

EAGLE – Europeana Network of Ancient Greek and Latin Epigraphy A Digital Bridge to the Ancient World

Giuseppe Amato¹, Vittore Casarosa¹, Philippe Martineau², Silvia Orlandi³,
Raffaella Santucci³ and Luca Marco Carlo Giberti³

CNR-Isti¹, Eureva², Università di Roma La Sapienza³
giuseppe.amato@isti.cnr.it, casarosa@isti.cnr.it, philm@eureva.fr,
silvia.orlandi@uniroma1.it, raffaella.santucci@uniroma1.it, luca.giberti@stx.oxon.org

Abstract This paper discusses the experience of developing a mobile application to increase the use and visibility of EAGLE, the Europeana network of Ancient Greek and Latin Epigraphy. EAGLE is a project financed by the European Commission. Its main aims are to bring together the most prominent European institutions and archives in the field of Classical Latin and Greek epigraphy and to provide Europeana with a comprehensive collection of unique historical sources related to inscriptions. EAGLE also aims to provide a single, user-friendly portal by which to search and browse the majority of surviving inscriptions from the Greco-Roman world. In order to increase the usefulness and visibility of EAGLE, two applications are being developed: a storytelling application to allow teachers and experts to assemble epigraphy-based narratives for the benefit of less experienced users, and a mobile application to enable tourists and scholars to obtain detailed information about the inscriptions they are looking at by taking pictures with their smartphones. In this paper, we will focus on the EAGLE mobile application and give an outline of its architecture and design.

1 Introduction

The cultural identity of the Western world is rooted in the Greco-Latin tradition. From philosophy to architecture, geometry to law, and rhetoric to literature, the presence of the ancients lingers still in the way we think, live, and express ourselves. Only a small fraction of all ancient Greco-Roman texts has survived to modern times, leaving sizeable gaps in the historiographic record. A precious alternative source of historical evidence can be found in the form of ancient inscriptions. These are invaluable ‘time capsules’ that provide a myriad of useful facts and allow us to cast light on otherwise undocumented historical events, laws, and customs. EAGLE (Europeana network of Ancient Greek and Latin Epigraphy) is a Best Practice Network co-funded by the European Commission, which will allow for the virtual reconstruction of the inscriptions’ original archaeological and historical context. It will make accessible the majority of the surviving inscriptions of the Greco-Roman world, complete with the essential information about them.

Within EAGLE, the word “inscription” is used to indicate the document engraved on non-perishable materials whose text, along with a rich set of metadata (on its history, dating, preservation, place of finding, present location,

etc.), is stored in the collections of 15 partners, the so-called Content Providers. The inscriptions collected by EAGLE come from 25 EU countries, providing more than 1.5 million images and related metadata, which includes translations of selected texts for the benefit of the general public. Taken altogether, the EAGLE collections represent approximately 80% of the total number of inscriptions in the Mediterranean area.

The descriptions (i.e. the metadata) of the various inscriptions are being integrated in an online repository, both for ingestion into Europeana and for the provision of services to interested users and scholars directly from the EAGLE portal.

The main service provided by the portal is the ability to find information about inscriptions through a sophisticated search interface. It will be possible to perform full-text searches over the entire body of text associated with the inscriptions, using a simple interface “à la Google”, or more complex queries that will make use of the various fields describing an inscription. For the advanced search EAGLE has developed seven controlled multilingual vocabularies related to the main fields used to describe an inscription (i.e. type of inscription, type of object, material, writing, decoration, state of preservation, and dating criteria). Users registered to the EAGLE system will be able to store the results of a search in their ‘private area’, to avoid having to retrieve them again in subsequent visits, and will also be able to store comments and notes about the retrieved material.

In addition to textual search, an important feature offered by EAGLE is the ability to find information about an inscription by providing an image as a ‘query image’. Taking advantage of advanced image recognition technology provided by EAGLE partner CNR-ISTI, which has extensive experience in this field, EAGLE will search in its data base of images and will return the information associated with the recognised image.

2 EAGLE services

A high level overview of the Aggregation and Image Retrieval system (AIM), the main component of EAGLE, is depicted in Fig. 1, which shows its main two subsystems, namely the Metadata Aggregation System (MAS) and the Image Retrieval System (IRS).




Figure 1. High level view of the EAGLE services

The Metadata Aggregation System collects the description of the inscriptions held by the EAGLE Content Providers, consisting of metadata expressed in

XML format. Each collected metadata record needs to be transformed conforming to a homogeneous structural format (the EAGLE Metadata Format) and stored for further processing. Digital image resources possibly pointed to by metadata records are fetched following their URLs and stored as well in a component of the MAS, for subsequent use by the IRS system.

The metadata needs to be cleaned (i.e. harmonised by applying vocabularies for allowed values) and possibly curated (i.e. selectively edited). The set of cleaned and curated metadata records constitutes the EAGLE Information Space. Once the information space reaches the desired quality, its content is ready to be indexed, ingested into Europeana and disseminated to the public and to external applications (e.g. web portals and mobile applications). More specifically, the information space is i) used as input for the indexing process to make it accessible from portals via web searches, and ii) exposed record-by-record toward Europeana or any other third-party system that might be interested in such a content.

The Image Retrieval System will receive as an input the image of an inscription and will provide effective and efficient image similarity search and image content recognition of inscriptions. The images provided by the Content Providers and collected by the MAS will be processed by the IRS components (Image Feature Extractor, Image Indexer, CBIR Index; see next paragraph for explanations) in order to build the index that will allow a fast and efficient similarity search during the query phase.

In addition, for those inscriptions for which a set of images is available (training set), each training set will be processed to extract the main features characterising the inscription. The training sets and the characterising features will be the base for building another index, to be used by the image recogniser in order to decide if an image received during the query phase can be classified as belonging to one of the existing sets or not. In this way, in many cases, the recogniser is able to associate a query image with the correct inscription even if the image given in the query was never stored in the database.

3 The Image Retrieval System

The Image Retrieval System (IRS) will provide two modes of recognition of an image provided as an input for the query. In the first mode, called Similarity Search, the query result will be a list of images contained in the data base, ranked in order of similarity to the query image. In the second mode, called Recognition Mode, the result of the query will be the information associated with the recognized image, or a message indicating that the image was not recognised. In this second mode the system will be able to treat query images in which the inscription appears in any position, even as a smaller part of a more general picture (e.g. the photo of an archaeological site, or the scanned image of the page of a book).

One of the main components of the IRS is the Image Feature Extractor, which analyses the visual content of an image (to be added to the database or provided as a query) in order to generate a visual descriptor. Visual descriptors are mathematical descriptions of the features in an image. A feature captures a certain visual property of an image, either globally for the entire image or locally for a small group of pixels. Global features are computed to capture the overall characteristics of an image, such as those reflecting colour, texture, shapes, and interest (or salient) points in an image. Local features are low-level descriptions of keypoints in an image. Keypoints are interest points in an image that are invariant with respect to changes in scale and orientation, and therefore are very useful for determining the similarity of two images in all those cases

where the second image has a different orientation or scale. The result of global and local extraction of visual features is a mathematical description of the image visual content that can be used to compare different images, judge their similarity, and identify common content. The IRS will use global visual descriptors for the Similarity Search task, and local visual descriptors for the Recognition task.

As stated before, global features are based on the overall properties of an image, such as the histogram of colours present in the image, the relative position of shapes identified in the image, the texture (information about the spatial arrangements of colours) of the image or of selected sub-regions. A number of mathematical libraries already exist for extracting global features from an image. Figure 2 shows the local features extraction process. First a number of keypoints (usually of the order of one thousand) are identified and selected in the image (coloured dots in the second image), and then for each point its coordinates, scale and orientation are computed. In this case as well there are already existing mathematical libraries that can be customised by taking into account the specific features of inscription images.

In order to efficiently and effectively execute the retrieval and recognition process, visual descriptors extracted from images have to be inserted in an index for efficient similarity search and ranking. The Image Indexer component of the IRS analyses the visual descriptors and processes them for insertion in the index. State of the art technology in CBIR (Content Based Image Retrieval) means that the visual features of an image (global or local) are encoded into a text string, so that in the search and retrieval phase the mature technology of text-based search engines can be exploited. The IRS will use two different approaches to generate a convenient text encoding for local and global visual features.

Local features will be encoded using an approach called “Bag of Features”, where a textual vocabulary of visual words will be created starting from all the local descriptors of the whole dataset. The set of all the local descriptors of all the images is divided into a number of clusters (usually in the order of 100 thousand clusters) and to each cluster is assigned a (usually random) textual tag. The set of all tags becomes the “vocabulary” of visual words related to the whole set of images. At this point each image can be described by a set of “words” in the vocabulary (the textual tags), corresponding to the clusters containing the visual features of the image.

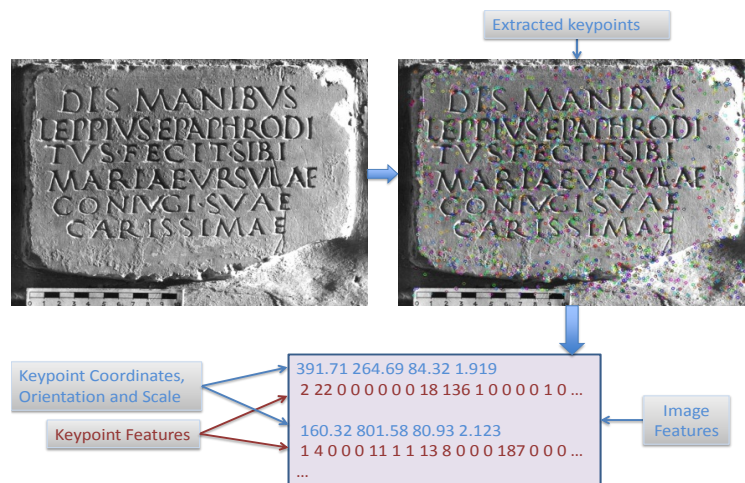


Figure 2. Local features extraction process

Global Features will be encoded using an approach called “perspective based space transformation”. The idea driving this technique is that when two global descriptors (GDs) are very similar (with respect to a given similarity function), they ‘view the world’ around them in a similar way. To capture this ‘view’ a set of selected images, called anchor descriptors (usually in the order of a few thousands), is defined (usually in a random fashion) and the similarity function is defined to represent a ‘distance’ between two images. In this way it is possible to encode an image by first defining the order in which the image ‘sees’ the anchor descriptors, ranked by their distance to the image itself and then encoding with a textual string the specific permutation of the anchor descriptors. In both cases it is then possible to build an index of ‘words’ using the consolidated technologies of the Image Retrieval field, and to use also Image Recognition technologies for searching and ranking the results of a query.

The Similarity Search will be based on the well-known “tf-idf” formula for weighting the “words” (visual features) within an image and on the cosine formula for ranking similar images. The tf-idf formula assigns a weight to each word in an image, and this weight is proportional to the number of occurrences of the word in the image (term frequency) and inversely proportional to the number of images in which this word appears (document frequency). In this way each image is represented by a vector of word weights and the similarity between an image in the database and the query image (to which the same processing has been applied at query time) will be measured by the cosine of the angle between the two vector representations.

The Image Recognizer (for Recognition Mode) is built for those inscriptions that have a set of different images associated with each one of them, so that each inscription can be associated with a ‘training set’. For each training set the recognizer will build a classifier that will be used to compute the probability that a given image belong to that set. The higher the number of images in the training set of a given inscription, the better the precision of the corresponding classifier will be. The Image Recognizer also builds a database containing the visual descriptors of the images contained in the training sets.

The Image Recognizer uses a recognition technology called “Single-label Distance-weighted kNN”. In simple words, the recognizer determines the k images that are closest to the query image (the Nearest Neighbours), then determines the training sets to which the nearest neighbours belong and finally, with mathematical algorithms based on weighted distances between the query image and the training sets, will determine the ‘single label’ (i.e. the textual tag associated with a training set) to be returned as result of the query, provided that the confidence level of the classification is above a certain threshold.

The functionality provided by the Image Recognizer System is essential for supporting also the Flagship Mobile Application, which will provide information to a user that sends to the EAGLE system the picture of an inscription taken with her mobile device.

4 The Flagship Mobile Application

Imagine the following scenario. A tourist is visiting Pisa, and in “Piazza dei Miracoli” (the cathedral square, with the famous Leaning Tower) is struck by an inscription on one of the walls of the cathedral. She wonders why it is there, what the meaning is, so she takes a picture with her smartphone and sends it to the EAGLE portal. In a few seconds she receives back the translation of the inscription, and a brief description, which explains that the marble stone is coming from the town of Ostia, near Rome, a summer resort for wealthy families in Roman times, with many villas and mansions. It becomes now clear

that in 1100 A.D. (the time when the Pisa cathedral was built), when Pisa was a powerful maritime republic, it was cheaper to take marble stones from ruined Roman villas in Ostia and carry them to Pisa by ships, rather than getting them from the real marble quarries in Carrara, which is less than 60 kilometres away from Pisa.

Other examples of this type can be found in Rome, where on the walls of many palaces of noble families appear stones with Latin inscriptions. The Coliseum was used as a ‘low cost’ marble quarry until mid-1700, when Pope Benedict XIV declared it a sacred place (because many early Christians had been martyred there) and prohibited further removal of marble and stones. In this case as well, a tourist visiting Rome and seeing a Latin inscription on the wall of a palace can send a picture to the EAGLE portal and discover that the inscription is coming from the Coliseum.

The two examples above illustrate the spirit of the “Flagship Mobile Application”. Using the camera of a mobile device, a user takes a picture of an inscription either from a monument she has access to, or from a printed or digital reproduction, and the picture is then sent to the EAGLE system, specifying which type of search is desired. Figure 3 shows a typical scenario for the mobile application

In the case of Similarity Search, the result is returned to the user as a list of images of different inscriptions, ranked in order of similarity with the picture taken. Each item in the list is a thumbnail and a very brief description of the inscription. The user can select a thumbnail from the list and can visualize either the full image, if the image is the only one related to the inscription, or a list of thumbnails of all the images related to the same inscription. Selecting the brief description, the user can visualise more detailed textual information about the inscription (title, type of inscription, type of object, ancient find place (region and city), present location, date, content provider).

In the case of Recognition Mode, if the image is recognized, the EAGLE system returns the 'representative image' of the recognised picture (the one indicated in the training set as best representing the inscription) together with detailed information, in the same format as in the case before. Figure 4 shows screenshots for this case. The layout and the functionality of the mobile application are capitalising on the experience of another EAGLE technology partner, EUREVA, which has already developed in the past mobile applications for commercial use.




Figure 3. Image recognition typical scenario

For Recognition Mode, the initial application will index a few thousand inscriptions, chosen so as to ensure maximum visibility and usefulness for the potential users, assumed to be mainly tourists and people who have a general interest in inscriptions. However, scholars as well (epigraphy specialists, historians, archaeologists, etc.) may also find the application interesting when visiting previously unknown places (cities, museums, archaeological sites), where new inscriptions can be found.



Figure 4. Snapshot of the Recognition Mode

The images sent to EAGLE and the information thus returned are saved in the application history on the mobile device and can be retrievable from the “History” tab appearing on the screen. In addition, a registered user, after logging in to the EAGLE system through the mobile device, is able to add a simple text note to the information displayed on the screen and save the screen and the note in her private area on the EAGLE server. The user can also take a picture, add a simple text note, and save both of them in this area.

5 Conclusions

EAGLE represents a community of research institutions documenting the most recent progress in the study of Classical Epigraphy. By aggregating digital content from unique and authoritative collections in its domain, EAGLE brings together a significant quantity of information about ancient writings on ancient artefacts to the vast user base of Europeana. This will in turn be useful for anyone who accesses Europeana for studying the classical world - amongst other things, these people will be able to explore primary-source materials never before available on such a large scale and from a single source. The visibility of Europe’s ancient history heritage will also be increased as a result and the richness of the collections will hopefully lead to improved acknowledgement of its importance as a resource for scientific research and as unique evidence of Europe’s past and identity.

As with many projects funded by the European Commission, there is a danger that the momentum and the impact of the initiative will decrease when the funding will end. We expect the mobile application to decrease this danger by virtue of its wide range of users, which comprises people that normally wouldn’t be interested in visiting the EAGLE or Europeana portals. These users, are, nonetheless, interested in learning a little bit more about an inscription that they may see on a wall, in an archaeological site, or in a book, given that this

information can be obtained simply taking a picture of the inscription. This will be our 'hook' to bring them onto the EAGLE portal.

For this reason, the mobile application is being developed with a multilingual interface, so as to facilitate its adoption throughout Europe. While initially it will only be available for Android, a release for iPhones and Windows phones is planned soon after. In addition, for its future sustainability, the application has been designed and engineered so as to enable the inclusion in its GUI of paid advertising, e.g. from restaurants and hotels in the neighbourhood of the inscription site.

Another opportunity for the adoption of the mobile application is given by museum curators, who could use it in the future as an easy and inexpensive way to get detailed information about their hidden assets: it is, in fact, well known by scholars in the field that many inscriptions are often held in museum galleries not open to the public, usually due to lack of resources or lack of adequate descriptions. This usage of the mobile application would also contribute towards increasing the overall value of the EAGLE collections, by providing possibly new and better images for its inscriptions.