



# ICONOGRAPHY ON ANCIENT COINS: DECIPHERING SYMBOLISM AND MEANING

Vilin U. Chauhan

*Department of Master of Computer Application,  
G H Rasoni University ,Amravati , India*

**Abstract:** In recent years, a range of problems within the broad umbrella of automatic, computer vision based analysis of ancient coins has been attracting an increasing amount of attention. Notwithstanding this research effort, the results achieved by the state of the art in the published literature remain poor and far from sufficiently well performing for any practical purpose. In the present, paper we present a series of contributions which we believe will benefit the interested community. Firstly, we explain that the approach of visual matching of coins, universally adopted in all existing published papers on the topic, is not of practical interest because the number of ancient coin types exceeds by far the number of those types which have been imaged, be it in digital form (e.g. online) or otherwise (traditional film, in print, etc.). Rather, we argue that the focus should be on the understanding of the semantic content of coins. Hence, we describe a novel method which uses real-world multimodal input to extract and associate semantic concepts with the correct coin images and then using a novel convolutional neural network learn the appearance of these concepts. Empirical evidence on a real-world and by far the largest data set of ancient coins, we demonstrate highly promising results.

**Index Terms:** Indian coins, certain mystery, punch marked coins, Semantic element, Machine learning

## I. INTRODUCTION

Numismatics is the study of currency, including coins, paper money, and tokens. This discipline yields fascinating cultural and historical insights, and is a field of great interest to scholars, amateur collectors, and professional dealers alike. Important applications of machine learning in this field include theft detection, identification and classification of finds, and forgery prevention. The present work focuses on ancient coins.

Individual ancient coins of the same type can vary widely in appearance due to centering, wear, pagination, and variance in artistic depiction of the same semantic elements. This poses a range of technical challenges and makes it difficult to reliably identify which concepts are depicted on a given coin using machine learning and computer vision based automatic techniques [4]. For this reason, ancient coins are typically identified and classified by professional dealers or scholars, which is a time consuming process demanding years of experience due to the specialist knowledge required. Human experts attribute a coin based on its denomination, the ruler it was minted under and the time and place it was minted. A variety of different characteristics are commonly used for attribution, including:—Physical characteristics such as weight, diameter, die alignment and colour.

- Obverse legend, which includes the name of the issuer, titles or other designations.
- Obverse motif, including the type of depiction (head or bust, conjunctional or facing, etc.), head adornments (bare, laureate, diademed, radiate, helmet, etc.), clothing (draperies, breastplates, robes, armour, etc.), and miscellaneous accessories (spears, shields, globes, etc.),0
- Reverse legend, often related to the reverse motif,
- Reverse motif (primary interest herein), e.g. person (soldier, deity, etc.), place (harbour, fortress, etc.) or object (altar, wreath, animal, etc.), and
- Type and location of any mint markings.

### 1.1. RELEVANT PRIOR WORK

Most existing algorithmic approaches to coin identification are local feature based which results in poor performance due to loss of spatial relationships between elements [3]. To overcome this limitation, approaches which divide a coin into segments have been proposed but these assume that coins are perfectly centered, accurately registered, and nearly circular in shape [2]. Recent work shows that existing approaches perform poorly on real world data because of the fragility of such assumptions . Although the legend can be a valuable source of information, relying on it for identification is problematic because it is often significantly affected by wear, illumination, and minting flaws [5]. Lastly, numerous additional challenges emerge in and from the process of automatic data preparation, e.g. segmentation, normalization of scale, orientation, and colour [6].

In the context of the present work it is particularly important to note that all existing work in this area is inherently limited by the reliance on visual matching and the assumption that the unknown query coin is one of a limited number of gallery types. However, this assumption is unrealistic as there are hundreds of thousands of different coin types. Hence herein our idea is to explore the possibility of understanding the artistic content depicted on a coin, which could then be used for subsequent text based matching, allowing access to a far greater number of coin types without requiring them to be represented in the training data. In short, our ultimate aim is to attempt to describe a coin in the same way that a human expert would – by identifying the individual semantic elements depicted thereon.

## II. PROBLEM SPECIFICATION, CONSTRAINTS, AND CONTEXT

Deep learning has proven successful in a range of image understanding tasks. Recent pioneering work on its use on images of coins has demonstrated extremely promising results, outperforming more traditional approaches by an order of magnitude. Major difference in the nature of the challenge we tackle here is that our data set is weakly supervised – that is, samples are labelled only at the image level by a corresponding largely unstructured textual description, even though the actual semantic elements of interest themselves occupy a relative small area of the image. This kind of data is much more abundant and easier to obtain than fully supervised data (pixel level labels), but poses a far greater challenge.

### 2.1 CHALLENGE OF WEAK SUPERVISION

Recall that our overarching goal is to identify the presence of various important semantic elements on an unknown query coin's reverse. Moreover, we are after a scalable approach – one which could be used with a large and automatically extracted list of elements. As this is the first attempt at solving the suggested task, we chose to use images which are relatively uncluttered, thus disentangling the separate problems of localizing the coin in an image and segmenting it out. Representative examples of images used can be seen in Fig 1.



Fig. 1. Typical images used in the present work (obverses on the left, reverses on the right)

To facilitate the automatic extraction of salient semantic elements and the learning of their visual appearance, each coin image is associated with an unstructured text description of the coin, as provided by the professional dealer selling the coin; see Fig 2. The information provided in this textual meta-data varies substantially – invariably it includes the descriptions of the coin's obverse and reverse, and the issuing authority (the person, usually emperor, on the obverse), but it may also feature catalogue references, minting date and place, provenance, etc.

```
83, Lot: 226. Estimate \ $100. Sold for \ $92.

MACRIANUS. 260-261 AD. Antoninianus (22mm, 3.40 gm). Samosata? mint. Radiate,
draped, and cuirassed bust right / Apollo standing facing, head left, holding
laurel-branch and leaning on lyre; star in left field. Cf. RIC V 6 (Antioch);
MIR 44, 1728k; cf. Cohen 2 (same). VF, light porosity. Rare.
```

Fig 2. Typical unstructured text attribution of a coin.

We have already emphasized the practical need for a scalable methodology. Consequently, due to the nature of the available data (which itself is representative of the type of data which can be readily acquired in practice), at best it is possible to say whether a specific semantic element (as inferred from text; see next section) is present in an image or not. Due to the amount of human labour required it is not possible to perform finer labelling i.e. to specify where the element is, or its shape (there are far too many possible elements and the amount of data which would require labelling is excessive), necessitating weak supervision.

### 2.2 DATA PRE-PROCESSING AND CLEAN-UP

As in most real-world applications unstructured, loosely standardized, and heterogeneous data, a major challenge in the broad domain of interest in the present work is posed by possibly erroneous labelling, idiosyncrasies, etc. Hence, robust pre-processing and clean-up of data is a crucial step for facilitating learning [12,14].

Image based pre-processing we begin by cropping all images to the square bounding box which contains just the reverse of the corresponding coin, and isotropically resizing the result to the uniform scale of 300 × 300 pixels. The obverse of a coin typically depicts the bust of an emperor, whilst the reverse contains the semantic elements of interest herein.

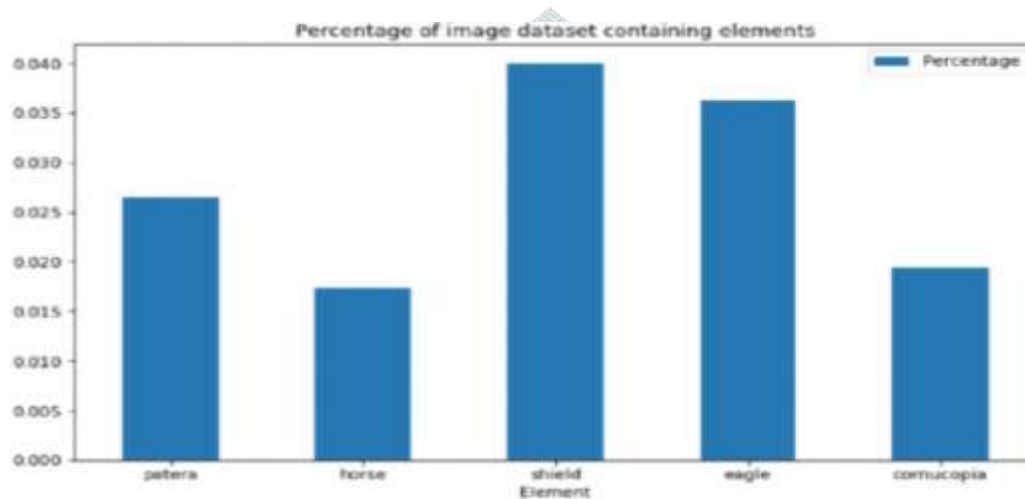
Text based extraction of semantics In order to select which semantic elements to focus on, we use the unstructured text files to analyse the frequency of word occurrences and, specifically, focus on the most common concepts, see Figs 3 and 4. Considering

that the present work is the first attempt at addressing the problem at hand, the aforementioned decision was made in order to ensure that sufficient training data is available.



Fig 3. Examples of depictions of reverse motif elements the present work focuses on

Due to the great diversity of different semantic elements which appear on ancient coins only a relatively small number of coins in the corpus as a whole contain the specific semantic elements which we attempt to learn herein, raising class imbalance issues which we address using stratified sampling.



Clean-up and normalization of text data to label the data we first clean the attribution text files to remove duplicate words and punctuation. Because the data included attributions in French, Spanish, German and English, we employ a translation API to generate translations of each keyword. Thus, when building the data set used to learn the concept of 'horse', images with associated text which contains the words 'horse', 'cabala', 'cheval' and 'pferd' are also included as positive examples, making better use of available data. Plurals, synonyms and other words strongly associated with the concept in question (e.g. 'horseman'), and their translations are also included.

Randomization and stratification We shuffle the samples before building training, validation and test sets for each of the selected elements ('horse', 'cornucopia', 'patera', 'eagle' and 'shield'), with a ratio of 70% training set, 15% validation set and 15% test set. To address under-representation of positive examples, we use stratified sampling to ensure equal class representation.

### 2.3 ERRORS IN DATA

In the context of the present problem as in most large scale, real-world applications of automatic data analysis, even after preprocessing, the resulting data will contain a certain portion of erroneous data. Errors range in nature from incorrect label assignments to incorrectly prepared images.

To give a specific example, we observed a number of instances in which the keyword used to label whether or not a given coin image contains the element in question was actually referring to the obverse of the coin, rather than the reverse, with which we are concerned. This is most prevalent with shields, and leads to incorrect labelling when building the corresponding data sets. Our premise (later confirmed by empirical evidence) is that such incorrect labelling is not systematic in nature (c.f. RANSAC) and thus that the cumulative effect of a relatively small number of incorrect labels will be overwhelmed by the coherence (visual and otherwise) of correctly labelled data. We also found that some of the original (unprocessed) images were unusually laid out in terms of the positioning of the corresponding coin's obverse and reverse, which thus end up distorted during preprocessing. However, these appear to be almost exclusively of the auction listing type described above, and as such would be labelled as negative samples and therefore cumulatively relatively unimportant.




Fig 4. Image and Pixel Labeling

The amount of labour which would be required to provide pixel level labelling (right) of different concepts is prohibitive, imposing the much more challenging task of having to perform learning from weakly annotated data (left).

### III. EXPERIMENTS

Experiments discussed in the present section were performed using Asus H110M-Plus motherboard powered by Intel Core i5 6500 3.2GHz Quad Core processor and NVidia GTX 1060 6GB GPU, 8GB RAM. Data comprised images and the associated textual descriptions of 100,000 auction lots kindly provided to us by the Ancient Coin Search auction aggregator

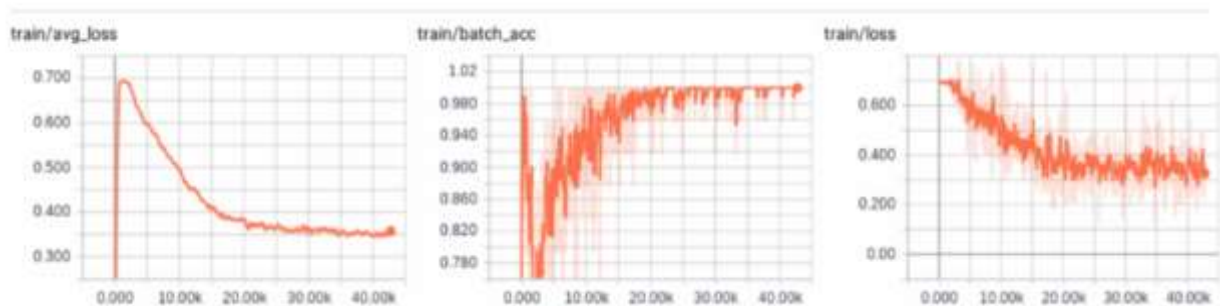


Fig. 5. Training a model to identify depictions of cornucopia

Table 1. Summary of experimental settings and results.

	Cornucopia	Patera	Shield	Eagle	Horse
Number of epochs	105	136	118	86	148
Training time (min)	58	30	82	51	106
Training accuracy	0.71	0.83	0.75	0.88	0.88
Validation accuracy	0.85	0.86	0.73	0.73	0.82
Validation precision	0.86	0.85	0.72	0.71	0.82
Validation recall	0.83	0.86	0.75	0.81	0.84
Validation F1	0.85	0.86	0.74	0.75	0.83
<b>Test accuracy</b>	0.84	0.84	0.72	0.73	0.82
<b>Test precision</b>	0.85	0.82	0.71	0.70	0.81
<b>Test recall</b>	0.83	0.87	0.74	0.81	0.82
<b>Test F1</b>	0.84	0.84	0.72	0.75	0.82

### IV. RESULTS AND DISCUSSION

As readily seen from Table 1, our approach achieves a high level of accuracy across training, validation, and test data. The validation and test scores are not significantly lower than the training scores, which suggests little evidence of overfit and indicates a well-designed architecture and sufficient data set. The models used to identify cornucopia and paterae are particularly successful, which is likely due to the low level of variance in artistic depictions of these elements – the orientation, position and style of shields, horses and eagles varies quite a bit, but cornucopia and paterae are fairly constant in their depiction.

#### 4.1 LEARNT SALIENT REGIONS

We use the occlusion technique to quantify the importance of different image regions in the context of the specific task at hand. In brief, the process involves synthetic occlusion by a uniform kernel and the quantification of the corresponding classification performance differential between unclouded and occluded inputs, with a higher difference suggesting higher importance. Previous work on computer vision based ancient coin analysis has demonstrated the usefulness of this technique in the interpretation of empirical results. In order to ensure robustness to relative size (semantic element to coin diameter), in the present work we adopt the use of three kernel sizes,  $32 \times 32$ ,  $48 \times 48$ , and  $64 \times 64$  pixels.



Fig. 6. Examples of automatically identified salient regions corresponding to a cornucopia, a patera, and a shield, respectively.

Typical examples of identified salient regions for different semantic elements are shown in Fig 8. It can be readily seen that our algorithm picks up the most characteristic regions of the elements well – the winding pattern of cornucopia, the elliptical shape of paterae, and the feather patterns of eagles. Since, as we noted before and as is apparent from Table 1, the performance of our algorithm is somewhat worse on the task of shield detection, we used our occlusion technique to examine the results visually in more detail. Having done so, our conclusion is reassuring – our hypothesis is that shields are inherently more challenging to detect as they exhibit significant variability in appearance and the style of depiction (shown from the front they appear circular, whereas shown from the side they assume an arc-like shape) and the least amount of characteristic detail. In this paper we made a series of important contributions to the field of computer vision based analysis of ancient coins. Firstly, we put forward the first argument against the use of visual matching of ancient coin images, having explained its practical lack of value. Instead, we argued that efforts should be directed towards the semantic understanding of coin images and described the first attempt at this challenging task.

Specifically, we described a novel approach which combines unstructured text analysis and visual learning using a convolutional neural network, to create weak associations between semantic elements found on ancient coins and the corresponding images, and hence learn the appearance of the aforementioned elements. We demonstrated the effectiveness of the proposed approach using images of coins extracted from 100,000 auction lots, making the experiment the largest in the existing literature. In addition to a comprehensive statistical analysis, we presented the visualization of learnt concepts on specific instances of coins, showing that our algorithm is indeed creating the correct associations. We trust that our contributions will serve to direct future work and open avenues for new promising research.

#### V. ACKNOWLEDGMENT

The authors gratefully acknowledge the support of NVIDIA Corporation for their donation of the GPU used for this research.

#### VI. REFERENCES

- [1]. Agarap, A.F.: Deep learning using rectified linear units (ReLU). arXiv p. 1803.08375 (2018)
- [2]. Anwar, H., Zambian, S., Kampel, M.: Supporting ancient coin classification by image based reverse side symbol recognition. In Proc. International Conference on Computer Analysis of Images and Patterns pp. 17–25 (2013)
- [3]. Anwar, H., Zambanini, S., Kampel, M.: Coarse-grained ancient coin classification using image-based reverse side motif recognition. Machine Vision and Applications 26(2), 295–304 (2015)
- [4]. Arandjelović, O.: Automatic attribution of ancient Roman imperial coins. In Proc. IEEE Conference on Computer Vision and Pattern Recognition pp. 1728–1734 (2010)
- [5]. Arandjelović, O.: Reading ancient coins: automatically identifying denarii using obverse legend seeded retrieval. In Proc. European Conference on Computer Vision 4, 317–330 (2012)
- [6]. Conn, B., Arandjelović, O.: Towards computer vision based ancient coin recognition in the wild – automatic reliable image preprocessing and normalization. In Proc. IEEE International Joint Conference on Neural Networks pp. 1457–1464 (2017)
- [6]. Fare, C., Arandjelović, O.: Ancient Roman coin retrieval: a new dataset and a systematic examination of the effects of coin grade. In Proc. European Conference on Information Retrieval pp. 410–423 (2017)
- [7]. Janocha, K., Czarnecki, W.M.: On loss functions for deep neural networks in classification. arXiv p.1702.05659 (2017)
- [8]. Kampel, M., Zaharieva, M.: Recognizing ancient coins based on local features. In Proc. International Symposium on Visual Computing 1, 11–22 (2008)
- [9]. Kinga, D., Adam, J.B.: A method for stochastic optimization. In Proc. International Conference on Learning Representations 5(2015)

- [10]. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* pp. 1097–1105(2012)
- [11]. Mattingly, H.: *The Roman Imperial coinage.*, vol. 7. Spink (1966)
- [12]. Rieutort-Louis, W., Arandjelovi´c, O.: Bo(V)W models for object recognition from video. In *Proc. International Conference on Systems, Signals and Image Processing* pp. 89–92 (2015)
- [13]. Rieutort-Louis, W., Arandjelovi´c, O.: Description transition tables for object retrieval using unconstrained cluttered video acquired using a consumer level handheld mobile device. In *Proc. IEEE International Joint Conference on Neural Networks* pp. 3030–3037 (2016)
- [14]. Schlag, I., Arandjelovi´c, O.: Ancient Roman coin recognition in the wild using deep learning based recognition of artistically depicted face profiles. In *Proc. IEEE International Conference on Computer Vision* pp. 2898–2906 (2017)
- [15]. Yue, X., Dimitriou, N., Arandjelovi´c, O.: Colorectal cancer outcome prediction from H&E whole slide images using machine learning and automatically inferred phenotype profiles. In *Proc. International Conference on Bioinformatics and Computational Biology* (2019)
- [16]. Zaharieva, M., Kampel, M., Zambanini, S.: Image based recognition of ancient coins. In *Proc. International Conference on Computer Analysis of Images and Patterns* pp. 547–554 (2007)
- [17] Usha Kosarkar, Gopal Sakarkar, Shilpa Gedam (2022), “Revealing and Classification of Deepfakes Videos Images using a Customize Convolution Neural Network Model”, *International Conference on Machine Learning and Data Engineering (ICMLDE)*, 7<sup>th</sup> & 8<sup>th</sup> September 2022, 2636-2652, Volume 218, PP. 2636-2652, <https://doi.org/10.1016/j.procs.2023.01.237>
- [18] Usha Kosarkar, Gopal Sakarkar (2023), “Unmasking Deep Fakes: Advancements, Challenges, and Ethical Considerations”, *4<sup>th</sup> International Conference on Electrical and Electronics Engineering (ICEEE)*, 19<sup>th</sup> & 20<sup>th</sup> August 2023, 978-981-99-8661-3, Volume 1115, PP. 249-262, [https://doi.org/10.1007/978-981-99-8661-3\\_19](https://doi.org/10.1007/978-981-99-8661-3_19)
- [19] Usha Kosarkar, Gopal Sakarkar, Shilpa Gedam (2021), “Deepfakes, a threat to society”, *International Journal of Scientific Research in Science and Technology (IJSRST)*, 13<sup>th</sup> October 2021, 2395-602X, Volume 9, Issue 6, PP. 1132-1140, <https://ijsrst.com/IJSRST219682>
- [20] Usha Kosarkar, Gopal Sakarkar (2024), “Design an efficient VARMA LSTM GRU model for identification of deep-fake images via dynamic window-based spatio-temporal analysis”, *International Journal of Multimedia Tools and Applications*, 8<sup>th</sup> May 2024, <https://doi.org/10.1007/s11042-024-19220-w>