Digital Mont'e Prama: Exploring large collections of detailed 3D models of sculptures

Marcos Balsa Rodríguez, CRS4, Italy
Marco Agus, CRS4, Italy
Fabio Bettio, CRS4, Italy
Fabio Marton, CRS4, Italy
Enrico Gobbetti, CRS4, Italy

We present and evaluate a scalable interactive system for the exploration of large collections of detailed 3D digital models of sculptures. The system has been applied to the valorization of the Mont'e Prama complex, an extraordinary collection of protostoric Mediterranean sculptures, which depict models of cone-shaped stone towers, as well as larger-than-life human figures. The software architecture is based on scalable components for efficient distribution and adaptive rendering of extremely detailed surface meshes with overlaid information. The user interface, based on a simple and effective interactive camera controller tailored for touch interaction, has been designed for targeting both small screens and large display systems. The system components have been integrated in different interactive applications, ranging from large-screen museum setups and low end mobile devices both with very high visual quality. The large scale system has been installed in a variety of temporal and permanent exhibitions, and has been extensively used by tens of thousands of visitors. We provide an early analysis in this paper of the data gathered during a 20 month period in the National Archaeological Museum in Cagliari and a 6 months period in the Civic Museum in Cabras, for a total of over 67K exploration sessions.

Categories and Subject Descriptors: I.3.3 [**Computer Graphics**] Picture and Image Generation; I.3.7 [**Computer Graphics**] Three-Dimensional Graphics and Realism

General Terms: Cultural Heritage

Additional Key Words and Phrases: massive model rendering, mobile rendering

ACM Reference Format:

Marcos Balsa Rodríguez, Marco Agus, Fabio Bettio, Fabio Marton, Enrico Gobbetti. 2016. Digital Mont'e Prama: 3D Cultural Heritage Presentations in Museums and Anywhere. *ACM J. Comput. Cult. Herit.* 0, 0, Article 0 (January 2016), 23 pages.
DOI:http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

The continuous improvement and proliferation of interactive digital media is reshaping our education and information landscape. Museums are evolving into one of the principal components of our leisure and education industry and, in doing so, are faced with many challenges arising from the compet-

Address: CRS4, POLARIS Ed. 1, 09010 Pula (CA), Italy; email: firstname.lastname@crs4.it www: http://www.crs4.it/vic/

DOI:http://dx.doi.org/10.1145/0000000.0000000

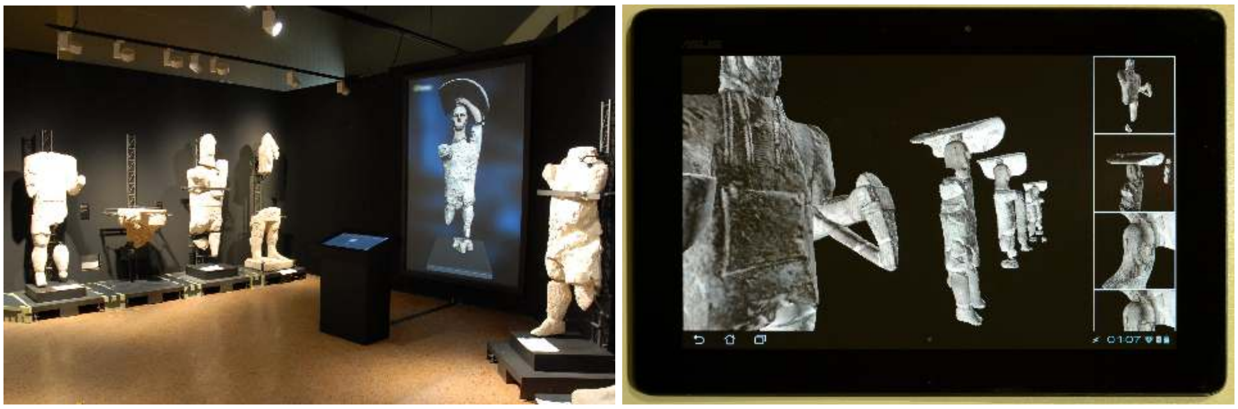

Fig. 1. **Digital Mont'e Prama project**. Left: permanent museal exhibition setup with a large projector display, together with a small group of real statues; Right: a collection of statues visualized on an ASUS Nexus 7 tablet.

ing needs of attracting large masses of visitors while preserving their specificity with an appropriate balance between leisure and learning. At the same time, people is continuously connected to a vast information landscape through mobile devices and pervasive high-speed Internet. In the context of Cultural Heritage, the use of highly detailed 3D models is often necessary to be able to fully represent and present the peculiarities of artifacts. In this sense, the idea of *e-tangible* collection is growing in the cultural heritage community. In order to fully cover the pre-visit (documentation), visit (immersion) and post-visit (emotional possession) phases [Economou and Meintani 2011] of cultural tourism, it is important to create integrated distribution and presentation ecosystems, capable to let users inspect cultural artifacts in both highly spectacular museum settings (e.g., large projection systems), as well as "on the move", using the capabilities of nowadays mobile devices. Thus, the challenges for digital technologies and technology design are to provide tools that will play a leading role in key issues such as providing access, increasing interaction, sharing knowledge and increasing the commercial viability of heritage institutions [Karran et al. 2015]. To this end, adaptive systems for easily and naturally exploring and discovering 3D high resolution Cultural Heritage scenes and models are strongly solicited. In this paper, we present the framework developed for the *Digital Mont'e Prama* project, which aims to present to the public a large unique collection of pre-historic statues of the Mont'e Prama complex, including larger-than-life human figures and small models of prehistoric buildings (cone-shaped stone towers). The overall complex is composed of 38 reassembled statues, and currently subdivided in two groups, available to visitors in the National Archeological Museums of Cagliari and Cabras. The project covered all 3D digital aspects ranging from acquisition and processing [Bettio et al. 2015] to exploration and rendering [Marton et al. 2014]. Here we describe our scalable exploration software architecture developed for the project which supports, in an integrated manner, distribution and rendering of massive annotated detailed models with high visual quality. The system enables ubiquitous distribution and visualization on mobile devices, as well as large scale screen projection-based visualization in museum setups. The main components of the system are:

—a framework for interactive exploration of massive highly detailed 3D models of sculptures, adapted for rendering on various scale display, ranging from mobile systems to large projection setups;

—an intuitive 3D control interface decorated with thumbnail-based point-of-interest selector, adapted to work with various scale displays;

—an infrastructure for authoring and presenting 2D annotation overlays over 3D models and scenes.

Thanks to our framework, object-aware user-interface components are capable to support models composed of various billions of primitives and large numbers of points of interest on either large scale projector systems and mobile devices, and have been used by tens of thousands of inexperienced users in real exhibitions showing the collection of Mont’e Prama sculptures. This work is an invited extended version of our *Digital Heritage 2015* contribution [Balsa Rodriguez et al. 2015]. In addition to supplying a more thorough exposition in this work, we also provide significant new material including an extensive analysis of usage of the large scale projector system in a real museum setting. The user study was carried out by recording the real time explorations performed by tens of thousands of visitors during almost 20 months of permanent exhibition at the Archaeological Museum of Cagliari and 6 months at Archaeological Museum of Cabras.

2. RELATED WORK

Creating interactive systems for the visualization of annotated massive 3D models, in particular for cultural heritage, requires handling large amounts of geometry while providing the user with a user interface that enables both global and in-detail inspection, and providing some additional information in order to better comprehend the artifact. Here, we discuss the works which most closely relate to our method.

Massive model rendering. There is a vast literature covering massive 3D model visualization [Gobbetti et al. 2008], but only recently real-time performance has been demonstrated on mobile devices for very large surface models. We have adopted the solution proposed by Balsa et al. [Balsa Rodriguez et al. 2013] which adapts the well-known Adaptive TetraPuzzles [Cignoni et al. 2004] to be used in networked environments using mobile devices.

Natural 3D interaction. In terms of complex scene exploration, most of the work in the area is connected to camera/object motion control [Jankowski and Hachet 2013]. Constrained viewpoint navigation has been proposed as an option to reduce the degrees of freedom in order to simplify the access to the user interface for novice users [Marton et al. 2014; Boubekeur 2014]. An automatic pivoting method has been presented [Trindade and Raposo 2011], but the method requires access to the depth buffer, which is not always available on mobile platforms, and suffers from discontinuities when considering models with sharp features. Our work extends to both co-located and not co-located interaction the approach recently presented by Balsa et al. [Balsa Rodriguez et al. 2014] on auto-centering virtual trackballs for small screens.

Image-assisted navigation. Image-assisted navigation has been implemented in a variety of ways in the literature. Images can be linked to viewpoints when navigating in a 3D scene, or can also be linked to additional information by using images as hot-spots [Andujar et al. 2012]. Some authors have proposed to hierarchically cluster images according to some image semantic, like combining time and space [Ryu et al. 2010], spatial image-distances [Jang et al. 2009], or a mixture of them [Mota et al. 2008]. Our approach is navigation-oriented, presenting the user with a selection of the best candidate views from the current view position. Thus, the user can at any time use the image-based navigation to travel to another point of view.

Information visualization. Using linked multimedia information to enhance the presentation of complex data has been long studied, mostly focusing on guided tours [Faraday and Sutcliffe 1997], text disposition and readability [Sonnet et al. 2005; Jankowski et al. 2010], usability of interaction paradigms [Polys et al. 2011], and the integration of interconnected text and 3D model information with bidirectional navigation [Götzelmann et al. 2007; Jankowski and Decker 2012; Callieri et al. 2013]. Most methods require precise picking to navigate through the information, thus presenting problems when targeting non co-located interaction setups (e.g., large projection displays), and often introduce clutter in the 3D view to display the pickable regions. The solution presented here is based on a method which presents contextual information associated to regions of interest in selected viewpoints, reachable with a thumbnail-based point-of-interest interface without requiring precise picking. A follow-up work, not described in this paper, concerns techniques useful to provide more structure (e.g., temporal ordering) among annotations, as well as implicit presentation methods based on recommendation systems [Balsa Rodriguez et al. 2015].

3. CULTURAL HERITAGE APPLICATION DOMAIN

The main motivation of this work was the valorization of the Mont'e Prama complex, which is a large set of extraordinary sandstone sculptures created by the Nuragic civilization in Western Sardinia [Tronchetti and Van Dommelen 2005]. The complex was casually discovered during hoeing in 1974, and it appeared soon evident that these were truly unique sculptures, displaying features related to both the indigenous Sardinia Nuragic civilization as well as the wider world of the Phoenician colonizers. Four archaeological excavation campaigns were carried out between 1975 and 1979, and revealed five thousands of sculpture fragments which, according to the most recent estimates, comes from a total of life-size 44 statues depicting archers, boxers, warriors and models of prehistoric buildings. With respect to dating, they can be traced to an as-yet undetermined period, which goes from the tenth to the seventh century BC. All the sculptures representing human figures conform to a single iconography: they are all standing upright and rendered in a flattened manner in the sense that both the front and back are elaborated but no attempt has been made to work them in the round. The faces are the distinctive element of the sculpture complex: they are schematic, defined by the heavily rendered eyebrows and a straight nose that make up a prominent T-shape(see figure 7(d)). The eyes are represented by two nested circles that have clearly been created with a pair of compasses, while the mouth is indicated by a simple shallow line [Lilliu 1997]. Restoration was a complicated and time-consuming process: it was carried out at the *Centro di Restauro e Conservazione dei Beni Culturali (CRCBC)* of Li Punti (Sassari, Italy) and resulted in the partial reassembly of 25 human figures with height varying between 2 and 2.5 meters, and 13 approximately one-meter-sized building models. Modern non-invasive restoration criteria were considered for the reassembly. No drilling or insertions were performed into the sculptures, fragments were glued together using a water-soluble epoxy resin, and all the gaps on the resin-filled surface were covered with lime-mortar stucco. Furthermore, custom external supports were designed to ensure stability to all the components of the statue, reduce contacts and maximize visibility. The digital valorization project consisted of a long term scientific and technological research plan on the 3D digitization of the complex, aimed at covering all the phases, from the acquisition and processing to the application of innovative methodologies for multiresolution storage, distribution and display. The acquisition process resulted in 37 quarter-millimeter resolution colored surface models [Bettio et al. 2015] (see figure 2 for snapshots and sizes of the collection models). The rest of the paper describes the architecture developed for the distribution and real time interactive exploration of the models on commodity platforms, and for museal exhibitions.

4. SYSTEM REQUIREMENTS

The design of our system has been guided by the requirements collected from domain experts, as well as our analysis of the related work (see Sec. 2). Additional requirements were added from our past experience developing interactive systems for cultural heritage [Marton et al. 2014]. Below we summarize the main requirements that were derived from our analysis for the design of the different system components. See also our previous work for an account of requirements gathered in the Digital Mont'e Prama Project concerning user interfaces [Marton et al. 2014; Balsa Rodriguez et al. 2014] and information presentation [Balsa Rodriguez et al. 2015]. The requirements are numbered with respect

Fig. 2. **Digital Mont'e Prama collection.** The acquisition process resulted in 37 colored triangle meshes at the resolution of 16 vertices/mm^2. Mesh size (T) in millions of triangles and height (H) in meters from the base level are indicated below each sculpture.

to the various system components: distribution and display (**R.DD.x**), user interface (**R.UI.x**), and information presentation (**R.IP.x**).

4.1 Distribution and display

[R.DD.1] High-resolution geometry. Cultural heritage artifacts generally present information at multiple scales (i.e. global shape and carvings). In particular, the micro-structure typically carries valuable information on the carving process. The Mont'e Prama statues, in fact, present millimeter-sized carvings, thus requiring sub-millimetric model precision.
[R.DD.2] Large-scale visualization. The imposing size of many artifacts (i.e. statues) has to be transmitted to the visitor for him to have a correct comprehension of the original work. In order to better represent the aura of the statue, in museal exhibitions large displays which reflect the size of the represented object are needed.
[R.DD.3] Display flexibility. In order to cover the whole pre-visit, visit, and post-visit phases, one should support a wide range of setups, including museum setups, as well as smartphone and tablet applications. The specific application case for our designs is the presentation of larger-than-life human statues, reaching over 2.5m of height. The statues were constructed at imposing scale, and this macro-structure information should be immediately conveyed to the visitor through a real-scale (or larger-than-real) presentation. This means, in particular, that in museum setups we need to support large (wall-sized) displays.
[R.DD.4] Public accessibility (ubiquity). Nowadays, smartphones and tablets have become affordable enough as to be present everywhere. Furthermore, the connection bandwidth and performance of those devices are continuously increasing at fast paces, thus providing a great platform for media distribution. So, not only to empower the public access to the information, but also to improve the visit experience (pre-visit and post-visit phases), the pervasiveness of mobile devices should to be exploited.
[R.DD.5] Low bandwidth usage. The interactive exploration of those huge models produces a substantial amount of data traffic. In order to provide an interactive and pleasing exploration experience to the user the size of the transmitted information has to be aware of the bandwidth capabilities of the targeted systems (i.e. WIFI and 3G/4G connections).

4.2 User interface

[R.UI.1] Single user control, multi-user fruition. Museums have large amounts of visitors willing to inspect the various collections exhibited. Therefore, enabling multiple visitors to observe the display while the user in control of the interface inspects the object, facilitates the flow of visitors to access visual information faster. In such a way, it will be improved both the visitor experience, since wait time is much less, and the total number of visitors who can access the museum installations.
[R.UI.2] Seamless interactive exploration and zooming. The high resolution geometry in the 3D representations requires the user to be able to reach viewpoints all around the model, both from far and from very close distances. The exploration, therefore, has to permit a seamless transition from global shape visualization to close detail inspection (i.e. carvings).
[R.UI.3] Fast learning curve and assisted navigation. Museums receive large amounts of visitors, so any kind of training taking long time is not practical. Hence, the user interface is required to be immediately accessible, allowing the user to inspect the artifact both by directly trying and understanding the user interface, or by using some assisted navigation that gives some hints and/or shortcuts to explore the object.

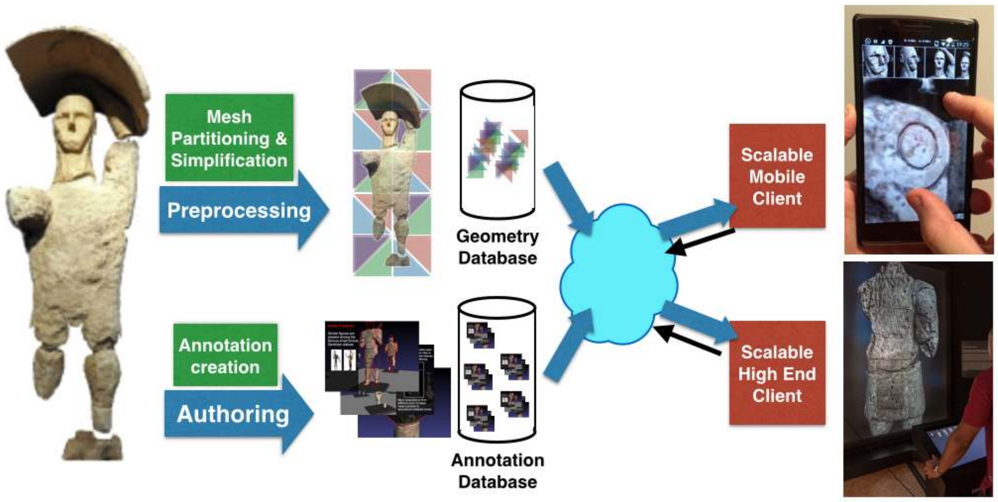

Fig. 3. **Architecture components.** Pre-processing and server components are shared between the mobile and high-end clients. User interface is implemented on touch devices.

4.3 Information presentation

[R.IP.1] Focus on work of art. The visitor intention is on improving his knowledge on the artifact being inspected. So, the user interface needs to be unobtrusive and avoid any distracting widgets, or occlusion from other visitors. At the same time, it should provide the user with an easy and direct way to get to the information without requiring complicated interactions.
[R.IP.2] Information authoring. Textual and visual information (drawings, images) should be supported. Editing should be made possible for museum curators and archaeologists without particular training. Adding annotations and linking them should not require intervention of specialized personnel.
[R.IP.3] Information spatially connected with 3D models. Most of the information, textual and visual, is spatially connected to a region of a 3D model. This implies that descriptive information should be associated to parts of the cultural objects, typically seen from a canonical point of view, or at least close to it. Examples are descriptions of carvings and decorations, reconstruction hypotheses, comparisons with other objects. Different macro-structural and micro-structural views should be associated with different kinds of information.
[R.IP.4] Engaging experience. In general, visitors do not want to be overloaded with instructional material, but to receive the relevant information, learn, and have an overall interesting experience, which should be be personal, self-paced, and exploratory. The user interface should provide guidance, while not being perceived as overly obtrusive.

5. SYSTEM COMPONENTS

Our system is based on a touch interface for the exploration of highly detailed 3D objects using a multiresolution framework for interactive rendering and networked components for streaming the required data. The main components of the framework are represented in Fig. 3.

5.1 Scalable rendering and distribution

Real-time rendering of large datasets at high resolution, that is billions of triangles and hundreds of points of interest per model, requires dealing with two main problems: the complexity of the model, and the amount of data to be transmitted and loaded into core memory. Dealing with such complex models requires specialized multiresolution structures and adaptive algorithms. Our rendering framework, and also the user-interface subsystem, exploit a multiresolution representation of the 3D models, based on a variation of Adaptive TetraPuzzles (ATP) [Balsa Rodriguez et al. 2013]. Specifically, the ATP method has been extended with all the geometric queries required to compute the automatic pivot, and kd-trees to organize the points of interest. In order to cope with the low resource availability on mobile devices, we have used a compact representation that encodes position, normal and color into 64 bytes, thus better scaling on devices with low memory available, while still providing high visual quality. In addition, this compact representation is further compressed for data streaming, to reduce bandwidth requirements and so improve the refinement speed, and the number of triangles per node has been increased from 8k triangles to 16k triangles in order to improve streaming performance by reducing the number of packets to transmit.

5.2 User interface

According to requirements described in section 4, we designed our user interface components in a way that they could be easily adaptable to various configurations, and easily usable with fast learning times by casual users. Furthermore, we wanted to keep the interface as less distracting as possible, in order to keep the main focus of users to the details of sculptures and the information associated. The main components of our user interface are a selection widget for easily navigating through the entire sculpture collection, an object-based assisted system for camera control tailored for both orbiting and proximal object inspection, and an image-based navigation component for navigating amongst a series of pre-defined viewpoints and displaying associated information in overlay.

Fig. 4. **Sculpture selection widget.** The main interface is organized as two level scroll interface hierarchies. The Mont'e Prama sculpture collection is subdivided in 4 categories: boxers, archers, warriors and building models(see left). Selecting one of the group, the sculptures of given category are presented (see right), and can be chosen.

Sculpture selection widget. In order to enable casual users to easily navigate the collection and select the sculptures to explore, we designed a simple selection widget based on a hierarchy of thumbnail-based scroll interfaces. The first level of hierarchy is a scroll bar containing images of the four statue categories: boxers, archers, warriors, and building models (see figure 4 left). The second level of hierarchy contains instead, for each group, the thumbnails associated to each sculpture model (see figure 4 right). Navigating and exploring the entire collection is immediate, and users can decide in any moment to explore a different sculpture model with a limited number of scroll and select operations.

Camera control. Interactive exploration of complex 3D models requires continuously switching between rotation, panning and zooming, which can be difficult for novice users to handle. For this reason, we have implemented an automatic centering pivot (see ACeViT [Balsa Rodriguez et al. 2014]) which performs a stochastic sampling of the visible surface every time the user pans or zooms in order to determine the pivot. In this way, the rotation pivot is always centered on the part of the object which is currently visible, thus providing a natural and understandable rotational behavior.

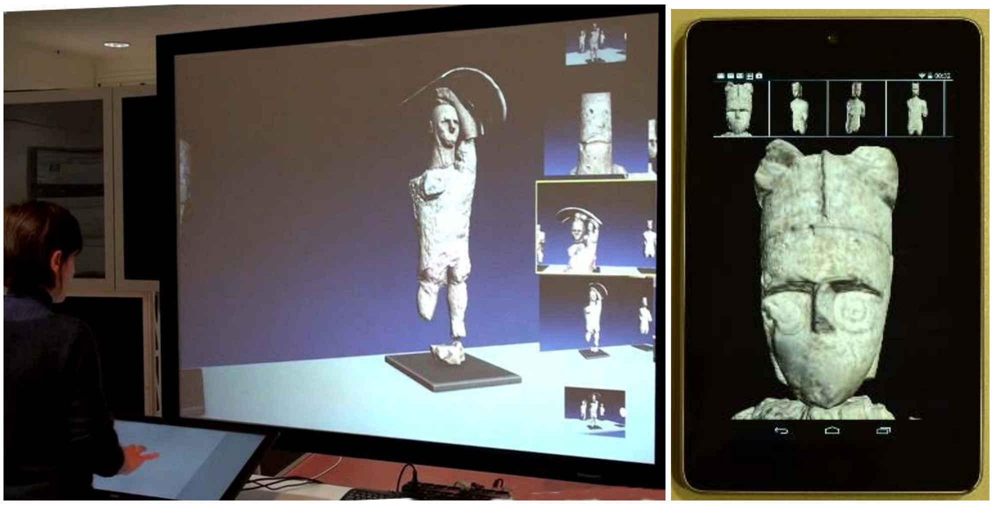

Fig. 5. **Image-based selection.** The user is presented with a set of interesting viewpoints near the current position. The same scroll view selector is employed for both the large display setup (see left), and the small screen mobile setup (see right).

Image-based navigation. Image-assisted exploration has been implemented in many ways in the literature, both as hot-spots to highlight interesting positions, or as 3D viewpoints linked to the scene providing real images integrated into the 3D representation. . Our approach is navigation-oriented, providing the user with a collection of context-based selection of best views that can be used to travel to near interesting view points without distracting her from the main exploration task. This selection is based on an image-space distance function which selects a small number of views which are most similar in image-space to the current one. Figure 5 shows examples on how point-of-interest view selectors are integrated in a large display setup and on a small screen setup.

6. INFORMATION PRESENTATION

Communicating information on cultural artifacts require the ability to mix highly photorealistic presentations that convey the aura of original objects, with illustrations that provide additional insights on these objects. One of the goals is to study how to best communicate auxiliary (geometric, conceptual and semantic) information that is attached to 3D models. A variety of systems and standards for annotating and linking metadata to artworks exist [Szekely et al. 2013; Dijkshoorn et al. 2014]. However, most of efforts are targeted to manage, link and spread information related to 2D artworks and images, by exploiting the capabilities of web and cloud computing [Isaac and Haslhofer 2013], and adding spatial annotations to 3D models is still not considered an easy task for curators and cultural heritage experts. Our system was designed aiming at making it easy for archaeology experts and curators to quickly edit various kinds of informative content spatially associated to the sculptures. A follow-up work, not described in this paper, concerns techniques useful to provide more structure (e.g., temporal ordering) among annotations, as well as implicit presentation methods based on recommendation systems [Balsa Rodriguez et al. 2015].

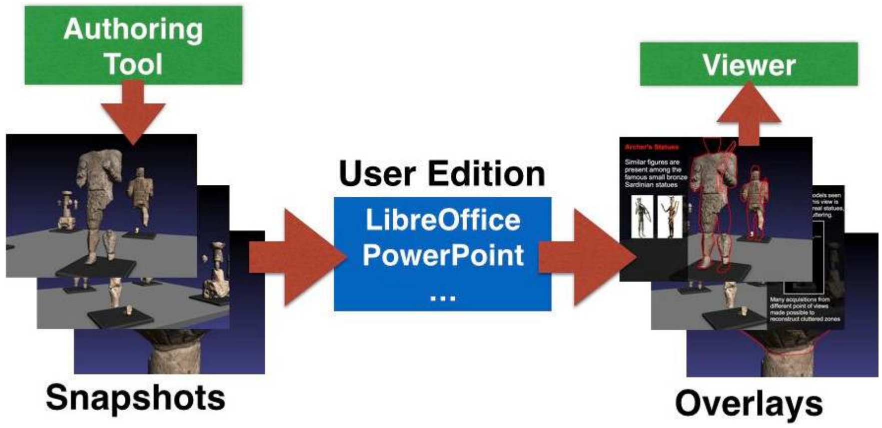

Fig. 6. **Authoring overlays.** Experts can create information overlays by selecting 3D points of interests, and by editing above associated views with popular image or presentation authoring software.

Information authoring. The authoring system is a version of the exploring application in which experts can navigate the sculpture collection and choose interesting views which they aim to comment, or where they want to draw some highlight or sketch, or add other kind of information (images, links). All selected views are saved in a database in form of pictures and the associated view transforms. Once the selection process is completed, the editing process consists of adding the annotations by employing the preferred image editing or presentation software to decorate the selected pictures. Each decorated images is then added to the database as overlay picture with transparent background (obtained with the removal of the original picture), and the information is ready to be displayed when users reach the associated view position. The authoring system is easy to use, fast and can be customized to add heterogeneous data as fields of database. Figure 6 shows a schematic diagram of the process of creating annotations, from selection of interesting view, to the information editing, and the resulting overlay picture to be displayed during the 3D navigation.

Overlay display. During the navigation, users can scroll a thumbnail-based point-of-interest selector showing the interesting viewpoints close to the area of the sculpture which is explored in that moment. The information overlays are displayed any time the user selects an interesting picture. Specifically, the system automatically leads the user to the selected view position and the information overlay appears to decorate that view, and it persists for a fixed amount of time (10-20 seconds according to the displayed information) or until the user decides to explore different parts of the sculpture or the collection by performing further interactions. In this way, users can easily get different kind of interesting information associated to the collection of artworks (see section 8).

7. SYSTEM INTEGRATION

All the components of the architecture were designed to be easily adapted to work with different kind of displays and devices, ranging from smartphones to large high-resolution screens. We tested them in two different setups: one high end employing a large projector, and one for mobile devices.

High end setup. For this setup, we have preferred to decouple interaction and rendering, in order to permit the use of large displays for rendering. The touch device in charge of the user interface is located at a distance from the display enough to grant the user controlling the inspection with a whole view of the display, see Fig. 1. At the same time, this setup enables multiple users to watch the display without occlusion problems, which wouldn't be possible with an interactive wall display or a large touch monitor. This setup is an ideal component of presentation systems designed for large audiences or small working groups sharing a common view. This is the case of interactive installations in museums. Another domain is visual simulation and rehearsal, already used to support the cooperative work of small research group in archaeology for presenting and assessing reconstruction hypothesis(e.g., for the assessment of virtual restoration results or inspection of degradation status).

According to our design, the user interacts with a non co-located touch-based interface that provides very little visual information in order to not distract the user from the inspection task. The 3D navigation interface doesn't require co-located interaction, thus the user can abstract interaction just by dragging and pinching without any need to look at the touch monitor. At any time the user can scroll the image-based view selector interface on the side, which shows the closest suggested point-of-interests (see Fig. 5).

Mobile setup. The majority of the requirements derived in the section 4 are common both for the high end setup, and the mobile application, with the main difference being with respect to the display size and user interaction, besides performance issues. In this case, the co-located 3D navigation interface maps typical touch-based gestures, like dragging and pinching, into a behavior that appears natural and coherent with other standard mobile applications. Since there is no need for precise touch input it is easy for the user to concentrate on the inspection task, getting used to the interface just by trying typical gestures. The computation cost is only a small fraction of the frame rendering time, enabling the use of this solution in low performance devices. Due to the small size of the screen we have opted for a scroll widget that occupies a small part of the screen, allowing the user to select the target view with a simple click at any time, see Fig. 5. Additionally, since the list of proposed views is updated only when user stops interacting (i.e. not touching the screen), there are no distracting elements while the user is exploring the object. The rendering and user interface frameworks are shared between the various systems, providing a unified platform for distribution and rendering of large 3D models, see Fig. 3.

8. IMPLEMENTATION AND PERFORMANCES

We have implemented the proposed approach both as a high-end system for museum installations, and as a mobile application, targeting mid-level hardware. Both systems share the data access components, using a NoSQL database (leveldb) for local data access and HTTP+Apache2 for networked access. The preprocessing step is also shared and enables processing with a performance of about 24K triangles/second on an Intel Core i7 960 with 24GB of RAM.

8.1 High-end system

The presented approach has been integrated into a reference system using OpenGL and GLSL for rendering, and Qt 4.7 for the interface. The system runs on a PC with Ubuntu Linux 12.10, with a Intel Core i7-3820 @ 3.6Ghz, 8GB of RAM and a NVIDIA GTX 680 GPU, sending the visual output to a 156cmx250cm back-projection screen through a 7500 ANSI Lumen DigitalProjection E-Vision WUXGA (1920x1200) projector. The user interface is controlled with a 27" DELL P2714T multi-touch screen, where very brief information is shown about the commands for controlling the navigation, only when there is no interaction, in order to attract the attention of the visitors. The multi-touch screen was designed to be placed 180cm away from the projection screen, allowing the user to see the whole screen in his field of view and, at the same time, to avoid obstacling other visitors to watch the display. The streaming performance under both high speed WIFI or wired networks provided refinement times below 5-10 seconds from cool start to a very close detail (i.e. an eye), requiring about 70MB download. An experimental setup for showing annotations has also been implemented: during a typical exploration of the complex, users can gather lot of information and easily learn, for example, about the casual dis-

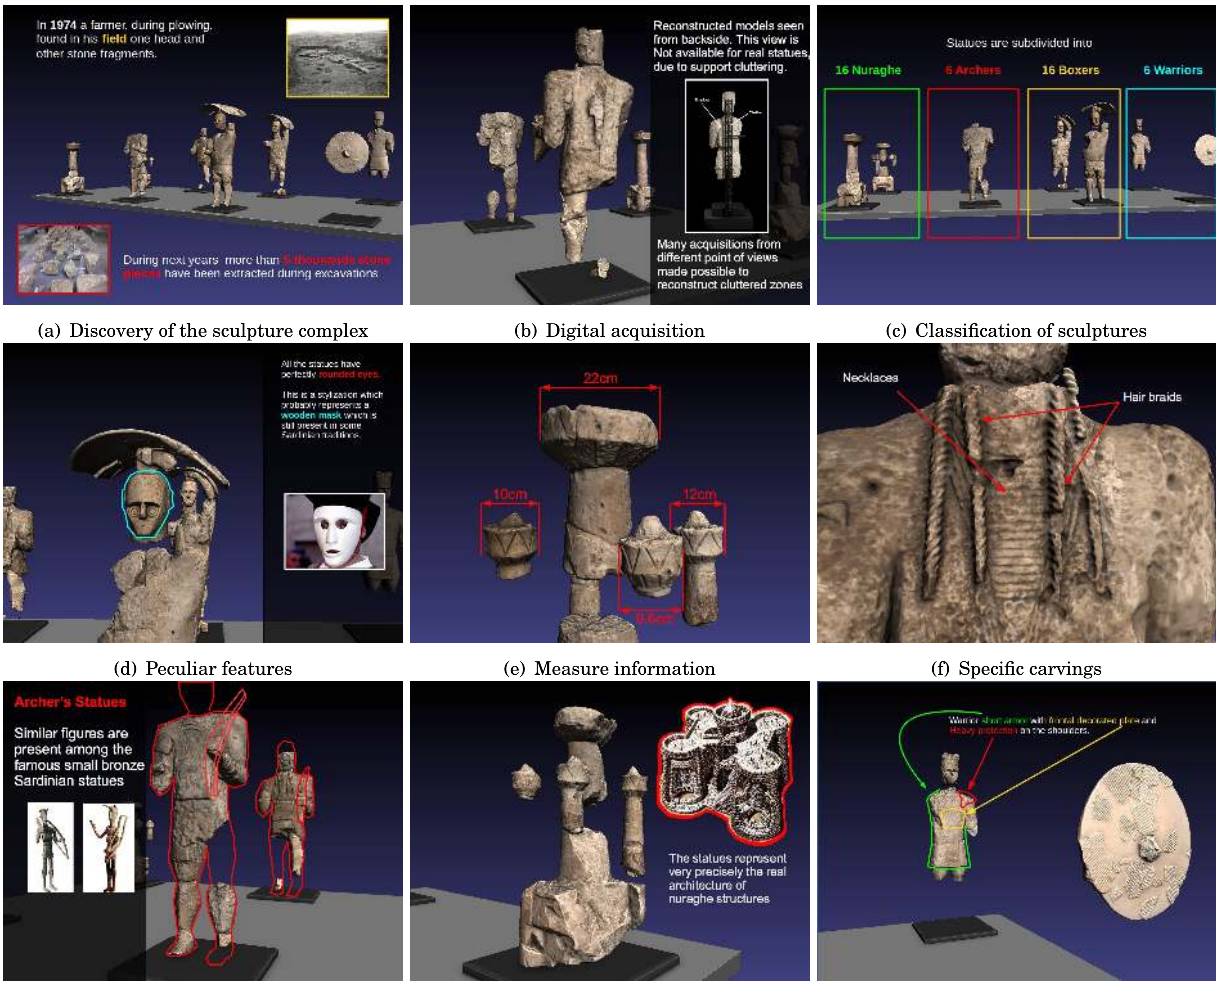

(a) Discovery of the sculpture complex (b) Digital acquisition (c) Classification of sculptures

(d) Peculiar features (e) Measure information (f) Specific carvings

(g) Connections to other artworks (h) Comparisons to other buildings (i) Highlights of sculpture parts

Fig. 7. **Information overlay examples.** Various kind of information are presented to users during the exploration of the sculpture collections, highlighting interesting artistic details or historical information.

covery of the sculptures (see Fig. 7(a)), the acquisition process (see Fig. 7(b)), their classification (see Fig. 7(c)), their peculiar and unique features (see Fig. 7(d)) as well as the similarities with respect to other artworks (see Fig. 7(g)) or to other existing building ruins (see Fig. 7(h)). Moreover, specific details can be highlighted on single sculptures, indicating measures (see Fig. 7(e)), specific carvings (see Fig. 7(f)), and the parts composing the models (see Fig. 7(i)).

8.2 Mobile system

We have also implemented our approach as a Android 2.3+ application, using OpenGL ES 2.0 and GLSL for the rendering through a thin wrapper unifying both ES and desktop OpenGL APIs. Qt 5.2 has been used for handling UI events and GL context creation, also providing good portability amongst Android, Windows, Linux or iOS. The application was tested in a variety of devices covering mid-class and high-class hardware, a number of display sizes ranging from 4” to 10”, and a wide range of screen

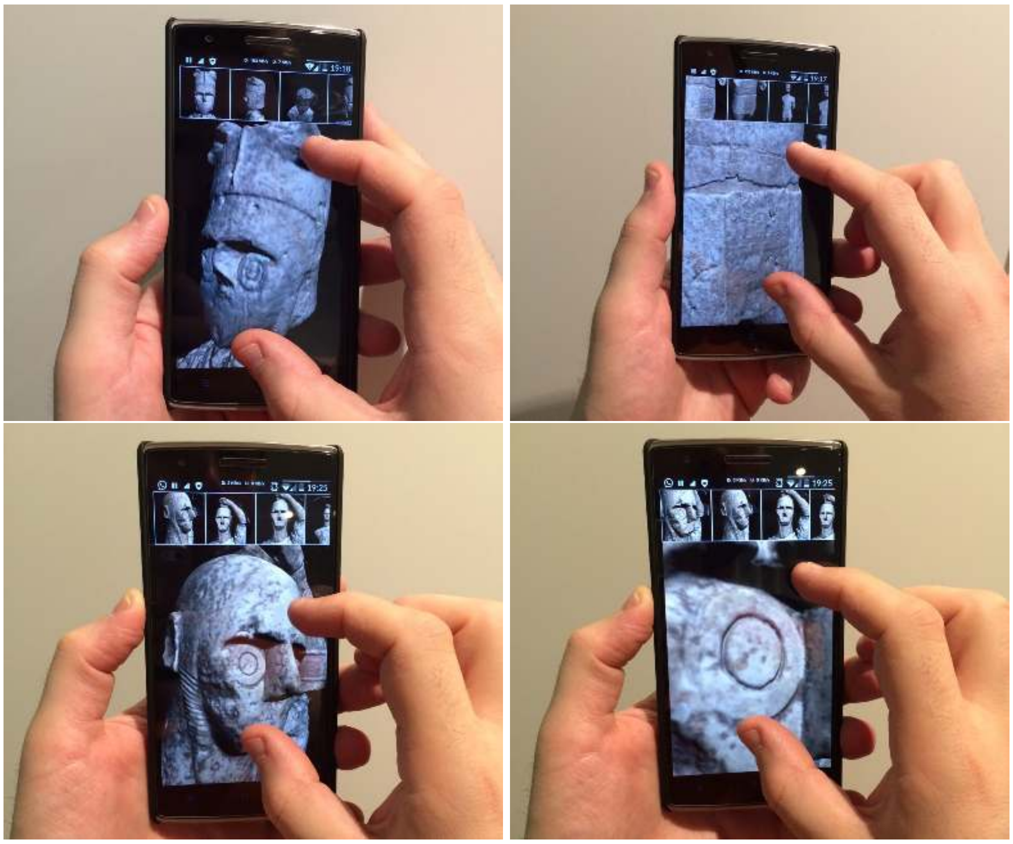

Fig. 8. **Mobile device application.** Exploration of highly detailed sculptures of the collection on a Android OnePlus One smartphone.

resolutions (i.e. 1024x768 to 1920x1080). The performance was always interactive with minor performance issues under maximum render load, that was palliated by adaptively lowering the maximum triangle budget when interacting in order to maintain interactive frame rates. On mobile devices, the WIFI performance is very heterogeneous varying heavily amongst devices, besides the different screen resolutions. Initial loading time is below 5 seconds, while additional 5-10s are required for a fully refined full view. In order to have a fully refined close view (i.e. an eye or some millimeter sized carvings) 40-60 seconds are required to download the required data (i.e. below 40MB for 1Mpixel resolution, below 70MB for 2Mpixel resolution). Under UMTS/HSPA connection, again the main factor is connection performance, so under a connection with a peak fetch performance of 3.4Mbps having a fully refined close view required 120-140 seconds. Nevertheless, thanks to the multiresolution approach within a few seconds a reasonable quality representation is available for inspection. A series of preliminary usability tests were performed in order to evaluate the system with positive results. Figure 8 shows the exploration of two high resolution sculptures of the complex on a OnePlus One smartphone, running Android OS, with a 5.5" FHD screen, 3GB RAM, and a Snapdragon 801 CPU, paired with an Adreno 330 GPU.

9. USER ANALYSIS

The high-end setup has been permanently installed in two museums, one at the Museo Archeologico Nazionale di Cagliari, one at the Museo Civico di Cabras, where a large amount of visitors have been able to interactively explore the 3D representations of the various objects (i.e. statues and models) included in the collection in the same rooms containing the original statues. Furthermore, various custom versions of the museum systems have been installed in successful purely virtual temporary exhibitions visited by tens of thousands of people: Fiera Internazionale di Cagliari (May 2013), Museo Nazionale Preistorico Etnografico "Luigi Pigorini" in Rome (November 2014-March 2015), and Museo Civico Archeologico in Milan (May-November 2015), Universal Expo Sardinian Booth in Milan (September 2015).

We carried out an extensive user analysis by recording the long period activity with the large projector exploration setup on the archaeological museums of Cagliari and Cabras, to interpret and understand the behavior of visitors during virtual sculpture explorations. Each of these two museums contains a permanent exhibitions of the collection of the restored Mont'e Prama sculptures, namely 26 sculptures in Cagliari, and 10 in Cabras. In both museums, visitors have the chance to use the 3D exploration system to visually compare the appearance of real displayed sculptures and the digitally reconstructed ones. Furthermore, users can use the system to visually explore the remaining part of Mont'e Prama sculpture complex, which is instead displayed at the other archaeological museum. Specifically, with the installation of the permanent visualization setup, we added a logger of the system usage, which gave chance to acquire input data from thousands of museum visitors. The data was collected in the Archeological Museum of Cagliari in the period between the inauguration of the "Mont'e Prama exhibition", March 23rd, 2014 and November 22nd, 2015, and in the Archeological Museum of Cabras in the period between March 21st, 2015 and October 18th, 2015. In the following paragraphs we provide an extensive analysis of these data, as well as a discussion of the main findings, about visitor behaviors.

Since we expected to collect a great quantity of user data, we decided to focus on the main exploration activities of users, and we recorded the following actions associated with timestamps:

- start interactive exploration of sculpture X at time t
- stop interactive exploration of sculpture X at time t
- inspect the sculpture X at time t from the viewpoint PV reached through free 3D interaction
- inspect the sculpture X at time t from the viewpoint PV reached through selection from a list of precomputed points of interest

With respect to the viewpoints acquisition, we limited it to the positions considered interesting by users, in the sense that we considered only the camera positions in which visitors kept observing the sculptures for at least 5 seconds.

General usage analysis. From the log files containing data collected during the observation periods, we were able to gather general statistics about the usage of the system. With respect to the exploration times, we analyzed the total usage time of the system, by integrating the sculpture exploration durations. Considering the days in which the exhibition was publicly available, we obtained that the system had a mean daily usage of 122.4 ± 99.2 minutes in the archaeological museum of Cagliari, and 149.1 ± 80.2 minutes in the archaeological museum of Cabras. The usage time computed in this way is clearly underestimated since it does not consider the time visitors employ for navigating the two level menus and for selecting the sculptures. We also carried out an analysis of daily usage with respect to the month and the week day. Figures 9(a) and 9(b) show the boxplots of the system usage daily times for Cagliari, per month and per week day, while figures 10(a) and 10(b) show the boxplots of the

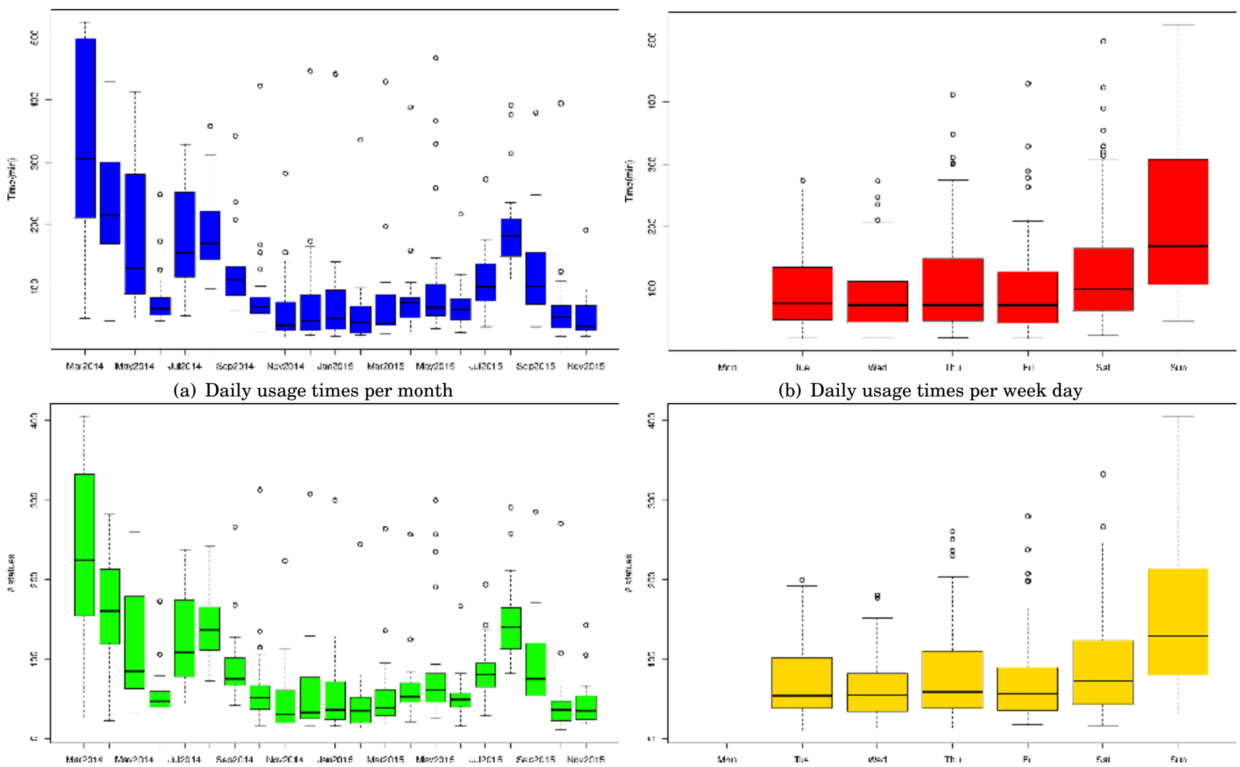

(a) Daily usage times per month

(b) Daily usage times per week day

(c) Daily sculptures count per month

(d) Daily sculptures count per week day

Fig. 9. **Archeological Museum of Cagliari exploration system daily usage times and sculptures explored.** a) Daily times are separated per month basis. b) Daily times are studied with respect to the week day. c) Number of sculptures daily explored per month basis. d) Number of sculptures daily explored studied with respect to the day of the week.

system usage daily times for Cabras, per month and per week day, as rendered by the *R* package [R Core Team 2013]. In these boxplots, as well as in the following ones, the bottom and top of each box are the first and third quartiles, the band inside the box is the second quartile (the median), and the ends of the whiskers extending vertically from the boxes represent the lowest datum still within 1.5 IQR (inter-quartile range) of the lower quartile, and the highest datum still within 1.5 IQR of the upper quartile. Outliers are indicated as small circles. From the boxplots, it appears evident that the system usage times are strictly correlated to the number of persons visiting the museum: with respect to the week days, the peaks are during the week ends, and especially during Sundays (mean daily usage of 211 ± 127.5 minutes at Cagliari). This is due to the fact that normally museums are visited during festive days, but the statistics are influenced from the fact that the entrance to the Archeological Museum of Cagliari is free on the first sundays of each month. Similar correlation can be observed to the usage total time with respect to the month: in this case, the system was more used during the first months of the exhibition in Cagliari (mean daily usage of 320.8 ± 170 minutes during March 2014), and other peaks can be observed during summer months since July and August are the periods with greater tourist presence in Sardinia (mean daily usage of 194.4 ± 70.6 minutes during August 2015 in Cagliari, and of 266.3 ± 64.1 in Cabras). In figure 9(a) the limited number of outliers represent the sundays in which the museum entrance is free. Similarly, we examined the number of statues daily explored with the system: in the archaeological museum of Cagliari, a total number of 43953 sculpture exploration sessions were carried out during the full period, with a mean of 89.7 ± 68.2 statues per

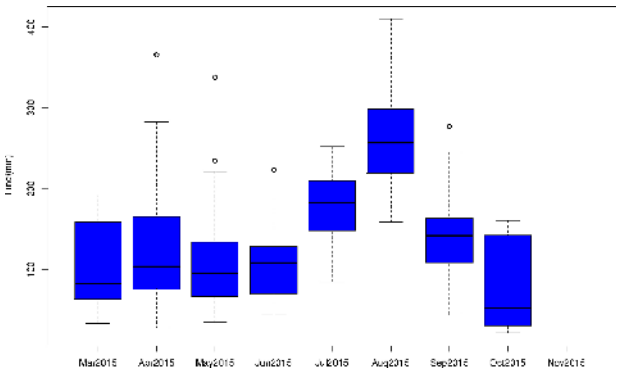

(a) Daily usage times per month

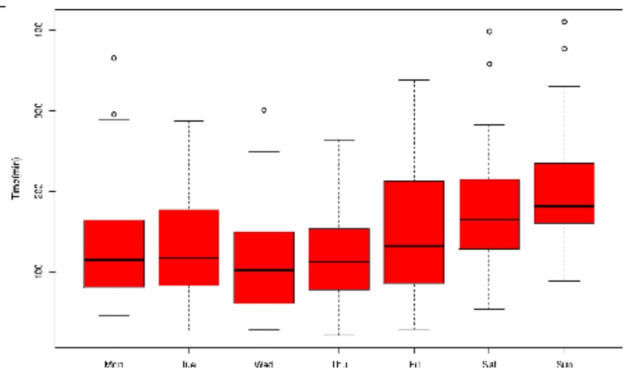

(b) Daily usage times per week day

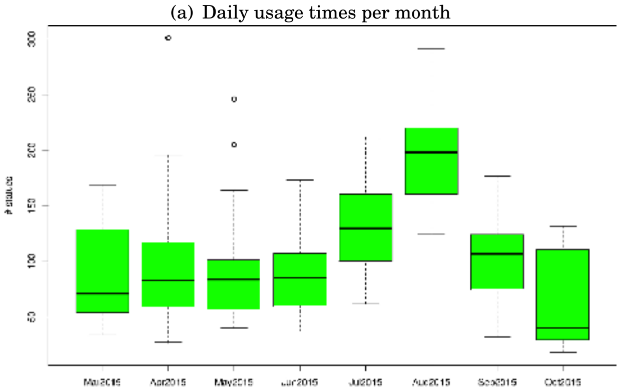

(c) Daily sculptures count per month

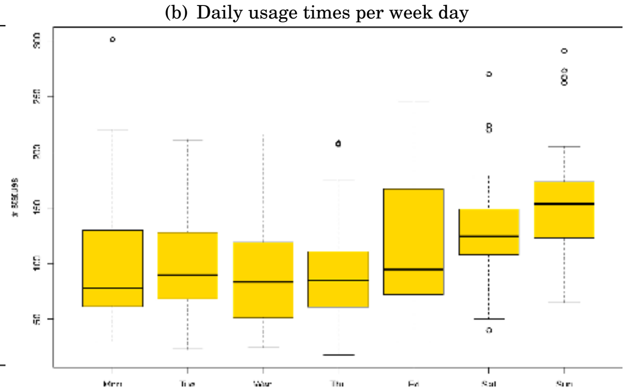

(d) Daily sculptures count per week day

Fig. 10. **Archeological Museum of Cabras: exploration system daily usage times and sculptures explored.** a) Daily times are separated per month basis. b) Daily times are studied with respect to the week day. c) Number of sculptures daily explored per month basis. d) Number of sculptures daily explored studied with respect to the day of the week.

day. In the archaeological museum of Cabras, the observation period was significantly shorter, hence the total number of exploration sessions was 23537, with a mean of 113.3 ± 58.2 statues per day. We carried out statistics with respect to the month and the week-day, reported in figures 9(c) and 9(d) for the museum of Cagliari, and figures 9(c) and 9(d) for the museum of Cabras. The distribution of the number of explored statues between week days and months show similar trends for both museums and direct correlation with respects to the system usage time: for the considered periods, in Cagliari the peak was in May 2014 (228.4 ± 113.3 statues per day explored), while in Cabras the peak was in August 2015(196.5 ± 45.2 statues per day explored). Furthermore, figures 11(a) and 11(b) show the ordered partial histograms of number of sculptures explorations (for space and clarity reasons we included only the six most visited statues) in Cagliari and Cabras. The diagram highlights the main preferences of visitors: as expected, in both museums users tend to prefer exploring the sculptures exhibiting the best conservation conditions and the most significant features. Specifically, in both museums the ideal podium is the same: the main boxer currently displayed in Cabras is the most explored with the system (5047 visits in the museum of Cabras in 6 months, and 10676 visits in the museum of Cagliari in 20 months), followed by the archer (2828 visits in Cabras and 4756 visits in Cagliari), and by the boxer currently displayed in Cagliari (2314 visits in Cabras and 3306 visits in Cagliari). Interestingly, in the six top explored ranking, all categories of the Digital Mont'e Prama collection are represented (boxers, archers, warriors, and nuraghe models). With respect to the exploration times, boxplots in figures 11(c)

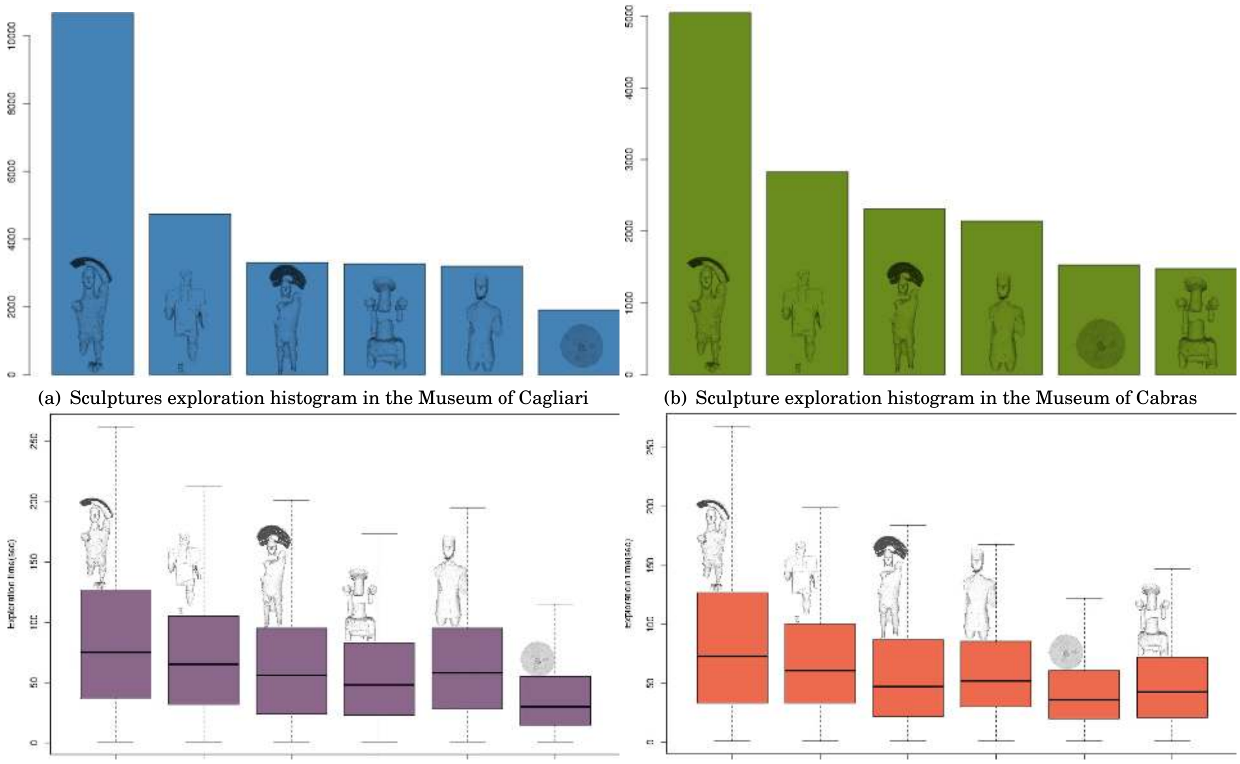

(a) Sculptures exploration histogram in the Museum of Cagliari
(b) Sculpture exploration histogram in the Museum of Cabras
(c) Sculptures exploration time in the Museum of Cagliari
(d) Sculpture exploration time in the Museum of Cabras

Fig. 11. **Sculpture preferences.** a) Partial histogram of number of sculpture explorations in the Museum of Cagliari. b) Partial histogram of number of sculpture explorations in the Museum of Cabras. c) Sculpture exploration duration at the Museum of Cagliari. d) Sculpture exploration duration at the Museum of Cabras.

and 11(d) show the exploration durations for the main sculptures in the museums of Cagliari and Cabras: the mean sculpture exploration time is of 63.8 ± 75.7 seconds in the museum of Cabras, and 66.8 ± 73.9 seconds in Cagliari. In both museums, the same boxer model is the most carefully explored sculpture (97.5 ± 96 seconds in Cagliari, and 98.1 ± 104.9 seconds in Cabras).

3D exploration analysis. Log files contain also the view positions from which visitors decided to observe the sculptures for more of 5 seconds. For the most explored sculptures of each category we carried out an analysis of the visitor behaviors, trying to understand which parts of the sculptures were considered the most appealing, and trying to individuate the main movements and actions. In general, there is a bunch of methods aiming at finding the regions of interest in scenes and models, based on measures of surface saliency [Lee et al. 2005] or importance based on geometric features [Leifman et al. 2012], and used for detecting the most informative viewpoints [Callieri et al. 2013], or for driving geometric processing algorithm, as simplification, or registration. In our case, in order to have a visual representation of visitors interest, we derived an **interest map**, in which, for each vertex v_k of a sculpture model, a focus value ϕ_k is computed with respect to the view points represented by projection and view matrices P_i and V_i, from which it was observed. The interest value is computed by considering the following contributions:

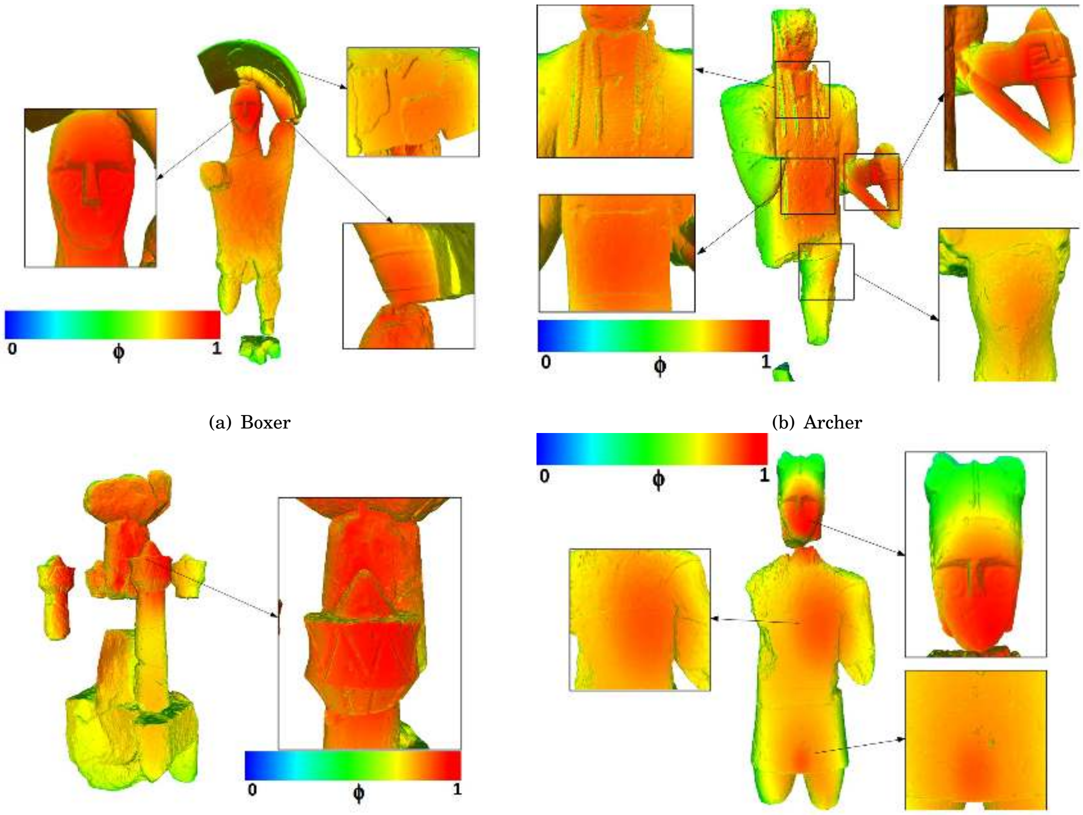

(a) Boxer (b) Archer

(c) Nuraghe model (d) Warrior

Fig. 12. **Interest maps according to visitor viewpoints for the most significant sculptures.** The colors are computed with respect to a log-normalized scale of focus values computed by using equation 1. From blue to red, blue indicates low interest while red indicates high interest. Insets show the the most appealing parts of sculptures according to these interest maps.

—a screen-space penalization factor based on the distance between the projected vertex and the center of screen, representing the fact that more central is the vertex according to the current view, more interesting is the same vertex for the viewer;

—a penalization factor based on the distance between the current view point and the vertex in world coordinates, representing the fact that closest viewpoints indicate a bigger visitor interest in specific features of the sculptures;

—a penalization factor depending on the angle between the direction from the viewpoint to the vertex and the normal at the vertex, representing the fact that surfaces directly facing the view point are more interesting for the user.

Hence, given a vertex and the normal (v_k, n_k) on the surface mesh and the list of view position matrices (P_i, V_i), the focus value ϕ_k is computed according to the following equation:

$$\phi_k = \sum_i max(0, n_k \cdot \frac{d_{ik}}{\|d_{ik}\|})(1 - \|P_i V_i v_k\|^2)^2 (1 - \frac{\|d_{ik}\|^2}{d_{max}^2})^2 \quad (1)$$

where $d_{ik} = e_i - v_k$ is the vector connecting the current view point $e_i = V_i^{-1} O$ to the vertex, and d_{max} is a distance threshold (in our results we considered 3 meters). The focus map can be used to colorize the surface models in a way to highlight the sculpture parts which are considered more interesting for the users. Figures 12(a),12(b),12(c),12(d) show the models colored according to a log-normalized representation of the interest map computed considering all the view position in both museums (from blue to red, passing from cyan, green, and yellow with interpolation factor $\tau = \frac{log(phi_k+1)}{log(phi_{max}+1)}$). As expected, the visitors regions of interest correspond to the most significant features and the finest decorations of the sculptures: specifically, figure 12(a) shows that for the main boxer, visitors are mainly attracted to the features of the face expression, the arm decoration and shield fragments reconstruction. Furthermore, figure 12(b) shows that for the archer visitors are mainly focused on the arch and hand details, as well as the braids, the leg and the metal plate protections, while figure 12(c) shows that for nuraghe model visitors are interested to the towers decorations. Finally, figure 12(d) shows that for the warrior visitors are attracted to the face, and to the vest decorations. We also computed different in-

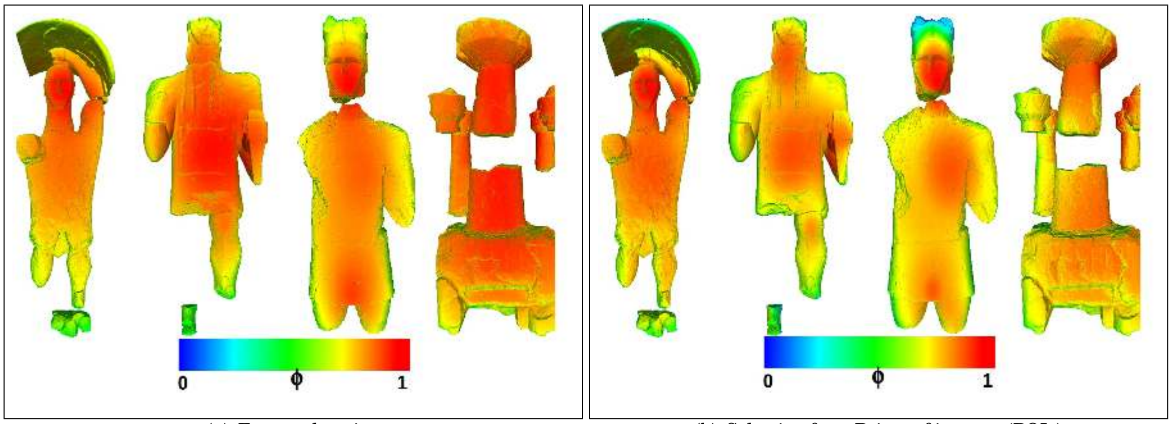

(a) Free exploration

(b) Selection from Points of interest (POIs)

Fig. 13. **Interest maps obtained by separating viewpoints.** From blue to red, blue indicates low interest while red indicates high interest: a) viewpoints from free 3D exploration. b) viewpoints from point of interests selection

terest maps separating the view positions reached through free 3D interaction and the view positions reached through selection of precomputed point of interests. As expected, as highlighted in figure 13, with free manipulation the interest maps appear slightly more uniform, indicating that users employ 3D interface to explore and discover all details of the statues. In any case, with respect to the color peaks representing the main focus parts, no evident differences can be found between the two interest maps in figures 13(a) and 13(b), indicating that the precomputed Points of Interest (POIs) mostly match with the visitors preferences, and that the 3D user interface is intuitive enough to make the most appealing parts of the sculptures easily reachable through free interaction.

Similar behavior considerations can be extracted by visualizing the visitors view positions reached through free 3D exploration and trying to extract main user movements.

Figure 14 shows a selection of the sculpture models (boxer, archer, warrior and nuraghe building) together with all the recorded view positions reached trough free 3D interaction and represented by small frustum tetrahedra. The colors of tetrahedra are chosen with respect to a density based spatial clustering DBSCAN[Ester et al. 1996], which was performed in order to try a preliminary identification of the main behaviors of visitors while exploring the sculptures (for all figures the parameters used for DBSCAN were $\rho = 200.0$ mm and $K = 20$ points). It appears evident how visitors tend to visualize the

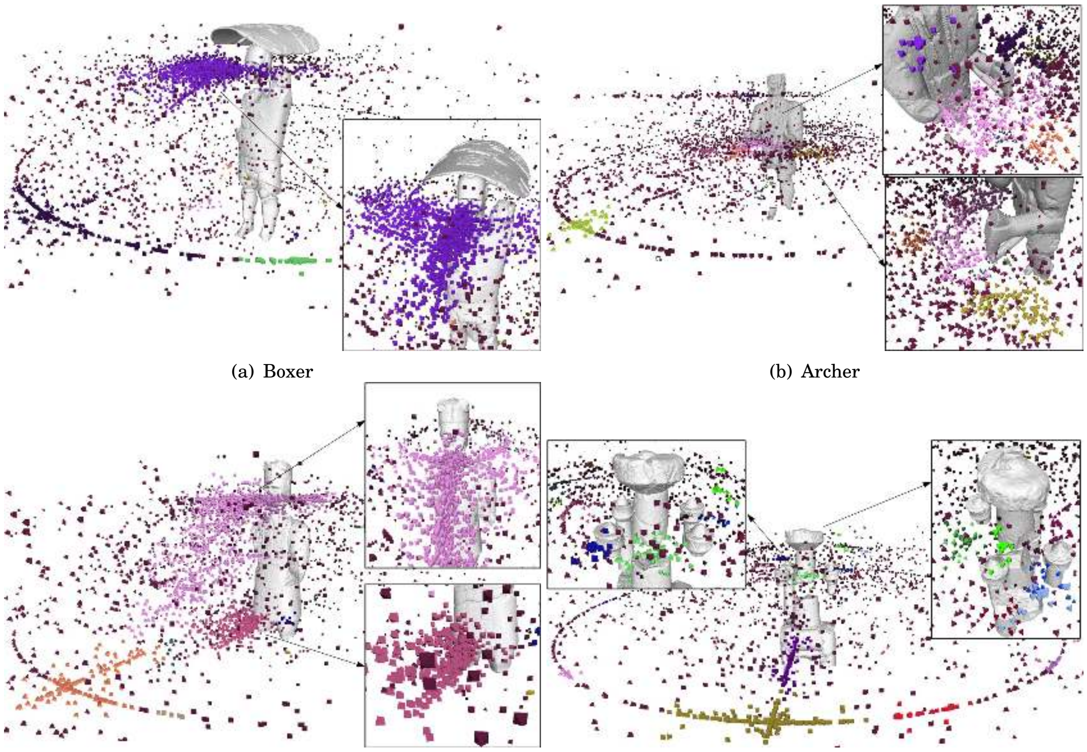

(a) Boxer (b) Archer

(c) Warrior (d) Nuraghe

Fig. 14. **Visualization of view positions.** Insets show the the most interesting sculpture parts according to visitors exploration activity. Main clusters are computed by employing Density-Based Spatial Clustering DBSCAN[Ester et al. 1996], with $\rho = 200.0$ mm and $K = 20$ points.

models from a number of different but limited zoom levels, and how do they focus on the main features of the sculpture: namely the arm decoration and the face in the case of the boxer in figure 14(a), the arm and the arch details, the braids and the protective plate in the case of the archer in figure 14(b), the face and the skirt decoration in the case of the warrior in figure 14(c), and the tower decorations in case of the nuraghe model in figure 14(d). Moreover, it is noticeable the tendency of users to rotate around a specific zoom level in order to enjoy the sculpture from different angles. This is evidently due to the exploring interface with automatic computed pivoting [Balsa Rodriguez et al. 2014].

Discussion. From the reports and logs obtained from permanent exhibitions, the system has been tested extensively by casual people with successful results. A comparison between the visitors attendance at exhibitions and the statistics about the usage of the system shows that the majority of visitors has at least a visual contact with the system (during the days with peaks of visitors the system appears to be used for the whole museum opening time). The analysis of 3D exploration logs reveals that the simplicity of the control metaphor, together with the high resolution rendering permit the visitors to inspect very close details of the objects that would be otherwise hard to observe on the real ones (i.e., small carvings). Additionally, the large display provides the correct sense of scale transmitting to the user the sensation of exploring an imposing artwork. Given the enormous success of the previous exhibitions, other custom versions of the museal setup are currently developed and ready to be installed for permanent exhibitions and for temporary exhibitions in highly relevant national and international Cultural Heritage events.

10. CONCLUSIONS

We have presented a visualization architecture for exploring extremely detailed annotated surface models, which has been developed for the valorization of an extraordinary collection of protostoric Mediterranean sculptures and it has been integrated both in mobile applications and in high end setups. The system successfully integrates network streaming, scalable multiresolution structures, and adaptive rendering techniques.

The rendering framework has demonstrated its scalability by enabling interactive rendering of huge 3D models both on high-end hardware and on embedded systems, with heavy restrictions both on rendering performance and resource availability. The 3D navigation user interface has been tested on a variety of display sizes permitting the user to naturally examine the artifacts, and smoothly transition from analyzing its shape to getting a very close view of a particular detail. Even if the project was focused on the valorization of the Mont'e Prama sculpture complex, either the desktop and mobile implementation have shown scalability on a variety of models and collections of different sizes, including the Gigatriangle datasets of the Digital Michelangelo project [Balsa Rodriguez et al. 2013]. Interactive performance was achieved for all testes datasets, with frame-rates constantly above 30fps on desktop platforms and 10fps on mobile platforms and negligible interaction delays. Similarly to the rendering component, even the interaction component have been demonstrated to be scalable on a variety of Giga-triangle models, as camera transformation computation, and pivot calculation, always took below 10% of the frame time. The user interface has been implemented in large displays as a non co-located touch interface separated from the rendering display in order to provide both full field of view to the user, and low occlusion of the display to other visitors. On mobile devices, instead, the user interface is a co-located touch interface that maps to the well-known touch-based user interface typically present on this devices. Thanks to the automatic centering pivot the interaction with the touch screen can be abstracted allowing the user to focus on the rendered 3D model. Furthermore, a thumbnail-based selector widget allows users to easily discover interesting views of the sculpture collection, decorated with 2D overlays containing information previously authored by archaeology experts and curators. The large display setup has been tested in a variety of museum installations and exhibitions, while a cluster of user tests have been performed for the mobile version. An extensive user analysis has been carried out in museal context on data gathered from thousands of visitors, revealing what are the main preferences, and the main behaviors during the 3D exploration of artworks on large display setups. It appears evident how the system has been extensively used by casual users to inspect and appreciate the artworks at different scales: from general views to very close details that would be otherwise hard to observe on the real sculptures. Our current work focuses on studying new ways to author and provide narrative content to users, as well as exploring the capabilities of autostereoscopic displays.

Acknowledgments

This work was partially supported by the EU FP7 Program under the DIVA project (290277), by RAS under projects BIGDATA and HELIOS, and by *Soprintendenza Archeologia Sardegna*.

REFERENCES

C. Andujar, A. Chica, and P. Brunet. 2012. Cultural Heritage: User-interface design for the Ripoll Monastery exhibition at the National Art Museum of Catalonia. *Computers and Graphics* 36, 1 (2012), 28–37.

Marcos Balsa Rodriguez, Marco Agus, Fabio Bettio, Fabio Marton, and Enrico Gobbetti. 2015. Digital Mont'e Prama: 3D cultural heritage presentations in museums and anywhere. In *Proc. Digital Heritage*. 545–552.

Marcos Balsa Rodriguez, Marco Agus, Fabio Marton, and Enrico Gobbetti. 2014. HuMoRS: Huge models Mobile Rendering System. In *Proc. ACM Web3D International Symposium*. 7–16.

Marcos Balsa Rodriguez, Marco Agus, Fabio Marton, and Enrico Gobbetti. 2015. Adaptive Recommendations for Enhanced Non-linear Exploration of Annotated 3D Objects. *Computer Graphics Forum* 34, 3 (2015), 41–50.

Marcos Balsa Rodriguez, Enrico Gobbetti, Fabio Marton, and Alex Tinti. 2013. Compression-domain Seamless Multiresolution Visualization of Gigantic Meshes on Mobile Devices. In *Proc. ACM Web3D*. 99–107.

Fabio Bettio, Alberto Jaspe, Emilio Merella, Fabio Marton, Enrico Gobbetti, and Ruggero Pintus. 2015. Mont'e Scan: Effective Shape and Color Digitization of Cluttered 3D Artworks. *ACM JOCCH* 8, 1 (2015), 4:1–4:23.

Tamy Boubekeur. 2014. ShellCam: Interactive geometry-aware virtual camera control. In *Proc. ICIP*. 4003–4007.

Marco Callieri, Chiara Leoni, Matteo Dellepiane, and Roberto Scopigno. 2013. Artworks narrating a story: a modular framework for the integrated presentation of three-dimensional and textual contents. In *Proc. ACM Web3D*. 167–175.

Paolo Cignoni, Fabio Ganovelli, Enrico Gobbetti, Fabio Marton, Federico Ponchio, and Roberto Scopigno. 2004. Adaptive tetrapuzzles: efficient out-of-core construction and visualization of gigantic multiresolution polygonal models. *ACM TOG* 23, 3 (2004), 796–803.

Chris Dijkshoorn, Lora Aroyo, Guus Schreiber, Jan Wielemaker, and Lizzy Jongma. 2014. Using Linked Data to Diversify Search Results a Case Study in Cultural Heritage. In *Knowledge Engineering and Knowledge Management*. 109–120.

Maria Economou and Elpiniki Meintani. 2011. Promising beginnings? Evaluating museum mobile phone apps. In *Proc. Rethinking Technology in Museums Conference*. 26–27.

Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise.. In *Kdd*, Vol. 96. 226–231.

Peter Faraday and Alistair Sutcliffe. 1997. Designing effective multimedia presentations. In *Proc. ACM SIGCHI*. 272–278.

Enrico Gobbetti, David Kasik, and Sung-eui Yoon. 2008. Technical Strategies for Massive Model Visualization. In *Proc. ACM SPM*. 405–415.

Timo Götzelmann, Pere-Pau Vázquez, Knut Hartmann, Andreas Nürnberger, and Thomas Strothotte. 2007. Correlating Text and Images: Concept and Evaluation. In *Proc. Smart Graphics*. Berlin, Heidelberg, 97–109.

Antoine Isaac and Bernhard Haslhofer. 2013. Europeana linked open data–data. europeana. eu. *Semantic Web* 4, 3 (2013), 291–297.

Chuljin Jang, Taijin Yoon, and Hwan-Gue Cho. 2009. A smart clustering algorithm for photo set obtained from multiple digital cameras. In *Proc. ACM SAC*. 1784–1791.

Jacek Jankowski and Stefan Decker. 2012. A dual-mode user interface for accessing 3D content on the world wide web. In *Proc. WWW*. 1047–1056.

Jacek Jankowski and Martin Hachet. 2013. A Survey of Interaction Techniques for Interactive 3D Environments. In *Eurographics STAR*.

Jacek Jankowski, Krystian Samp, Izabela Irzynska, Marek Jozwowicz, and Stefan Decker. 2010. Integrating Text with Video and 3D Graphics: The Effects of Text Drawing Styles on Text Readability. In *Proc. ACM SIGCHI*. 1321–1330.

Alexander J Karran, Stephen H Fairclough, and Kiel Gilleade. 2015. A Framework for Psychophysiological Classification within a Cultural Heritage Context Using Interest. *ACM TOCHI* 21, 6 (2015), 34.

Chang Ha Lee, Amitabh Varshney, and David W Jacobs. 2005. Mesh saliency. In *ACM transactions on graphics (TOG)*, Vol. 24. ACM, 659–666.

George Leifman, Elizabeth Shtrom, and Ayellet Tal. 2012. Surface regions of interest for viewpoint selection. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. IEEE, 414–421.

G. Lilliu. 1997. La grande statuaria nella Sardegna nura- gica. *Atti della Accademia Nazionale dei Lincei, Classe di Scienze Morali, Storiche e Filologiche* 9.9 (1997), 281–385.

Fabio Marton, Marcos Balsa Rodriguez, Fabio Bettio, Marco Agus, Alberto Jaspe, and Enrico Gobbetti. 2014. IsoCam: Interactive Visual Exploration of Massive Cultural Heritage Models on Large Projection Setups. *JOCCH* (2014).

João Mota, Manuel J. Fonseca, Daniel Gonçalves, and Joaquim A. Jorge. 2008. Agrafo: a visual interface for grouping and browsing digital photos. In *Proc. ACM AVI*. 494–495.

Nicholas F. Polys, Doug A. Bowman, and Chris North. 2011. The role of Depth and Gestalt cues in information-rich virtual environments. *International Journal of Human-Computer Studies* 69, 1-2 (Jan. 2011), 30–51.

R Core Team. 2013. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Dong-Sung Ryu, Woo-Keun Chung, and Hwan-Gue Cho. 2010. PHOTOLAND: a new image layout system using spatio-temporal information in digital photos. In *Proc. ACM SAC*. 1884–1891.

Henry Sonnet, Sheelagh Carpendale, and Thomas Strothotte. 2005. Integration of 3D Data and Text: The Effects of Text Positioning, Connectivity, and Visual Hints on Comprehension. In *Proc. INTERACT*. LNCS, Vol. 3585. 615–628.

Pedro Szekely, Craig A Knoblock, Fengyu Yang, Xuming Zhu, Eleanor E Fink, Rachel Allen, and Georgina Goodlander. 2013. Connecting the smithsonian american art museum to the linked data cloud. In *The Semantic Web: Semantics and Big Data*. 593–607.

Daniel R Trindade and Alberto B Raposo. 2011. Improving 3D navigation in multiscale environments using cubemap-based techniques. In *Proc. ACM SAC*. 1215–1221.

Carlo Tronchetti and Peter Van Dommelen. 2005. Entangled objects and hybrid practices: colonial contacts and elite connections at Monte Prama, Sardinia. *Journal of Med. Arch.* 18, 2 (2005), 183.