Introduction

Many encyclopaedic and thematic museums around the globe exhibit artefacts discovered in various parts of the world; for example, the Egyptian museum in Torino is the largest institution exhibiting Egyptian artefacts outside Egypt. Moreover, objects belonging to the ancient Egyptian culture can be found in museums in UK, Germany, France, etc. Many of these objects now populating the exhibit rooms or are filling the storage shelves of museums were brought in decades or even hundreds of years ago, either following purchases, excavations in their countries of origin, donations and so forth. Thus, homogenous collections of artefacts - for example all objects found in a mortuary complex, or all items found within an ancient habitation context or a temple- are now generally distributed among several institutions, either public or private, for exhibiting or educational purposes. It may also happen that different institutions held fragments belonging to the same object, for example one museum holding the head, the other one the torso and the third one the limbs of a large statue. Such objects may never be restored to their original physical appearance, museums being very reluctant to give up pieces from their collection. Also entire collections may never be exhibited together and the provenance of artefacts may get lost and remain orphaned from the original contextualisation point of view.

The picture presented above has another facet as well. Archaeologists working on understanding past societies and their socio-cultural and economic structures look at objects of material culture as a main source of research. They study the shape, production techniques and manufacture process of these objects, within their original context of discovery. Thus, the physical integrity of the object, its relations with other objects and the completeness of the collection are essential and instrumental for assuring a high-quality archaeological research. Thus, many musealised artefacts are not suitable for a systematic and scientific archaeological investigation, given their fragmented condition, the unclear archaeological context and the difficulty in their physical study because of the geographic distribution across continents.

The EU funded project GRAVITATE aims to provide a digital technology based solution to the problems detailed above, addressing the **Re**-unification of items belonging to the same collection, the **Re**-association of orphan objects to their initial cultural assemblage, and the **Re**-assembly of fragmented artefacts. The development of technological solutions is driven by real-world archaeological, conservation, restoration and museological questions and the adopted methodology integrates archaeological research with computer graphics, computer vision, natural language processing and semantic technologies, in order to develop a product that will provide sustainable solutions to the three **R**s challenges presented above. The solution proposed by GRAVITATE is a research platform that allows scientists and CH professionals alike to investigate objects looking at their 3D geometry, surface properties, colouring texture and related textual descriptions and semantic information within a single digital environment, where they can conduct 3D shape analysis, features comparison, semantic and 3D annotation, similarity search and so forth, in order to digitally respond to the three Rs challenges.

The main case study of the project is an archaeological collection of about 250 fragments of votive terracotta statues from the ancient city Salamis, on the south-east coast of Cyprus. This collection was unearthed in the 19th century by a British excavation team and is dispersed in various collections in different countries with the majority of the pieces currently being stored in the United Kingdom in the British Museum, the Ashmolean Museum and the Fitzwilliam Museum. The number and the typologies of the complete statues present in the original context has never been established and the task to study the material has been considerably hindered by the dispersion of the material. In the course of the project more datasets were added to represent better the heterogeneity of the archaeological material or to test and develop specific parts of the platform, reaching more than 450 artefacts in total. The two most important additional datasets, which also consist of dispersed material, are pottery fragments from the archaeological collection of Naukratis, an ancient Greek settlement in Egypt, and votive statuettes from the Ayia Irini collection, an important sanctuary on the north coast of Cyprus.

Scientific approach

The scientific approach of the project revolves around a tight integration of semantics-oriented data and descriptions (e.g., archaeological descriptions or catalogue metadata) with information that can be extracted by 3D geometry processing techniques. This focus on both aspects of cultural heritage artefacts – qualitative textual descriptions and quantitative data and measures – is reflected throughout the development of the platform, which can be used once the data in the collections have been properly prepared and organized. The description of the platform will discuss the functionalities and goals of two phases, which in parts ran alongside each other: data collection and preparation with the creation of a repository and the development of the platform itself with its toolsets (Fig.1). The latter includes a semantic and a geometric search-engine, 3D visualisation and analysis tools, 3D re-assembly and metadata enrichment tools.

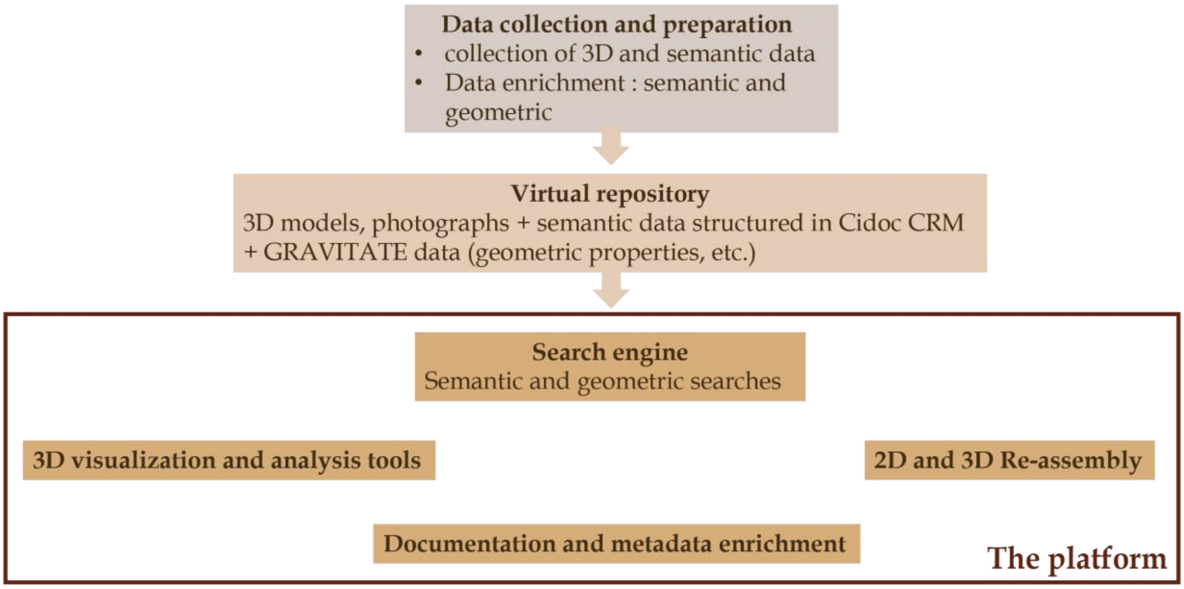

Figure 1 Structure of the GRAVITATE development process

Data collection and preparation

The data collection and preparation phase consisted in collecting museum records and archaeological descriptions of the GRAVITATE datasets on the semantic side, and in the 3D digitization of the artefacts on the geometric side. Since the metadata from the different museums did not follow the same scheme, it was necessary to bring all the information on the same level and into the same ontological system. All the available information was mapped into the British Museum scheme, which extends CIDOC Conceptual Reference Model (CRM) scheme (Doerr 2003), and then codified in RDF, a language that describes information in a form that can be processed by computers (McBride 2004).

The 3D digitization of the artefacts was carried out using photogrammetry and two different kinds of close range scanners: the NextEngine and the Aicon SmartScan. The colour information of the artefacts was encoded on the vertices of the 3D models by manual alignment of colour calibrated photographs using the software MeshLab (Cignoni et al. 2008). The 3D models were then run through a cleaning pipeline, named GRAVIfix, specifically designed for the GRAVITATE project to ensure that the models meet the requirements of the algorithms integrated in the platform (Mortara, Pizzi and Spagnuolo 2017).

Another important data preparation step was faceting (Fig.2), that is, the semi-automatic distinction of the skin of an artefact and its fracture (ElNaghy and Dorst 2017). This is particularly important for methods that address the three Rs problems.

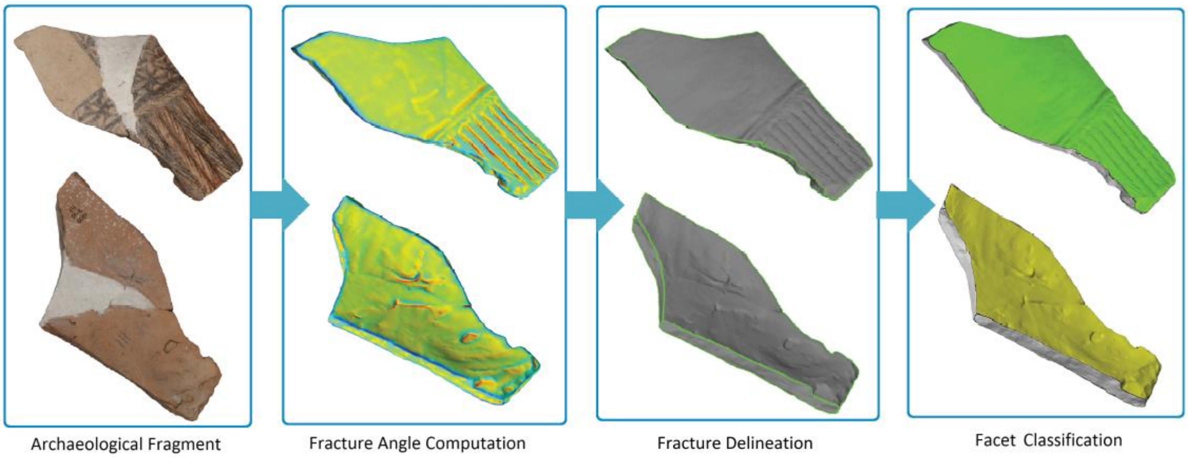

Figure 2 Automatic computation of the facets on an artefact. ElNaghy and Dorst 2017

Moreover, to reduce the computational effort for some of the algorithms employed in the next stage of data preparation and in the platform itself and for 3D web visualisation, several resolutions of the 3D models (1M, 100K and 50K) were generated. This required to develop a method to transfer the geometry of selected areas (for faceting and part-based annotation , see below) from one resolution to another in an automatic and accurate way (Scalas, Mortara and Spagnuolo 2017).

The next phase of data preparation consisted in semantic and geometric data enrichment which serves for different purposes, such as enhanced visualisation or search. Where and when this enriched data comes into play will be discussed with the single functionalities. An extensive description of the methods and results in the data preparation process is not within the scope of this paper and are elucidated in the respective publications.

(a) (b) (c)

Figure3 Automatic recognition of flower pattern based on curve approximation. (Torrente, Biasotti and Falcidieno 2018)

In terms of semantic data enrichment, free text descriptions of the single pieces in the collections were analysed with Natural Language Processing (NLP) techniques to extract new semantic information. The curated metadata were enriched by a novel algorithm based on Open Information Extraction (Christensen et al. 2011), extracting additional metadata from artefact physical descriptions. For the geometric enrichment, various properties of the 3D models were calculated, which include the prevalent thickness based on the shape diameter function (SDF) (Shapira, Shamir and Cohen-Or 2008) or different ways to quantify the curvature, for example through the shape index (Koenderink and van Doorn 1992) or the mean curvature. Some of these geometric properties were also processed further to define collection-specific feature extraction modules, such as eyes or relief decorations (Biasotti et al. 2017; Torrente, Biasotti and Falcidieno 2018).

The collected and newly generated information is being stored in a virtual repository based on RDF triples, which contains: 3D models and photographs, semantic metadata in the CIDOC CRM scheme and the new semantic and geometric information from the data enrichment phase.

Development of GRAVITATE toolsets

The search engine

All this data needs to be made accessible and searchable in order to be able to retrieve information of interest and make connections between different objects in the database. Thus, the GRAVITATE platform was equipped with a sophisticated search engine that allows to perform both semantic and geometric queries, which is particularly novel and interesting.

The starting point of such a query is an artefact of interest and the aim is to find other pieces that share with the query one or a combination of similar characteristics chosen by the user. On the semantic side, the similarity search is based on graph matching. We used a RDF2VEC graph embedding technique (Ristoski and Paulheim 2016) to perform query by example, using the query artefact metadata as the basis for similarity. A user is able to look for objects that have similar material, type, decoration, production time and place, symbol and physical features.

The geometric search is based on a set of descriptors computed for each artefact and distances that evaluate the similarity between two artefacts as the distance between their descriptors. Given the query artefacts, the geometric search permits to look for objects that have a similar overall shape, size, colour, or roughness as well as a similar 2D or 3D decoration (for the retrieval of 3D patterns (see Biasotti et al. 2018). As most of the fragments exhibit a different appearance on the external and internal facets (e.g. colour, roughness) it is possible to improve search results by restricting the search to the pertinent side of a fragment (see Fig.4).

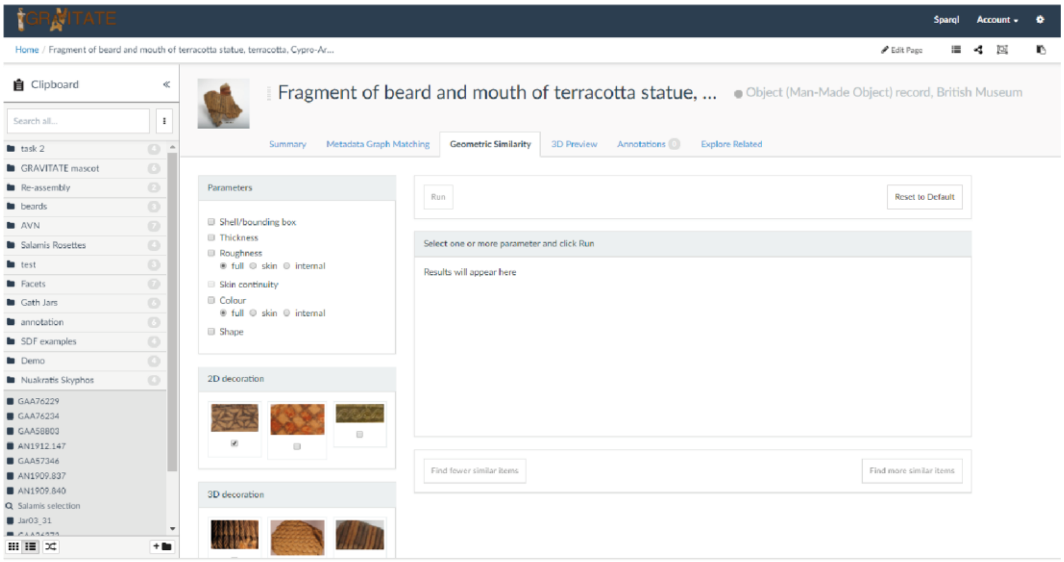

Figure 2 Geometric Similarity search in the web-client.

3D visualisation and analysis tools

Most of the geometric information computed during the geometric enrichment phase are used to build the descriptors that are needed to run the geometric search. However, their value is not limited to the search engine: the geometric data are indeed useful also for direct inspection, visual comparison and analysis, using 3D models as simulacra of the real artefacts. Thus, 3D visualisation and analysis tools for the comparison of 3D models and the visualisation of geometric properties were developed in the platform (see Fig.8).

3D-reassembly

In the case of a discovery of potential matches between fragments, identified for example with the combined search engine, a semi-automatic 3D re-assembly tool was developed building on the experience of the PRESIOUS (Predictive digitization, restoration and degradation assessment of cultural heritage objects) project (Savelonas et al. 2017). In our approach the fracture facet of a fragment is detected and principal directions and curvatures are computed. Then, the algorithm looks for compatibility of curvatures between a set of previously selected fragments, attempting re-assembly and suggesting aligned pairs to the user. A fine alignment of this initial rough positioning of pairs of fragments is being developed based only on the geometry of the fracture surfaces. The problem of potentially missing geometry due to abrasion or other type of damages to the fractures is being addressed through a simplification of the fracture surfaces using mathematical morphology operations (ElNaghy and Dorst 2018).

Metadata enrichment

In order to make the knowledge and information content of the system dynamic, the platform contains also different possibilities to enrich the metadata in various ways. Observations that concern directly the geometry of artefacts can be annotated in 3D through the selection of the area of interest and tagging it with an appropriate term from a specifically designed Cultural Heritage Artefact Partonomy (CHAP) (see Fig.5). Such vocabulary has been defined starting from the archaeological description of the artefacts of the Salamis, Naukratis and Ayia Irini datasets studied in the project. Moreover, the annotated features can be suitably measured, enriching the semantic annotation with quantitative values to be used in archaeological research. The geometry of this area receives thus a direct semantic link and meaning, codified in RDF, as the metadata of the respective artefact (Catalano, Repetto and Spagnuolo 2017). There is also the possibility to enrich the metadata by adding information to the traditional CIDOC CRM metadata scheme for observations that are not directly linked to the geometry (e.g. provenance, dating, material, etc.). Every intervention is documented and has to be motivated making the process transparent and traceable, enabling a dialogue between different users of the platform.

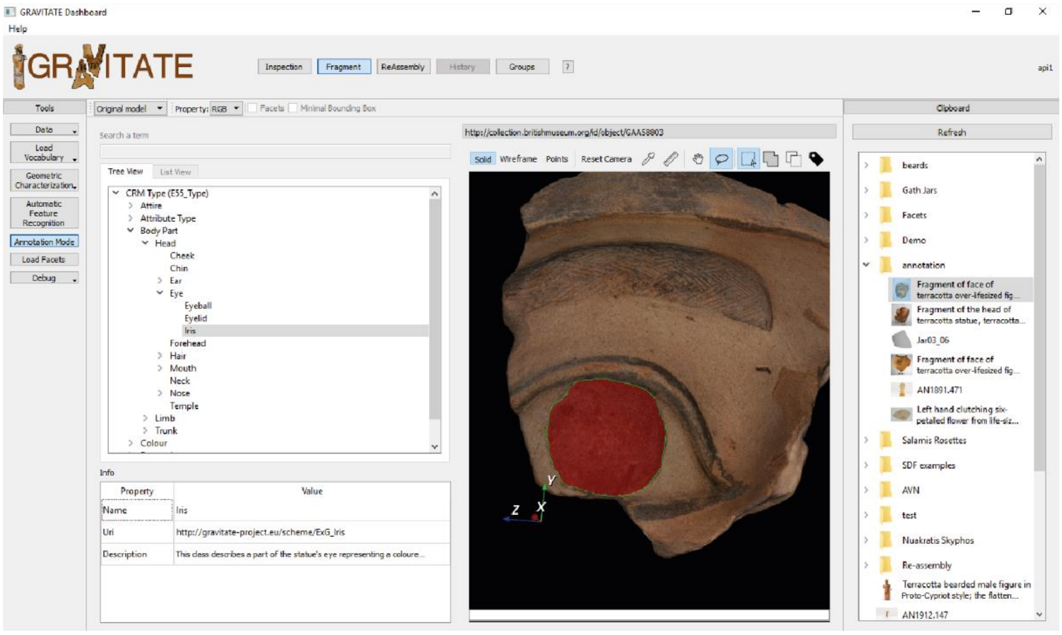

Figure 3 3D annotation of the iris of a fragment in the desktop-client.

The platform interface

The interface of the platform is divided in two clients: a web-client and a desktop-client. This division is the consequence of the intrinsic limitations of web-browsers to deal with high resolution 3D models in a fast and efficient way and the need to use the web for the GRAVITATE semantic search-engine. The latter is in fact largely based on ResearchSpace, a collaborative semantic web environment developed by the British Museum that works entirely from the web (Tanase and Oldman 2018). It was adopted by GRAVITATE and modified to adapt it to its needs. The main purpose of the web-client is to explore and browse the database for objects of interest and to provide useful information on the pieces. The desktop-client on the other hand serves mainly for the manipulation and visualisation of high resolution 3D models (Catalano, Repetto and Spagnuolo 2017). The two clients are synchronised through a clipboard, in which objects of interests can be stored and grouped by the user.

The Web-Client:

More specifically the web-client contains the search engine that permits access to the repository. It is divided in a work-space and the clipboard where objects of interest can be saved and grouped. In the work-space queries can be composed in a versatile highly adaptive manner using concepts derived from CIDOC CRM and can be further refined through filters (Fig.6). The user can select two key concepts (Actor, Place, Event, Thing, Time, Concept and Class) and define the relationship between them (e.g. search for a Thing *from* a specific Place).

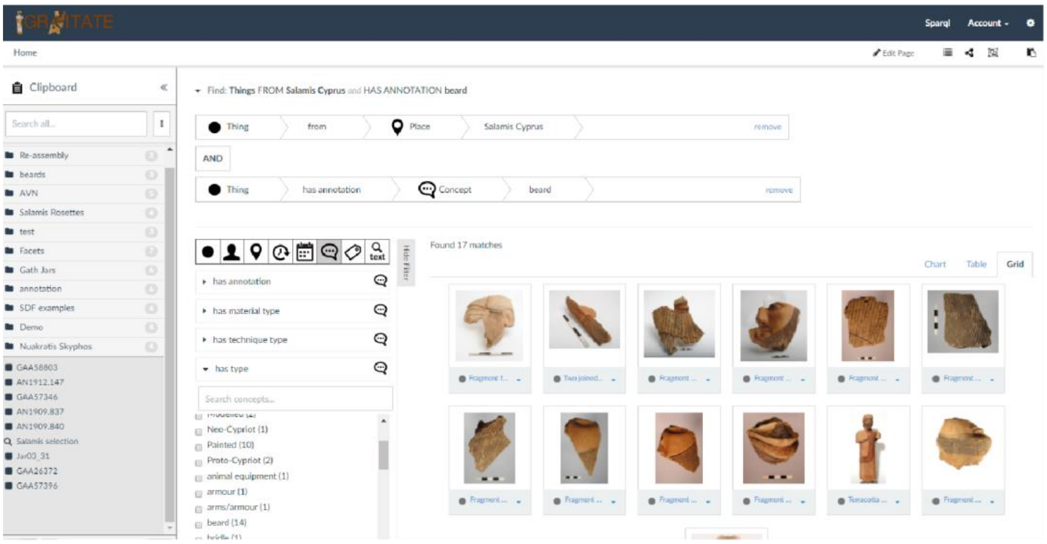

Figure 4 Query in the web-client looking for objects from Salamis that are or have a beard.

Search results can be browsed and single objects explored in more detail by looking at their metadata, photographs and a preview of its 3D model (Fig.7). In the metadata display the previously mentioned metadata enrichment that is not based on geometry takes place. A user can add more information to the different CIDOC CRM metadata categories (e.g. material, type) and express an opinion on the already available information. Once one object is selected, it is also possible to look for similar objects by conducting the semantic or geometric similarity searches.

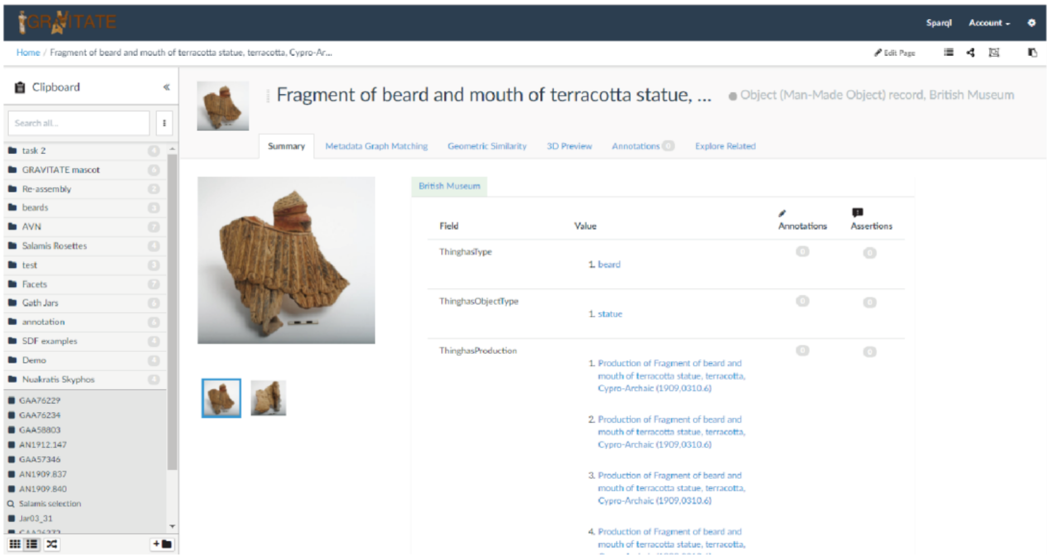

Figure 5 Looking at the metadata of an artefact in more detail in the web-client.

The Desktop-Client:

If a user wishes to explore selected artefacts in more depth, (s)he can move to the desktop-client, where the clipboard contains the same objects that were selected and grouped together in the web-client. The desktop-client is also divided in the clipboard and a workspace as well as a toolbar and a viewbar. The latter determines the kind of activity a user can perform and what tools are available and consists in: Inspection View, Fragment View, Re-Assembly View, History View and Groups View.

The Inspection and Fragment View contain a standard visualisation and manipulation tool for rotation, zoom in and out, measurement, changing the light position and change the visualisation mode between points, wireframe and solid. The metadata and paradata of each artefact are displayed as an expandable tree below the interactive 3D canvas. In the Inspection view, geometric properties, facets and annotations of the displayed artefacts can be loaded as well and being inspected in parallel for two or more artefacts (Fig. 8).

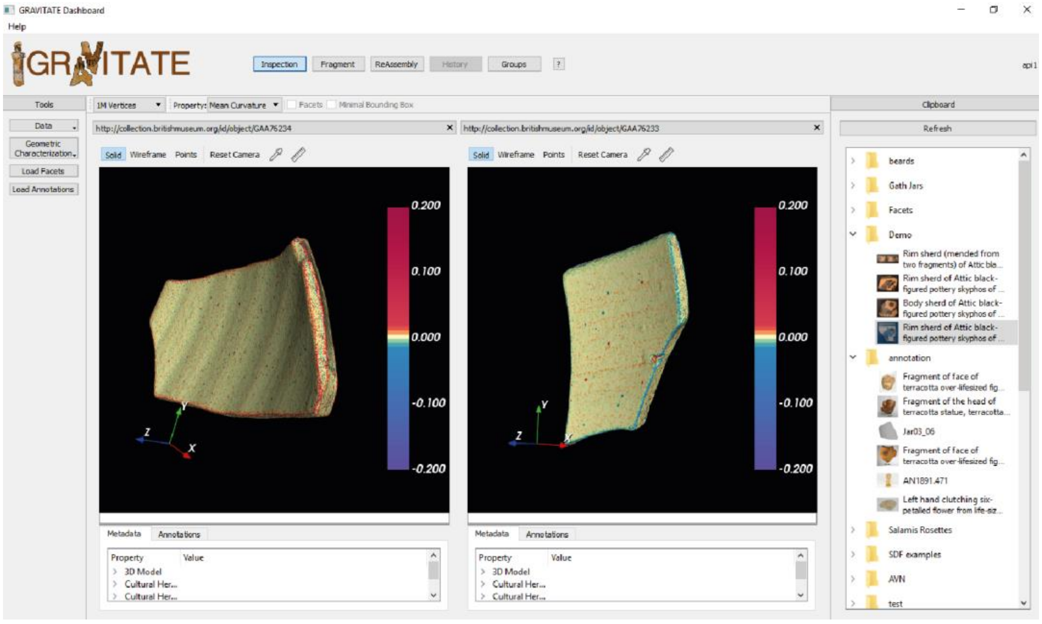

Figure 6 Visualisation and comparison of the mean curvature on two fragments in the desktop-client.

The Fragment View has the annotation tool, which allows a user to select a specific area on the model and create a new 3D part-based annotation with the developed CHAP vocabulary. If the user has discovered artefacts that might match physically and can be reassembled, (s)he can move to the Re-assembly View. The potentially matching objects are put in a workbench by the user and the reassembly algorithms will attempt to find matching pairs, which the user can approve or discard. All the steps taken and observations made can be documented in the History view. Finally, the Groups View is used to manage the groups of artefacts made in the clipboard.

The 3 Rs and the platform – workflows and use of the system

The Re-unification of dispersed collections in a digital repository, made accessible through a platform that allows to perform also Re-associations and Re-assemblies, enables to ask archaeological questions that previously could only be addressed with difficulty. The research questions and objectives we identified relate to the exploration and classification of a dataset, the study of style and of production techniques within or across collections and the understanding of archaeological context. Furthermore, the system can be used for a systematic enrichment of the metadata.

The Salamis collection provided a good example where the platform helped to understand the material (what type of statues are present in the collection) and attempt a classification. In this process the web-client was used for a semantic exploration of the dataset in order to create meaningful groups of fragments that can be associated. The semantic and geometric similarity searches play an important role here, allowing a user to find similar artefacts faster. On some occasions it might even be possible to attempt a reassembly of a few fragments to understand better the shape and dimension of the statue the pieces belonged to. In other cases, re-assembly can help understanding the archaeological context: by reassembling a group of fragments found in the same context, even if their typology is known, we may understand how many complete objects there were originally. Knowing how many objects of a certain type were in a context facilitates its interpretation.

The study of style uses the re-association capabilities of the platform. In the first step this happens in the web-client through the different possibilities to conduct queries and then, in a second step, in the desktop-client, taking measurements and observing subtle differences between elements and features of interest that point to stylistic influences and developments. In this context, it may be particular useful to create also 3D annotations to enrich the metadata together with the exact measurements of the interesting areas, and make future queries more effective and accurate.

Also the investigation of production techniques may be accomplished through the 3D visualisation and analysis tools in the desktop-client. The mean curvature and shape index calculated on the high resolution models as well as the enhanced visualisation of features and the possibility to make measurements allow to study signs of production on artefacts, such as manufacturing lines in the internal face of a sherd (see Fig.8).

The system was designed to be open only to people authorized by an administrator and requires a login to access and use it. This is a deliberate choice to address right issues revolving around data, since the system allows both the access and manipulation of it. In terms of use, there is no predefined workflow of the platform, allowing a user to begin, stop and continue working at any time and place in the system. To illustrate the usage of the platform, the typical workflow begins in the web-client, where the user looks for artefacts relevant for a particular archaeological question. This reasoning process can be accompanied by a deeper analysis and annotation of the high-resolution 3D models in the desktop-client. It may or may not include the use of the Re-assembly tool, and may be followed by the return to the web-client either to look for other artefacts or to conduct a different semantic or geometric search.

Conclusions

The aim of the project was to propose solutions to three cultural heritage challenges: Re-unification, Re-association and Re-assembly. The problem of Re-unification was addressed by bringing dispersed material together in one digital repository. The different tools allowing to explore the repository as well as the tools dedicated to the analysis of the 3D models enable to create and discover meaningful connections between artefacts in the repository and thus answer to the issue of Re-association. Finally, the problem of Re-assembly was addressed through the development of a semi-automatic tool for matching fragments.

There are still many potential developments for the platform. Some elements, such as the 3D reassembly or the semantic graph matching, can still be improved and deserve further attention. Other research lines which were pursued in the project did not reach a stage in which it was possible to integrate them in the platform, but show great potential and merit future work. These include tools for 2D re-assembly (Paikin and Tal 2015), or a 3D re-assembly based on a template for the refitting of fragments that belong to the same object but do not present common fracture surfaces.

A platform like GRAVITATE is capable of hosting the entire workflow necessary when studying a cultural heritage collection, from browsing a museum catalogue in search of pieces of interest to the inspection and study of the objects themselves. The environment has been designed to maintain the connection between the semantic and the geometric data and to be transparent, open and dynamic, enabling the exchange of ideas and interpretations through the continuous enrichment of the underlying metadata and the accurate measurements of geometric and colorimetric properties of the 3D high-res models. The biggest challenge for the adoption of such a system lies in the ingestion phase, that is the creation and preparation of the necessary data, which requires the mapping of semantic data into the CIDOC CRM scheme and the availability of the dataset in high quality 3D models. In a wider perspective, the issue of the federation of the repositories of different museums should be tackled to guarantee the share of knowledge among the institutions, while preserving the specific organisation of the individual digital archives. Indeed, the Cultural Heritage sector has already made considerable efforts towards a common way of structuring and storing of metadata and a mass digitization of cultural heritage objects. It is possible to observe the latter in initiatives such as CultLab3D, led by the Fraunhofer Institute for Computer Graphics Research IGD, which developed an automated 3D capture pipeline (Santos et al. 2017). In terms of communal data structure and storage, we can mention, among others, the ARIADNE project, which aims at providing e-infrastructure to collect, share and access heterogeneous data created across various institutions and countries (Meghini et al. 2017; Niccolucci and Richards 2019) and was recently extended within the ARIADNEplus initiative.

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement 665155. We thank and acknowledge all members of the Consortium of the GRAVITATE project.

References

Biasotti, S, Moscoso Thompson, E, Aono, M, Ben Hamza, A, Bustos, B, Dong, S, Du, B, Fehri, A, Li, H, Limberger, F A, Masoumi, M, Rezaei, M, Sipiran, I, Sun, L, Tatsuma, A, Velasco-Forero, S, Wilson, R C, Wu, Y, Zhang, Y, Zhao, T, Formasa, F and Giachetti, A 2017 Shrec'17 Track: Retrieval of surfaces with similar relief patterns. In: *10th Eurographics Workshop on 3D Object Retrieval*. April 2017 Lyon, France. p.

Biasotti, S, Thompson, E M, Bathe, L, Berretti, S, Giachetti, A, Lejemble, T, Mellado, N, Moustakas, K, Manolas, I, Dimou, D, Tortorici, C, Velasco-Forero, S, Werghi, N, Polig, M, Sorrentino, G and

Hermon, S 2018 SHREC'18 track: Recognition of geometric patterns over 3D models. In: *Eurographics Workshop on 3D Object Retrieval*. 2018 p. 9.

Catalano, C E, Repetto, A and Spagnuolo, M 2017 A Dashboard for the Analysis of Tangible Heritage Artefacts: a Case Study in Archaeology. In: *15th EUROGRAPHICS Workshop on Graphics and Cultural Heritage*. 2017 The Eurographics Association. p. DOI: https://diglib.eg.org:443/handle/10.2312/gch20171307.

Christensen, J, Mausam, Soderland, S and Etzioni, O 2011 An analysis of open information extraction based on semantic role labeling. In: *K-CAP 2011*. 2011 ACM Press. p. 113. DOI: https://doi.org/10.1145/1999676.1999697.

Cignoni, P, Callieri, M, Corsini, M, Dellepiane, M, Ganovelli, F and Ranzuglia, G 2008 *MeshLab: an Open-Source Mesh Processing Tool*. In: 2008 pp. 129–136.

Doerr, M 2003 The CIDOC Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata, *AI Magazine*, 24(3): 75.

ElNaghy, H and Dorst, L 2017 *Geometry Based Faceting of 3D Digitized Archaeological Fragments*. In: October 2017 IEEE. pp. 2934–2942. DOI: https://doi.org/10.1109/ICCVW.2017.346.

ElNaghy, H and Dorst, L 2018 Using Mathematical Morphology to Simplify Archaeological Fracture Surfaces. In: *Symposium on Geometry Processing 2018- Posters*. 2018 The Eurographics Association. p. DOI: https://doi.org/10.2312/sgp.20181179.

Koenderink, J J and van Doorn, A J 1992 Surface shape and curvature scales, *Image and Vision Computing*, 10(8): 557–564. DOI: https://doi.org/10.1016/0262-8856(92)90076-F.

McBride, B 2004 The Resource Description Framework (RDF) and its Vocabulary Description Language RDFS. In: *Handbook on Ontologies*. International Handbooks on Information

Systems. Springer, Berlin, Heidelberg. pp. 51–65. DOI: https://doi.org/10.1007/978-3-540-24750-0_3.

Meghini, C, Niccolucci, F, Felicetti, A, Ronzino, P, Nurra, F, Papatheodorou, C, Gavrilis, D, Theodoridou, M, Doerr, M, Tudhope, D, Binding, C, Scopigno, R, Vlachidis, A, Richards, J, Wright, H, Geser, G, Cuy, S, Fihn, J, Fanini, B and Hollander, H 2017 ARIADNE: A Research Infrastructure for Archaeology, *Journal on Computing and Cultural Heritage*, 10(3): 1–27. DOI: https://doi.org/10.1145/3064527.

Mortara, M, Pizzi, C and Spagnuolo, M 2017 Sreamlining the Preparation of Scanned 3D Artifacts to Support Digital Analysis and Processing: the GRAVITATE Case Study. In: *15th EUROGRAPHICS Workshop on Graphics and Cultural Heritage*. 2017 p. 5.

Niccolucci, F and Richards, J 2019 *The ARIADNE Impact*. Hungary: ARCHAEOLINGUA FOUNDATION. DOI: https://doi.org/10.5281/zenodo.3476712.

Paikin, G and Tal, A 2015 Solving multiple square jigsaw puzzles with missing pieces. In: *The IEEE Conference on Computer Vision and Pattern Recognition*. June 2015 IEEE. pp. 4832–4839. DOI: https://doi.org/10.1109/CVPR.2015.7299116.

Ristoski, P and Paulheim, H 2016 RDF2Vec: RDF Graph Embeddings for Data Mining. In: Groth, P., Simperl, E., Gray, A., Sabou, M., Krötzsch, M., Lecue, F., Flöck, F., and Gil, Y. (eds.) *The Semantic Web – ISWC 2016*. Cham: Springer International Publishing. pp. 498–514. DOI: https://doi.org/10.1007/978-3-319-46523-4_30.

Santos, P, Ritz, M, Fuhrmann, C and Fellner, D 2017 3D mass digitization: a milestone for archeological documentation, *Virtual Archaeology Review*, 8(16): 1. DOI: https://doi.org/10.4995/var.2017.6321.

Savelonas, M A, Andreadis, A, Papaioannou, G and Mavridis, P 2017 Exploiting Unbroken Surface Congruity for the Acceleration of Fragment Reassembly. In: *Eurographics Workshop on Graphics and Cultural Heritage*. 2017 The Eurographics Association. p. DOI: https://doi.org/10.2312/gch.20171305.

Scalas, A, Mortara, M and Spagnuolo, M 2017 3D Annotation Transfer. In: *15th EUROGRAPHICS Workshop on Graphics and Cultural Heritage*. 2017 The Eurographics Association. p. DOI: https://doi.org/10.2312/gch.20171295.

Shapira, L, Shamir, A and Cohen-Or, D 2008 Consistent mesh partitioning and skeletonisation using the shape diameter function, *The Visual Computer*, 24(4): 249–259. DOI: https://doi.org/10.1007/s00371-007-0197-5.

Tanase, D and Oldman, D 2018 Reshaping the Knowledge Graph by connecting Researchers, Data and Practices in ResearchSpace. In: *Proceedings of the 17th International Semantic Web Conference*. 2018 p.

Torrente, M-L, Biasotti, S and Falcidieno, B 2018 Recognition of feature curves on 3D shapes using an algebraic approach to Hough transforms, *Pattern Recognition*, 73: 111–130. DOI: https://doi.org/10.1016/j.patcog.2017.08.008.