest Neighbor classifier; and Wei et al. [15], who presented an approach based on image textures, using Ant colony optimization (ACO) for optimal threshold segmentation and Tree-structured Wavelet Transform (TWT) for the Textural Characteristics Extraction.

An interesting survey on Techniques of Coin Detection and Recognition is presented by Mehta et al. in [16]. But, by that time, with all the previous studies, and optimistic success rates, the problem of coin recognition seemed all but solved, to most researchers.

Yet, these same methods provided very frustrating success rates when applied to ancient coins, proving that the techniques developed for modern coin classification are not sufficient for ancient coin classification [17], [18], [19]. The reasons were already discussed in the introduction, but some basic assumptions also changed. For instance, the generally accepted use of Hough transforms for segmentation of the image is no longer appropriate because very frequently the coins are not perfectly round. Also the controlled lighting conditions are harder to achieve due to the degradation and deformation of the ancient coins. And, of course, the use of sensors does not make much sense when dealing with just a few specimens instead of thousands, or in industrial/commercial applications.

2.2.2 Of ancient coins

Some attention to ancient coin recognition came from the EU sponsored COINS (Combat On-Line Illegal Numismatic Sales) project [1], [2], focusing on fake and stolen coins identification. Its main approach relied on individual, unique features, which make a specimen different from all other individuals in the same class. For over two years, a good number of studies came from this project. The team started with a good analysis of the problem [20], provided a dataset of ancient coins from the collection of Fitzwilliam Museum in Cambridge, and have shown again the potential of SIFT classification. The various problematics were approached in distinct papers. The segmentation problem was discussed in [21] and [22] suggesting the use of local entropy and a local range of grey values, but still facing the big problem of the coin shadows. The image acquisition was also tested in [23], towards an optimized acquisition process. Several books [24], [25], and internet articles already existed on the subject, but the study made sense since these

are focused on the aesthetics and not on the recognition optimization. Nevertheless the investigation continued to be on controlled images. Some possible numismatic research fields were debated in [26], paving the way for more investigation in the area.

The team presented in [27] an end-to-end coin identification workflow for ancient coins, with decent results. They cross-evaluated the performance of several Interest Point Detectors [Difference-of-Gaussian (DoG), Harris-Laplace, Harris-Affine, Hessian-Laplace, Hessian-Affine, Fast-Hessian, Geometry-based region (GBR), IBR, and Maximally Stable Extremal Regions (MSER)] with different local image feature descriptors for coin classification and recognition [SIFT, Gradient Location and Orientation Histogram (GLOH), shape context and Speeded Up Robust Features (SURF)]. Based on the irregular shape of ancient coins, in [17] they introduced the use of a deviation from circular shape matching (DCSM) as a form of identification (not classification), a method they would use in several of their later papers, under the argument that the outline is a unique characteristic of a coin. While this is true and useful when identifying stolen coins, relying on it for the identification of fake coins shows some unawareness on the refined methods used by the fakers. Other problems with the concept are the lighting conditions, which admittedly could influence the results, and the computation time which, according to a newspaper article³, took a few minutes.

The fact that the online tools from COINS project come public provided much more feedback about its effectiveness. Some reviews claimed that the results were variable⁴, and the segmentation was not impressive, exposing the difficulties of using “real world” photos. By now the internet site along with the online tools have vanished⁵, leaving the big public with nothing to work with once again.

Meanwhile a limited number of researchers add their contribution, too. Arandjelović [28] introduced a new concept based on localized analysis rather than taking the coin as a whole. His method is based on a feature he called Locally Biased Directional Histogram (LBDH). For each interest point found by the Difference-of-Gaussian (DoG) detector, a set

³ The article can be read online in diepresse.com/home/techscience/wissenschaft/364117/Keine-zwei-Munzen-gleichen-einander-zu-100-Prozent?_vl_backlink=/home/techscience/index.do.

⁴ e.g. digitalhn.blogspot.pt/2009/06/software-from-coins.html.

⁵ Although there is still a page about the project in oldwww.prip.tuwien.ac.at/research/completed-projects/coins.

of weighed and directed histograms is computed. These features aim to capture geometric relationships between interest points. This proved a promising kind of feature as it offers a powerful representation able to capture the class-specific coin appearance. The method achieved 52.7% classification accuracy, largely outperforming a histogram of SIFT representation, but mainly it opened the eyes of the community for the need to look at special relations between the interest points.

Allahverdi et al. [29], [30] tested some already known methods on Sassanid coins, which are relevant for this project for being closer to medieval coin's style than the Roman coins used in most investigations. One of the papers explored the Discrete Cosine Transform (DCT), the other used Principal Component Analysis (PCA), plus Bhattacharyya distances between the coefficients vector and those representing each training coin. Both methods presented interesting results but were meanwhile outperformed by later studies. The same kind of coin was later approached by Parsa et al. [31], using a representation of the coin image based on the phase of the 2-D Fourier Transform (FT) of the image so that the adverse effect of illumination was eliminated. Then, a Bi-Directional PCA (BDPCA) approach was used and an entry-wise matrix norm calculated the distance between two feature matrices so as to classify coins.

From 2011 to 2014 the ILAC⁶ [32] research project joined again (at least) some of the researchers from COINS project for a new batch of studies. The previous approaches were extended by Huber-Mörk et al. [33] using a preselection step based on the coin's contour. In this step, equally spaced rays are cast from the coin's center of gravity and intersected with its contour. The distances along the rays between these intersection points and the hypothetical perfect circle fitted to the coin area are measured and form a descriptor that can be computed quickly. This descriptor can be quickly matched and allows for fast pruning of large coin databases when attempting to identify a specific coin from an image. The second stage uses preselection by the first stage in order to refine the matching using local descriptors, and the results are combined using naïve Bayesian fusion.

The multi-stage method was also used in classification in [34]. The hierarchical sub-selection scheme showed that the classifier-free classification time could be reduced to

⁶ The project site can be found in www.caa.tuwien.ac.at/cvl/project/ilac/.

one-seventh without a loss of classification accuracy, a crucial gain. This was an improvement to a previous paper [35] defending the use of a classifier-free approach, like SIFT flow matching, with the main benefit of making us less dependent on the availability of a large and representative set of training images. This makes total sense when talking about scarce ancient coins. On the other hand the computation times, even using the hierarchical sub-selection scheme, are big and get even bigger as the dataset grows. In their tests, with 180 images, from 60 classes, the average classification time varies between 7 and 472 seconds. With a bigger dataset the times would become too uncomfortable for general public use. The classifier-free and multi-stage viewpoint was maintained later in [36], with computation times around 22s in the same 60-class problem. This time their approach was a data-driven first-order matching and used geometric constraints afterwards to reason about the geometric plausibility of the correspondences found. They also opted out SIFT descriptor in favor of Local Image Descriptor Robust to Illumination Changes (LIDRIC), a descriptor presented in [37] by the same team.

Again Huber-Mörk et al. [38], brought back edge features, in a comparison analysis between two approaches for classification and identification of coins: a method based on matching edge features in polar coordinates representation (as in [4]) and a method for matching based on an Eigenspace representation (as in [5]). Interesting results were achieved for identification with the Eigenspace method using deviation from circular shape matching (DCSM) and SIFT (again, with the limitations already commented on).

Anwar et al. [39] and [40], used Bags Of Visual Words (BoVW) – also called Bags of Features (BoF) – based on densely extracted local features such as SIFT, with spatial information, to propose a new method for classification by recognizing motifs minted on their reverse sides. The dense sampling for BoVW results in a better classification rate as it is capable to capture the underlying geometry of the motif even if some of its parts are missing.

Inevitably, the glamour of Convolutional Neural Networks (CNN) would prove irresistible to try on this field too, a good example is the work of Kim et al. [41], concluding that CNN outperforms Support Vector Machine (SVM), but presenting success classification rates no better than the previous papers. There are hundreds of new papers on the subject

every year (unfortunately not explicitly on coins) which makes it hard to keep track, but a very good example of what is being made is presented by Szegedy et al. [42].

2.2.3 Other approaches

Different approaches have been tried which, even if not applicable in the current study, are very interesting and show that creativity in new methodologies shall not be restrained.

Facial recognition

The fact that a great number of Roman coins display the rulers' bust or some deity inspired a few authors to use some form of adapted facial recognition approaches in coin recognition. Kim et al. [43], [44] approaches the subject with a method based on discriminative Deformable Part Models (DPM).

Character recognition

Most coins have legends, so the use of character recognition to classify a coin or at least to help in the process seems a logical next step. Arandjelović [45] concentrated on Roman Imperial denarii, which have uniform legends, and after a geometric normalization of the text (through the use of polar coordinates) applied a HOG-like descriptor for letters. In spite of being a very interesting approach, it is limited to a very specific niche of coins, presenting difficulties in the presence of Roman numerals and being unable to deal with legends which are not arranged along the border.

The ILAC project team addressed the same problem in [46], [47] and [48], both based on the work of Wang et al. [49]. The latter allows both straight and curved words (an important feature in coins), but requires the text to have a certain size relative to the coin image size, as the SIFT descriptors are only computed at one relative size. In [50] the same team finally combines image matching and the recognition of the coin legends (using the same methods already discussed), in order to improve the robustness of image-based coin classification.

In any case, and as interesting as the ability of character recognition may seem, we can extract information from the legend only if we are given a very well-preserved coin, with very well defined legends and no other similar features complicating the "reading".

When we deal with medieval coins, we know that it is hardly an option. In fact the characters are usually so rough and worn that they are difficult to read even by the human eye. Another issue is the evolution of the alphabet styles (Latin, Uncial, Gothic, Medieval, among others) and all its regional and period variants, and sometimes the use of several of these styles in one single coin, which gives the problem a much higher dimensionality than the above approaches suggest.

3D models

The use of 3D images for recognition or identification was approached in a hybrid way by Marchand et al. [51], [52] in which several photos were taken of the same coin with different lighting directions, in order to make a model resilient to the lighting conditions.

Huber-Mörk et al. [38] discussed the advantages of using 3D data for surface analysis. Once again, they aim to avoid the interference caused by lighting variations, like shadows, or highlights due to specular reflections that distort the features of the coin. However 3D acquisitions are more laborious and expensive and, to our knowledge, 3D vision approaches applied to 3D databases of coins do not exist at the moment.

Measurement data

We have already seen that some works used sensors to retrieve more valuable data, helping in the classification process and computation time. Yet the use of sensors is not the only way to go: Herrmann et al. [53] showed a method for retrieving measurements on coins by means of a ruler placed next to the coin when taking the photo. As interesting as it is, it implies that we have control over the photo taken and thus it is not applicable in the current study.

Iconography

The project DIANA [54] avoids the usual emphasis on classification or identification, and gives us an interesting new function: it recognizes the iconography on the coin in order to map its origin. In fact the iconography on the (loosely labeled) Greek coins are very specific to the states which minted the coins, thus an analysis of its figures can give us its origins.

2.3 OTHER RELEVANT STUDIES

Even if not specifically related to coin recognition or identification, some state of the art studies deserve to be mentioned for the potential they present to new approaches applied to coins in this or future works.

The light/shadow variation problems are addressed by Kwatra et al. [55], who present an apparently very good method to remove shadows from images. Unfortunately the process appears to be closed and patented. Guo et al. [56] presented another method, but the tests carried out in coins revealed a computation time unaffordable for this work.

Chen et al. [57] introduce the Logarithmic Total Variation (LTV) model and explain the way it removes varying illumination for face images. Although it is applied to face recognition, it could be applied in the recognition of busts in ancient coins, or major coin features.

The segmentation problem is approached in [58], where Arbeláez et al. examine the effect of multiple local cues combined into a globalization framework based on spectral clustering. Rother et al. [59] launches the concept of GrabCut, which extends the graph-cut approach by means of an iterative version of the optimisation and a robust algorithm for “border matting”.

Edge detect advances were made by Dollár et al. [60] with the use of Structured Random Forests, capable of real time frame rates (faster than most competing state of the art methods) while achieving state of the art accuracy.

Law’s Texture Energy Measure (TEM) was used in [61] on the subject of butterflies, but could as well be explored on coins.

Even if SIFT and Speeded Up Robust Features (SURF) (SIFT’s faster version) are the most well-known descriptors around, every year several more algorithms appear competing for accuracy or faster performance. A good example is given by Takacs et al. [62], presenting a Radial Gradient Transform (RGT) and a fast approximation: the approximate RGT (ARGT), which is incorporated in Rotation-Invariant Fast Feature (RIFF). They demonstrate that using the ARGT, RIFF extracts features 16× faster than SURF, while achieving a similar performance for image matching and retrieval. Other more generalized methods shall be presented in a later chapter.

THE PROBLEM AND ITS CHALLENGES

3.1 THE PROBLEM IS THAT THERE ARE MANY PROBLEMS

By recognition we usually mean the acknowledgement of something relevant, either by realizing something as existing or previously known, or by finding some useful relation to some other known thing. So when we say coin recognition, we are entering a multiple war front (or at least we should).

3.1.1 Classification

This is the most widely studied problem and probably the one that once presented to the public would be more acclaimed and criticized.

The act of classification consists of putting the coin successfully into a group of similar coins. These distinct groups are already established by generations of numismatists and sometimes it is not very easy to understand the criteria on which they are based. Taking into account that a group (or class) of ancient/medieval coins is not as uniform as a class of modern coins, an automatic classifier based on pictures of one's coin would be gratefully received by the whole numismatic and archeological community (as it would be lapidated for its lapses).

The challenges

The fact that they are hand minted and submitted to centuries of degradation and deformation makes it virtually impossible to find two equal coins. Their differences may be bigger due to malformation than due to class distinction. In fact the intra-class variations may sometimes stand out more than inter-class variations, it is all about conventions. To that we must add the abundant noise and deformation caused by degradation, or improper minting. So grouping coins in classes is more an approximation process than a complete match, more about the symbolism present on the coins than the visual aesthetics.

3.1.2 Identification

The uniformity and high levels of quality control on modern coins make the task of identifying a fake very hard for image recognition, unless we know beforehand which inaccuracies we are looking for. But the same factor that makes the ancient and medieval classification so hard - there are no two equal coins – is a big advantage on ancient coin identification. Identification of a specific stolen coin, for example, is a matter of simple matching. So theoretically, we are technologically fit to pick an image from eBay or some other auction house and compare it to a database of stolen coins (that is, if every museum and police force could be persuaded to cooperate). Moreover a fake coin is often made from a mold of an existing coin, so sometimes it is also possible to identify it. Yet, in this case it is not as easy as it sounds, as we will see below.

The challenges

In practice, the problem of identification is the same as classification, except that in this case each class is an individual coin, thus the volume of information and computation time will be inevitably bigger. The team from COINS project [1], [2], and later ILAC project [32] found a shortcut: instead of processing the whole coin, they concentrate only on the coin's edge. The principle is that if we could have a perfect enough segmentation from the coin's edge, we would have a faster, smaller and still accurate descriptor. The challenge is that a perfect enough outline has proven sometimes hard and it takes a long time. But the really tricky part is that the lighting conditions can severely influence the resulting extracted coin shape and the perfectly controlled lighting conditions of the dataset images can hardly be found in the real world.

Relying on the coin's edge for identification of fake coins has a very limited effectiveness since modern fakers use much more refined methods than simple copying.

Fake identification becomes a different challenge from stolen identification, in fact, three different challenges:

Falsifications by casting, the easiest and most common, are made using a mold from an existing coin, so they present a large similarity with the original coin. Yet, a mold has to have an opening to pour the metal in, in fact it has a second opening on the opposite

side so the air doesn't get trapped. So we have already two discontinuities that can represent around 10% to 30% of the coin edge (depending on the coin size). In some cases, instead of one mold, they use two, one for each side of the coin, and then glue both halves and all the edge of the coin is filled off to smooth it (so much for coin edge identification). In both cases, as the molds are used a second, a third or even more times, some imperfections start to appear shifting the fake from the original even more.

Then, there is the falsification by die-struck. In this case the faker hand-cuts a die, pretty much with the same kind of process as the ancients did, and mints a very credible coin. Even if these fakes are less common, for they require a very talented engraver, these are very dangerous fakes, often hard to spot even by experts. In this case we cannot compare the fake coin with an existing coin, so the analysis must be on the metal (like searching for silver crystallization, a sign of oldness), the style of the coin, and the search for artificial ageing processes. All these are very hard, if not impossible, to spot from a photo.

And then, there is the (so-called) fantasy coin, one that has all the correct style from the time and place it intends to imitate, but that never existed. These present day coinages mostly imitate pre-classical coins, or classes not much documented. In these cases there is absolutely no comparison point, so the identification shall be made by the same methods as the die-struck fakes, except that the stylist analysis is harder even for humans.

On top of all that, in most cases the fakers usually strive to give an ancient look, they disguise every revealing mark and force some degradation to keep it credible or to be different enough between copies so that they can sell several without suspicion. They also apply chemical patina (the characteristic colors from ancient coins) and bury the specimens to get dirty.

In conclusion, usually the fakes are different enough from the original to restrain us from using simplistic methods like direct matching or the coin's outline matching, although these can be used as a supporting process if they are efficient. The strategy should include occlusion resilience (to ignore the disguised parts of the coin) and a probabilistic approach (the more features of the coin are equal to another one, the more probable it is a fake). That said, for humans comparing a coin to every known fake is very hard to be done (as new fake types appear every day) and so the trained eye can use

other signs, like too regular holes in the coin field, or small metal bubbles, or the lack of stress marks from the mint (all possible indicators of casting), or an abstract style evaluation. Even if this may seem next to impossible to do with current technology, it would be a good future course of action.

3.1.3 Style

While nearly every single study on the field seems to overlook it, computer vision applied to coins doesn't have to be all about classification and identification. The stylistic aspects also have a tremendous importance in numismatics research and though it is hard, with current knowledge, to evaluate a specific aesthetic facet, it is possible to compare the symbolic features between specimens. This analysis may be used for clue searching from mint origins to dating the coin relatively to others. We could go even further and try to correlate icons from geographically distant coins that could point us to unsuspected commercial or political relations.

The challenges

How do we deal with an abstract concept like style? We could evaluate the lines and motifs struck (in general, the fineness or coarseness of the struck are a good indicator of the quality of the artisan), but the edge extraction is too prone to lighting variations to make a correct analysis. The concept of aesthetics is also too hard (by now) to explain to a machine. That leaves us to symbolism. Symbols are something we can deal with (just think of characters), we can identify, compare and scrutinize the spatial relations between them. If we tried to build classifiers for symbols and their positioning instead of classifiers for whole coins, we would probably be surprised by the information that would give us, not only for coin classification, but also for coin class relationships.

We could also get valuable information from legend reading, but the abbreviations were so common that we would need a huge and indecisive dictionary to be able to deal with them. Besides, the medieval characters were so unpolished that it is much easier for now just to deal with them as symbols.

The big challenge is that, given the malformations of ancient and medieval coins, the current edge detectors are not very precise in these cases. We would have to assume

occluded symbols and the fusion between them, and to learn which ones are real symbols and which are noise. Either we need better edge detection or better ways of learning from the current ones. This is a case that could fit well in convolutional networks, if only we had enough images to make the method accurate.

3.2 AND EVEN MORE CHALLENGES

When we look at the studies in the field, some more general challenges become evident:

3.2.1 Medieval coins are evil

Although several researchers experienced coin recognition on classical coins (usually on Roman coins, a few on Greek coins), from all the papers searched I could find only two teams approaching the problem with medieval coins, and both in a sideways, comparing with modern or Roman coins and achieving modest results. The difference is not as subtle as one may imagine. In fact, after the fall of the Roman Empire, the artistic quality standards and the artisans' talent gave a massive leap back. The coins from medieval times were as coarse as the ones thirteen centuries before (this is particularly true for peripheral European states). So the task of recognizing medieval coins is expected to be even more difficult than the already approached classical coins.

It gets worse: the chosen datasets, almost always the same ones, come from museums' collections and represent coins generally with a very good grade (conservation status). If we try to deal with the coins generally found in the market, private collections, or minor museums, we will confirm that such a grade of coins represent only a very small fraction of the specimens available. Most coins present a modest conservation status, which means an even harder task on relevant features extraction.

3.2.2 Ordinary images are not role models

While on laboratory we can build a dataset in a controlled environment with uniform conditions. When we get out to the general public, that's an impossible premise to maintain. The common images shared on the internet present a dreadful diversity of: light colors, directions and intensities; shadows, highlights and reflections; complex

backgrounds; sometimes the picture is taken with the coin inside its protective blister; the image resolution is at times too low, the noise too high; and a lot of other variations we cannot control. So, in any problem we shall focus on, the first step must categorically be trying to find the most resilient algorithms and homogenize as far as we can all these variations.

3.2.3 Big data is hardly an option (but big contribution is not)

When we read in a scientific magazine, or on the internet, on the subject of computer recognition, one concept is omnipresent: Deep Learning. Big Data, Convolutional networks, and alike notions seem to be putting a test on every imaginable problem around. Yet, how can we speak of big data when so many coin classes have no more than a handful of existent specimens, and often only one or two photos around? We may argue that if they are so rare the problem of classifying is not so hard, but the variations between classes are often so subtle that it is problematic to distinguish if we have two coins of different classes or if it is some variation inside the same class. If we want to classify a coin, we must have classifiers; however we can hardly have a complete dataset, and even that will be scarcely populated for each class. Therefore, instead of talking about Big Data, in this context we should be discussing Small Data.

That said, it doesn't mean we shouldn't use the internet communities to help feed the learning machine, on the contrary, their help is crucial in obtaining a big database and correct its mismatches, towards a recognition system both robust and helpful. That is why it is so important to bring these tools to the public instead of keeping them in the laboratory.

3.3 PROPOSED APPROACH – SOLUTION

3.3.1 The goal

The aim of this work is to contribute to a model for automatic recognition of medieval and ancient coins that could lead in the future to a system opened to the big public and helpful in classification and identification problems.

It is not the purpose to create new algorithms or methods, but to use or adapt the current state of the art, for only that way it is possible to cover such a large scope. That said, some innovations were needed in order to cope with the rough initial premises.

Since the goal is to orient the model to big public usage, there are two premises:

- The software, methodologies and algorithms are preferred to be of free access, or open source. This choice implies the dismissal of methods which hold a patent, closed, or with no available implementation, as promising as they may seem. It doesn't mean proprietary software cannot be used during the tests, but the final model must be free to use or be able to be adapted for that purpose.
- It must be oriented to the use of non-controlled images, as it is the expected reality. Nevertheless it was needed to include controlled images during the tests in order to understand the most critical variances and how to deal with them.

3.3.2 The approach

As the chosen subject and conditions are very hard, this work concentrates more on the initial phases of the recognition process – dealing with non-controlled images, segmentation, edge detection and feature extraction – since these present the biggest challenge and the most ill-studied one.

As discussed before, I believe that both identification and classification should be made by a probabilistic approach, being the big difference a matter of building class descriptors or specimen descriptors. In that sense both methodologies can be approached in the same way.

As for the style analysis (in the previously explained sense), this is a virtually unstudied concept, with its own challenges demanding new approaches, so it deserves its own focused research and it is not to be covered here.

Neither character recognition nor facial detection were used to assist the recognition process because, as discussed above, their contribution would be too limited in these particular coins and uncontrolled images.

3.3.3 The framework and test environment

Even if Matlab is the most widely used framework in the research community, this work was based on OpenCV 3.1 over C++. This is an open source framework, that besides going along with the initial premises of free access for all, it makes much easier if someone wishes to import the models here defined and programmed to end-user application (either web, desktop or even smart-phone based). As a bonus, some studies suggest that OpenCV tends to be faster than Matlab dealing with image processing, as shown in [63].

As for the tests, a controlled environment image dataset of medieval coins, detailed in chapter 4, was composed in order to represent the most common variations that interfere with the segmentation or edge detection, as: coin color and texture; lighting color, direction and intensity; noise and poor resolution; shadows; and different complex backgrounds. To this set some challenging images chosen from the internet were added. This dataset was the basis for the tests over Segmentation and Edge Detection.

For classification/identification using machine learning algorithms, a known existent image dataset of roman coins was used.

THE METHOD: I – SEGMENTATION

4.1 INTRODUCTION

Segmentation, when applied in the context of coins, usually refers to the separation of the coin from the background on an image and it is expectedly the first step in any classification method. Yet the researchers keep basing their work on little more than a couple of available image datasets, with grayscale pictures taken in carefully controlled light conditions, and contrasted backgrounds in uniform tones, in order to maximize detection results. This presents a serious problem in two ways. Firstly, we are wasting vast repositories of coin images all over the internet that could be used to make our algorithms more robust, and secondly, the same algorithms that achieve great results on these controlled images, when applied to “real world” coin images, tend to be very disappointing.

It is undeniable that dealing with internet or uncontrolled images and their low quality standards brings great challenges. Dark cast shadows or bright reflections severely interfere with most algorithms. Poor image resolution or too high compression lead to pixilation or “block” effect, thus adulterating the shapes. The complex backgrounds are presently very hard to remove. Ancient and medieval coins provide even more difficulties, like the irregularity of its borders that very rarely are a perfect circle; some coins are so thin that they barely contrast with the background; some are so thick that their shadows extend their border on segmentation algorithms; some coins are square, some are cut in half, some have holes. To make it even worse many images display the coin inside protection holders, making it very hard for a segmentation algorithm to distinguish between the coin border and the holder border. On the upside, the computation power is now much higher than a decade ago, so there is no more reason not to use color information on our algorithms.

Although the homogeneity of the controlled datasets facilitates higher (apparent) success rates, this is a false premise. Therefore, the segmentation of the public coin

images, overcoming big problems as bad illumination and terrible background choices, is the ground to more robust approaches on classification, not only narrowing the gap between researchers and public applications, but also allowing easy access to huge image datasets.

As this is the most ill-studied step on the present context, it is also where this research may give the most significant contribution, and so it is where the biggest share of time and effort were thrown in order to achieve a new efficient method.

A primary approach of the present method was accepted in the 2016 International Conference on Autonomous Robot Systems and Competitions (ICARSC 2016) and gave way to a paper to be published by IEEE [64]. Meanwhile the research continued and the method has evolved to include edge information on the detection, as long as some optimizations.

4.2 RELATED WORK

The research on coin segmentation has seldom been approached in isolated studies, being rather associated to a whole classification process.

The already mentioned initial researches with big acclamation, like Dagobert Project [4] and Coin-O-Matic [8], set the tone to the use of a massive dataset of modern coins photographed in an extremely controlled environment. In these conditions, they need only to rely on a global threshold and some basic edge detection to provide segmentation. Many researches, as in [7], frequently working over the same dataset, made use of generalized Hough Transform (HT). The HT limitation of detecting only circles is not a problem on modern coins, so this method was broadly accepted in that context. Unfortunately ancient and medieval coins do not provide perfect circles (or no circle at all) so HT is clearly insufficient. Some pioneer ancient coins researches, as [18], proposed the use of Canny Edge, since Sobel filters provided inaccurate edge information. Even GrabCut was deprecated in favor to Canny Edge in [43], which was shown to provide better accuracy.

It 2009 Zambanini and Kampel [22] concentrated exclusively on the segmentation problem over ancient coins and came up with a method combining Local Entropy and

Local Range of Gray Values. The method, still applied on grayscale images, provided good results, except when addressing the border shadows. My tests also revealed that Local Entropy calculation tends to be very computational intensive, even more adding a gray range calculation according to the paper description. More recently, Huber-Mörk, Zambanini, Zaharieva, and Kampel [33] presented a coin identification method heavily based on the border shape, and thus in segmentation, and suggested the use of a connected components analysis and the same Local Range of Gray Values. The method is not detailed but seems a lot similar to [22], proving the team was very confident on its accuracy. Still, in the presence of complex backgrounds, both entropy measure and local range if gray tend to highlight the background as much as the coin, making these methods unhelpful in the present context.

All these methods were using grayscale and controlled images, at least to some degree, and surely with uniform backgrounds.

4.3 PROPOSED METHOD

Despite its reliefs, a coin usually presents a narrow range of colors. Even if on occasion ancient and medieval coins may present areas of oxidation or dust, usually the most dramatic color variations are on its intensity. Thus, this approach begins with isolating the coin from the background based on a probabilistic comparison of color hue and saturation using Histogram Backprojection (described in A.1). Ignoring the intensity value has the additional advantage of diminishing the shadows noise on the border detection. We can then calculate a border approximation by means of a Convex Hull over the previous results (A.2).

A second stage uses edge information in order to refine the detected border. A Canny edge operator is calculated over the area previously calculated (B.1), followed by a second Convex Hull (B.2) resulting in a mask to apply over the original image hopefully removing all the background.

Some optimizations are made along the process in order to reduce noise and enhance the borders. The edge operators are particularly useful in the presence of pixelation caused by less than optimal compression of the original image.

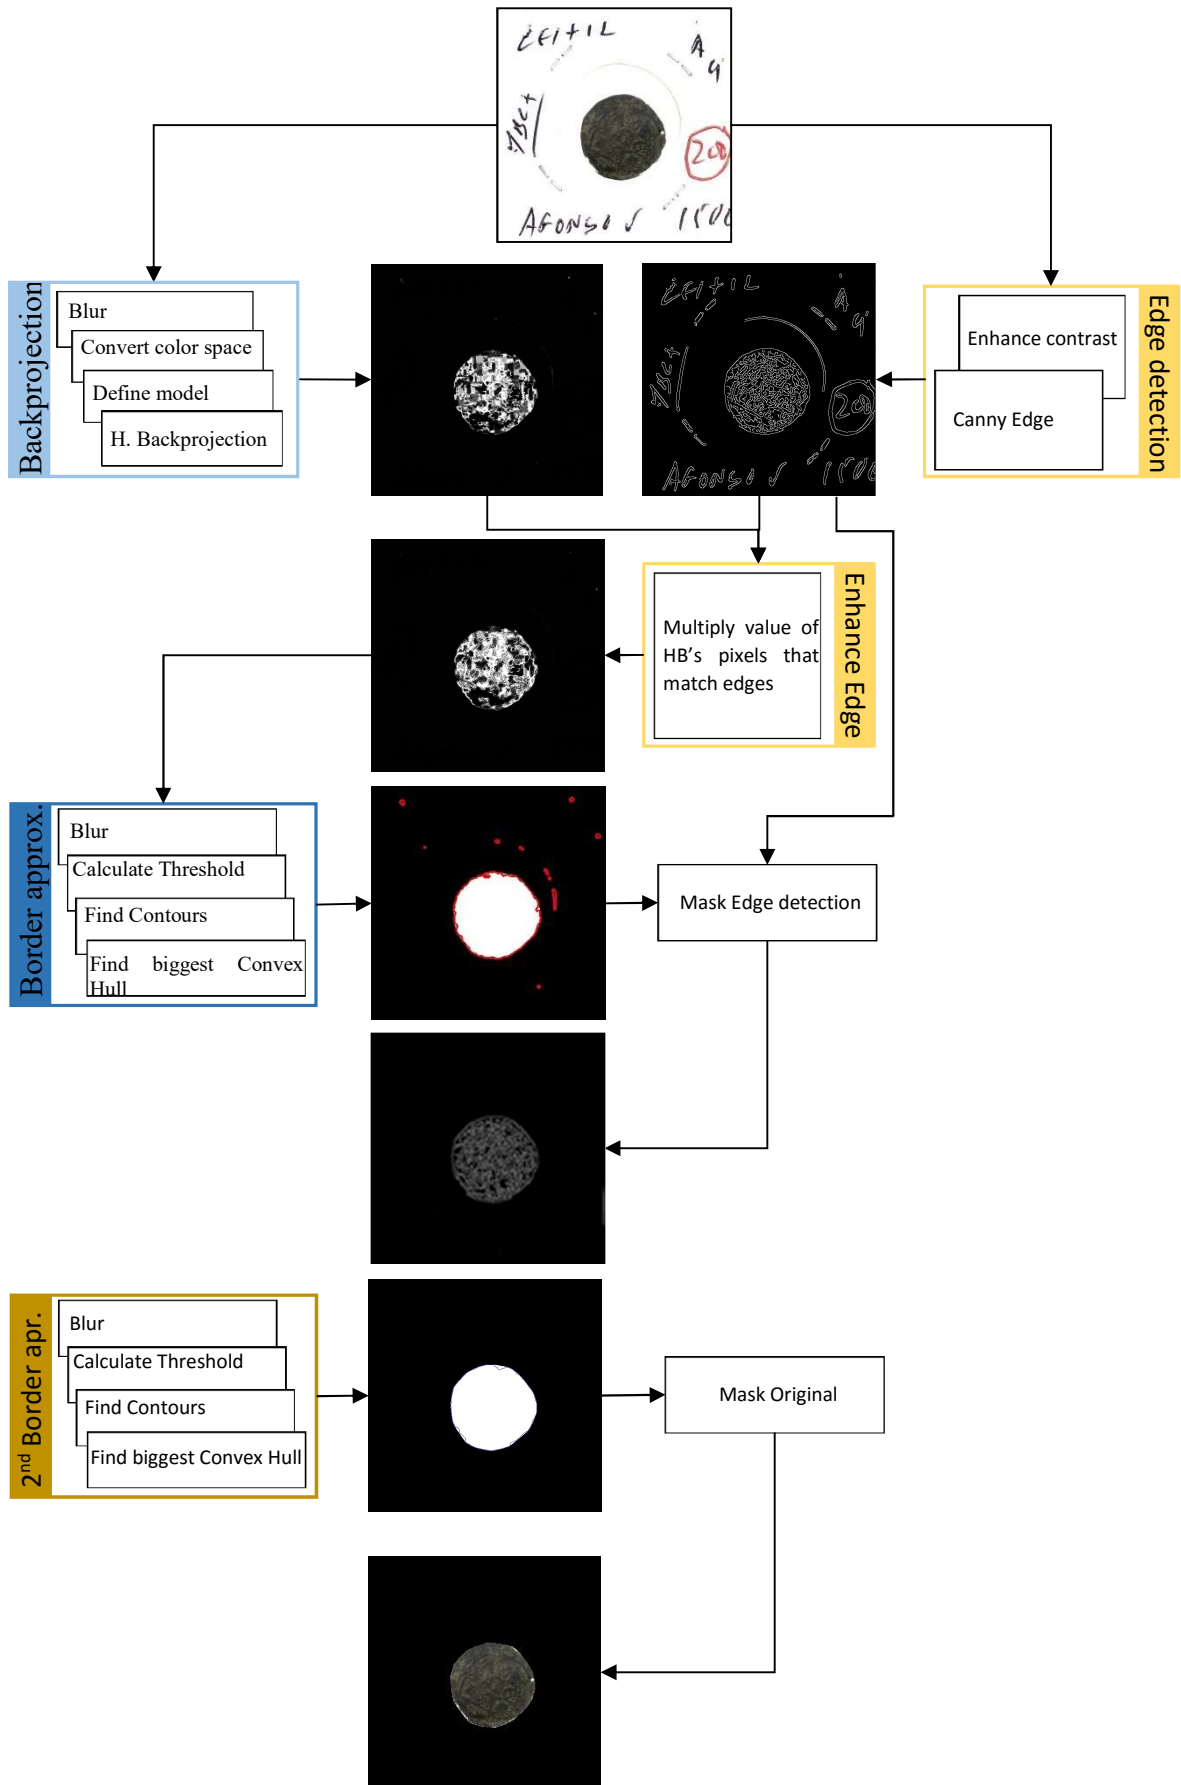


Figure 1 - Proposed method for segmentation.

The stages of the process are outlined in Figure 1. From top to bottom, the images show: original image; result from Histogram Backprojection and result from Canny Edge; contours in red and the biggest convex hull as white area; masked Canny Edge's result; result from second convex hull mask; final segmented image.

4.3.1 Step A.1 – Histogram Backprojection

Histogram Backprojection (HB) has been proposed by Swain and Ballard [65] back in 1992. Given a model image, the HB algorithm computes its histogram and then back-projects it in the test image's histogram in search for fitting pixels. The result is a single channel matrix with the probability of each pixel belonging to our ground model. In other words it gives us the probabilistic prediction of our object of interest.

Commonly HB is used with a previous model image like a section of skin to detect faces or hands in a test image. The test image may also be used as the model to perform some kind of self-HB, assuming that the object of interest fills most of the image. Given the variability of the coins, uncontrolled images and the complexity of possible backgrounds, neither of these approaches by itself is adequate in the present context.

Pre-processing and Color space conversion

A previous filtering with a Gaussian blur has shown to improve the results by removing noise. Gauss blur proved to be more successful than edge-preserving filters like Median blur or Meanshift.

Although a following addition to the image of its Laplacian, thus sharpening the image edges, may sometimes ameliorate outcomes, this was deprecated in favor of an edge refinement over the HB discussed later.

Converting the image from the usual RGB color space (Red, Green, Blue) to HSV color space (Hue, Saturation, Value or intensity) allows us to ignore the intensity value, more prone to lighting variations, and rely solely on the hue and saturations to define our object of interest. Obviously, a grayscale image, which has only one channel based on the

intensity – the channel we desire to ignore -, provides us low discriminating capabilities and it is not a good candidate for this algorithm.

Histogram Backprojection

The HB is calculated having the defined model as seed (see section: *Automatizing the method*) and its resulting probability matrix may be shown as an image, as in Figure 2, and has proven very efficient in attenuating complex backgrounds. However, sometimes the border contours may seem roughly defined, thus the need to refine the result.



Figure 2 - Original image with rectangle model (left), result from HB (right).

Adding Edge Info

Hue channel is particularly sensitive to artifacts caused by less than optimal image compression (in jpeg for instance). This may lead to an extravagant aliasing on the HB result. In these circumstances the border may be refined using edge information from a Canny edge operator over a grayscale version of the original image. We cannot simply add both results (HB's and Canny's) because it would retrieve back all the undesired features or noise. A *bitwise And* between the results is a better option, but still less successful than the chosen one: inflate the HB's pixel values that coincide to the edge detection by a multiplication factor. For instance: if some pixel in the HB result has the value 15 (dark, meaning low probability of belonging to the object of interest) and this pixel is in the edge detected by Canny operator, then we multiply that value by, let us say 4 (tests showed 2 to 4 to be good multipliers, more than that and the noise becomes too significant), getting a new value of 60. Our pixel will now for sure be kept once we do the binarization by thresholding on step A.2.

This process does not add border information where it absolutely was not found by HB, on the other hand it will also not highlight noise or undesired features to a degree that will interfere with the thresholding. In fact the tests showed that even if the improvements

are very slight, they are relevant enough to include this step in the method, and the increased processing time is not significant, since we will reuse the Canny edge calculation on step B.1.

4.3.2 Step A.2 – Border approximation by Convex Hull

Pre-processing and Thresholding

A Gaussian blur is applied to the output from the previous step in order to remove noise peaks, followed by an automatic threshold (see section: *Automatizing the method*) to eliminate the residual noise from the background that still populates the image. At the end of these operations the background should have been completely replaced by a black level.

Convex Hull

By then whatever remains must belong to our coin, although it is presented as an unconnected point distribution and not a uniform shape. To find the outline we calculate the Convex Hull, which is defined as the smallest convex set that contains a set of neighbor points. If more than one shape is found, either inside the coin (from shadows or colored dirt on the coin), or outside (from remnants from the background), we choose the biggest one and take it as the outline (border) of the coin.

As a failure recognition system, if the biggest shape's area is smaller than 20% of the image area (the size of the model rectangle) then it is marked as an assured failure and the chosen shape is set back to whole-image, based on the consideration that it is better to have a non-segmented image than just a small fragment of the coin. This possibility is interesting since it allows an automatic process to try again with the second stage, or a completely different method, without the need for human interaction.

We can now fill the outline calculated and use it as a mask over the original image to complete the first stage of the segmentation process.

4.3.3 Step B.1 – Second stage – Canny edge

If we really need real-time performances, we can end the method here (and deactivate the previous *Adding Edge Info* step). But we can significantly improve accuracy by refining a bit more the segmentation, especially in the cases where the first stage failed or had less than optimal results. As we based the first stage on color/texture information, it makes sense to base this stage on edge information.

Canny Edge

The algorithm for this operator was proposed by John Canny [66], and intends to find the edges of an image by a multi-stage algorithm consisting in: applying a Gaussian filter to remove noise; finding the intensity gradients; applying non-maximum suppression; applying a double threshold to determine the edges; and tracking the edges by hysteresis for suppression of the edges that are weak and not connected to strong edges. This is still one of the most robust and widely used algorithms for edge detection.

Masking the calculated Canny edges

As good as Canny's algorithm may be calculating the edges, it provides no means of distinguishing between our coin and the unintended features or background patterns. Thus we apply the mask calculated in the first stage.

4.3.4 Step B.2 – Refined border

Morphological operator: Close

After we apply the Canny edge operator, the result will be an image with black background and white features' edges. Still, rather than lines delineating the edges, we have a group of unconnected pixels. We can apply a common *Find edges* procedure to connect these pixels, but there is a problem: on uncontrolled images, we want to choose Canny's threshold values (see section: *Automatizing the method*) that maximize the edges in order not to end up with an open contour of the coin, but this also means that the inner edges from the coin faces will also be maximized. As these inner edges often touch the

border of the coin, the *Find edges* procedure often follows them causing a contour with canals cutting the outer border into the inner coin.

Applying a *Close* operator before finding the edges minimizes to a great extension the mentioned problem. On the downside: sometimes some small holes or dents persist on the border, inhibiting us from using the direct edge found.

A second Convex Hull

Consequently we need to recalculate a Convex Hull over the found edges. The process is the same as before, except this time we need not to binarize the image.

Final result

The result is a mask that, when applied to the original image, will show us a nice segmentation of our coin.

4.3.5 Automating the method

The automation of the process implies being able to calculate a good model for the HB algorithm and the most favorable threshold values for Convex Hull and Canny edge calculations. Many tests have been made and these are the elected methods:

Defining a model for Backprojection

The tests revealed that using a partial section of the coin in the test image results in a good model for the segmentation. The larger the section model chosen, the better it covers the tone and shadow variances of the coin, thus tending to provide a more exact result. Nevertheless we cannot risk to overflow too much over the background or to choose the wrong section. In order to automatize the process, a rigid rectangle model with the following characteristics is established:

- centered in the image;
- sized on 20% of the image dimensions.

These conditions may be changed according to the needs of each project, yet they are not arbitrary: Even if most coins are round, using a round model would cover only its

center area, thus leaving out most nuances caused by lateral lighting (shadows and highlights). The rectangular shape allows us to cover the center and to reach out a part of the outer sections of the coin with the model's corners. Also the position and dimensions of the established model force that the coin is also centered in the image and occupies at least 20% of it. Most images will easily satisfy this requirement, yet, in the cases where the coin is smaller, or there are several coins, or (the most common case) both sides of the coin are collated on the same image, we can use Cascade Classifiers, or any other detection method, to identify each coin and treat them separately, one at a time (this detection step is not covered in this work).

Threshold calculation over HB

A good estimation of the most favorable threshold value possible over the HB results is a crucial point in the presented method, but traditional automatic methods like Otsu's or Adaptive threshold failed miserably in this task. On empirical analysis, a relation was found between the ideal value and the end of the gradient from the biggest histogram peak: we reach an ideal value being equal to the first histogram bin, after the biggest peak, that is lower than 2.5% the maximum histogram value (i.e. if the maximum histogram value is 400, the ideal threshold value shall be the following bin number with a value lower than 10). To prevent remaining noise, the next few pixels are also required to remain below that percentage. Even if this value fails in some cases, it was the best approximation found, beating all statistical calculations (like averaging and standard deviations) and the mentioned traditional methods. In Figure 3 we can see an example of HB's histogram (flatted for display purposes): the maximum peak usually occurs in the first bins, representing the low probability assigned to background pixels. The secondary big peaks shown in this example suggest a somewhat troubled result with remaining noise.



Figure 3 - Histogram from HB result with a red line on the estimated ideal threshold value.

Canny Edge Thresholds

Fang, Yue and Yu [67] suggested the use of the Otsu's threshold value as Canny's high threshold (T1) and half that value for the low threshold (T2). This choice proved particularly efficient in images whose histogram presents the characteristic of two-extremum, which is rarely the case in pictures from coins. Nevertheless, the tests carried out in this research's context deducted that Otsu's value is a good reference, yet more detailed results came from using $T1 = Otsu/2$ and a $T2 = Otsu/4$. This values were later corroborated by [68] that presents a very proximate conclusion (for T2 they use $0.3*Otsu$).

In another line of thought, we could also use the average (avg) of the image values and its standard deviation (std), in the form of: $T1 = avg/2$, and $T2 = (avg/2) - std$. Both estimation methods result in very similar values, and show no perceptive differences.

4.4 EXPERIMENTS AND RESULTS

As the segmentation over non-ideal images is the least studied step of the classification of coins, this research invested a vast amount of time and effort being thorough on studying and testing the existing methods (always respecting the initial premises). The proposed method evolved over the limitations found along those tests. In the hope that it may be useful for other researchers to avoid the same errors, a list of other main methods and variants tested is presented in the Support Material chapter, along with some simple notes over their advantages and disadvantages. It is also worth to note that, besides the methods presented here, several other state of the art approaches were tested, proving disappointing results in this specific context, and so they were excluded from this analysis.

4.4.1 The datasets

Development dataset

In order to identify the variants that interfere with the coin segmentation, a dataset of "bad cases" was constructed over 17 medieval coins, chosen to represent the combinations of their usual relevant characteristics: color, brightness and field texture. Over these coins we constructed a controlled data set contemplating hard conditions, with combinations of the following variations: lighting directions, shadow intensities, shadows

over fractures (over several directions), textured backgrounds (colored coin trays, wood, marble, textured metal and newspaper) and noise (either caused by bad exposure, low resolution, or bad white balance). A small collection of 18 images of unusual specimens, like oddly shaped coins (square, oblong, holed, or concave), half coins, or unusually bright ones, was combined to the dataset. Over the refining tests the images that did not represent real variations, or repeated similar cases, were removed. In order to cover non-predicted variants, a random set of 32 coins was selected from the internet, regarding their complex backgrounds or challenges.

The resulting dataset, although partially having controlled conditions (in a harmful way), represents a collection of harsh cases that were used to study the challenges and to refine the proposed method, hoping to build a system as robust as possible.

Test dataset

To avoid the common pitfall of over-tuning a method to a dataset, a second collection was built of random internet images that provided a final test dataset.

A spider was used to transfer images from a google search on “medieval OR ancient coin”. After removing all unsuitable images (like images with several coins, drawings, or unrelated images) and the ones that didn’t follow the premises for both methods (proposed and GrabCut), the first 100 images were chosen to compose a test dataset. This provided a random collection of “real world” images, used for a final testing.

4.4.2 Comparing method - GrabCut

As a comparison method we used the general purpose GrabCut algorithm [59], an approach based on optimization by graph-cut. This interactive method relies on a user defined bound box around the object of interest and estimates the color and contrast distributions from both the object and the background, isolating the connected regions inside the bound box. The process may suffer multiple iterations until the user finds a satisfactory result. Some interesting variations of this algorithm appeared meanwhile, including “GrabCut in One Cut” [69], which is interactive just the same but, as the name indicates, tries to achieve the best result in only one iteration.

In order to automate the process an arbitrary bound box was chosen, 10 pixels from the image edge, on the assumption that people always leave a margin around the coin. If one wants to limit the method to round coins, one could also use the corners of the image.

4.4.3 Results

The results of both methods applied over the Test Dataset (100 random google images) are summed up in Table 1 and are grouped according to four criteria:

- Good: when the segmentation follows closely the border of the coin.
- Acceptable: near miss, leaves a small margin or small bits of the background around the coin, but still provides an acceptable segmentation.
- Non-acceptable: near miss, similar in magnitude to Acceptable, but crops minor segments of the coin possibly limiting a future classification process. In other contexts both Acceptable and Non-acceptable may be joined into one acceptable/failed group.
- Failure: when the detected outline crops significant segments of the coin, or leaves large areas of the foreground.

A human evaluation inevitably falls on some degree of subjectivity, yet this scale was preferred to an exact quantification of the deviation for the reason that it would not improve the usefulness of the evaluation. In fact, if we have a margin of a couple of pixels along all the coin border, the deviation would be the same as an equivalent area cropped from the coin. In the first case the segmentation would still provide a very valid object for classification, in the latter it would seriously damage the attempt of classification.

Method	Segmentation results			
	Good	Near miss		Failure
		acceptable	non-acceptable	
GrabCut	30%	19%	07%	44%
Proposed	51%	25%	09%	15%

Table 1 - Summary of empirical evaluation results

The average computation times are presented in Table 2. For a better perspective, the images were also separated into Small (file size below 400KB, average size of these images: 83KB), Medium (between 400KB and 1MB, average size: 580KB) and Big (above 1MB, average size: 2,1Mb) and time averages were calculated for each of the groups.

Method	Segmentation times			
	<i>Average</i>	<i>small images</i> <i>average: ~80KB</i>	<i>medium images</i> <i>average: ~580KB</i>	<i>big images</i> <i>average: ~2,1MB</i>
GrabCut	19,58 s	9,63 s	82,66 s	148,65 s
Proposed	0,29 s	0,16 s	0,80 s	2,62 s

Table 2 - Computation times on empirical evaluation

4.5 DISCUSSION

The chosen GrabCut approach presents reasonable results, and it easily addresses coins' fractures and holes. Surprisingly, while in the development dataset (the "hard cases"), GrabCut results fall mostly in the class of Near misses, in the test dataset the results tend to be more decisive and fall in Good or Fail. One limitation is that, although it works well with complex backgrounds, such a large initial bound box may embrace undesired elements that will not be excluded from the result image. In general terms, this method is very sensitive to shadows and tends to include them in its segmentation misstating the results. But the biggest deal breaker is the elevated computational times, typically around 10 seconds for small images, and 80-150 seconds for larger ones, which may stand prohibitive in most applications.

The proposed method achieves a good segmentation in half of the internet images and a low failure rate of 15%. Its biggest challenge is when the coin tones are very similar to the background. It also does not address the concave curves on the coin border, including fractures, due to the use of the convex hull method.

Both GrabCut and proposed method results may be affected by extreme cases of compression artifacts, the blocky effect caused by some algorithms like jpeg when too high compression is applied. This is something hard to attenuate and no process was found to overcome this problem.

The proposed method reveals to be very efficient dealing with shadows and highlights, and excluding them from the segmentation result. Besides accuracy, a very significant advantage of the proposed method is the computation time, in average more than fifty times faster than GrabCut.

4.6 CONCLUSIONS

None of the methods provide next to perfect accuracy in border description, so their use for cosmetic purposes is limited. Nevertheless the most usual need for coin segmentation is to be followed by a classification or identification process, and in that perspective the outcomes are very satisfactory.

The proposed method shows convincing results and proves that HB is able to give a convenient estimate of the actual coin region in a simple and fast way. Yet there is space for future improvements that may include: an HB model estimation, more flexible and comprehensive than the centered rectangle; a correction for concave curves; addressing image's high compression issues. The previous use of a detection method, like Cascade Classifiers, could allow every image to fulfil the premises for this method and increase even more the number of internet images fit for a classification/identification process.

THE METHOD: II – FEATURE EXTRACTION & RECOGNITION

5.1 INTRODUCTION

The previous chapter helped purge our tests images from everything but the coin itself, but since the beginnings of machine vision we know it is not possible to match two distinct images (even from the same object) by a simple bitwise compare. There are too many variants – like lighting direction or color, or a variation on the angle of the photo – for that to be a valid solution.

As humans, in order to classify a coin, we find its meaningful elements (or features) and compare it to previously known elements from other coins. The classification process is a matter of finding the class with all the same elements (or at least most of them in case there is occlusion or poor conservation). Dealing in ancient coins, we tend not to be too picky about the shape similarity. In an over-simplified example of a common symbol on coins, a cross is a cross, no matter if its branches are a little thicker, or if it is damaged in a corner.

Machine recognition shall follow the same principles: we need a higher level description of our object. More than that, we need a way to deal with all the variances that description may have, either environmental inconsistencies (again: lighting, angle, among others) or inter-class variations.

Having these feature descriptors, we then need to divide groups of features into classes. This task may be fuzzier than it seems. In the above example, we must admit that there are many different types of crosses (Latin cross, Christ cross, Templar cross, and tens of others) and sometimes it is not very easy (even for humans) to discriminate between them, especially if they are not in perfect shape. Although initial researches used a simple shape compare (suitable for modern coins), that proved too rigid for ancient coins.

We need some advanced techniques of machine learning in order to be able to separate sets of features (with many shared elements) into different classes, and then be able to apply that skill into new images to classify them.

The difference between a classification problem and an identification problem is in the magnitude of the classes: in the identification problem each coin is a class. A direct implication of this remark is that the identification problem is in fact a simplification of a classification problem, whether on the features to consider (like the simple use of outline shape on [38]), or on the matching of specimens (one can use a direct feature match in this case). Consequently, in this research we will focus on the classification problem and assume that identification can be achieved just as successfully with the same method, or, as an alternative, may be derived by a simplification of that method.

5.2 PROPOSED METHOD

5.2.1 Overview

As explained, bits are not enough to define features, so we need a higher level description of our object, and one that follows at least these two conditions:

- Focus only on the meaningful features from the object;
- It is able to translate that description into a mathematical language, so a machine can deal with it.

Many visual descriptors have been proposed in the past, from simple shape or color to more elaborated ones like texture, energy or more abstract forms. In this research we test some of the latest and more successful feature descriptors, as long as some promising new-comers: SIFT, SURF, ORB, A-KAZE and DAISY. These shall be further reviewed in the following sub-chapters.

Just as letters have a much stronger meaning when grouped into words, these descriptors will then be grouped into Bags of Visual Words (BoVW) to increase their significance.

After achieving a robust description for each image, we need a way to teach the machine how to distinguish to which class it belongs to. That task is well suited for a supervised learning approach, one that conjectures a function from labeling training data (in opposition to unsupervised learning where data is unlabeled, thus excluding the possibility of evaluating the solution). For that we need a set of labeled template images, the labeling being the class which each image represents. That set is then “trained” by a chosen supervised training algorithm in order to provide us with the function (or model) that will allow us to classify future test images.

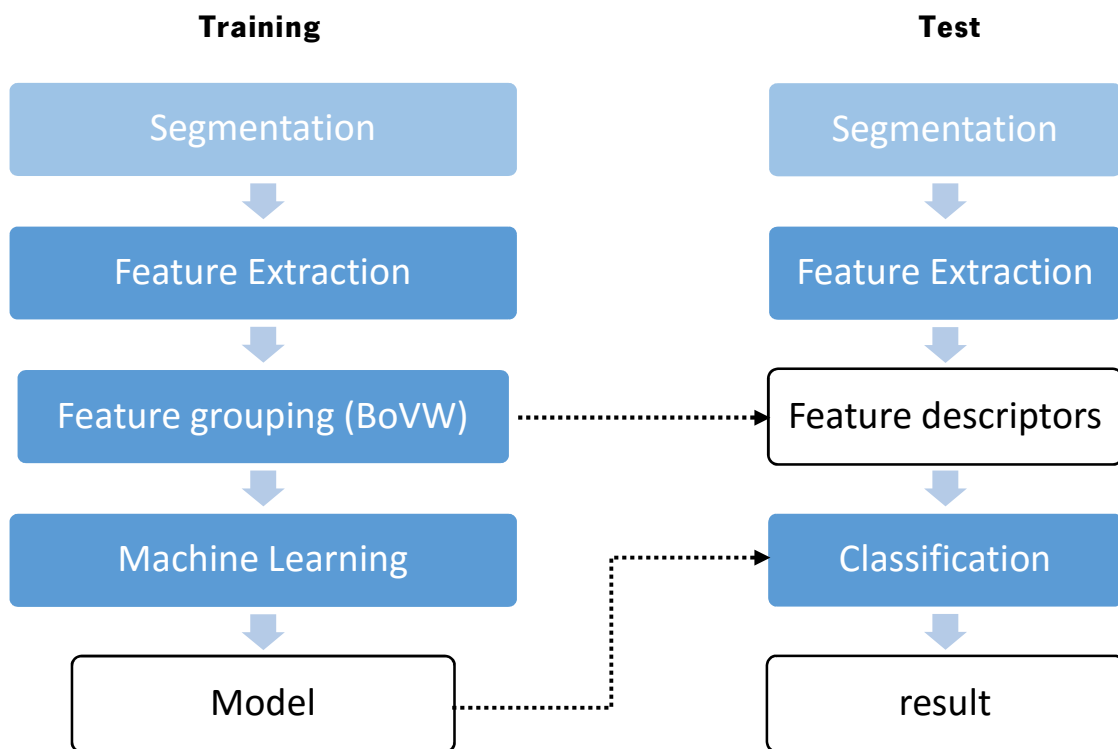


Figure 4 - Overview of the proposed classification method

The following sections shall explain and justify the choice of the methods that were included in the tests in each of the steps presented here.

5.2.2 Feature Extraction

The previous work on this area, summed up in the chapter State of the art, shows that the latest and most successful articles tend to use SIFT and alike for detection and description in feature extraction, thus this is also the approach in this research. In the last few years many alternatives for SIFT were presented, some of these most hailed new contenders are also interesting to examine in the present context. The following methods were chosen to be tested, taking into account their implementation availability, their success reports, and for being both scale-invariant and rotation-invariant (critical capabilities when dealing with coin images):

SIFT

Scale-Invariant Feature Transform was presented in [3] and is arguably the most acclaimed feature descriptor. It is both scale-invariant and rotation-invariant and its implementations include the detection of the key points. Although it is erroneous to generalize to all contexts, it is usually seen as the most (or at least one of the most) precise descriptor around. On the low side it is a little slow (for today's standards), it is somewhat sensitive to orientation estimation error. This method is patented though free to use for research purposes.

SURF

Speeded Up Robust Features [70] is viewed as the fast brother of SIFT, though the fastness often comes with the price of a reduced accuracy. It retains the same invariance properties, same sensitiveness, and it is also protected by patent but free to use for research purposes.

ORB

Oriented FAST and Rotated BRIEF [71], as its name discloses is a fusion of *Features from Accelerated Segment Test* (FAST) key point detector and a scale-invariant and rotation-invariant version of the *Binary Robust Independent Elementary Features* (BRIEF) descriptor. ORB claims to be a free, fast and efficient alternative to SIFT/SURF and in fact it was presented in [72] to be the best of a group of contenders of SIFT/SURF for image

recognition and as being the descriptor of choice for large illumination changes or perspective distortions. Nevertheless, SIFT still usually presents slightly better accuracy and the same paper is surprised to note that ORB descriptor is less efficient to compute than SURF.

A-KAZE

KAZE, the Japanese word for *wind*, was presented in [73] and operates in a nonlinear scale space, in opposition to previous methods as SIFT that find features in the Gaussian scale space. Accelerated-KAZE (A-KAZE) is a faster version of the method using *Fast Explicit Diffusion*, and claims to obtain comparable results to KAZE. The methods are perhaps even more robust than SIFT/SURF/ORB in terms of rotation and scale invariance, and also to brightness and blur invariance. Some tests present an accuracy comparable to SIFT or even better. This method is open source.

DAISY

DAISY [74] does not incorporate a key point detector, thus one must use another detector in combination with this descriptor. It was designed specifically for dense wide baseline matching purposes and some tests show comparable accuracy to SURF and even SIFT at a fraction of time. It is supposed to be scale and rotation invariant, nevertheless it usually shows lower performance in comparison with other methods like SIFT and ORB [72]. It is also somewhat sensitive to orientation estimation error.

Whatever detector/descriptor is chosen, the previous segmentation (as the one presented in part 1 of this research) shall exclude any key point detected outside the coin thus preventing the background interference in the classification process. Depending on the detectors, some pre-processing, like smoothing the image to reduce noise, may also improve the success rate.

5.2.3 Feature organization

The direct matching of features may be applied in an identification problem, but for classification it has many limitations (remember the inter-class variability explained in the

beginning of this report?). In fact, the features detected are not complete human readable symbols, but somewhat rigid descriptions of small areas. It is not a surprise that the first attempts to match features in ancient coins had promising results [27], but still too inaccurate. Some model of organization or grouping may be helpful in this task.

A strategy that is gaining adepts is inspired on the field of natural language processing:

Bag of Features (BoF), aka Bag of Visual Words (BoVW)

In the Bag of Words (BoW) model a text is represented as a set (bag) of words. The model neglects the order of the words, or even grammar, but focus on the multiplicity of words. The concept is that a letter by itself has little discriminative value, but when grouped in words it provides a strong meaning. And the counting of words may well identify a text. No wonder this ended up being used for computer vision [75].

In images, instead of letters we have features, thus the name of the method changing to Bag of Features or Bag of Visual Words. We can cluster those key points (using k-means for instance) and the cluster centers will act as visual words (or features, but in a higher level sense, the nomenclature may be confusing here). We now have a Bag of Visual Words.

Taking a set of template images we can compile a collection of visual words and establish a reference dictionary. For each image we can now represent it by comparing each feature and its surroundings with its nearest neighbor in the dictionary and computing a histogram of the words frequency.

Recent articles with the highest successful rates stand out this strategy as a very efficient one in the context of Roman coins, especially when including spatial information [39][40][44], making it strong candidate for this research. Unfortunately most implementations of BoVW, including Opencv v.3.1 used in this research, cluster the key points based in Euclidean distance, which is an inappropriate way to cluster binary descriptors like ORB and A-KAZE. An work-around could be implementing the clustering process based in Hamming distance, like suggested in [76], but that is beyond the scope of this research so ORB and A-KAZE will be left out of BoVW approach.

5.2.4 Machine Learning

In a supervised learning approach applied to image classification, a set of images (templates) carefully labeled with its corresponding classes is trained in order to establish a model able to distinguish each class. That model will later be used to predict the class for a giving test image. The mathematical way of discriminating between classes is what distinguishes each algorithm, making them more appropriate for some scenarios and less to others. Unfortunately there is no good-for-all algorithm, so the approach in this research is to test several of the most promising methods.

A side note: in the last few years we have seen a lot of hype about Convolutional Neural Networks for its good results in many areas, even if it often needs hard-to-rationalize parametrizations. The problem is that these algorithms perform better with training sets in the order of the tens or hundreds of thousands, and up, the so-called Big-Data. As exposed in the beginning of this report, Big Data is not an option in the present context. Even if some work-arounds are starting to be explored, as artificially expand the “small data” for instance, this is still a very shady area, deserving its own research. Thus, even if I have no doubt that the application of Convolutional Networks on classification of ancient coins is a subject to be watchful, it is not tested in this research.

On the other hand, some other favorite techniques are very suitable in this context and shall be tested. It is the case of SVM, Random Forests and Naive Bayes. The older and not so promising K-Nearest neighbors is also included for comparison purposes.

SVM

Support Vector Machine [77], has been around for some time, used in many contexts, and still shows up as an excellent alternative. This discriminative classifier takes the labeled training data (the extracted features for each class) and calculates a separating hyperplane that can be used to mark new images into one class or the other, based on which side of the hyperplane they fall, thus making SVM a non-probabilistic binary linear classifier. In Figure 5 we can see a simplified example of the discriminative hyperplane calculated, the red line, H3, marks the larger minimum distance between the two classes (black dots and white dots).




Figure 5 - Simple example of discrimination on SVM

Of course this example is an oversimplification. Besides SVM being able to deal with multi-class discrimination, the features often intersect between classes and the algorithm can perform non-linear classification mapping spaces whose dimension is higher than two.

We can have many ways of applying SVM to a classification task, these are the most common strategies:

- Multiclass ranking: SVM tries to separate all classes with a single mathematical function. This is the fastest and most compact approach, yet in many cases a single function doesn't exist so the approximations found may have a low accuracy.
- One-Against-All: Rather than one n-Class problem, we deal with n binary problems. One class remains on one side of the hyperplane (as a positive), all the other classes are on the other side (as negatives). In practice, this means having to learn a function for each class known. Although more computational intensive, this strategy tends to be more accurate.
- Pairwise (aka One-Against-One): Again, the big problem is reduced to multiple binary problems, but in this case a function is computed for each pair of classes, leaving us with $n(n-1)/2$ classifiers. As the dataset grows this strategy may become computational prohibitive for end-user applications thus this strategy will not be tested in this work.

SVM is frequently the first try (and often the last) in many classification problems. In the coin context, it is still part of the most successful methods, so it is no surprise that it is tested in this research too.

Random Forests (aka Random Trees, or Random Decision Forests)

Random Forests [78] is a method that can deal both with classification and regression, among other tasks. It is based on *Decision Tree*, a predictive model which maps binary decisions about each node (in our case it can be a feature or group of features), where each non-leaf node has two child nodes. The process leads to leaves that represent conclusions about the whole object (in our case, the class). The same class may be represented in several leaves, which represent several paths to the same conclusion. Figure 6 shows a fantasy decision tree that tries to figure the predominant color in an image. Notice that it is imperfect and very incomplete, but it is the representation the algorithm achieved based on the training data supplied.

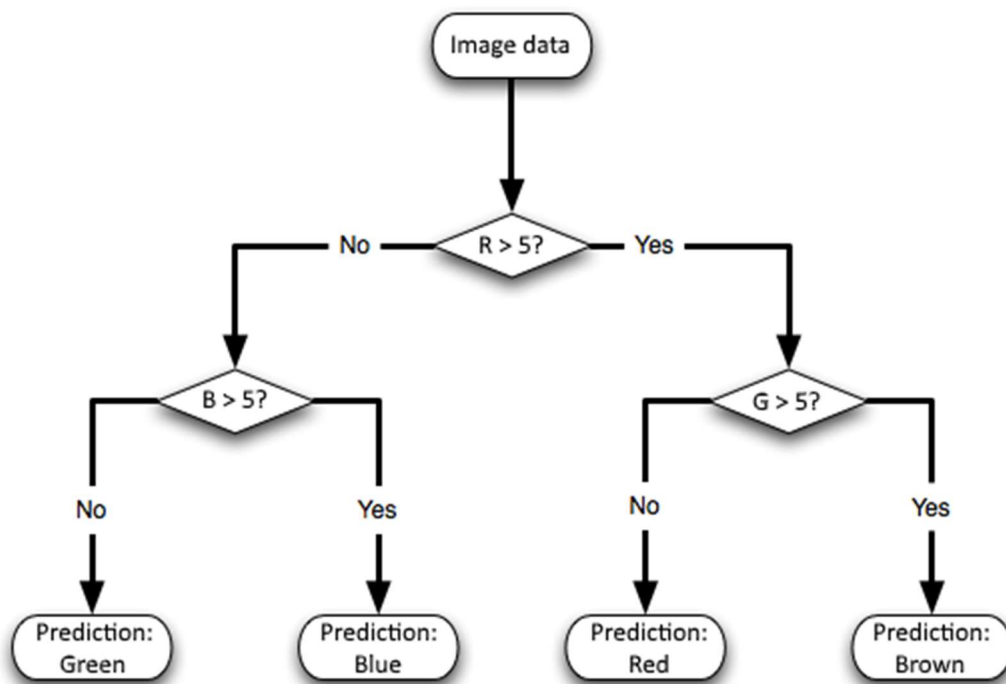


Figure 6 - A Decision Tree

A big problem about Decision Trees is that, as we allow it to grow, they tend to become over-trained, thus not being able to classify any sample we give. An intuitive explanation,

using our example: as we provide more templates and allow the tree to grow, it will achieve very specific colors as navy-blue and turquoise-blue, making it more improbable to be able to identify many other tones or the ability to just answer: “It is blue”.

Random Forests overcame this problem by providing an aggregation of different Trees, each trained with the same parameters but with random sub-sets of the features. The output is the prediction that receive the majority of votes. The random sampling feeding multiple trees (formally *bootstrap aggregating*) decreases the variance of the model by preventing strongly correlated trees, thus conducting to higher accuracy.

In many problems, Random Forests tend to achieve better results than SVM, it is usually faster to train, easier to parametrize, and may provide intelligible models showing the most relevant features. It also tends to work better than SVM when there are complex interactions between features. On the downside, some authors claim it tends to perform worse on high-dimensional sparse data, which is the case with *Bag of Words*.

Naïve Bayes

Naïve Bayes [79] is in fact a whole family of simple classification techniques, that represents classes as vectors of feature values, assumes they are normally distributed and usually (but not necessarily) the features values are taken as independent of the value of any other feature. Despite being considered a probabilistic classifier, it can be used without accepting the Bayesian probability. It remains a very popular method in many areas, and in some contexts is still competitive with more advanced methods as SVM, or Random Forests, thus the reason to be included in this research. A strong benefit of Naïve Bayes is that it only requires a small amount of training data, which may be interesting in our context.

OpenCV’s implementation, the one tested, assumes that the continuous values associated with each class are distributed according to a Gaussian distribution. The vectors of feature values are calculated for all the training data, and for each class, a mean vector is estimated as well as covariance matrices. Those will later be used for classification of new images.

k-NN

K-Nearest neighbors is about the simplest as a classification algorithm can be. It matches the raw training data and the class of a new sample is predicted by a majority vote, i.e. the most common, of a number (k) of its nearest neighbors. To increase accuracy a weight may be assigned to each neighbor according to its distance, and that factor is taken into account during the voting. The test here presented considers this weight factor.

The computation is deferred until classification, but the process is so simple that it does not imply necessarily higher computation times in relation to other more advanced methods.

k -NN may hardly be expected to be as accurate as more advanced methods, nevertheless it is included in this research to serve as a base line for comparison purposes.

5.3 EXPERIMENTS AND RESULTS

5.3.1 The dataset

For the tests an image dataset of Roman republican coins, belonging to the Museum of Fine Arts in Vienna, Austria, was provided by Computer Vision Labs. This dataset was used for evaluation in [34] and it is publicly available for research purposes in [80].

The dataset is composed of 180 images, distributed in 60 classes, 3 samples per class.




Figure 7 - Sample images on the dataset

For each class:

- The first samples were used to compose the dictionary on the Bag of Visual Words model.
- Both the first and second samples were used as training images in all the machine learning algorithms.

- The third samples were used as test images to which correspond the results presented.

5.3.2 The tests

The method is divided into 3 stages:

- Dictionary Creation

It is a one-time operation. The template images (first sample of each class) have its features extracted, and these are clustered into a pre-defined number of Visual Words that will establish the dictionary (also called vocabulary). This dictionary is saved into an “yml” file in order to be read and used in the other stages.

- Training

Also a one-time operation. The template images (first and second samples of each class) have their features extracted and, in the BoVW approach, these are also grouped according to the dictionary. Then the descriptors and corresponding labels are fed to the Machine Learning algorithm that will generate a Model which can be saved as an “xml” file to be used later.

- Testing

Applying the prediction model to unknown images is the real goal of all the process, and is supposed to be used as needed. Again, the features of the image are extracted (with the same technique used in the previous stages). In the case of the BoVW these are also grouped according to the dictionary. Then, based in the model created in the training stage, the algorithm outputs a prediction of the coin class.

Although several feature detectors/extractors were tested, there was no combination between them and the same method was used across all stages. For instance, when SIFT is used as a detector, it is also used as a descriptor, and it is used for dictionary creation, for training, and for testing. The whole process is separately tested for each of the remaining feature descriptors (SURF, ORB, A-KAZE, DAISY).

All the tests were coded over OpenCV v.3.1 framework.

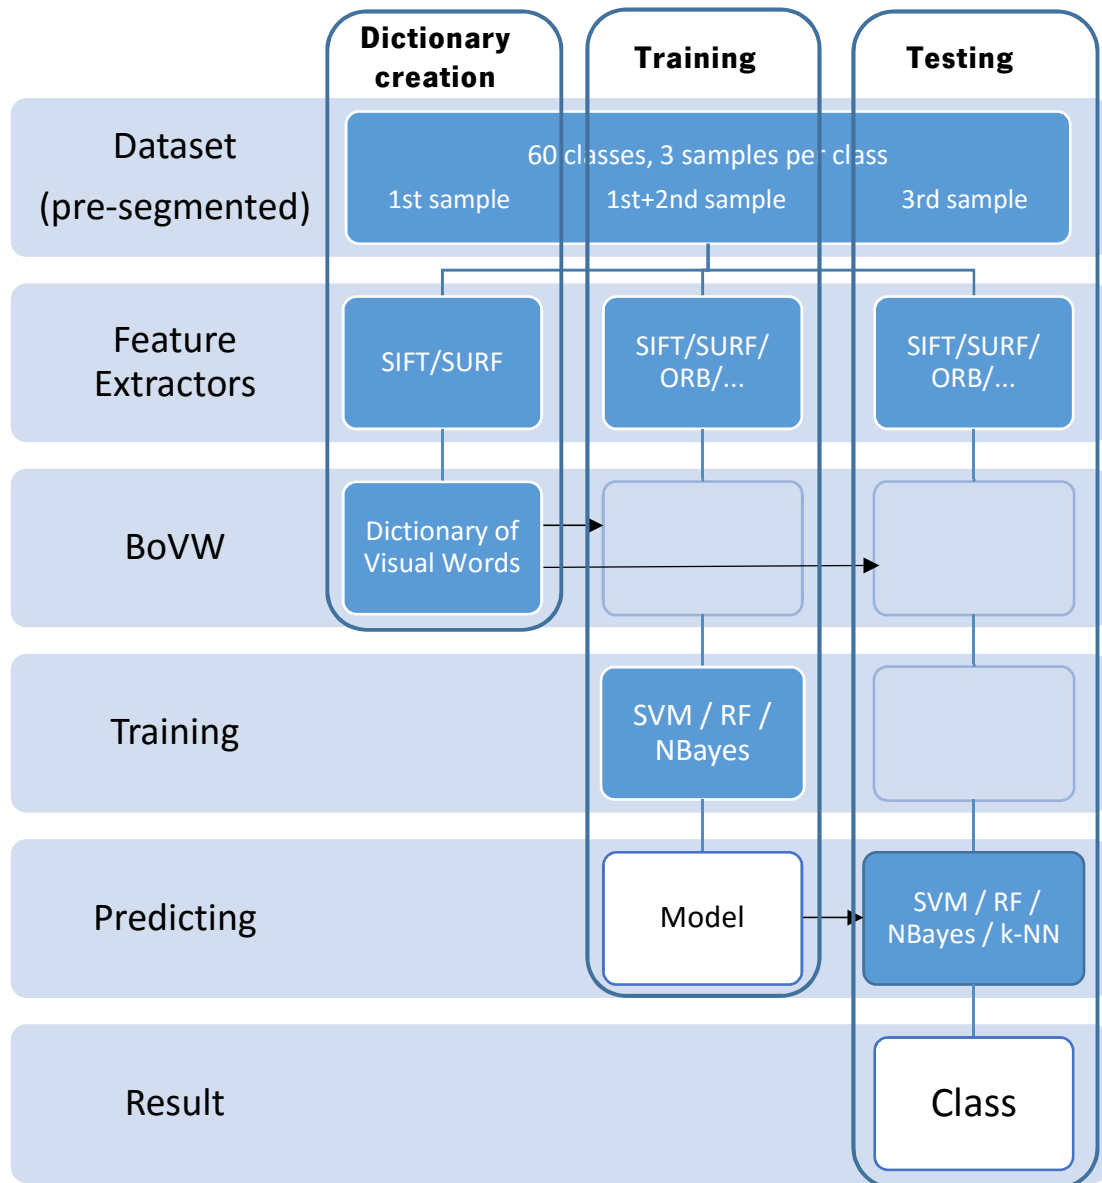


Figure 8 – Chart of the classification tests

5.3.3 Results and Discussion

The initial tests, using a BoVW dictionary with size 200, are summarized in Table 3. They show that the most promising methods are SVM, with the One-against-All strategy, and Naïve Bayes. Random Forests tends to have worse results than the above and does

not present much improvement when adjusting the parameters or applying pre-processing to the images. In fact R. Forests consistently behaves worse when the images are pre-processed. The results shown here for SVM consider a linear kernel, since both Polynomial and Radial Basis Function always presented worse accuracy. The value of parameter C was kept very low (0.05 or below) for the best accuracy. The advisable strategy for Non-linearly separable data suggests that besides keeping C low, one should increase considerably the maximum number of iterations in order to correctly solve the problem. Nevertheless, increasing the iterations conserved or lowered the accuracy during the tests.

	SVM <i>(1-vs-All)</i>	SVM <i>(multiclass)</i>	RForests	Nbayes	k-NN
SIFT	15,00%	13,33%	16,67%	15,00%	8,33%
SIFT+Gauss	28,33%	16,67%	6,67%	21,67%	8,33%
SURF	18,33%	16,67%	15,00%	23,33%	8,33%
SURF+Gauss	25,00%	18,33%	15,00%	16,67%	11,67%

Table 3 - Summary of Machine Learning results

Pre-processing tests were tried, always applied to all stages: dictionary creation, training and detecting. Although the Gauss filter seems to show the best accuracy improvement, there are cases where the results deteriorate. Other strategies included: normalizing the image's histogram; equalizing the image; resizing to 480x480 pixels in order to minimize noise and minor undesired features; and contrast enhancing. Table 4 summarizes the results, considering the use of SIFT descriptors.

	SVM <i>(1-vs-All)</i>	SVM <i>(multiclass)</i>	RForests	Nbayes	k-NN
No pre-proc.	15,00%	13,33%	16,67%	15,00%	8,33%
Histogram norm.	20,00%	13,33%	6,67%	16,67%	6,67%
Resize to 480px.	10,00%	21,67%	5,00%	18,30%	6,67%
Gauss filter	28,33%	16,67%	6,67%	21,67%	8,33%
Equalization	28,33%	16,67%	6,67%	20,00%	10,00%

Table 4 - Impact of preprocessing over de training methods with SIFT descriptors

As for training times R. Forests is very fast, N. Bayes is by far the slowest method and SVM is fast but the One-against-All implies multiplying the training time by the number of classes and in this case it becomes the most time consuming training. k-NN does not have a training stage. All methods are very fast in predicting stage.

5.3.4 Conclusion

The results are quite disappointing, considering the optimistic results shown by a couple of papers using the same methods. Some more work needs to be done in this area both in the search for pre-processing filters that enhance the final results and in the search for optimal parameters. R. Forests has generally shown worse results than N. Bayes, which comes as a surprise. An interesting approach to study could be the use of a One-against-All strategy similar to SVM.

FINAL CONCLUSIONS

6.1 CONCLUSIONS

Detailed conclusions for each part of the method were already presented in the corresponding sections. In general terms the segmentation method here proposed reveals itself as a very promising way of making coin datasets step out of the labs and allowing the internet to provide material for future researches. As discussed, a coin detector method could improve even more the number of available images.

Although this work was focused on such a specific subject as medieval and ancient coins, the same segmentation method can be easily applied or adapted to other flat and uniform-colored subjects.

On the machine learning part, there is still some work to be done. None of the tried strategies provided acceptable results for end-use application, in spite of the optimistic results shown on some papers using the same strategies. The work started here must continue, trying other approaches and combinations, and keeping alert to new findings in the area.

6.2 PROSPECT FOR FUTURE WORK

The focus on this research was always to provide a basis for getting the coin recognition out of the lab to the public usage. Thus, as long as we can reach interesting results in the machine learning part, it would be interesting to really implement these methods in an end-user application.

On the research point of view, a different approach to overcome the difficulties of contours detection could be the use of video in order to infer a 3D model of the coin face. With this model we would eliminate the effects of shadow and highlights and would have an accurate representation of saliencies, thus the segmentation would be a simpler task and the contours would be much more precise. Ideally this method would allow an end-

user to simply wave a smartphone over a coin, to get the 3D representation of its face and get a more accurate classification, or an alert on an existent identical coin (perhaps a fake or a stolen coin).

A very interesting follow-up to this research would be to try to adapt traditional Convolutional Networks to the use of “small data” (instead of the more fitted concept of “big data”). Besides auspicious accuracy rates, this could provide the tools to change the focus from random feature extraction to symbol recognition, thus paving the way to style analysis (whose benefits were discussed in the introduction).

BIBLIOGRAPHY

- [1] M. Zaharieva, M. Kampel, and S. Zambanini, "Image based recognition of coins – An overview of the COINS project," *Perform. Eval. Comput. Vis. 31st AAPR/OAGM Work.*, pp. 57–64, 2007.
- [2] J. Jarrett, R. Hüber-Mork, S. Zambanini, and A. Felicetti, "Coinage, digitization and the world-wide web: numismatis and the COINS project," *New Technol. Mediev. Renaiss. Stud.*, vol. 3, pp. 459–489, 2011.
- [3] D. G. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, Nov. 2004.
- [4] M. Nölle, H. Penz, M. Rubik, K. Mayer, I. Holländer, and R. Granec, "Dagobert – A New Coin Recognition and Sorting System," *Proc. VIIIth Digit. Image Comput. Tech. Appl.*, pp. 329–338, 2003.
- [5] R. Huber, H. Ramoser, K. Mayer, H. Penz, and M. Rubik, "Classification of coins using an eigenspace approach," *Pattern Recognit. Lett.*, vol. 26, no. 1, pp. 61–75, 2005.
- [6] M. Reisert, O. Ronneberger, and H. Burkhardt, "An efficient gradient based registration technique for coin recognition," *Proc. Muscle CIS Coin Compet. Work.*, vol. 19, p. 31, 2006.
- [7] M. Reisert, O. Ronneberger, and H. Burkhardt, "A fast and reliable coin recognition system," *Proc. 29th DAGM Conf. Pattern Recognit.*, pp. 415–424, 2007.
- [8] L. J. P. Van Der Maaten and P. J. Boon, "COIN-O-MATIC: A fast system for reliable coin classification MUSCLE CIS benchmark," *Proc. Muscle CIS Coin Compet. Work.*, 2006.
- [9] L. J. P. van der Maaten and P. J. Boon, "Reliable Classification of Partially Occluded Coins," *Proc. 2007 MUSCLE CIS Work.*, 2007.
- [10] M. Zaharieva, "To SIFT or Not to SIFT? Exploring the discriminating power of SIFT features for coin classification," *Proc. MUSCLE CIS Coin Recognit. Compet. Work.*, 2007.
- [11] K. Adnan, B. Sekeroglu, and K. Dimillier, "Coin Identification Using Neural Networks," *Proc. 5th WEAS Int. Conf. Signal Process.*, vol. 2006, pp. 88–92, 2006.
- [12] T. M. Ghanem, M. N. Moustafa, and H. I. Shahein, "Gabor wavelet based automatic coin

- classification," *Icip 2009*, pp. 305–308, 2009.
- [13] A. Gavriaseva, A. Molder, O. Martens, C. Kyrkou, and T. Theocharides, "Cross-correlation-based image matching of coins," in *2012 13th Biennial Baltic Electronics Conference*, 2012, pp. 319–322.
- [14] A. Vadivelan and K. Praveen, "Implementation of Coin Recognition by Extracting the Local Texture Features with Gabor Wavelet and LBP Operator," *Int. Conf. Electr. Eng. Comput. Sci.*, no. 1, pp. 135–139, 2013.
- [15] K. Wei, B. He, F. Wang, T. Zhang, and Q. Ding, "A novel method for classification of ancient coins based on image textures," *Proc. - 2nd Work. Digit. Media its Appl. Museum Heritage, DMAMH 2007*, no. 1, pp. 63–66, 2007.
- [16] D. Mehta and A. Sagar, "A Survey on Various Techniques of Coin Detection and Recognition," *Int. J. Comput. Appl.*, vol. 69, no. 5, pp. 29–32, 2013.
- [17] M. Kampel, R. Huber-Mörk, and M. Zaharieva, "Image-Based Retrieval and Identification of Ancient Coins," *IEEE Intell. Syst.*, vol. 24, no. 2, pp. 26–34, 2009.
- [18] M. Kampel and M. Zaharieva, "Optical Recognition of Modern and Roman Coins," *Layers Perception. Proc. 35th Int. Conf. Comput. Appl. Quant. Methods Archaeol. (CAA), Berlin, 2.-6. April 2007*, p. 172, 2008.
- [19] L. Van Der Maaten and E. Postma, "Towards automatic coin classification," *Proc. of the EVA-Vienna*, 2006.
- [20] M. Zaharieva, R. Huber, M. Nölle, and M. Kampel, "On ancient coin classification," *8th Int. Symp. Virtual Reality, Archaeol. Cult. Herit.*, pp. 55–62, 2007.
- [21] S. Zambanini and M. Kampel, "Segmentation of Ancient Coins Based on Local Entropy and Gray Value Range," *Proc. 13th Comput. Vis. Winter Work.*, 2008.
- [22] S. Zambanini and M. Kampel, "Robust Automatic Segmentation of Ancient Coins," *4th Int. Conf. Comput. Vis. Theory Appl.*, pp. 273–276, 2009.
- [23] S. Z.-M. KAMPEL, "Coin data acquisition for image recognition," *Proc. 36th CAA Conf.*, pp. 163–170, 2008.
- [24] M. Goodman, *Numismatic Photography*. 2008.
- [25] G. Hobermann, *The Art of Coins and their Photography*. H.N. Abrams, 1982.

- [26] S. Zambanini, M. Kampel, and M. Schlapke, "On the Use of Computer Vision for Numismatic Research," *9th Int. Symp. Virtual Reality, Archaeol. Cult. Herit.*, vol. 45, no. 7, pp. 17–24, 2008.
- [27] M. Kampel and M. Zaharieva, "Recognizing Ancient Coins Based on Local Features," *Adv. Vis. Comput. Part I*, pp. 11–22, 2008.
- [28] O. Arandjelović, "Automatic attribution of ancient Roman imperial coins," *Comput. Vis. Pattern Recognit. (CVPR), 2010 IEEE Conf.*, pp. 1728–1734, 2010.
- [29] R. Allahverdi, M. Dehshibi, A. Bastanfard, and D. Akbarzadeh, "Sasanian Coins Classification Using Discrete Cosine Transform," *Artif. Intell. Signal Process. (AISP), 16th CSI Int. Symp.*, pp. 278–282, 2012.
- [30] R. Allahverdi, M. Mahdi, A. Bastanfard, and D. Akbarzadeh, "EigenCoin: Sassanid coins classification based on Bhattacharyya distance," *WCIT-2011, AWERProcedia Inf. Technol. Comput. Sci.*, vol. 1, pp. 1151–1160, 2012.
- [31] S.-S. Parsa, M. Rastgarpour, and M. M. Dehshibi, "Statistical Feature Fusion for Sassanian Coin Classification," in *Recent Advances in Information and Communication Technology, Advances in Intelligent Systems and Computing*, vol. 265, Springer, 2015, pp. 75–84.
- [32] A. Kavelar, S. Zambanini, M. Kampel, K. Vondrovec, and K. Siegl, "The ILAC-project: Supporting ancient coin classification by means of image analysis," *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci. - ISPRS Arch.*, vol. 40, no. 5W2, pp. 373–378, 2013.
- [33] R. Huber-Mörk, S. Zambanini, M. Zaharieva, and M. Kampel, "Identification of ancient coins based on fusion of shape and local features," *Mach. Vis. Appl.*, vol. 22, no. 6, pp. 983–994, 2011.
- [34] S. Zambanini and M. Kampel, "Coarse-to-Fine Correspondence Search for Classifying Ancient Coins," *Comput. Vis. – ACCV 2012 Work.*, vol. 7729, pp. 25–36, 2013.
- [35] S. Zambanini and M. Kampel, "Using Image Analysis to Match a Coin to a Database," *Archaeol. Digit. Era – 40th Annu. Conf. Comput. Appl. Quant. Methods Archaeol.*, pp. 194–198, 2012.
- [36] S. Zambanini, A. Kavelar, and M. Kampel, "Classifying Ancient Coins by Local Feature Matching and Pairwise Geometric Consistency Evaluation," *22nd Int. Conf. Pattern*

- Recognit.*, vol. 1, 2014.
- [37] S. Zambanini and M. Kampel, "A Local Image Descriptor Robust to Illumination Changes," *Image Anal. - 18th Scand. Conf.*, pp. 11–21, 2013.
- [38] R. Huber-Mörk, M. Nölle, M. Rubik, M. Hödlmoser, M. Kampel, and S. Zambanini, "Automatic Coin Classification and Identification," in *Advances in Object Recognition Systems*, no. 2003, InTech, 2012.
- [39] H. Anwar, S. Zambanini, and M. Kampel, "Coarse-grained ancient coin classification using image-based reverse side motif recognition," *Mach. Vis. Appl.*, vol. 26, no. 2–3, pp. 295–304, 2015.
- [40] H. Anwar, S. Zambanini, M. Kampel, and K. Vondrovec, "Ancient Coin Classification Using Reverse Motif Recognition: Image-based classification of Roman Republican coins," *IEEE Signal Process. Mag.*, vol. 32, no. 4, pp. 64–74, 2015.
- [41] J. Kim and V. Pavlovic, "Discovering Characteristic Landmarks on Ancient Coins using Convolutional Networks," Jun. 2015.
- [42] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going Deeper with Convolutions," *arXiv Prepr. arXiv1409.4842*, pp. 1–12, 2014.
- [43] J. Kim and V. Pavlovic, "Ancient Coin Recognition Based on Spatial Coding," in *2014 22nd International Conference on Pattern Recognition*, 2014, pp. 321–326.
- [44] J. Kim and V. Pavlovic, "Improving Ancient Roman Coin Recognition with Alignment and Spatial Encoding," *ECCV Work.*, pp. 1–17, 2014.
- [45] O. Arandjelović, "Reading Ancient Coins: Automatically Identifying Denarii Using Obverse Legend," *ECCV*, pp. 317–330, 2012.
- [46] A. Kavelar, S. Zambanini, and M. Kampel, "Word Detection Applied to Images of Ancient Roman Coins," *Proc. VSMM 2012 - Virtual Syst. Inf. Soc.*, pp. 577–580, 2012.
- [47] A. Kavelar, S. Zambanini, and M. Kampel, "Reading Ancient Coin Legends: Object Recognition vs. OCR," *OAGM/AAPR Work.*, vol. 2013, no. 1, pp. 1–9, Apr. 2013.
- [48] A. Kavelar, S. Zambanini, and M. Kampel, "Reading the legends of Roman Republican coins," *J. Comput. Cult. Herit.*, vol. 7, no. 1, pp. 1–20, Apr. 2014.

- [49] Kai Wang, B. Babenko, and S. Belongie, "End-to-end scene text recognition," *2011 Int. Conf. Comput. Vis.*, pp. 1457–1464, 2011.
- [50] S. Zambanini, A. Kavelar, and M. Kampel, "Improving Ancient Roman Coin Classification by Fusing Exemplar-Based Classification and Legend Recognition," *Lect. Notes Comput. Sci.*, vol. 8158, pp. 149–158, 2013.
- [51] S. Marchand, "IBISA 3D: Image-Based Identification/Search for Archaeology Using a Three-Dimensional Coin Model.," *Archaeol. Digit. Era – Pap. from 40th Annu. Conf. Comput. Appl. Quant. Methods Archaeol.*, pp. 186–193, 2012.
- [52] S. Marchand, "IBISA: Making Image-Based Identification of Ancient Coins Robust to Lighting Conditions," *Proc. 12th EUROGRAPHICS Work. Graph. Cult. Herit.*, pp. 13–16, 2014.
- [53] M. Herrmann, S. Zambanini, and M. Kampel, "Image Based Measurement of Ancient Coins," *Comput. Appl. to Archaeol.*, no. December, pp. 1–7, 2009.
- [54] M. Caltabiano, G. Salamone, M. Puglisi, B. Carroccio, B. Sisalli, A. Celesti, and A. Nucita, "DIANA: an Approach to Coin Iconography According to Time and Space Through Digital Maps," *2013 Digit. Herit. Int. Congr.*, pp. 1–4, 2013.
- [55] V. Kwatra, M. Han, and S. Dai, "Shadow removal for aerial imagery by information theoretic intrinsic image analysis," *2012 IEEE Int. Conf. Comput. Photogr.*, pp. 1–8, 2012.
- [56] R. Guo, Q. Dai, and D. Hoiem, "Paired regions for shadow detection and removal," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 12, pp. 2956–2967, 2013.
- [57] T. Chen, W. Yin, X. S. Zhou, D. Comaniciu, and T. S. Huang, "Total variation models for variable lighting face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 9, pp. 1519–1524, 2006.
- [58] P. Arbeláez, M. Maire, C. Fowlkes, and J. Malik, "Contour Detection and Hierarchical Image Segmentation," *Tpami*, vol. 33, no. 5, pp. 898–916, 2011.
- [59] C. Rother, V. Kolmogorov, and A. Blake, "'GrabCut': interactive foreground extraction using iterated graph cuts," *ACM Trans. Graph.*, vol. 23, no. 3, p. 309, 2004.
- [60] P. Dollár and C. Zitnick, "Fast Edge Detection Using Structured Forests," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 8, pp. 1–13, 2014.
- [61] E. Ö. Faruk, K. Yılmaz, K. Lokman, and T. Ramazan, "A Vision System for Classifying Butterfly

- Species by using Law's Texture Energy Measures," *Int. J. Comput. Vision, Mach. Learn. Data Min.*, vol. 1, 2015.
- [62] G. Takacs, V. Chandrasekhar, S. S. Tsai, D. Chen, R. Grzeszczuk, and B. Girod, "Fast computation of rotation-invariant image features by an approximate radial gradient transform," *IEEE Trans. Image Process.*, vol. 22, no. 8, pp. 2970–2982, 2013.
- [63] S. Matuska, R. Hudec, and M. Benco, "The comparison of CPU time consumption for image processing algorithm in Matlab and OpenCV," *Proc. 9th Int. Conf. ELEKTRO 2012*, pp. 75–78, 2012.
- [64] L. Salgado and C. P. Santos, "Segmentation of Medieval and Ancient Coins Over Complex Backgrounds (a fast automatic method)," in *ICARSC*, 2016.
- [65] M. J. Swain and D. H. Ballard, "Indexing via Color Histograms," in *Active Perception and Robot Vision*, Berlin, Heidelberg: Springer Berlin Heidelberg, 1992, pp. 261–273.
- [66] J. Canny, "A Computational Approach to Edge Detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, no. 6, pp. 679–69, 1986.
- [67] M. Fang, G. Yue, and Q. Yu, "The Study on an Application of Otsu Method in Canny Operator," *Int. Symp. Inf. ...*, vol. 2, no. 4, pp. 109–112, 2009.
- [68] S. N. Tica, B. Costin-Anton, and I. Bucur, "Metallic Coin Analyzer System for Automatic Identification and Classification," no. December, 2015.
- [69] L. Gorelick and Y. Boykov, "GrabCut in One Cut," no. *iccv*, pp. 1–8, 2013.
- [70] H. Bay, T. Tuytelaars, and L. Van Gool, "SURF: Speeded Up Robust Features," *Comput. Vis. Image Underst.*, no. 110, pp. 346–359, 2008.
- [71] E. Rublee, V. Rabaud, K. Konolige, and G. R. Bradski, "ORB : an efficient alternative to SIFT or SURF," *ICCV IEEE Comput. Soc. 2011*, pp. 2564 – 2571, 2011.
- [72] N. Khan, B. McCane, and S. Mills, "Better than SIFT?," *Mach. Vis. Appl.*, vol. 26, no. 6, pp. 819–836, 2015.
- [73] P. F. Alcantarilla, A. Bartoli, and A. J. Davison, "KAZE Features," in *Eur. Conf. on Computer Vision (ECCV)*, 2012.
- [74] E. Tola, V. Lepetit, and P. Fua, "DAISY: An Efficient Dense Descriptor Applied to Wide-Baseline Stereo," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 5, pp. 815–830, May

2010.

- [75] J. Sivic and A. Zisserman, "Efficient visual search of videos cast as text retrieval," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 4, pp. 591–606, 2009.
- [76] C. Grana, D. Borghesani, M. Manfredi, and R. Cucchiara, "A fast approach for integrating ORB descriptors in the bag of words model," *IS&T/SPIE Electron. Imaging. Int. Soc. Opt. Photonics*, pp. 866709–866709, 2013.
- [77] C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, 1995.
- [78] L. Breiman, "Random Forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
- [79] K. Fukunaga, *Introduction to Statistical Pattern Recognition (2nd. ed.)*. Academic Press Professional Inc., 1990.
- [80] Computer Vision Lab - image dataset, "Coin image dataset," 2014. [Online]. Available: <https://www.caa.tuwien.ac.at/cvl/research/cvl-databases/coin-image-dataset/>. [Accessed: 17-Jul-2016].

SUPPORT MATERIAL

A.1 NOTES ON EXAMINED METHODS ON SEGMENTATION CHAPTER

According to the explained approach, the following is a light summary of the most relevant methods scrutinized from proven state of the art papers or from basic computer vision standards, and from tests conducted over the present context. Some examined methods were excluded for not following the established criteria, some provided such weak results (in this context) that dispensed further analyses, others were thoroughly explored in a detail that couldn't fit in such a table.

This is presented here in the hope of being useful as a starting point for future investigations on the subject.

Method	Advantages / Notes	Disadvantages
<u>Segmentation</u>		
Local entropy + Local range of gray values	<ul style="list-style-type: none"> > outperforms adaptive thresholding > outperforms mean shift method > no parameter adjustment 	<ul style="list-style-type: none"> > shadows still pose a problem > border tracing methods can (should?) be used to determine the exact border of the coin.
Structured Random Forests	<ul style="list-style-type: none"> > define a better outline than entropy filter > faster than entropy filter, but still slow 	<ul style="list-style-type: none"> > may not deal so well with complex backgrounds > slow (in OpenCV implementation) > results apparently worse than canny
Entropy filter		<ul style="list-style-type: none"> > slow > often very light and unpredictable results > useless on highly textured backgrounds
Canny edge	<ul style="list-style-type: none"> > outperforms GrabCut in segmentation > Increasing the contrast first (ex. With CLAHE) can enhance image details and facilitate the coin detection 	<ul style="list-style-type: none"> > low resilience to highlights on the edge of the coin > weak results in conjunction with shadow removal methods > applying Canny edge over the 3 RGB channels (and add-weighted) has similar results when applied to gray image (even worse in HSV color space) > Canny doesn't connect pixels into chains or segment.
GrabCut	<ul style="list-style-type: none"> > good results after applying constant L in Lab color space 	<ul style="list-style-type: none"> > outperformed by Canny edge on same papers (ideal conditions) > extremely slow (4-50s) > low resilience to highlights and shadows
Gaussian Mixture Models (GMM)		<ul style="list-style-type: none"> > slow > seems to have low resilience to shadows

		<ul style="list-style-type: none"> > GrabCut (and similar methods) use GMM as a 'color' clustering step > GraphCuts is patented
Hough transform		<ul style="list-style-type: none"> > assumes a perfectly round coin, thus not applicable on ancient coins > very unreliable
Adaptive Thresholding		<ul style="list-style-type: none"> > outperformed by Local entropy + Local range of gray values
Sobel + threshold	<ul style="list-style-type: none"> > better and faster than Scharr > in practice seems better than Laplacian in shadows and small cracks > better resilience to shadows than Laplacian > better results using "Total Gradient absolute (approximate)" than "Total Gradient (approximate)" 	<ul style="list-style-type: none"> > theoretically should be outperformed by Canny edge > Sobel slower than Laplacian > Laplacian over 3 channels seems to have better contours (but less resilient to shadows) > Sobel over the 3 channels (RGB) takes the double time and has similar results
Scharr + threshold	<ul style="list-style-type: none"> > better results using "Total Gradient absolute (approximate)" than "Total Gradient (approximate)" 	<ul style="list-style-type: none"> > low resilience to light gradients > too sensible to background even after fixed Y > usually slower than Sobel and worse results
Laplacian	<ul style="list-style-type: none"> > contours as good as or better than Gradient > Laplacian over 3 channels much better than just over grey image > very fast > Laplacian over 3 channels is faster than Sobel; and better than Sobel for coin outline, but shows more noise. 	<ul style="list-style-type: none"> > low contrast may difficult thresholding (over gray) > low resilience to shadows (worse than Sobel)
Gradient computing	<ul style="list-style-type: none"> > better outlines than Canny edge 	<ul style="list-style-type: none"> > sacrifices shadow zones > highlights on the outer edge remain a problem
Morphological operator: Opening	<ul style="list-style-type: none"> > more resilience to highlights on the edge of the coin 	<ul style="list-style-type: none"> > rounds up the bulges
OpenCV: FindContours()		<ul style="list-style-type: none"> > weak results
Watershed		<ul style="list-style-type: none"> > unhelpful segmentation, it is hard to guess which segments should be joined to form the coin > deforms the border too much > bad results also with shadow removal methods
Mean Shift Filtering		<ul style="list-style-type: none"> > very slow (15 to 200s) when high spatialrad > low resilience to highlights & shadows > very susceptible to compression artifacts > no improvement for fixed Y images (shadow removal)

Saliency Detection Algorithms		<ul style="list-style-type: none"> > usually be slow in big images > hard to discriminate the ROI
Connected Components	<ul style="list-style-type: none"> > good selection of connected components > nice segmentation from background / other features 	<ul style="list-style-type: none"> > too sensitive to shadows and highlights > inconsistent “ideal” threshold (could not find an auto estimative) > CannyEdge defines better edges > still needs border approximation (with convex hull or similar)
Image feature detection using Phase Stretch Transform		<ul style="list-style-type: none"> > US Patented > results no better than canny edge
<u>Edge detection</u>		
Sobel	<ul style="list-style-type: none"> > slightly better than LoG, worse than Canny edge 	<ul style="list-style-type: none"> > uses generic operators and does not consider the image characteristics. > tends to provide inaccurate edge information in the presence of noise > theoretically should be outperformed by Canny edge
Scharr	<ul style="list-style-type: none"> > theoretically as fast as but more accurate than the standard Sobel 	<ul style="list-style-type: none"> > low resilience to light gradients
Prewit		<ul style="list-style-type: none"> > uses generic operators and does not consider the image characteristics > worse than LoG
Laplacian of Gaussian (LoG)		<ul style="list-style-type: none"> > uses generic operators and does not consider the image characteristics. > low resilience to shadows > worse results than Canny
Canny edge	<ul style="list-style-type: none"> > seems to be the best of the basic methods 	<ul style="list-style-type: none"> > uses generic operators and does not consider the image characteristics.
Deriche edge detector	<ul style="list-style-type: none"> > theoretically better than Canny edge > Deriche uses two IIR filters: one for blurring and another for derivative. As it is, IIR filter edge localization is better than Sobel filter but it is time consuming 	<ul style="list-style-type: none"> > slower
Roberts		<ul style="list-style-type: none"> > worse results than Canny Edge
ZeroCross		<ul style="list-style-type: none"> > similar results to LoG
Ant colony optimization (ACO)		<ul style="list-style-type: none"> > based on threshold, just searches for optimal threshold value > presented on modern coins, good illumination > ACO might not be the fastest algorithm available

Structured Random Forests	> capable of real time frame rates (faster than most competing state of the art methods) while achieving state of the art accuracy	> Results similar to Canny edge but with double lines
Laplacian	> good results when the coins has good contrast	> bad contrast = bad results
Morphological operator: Gradient	> better outlines than Canny edge	> sacrifices shadow zones > highlights on the outer edge remain a problem
Gradient computing		> outlines sometimes worse than Canny edge sometimes better > sacrifices shadow zones
Difference of Gaussians (DoG)		> results similar to Gradient but less well defined and penalizing in the shadows > generally bad with textures
Morphological operator: Top Hat	> outline very well defined	> too penalizing in the shadows, even more than Gradient > in dark images may lose detail
Wavelet Transform (Wavelet-Based Edge Detection)		> weak results > blurs instead of lines > too variable
<u>Shadow Removal</u>		
Shadow removal by information theoretic intrinsic image analysis		> patented
Paired Regions for Shadow Detection and Removal		> too slow (several minutes) > needs other related images > too shallow effect
Equalize shadow		> can infatuate noise
Through RGB and HSB values pixel to pixel		> very weak results
Fix V value in HSV	> helpful in some cases	> weak results for most cases > sometimes there is severe "blocky" effect in the result (jpg losses) > does not help for GrabCut > may help for borders with Canny but not for edges
Fix V value in HSV and equalize S	> may improve color frontiers (careful with coins with dust similar to background)	> tested following by Canny and GrabCut with poor results > in general following Canny appears to have worse results than without equalization
Fix L in Lab channels	> very good results when well contrasted image	> too low variance in many cases

	> not as strong "blocky" effect as in HSV	> bad results when similar coin and background colors are present in image
Fix Y value in YCrCb	> results similar to Lab, some better, some worse	