EUROGRAPHICS Workshop on Graphics and Cultural Heritage (2016)
C. E. Catalano and L. De Luca (Editors)

GRAVITATE: Geometric and Semantic Matching for Cultural Heritage Artefacts

S. C. Phillips[1], P. W. Walland[1], S. Modafferi[1], Leo Dorst[2], Michela Spagnuolo[3], C. E. Catalano[3], Dominic Oldman[4], Ayellet Tal[5], Ilan Shimshoni[6] & Sorin Hermon[7]

[1]University of Southampton IT Innovation Centre, UK
[2]Universiteit van Amsterdam, Netherlands
[3]Istituto di Matematica Applicata e Tecnologie Informatiche "E. Magenes", CNR, Genova, Italy
[4]British Museum, London, UK
[5]Technion Israel Institute of Technology
[6]University of Haifa, Israel
[7]STARC - The Cyprus Institute, Cyprus

Abstract

The GRAVITATE project is developing techniques that bring together geometric and semantic data analysis to provide a new and more effective method of re-associating, reassembling or reunifying cultural objects that have been broken or dispersed over time. The project is driven by the needs of archaeological institutes, and the techniques are exemplified by their application to a collection of several hundred 3D-scanned fragments of large-scale terracotta statues from Salamis, Cyprus. The integration of geometrical feature extraction and matching with semantic annotation and matching into a single decision support platform will lead to more accurate reconstructions of artefacts and greater insights into history. In this paper we describe the project and its objectives, then we describe the progress made to date towards achieving those objectives: describing the datasets, requirements and analysing the state of the art. We follow this with an overview of the architecture of the integrated decision support platform and the first realisation of the user dashboard. The paper concludes with a description of the continuing work being undertaken to deliver a workable system to cultural heritage curators and researchers.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems; H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval—Clustering

1. Introduction and objectives

The vast majority of archaeological objects are discovered in a fragmentary state, and the poor state of preservation of these pieces further hampers the extent of the archaeological research possible. Moreover, pieces of historical importance and interest may be dispersed across different collections making their reassembly difficult and even impossible due to missing or eroded parts. Much effort and time is expended in trying to reassemble these fragments as accurately and completely as possible, whether for public display or historical research. The artefacts might be the relatively recently excavated broken terracotta votive statues from Salamis being investigated by GRAVITATE, or long-lost antiquities such as the missing parts of the basalt slab of Nectenebo I, held by the British Museum, but in either event, digital procedures may greatly improve the speed and effectiveness of the work done by researchers and curators to discover, identify and reassemble the pieces.

Recent technology developments in the area of 3D scanning and digital shape analysis and matching have shown that it is possible to scan recently fractured surfaces and reassemble them using visualisation tools and geometric matching. However, most of the applications which are encountered in reality in the reconstruction of cultural heritage objects call for non-exact matching techniques, since surfaces are generally heavily abraded or damaged, and exact or near-exact matches are not possible. The three year GRAVITATE project, started in June 2015, brings together diverse expertise from IT Innovation, UVA, CNR-IMATI, British Museum, Technion, University of Haifa and Cyprus Institute along with the assistance of the Cyprus Department of Antiquities, the Fitzwilliam and Ashmolean museums. It aims to address this issue through a combination of geometrical and semantic approaches to discover similarity between fragments and therefore associations and even geometric matches. In addition to the development of improved matching algorithms more suited to this specific task, the project partners are working together to create a pipeline specifically targeted at the similarity assessments of 3D artefacts found in the real world, concurrently evaluating heterogeneous properties such as geometric aspects (e.g. curvature, size, roundness or mass distribu-

DOI: 10.2312/gch.20161407

tion), photometric aspects (e.g. texture, colour distribution, surface patterning or reflectance) and semantic annotation, enriching the curated knowledge found in the documentation.

By combining the semantic and geometric data we will enable researchers to discover matches and similarities that may not have been obvious in the past and go beyond the manual or semi-manual matching techniques currently used. This will lead to further research challenges around the use of these matching techniques to support the digital reunification of pieces of artefacts that may have been dispersed across different collections (and which often cannot be moved) and the re-association of artefacts having common origins which will provide new understanding and insight into cultural and social history. Digital reassembly of artefacts is a step forward from (semi-permanent) physical reassembly or skilled 2D illustration as it allows for different hypotheses to be freely explored using fragments potentially located in different museums.

2. Baseline for the technical work

2.1. Semantic and geometric data sets

The primary use case of the project is a collection of fragmented terracotta funerary statues from Salamis in Cyprus which are dispersed across the British, Ashmolean, Fitzwilliam and Cyprus museums. The collection forms a good starting point for two key reasons: from the geometrical side because of a pre-existing agreement between the museums to scan the artefacts in 3D, and from the semantic side because of the British Museum's ResearchSpace project [Old] which leads the way in the semantic documentation of artefacts for research purposes. Access to these evolving datasets offers GRAVITATE a unique opportunity to demonstrate how the combination of geometric and semantic techniques, extending the state of the art, can revolutionize digital search and matching within our cultural heritage institutions.

In total there are 221 3D models and over 2000 scan images of Salamis artefacts available to the project [NOR*16]. The majority of scans are in PLY format and of tens of MB in size but some are up to 0.5 GB. Scans of other artefacts are also available to the project and further multi-spectral scans will be performed once the particular quality attributes (resolution etc.) required for the matching and mating algorithms are understood.

The British Museum exposes its catalogue records as Linked Open Data using the CIDOC CRM family of ontologies [Doe03]. The dataset, describing 2.5 million objects, has been extended in GRAVITATE to incorporate the catalogue records of the three other museums holding artefacts from the Salamis case study, mapping each institution's catalogue schema into the CRM.

2.2. User requirements

The project has identified four distinct classes of users for the planned re-association, reunification and reassembly features: researchers, whose generic aim is to solve some scientific questions; conservators, who aim to restore broken artefacts; curators who manage, document and display archaeological collections; and illustrators who work primarily with curators and visually interpret pieces for the public and for researchers.

Through a questionnaire and through face to face interviews with experts, the requirements of these (often overlapping) users have been elicited. The project aims to satisfy the requirements by building a multi-modal platform incorporating a graphical user interface to explore and manipulate the data. Here we focus on key requirements that are specific to the geometric aspects or which can be solved by the combination of geometric and semantic data.

Re-Assembling: such as matching of two fragments based on 2D or 3D pattern alignment; matching of two fragments based on the complementarity of the fracture facets; from a set of fragments, proposing reconstructions; and placing one or more fragments automatically on the most plausible area of a virtual mannequin.

Annotation/Analysis: such as automatic detection of fragment's facets (external, internal, fracture); automatic detection of anatomical or ornamental features; automatic detection of coloured paintings/drawings; and manual selection and annotation of a region of interest on the fragment.

Selection: such as selecting a set of fragments using a search with semantic criteria and by similarity to one or more other fragments, taking into account both the appearance and semantic data.

Visualization: the visualization of 3D data as well as textual or conceptual data. Such as graphical interaction and displaying the meta-data and part-based annotations related to the fragment;

Publishing: documentation and sharing of new knowledge. Such as adding to the semantic data for a fragment and creating a new record associating one or more fragments, labelled as being a belief of the user.

2.3. State of the art

A full state of the art review has been conducted covering both the geometric and semantic aspects of the work in the context of cultural heritage [BCC*16]. For this paper we limit ourselves to a brief discussion of the main prior work in the geometric field.

Using 3D scans of potsherds, an ancient vase was digitally reassembled with the positions being chosen by hand [HPI*]. The complete vase was then virtually reconstructed by semi-automatic extrapolation.

Shape descriptors can be categorised as local [SPS15], region-based [IT11] and global [GG15] descriptors. Multiple descriptors can then be combined into a feature vector and pairwise similarity can be computed [BCFS15].

The shape descriptors (along with semantics) can be used in a similarity search to help select a small set of fragments which may fit together, extending the approach of the ICON project [BCS*11]. The complexities such as physical degradation inherent in working with real-world fragments [WC08] must be taken into account. For instance, the work of Kleber et al. [KDSK10] combines features extracted from geometry, marble texture, thickness of marble fragments and location at the excavation, to reassemble the eroded Ephesos marble plates.

Geometric mating is described by Huang et al. [HFG*06] and has recently been investigated in the PRESIOUS project [PGS*15]. In addition, in the mathematics of ideal shapes in exact contact, called "Mathematical Morphology" [DvdB00], there exists a morphological scale space in which objects can be described at various resolutions in a manner that hierarchically preserves their potential contact in a well-understood way which may well be applicable to the project's case of abraded contact surfaces.

It would not be practical to attempt a mating of every pair of pieces (or even pairs related semantically) and so we will draw on work done in the field of 2D jigsaw puzzles to reduce the combinatorial space to just likely matches. Here, the most relevant work is that of Paikin and Tal [PT15] which addresses the problem of multiple incomplete puzzles mixed together based on their colorimetric properties.

3. Technical progress

3.1. Architecture

To meet the requirements, the GRAVITATE platform needs to provide a multi-user system integrating a diverse set of tools for 3D geometric and surface analysis, 3D matching and a rich query and update mechanism incorporating coordinated access to three distinct data types:

1. semantic: primarily describing the artefacts in curatorial terms;
2. numeric: computed shape descriptors and shape properties;
3. 3D: the models and parts of the models.

The platform architecture is shown in Figure 1. The main client interface will be through a rich web application, with manipulation of high resolution models being delegated to a native desktop application. The service interface will be protected by an authentication and authorisation system. Federation of data between GRAVITATE services (external clients) is also envisaged through the same RESTful service interface. The platform will integrate with the existing ResearchSpace software which itself builds on the Metaphacts platform and the BlazeGraph semantic database.

Simple data queries and updates are dealt with directly by the metadata manager which encapsulates the logic required to interface to both the semantic database (such as Blazegraph) and the numeric database (such as PostgreSQL). Queries touching on both data types are executed on the semantic database using federated SPARQL which in turn queries the numeric database via a SPARQL custom service.

More complex operations requiring computation are handled by the workflow manager. For instance, when a new 3D model is ingested into the system, surface faceting (i.e. finding the front, back and fracture regions of the fragment) and many shape descriptor calculations must be performed and the results stored in the appropriate databases. Operations such as faceting are potentially slow and may also require review by the user. The workflow manager manages the state of these operations and of the extended user interactions.

The job service provides a unified interface to the multitude of algorithms required and is able to schedule jobs on a computational cluster. The event bus collects events describing the operations being performed by the user. This is important because if a researcher uses the platform and, for instance, believes she has discovered two fragments which appear to fit together then that new knowledge must be stored in the semantic database alongside the supporting

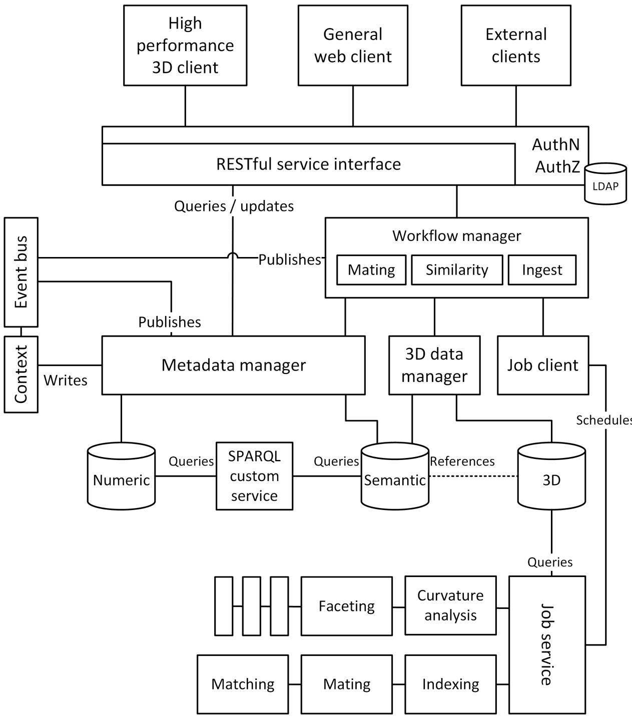

Figure 1: *GRAVITATE platform architecture.*

evidence for the researcher's belief. This context or provenance information describing what queries were performed and algorithms executed then allows other cultural heritage researchers to evaluate the validity of the assertion.

3.2. User interface

An extensive interactive dashboard mock-up has been created to help potential users understand the project's vision and to elicit feedback and further requirements. The dashboard will support the user in the interaction with and manipulation of the data, centred around fragment similarity. The interface has been divided into different functional areas, linked through a clipboard (and associated provenance context). The views identified are the inspection view for viewing the 3D objects in the repository along with their metadata; the fragment view to analyse geometrically a single object and perform part-based annotation; the reassembly view for reassembling a fragmented object; and the exploration view for exploring data using similarity measures based on the semantics and appearance of fragments. Figure 2 shows the fragment view with several regions of a fragment identified semantically as parts of the body and the left-ear labelled.

The main user interface for all classes of user will be a rich web application which has many advantages including the reuse of existing ResearchSpace components. The manipulation of 3D models in a web browser is now easily achieved through technologies such as x3dom and xml3d and, where practical, 3D models will be integrated into the web interface. However, we have found that (with many of the models containing several million vertices) the current

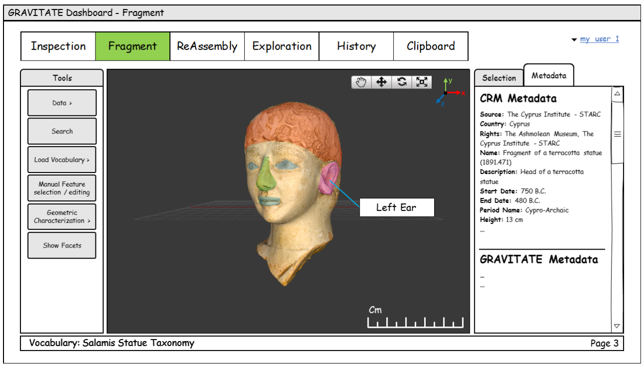

Figure 2: *GRAVITATE web interface mock-up.*

web browsers cannot be used in all cases and so a linked native desktop client for model visualisation and manipulation will also be provided.

4. Future work

The project is now working on implementing the combination of semantic and geometric technologies required to make the platform a reality whilst refining the requirements through a continuing dialogue with cultural heritage professionals.

The development of an advanced similarity search is key: providing a small set of fragments for reassembly and providing the means for digital reunification and re-association. An interface allowing control of the search in multiple similarity dimensions will be developed and as input to this many geometric characteristics will be computed from the 3D scans and the semantic data will be enriched using natural language processing of the grey literature related to the artefacts.

A new faceting algorithm will underpin many other shape descriptors and is key to the reassembly pipeline in which a set of related fragments is first found using similarity search, the local fracture surfaces will then be characterised and likely matches found in an extension of the 2D puzzle approach, and pairwise mating of pieces will be attempted building on the mathematical morphology approach.

A semantic partonomy integrated with the CIDOC CRM data is being developed to describe parts of statues (e.g. "nose" is part of a "face") and will later be broadened out to describe fragments of other artefacts. Fragments semantically described so will inform geometric searches (e.g. searching for the eye region in 3D models where the semantics indicates an eye exists) and vice versa the semantics will be enriched by geometric findings.

5. Conclusion

This paper has outlined the aims and objectives of the GRAVITATE project, described the progress to date and briefly set out the project's ambitious programme of further work. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 665155.

References

[BCC*16] Biasotti S., Cerri A., Catalano C. E., Falcidieno B., Torrente M.-L., Middleton S. E., Dorst L., Shimshoni I., Tal A., Oldman D.:. Report on shape analysis and matching and on semantic matching [online]. 2016. URL: http://gravitate-project.eu/ [cited 2016-05-01]. 2

[BCFS15] Biasotti S., Cerri A., Falcidieno B., Spagnuolo M.: 3d artifacts similarity based on the concurrent evaluation of heterogeneous properties. *Journal on Computing and Cultural Heritage (JOCCH) 8*, 4 (2015), 19. 2

[BCS*11] Beales R., Chakravarthy A., Selway M., Stapleton M., Kuhn S., Stevenson J., Luther S.: Icon: authentic 3d cultural heritage models for the creative industries. In *Electronic Visualisation and Arts conference (EVA), 2011* (2011). 2

[Doe03] Doerr M.: The cidoc conceptual reference module: an ontological approach to semantic interoperability of metadata. *AI magazine 24*, 3 (2003), 75. 2

[DvdB00] Dorst L., van den Boomgaard R.: The systems theory of contact. In *Algebraic Frames for the Perception-Action Cycle*. Springer, 2000, pp. 22–47. 2

[GG15] Garro V., Giachetti A.: Scale space graph representation and kernel matching for non rigid and textured 3d shape retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2015). 2

[HFG*06] Huang Q.-X., Flöry S., Gelfand N., Hofer M., Pottmann H.: Reassembling fractured objects by geometric matching. *ACM Transactions on Graphics (TOG) 25*, 3 (2006), 569–578. 2

[HPI*] Hermon S., Pilides D., Iannone G., Georgiou R., Amico N., Ronzino P.: Ancient vase 3d reconstruction and 3d visualization. 2

[IT11] Itskovich A., Tal A.: Surface partial matching and application to archaeology. *Computers & Graphics 35*, 2 (2011), 334–341. 2

[KDSK10] Kleber F., Diem M., Sablatnig R., Kampel M.: Proposing features for the reconstruction of marble plates of ephesos. In *Virtual Systems and Multimedia (VSMM), 2010 16th International Conference on* (2010), IEEE, pp. 328–331. 2

[NOR*16] Norton B., Oldman D., Rychlik A., Hermon S., Middleton S. E., Repetto A., Torrente M.-L.:. Report on existing 3d scans and metadata [online]. 2016. URL: http://gravitate-project.eu/ [cited 2016-05-01]. 2

[Old] Oldman D.:. Researchspace - a digital wunderkammer for the cultural heritage knowledge graph [online]. URL: http://www.researchspace.org [cited 2016-05-01]. 2

[PGS*15] Papaioannou G., Gregor R., Sipiran I., Schreck T., Savelonas M., Arnaoutoglou F., Zagoris K., Theologou P.:. Reassembly and object repair methodology report [online]. 2015. URL: http://presious.eu/ [cited 2016-05-01]. 2

[PT15] Paikin G., Tal A.: Solving multiple square jigsaw puzzles with missing pieces. In *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on* (2015), IEEE, pp. 4832–4839. 3

[SPS15] Savelonas M. A., Pratikakis I., Sfikas K.: Partial 3d object retrieval combining local shape descriptors with global fisher vectors. In *Proceedings of the 2015 Eurographics Workshop on 3D Object Retrieval* (2015), Eurographics Association, pp. 23–30. 2

[WC08] Willis A. R., Cooper D. B.: Computational reconstruction of ancient artifacts. *Signal Processing Magazine, IEEE 25*, 4 (2008), 65–83. 2