A deep learning approach to classify country and value of modern coins

(Article begins on next page)

SOURCE:Neural Computing and Applications

DOI: https://link.springer.com/article/10.1007/s00521-023-09355-6

A deep learning approach to classify country and value of modern coins

Stefano Cirillo[1†], Giandomenico Solimando[1†] and Luca Virgili[2*†]

[1]Department of Computer Science, University of Salerno, Italy.
[2]Department of Information Engineering, Polytechnic University of Marche, Italy.

*Corresponding author(s). E-mail(s): luca.virgili@univpm.it;
Contributing authors: scirillo@unisa.it; gsolimando@unisa.it;
[†]These authors contributed equally to this work.

Abstract

The use of Artificial Intelligence (AI) to preserve and promote cultural heritage has experienced significant growth in recent years. Among the various areas of cultural heritage, numismatics has emerged as a particularly promising field where we can develop AI solutions. Numismatics refers to the study of coins, tokens, paper money, and medals, which play a critical role in understanding human history and culture. However, there are still limited resources available to help researchers and collectors in the identification of coins. This is due to the vast number of coins in circulation, which presents a significant challenge in developing smart tools for classification tasks. This paper aims to provide a contribution to this setting. In particular, we start by creating a new dataset called EURO-Coin, which consists of images showing the side of coins with reliefs and is designed to facilitate the training and testing of AI models for euro coin classification. Then, we propose two approaches that leverage Convolutional Neural Networks and self-attention layers to classify the country and value of the coins. In our experiments, we obtain an accuracy of 86.9% for country classification and an accuracy of 96.4% for value classification. Finally, we conduct an ablation study to evaluate the impact of the preprocessing activities and attention layers in our approaches.

Keywords: Cultural Heritage, Numismatics, Convolutional Neural Network, Deep Learning

1 Introduction

The use of Artificial Intelligence (AI) technologies to preserve, rediscover, and promote historical cultural heritage is a multidisciplinary topic of great interest that involves both public institutions and private companies [1–3]. In recent years, this phenomenon has led to the development of several tools for conserving and transmitting the cultural heritage of different countries. As an example, one of the tools developed in the last few years is *AI for Cultural Heritage*[1] by Microsoft, which exploits artificial intelligence to provide effective tools for people and organizations aiming at the preservation and enrichment of cultural heritage. For instance, the Arts and Humanities Research Council in the United Kingdom (UK) has funded several projects for allowing online access to the cultural collections of the UK using innovative AI-based technologies [4].

The protection of artistic and historical heritage is an ethical duty, but it also has strong economic implications. This was highlighted by the European Commission, by publishing a recommendation on a common European data space for cultural heritage, with the aim of accelerating the digitization of cultural heritage [5].

Among the different areas of cultural heritage, numismatics is one of the areas that is trying to emerge in this digital transition scenario [6, 7]. Numismatics deals with the study and collection of coins, tokens, paper money, medals, and similar objects, playing an important role in cultural heritage, since coins and paper money have been and continue to be an integral part of human history. Coins provide relevant information about a society's economy, politics, religion, and art. Through the study of numismatics, historians, archaeologists, and other researchers can gain insights into the past and enhance our understanding of cultural heritage [8, 9]. Coins can also serve as a powerful tool for promoting cultural heritage by inspiring interest and curiosity in history, art, and archaeology. However, as in the case of most areas of collecting, recognizing a coin, appreciating its history, and the journey it has traveled to be in people's hands requires considerable knowledge of domain experts. Nevertheless, there are still few tools able to support people, but also researchers and collectors (i.e., Numismatists) in the automatic identification of coins and in the management of new private collections. In fact, the huge amount of coins in circulation, due to the different types and currencies, makes the definition of intelligent tools a complex task. In fact, existing tools only identify the type of currency, such as euro, dollar, pound, or rupee, providing information that is easily deductible by seeing the coin.

In this study, we address the lack of research on modern coins and AI. To this end, we propose a new dataset of euro coin images of their relief sides that differs from existing datasets and allows the research community to work with modern coins from country and value perspectives. Indeed, while previous works have focused on classifying coins by their value [10, 11], to the best of our knowledge our approach is one of the first attempts to classify coins by

[1] https://www.microsoft.com/en-us/ai/ai-for-cultural-heritage

the country that minted them. This introduces challenges such as recognizing similar drawings used for different coin versions and variations in relief designs over time.

Following related works for the classification of ancient coins, our approaches include the use of Convolutional Neural Networks [10, 12, 13] and attention layers [14] for addressing the country and value classification tasks of modern coins. In this setting, we also introduce the analysis of motifs of the relief sides of the coins [13–16] thanks to the application of image filters such as Sobel or Prewitt. To the best of our knowledge, the analysis of motifs present in the relief side has not been explored so far in modern coin classification through machine learning and, as will be clear in the next sections, it retrieves promising results for both classification tasks. Finally, our approaches consider the relief side of coins and do not use the front side (which is typically used in ancient coins classification [14, 17]), as it does not allow the classification of countries since it is the same for all of them, while the value classification could be too straightforward due to the presence of clear inscriptions. On the other hand, the relief side allows for both country and value classification since each country could have unique drawings for each coin value.

To summarize, the contributions of this paper are:

- A new dataset, namely *EURO-Coin dataset*, containing the images of the euro coins in circulation. It represents the first dataset containing images of euro coins on the side with reliefs and has the potential to increase the development of AI approaches to preserve, organize and transmit a part of the cultural heritage.
- A new approach for identifying the country of coins through their side with reliefs. This is a complex task due to both the high number of countries that adopt the euro as their official currency and the large number of types of coins in each country having different reliefs minted over time.
- A new approach for identifying the value of coins through the side of the coin with reliefs. This is a complex task due to both the large number of types of coins in circulation with different reliefs and to the fact that many countries have the same reliefs on coins with different values;
- An ablation study for evaluating the classification performances of our approaches with different preprocessing activities and without the attention layers.

The rest of the paper is organized as follows: Section 2 surveys relevant related works and datasets in the literature; Section 3 analyzes and shows the problem of the euro coin classification in numismatics; Section 4 introduces the proposed dataset and the procedures for collecting data; Section 5 shows the experimental evaluation and provides a discussion of results; Conclusions and future directions are discussed in Section 6.

2 Related Literature

In this Section, we provide a review of related works and datasets relevant to the application of AI to numismatics. Specifically, in Section 2.1, we describe articles regarding the application of AI algorithms to the numismatic scenario, while in Section 2.2, we report the numismatic datasets for ancient and modern coins available in the literature.

2.1 Related Works

Artificial intelligence has been making significant advancements in recent years and has been applied in various fields, including healthcare, finance, manufacturing, and many others [18–20]. However, AI has also started to gain importance in cultural heritage scenarios. For instance, AI is used in enriching museum experiences [1, 21, 22], where it provides personalized information and recommendations to visitors and creates more immersive experiences. Another application regards the detection of fake artwork [2, 3]. Machine Learning algorithms can analyze patterns and features of a painting or sculpture to detect any anomalies that may suggest a forgery. AI is also being used to preserve cultural heritage sites [23] and to solve many tasks in the field of numismatics [6, 7, 24, 25]. For instance, AI algorithms can be trained to identify and analyze different features of coins to determine their physical condition and assign them a score [26]. Similarly, predicting the value of a coin [27, 28] and possible counterfeits [29] are challenging tasks that can benefit from AI. A further application concerns the classification of ancient and modern coins, which can help collectors and experts quickly identify and catalog coins [12, 30–33].

According to the literature on the classification of ancient and modern coins in the numismatic field, two main groups of approaches can be identified: *(i)* those using statistical and computer graphics techniques, and *(ii)* those using machine learning and deep learning models.

As for the former group of approaches, in [34], the authors propose an improvement of the Histogram of Oriented Gradients (a commonly used method for object recognition) called Dynamic-HOG. It locates and analyzes objects in images to define a dynamic window size based on each object's size. Dynamic-HOG is used to create a method for coin recognition by focusing on the characters present in the image. In [15] and in [35], the authors propose an approach to classify ancient coins based on the motifs present on their relief sides using a bag-of-visual-words model. The authors also incorporate specific geometric structures of the motifs present in the coins to improve the classification rate. In [36], the authors use a multistep approach, where the first computes a translationally and rotationally invariant descriptor, and the second computes an illumination-invariant eigenspace that retrieves the probabilities for coin classes for both the obverse and relief sides. In the final step, the probabilities are combined and a coin is accepted to be of a certain class if these probabilities are higher than a threshold. In [37], the authors propose a coin classification method using local texture features for ancient Roman coins. Also in [38], the authors introduce an automated system for coin recognition using spatial-appearance relationships. The authors of [39] propose a classification framework that performs an image-matching task between reference and input images, and then identifies words present in a coin image. These two steps are performed on both coin sides and the resulting class probabilities are then combined. [16] presents a method for estimating the visual similarity between two coin images that is rotation-insensitive. [40] presents a method for recognizing words in ancient coins images. The authors first select character candidate locations in the image, then compute the probabilities for each of the 18 considered characters, and finally check if these are present in their lexicon. [41] introduces an automatic system based on a set of features representing characteristics of local image portions. Finally, in [42] the authors propose a method that computes gradient magnitudes in a coin image, and extracts rotation-and-flipping-robust patterns present in local portions. Their solution is then tested on the MUSCLE CIS dataset [43].

As for the latter group of approaches, we can further divide the corresponding articles by the type of coins under consideration, i.e. ancient or modern coins. In the ancient coin scenario, [44] proposes a framework for the classification of ancient Roman coins using a hierarchical knowledge structure and a convolutional neural network classifier. The framework comprises feature-based identification extracted from image matching and the detection of points of interest in local image portions. In [17], the authors introduce a semi-supervised learning model called Graph Transduction Games that uses both sides of the coins for the classification. In [14], the authors introduce CoinNet, a model for the classification of ancient Roman coins that recognizes their motifs like faces, scenes, animals, and buildings. CoinNet is made up of bilinear pooling, residual convolutional blocks, and attention layers. Moreover, in [45], the authors leverage convolutional neural network models and transfer learning to classify the images of Greek ancient coins for identifying their issuing authority. Furthermore, [13] classifies Roman Imperial coins through a profile face recognition of the Roman emperors, while [9] trains two CNN models, one that works with a specific label of the Roman Imperial Coinage on the relief side, and the other that looks at the Roman emperor. In [8], the authors describe a method using multimodal input to extract and associate semantic concepts with the correct coin images and then train an AlexNet to learn the appearance of these concepts in the images. As for the modern coin scenario, in [10], the authors use a convolutional neural network to classify euro coins automatically and create a smartphone app so that users can identify coins by taking a photo with a smartphone. In [12], the authors employ convolutional neural networks for classifying Indian coins and then investigate the impact of design parameters, including the dataset size, number of classes, image quality, and so on. [46] proposes a dense neural network and pattern averaging method to distinguish the Turkish 1 Lira coin from the 2 euro coin. [47] presents a coin recognition system based on the alignment of two coins by searching for a rotation that maximizes the number of collinear gradient vectors, which is then used as input to a nearest neighbor classifier. In [11], the authors propose a coin recognition system that involves a training phase where the system identifies coin positions in the image, extracts texture features, generates dictionaries based on a bag of words approach, and trains a Support Vector Machine model for each coin class. In the testing phase, the coin image is classified based on the models obtained previously.

Considering the numismatic and AI scenario described above, our paper has some similarities and unique contributions. First of all, there is a lack of research on modern coins, which we address by proposing a new dataset comprising images of euro coins minted over the years. Our dataset has distinct characteristics compared to the limited datasets currently available to the research community as will be clear in the next section. Furthermore, while some previous works [10, 11] classify coins based on their value, our paper represents, to the best of our knowledge, the first effort to classify coins according to the country that minted them. This perspective introduces novel challenges, such as countries using similar drawings for different coin versions and variations in relief designs over time. In our approaches, we use Convolutional Neural Network models as other works [10, 12, 13] due to its effectiveness in working with images, but we also perform an analysis of motifs typical of ancient coins works [13–16] thanks to the application of image filters such as Sobel or Prewitt. We also decided to include attention layers into our model similar to [14], which were not used before for modern coin classification. Furthermore, to the best of our knowledge, our approach is one of the first attempts to classify the countries and the values of the coins from their relief sides. Lastly, we exclusively work with the relief side of coins and exclude the use of both sides as in [14, 17]. In the modern coin scenario, the front side of coins does not allow country classification since it is the same for all countries, and the value classification could be straightforward due to the clear inscriptions. On the other hand, the relief side enables both classification tasks, as different countries employ different drawings, and the value inference is not obvious apriori. Overall, our paper contributes to the field by introducing a new dataset and proposing country and value classifications of modern coins using relief side motifs obtained through filters.

2.2 Related Datasets

As we introduce a new dataset called EURO-Coin, we report the other numismatic datasets that can be found in the literature. Since the availability of datasets is fundamental for the research community, we considered only the ones that are publicly available. The results of this research are reported in Table 1.

The datasets present in Table 1 can be divided into two groups: the one regarding ancient coins (such as Roman coins dating back to the republican or imperial period) and the one concerning modern coins (such as euro or rupee).

As for the ancient coins, [16] introduced RRC-60 (1), a dataset of Roman Republican coins from the Museum of Fine Arts in Vienna, which contains

Dataset	Brief description	Articles	Objective	# of images
RRC-60 (1)	Roman Republican coins	[13, 16, 17, 37, 39, 44]	Matching, Classification, Segmentation	3900
RRC-60 (2)	Roman Republican coins	[17]	Classification	6000
RRCD	Roman Republican coins	[14, 44]	Classification	18285
AnCoins-12	Greek coins	[45]	Classification	565
MUSCLE CIS [43]	Euro coins	[11, 42, 47]	Classification	60000

Table 1: Datasets available for ancient and modern coins

3900 images representing 515 different coin types. However, only for 237 of them, there are more than three images available. The authors of [16] suggest grouping them into 60 classes according to a taxonomy of Roman Republican coins [48] in order to create a coin classifier algorithm. Afterward, the authors of [17] extend RRC-60 (1) thanks to the integration of new images from popular websites[2], and obtain a new dataset of 6000 images and 60 classes [48]. Then, they develop a classification algorithm for both RRC-60 (1) and RRC-60 (2). In [14], the authors propose a dataset comprising 18285 images corresponding to 228 different motif classes. Images are collected from the Vienna Museum of Fine Arts, the British Museum, and popular websites. Then, the authors use a deep learning model to classify the ancient Roman Republican coins through their motifs. Besides Roman coins, [45] introduces a dataset called AnCoins-12 that contains 565 images of 12 different classes of Greek ancient coins from the area of ancient Thrace. The authors of [45] use a convolutional neural network model to identify the issuing authority of those coins.

As for the modern coins, very few datasets are publicly available and used by the articles in the literature. To the best of our knowledge, only [43] proposed a dataset of euro coins, which was used for the MUSCLE Coin Images Seibersdorf Benchmark Competition in 2006. This dataset contains 60000 images of euro coins corresponding to 12 European countries. It contains both front and relief sides and hundreds of coin classes depending on the coin types and motifs. With respect to MUSCLE CIS, our dataset contains more countries (12 against 23), and the images of coins minted over the years. MUSCLE CIS is used in [11] for a coin recognition system based on the training of a Support Vector Machine model for each class. [47] leverages MUSCLE CIS for classifying coins based on the image alignment and on a nearest neighbor classifier. Finally, in [42] the authors compute gradient magnitudes and extract robust patterns from MUSCLE CIS coin images for their classification.

In conclusion, there are some datasets available for the classification and recognition of ancient coins, but the literature lacks datasets concerning modern coins. For this reason, we decided to introduce EURO-Coin, which contains the euro coins minted by European countries over the years. As will be explained in the next sections, EURO-Coin is in line with the other datasets in terms of the number of images and classes and provides labels of the country and value of a coin, which allows the researcher to perform classification and segmentation tasks along these two perspectives.

[2]such as acsearch.info and numismatics.org/crro/

3 Euro coins classification in numismatics

The classification of euro coins for numismatic experts is a very challenging task due to the large number of different coin series in circulation, issued in different European Union countries (i.e., Member States). Each coin series can have numerous variations, such as different legends, mint marks, symbols, and designs, making classification a challenge even for numismatic experts. Furthermore, the problem of classifying euro coins has been further complicated by the issue of commemorative coins coined in limited editions, which has significantly increased the number of coins in circulation. In fact, actually, there are eight different types of euro coins in circulation among all Member States: 1 cent, 2 cents, 5 cents, 10 cents, 20 cents, 50 cents, 1 euro, and 2 euros. Each Member State has its own version of coins, with one side representing European unity and the other side featuring a national design. This means that there are currently over 120 different variants of euro coins in circulation. Figure 1 shows an example of 1 euro coin currently in circulation in 12 EU countries issued in 2002.

Fig. 1: An extract of 1 euro coins issued in 2002.

Other than the number of different variants, it is necessary to consider that some countries have changed the reliefs on the coins over time, such as the Netherlands, Vatican City, and Belgium, or that some states mint multiple coins every year with different mint marks, such as German euro coins that have the letters “A”, “D”, “F”, “G”, or “J” for identifying the mint that minted the coin (see Figure 2).

Fig. 2: An extract of 1 cent coins with “A”, “D”, “F”, “G”, and “J”.

Despite the challenging problem, in this paper, we aim to address two different issues related to the classification of euro coins: identifying the value of a coin and the state member that has issued it considering the side with reliefs of the coin.

Identifying the countries that minted a coin based on its relief is an extremely complex task since many countries issue coins with images of historical characters or symbolic figures, which can be shared by multiple countries or even different eras. Furthermore, the artistic styles of the reliefs can be similar among countries, making it difficult to distinguish between them. In some cases, the presence of a mint mark or an abbreviation, such as in the case of Germany, can help identify the issuer of the coin. However, these indicators are not always present, and in some cases, they can be difficult to interpret without specialized knowledge.

On the other hand, identifying the value of a coin based on relief alone adds new complexity to the problem. Coins often have the same design or picture, regardless of value, making it difficult to distinguish between them. For example, several countries, including Belgium, Malta, and Luxembourg, have the image of a common face or object on most coins, regardless of their value. To the best of our knowledge, we are one of the first to face the problem of recognizing the value of a euro coin from the relief side.

Numismatic experts generally use specialized catalogs to identify the different variants of each coin versions, as well as key indicators that aid in their identification. However, manual identification is time-consuming and requires specialized knowledge, hence it is necessary to explore the use of technologies related to artificial intelligence and machine learning to design and develop automated systems for euro coin classification. These approaches would allow to analyze a large number of coin images and automatically identify different variants, helping numismatic experts classify coins in a more efficient and accurate way.

In the following sections, we will discuss new approaches designed to deal with these problems, which are based on our EURO-Coin dataset created thanks to a numismatic expert.

4 Proposed Solution

In this Section, we describe the numismatic scenario and the classification tasks we address. Specifically, in Section 4.1 we propose a generic model to describe the numismatic scenario. In Section 4.2, we outline the dataset for our scope and its statistics. In Section 4.3, we report the considered preprocessing activities for our dataset, and in Section 4.4 we describe our Convolutional Neural Network architecture to address the country and value classification tasks.

4.1 Model Definition

In this paper, we aim to classify images of coins in the numismatic context. To achieve this, we need a model to represent the entire scenario. Let us consider a dataset D of images, where each $d \in D$ represents a coin. Each d is an RGB image of specific width w_d and height h_d, and can also contain some background noises depending on the creator of that image. For this specific study, we are only interested in the images of the relief sides of the coins, and only in the modern coins coined in Europe countries.

Therefore, every relief of the coins, represented by an image in $d \in D$, is minted in a particular country s_d and has an associated current value v_d. We denote the set of countries of all the coins corresponding to the images in D as S and the set of all possible values of these coins as V. Starting from this setting, we are able to identify two possible classification tasks: *(i)* country classification, and *(ii)* value classification. The country classification task involves identifying the country s_d of each coin d based on the side with relief of the coins, while the value classification task involves predicting the current value v_d of d.

Considering the definition of D, the cardinality of S and V are greater than 2, therefore we are not dealing with a binary classification task but with a multiclass classification task. In the former case, we have $|S|$ labels while $|V|$ labels in the latter.

In order to standardize the images $d \in D$, we can define a set of transformation functions T. The set T could contain any function typical of image preprocessing, such as cropping, resizing, and filtering. The objective of T is twofold. The first concerns the standardization of the width and height of the images of D to w_s and h_s in order to obtain a uniform input for the machine learning model. The second regards the extraction of relevant features to improve the classification performance.

To predict the country and value of the features in dataset D, a machine learning model C capable of multiclass classification is required. While any classification model could be used, we have chosen to use a Convolutional Neural Network as our C due to its success in handling complex image classification tasks [9, 10, 14]. The training of our C consists of the minimization of the cross-entropy loss. Let $\hat{y}_d$ be the softmax output of C obtained from an image d whose true label is a one-hot encoding y_d of the S. The cross-entropy loss for a sample is defined as:

$$\mathcal{L}_{CE_S} = -\sum_{i=1}^{|S|} y_{d_i} \cdot \log(\hat{y}_{d_i}) \quad \mathcal{L}_{CE_V} = -\sum_{i=1}^{|V|} y_{d_i} \cdot \log(\hat{y}_{d_i}) \tag{1}$$

where $\mathcal{L}_{CE_S}$ corresponds to the cross-entropy loss for the country classification task, while $\mathcal{L}_{CE_V}$ corresponds to the cross-entropy loss for the value classification task.

In order to evaluate the performance of M, we leverage four common metrics in the machine learning field, which are Accuracy, Precision, Recall, and

F1-Score. These metrics are defined in terms of the number of true positives (i.e., TP), false positives (i.e., FP), true negatives (i.e., TN), and false negatives (i.e., FN) that the model generates when making predictions. The formulas to calculate them are the following:

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$
$$Precision = \frac{TP}{TP + FP}$$
$$Recall = \frac{TP}{TP + FN}$$
$$F1 - Score = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

The Accuracy measures the proportion of correct predictions made by the model over all the predictions. Precision evaluates the proportion of true positives out of all the predicted positives, while Recall measures the proportion of true positives out of all the actual positives. Finally, F1-Score is the harmonic mean of precision and recall that provides a balance between them. In order to properly estimate the classification performances of our models, these metrics are computed on a testing set that is not used for the training.

4.2 EURO-Coin dataset

One of the main challenges in defining an approach for coin classification concerns the creation of ad-hoc datasets to be used in predictive processes. In fact, most of the datasets proposed in the literature contain images of coins that are not euro and/or collect images of euro coins from the value side (e.g., 1c, 2c, 5c, 10c, 20c, 50c, 1€, 2€).

This makes all datasets not suitable for our study since our aim is to provide an approach for classifying euro coins from the relief side. To the best of our knowledge, there are no datasets containing images of euro coins of the Member States from the side with relief. To this end, we have created a new dataset, named *EURO-Coin dataset*[3], by using some of the best-known search engines for extracting images from the web.

The process of creating a dataset of coin images opens different challenges and issues. In particular, to create a dataset of coin images, it is important to consider the potential challenges that can arise:

- **Coin variability**: Coins can vary in size, shape, color, lighting, and other factors. All these factors can make it difficult to standardize the images for the dataset, which can make it more challenging to train artificial intelligence models to recognize the coins;
- **Image quality**: Coin images can be of low quality or blurry, which can make it difficult to extract important features of the coins;

[3]The dataset is publicly available on https://github.com/cciro94/EuroCoinDataset.

- **Amount of data**: To create a representative dataset of coins, it is necessary to acquire a large number of coin images. However, it is not always easy to collect a large number of coin images under different conditions. In addition, creating a large dataset requires a lot of time and resources;
- **Data labeling**: Labeling coin images takes time and skills, since it is necessary to accurately identify the coin represented in each image, also involving a domain expert in the labeling step.

Figure 3 provides an overview of the process for creating the EURO-Coin dataset. As we can see, the process has been divided into two steps, automatic extraction and manual classification, respectively. The automatic extraction step aims to collect images of coins from the web through some of the most known web search engines, such as Google, Bing, and Baidu. Although there exist several other web search engines, we choose to limit our analysis only to these, so as to reduce the number of duplicate images extracted from them. This first step required defining a module capable of automatically connecting to each web search engine and composing queries for searching images according to all countries and all coin values. Specifically, we have defined two different types of queries to extract coin images from search engines that start by considering a coin value $[v]$ and a country $[s]$. For each query, we defined a template of the sentence that has been automatically filled from the module for composing queries with both the coin value and the country to search for. In particular, the first query defined for our study was $q_1 =$ "*Images of the* $[v]$ *euro cent coin* $[s]$" while the second query was $q_2 =$ "$[v]$ *euro cent coin* $[s]$". It is important to notice that, similar to the limitation of the number of web search engines involved in our study, we have considered only two different types of queries in order to avoid extracting a large set of duplicate images. Both modules for creating queries and extracting information exploit a parallelization paradigm capable of performing multiple searches simultaneously. This strategy allowed the image extraction to be performed efficiently, despite the large number of queries made. In fact, if we consider the number of European countries using the Euro, the different coin values, the different search engines considered, and the queries defined, we can see that the number of searches performed is extremely large. Therefore, the parallelization paradigm allowed us to significantly speed up the extraction of images from the web. Also, to extract good quality images for our dataset, we configured search engines to only extract images with dimensions of at least 100px in both width and height. This has been made possible thanks to the configuration settings of each search engine, which can be easily customized through the parameters of their APIs.

After the automatic extraction step, we collected several sets of images for each country and coin value. Although the image search and storing strategy provided us with a preliminary division of the images for each country and value, it was necessary to proceed with a step of manual selection and removal of any outliers. To this end, we have requested a domain expert, i.e., a numismatic, to manually filter and label the images in the dataset. In our study, the

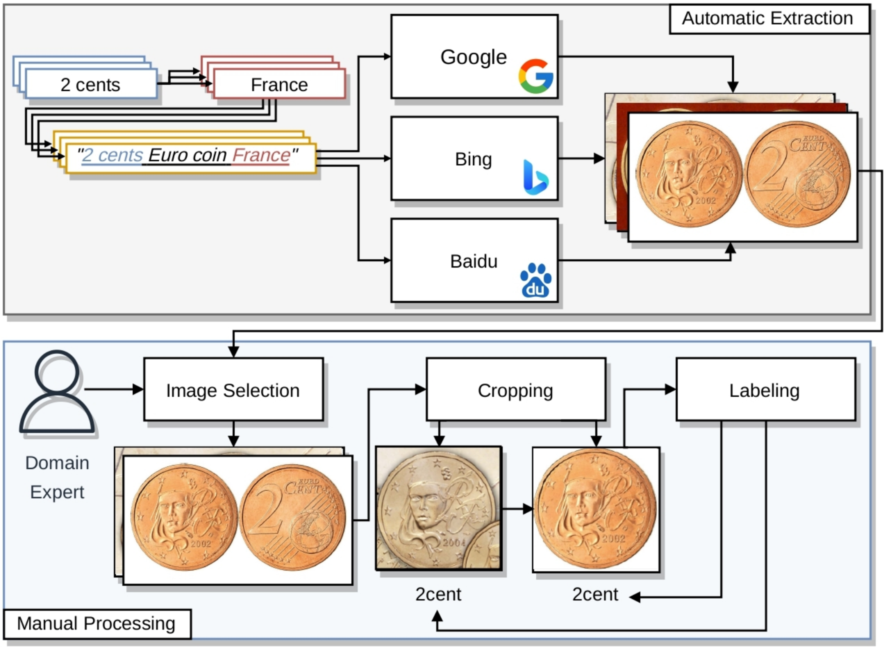

Fig. 3: Dataset creation and classification processes.

domain expert is one of the authors that has twenty years of experience in the collection and classification of euro coins.

The manual classification step of the domain expert has been divided into three different activities: *i*) the selection of images related to the coins of the member states and the consequent removal of the incorrect ones (e.g., extracted incorrectly from search engines or duplicate images); *ii*) manual cropping of images when they contained more than one coin, and *iii*) coin labeling and dataset definition. During the selection and labeling steps, the domain expert removed photos of coins in which the reliefs on the coins were unclear, such as those taken at particular angles or where the reliefs of the coin were not fully visible. After the manual classification step, we obtained the images of our EURO-Coin dataset, whose statistics are reported in Table 2.

Regarding the quality of these images, the expert manually cropped many of the images extracted from the web, since they showed different coins or had a very large background that could harm the classification task. The statistics on the quality of the images after the manual classification step are shown in Figure 4. As we can see, the dimensions of the images in pixels, i.e. height and width (blue and red bars, respectively), are almost always the same for all percentiles, which means that all the images tend to have an aspect ratio close to 1:1. Moreover, most of the images have dimensions larger than 300 pixels, which is already visible from the 25th percentile of each coin value. Only less than 1% of the images for each coin are smaller than 100 pixels. This

	1c	2c	5c	10c	20c	50c	1€	2€	Total per country
Andorra	42	46	71	43	49	57	60	220	588
Austria	83	64	76	72	49	57	73	94	568
Belgium	86	29	99	98	112	61	78	144	707
Cyprus	13	23	38	40	25	28	84	107	358
Estonia	58	44	45	41	44	30	61	102	425
Finland	49	38	61	36	72	58	99	201	614
France	34	57	67	52	59	73	57	128	527
Germany	163	131	60	163	80	62	90	178	927
Greece	121	52	58	63	69	59	117	169	708
Ireland	56	65	50	83	76	45	66	124	565
Italy	121	36	67	59	72	79	70	119	623
Latvia	71	30	45	17	21	45	91	232	552
Lithuania	94	15	25	28	35	31	90	181	499
Luxembourg	43	29	35	31	52	35	90	173	488
Malta	74	52	64	76	83	79	76	154	658
Monaco	14	3	33	48	57	47	110	107	419
Netherlands	35	32	64	76	72	69	62	166	576
Portugal	55	57	69	80	40	37	101	132	571
Sant Marino	47	56	57	60	62	77	109	114	582
Slovakia	34	26	21	20	19	32	39	156	347
Slovenia	56	45	42	50	52	45	65	136	491
Spain	23	41	46	48	52	55	97	174	536
Vatican City	52	30	34	58	50	91	62	172	549
Total per Value	1424	1001	1227	1342	1302	1252	1847	3483	**12878**

Table 2: Statistics about coin images in the EURO-Coin dataset.

is mainly due to the fact that the manual cropping step of many of the figures has led to the reduction of the dimensions of some images in order to remove noisy elements or to isolate coins when others were contained. This especially happened in countries where the spread of coins is rarer, such as Andorra, San Marino, Monaco, and Vatican City, since it is often difficult to find images of single coins.

Then, we analyzed the aspect ratio of the images in order to set up the preprocessing task and thus to uniform the input to our deep learning model. To this end, we computed the aspect ratio of the images in our dataset and reported the corresponding results in Table 3.

From Table 3, we noted that the 90.25% of the images had a standard 1:1 aspect ratio, while the remaining 9.75% had a different ratio. Therefore, we modified the last set of images to obtain a dataset having the same aspect ratio of 1:1. Finally, we point out that the whole process of creating the EURO-Coin dataset took about 15 days to complete.

4.3 Image Preprocessing

The process of creating the dataset has already foreseen some preliminary operations of preprocessing on the images, such as cropping. Although these

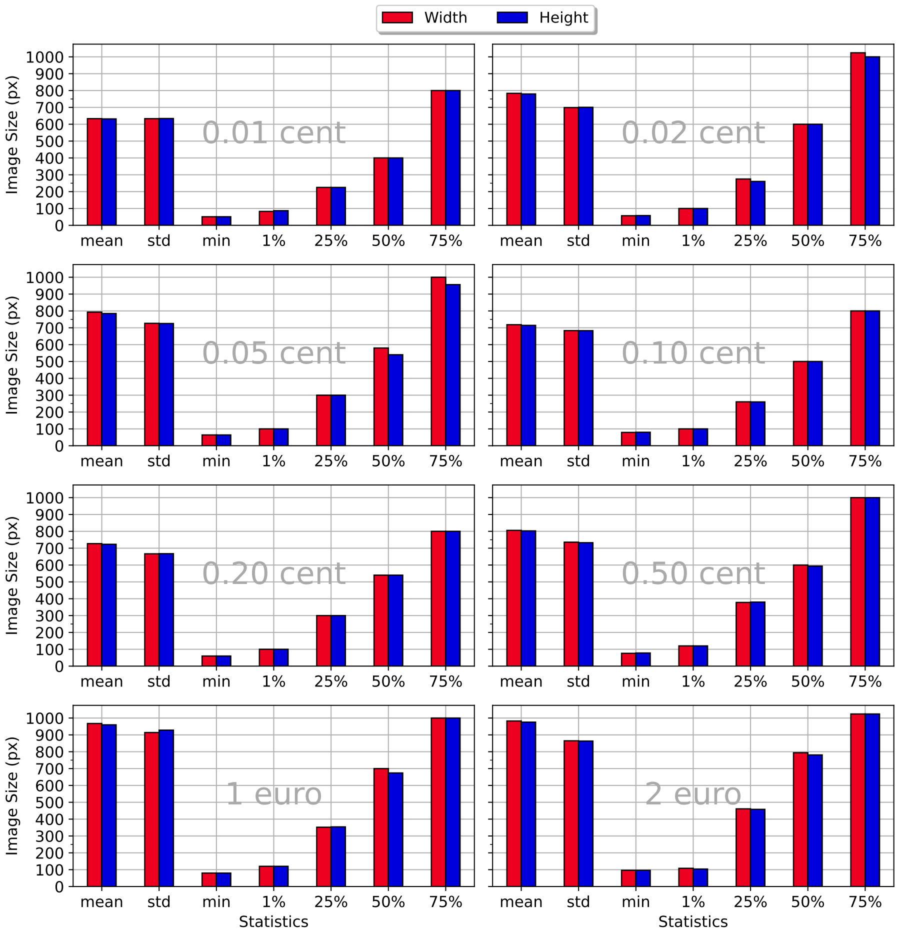

Fig. 4: Statistics about the quality of the images in the *EURO-Coin dataset*.

operations have partially improved the images, it was necessary to use further preprocessing and filtering techniques on the images to make them suitable for deep learning models. To this end, we have chosen to use some filters on the images with the aim of emphasizing the reliefs of the coins and preparing the images for the training and testing steps.

In our experimental evaluation, we have chosen to test Sobel, Laplace, Uniform, and Prewitt filters to detect edges in order to improve the exposure of the reliefs on the coins. Figure 5 shows an example of filters applied on the image of 10 cents.

The Sobel filter uses a 3×3 or 5×5 convolution matrix to calculate the luminosity derivative at each point in the image, which is often used for image segmentation [49]. Specifically, there are two Sobel matrices, one for detecting horizontal edges and the other for vertical edges. Applying this filter provides

Aspect		1c	2c	5c	10c	20c	50c	1€	2€	Total per aspect
0.4	5:12	-	1	-	-	-	-	2	1	4
0.5	1:2	-	5	13	5	4	6	7	15	55
0.6	5:9	4	2	4	5	4	6	36	37	98
0.7	2:3	21	7	5	17	18	9	26	13	116
0.8	3:4	28	13	28	26	34	32	91	91	343
0.9	6:7	119	40	55	69	67	45	119	126	640
1.0	1:1	1252	933	1122	1220	1175	1154	1566	3200	11622
Total per Value		1424	1001	1227	1342	1302	1252	1847	3483	**12878**

Table 3: Statistics about the aspect ratio of coin images in the dataset.

two images, each representing the horizontal or vertical edges of the original image. Figure 5b shows an example of Sobel filter applied on the image of 10 cents.

The Laplace filter uses a 3×3 or 5×5 convolution matrix to highlight image details. This filter calculates the second derivative of the image luminosity and is generally used to detect edges, improve image sharpness, and emphasize texture features. Figure 5c shows an example of Laplace filter applied on the image of 10 cents.

The Uniform filter is mainly used to reduce image noise and imperfections. This filter calculates the average of the surrounding pixels for each pixel in the image, creating a more uniform and homogeneous image. It is generally used to improve image quality and prepare it for subsequent computer vision processing. Figure 5d shows an example of Uniform filter applied on the image of 10 cents.

Finally, the Prewitt filter uses a 3×3 convolution matrix to calculate the luminosity derivative of the image and detect edges. This filter is similar to the Sobel filter and is used to detect edges in an image and improve its sharpness [49]. Figure 5e shows an example of Prewitt filter applied on the image of 10 cents. Generally, it is used in combination with other filters to improve image quality and detect edges.

In our study, we evaluate the effectiveness of these filters individually or in combination with others, with the aim of improving image quality and more accurately detecting edges and reliefs.

4.4 Network Architecture

The *EURO-Coin dataset* is used to train a Convolutional Neural Network which consists of four blocks: the first two contain two convolutional layers; the third is a self-attention layer; the last block consists of one linear layer.

In particular, this model accepts images having three channels and a specific dimension of 128×128. Then, we have a two sequence of two convolution layers consisting of 8 filters each one with a kernel size equal to 3. On the output of the second convolutional layer, we compute a ReLU activation function


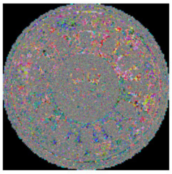

(b) Sobel filter.

(c) Laplace filter.


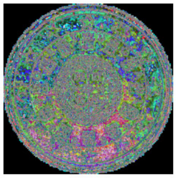

(a) Original coin. (d) Uniform filter. (e) Prewitt filter.

Fig. 5: Example of filters applied on the image of 10 Portuguese cents.

and then apply a max pooling with a window 2×2 to reduce the image size and keep only the most important features.

We use two other convolutional layers with 16 filters each one with a kernel size equal to 3. As in the previous case, we compute a ReLU activation function and a max pooling with a window 2×2. After this computation, we have a self-attention layer with an embedding dimension for each input token of 16, and with a multi-head attention of 4 heads.

For each image, we obtain a feature vector that is the input to a linear layer and a softmax activation function. The output dimension of this last layer depends on the classification task. In the country case, the output is 23 ($|S| = 23$), while in the value case, the output is 8 ($|V| = 8$). Except for the last layer dimension, we have used the same CNN model for both classification tasks. The network architecture is reported in Figure 6.

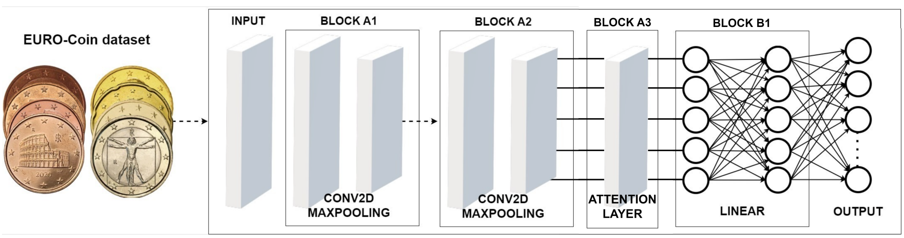

Fig. 6: Architecture of our Convolutional Neural Network

5 Experimental Evaluation

In this Section, we outline the experimental evaluation performed to address the country and value classification tasks of coins from their relief sides. In particular, in Section 5.1, we provide details about the experimental settings for our evaluation. In Section 5.2, we describe the results achieved in the two classification tasks, and in Section 5.3 we perform an ablation study to understand the impact of our preprocessing activities on the model performance.

5.1 Experimental Settings

The experimental evaluation was conducted using the EURO-Coin dataset introduced in Section 4.2. For the identification of the country through the currency, 23 classes were identified (i.e., one for each country that uses the euro as the current currency), while for the task for value identification, 8 classes were identified (i.e., one for each type of coin: 2 euros, 1 euro, 50 cents, 20 cents, 10 cents, 5 cents, 2 cents, and 1 cent).

To evaluate the classification performances of our models, we decided to split the dataset into training and testing sets. Our EURO-Coin dataset contains 12878 images and many examples for each country and value class, thus representing the scenario thoroughly. Therefore, we decided to make a stratified 80/20 training/test split, where 80% of images are used for training and the remaining 20% are for estimating the classification performances. In the ablation study, we keep the same training/test split for all the experiments.

The network was trained using a Stochastic Gradient Descent optimizer with a mini-batch size of 32. We use *Reduce Learning Rate on Plateau* to monitor training performances at each epoch and automatically reduce the learning rate when performance does not improve for a certain number of consecutive epochs. This strategy enables the neural network to better fit the data and avoid getting stuck in the local minima of the loss function.

We used the functions of the *scipy* library for multidimensional image processing, and its package *scipy.ndimage*, which contains several functions and filters for multidimensional image processing. The images have been transformed through the *transform* method, which enables to combine multiple transformations to be applied to an image. Moreover, each image has been transformed by applying a resize to 128×128 and then converted into a tensor.

The CNN has been implemented using Python version 3.9 and with the support of PyTorch 1.13.1 and CUDA 11.9. All the experiments have been executed on a workstation with an Intel i7 CPU at 2.30 GHz, 6-core, and 32GB of memory, equipped with a GPU NVIDIA 3070 GPU with 8GB of RAM dedicated.

5.2 Evaluation

As for the country classification task, our dataset consists of 12878 images of the side with reliefs of coins divided into 23 classes. We first split the dataset into training and test, and then preprocess the training images by resizing them to $w_s = 128$ and $h_s = 128$ and applying a Sobel filter for edge detection (i.e., $T = \{Resize, Sobel\}$). Then, we train our CNN model in order to predict the countries that coined them. Finally, the classification metrics of our CNN model are computed on the testing images that are preprocessed as those of training, and the corresponding results are reported in Table 4.

Accuracy	**Precision**	**Recall**	**F1-Score**
0.869	0.871	0.869	0.870

Table 4: Performance of our CNN model in the country classification task.

From the analysis of Table 4, we observe that our CNN model achieves good performances. Indeed, we obtain an Accuracy of 0.869, a Precision of 0.871, a Recall of 0.869, and an F1-Score of 0.870. These results are promising also due to the fact that we are facing a multi-class classification problem with 23 classes that is not trivial. Indeed, some countries could have the same design for a coin or could have changed the relief over the years having therefore more than one drawing for each class. We can conclude that our model can identify interesting patterns useful to predict the country of each coin. In order to better investigate the performance of our model, we report the related confusion matrix in Figure 7.

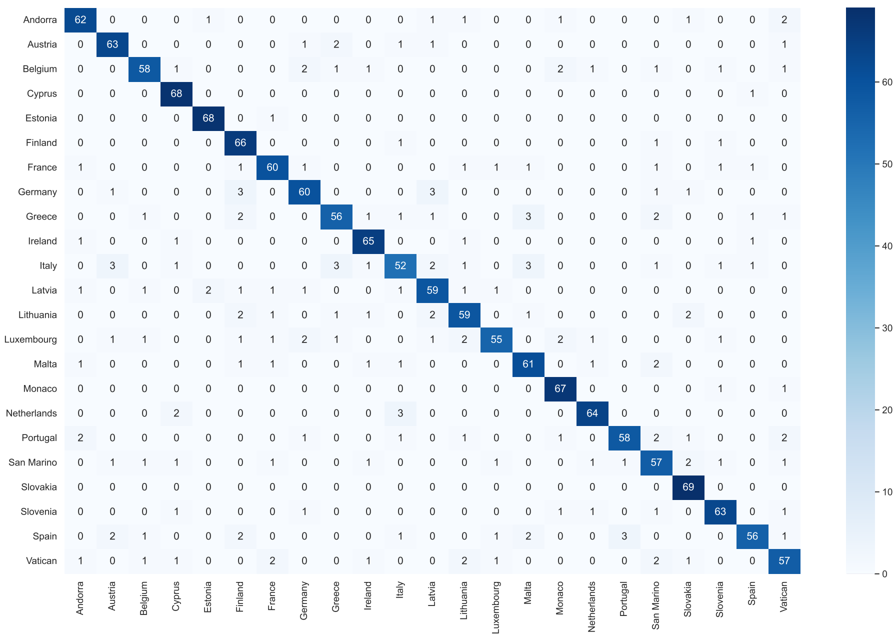

Fig. 7: Confusion matrix of the country classification task

In Figure 7, we note that our model is able to classify more than the majority of the coins for each country. Indeed, the diagonal of the confusion matrix is the one reporting the highest values. However, we can also observe that in some cases, the model misclassifies some images, which could be due to several factors. For instance, think of Belgium, Netherlands, and Monaco have coined similar coins over the years. Another example regards the 1€ coins of the Netherlands and Luxembourg that are also very similar.

As for the value classification task, we have 12878 images divided into 8 classes. Similar to the previous task, we split the dataset into training and testing. Then, we resize the training images to $w_s = 128$ and $h_s = 128$ and apply Sobel filters for edge detection (i.e., $T = \{Resize, Sobel\}$). Our CNN model is trained on the preprocessed training images in order to predict the value of the coins. The classification performances of our CNN model are computed on the testing images, and the obtained results are reported in Table 5.

Accuracy	Precision	Recall	F1-Score
0.964	0.964	0.964	0.964

Table 5: Performance of our CNN model in the value classification task.

From the analysis of Table 5, we observe that our CNN model achieves an Accuracy of 0.964, a Precision of 0.964, a Recall of 0.964, and an F1-Score of 0.964. Even in this case, our model obtains interesting performance in this multiclass classification task. We compute the corresponding confusion matrix in Figure 8.

From Figure 8, we observe that our model is capable of identifying the right classes for the majority of coins. However, there is an interesting pattern in the cases in which the model does not classify correctly. Indeed, there are two squares in the confusion matrix that corresponds to classes that are similar to one another, which are *(i)* 1-cent, 2-cent, 5-cent, and *(ii)* 10-cent, 20-cent, 50-cent. These groups of coins contain images that have similar patterns due to some decisions of the country that minted them. For instance, Malta and Luxembourg have similar drawings printed on the 1-cent, 2-cent, and 5-cent, while Belgium typically uses similar patterns for all those coins.

Finally, we point out two interesting aspects of our CNN model. The first regards the fact that we used edge detection filters to classify images. In this way, we are able to remove part of the background in the images, and highlight the patterns of the drawings present on the sides with relief of the coins. As will be clear in Section 5.3, the natural images are not suitable for training the CNN even if they contain color features. Indeed, we show that the extraction of the patterns from the drawings of the coins is fundamental for the country and value classification tasks. The second interesting aspect concerns the architecture of the CNN for these two tasks. In fact, we have employed the same CNN model and the same filter for both cases. This result highlights the

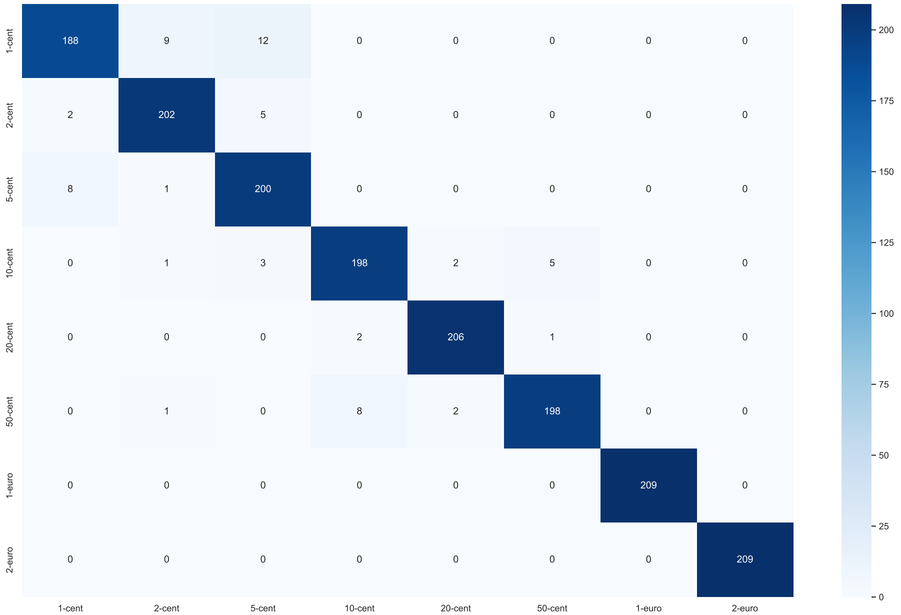

Fig. 8: Confusion matrix of the value classification task.

fact that the same CNN architecture can perform well in these two tasks that have completely different objectives.

5.3 Ablation study

In this section, we conduct an ablation study to assess the impact of different preprocessing activities on the performance of our CNN model in the country and value classification tasks. As in the previous case, we use a stratified 80/20 training/test split and keep the same resulting image sets for all the experiments in the ablation study. We train our CNN model using the images in the training set that have been preprocessed with a variety of edge detection and noise reduction filters, including the Sobel, Laplace, and Prewitt filters for edge detection and the Uniform filter for noise reduction, and then measure the corresponding performance metrics on the images in the testing set preprocessed as the training ones [50]. Through this analysis, we aim to determine the effectiveness of the filters and identify the most suitable filter(s) for optimizing the classification accuracy of our model. Furthermore, we tested our CNN model with and without the self-attention layer.

As for the country classification task, we report the obtained results from our ablation study in Table 6.

From the analysis of Table 6, we observe that the best performance in terms of Accuracy, Precision, Recall, and F1-Score is achieved with the attention layer and the application of the Sobel filter. The other edge detection filters

	Filter(s)	Accuracy	Precision	Recall	F1-Score
Without Attention Layer	Natural	0.067	0.009	0.067	0.016
	Sobel	0.739	0.740	0.739	0.738
	Laplace	0.534	0.604	0.534	0.540
	Uniform	0.063	0.040	0.063	0.034
	Prewitt	0.453	0.723	0.453	0.477
	Sobel + Laplace	0.700	0.703	0.700	0.700
	Sobel + Prewitt	0.640	0.680	0.640	0.643
	Sobel + Uniform	0.578	0.708	0.578	0.600
	Sobel + Prewitt + Laplace	0.394	0.394	0.394	0.431
	Sobel + Uniform + Laplace	0.590	0.714	0.590	0.612
	Sobel + Uniform + Prewitt	0.368	0.663	0.368	0.387
With Attention Layer	**Filter(s)**	**Accuracy**	**Precision**	**Recall**	**F1-Score**
	Natural	0.139	0.142	0.132	0.150
	Sobel	**0.869**	**0.871**	**0.869**	**0.870**
	Laplace	0.553	0.637	0.570	0.602
	Uniform	0.142	0.118	0.324	0.132
	Prewitt	0.178	0.17	0.17	0.17
	Sobel + Laplace	0.721	0.721	0.766	0.745
	Sobel + Prewitt	0.722	0.722	0.722	0.722
	Sobel + Uniform	0.443	0.429	0.500	0.450
	Sobel + Prewitt + Laplace	0.450	0.443	0.466	0.447
	Sobel + Uniform + Laplace	0.451	0.429	0.317	0.432
	Sobel + Uniform + Prewitt	0.203	0.202	0.212	0.210

Table 6: Ablation study on the country classification task.

obtain lower results, such as the Laplace or the Prewitt. Neither the sequential application of two or three detection filters (such as Sobel + Laplace, Sobel + Prewitt, or Sobel + Uniform + Laplace) reaches the same performance as the Sobel case. It is interesting to note that all the metrics are very low in the Natural case, where the images are only resized. The CNN fails to identify interesting patterns for the classification, which affects its performance. Finally, we note that the usage of an attention layer increases all the classification metrics in most of the cases.

Afterward, we conducted the same ablation study for the value classification task, and the corresponding results are reported in Table 7.

From the analysis of Table 7, we observe that our CNN model achieves the highest classification accuracy with the attention layer and the application of Sobel filter. The Sobel filter is performing well also without the attention layer, but this last allows us to obtain higher results. In contrast, the single application of the Uniform and Prewitt filters did not improve much the classification performances compared to the Natural case. Finally, we note that while the Natural case resulted in higher classification accuracy in the value prediction task with respect to the country prediction, also in this case the application of edge detection filters allows our model to identify useful patterns and significantly improve performance across all evaluation metrics.

	Filter(s)	Accuracy	Precision	Recall	F1-Score
Without Attention Layer	Natural	0.445	0.269	0.316	0.259
	Sobel	0.796	0.786	0.771	0.777
	Laplace	0.532	0.512	0.416	0.359
	Uniform	0.308	0.268	0.175	0.110
	Prewitt	0.478	0.534	0.397	0.415
	Sobel + Laplace	0.818	0.810	0.789	0.796
	Sobel + Prewitt	0.817	0.803	0.798	0.800
	Sobel + Uniform	0.550	0.576	0.471	0.487
	Sobel + Prewitt + Laplace	0.823	0.804	0.801	0.802
	Sobel + Uniform + Laplace	0.565	0.595	0.500	0.512
	Sobel + Uniform + Prewitt	0.769	0.773	0.738	0.748
	Filter(s)	**Accuracy**	**Precision**	**Recall**	**F1-Score**
With Attention Layer	Natural	0.460	0.455	0.460	0.405
	Sobel	**0.964**	**0.964**	**0.964**	**0.964**
	Laplace	0.363	0.384	0.415	0.357
	Uniform	0.244	0.242	0.326	0.259
	Prewitt	0.179	0.194	0.239	0.189
	Sobel + Laplace	0.941	0.942	0.941	0.941
	Sobel + Prewitt	0.959	0.959	0.959	0.959
	Sobel + Uniform	0.922	0.923	0.922	0.922
	Sobel + Prewitt + Laplace	0.930	0.932	0.930	0.930
	Sobel + Uniform + Laplace	0.962	0.962	0.962	0.962
	Sobel + Uniform + Prewitt	0.962	0.962	0.962	0.962

Table 7: Ablation study on the value classification task.

6 Conclusion

In this paper, we have presented two approaches that exploit Convolutional Neural Networks to automatically identify the country and value of euro coins from the side with reliefs. These approaches aim to provide a useful tool for amateur and expert numismatists to improve the time-consuming processes of manual coin identification through AI approaches. As we have previously discussed, these approaches have proven to be efficient on the euro coins minted by the Member States of the European Union, but are easily generalizable to other coin classification scenarios. In order to prove the validity of our proposals, we have tested them on the new EURO-Coin dataset, by applying different manual and automatic preprocessing steps on the coins with the aim to improve the classification performances. Finally, we have conducted an ablation study to identify the most important preprocessing activities to be applied in our approaches.

In the future, we plan to extend this work in several directions. For instance, we can consider an ensemble of multiple Convolutional Neural Network models to make predictions. Each model can be trained with a slightly different architecture or hyperparameters, and the predictions can be combined using a weighted average or other ensemble methods. Finally, we plan to test our approaches to other coin types of different continents using the side with reliefs. This could involve retraining the model with additional data or fine-tuning the existing model to classify the new types of coins.

Data Availability

The dataset used for our study is publicly available at the link https://github.com/cciro94/EuroCoinDataset.

Conflict of interest

The authors declare that they have no conflict of interest.

References

[1] Ioannakis, G., Bampis, L., Koutsoudis, A.: Exploiting artificial intelligence for digitally enriched museum visits. Journal of Cultural Heritage **42**, 171–180 (2020)

[2] Fiorucci, M., Khoroshiltseva, M., Pontil, M., Traviglia, A., Del Bue, A., James, S.: Machine learning for cultural heritage: A survey. Pattern Recognition Letters **133**, 102–108 (2020)

[3] Fontanella, F., Colace, F., Molinara, M., Di Freca, A.S., Stanco, F.: Pattern recognition and artificial intelligence techniques for cultural heritage. Elsevier (2020)

[4] Levitt, R., Celia, C., Diepeveen, S., Chonaill, S.N., Rabinovich, L., Tiessen, J.: Assessing the impact of arts and humanities research at the university of cambridge. technical report. RAND Corporation (2010)

[5] Lähdesmäki, T.: Politics of tangibility, intangibility, and place in the making of a european cultural heritage in eu heritage policy. International Journal of Heritage Studies **22**(10), 766–780 (2016)

[6] Zambanini, S., Kampel, M., Schlapke, M.: On the use of computer vision for numismatic research. In: Proceedings of the 9th International Conference on Virtual Reality, Archaeology and Cultural Heritage, Goslar,Germany, pp. 17–24 (2008)

[7] Feng, B.-Y., Sun, K., Atighechian, P., Suen, C.Y.: Computer recognition and evaluation of coins. In: Handbook of Pattern Recognition and Computer Vision, pp. 141–158. World Scientific, Wordwide (2016)

[8] Cooper, J., Arandjelović, O.: Understanding ancient coin images. In: Recent Advances in Big Data and Deep Learning: Proceedings of the INNS Big Data and Deep Learning Conference (INNSBDDL2019), pp. 330–340. Springer, Genova, Italy (2020)

[9] Kim, J., Pavlovic, V.: Discovering characteristic landmarks on ancient coins using convolutional networks. Journal of Electronic Imaging **26**(1), 011018–011018 (2017)

[10] Capece, N., Erra, U., Ciliberto, A.V.: Implementation of a coin recognition system for mobile devices with deep learning. In: Proceedings of the 12th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS), pp. 186–192. IEEE, Naples, Italy (2016)

[11] Lin, W.-J., Jhuo, S.-S.: Coin recognition based on texture classification on ring and fan areas of the coin image. In: IEEE International Instrumentation and Measurement Technology Conference, pp. 1–6. IEEE, Taipei, Taiwan (2016)

[12] Chauhan, V., Joshi, K.D., Surgenor, B.: Machine vision for coin recognition with anns: Effect of training and testing parameters. In: Proceedings of the Engineering Applications of Neural Networks, pp. 523–534. Springer, Athens, Greece (2017)

[13] Schlag, I., Arandjelovic, O.: Ancient roman coin recognition in the wild using deep learning based recognition of artistically depicted face profiles. In: Proceedings of the IEEE International Conference on Computer Vision Workshops, pp. 2898–2906 (2017)

[14] Anwar, H., Anwar, S., Zambanini, S., Porikli, F.: Deep ancient roman republican coin classification via feature fusion and attention. Pattern Recognition **114**, 107871 (2021)

[15] Anwar, H., Zambanini, S., Kampel, M.: A rotation-invariant bag of visual words model for symbols based ancient coin classification. In: Proceedings of the IEEE International Conference on Image Processing (ICIP), pp. 5257–5261. IEEE, Paris, France (2014)

[16] Zambanini, S., Kampel, M.: Coarse-to-fine correspondence search for classifying ancient coins. In: Computer Vision-ACCV 2012 Workshops: ACCV 2012 International Workshops, pp. 25–36. Springer, Daejeon, South Korea (2013)

[17] Aslan, S., Vascon, S., Pelillo, M.: Two sides of the same coin: Improved ancient coin classification using graph transduction games. Pattern Recognition Letters **131**, 158–165 (2020)

[18] Lu, Y.: Artificial intelligence: a survey on evolution, models, applications and future trends. Journal of Management Analytics **6**(1), 1–29 (2019)

[19] Sood, S.K., Rawat, K.S., Kumar, D.: Scientometric analysis of ict-assisted intelligent control systems response to covid-19 pandemic. Neural Computing and Applications, 1–21 (2023)

[20] Celard, P., Iglesias, E., Sorribes-Fdez, J., Romero, R., Vieira, A.S., Borrajo, L.: A survey on deep learning applied to medical images: from simple

artificial neural networks to generative models. Neural Computing and Applications **35**(3), 2291–2323 (2023)

[21] Pisoni, G., Díaz-Rodríguez, N., Gijlers, H., Tonolli, L.: Human-centered artificial intelligence for designing accessible cultural heritage. Applied Sciences **11**(2), 870 (2021)

[22] Peng, X., Zhao, H., Wang, X., Zhang, Y., Li, Z., Zhang, Q., Wang, J., Peng, J., Liang, H.: C3n: content-constrained convolutional network for mural image completion. Neural Computing and Applications **35**(2), 1959–1970 (2023)

[23] Yu, T., Lin, C., Zhang, S., Wang, C., Ding, X., An, H., Liu, X., Qu, T., Wan, L., You, S., *et al.*: Artificial intelligence for dunhuang cultural heritage protection: The project and the dataset. International Journal of Computer Vision **130**(11), 2646–2673 (2022)

[24] Modi, S., Bawa, D., et al.: Image processing based systems and techniques for the recognition of ancient and modern coins. arXiv preprint arXiv:1312.6599 (2013)

[25] Zambanini, S., Kavelar, A., Kampel, M.: Classifying ancient coins by local feature matching and pairwise geometric consistency evaluation. In: Proceedings of the 22nd International Conference on Pattern Recognition, pp. 3032–3037. IEEE, Washington D.C., United States (2014)

[26] Pan, X., Tougne, L.: Image analysis and deep learning for aiding professional coin grading. In: Proceedings of the International Conference on Image and Video Processing, and Artificial Intelligence, vol. 10836, pp. 21–25. SPIE, New York, United States (2018)

[27] Alcázar-Blanco, A.C., Paule-Vianez, J., Prado-Román, M., Coca-Pérez, J.L.: Generalized regression neuronal networks to predict the value of numismatic assets. evidence for the walking liberty half dollar. European Research on Management and Business Economics **27**(3), 100167 (2021)

[28] Conn, B., Arandjelović, O.: Towards computer vision based ancient coin recognition in the wild—automatic reliable image preprocessing and normalization. In: International Joint Conference on Neural Networks (IJCNN), pp. 1457–1464. IEEE, Anchorage, Alaska (2017)

[29] Hmood, A.K., Suen, C.Y.: Statistical edge-based feature selection for counterfeit coin detection. Multimedia Tools and Applications **79**(39-40), 28621–28642 (2020)

[30] Chang, H., Wei, Z., Huang, L., Nie, J., Zhang, W., Wang, L.: Coin classification in a complex environment. In: Proceedings of the 10th International

Conference on Graphics and Image Processing (ICGIP 2018), vol. 11069, pp. 453–461. SPIE, Chengdu, China (2019)

[31] Heinecke, A., Mayer, E., Natarajan, A., Jung, Y.: Unsupervised statistical learning for die analysis in ancient numismatics. arXiv preprint arXiv:2112.00290 (2021)

[32] Kampel, M., Zaharieva, M.: Recognizing ancient coins based on local features. In: Advances in Visual Computing: 4th International Symposium, ISVC 2008, pp. 11–22. Springer, Las Vegas, NV, USA (2008)

[33] Xiang, Y., Yan, W.Q.: Fast-moving coin recognition using deep learning. Multimedia Tools and Applications **80**, 24111–24120 (2021)

[34] Hmood, A.K., Suen, C.Y., Lam, L.: An enhanced histogram of oriented gradient descriptor for numismatic applications. Pattern Recognition and Image Analysis **28**, 569–587 (2018)

[35] Anwar, H., Zambanini, S., Kampel, M.: Coarse-grained ancient coin classification using image-based reverse side motif recognition. Machine Vision and Applications **26**, 295–304 (2015)

[36] 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)

[37] Munipalle, V.K., Nelakuditi, U.R.: Performance analysis of coin matching using local texture features. In: Proceedings of the International Conference on Cutting-edge Technologies in Engineering (ICon-CuTE), pp. 69–74. IEEE, Lucknow-Deva Road, India (2019)

[38] Kim, J., Pavlovic, V.: Improving ancient roman coin recognition with alignment and spatial encoding. In: Proceedings of the European Conference on Computer Vision, pp. 149–164. Springer, Zurich, Switzerland (2014)

[39] Zambanini, S., Kavelar, A., Kampel, M.: Improving ancient roman coin classification by fusing exemplar-based classification and legend recognition. In: New Trends in Image Analysis and Processing ICIAP International Workshops, pp. 149–158. Springer, Naples, Italy (2013)

[40] Kavelar, A., Zambanini, S., Kampel, M.: Word detection applied to images of ancient roman coins. In: Proceedings of the 18th International Conference on Virtual Systems and Multimedia, pp. 577–580. IEEE, Milan, Italy (2012)

[41] Arandjelović, O.: Automatic attribution of ancient roman imperial coins.

In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 1728–1734. IEEE, San Francisco, California (2010)

[42] Kim, S., Lee, S.H., Ro, Y.M.: Image-based coin recognition using rotation-invariant region binary patterns based on gradient magnitudes. Journal of Visual Communication and Image Representation **32**, 217–223 (2015)

[43] Nölle, M., Rubik, M., Hanbury, A.: Results of the muscle cis coin competition. In: Proceedings of the Muscle CIS Coin Competition Workshop, Berlin, Germany, pp. 1–5 (2006)

[44] Manzoor, S., Ali, N., Raees, M., Khan, K.A., Usama, M., Ayub, A.A.: Ancient coin classification based on recent trends of deep learning. In: Proceedings of the 1st International Virtual Conference on Visual Pattern Extraction and Recognition for Cultural Heritage Understanding, pp. 12–20. CEUR-WS.org, Pescara, Italy (2022)

[45] Kiourt, C., Evangelidis, V.: Ancoins: Image-based automated identification of ancient coins through transfer learning approaches. In: Proceedings of the International Conference on Pattern Recognition (ICPR), pp. 54–67. Springer, Virtual Event (2021)

[46] Khashman, A., Sekeroglu, B., Dimililer, K.: Intelligent coin identification system. In: Proceedings of the IEEE International Conference on Control Applications, pp. 1226–1230. IEEE, Munich, Germany (2006)

[47] Marco, R., Olaf, R., Burkhardt, H.: A fast and reliable coin recognition system. In: Proc. of the Joint Pattern Recognition Symposium, pp. 415–424. Springer, Heidelberg, Germany (2007)

[48] Crawford, M.H.: Roman Republican Coinage vol. 1. Cambridge University Press, Cambridge, United Kingdom (1974)

[49] Zhou, W., Du, X., Wang, S.: Techniques for image segmentation based on edge detection. In: IEEE International Conference on Computer Science, Electronic Information Engineering and Intelligent Control Technology, pp. 400–403. IEEE, Fuzhou, China (2021)

[50] Shrivakshan, G., Chandrasekar, C.: A comparison of various edge detection techniques used in image processing. International Journal of Computer Science Issues (IJCSI) **9**(5), 269 (2012)