

# Reading Ancient Coins: Automatically Identifying Denarii Using Obverse Legend Seeded Retrieval

Ognjen Arandjelović

Swansea University, Wales, UK  
ognjen.arandjelovic@gmail.com

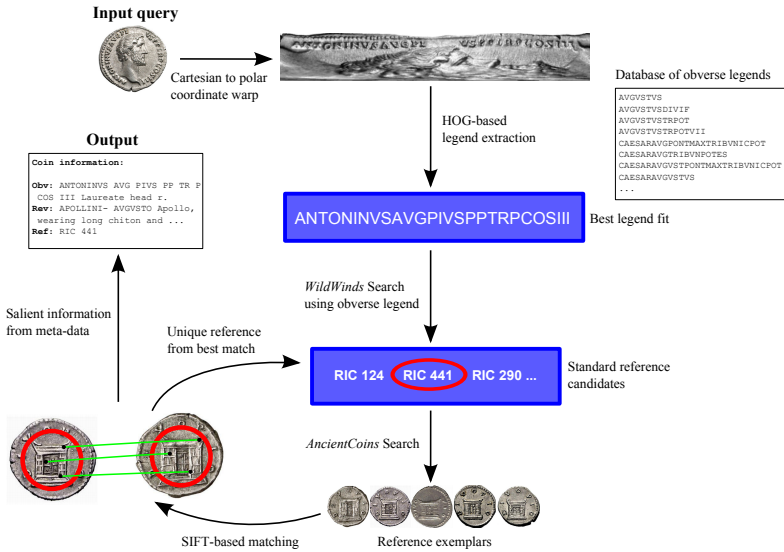
**Abstract.** The aim of this paper is to automatically identify a Roman Imperial denarius from a single query photograph of its obverse and reverse. Such functionality has the potential to contribute greatly to various national schemes which encourage laymen to report their finds to local museums. Our work introduces a series of novelties: (i) this is the first paper which describes a method for extracting the legend of an ancient coin from a photograph; (ii) we are also the first to suggest the idea and propose a method for identifying a coin using a series of carefully engineered retrievals, each harnessed for further information using visual or meta-data processing; (iii) we show how in addition to a unique standard reference number for a query coin, the proposed system can be used to extract salient coin information (issuing authority, obverse and reverse descriptions, mint date) and retrieve images of other coins of the same type.

**Keywords:** Recognition, Text, Image, Reverse, Motif, Inscription.

## 1 Introduction

The aim of this paper is to automatically identify a Roman Imperial denarius from a single query photograph of its obverse (“front”) and reverse (“back”). Specifically, we wish to infer the issuer of the coin (usually the emperor depicted on the obverse), textual descriptions of its obverse and reverse, its reference identifier in the standard reference work “Roman Imperial Coinage” (RIC) [1] and the year it was minted.

**Motivation.** Our primary motivation comes from the immense value that this functionality would bring to such projects as the “Portable antiquities scheme” [2]. This scheme, pioneered in England and Wales, encourages the general public (primarily metal detectorists) to report their archeological findings to local museums for the sake of obtaining a record of the relevant and potentially valuable details of the find, without confiscating the find. It has been an immense success. In fact, the scheme has been so popular that the major limitation at present is the ability to process the large volume of finds, most of which are coins, and which are individually identified by an expert. Identification by the



**Fig. 1.** (i) After the original image of a coin’s obverse is mapped from Cartesian to quasi-polar coordinates, the corresponding legend is extracted using a HoG-like descriptor and an exhaustive database of possible legends; (ii) the extracted legend is used in a *WildWinds* search to select potential candidate coin types, whose unique reference IDs are used to retrieve matching images using the *AncientCoins* search; (iii) the correct coin type is determined by SIFT matching the query coin’s reverse with the reverses of the retrieved exemplars for each candidate type; (iv) finally, salient coin information is extracted by integrating meta-data of all *AncientCoins* search results for the correct coin type.

finder is unsatisfactory: most of them are laymen, without the necessary expertise or access to specialist literature, and the risk of erroneous data entry would be unacceptably high. Our goal is to develop an automated online system which could process submitted images of coins. Such system would greatly reduce the burden on the experts, while at the same time making the “Portable antiquities scheme” even more widely accessible. Indeed, even a simple Google search readily reveals a plethora of requests to help identify a Roman coin.

**Previous Work and Its Limitations.** Computer vision work on the analysis of ancient coins is still scarce, with most of the previous work focusing on modern coins instead [3–6]. It is only in recent years that ancient coins in particular have started attracting attention of the community, through collaborations with museums and organizers of programmes such as the “Portable antiquities scheme” described previously. All of the published computer vision work focused on the analysis of ancient coins – by Zaharieva *et al.* [7], Kampel and Zaharieva [8] and Arandjelović [9] – aim to match coins using a variation of SIFT-like local features, the results universally demonstrating the challenge involved in this task. Another similarity, and we argue limitation, of these methods is that they treat

the entire area of a coin in exactly the same manner, as an appearance pattern. However, in doing so they fail to optimally exploit what is a characteristically very rich source of information on Roman Imperial coins: the legend (i.e. textual inscription) around the coin edge. The method described in the present paper makes the extraction of the legend its first step, which is followed by a sequence of retrievals, each of which is used to gather further visual or meta-data, until a unique coin type is identified.

**Overview of the Proposed Method.** Our system starts by extracting the obverse legend of a coin from its image. We select the legend from a database of 1478 legends<sup>1</sup> using a HoG-like descriptor to describe the appearance of an individual letter and a spatial model which constrains the relative locations of neighbouring letters. The extracted legend is used as the initial seed for a sequence of retrievals. The results of each retrieval, some visual some textual (meta-data), are used to further constrain the range of possible coin identities. First, we use the obverse legend to perform a *WildWinds* [10] search which explicitly retrieves all references in RIC with the same legend. Next, a new retrieval for each candidate reference is performed using the *AncientCoins* search [11], which indexes a greater number of entries and coin exemplar images. The correct type is chosen by visually matching the query coin against the retrieved reverse motifs. Finally, the meta-data of the matching type is processed and salient coin information extracted.

## 2 Method Details

We start with the problem of extracting the legend from an image of the obverse of a coin. As illustrated in Fig. 2 (a), the legend on Roman Imperial coins runs circularly around the edge of the coin. In this example it reads ANTONINVS AVG PIVS PP TR P COS III and it comprises a series of words or abbreviations. There are several features of obverse legends which are important to observe. Firstly, the legend is extraordinarily rich in information. ANTONINVS designates the issuer (here Antoninus Pius), usually depicted on the central part of the obverse; AVG (Augustus), PIVS, PP (father of the country), TR P (tribune of the people) and COS III (consul for the third time) all designate different titles of the issuer which can be used to constrain the coin's mint date. For example, in this case COS III can be used to constrain the mint date to 143-144 AD. Lastly, note the “-” symbol in the string used to describe the legend. It denotes a so-called “legend break” which is a point at which there is a gap in the physical inscription, usually due to a feature of the coin's design. Two inscriptions are usually considered as corresponding to the same legend even if their legend breaks are differently positioned.

### 2.1 Geometric Normalization

The obverse legend on a coin runs circularly near the coin's edge. This means that the orientation of letters varies across the entire range from 0° to 360°

<sup>1</sup> The database is available for download from <http://mi.eng.cam.ac.uk/~oa214/>

which is not ideal from the point of view of matching efficiency. Instead, we estimate the location  $\mathbf{x}_c = (x_c, y_c)$  of the coin's centre and its radius  $r$  using the Hough transform as described in [9], and then transform the image from Cartesian to quasi-polar coordinates, as illustrated in Fig. 2. A point at the location  $\mathbf{x} = (x, y)$  in the original image  $J$  is thus mapped to the point  $\mathbf{x}' = (x', y')$  in the geometrically normalized image  $I$ :

$$y' = s \cdot \sqrt{(y - y_c)^2 + (x - x_c)^2} \quad (1)$$

$$x' = \begin{cases} s \cdot 2r \arccos((y_c - y)/|\mathbf{x} - \mathbf{x}_c|) & x < x_c \\ s \cdot 2r (2\pi - \arccos((y_c - y)/|\mathbf{x} - \mathbf{x}_c|)) & x \geq x_c \end{cases} \quad (2)$$

The scaling factor  $s$  is used to ensure the uniform scale of  $942 \times 150$  pixels. Lastly, note that the  $x$  coordinate in the processed image starts at the line extending from the coin's centre downwards, just as the legend does.



(a) Query obverse



(b) Normalized query obverse

**Fig. 2.** (a) Original image of a query coin's obverse and (b) the corresponding geometrically normalized image we use to extract the legend

## 2.2 Legend Extraction

We treat the problem of extracting the obverse legend of a coin as one of optimal hypothesis choice, each hypothesis corresponding to a particular legend of  $n_l$  legends in our database. Thus we wish to find the index  $i^*$  of the hypothesis such that:

$$i^* = \arg \max_{1 \leq i \leq n_l} p(l_1^{(i)}, \dots, l_{n_l}^{(i)} | I) \quad (3)$$

where  $n_l$  is the number of letters in the  $i$ -th legend and  $l_1^{(i)} \dots l_{n_l}^{(i)}$  are the corresponding letters in order.

It is certainly not the case that different legends are equally common. Some legends are shared amongst more coin types; certain coin types are also rarer than others. Indeed, RIC provides a rarity guide (C = common, S = scarce, R = rare, R2 = very rare, and R3 = extremely rare), estimated from a range of museum collections. However, we argue that this would not form a good basis for a prior in the present work: museums have biases in their interests and the relative frequencies of different types of coins likely to be submitted by lay users is difficult to predict. Thus we adopt an uninformative prior which makes our choice a maximum likelihood test:

$$i^* = \arg \max_i p(l_1^{(i)}, \dots, l_{n_i}^{(i)} | I) = \arg \max_i p(I | l_1^{(i)}, \dots, l_{n_i}^{(i)}) \quad (4)$$

where also:

$$p(I | l_1^{(i)}, \dots, l_{n_i}^{(i)}) = \max_{\substack{x_1, \dots, x_{n_i} \\ y_1, \dots, y_{n_i}}} p(I | l_1^{(i)}, \dots, l_{n_i}^{(i)}, x_1, \dots, x_{n_i}, y_1, \dots, y_{n_i}) \quad (5)$$

The estimation of the likelihood  $p(I | l_1^{(i)}, \dots, l_{n_i}^{(i)})$  is computationally complex in no small part because of the potential presence of legend breaks which can have a range of widths and which can in principle occur between any two consecutive letters. To make the problem tractable, we propose a two stage approach: (i) first, we estimate the optimal placement of the legend letters using only the evidence from the corresponding image patches and a spatial constraint on consecutive letters, and then (ii) evaluate the likelihood for the entire image strip of the legend, taking into account how well the appearance of legend breaks is explained too. Thus, to find the optimal placement of letters  $(x_1, y_2), \dots, (x_{n_i}, y_{n_i})$  we maximize the following likelihood:


$$\hat{P}(x_1, \dots, x_{n_i}, y_1, \dots, y_{n_i}) = p(I_{x_1, y_1} | l_1^{(i)}) \prod_{j=1}^{n_i-1} p(I_{x_{j+1}, y_{j+1}} | l_{j+1}^{(i)}) p(x_{j+1}, y_{j+1} | x_j, y_j) \quad (6)$$

where  $I_{x_j, y_j}$  is a letter-sized image patch centred at  $(x_j, y_j)$ . Our spatial prior on the locations of consecutive letters is given by:

$$p(x_{j+1}, y_{j+1} | x_j, y_j) \propto \begin{cases} 1 & t_{x1} < x_{j+1} - x_j < t_{x2} \text{ and } |y_{j+1} - y_j| < t_y \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

The primary function of the the thresholds  $t_{x1}$  and  $t_y$  is to eliminate implausible relative letter placements. In contrast, the threshold  $t_{x2}$  is chiefly used for computational reasons, i.e. to restrict the image search area. We set  $t_{x1}$  to 80% of the letter width,  $t_y$  to 20% of the letter height and  $t_{x2}$  to six times the letter width. Our appearance model used to estimate individual letter likelihoods  $p(I_{x_j, y_j} | l_j^{(i)})$  is explained next.

**Letter Appearance Model.** The appearance of a particular letter in a legend can exhibit great variability. Firstly, legend letters are small features which were manually engraved without the use of any magnifying instruments. This means that both their shape and orientation can change significantly from instance to instance. Letter appearance is also affected by illumination, wear, strike in the minting process etc. We experimented with a number of possible representations including raw and filtered appearance, and wavelet based features, with limited success. An approach based on HoG features [12] was found to be the most successful one and it is what we adopt henceforth.



**Fig. 3.** (a) Letter sized patch is divided into nine overlapping sub-regions each covering a quarter of the area of the entire patch. Shown is a sub-region (red) and two horizontally and vertically adjacent sub-regions. (b) Initially, the optimal placement of the letters for a particular hypothesized legend is determined using image evidence corresponding to letter regions only (shown in red). The likelihood of the hypothesis is subsequently estimated using evidence both from letter image regions as well as legend breaks found between letters (shown in blue). Thus a legend hypothesis which explains letter regions well but produces gaps which do not look like actual breaks, does not result in a high likelihood value.

We represent each patch as a feature vector obtained by concatenating nine weighted histograms corresponding to different sub-regions of the patch. Each histogram comprises nine bins (gradient directions) over  $180^\circ$ , gradients at  $180^\circ$  from one another contributing to the same bin. We make each patch sub-region half as wide as the entire patch, with two vertically or horizontally adjacent sub-regions sharing a 50% overlap, as shown in Fig. 3 (a). That gives three possible vertical and horizontal displacement and hence  $3 \times 3 = 9$  sub-regions and a  $9 \times 9 = 81$  dimensions for the concatenated descriptor.

In comparison to the original descriptor proposed by Dalal and Triggs, our descriptor contains several modifications which make it more suitable for the problem at hand. Firstly, we do not weight gradient contributions in proportion to their magnitude. We found that due to small irregularities on the coin's surface and the small physical size of the surface represented by a patch, gradients of small magnitude can build up and obscure the main features in the patch that we seek to describe. Instead, we use weights obtained by transforming gradient magnitudes using a sigmoid function of the following form:

$$\hat{g} = \frac{1}{1 + \exp\{-0.5(g - 0.5g_{max})\}} \quad (8)$$

where  $g_{max}$  is the maximal gradient magnitude of the entire patch (not a specific sub-region),  $g$  the original gradient magnitude and  $\hat{g}$  the transformed value used to weight bin contributions. Another difference introduced here is that unlike Dalal and Triggs, we do not normalize histograms of different sub-regions. Rather, sub-region histograms are concatenated in their raw form. We found that this produced superior results on our problem.

Using a manually annotated corpus of 30 images of coins with the associated legends, we learn the range of variation of each letter's descriptor as a

multivariate Gaussian distribution. The distribution is over an 81D space in which we take the first 5 eigenvector directions as the basis of the principal subspace. The corresponding largest eigenvalues (variances) are left unchanged. Since the remaining eigenvalues are assumed to come from random noise sampling they are averaged, thus ensuring that the Kullback-Leibler divergence between the true distribution and its estimate is minimized [13]. The likelihood  $p(I_{x_j, y_j} | l_j^{(i)})$  is then evaluated by first computing the HoG-like descriptor of the image patch  $I_{x_j, y_j}$  and then the corresponding value of the Gaussian representing letter  $l_j^{(i)}$ . For each locus  $(x_j, y_j)$  we compute the likelihood at three scales (letter heights of 18, 22 and 26 pixels) and assign the largest of these to  $(x_j, y_j)$ .

**Inferring Optimal Letter Placement.** Using the introduced appearance and spatial models, the maximum likelihood solution to Equation 6 can be computed exactly and efficiently using dynamic programming. If  $L_{k+1}^{(i)}(x, y)$  the maximum likelihood of the  $i$  partial legend up to its  $(k + 1)$ -st letter:

$$k = 0 : L_{k+1}^{(i)}(x, y) = p(I_{x_1, y_1} | l_1^{(i)}) \tag{9}$$

$$k > 0 : L_{k+1}^{(i)}(x, y) = \max_{\substack{x_1, \dots, x_k \\ y_1, \dots, y_k}} p(I_{x_1, y_1} | l_1^{(i)}) p(I_{x, y} | l_{k+1}^{(i)}) p(x, y | x_k, y_k) \times \tag{10}$$

$$\times \prod_{j=1}^{k-1} p(I_{x_{j+1}, y_{j+1}} | l_{j+1}^{(i)}) p(x_{j+1}, y_{j+1} | x_j, y_j) \tag{11}$$

then the following recurrence holds:

$$k = 0 : L_1^{(i)}(x, y) = p(I_{x, y} | l_1^{(i)}) \tag{12}$$

$$k > 0 : L_{k+1}^{(i)}(x, y) = p(I_{x, y} | l_{k+1}^{(i)}) \times \tag{13}$$

$$\times \max_{\Delta x, \Delta y} L_k^{(i)}(x - \Delta x, y - \Delta y) p(x, y | x - \Delta x, y - \Delta y) \tag{14}$$

In other words, the maximal likelihood of a partial legend which places its last letter at a particular location in the image can be computed by scanning the area of possible loci for the preceding letter, and updating the corresponding maximal likelihood value.

**Estimating Legend Likelihood.** The proposed dynamic programming based approach to estimating the likelihood in Equation 6 accounts only for evidence of image patches which correspond to the loci of the legend letters, as illustrated by red rectangles in Fig. 3 (b). There are two key reasons why this likelihood is not a good approximation to the likelihood in Equation 5 and thus not a good criterion for selecting the best fitting legend:

- generally it tends to penalize legends with a greater number of letters, and
- it fails to account for the appearance of spaces between letters, which should look as legend breaks.

Consequently, our approach continues as follows. After the optimal letter placements of a legend are estimated using Equation 6, we fill any significant gaps

between consecutive letters (greater than 80% of the letter width) with letter sized patches, as illustrated with blue rectangles in Fig. 3 (b). These should contain the appearance of legend breaks. Unlike in the case of letters, we do not learn the appearance model for the legend break because we know that in the idealized coin specimen they should be blank. In other words, considering that we do not perform the block normalization of Dalal and Triggs, all of the entries in the corresponding HoG-like feature should be close to zero. This observation allows us to compute the likelihood of a hypothesized legend break patch  $I_{x,y}$  at  $(x,y)$  using a zero mean isotropic Gaussian whose covariance we estimate as the mean noise covariance across the distributions representing individual letters. The likelihood of a particular legend thus becomes:

$$P(x_1, \dots, x_{n_i}, y_1, \dots, y_{n_i}) = \hat{P}(x_1, \dots, x_{n_i}, y_1, \dots, y_{n_i}) \prod_{j=1}^{\hat{n}_i} p_b(I_{\hat{x}_j, \hat{y}_j}) \quad (15)$$

where  $\hat{n}_i$  is the number of breaks in the  $i$ -th legend (note that  $n_i + \hat{n}_i = \text{const.}$ ),  $p_b$  the likelihood of a break corresponding to a letter sized patch, and  $I_{\hat{x}_j, \hat{y}_j}$  a letter sized patch at a hypothesized legend break location  $(\hat{x}_j, \hat{y}_j)$ .

### 2.3 Making a Shortlist of RIC Identifiers

The free *WildWinds* coin search engine allows the user to retrieve RIC identifiers of coin types that match a particular legend fragment, disregarding the positions of legend breaks. This means that a search using the legend ANTONINVS AVG PIVS P P T R P C O S I I I I (see Fig. 2) correctly finds the types RIC 70, RIC 612, RIC 660 and RIC 716, all which have the correct query legend at the obverse. However, it also finds the type RIC 415 with the legend ANTONINVS AVG PIVS P P T R P C O S I I I I I (note the extra “I”, signifying the fourth consulship year). To overcome this limitation of the search engine, we perform retrieval using multiple queries. First, we use the extracted legend as the query and obtain the set of possible matches,  $\mathcal{S}_0$ . In addition we also search using each of the  $d_{i^*}$  entries in our legends database which contain the extracted legend as a sub-string, obtaining further sets of matches,  $\mathcal{S}_1, \dots, \mathcal{S}_{d_{i^*}}$  say. These results allow us to infer the correct shortlist of identifiers as the set difference  $\mathcal{S}^* = \mathcal{S}_0 \setminus \bigcup_{j=1}^{d_{i^*}} \mathcal{S}_j$ .

### 2.4 Visual Sifting by Matching Reverse Motifs

As we explained earlier, the obverse legend of a Roman coin is typically very rich in information content. However, it is also seldom sufficient to uniquely identify a coin. Indeed, a particular legend is usually found on many different types; the legend ANTONINVS AVG PIVS P P T R P C O S I I I I I, for example, occurs on over twenty.

A coin type is characterized by particular obverse and reverse legends and central motifs. Two coins which match in these features are considered to be of the same type. The obverse motif on Roman denarii and aurei is universally a portrait of the coin’s issuer, shown in profile, and it provides little additional information over the corresponding legend regarding the coin’s type<sup>2</sup>. Thus, we

<sup>2</sup> A more detailed treatment of this issue is out of scope of the present paper.





**Fig. 4.** A random sample of six reverses retrieved in an *AncientCoins* search using the automatically generated query *Antoninus Pius* ("RIC 441" "R.I.C. 441"). The reverse motif of the query coin is matched against the set of retrieved reverses. The overall matching score of the query coin with the type is estimated as the highest of the corresponding individual matching scores.

focus on the content shown on the reverse to disambiguate the matching of our query coin against the shortlist  $S^*$  of its possible types. Specifically, we match reverse motifs, disregarding the reverse legend. Unlike in the case of obverse legends, the list of possible reverse legends is far greater and to the best of the knowledge of these authors, no such list has been compiled, which prohibits us from applying the approach described in Sec. 2.2.

To obtain exemplars of reverses of a particular coin type we employ the free *AncientCoins* search engine [11], which retrieves coins from a wide range of coin dealers' web sites and past auctions by matching a textual query with the text associated with each coin. Using this search engine with a simple query comprising the name of the coin's issuer (determined from the obverse legend, as explained in Sec. 2.2) and a particular RIC reference from the shortlist  $S^*$  we retrieve exemplar images of coins of the corresponding type. An example of six retrieved reverses is illustrated in Fig. 4. Note the variability in both the style and positioning of the legend, as well as the central motif (an altar in this case).

**Registration.** Our approach to matching the reverse of our query coin with each of the retrieved reverses comprises two stages. First, we register the motifs of the two reverses which are being compared. This is necessary because the precise positioning of the motif can vary significantly across different dies of the same type, as can be readily observed in Fig. 4. We use Euclidean registration and estimate its two parameters by matching SIFT descriptors. Following Lowe's recommendation [14], we accept a SIFT keypoint match in the query reverse with its closest (in terms of feature similarity) keypoint in a retrieved reverse, if the distance of the second closest keypoint is at least 1.5 greater. We apply this keypoint matching in a RANSAC framework so as to eliminate the effects of spurious matches and pool the estimates of correctly matched keypoints to achieve more robust registration.

**Appearance Matching.** After the two reverses are registered and their reverse motifs aligned, they are compared in appearance. Here too we employ SIFT features. We try to match each detected feature in the query coin's reverse with a feature detected in the reverse of the coin it is compared with, subject to appearance and spatial criteria. First, we require that the similarity of two features (as

a normalized dot product of the corresponding feature vectors) exceeds a threshold. Also, we require that the two features are within a specific distance from each other (in our implementation the maximal distance is set to 20 pixels), and in agreement in scale (within 20%) and direction (within 30°). The similarity of two reverses is then measured by the number of matched feature pairs.

After each of the reverses retrieved using a search for a particular RIC type is compared against the query reverse, we compute the overall matching confidence for the type as the maximum of all the computed similarities. Finally, the correct RIC type match is chosen as the one with the highest matching confidence.

## 2.5 Extraction of Salient Coin Information from Textual Meta-data

The first aim of the present paper was to uniquely identify the query coin's reference in RIC. To a proficient numismatist, this reference contains sufficient information which can be used to look up further relevant details, such as the coin's mint date. However, there are several reasons why it is advantageous to do this automatically, which we set out to do here. First, it saves time needed to look up a reference and then manually enter relevant detail. It also gives immediate and more readily understandable feedback which can be used to check for the correctness of the result. Lastly, it provides the lay user, who may be submitting his/her find online, a more satisfying and meaningful description of the find.

We specifically seek to extract textual descriptions of the obverse and reverse motifs, as well as the mint date of the coin. For this we use textual meta-data associated with the coins already retrieved using the *AncientCoins* search with the correct RIC reference. Any retrieved text which is not in English, we translate into English using Google's automatic translator and replace various delimiters by "white space". Examples include hyphens, which are used to denote legend breaks, and square brackets which signify that the enclosed part of the legend is missing, for example because it has been damaged or because it is off the flan of the coin.

Obverse and reverse descriptions are localized in text using explicit rules which reflect a number of standard conventions used in describing ancient coins. For example, the obverse description may be located as the sentence which contains the obverse legend extracted in Sec. 2.2, or the sentence which follows the word "obverse", its abbreviation "obv" or indeed "av", the abbreviation for obverse used in German and French (for "avers") and which is not automatically translated by Google. A description of the obverse and reverse thus may be extracted from every retrieved coin record. However, some of these may be incorrect as the search string comprising the RIC type of the query coin may in some instances occur even in records of coins of a different type. For example this may be because a coin is in some sense compared with the query coin type (rarer, similar, and so on). Thus, we wish to choose the best of extracted descriptions. We achieve this by creating a histogram of words across the corpus of all extracted descriptions after eliminating indiscriminative words (e.g. "in", "left", "emperor", "head", "bust"), and then selecting the best fitting legend as the one with the words of the highest average frequency in the corpus.

The manner in which we obtain the mint date (or more generally, mint period) of the coin involves a different strategy. There are two key problems that we had to address here. The first is that some records contain incorrect mint dates. The second is that some coin records do not contain the most precise (narrowest) mint period found in specialist literature, but a broader one. For example, some entries will simply have the entire period of the issuer's reign as the mint period (e.g. 138-161 AD). We solve both of the aforementioned problems as follows. Initializing the algorithm with the issuer's rule period, each time a candidate period is extracted from a coin record we fragment the range of possible periods, so that each fragment begins and ends at the beginning or the end of an extracted period. Then, we choose the fragment with the most votes (most overlapping periods extracted from coin records), as the correct one.

### 3 Results

The coin identification system described in this paper was evaluated on 25 coins. These coins were identified by an expert. Relevant ground truth information – the coin's issuer, its obverse and reverse legends, the descriptions of its obverse and reverse motifs and the minting date – was obtained from RIC.

#### 3.1 Legend Extraction

We first examined the performance of our method for extracting the obverse legend, described in Sec. 2.2. For all but one coin, the correct legend was inferred. The one incorrect result was caused by a particularly challenging relative placements of two letters. Specifically, the letter I representing the Roman numeral one, was engraved unusually close to the preceding letter. Consequently, the appearance of the preceding letter contributed to the feature vector extracted from a letter sized patch centred at I, producing a low likelihood score at that location for all letters. Since there is a valid legend identical to the correct one in all respects except that it does not contain the problematic I (i.e. the same form of the legend for the previous consulship year), this legend was selected as the highest likelihood one.

#### 3.2 RIC Types Shortlisting

Providing that the correct obverse legend was extracted in the previous stage of the algorithm and that the correct RIC type is not so rare as to be absent from the *WildWinds* database, our method of creating a shortlist of possible types is guaranteed to include the correct type. Thus, as expected, for the 24 test coins for which the obverse legend was correctly extracted, the ground truth RIC type was amongst the shortlisted ones. Equally, the correct type of the coin for which the legend was not correctly extracted, was not amongst the shortlisted types.

### 3.3 Visual Sifting by Matching Reverse Motifs

Of the 24 coins for which the obverse legend extraction and shortlisting produced correct results, 22 were matched with the correct RIC type based on the appearance of the reverse motif. A few representative examples are shown in Fig. 6. An example of an incorrect match is shown in Fig. 5. It can be readily seen that the matched motifs, although not the same, bear a high degree of resemblance – both feature a standing figure, holding a small object (patena and wand respectively) in the extended right arm and a long stick-like object in the left (spear and sceptre), with a further object at feet (altar and globe). It is equally interesting to notice that the two types are readily differentiated from one another by their reverse legends. While the query reverse legend reads RESTITVTOR VRBIS, that of the incorrectly matched type is PROVID AVGG.








**Fig. 5.** An example of an incorrect type match. Shown is (a) the correct type RIC 166 and (b) the incorrect type RIC 167 that the query coin was matched to instead

Considering that the correct RIC type was not in the shortlist of possible types for the one coin whose obverse legend was not correctly extracted, the end type it was matched to could not be correct. However, the coin was matched to the correct reverse motif, which means that in every respect except for the one missing letter of the obverse legend, our method was successful. Indeed, the obverse and reverse motifs were commonly repeated across different consulship years, which means that in most cases in which the obverse legend extraction fails due to unintelligibility of Roman numerals, a nearly identical if not entirely correct type will be found. This is highly comforting as it is reasonable to expect that most errors in our legend extraction algorithm will be caused precisely in the matching of numerals because the contextual constraints are much looser in comparison with, say, the name of the emperor in the legend.

### 3.4 Meta-data Parsing

Examples of automatically extracted textual descriptions of the key coin facts can be seen in the central column of Fig. 6. In all cases, the extracted information correctly matched the identified coin type. The only problem we observed with this stage of our system pertains to limitations of Google’s automatic translator when dealing with words which are rarely used in everyday speech but are

Query	Description	Example specimen
	<p>Issuer: Antoninus Pius                      Obverse: DIVVS ANTONINVS Bare head of Antoninus Pius to right.                      Reverse: CONSECRATIO Altar with two closed doors.                      Minted: 161 AD                      Reference: RIC 441</p>	
	<p>Issuer: Septimius Severus                      Obverse: SEVERVS PIVS AVG, bust right belorbeerte.                      Reverse: PM TRP XVII COS III PP, Jupiter stands left between two children.                      Minted: 209 AD                      Reference: RIC 226</p>	
	<p>Issuer: Faustina I                      Obverse: DIVA FAVSTINA, bust draped right.                      Reverse: AETERNITAS, draped and veiled female figure standing right, head left, raising right hand and holding scepter in left.                      Minted: 141 AD                      Reference: RIC 344</p>	

**Fig. 6.** Examples of typical end results of our system. The left-hand column shows query coins, the central column its RIC type and automatically extracted obverse and reverse descriptions, and the right-hand column a further example of the same type obtained using the free *AncientCoins* search engine.

frequent in numismatics. For example, note the German word “belorbeerte” (laureate) which was not translated in the description of the obverse of the second coin in Fig. 6. That happens if the coin entries of a specific type are predominantly in a foreign language – an untranslated word may feature in the majority of extracted descriptions and thus be included in the description which best matches the entirety of the retrieved meta-data.

## 4 Conclusions and Future Work

This paper introduced the first automatic system which can identify a Roman denarius from a single photograph. The system comprises a cascade of steps, each aimed at extracting additional information which allows the range of possible coin types to be reduced further. The extraction of the obverse legend, a problem also addressed here for the first time, is crucial as the legend is used to initiate a series of public search engine retrievals, each of which is used to harness new information. The first search is used to create a shortlist of possible types based on the obverse legend alone. The second search is used to obtain images of exemplar coins for each type. The reverse motifs of these coins are matched with the reverse of the query coin, the best matching type eventually being selected as the correct match. The

associated textual meta-data is further used to extract salient coin information: descriptions of its obverse and reverse motifs, and mint date period.

Our experiments demonstrated highly encouraging results and highlighted the most promising directions for further improvement. We first aim to investigate different letter appearance representations, which would allow to extract not only the obverse legend but also the highly discriminative reverse legend too. This would also allow us to extend our statistical model used to match obverse legends to handle more robustly partially damaged legends, which the method proposed in this paper does not do. Lastly, the occasional failure of our approach in matching reverse motifs and its sensitivity to the precise coin specimens retrieved, add to the corpus of evidence of previous research that the development of features more specific to the particular problem at hand, rather than generic SIFT features, is another promising research avenue.

## References

1. Webb, P.H. (vol. I), Mattingly, H., Sydenham, A., Sutherland, C.H.V. (vol. II-III), Sutherland, C.H.V., Carson, R.A.G. (vol. VI-IX), Carson, R.A.G., Kent, J.P.C., Burnett, A.M. (vol. X) (eds.): *Roman Imperial Coinage*, vol. I–X. Spink, London (1923-1994)
2. The portable antiquities scheme, <http://finds.org.uk/> (last accessed July 2012)
3. Davidsson, P.: Coin classification using a novel technique for learning characteristic decision trees by controlling the degree of generalization. In: Proc. IEA/AIE, pp. 403–412 (1996)
4. Mitsukura, Y., Fukumi, M., Akamatsu, N.: Design and evaluation of neural networks for coin recognition by using GA and SA. In: Proc. IJCNN, vol. 5, pp. 178–183 (2000)
5. Huber, R., Ramoser, H., Mayer, K., Penz, H., Rubik, M.: Classification of coins using an eigenspace approach. *Pattern Recognition Letters* 26(1), 61–75 (2005)
6. van der Maaten, L., Boon, P.: COIN-O-MATIC: A fast system for reliable coin classification. In: Proc. MUSCLE CIS Coin Recognition Competition Workshop, pp. 7–18 (2006)
7. Zaharieva, M., Kampel, M., Zambanini, S.: Image Based Recognition of Ancient Coins. In: Kropatsch, W.G., Kampel, M., Hanbury, A. (eds.) CAIP 2007. LNCS, vol. 4673, pp. 547–554. Springer, Heidelberg (2007)
8. Kampel, M., Zaharieva, M.: Recognizing Ancient Coins Based on Local Features. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Remagnino, P., Porikli, F., Peters, J., Klosowski, J., Arns, L., Chun, Y.K., Rhyne, T.-M., Monroe, L. (eds.) ISVC 2008, Part I. LNCS, vol. 5358, pp. 11–22. Springer, Heidelberg (2008)
9. Arandjelović, O.: Automatic attribution of ancient Roman imperial coins. In: Proc. CVPR, pp. 1728–1734 (2010)
10. WildWinds graphical partial legend search engine, <http://www.wildwinds.com/coins/findstr.html> (last accessed July 2012)
11. Ancient coins search engine, <http://www.acsearch.info/> (last accessed July 2012)
12. Dalai, N., Triggs, B.: Histograms of oriented gradients for human detection. In: Proc. CVPR, vol. 1, pp. 886–893 (2005)
13. Tipping, M.E., Bishop, C.M.: Probabilistic principal component analysis. *Journal of the Royal Statistical Society* 61(3), 611–622 (1999)
14. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *IJCV* 60(2), 91–110 (2003)