

Data-driven analysis of virtual 3D exploration of a large sculpture collection in real-world museum exhibitions

Marco Agus, King Abdullah University of Science and Technology, Saudi Arabia – CRS4, Italy
Fabio Marton, CRS4, Italy
Fabio Bettio, CRS4, Italy
Markus Hadwiger, King Abdullah University of Science and Technology, Saudi Arabia
Enrico Gobbetti, CRS4, Italy

We analyze use of an interactive system for the exploration of highly detailed 3D models of a collection of protostoric Mediterranean sculptures. In this system, when the object of interest is selected, its detailed 3D model and associated information are presented at high resolution on a large display controlled by a touch-enabled horizontal surface at a suitable distance. The user interface combines an object-aware interactive camera controller with an interactive point-of-interest selector and is implemented within a scalable implementation based on multiresolution structures shared between the rendering and user interaction subsystems. The system was installed in several temporary and permanent exhibitions, and was extensively used by tens of thousands of visitors. We provide a data-driven analysis of usage experience based on logs gathered during a 27 months period at four exhibitions in Archeological museums, for a total of more than 75K exploration sessions. We focus on discerning the main visitor behaviors during 3D exploration by employing tools for deriving interest measures on surfaces, and tools for clustering and knowledge discovery from high-dimensional data. The results highlight the main trends in visitor behavior during the interactive sessions. These results provide useful insights for the design of 3D exploration user interfaces in future digital installations.

Categories and Subject Descriptors: H.5.2 [Information Interfaces And Presentation (HCI)] User Interfaces – Interaction styles, Input devices and strategies; I.3.6 [Computer Graphics] Interaction Techniques; I.3.7 [Computer Graphics] Three-Dimensional Graphics and Realism; I.3.8 [Computer Graphics] Applications

General Terms: Cultural Heritage

Additional Key Words and Phrases: digital heritage, virtual museums, data-driven analysis, user study

ACM Reference Format:

Marco Agus, Fabio Marton, Fabio Bettio, Enrico Gobbetti, Markus Hadwiger. 2017. Data-driven analysis of virtual 3D exploration of a large sculpture collection in real-world museum exhibitions *ACM J. Comput. Cult. Herit.* 0, 0, Article 0 (January 2017), 21 pages.

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

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

Address: Visual Computing Center, King Abdullah University of Science and Technology, Building 1, 23955-6900, Thuwal, Kingdom of Saudi Arabia; email: firstname.secondname@kaust.edu.sa www: <http://vcc.kaust.edu.sa>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2017 ACM 1556-4673/2017/01-ART0 \$15.00

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

1. INTRODUCTION

Digital acquisition technologies, as well as 3D modeling methods, have reached a level of maturity such that highly detailed and accurate 3D representations of artifacts can be created within acceptable times and costs. Applications of this digitization process include communication, archiving, study, restoration, and fabrication of high-quality digital replicas.

Visual communication remains the most common application of 3D digital replicas, and, in recent years, cultural institutions have started to invest in creating high-quality accurate digital contents to be presented as online resources [Potenziani et al. 2015] or through interactive museum installations [Andujar et al. 2012; Marton et al. 2014]. Whereas early uses of high-quality 3D replicas have been the creation of passive media (e.g., movies [Debevec 2005]), the current trends are towards flexible and active presentation modalities, such as virtual systems that allow users to navigate and inspect 3D digital artifacts. Since the overall visit experience tends to be personal, self-motivated, self-paced, and exploratory [Falk and Dierking 2000], these highly interactive presentation approaches centered around high-quality digital replicas are known to engage museum visitors better than passive media, as well as to enhance the overall user experience.

The current proliferation of interactive systems for presenting accurate representations of cultural objects creates the possibility of analyzing the behavior of interacting users to gather feedback for human-centered design of interfaces and for classifying and ordering the 3D information that is presented. However, little research has been done on the behaviors of interacting users in real museum settings.

Here, we analyze the exploration behaviors of thousands of casual visitors interacting with a digital collection of highly detailed 3D representations of protostoric Mediterranean sculptures through large-display visualization systems installed in various exhibitions. The visual presentation system we study is based on a scalable exploration software architecture that supports, in an integrated manner, distribution and rendering of massive, annotated and detailed models with high visual quality. In this system, after the object of interest is selected by the user, its detailed 3D model and associated information are presented at high resolution on a large vertical display controlled by a touch-enabled horizontal surface placed at a suitable distance. The user interface comprises an interactive auto-centering camera controller [Balsa Rodriguez et al. 2014], and a thumbnail-based point-of-interest selecting widget, that are integrated in a scalable system based on multiresolution structures. This system was installed in permanent and temporary exhibitions around Europe (Cagliari, Cabras, Milan, Rome, and Zurich) for exploration of the Mont'e Prama sculpture collection, a large set of extraordinary sandstone sculptures created by the Nuragic civilization in Western Sardinia [Tronchetti and Van Dommelen 2005].

The system anonymously records user interactions, which allowed us to gather a very large body of data. Here we present results from analysis of massive amounts of data acquired during explorations performed by tens of thousands of visitors during 24 months at the National Archaeological Museum of Cagliari, 10 months at the Civic Archaeological Museum of Cabras, six months at the National Pre-historic Ethnographic Museum Luigi Pigorini in Rome, and six months at the Civic Archaeological Museum in Milan. We focus on analyzing the visitor behavior with respect to the exploration interface to understand how interface choices affect user performance, and how visitors use the various interface components. We individuate main visitor behaviors during 3D exploration by employing tools for deriving interest measures on surfaces [Dutagaci et al. 2012], and tools for clustering and knowledge discovery from high-dimensional data [Liu et al. 2015], to codify visual and quantitative analysis methods with large-scale datasets.

This work is an extended version of our contribution to the *Eurographics Workshop on Graphics for Cultural Heritage* [Agus et al. 2016]. We add new material, including an evaluation of the possibility of exploiting user-based interest measures to classify model types, a comparison of user-based interest measures with respect to automatic curvature-based saliency measures [Limper et al. 2016], and a scheme for generating best views and automatic sculpture exploration trajectories from cluster analysis of view transform data.

2. RELATED WORK

In this work, we analyze 3D exploration behavior by users in the context of a museum installation based on a dual display setup that combines guided navigation through thumbnail selection with free viewpoint navigation. Although we focus on interactive system designs, 3D navigation, and interest point analysis, fully covering the related work in these areas is beyond the scope of this paper. We briefly summarize only the most closely related work.

Interactive systems for 3D model exploration. Recently, visualization systems that exploit two display surfaces have been described [Weiss et al. 2010; Wimmer et al. 2010]. These systems employ an horizontal interactive touch table and a large display wall for visualizing the data, sometimes combined in a continuous curved display. Our system layout is similar to the one presented by Coffey et al. [2012], which, however, simultaneously displays a large-scale detailed data visualization and an interactive miniature. We use the horizontal surface exclusively for indirect interaction, such that users keep their focus on the main display of the statue. Multi-touch interaction on the horizontal surface controls camera motion on the large display. Constrained viewpoint navigation has often been proposed as a way to simplify the user interface for novice users [Marton et al. 2012; Marton et al. 2014; Malomo et al. 2016; Boubekeur 2014; Trindade and Raposo 2011]. Our work employs the approach recently presented by Balsa Rodriguez et al. [Balsa Rodriguez et al. 2014] on an auto-centering virtual trackball controlled by decomposition of motion into pan, zoom, and orbit. The virtual trackball is combined with a selection of precomputed viewpoints presented in a thumbnail bar. This dual interface setup allows us to compare interactive behaviors obtained by using a weakly constrained camera controller with those of a fully constrained one.

Interest points detection. The methods dedicated to the detection of interest points on 3D mesh models are relatively recent. Most rely on local surface descriptors, such as curvature, extrema of which are assumed to correspond to candidate interest points. It is common practice to employ a multi-scale approach, in which the algorithm analyzes the 3D surface at successive scales to search for interest points at various levels of detail. Various saliency measures have been proposed to individuate these interest points, including integral volume descriptors [Gelfand et al. 2005], differences in Gaussian-weighted mean curvatures at successive scales [Lee et al. 2005], or difference of gaussians applied to vertex positions [Castellani et al. 2008]. Other saliency measures have been derived by extending popular 2D detection operators, like SUSAN [Walter et al. 2008], or Harris [Sipiran and Bustos 2011], to 3D or by considering the Laplace-Beltrami spectral domain [Hu and Hua 2009; Sun et al. 2009]. Recently, [Song et al. 2014] used the log-Laplacian spectrum to develop a spectral method, in which saliency is determined via deviations from a local average in the frequency domain, whereas Limper et al. [Limper et al. 2016] considered the information theory to classify regions on a 3D surface according to the local curvature entropy. Finally, Dutagaci et al. [Dutagaci et al. 2012] proposed an evaluation strategy based on human-generated ground truth to measure the performance of 3D interest point detection techniques. We propose here a novel metric for computing interest-based measures based on user exploration activity, that aims to identify human-centred saliency points on 3D representations of works of art. A comparison of our user-based metric with respect to the curvature-based mesh

saliency method [Limper et al. 2016] revealed that automatic methods, even if accurate in detecting most of salient features, are not efficient in dealing with archeological artifacts, since they cannot distinguish between details of the sculpture and degraded parts of the sculpture that do not contain meaningful information.

Visual analytics of exploration activity. Using visual analytics techniques in cultural heritage applications is a relatively recent field of research [Patterson et al. 2014] that has already provided interesting results on architectural documentation [De Luca 2014], investigative analysis across documents and drawings [Deufemia et al. 2012], analysis of wall painting degradations [Zhang et al. 2013], as well as risk assessment of cultural heritage sites [Qian et al. 2015]. With respect to analysis of user behaviors, mining methods have been applied to tracing tourist activity for discovering landmark preferences from photo maps [Jankowski et al. 2010], recommending travel routes using geotags [Kurashima et al. 2010], and analyzing routing preferences in tourist areas [Torrise et al. 2015]. Here, we analyze the virtual 3D exploration activity of museum visitors inspecting a digital collection of sculptures, by considering histogram analysis of parametrization of view transforms and by employing classical cluster analysis methods [Jain 2010].

3. OVERVIEW

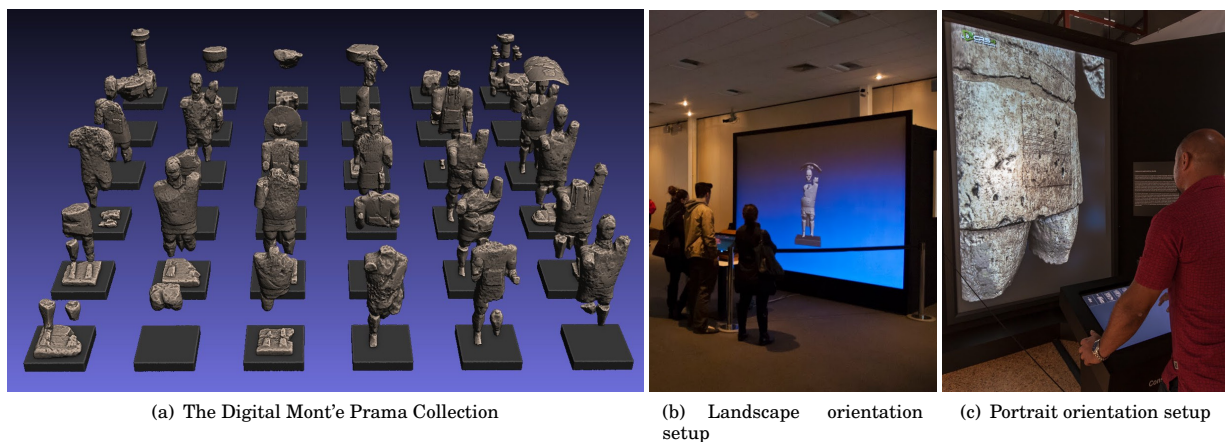


Fig. 1. **The Digital Mont'e Prama setup.** Left: the collection consists in 37 colored surface models at a resolution of 16 points / mm². [Bettio et al. 2015; Bettio et al. 2013]. Center: back-projection-based setup in landscape orientation (Rome and Milan exhibits); Right: back-projection-based setup in portrait orientation (Cagliari and Cabras exhibits). Additional setups based on monitors have been employed in the Milan/EXPO and Zurich exhibits.

The methods and results presented in this paper are outcomes of a broad cultural heritage digital acquisition and presentation project, motivated by the valorization of the Mont'e Prama sculpture complex, a large set of extraordinary sandstone sculptures created by the Nuragic civilization in Western Sardinia [Tronchetti and Van Dommelen 2005].

3.1 Digital sculpture collection

The collection comprises 25 life-size human figures varying in height between 2 and 2.5 meters, depicting archers, boxers and warriors, and 13 approximately one-meter-sized building models representing typical Nuragic towers. The sculptures were restored by non-invasive techniques starting from reassembling five thousands stone fragments. The collection contains various distinctive elements that make the sculptures attractive examples of human creativity in Mediterranean protostoric

civilizations. They are currently accessible to the public in the National Archaeological Museum of Cagliari (19 human sculptures and 9 Nuragic towers), and the Civic Museum Giovanni Marongiu of Cabras (6 human sculptures and 4 nuraghe models). The digital valorization project consisted of 3D high-resolution digitization of the complex, resulting in 37 quarter-millimeter resolution colored surface models [Bettio et al. 2015; Bettio et al. 2013] (see Fig. 1). We also developed a visualization architecture for distribution and real-time interactive exploration of the models on commodity platforms for museum exhibitions.

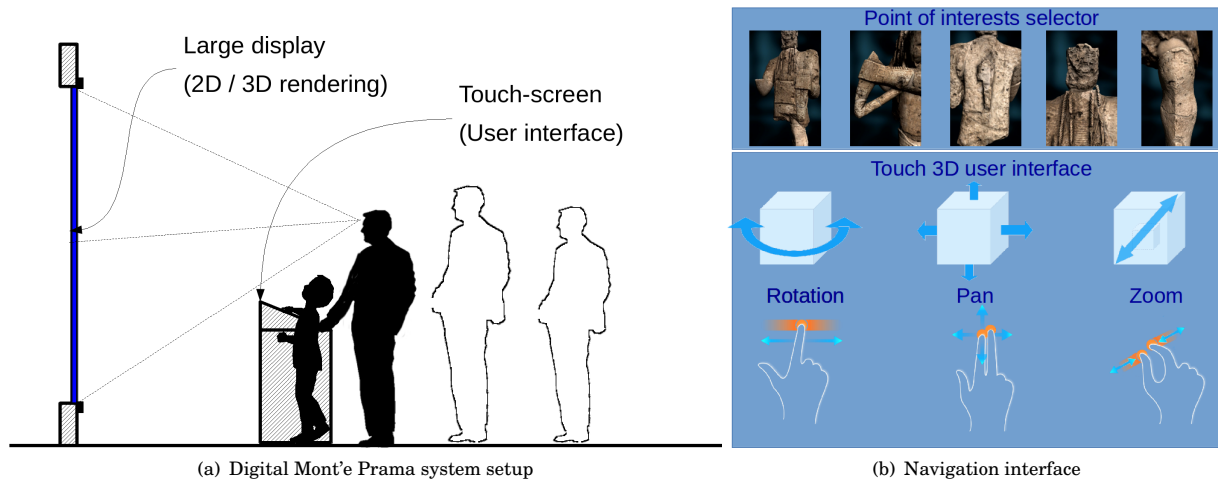


Fig. 2. **Digital Mont'e Prama Setup.** Left: the general concept of the application setup, with a large screen controlled by a touch-enabled surface placed in front of the screen. Right: the navigation interface, composed by an object-based assisted system based on an auto-centering virtual trackball [Balsa Rodriguez et al. 2014] and a thumbnail-based selection component for navigating among a series of pre-defined viewpoints [Balsa Rodriguez et al. 2015]

3.2 Visual exploration platform

The sculpture collection was presented in museums using an approach based on indirect touch-based interactive control of a large display surface through an assisted user interface (see Fig. 2 left). The use of a touch device for the user interface favored the *walk-up-and-use* approach, since users are nowadays accustomed to the now ubiquitous touch screens. However, instead of using just a large touch screen as the input/output device for direct manipulation, we preferred to decouple the devices for interaction and for data presentation, to allow for very large projection screens (see Fig. 1 center and right), and to enable multiple users to watch the whole screen without occlusion to permit group visits during periods in which exhibitions was particularly crowded.

The touch device guiding the user interface was placed at a distance far enough from the display to grant the user a full view of the display, which was large enough to display the statues at their natural scale.

The user interface integrated a sculpture selection system, which allowed the user to select the statue to be inspected, and a 3D navigation system which provided ways to explore a single statue in 3D. In order to achieve real-time performance in the exploration of statues, the rendering subsystem employs a scalable adaptive multiresolution rendering approach. The technique is based on a compressed version of the Adaptive TetraPuzzles framework [Cignoni et al. 2004; Balsa Rodriguez et al. 2013], implemented on top of a deferred rendering system that enhances visualization through screen-space ambient occlusion and radiance scaling [Vergne et al. 2010].

Sculpture selection was accomplished by presenting a two-level hierarchy of thumbnail-based images, displayed on the large display and controlled through a standard scroll-and-select approach. The first level of hierarchy was a scroll bar containing images of the four statue categories: boxers, archers, warriors, and building models (nuraghe). The second level of hierarchy contained, for each group, the thumbnail images associated with each sculpture 3D model. Other actions, such as viewing illustrative videos could also be executed through this menu system.

With the selection of a sculpture, a 3D exploration application was launched. The main functionality provided by this application was detailed 3D exploration of a high-resolution model of the statue. During exploration, the 3D model was displayed on the large display, while the console with the touch screen was used to control the position of the camera.

Two different navigation systems were provided: an object-based assisted system based on an auto-centering virtual trackball [Balsa Rodriguez et al. 2014] and a thumbnail-based selection component for navigating among a series of pre-defined viewpoints [Balsa Rodriguez et al. 2015] (see Fig. 2 right).

The continuous camera controller was tailored for continuous navigation. A single finger panning gesture was used to control rotation around an automatically positioned barycenter of the visible part of the models [Balsa Rodriguez et al. 2015]; two-finger panning controlled camera panning in the image plane, and a two-finger pinching controlled zooming (see Fig. 2 right). The up-vector remained fixed during camera motion.

The continuous camera controller was complemented by an interface that offered the possibility of selecting predefined thumbnails. The user could navigate by clicking on the various thumbnails, thus starting an animation that would take the observer from the current position to the target view point. The approach was extended with static overlays, that gave annotations and presenting heterogeneous enriched multimedia information [Marton et al. 2014].

Table I. **Digital Mont'e Prama museum setups:** main characteristics of the permanent and temporary exhibitions.

Museum	Exhibition	Period	Display orientation	Display size
National Museum (Cagliari)	Permanent	Mar 2014 - now	Portrait	2.5 m height
Museum Marongiu (Cabras)	Permanent	Mar 2015 - now	Portrait	2.5 m height
Museum Pigorini (Rome)	Temporary	Nov 2014 - Apr 2015	Landscape	3m x 2.5m
Civic Museum (Milan)	Temporary	May 2015 - Feb 2016	Landscape	3m x 2.5m
Universal EXPO (Milan)	Temporary	11 - 17 Sep 2015	Portrait	85 inch
Arch. Museum (Uni. Zurich)	Temporary	Apr 2016 - Sep 2016	Portrait	85 inch

4. USER ANALYSIS OVERVIEW

The setup described in Section 3 was customized to be installed in various temporary and permanent exhibitions (see Fig. 2 and Table I). The two permanent installations, at the National Archaeological Museum in Cagliari and the Civic Museum G. Marongiu in Cabras, are placed in the same rooms that host the sculpture collection (28 statues in Cagliari and 10 in Cabras). In this paper, we analyze the data recorded by the navigation system in the permanent exhibitions in Cagliari and Cabras, built around a portrait-mode projection system, as well as the two temporary exhibitions in Rome and Milan, built around a landscape-mode projection system. Many tens of thousands of people used the system, resulting in a very large number of exploration sessions (see Table II). The focus of our analysis is on the 3D exploration behavior of users, and on synthesis of novel information from large amounts of 3D exploration data.

Table II reports the number of exploration sessions and the total usages time. On average visitors actively used the system for exploring statues for more than 2 hours a day (over one hundred explo-

Table II. **Virtual exploration in museum setting:** main system usage statistics for the permanent and temporary exhibitions.

Museum	Exhibition	Acquisition Period	# Explored sculptures	Usage time (hours)
National Museum (Cagliari)	Permanent	Mar 2014 - Nov 2015	43943	816.6
Museum Marongiu (Cabras)	Permanent	Mar 2015 - Oct 2015	23537	416.3
Museum Pigorini (Rome)	Temporary	Nov 2014 - Apr 2015	2879	42.5
Civic Museum (Milan)	Temporary	May 2015 - Feb 2016	4378	61.3

ration sessions/day). This usage time does not include the time used for navigating the two-level menus and for selecting sculptures.

4.1 Qualitative evaluation

To understand visitor behaviors and their impressions of the 3D exploration system, we made observations and recording of users and we conducted informal interviews to collect opinions and suggestions. In general, we observed two main usage scenarios. The first was individuals or small groups. A single visitor, or a small group composed of a family and group of friends interacted with the system. In general one person, normally the most confident with touch-based interface and 3D interactions would take the control of the console, and the rest of the group would observe the screen and eventually give instructions on which sculpture to display and which detail of the model to focus on. The second usage scenario comprised medium/large groups. Given the number of visitors in the exhibition, heterogeneous groups could naturally form in front of the interactive station. In this case, one person took control of the console, and the rest of the group observed the screen. In many cases, the person controlling the console was a museum guide who used the system to describe the history and main features of the digital collection.

4.2 Usage data gathering

To understand the behavior of visitors beyond these basic usage scenarios and to extract large amounts of 3D navigation information, we added a log to capture interactive 3D exploration system usage. Given the long durations of the exhibitions and the high-frequency use of the system, we limited data collection to the following events and associated times :







- (1) the starting time of an interactive exploration of a sculpture;
- (2) the ending time of an interactive exploration of a sculpture. Because the 3D exploration application automatically quits when it is idle for more than 30 seconds; note, the ending time corresponds to either the time of the explicit closing action by a user, or to the last interactive event before the automatic quit;
- (3) the view positions selected by the user from the list of preselected points of interest;
- (4) the view positions at which the user stayed for at least 5 seconds during exploration. This interval was selected to be long enough to leave out the locations where the user briefly stops during an interaction without the intention of actually looking at the statue. It should be noted that, while this could slightly bias the identification of interest position by favoring users making frequent short stops, considering the look-at point during movement as “interesting” would bias the analysis towards places navigated by users making slow motions (i.e., very novice ones).

4.3 General statistics

From the log files containing the data collected during the observation periods, we gathered general statistics about the use of the system. We focus on analyzing the visitor behavior with respect to the 3D exploration interface. Specifically we seek to understand how interface choices affected user

performances, and how visitors used the various interface components. In addition, we also investigate how to extract information from large numbers of recorded 3D views to detect the best views, the most salient areas of the statues, and the most meaningful exploration paths. Table III reports the main

Table III. **Main statistics on the most visited sculptures.** For each museum installation, we indicate the number of times the specific statue was explored with the system (V) as well as the average exploration time in seconds.

						
Cagliari	V : 9479 T(s) : 98.4	V : 4254 T(s) : 79.53	V : 2912 T(s) : 71.56	V : 2892 T(s) : 61.52	V : 2851 T(s) : 73.41	V : 1694 T(s) : 42.47
Cabras	V : 5037 T(s) : 97.61	V : 2827 T(s) : 74.71	V : 2315 T(s) : 67.01	V : 1478 T(s) : 56.12	V : 2142 T(s) : 65.73	V : 1536 T(s) : 45.5
Rome	V : 543 T(s) : 70.48	V : 272 T(s) : 64.94	V : 239 T(s) : 60.52	V : 284 T(s) : 52.73	V : 218 T(s) : 64.59	V : 142 T(s) : 30.65
Milan	V : 986 T(s) : 69.06	V : 399 T(s) : 52.93	V : 367 T(s) : 57.9	V : 386 T(s) : 52.39	V : 310 T(s) : 51.61	V : 223 T(s) : 29.27

statistics on the most explored sculptures during each museum installation. Specifically, it reports the number of times each sculpture is explored (V) and the average exploration time. Of the six top most explored sculptures, all categories of the sculpture collection were represented (boxers, archers, warriors, and nuragic towers). The average exploration time varies from 30 seconds for the Nuragic tower to 70 seconds and more for the most decorated boxer and archer sculptures.

5. INTEREST ANALYSIS

We sought to understand which parts of the sculptures were considered the most appealing for visitors using the 3D exploration system.

In our case, in order to have an immediate visual representation of visitors interest to be mapped to the sculpture models, we extended the interest function defined by [Balsa Rodriguez et al. 2016]. Specifically, we define a **focus map**, in which, for each vertex v_k of a surface model, a value ϕ_k is computed with respect to the visitor view points represented by projection and view matrices P_i and V_i . Specifically, given a ray d_{ik} connecting the current viewpoint $e_i = V_i^{-1}O$ and the vertex v_k , the following quantities were considered:

- (1) **World coordinates distance:** a factor δ_{ik} based on the distance $\|d_{ik}\|$ in world coordinates, capturing the idea that closest viewpoints indicate greater visitor interest in specific features of the sculptures (in our settings d_{max} is 3 meters)

$$\delta_{ik} = \left(1 - \frac{\|min(d_{max}, d_{ik})\|^2}{d_{max}^2}\right)^2 \quad (1)$$

- (2) **Screen coordinates distance:** a screen-space factor σ_{ik} based on the distance between the projected vertex and the center of the screen, capturing the idea that the more central the vertex v_k

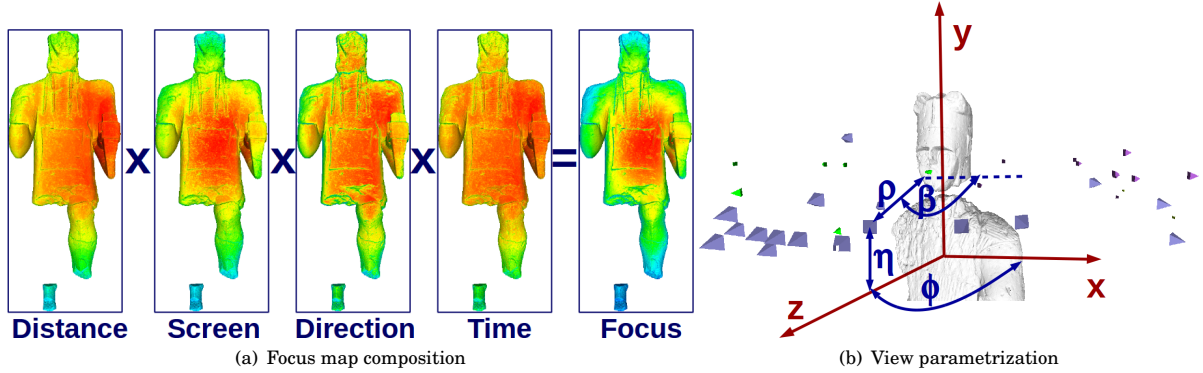


Fig. 3. **User exploration analysis.** Left: a focus value is computed by multiplying various contributions depending on vertex position and the viewpoints from which it is observed. We computed penalization factors depending on the distance from the vertex to the viewpoint, the screen projected distance, and orientation and observation duration are considered. The colors are computed by employing the standard jet color map with respect to a log-normalized scale of contribution values. Right: Recorded view transformations obtained with an auto-centering virtual trackball are parametrized by considering a normalized distance ρ , a normalized height η , an orientation angle ϕ and a view direction angle β on the horizontal plane.

is according to the current view, the more interesting the surface element for the view position represented by transform $P_i V_i$ is:

$$\sigma_{ik} = (1 - \|P_i V_i v_k\|^2)^2 \quad (2)$$

- (3) **Orientation:** a factor ν_{ik} depending on the angle between the direction $\hat{d}_{ik} = \frac{d_{ik}}{\|d_{ik}\|}$ from the viewpoint to the vertex and the normal at the vertex \hat{n}_k , representing the fact that surface portions directly facing the view point are more interesting for the current viewpoint

$$\nu_{ik} = \max(0, \hat{n}_k \cdot \hat{d}_{ik}) \quad (3)$$

- (4) **Time:** a factor depending from the duration t_i of the current observation (t_{max} is the maximum duration, and set to 20 seconds), representing the fact that more time a view position is kept, more interesting is that view

$$\tau_i = (1 - \frac{\min(t_{max}, t_i)^2}{t_{max}^2})^2 \quad (4)$$

The focus value for all sculpture mesh vertices is computed by accumulating the quantities coming from all view points through ray casting and by simulating depth buffering in order to avoid considering invisible vertices. Hence, for each vertex v_k , the focus value is


$$\phi_k = \sum_i \delta_{ik} \cdot \sigma_{ik} \cdot \nu_{ik} \cdot \tau_i. \quad (5)$$

Figure 3 left shows how the various quantities create a focus map. In this case, the colors are computed by employing the standard jet color map with respect to a log-normalized scale of values (with $s = \frac{\log(\phi_k+1)}{\log(\phi_{max}+1)}$).

5.1 Museum installations and user interface analysis

We used the focus maps to find differences in visitor behaviors at the various museum installations and between the two user interfaces employed in the interactive stations (the 3D exploration interface and the image-based points-of-interest selection interface). For brevity, we restricted our analysis to the focus maps of the four most explored sculptures, including a boxer, an archer, a warrior and a Nuragic

Table IV. **Comparison of focus maps for the most explored sculptures.** From top to bottom, the boxer, the archer, the warrior and the Nuragic tower. From left to right, the maps computed for the exhibitions of Cagliari, Cabras, Milan and Rome, separated in free 3D exploration and points of interest image-based selection.



tower. We found that most appealing features of the most explored sculptures did not depend on the user interface or the location of the exhibition. Table IV shows the different focus maps computed for the different sculptures: from top to bottom, the boxer, the archer, the warrior and the Nuragic tower, separated by views reached through free exploration and through image selection. There are slight differences between the maxima highlighted in the various focus maps, indicating that the most significant features of the sculptures attracted users in all exhibitions and that these features were reached easily and in a similar way through both user interfaces. The differences between the focus maps computed for the permanent exhibitions (Cagliari and Cabras on the left side) and the temporary exhibitions (Milan and Rome on the right side) are due to the significantly different number of views

used in the computation. As expected, the most distinctive elements attracted visitors. These included the schematic faces of the boxer and the warrior, as defined by their heavily rendered eyebrows and T-shaped noses as well as their magnetic eyes created by two nested circles, the archer's plate shield protection and bow and arm decoration, the warrior's vest, and the decorations on the Nuragic tower.

5.2 Sculpture focus maps

We expected that “similar” statues, in terms of kind, shape, and color, would have similar exploration patterns. We therefore considered all data gathered from all the exhibitions and we computed the complete focus maps according to equation (5) for all the sculpture models. For visual presentation, focus values were mapped to a log-normalized scale according to a scheme chosen in the Color Brewer system [Harrower and Brewer 2013], ranging from very light gray to dark magenta. Fig. 4 shows the focus maps computed in this way for all the 37 models composing the Digital Mont'e Prama collections over all gathered data. Inspection of these focus maps revealed the main points of interest, leading to a preliminary comparison of the sculptures of the collection. We inferred that human figures follow similar focus patterns: generally the main focus is on face details, followed by trunk of the body, as well as vest and weapon decorations. Similarly, Nuragic towers also follow similar focus patterns: the main focus in this case is on the main columns or decorations. Other insights were gathered by considering the restoration and completeness of the sculptures. All human figures with restored heads attracted focus on facial details. If the heads were not restored, users focused on other features, including weapons, vest decorations, arms or other anatomical details. Interestingly we noticed that users focused on vest details of boxers, whereas they focused on the protective armor shield of archers. Finally, incomplete sculptures missing important anatomical parts had uniform focus maps, indicating that users did not focus on particular areas of these statues. As future work, we plan to explore the focus maps in the context of automatic similarity assessment methods. They could integrate, for instance, the photometric channel in methods such as the one of Biasotti et al. [2015], in order to provide hints on the underlying semantics of the statues.

5.3 Comparison with curvature-based saliency maps

In order to evaluate the characteristics of usage-based focus maps, we carried out a comparison with respect to state of the art curvature-based mesh saliency methods, which are commonly used to guide mesh simplification and automatic viewpoint computation. To this end, we considered a recent and popular technique based on local curvature entropy (LCE) [Limper et al. 2016]. We compared LCE with our interest-based focus maps of the nine most explored and significant sculptures of the Mont'e Prama collection, employing the Color Brewer scheme ranging from very light gray to dark magenta on a log-normalized scale (the same normalization functions and color maps were used for both scalar functions). Table V presents results of the comparison between the two scalar functions, together with the difference between the two values mapped to a color scheme ranging from blue to very light gray to blue. In general, we found significant differences between the two schemes. In a number of cases, the curvature-based LCE method interprets sculpture degradations or defects as main saliency points, which instead are not considered as focus points by users. On the contrary, features that are considered extremely interesting by users (like facial details, protective plates, leg and vest decorations), are not clearly detected by the automatic LCE method. From this evaluation, we can infer that automatic curvature-based methods suffer in the presence of surface degradations, and they are not particularly effective in dealing with archaeological artifacts that contain many defects or missing parts.





Fig. 4. **Focus maps of all sculptures for visual comparison.** The color map is based on the Color Brewer system [Harrower and Brewer 2013], ranging from light gray to dark magenta. The main focus area is indicated below the name of each statue.

6. 3D EXPLORATION ANALYSIS

The analysis of focus maps provides an immediate visual tool to capture visitor preferences with respect to various parts of the sculptures, but it does not provide information on visitor behavior or on the

Table V. **Focus maps comparison with curvature-based mesh saliency (LCE).** The nine most viewed sculptures are considered (at left the focus map, at center the LCE, at right the difference).



main exploration paths they follow. The visitor view positions should be analyzed as high-dimensional data to identify position clusters and to classify fundamental motions. To this end, we performed a visual analysis of the clouds of visitor positions by applying classic visual analytics techniques employed for knowledge discovery and data mining in various application domains [Liu et al. 2015]. Our goal was to find differences and connections between the two user interfaces employed in the interactive stations, and any differences in behaviors with respect to the exhibitions. We carried out this analysis on the two most-explored sculptures in the collection, the boxer and the archer.

6.1 View transform parametrization

Since the 3D user interface employed for free exploration in all interactive stations was composed of a servo-assisted virtual trackball with automatic pivot computation [Balsa Rodriguez et al. 2014], and with fixed up direction, , we parametrized the recorded view transformations V_i by reducing the number of degrees of freedom to four components:

- (1) a distance, ρ_i , from the view position e_i , to the pivot on the surface, computed by employing the same stochastic sampling considered in the Automatic-Center Virtual Trackball (ACeViT) interface (the normalization was carried out with respect to the double diagonal R of the bounding box of the sculpture model) [Balsa Rodriguez et al. 2014];
- (2) a normalized height, η_i , with respect to the height, H, of the bounding box of the sculpture model;
- (3) an orientation, ϕ_i , on the horizontal plane, computed according to the cylindrical coordinates of the view position, e_i ;
- (4) an orientation, β_i , on the horizontal plane, representing the view direction.

A schematic representation of the view parametrization is shown in right side of figure 3.

6.2 Histogram analysis

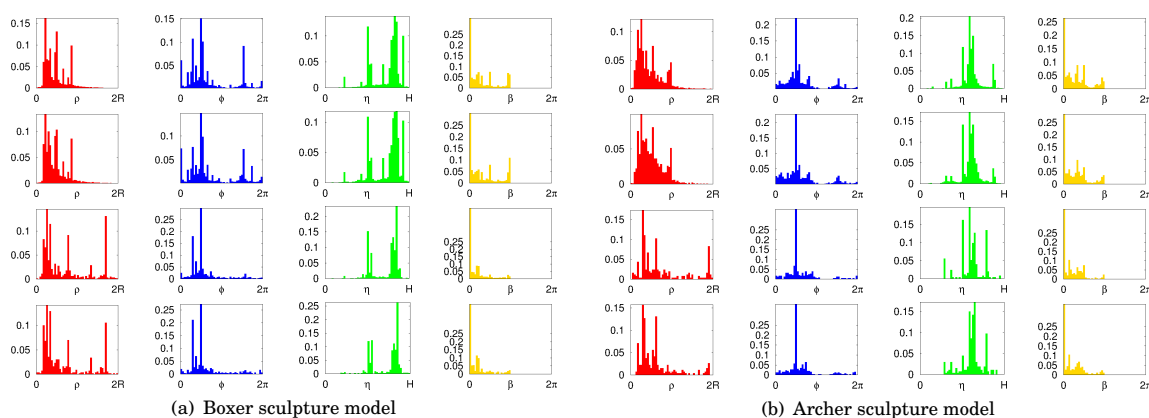


Fig. 5. **Frequency histograms of parametrized views.** From top to bottom histograms of view parameters for exhibitions in Cagliari, Cabras, Milan, and Rome. From left to right, the represented parameters are the distance, ρ , the orientation, ϕ , the height, η , and the view direction orientation, β .

In order to individuate the main behaviors of visitors during interactive exploration, we computed and analyzed frequency histograms of the view parameters. Figure 5 shows the comparison with respect to the various installations of parameter histograms for the boxer and archer sculpture. From left, the first columns in red show the distance, ρ , histograms. They distinguish three main peaks corresponding to the three main exploration scales: from macro-structure to meso-structure to micro-structure. Differences in distance histograms between the permanent exhibitions (Cagliari and Cabras in the two top rows) and the temporary exhibitions (Milan and Rome in the two bottom rows) are due to the different display setups (portrait display orientation versus landscape display orientation) leading to different exploration workspace sizes. The second columns in blue show the orientation, ϕ , histograms and highlight that visitors mostly prefer to explore the sculptures frontally ($\phi = -\frac{\pi}{2}$), even if a smaller peak can be individuated in the back ($\phi = \frac{\pi}{2}$). The third columns in green show the height, η , histograms and highlight how visitors tend to explore the top parts of the sculptures since they contain the most distinctive features: face decorations and shield in the case of the boxer, plate decoration, breeds and bow in the case of the archer. Finally, the histograms in yellow represent the view direction orientation, β . Apart from the favorite straight direction, various small peaks provide indicate the clusters around the positions of interest.

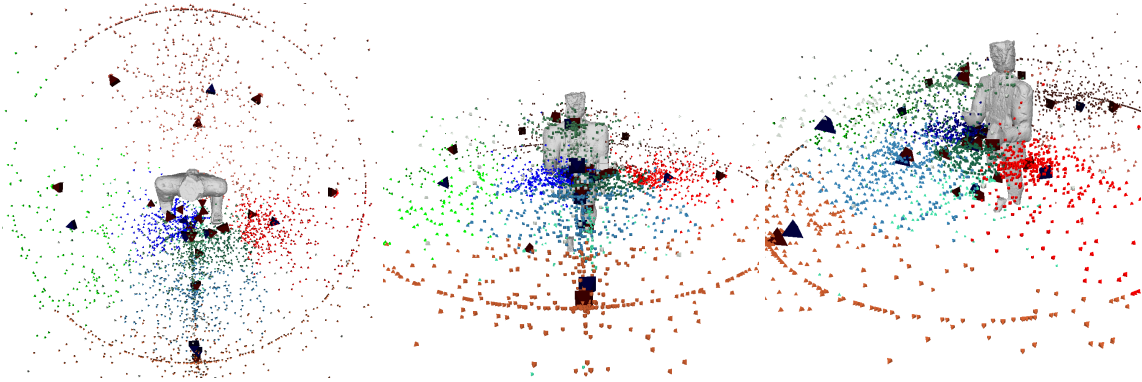


Fig. 6. **View position clusters for the archer in the permanent installation at Cagliari.** From left to right, top, front and side views. Clusters are computed with the K-means algorithm [Jain 2010] by considering 10 clusters with the help of the main peaks of histograms in figure 5. Cluster centroids are indicated in dark blue. Comparisons to the precomputed point-of-interest view positions employed in the thumbnail-based selection interface are indicated in dark red.




Fig. 7. **View positions clusters for the boxer at the permanent installation at Cabras.** From left to right, top, front and side views. Clusters were computed with K-means algorithm [Jain 2010] by considering 10 clusters with the help of the main peaks of histograms in figure 5. Cluster centroids are indicated in dark blue. Comparisons to the precomputed point-of-interest view positions employed in the thumbnail-based selection interface are indicated in dark red.

6.3 Clustering view transformations

The histogram analysis indicates that visitors views tend to cluster around the points of interest. To highlight the main visitor behaviors, we carried out a supervised cluster analysis, by employing the K-means algorithm [Jain 2010] over the view transforms according to the same parametrization employed in the histogram analysis. Since we noticed similar behaviors in different installations, we report here as examples the results for the archer sculpture in the installation at Cagliari (Figure 6), and the boxer sculpture in the installation at Cabras (Figure 7). Since our goal is to find relations and differences between the two interfaces employed in the interactive stations, we compare the exploration cluster centroids with the view transforms of the points of interests employed for the image-based select interface. Specifically, figures 6 and 7 show the top, front, and side views of the clustered view transforms rendered as small pyramids, together with the cluster centroids rendered as larger dark blue pyramids and the precomputed points of interest rendered as larger dark red pyramids. Both figures show that most of the cluster centroids are very close to the precomputed points of interest, confirming the visual impression derived from focus maps in Table IV. Furthermore, the densest clus-

ters match with the maxima in the focus maps (the protective plate and the bow in the case of the archer, and the face in the case of the boxer). Finally, circular patterns are perceptible both at macro and meso scale, confirming that visitors have the tendency to explore the sculpture by rotating around it in order to get an overall idea of the shape.

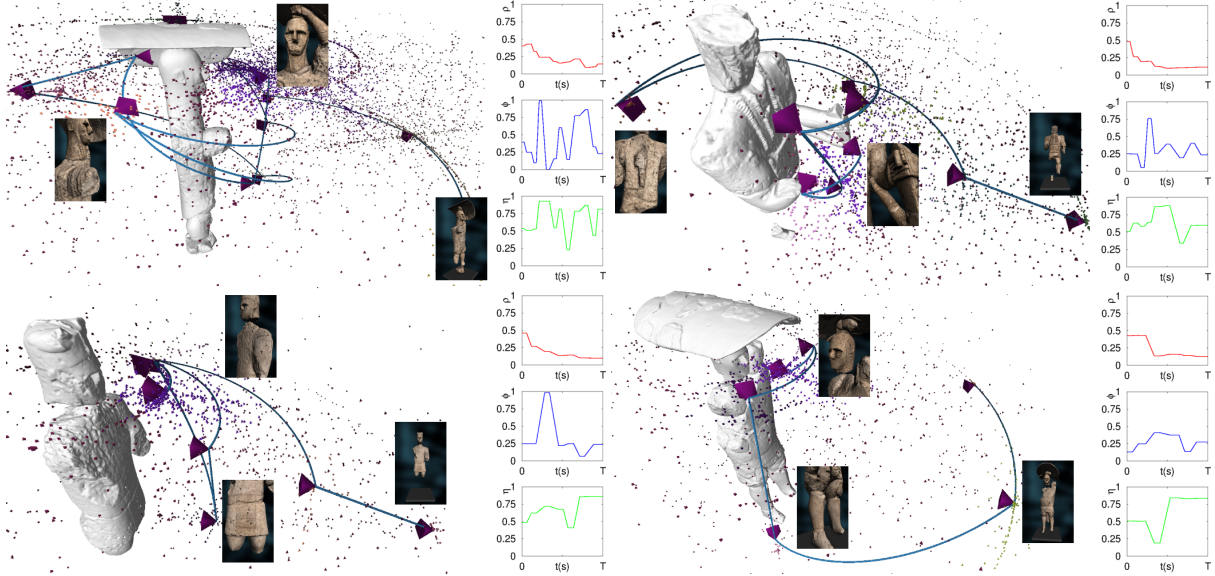


Fig. 8. **Automatic exploration trajectories.** Example of automatically computed trajectories obtained after cluster analysis from exploration data gathered in the museum of Cagliari for the boxer 16 (top left), archer 5 (top right), warrior 3 (bottom left), and boxer 15 (bottom right) sculptures. Time charts of view parameters (normalized distance ρ in red, normalized angle ϕ in blue, normalized height η in green).

6.4 Automatic trajectory generation

An important application of viewpoint analysis is the automatic generation of trajectories for inspecting the models. To this end, various schemes for deriving exploration paths can be employed, and different criteria can be exploited to mimic user inspection activity.

Since cluster analysis already identifies viewpoints that are considered important by users, as they are created from the most popular viewpoints where users statically looked at the models for extended times, we explore an approach that creates navigations by smoothly connecting the barycenters of view clusters with animation paths.

We considered as main guidelines for trajectory creation the parameters derived from clustering the view positions. We observed that visitors inspect the sculptures with a generally monotonically decreasing distance, from macro-scale to micro-scale resolution. Specifically, starting from a set of view position clusters and given an exploration path duration, T , which according to statistics derived from gathered data can vary from 30 to 90 seconds (see Table III), we generate the trajectory as a sequence of time intervals, t_{if} , in which users look at the statue from a given position represented by the center of the clusters C_i , and time intervals, t_{im} , in which users move linearly from the focus position C_i to the next position C_{i+1} , in a way to satisfy the following relationship:

$$\sum_i t_{if} + t_{im} = T_f + T_m = T. \quad (6)$$

The single time intervals are computed with respect to the cluster parameters; the focus time intervals t_{if} are proportional to the count of view positions contained in the cluster, while the transition time intervals t_{im} are proportional to the density of the two clusters which are to be connected, meaning that the less sparse are the two clusters, the slower is the transition between the represented positions.

The animation track starts from the viewpoint with the largest focus distance and visits the others to minimize the overall cost of the visit, defined as the sum of the costs associated with each individual segment. We heuristically compute the cost of a connecting segment by $c = c_f \cdot c_d$, where $c_f = \text{median}(0.25, 4.0, f_{origin}/f_{destination})$ is a weighting factor that increases if the focus distance f to the model decreases when moving from the origin viewpoint to the destination viewpoint, and $c_d = \text{distance}(origin, destination)$ is the linear distance from origin to destination. Only segments that do not intersect the model are considered in the optimization. Even though we plan in the future to exploit more elaborate solutions for finding smooth collision free curved paths, as proposed, e.g., by [Di Benedetto et al. 2014], we have not found problems with this simple solution in the database, mostly because navigation occurs reasonably far from the surface and many collision-free linear paths are available.



Fig. 9. **Automatic trajectories.** Frame sequences extracted from automatic computed trajectories in Fig. 8. From top to bottom: boxer 16, archer 5, warrior 3, and boxer 15.

Examples of trajectories generated in this way are provided in Figure 8. In this case, paths are computed for the main human figures in the collection (boxer 16, archer, warrior, and boxer 15) of the collection on clusters computed on exploration data gathered on the Archeological Museum of Cagliari according to DBSCAN algorithm [Ester et al. 1996]. For each trajectory, the time plots of the parameters computed according to the same view transform parametrization employed for histogram analysis are provided: normalized distance to surface, ρ , in red, normalized angle, ϕ , in blue, and

normalized height, η , in green. Figure 9 contains instead a selection of significant frames extracted from trajectories represented in Figure 8. These automatically generated trajectories confirm the tendency of users to explore the models at different scales and to focus on the distinctive features of the sculptures.

6.5 Discussion

From the focus map analysis and the 3D exploration analysis, we were able to identify the following visitor behaviors during the interactive virtual exploration of the collection of digital sculptures:

- (1) **Uniform exploration workspace:** especially for the most explored statues, there is the tendency of users to freely inspect the models by using different scales and by performing circular movements around the distinctive features. For example, the head or the trunk of human figures (see Figure 6). We also noticed the tendency for exploring the front side of the sculptures, and using the back side mostly for macro-scale and meso-scale inspection. Similar behaviors were observed on most sculptures. Information about the exploration workspace of sculptures can be also employed for deriving optimal exhibition layouts.
- (2) **Influence of display orientation:** the analysis of the distance histograms revealed significant differences with respect to the exhibition due to the display orientation (portrait orientation in the permanent exhibitions versus landscape orientation in the temporary exhibitions). We infer that portrait displays are more adequate for sculpture exploration since the zoom range is reduced either for macro-scale inspection (a lower zoom level is sufficient for viewing the entire sculpture) and micro-scale inspection (a lower zoom level is sufficient for observing the finest details).
- (3) **Attractiveness of distinctive features:** independently of the user interface employed and the exhibition considered, both focus map analysis and cluster analysis revealed that the most decorated landmarks are considered the most attractive to visitors, and both the control metaphors (free 3D interface and image-based selection widget) are effective enough to permit the visitors to inspect very close details of the objects that would be otherwise hard to observe on the real sculptures (i.e., small carvings). A more detailed visual analysis of focus maps revealed recurrent patterns in user interest depending on the subjects depicted as well as on the degradation of the sculpture. It is interesting to note that these patterns do not completely correspond to those emerging from geometric and photometric saliency computation, since they are influenced by other semantic information. It would be interesting to derive from these patterns some data-specific models of viewpoint preference tailored to specific models.
- (4) **View coverage and exploration paths:** each statue has a distinctive set of preferred viewpoints from which the statue is preferentially seen. These optimal viewpoints cover both overall views and zooms towards interesting details. Paths generated by connecting these viewpoints favoring a coarse to fine exploration generate interesting navigations. Interesting avenues of future research would include preferred navigation directions and clustering of user paths to generate these motions, in order to make the animations even more informative. The paper shows the possibility of generating trajectories given a vast amount of data. While it is true that "authored" solutions are obviously feasible and simpler, it is in our opinion worth to also explore data mining approaches to discover user preferences (possibly linking them to various user kinds). In particular, given the vast amount of data that a single museum setup can provide in reasonable short time, and the fact that the possibility of sharing models over the internet makes it possible to gather information from many users in parallel, such solutions could be of practical interest in the future.
- (5) **Interconnection between user interfaces:** the similarities between the focus map analysis and cluster analysis let us deduce that there is a tight connection between the two control interfaces.

Moreover, from direct observation of visitor behaviors, we realized that users tend to use the image-based selection widget for reaching the most attractive points-of-interest and successively they move around locally with the free exploration interface to further find appealing surrounding views or to discover other details. We think that the interconnection between the two interfaces greatly reduced the visitor effort, enabling them to perceive the aura of the work-of-art in less time and with greater satisfaction.

In general, we found that most of the choices for the setup and the user interface had influence on the exploration activity. For instance, some particular focus maps for free navigation could have been biased by the selection of thumbnails (since they act as starting points for the free navigation). We reduced this effect by not considering in the analysis of free motion the starting position after each goto operation, even if the user stays there for more than 5 seconds. Moreover, this effect is amortized over the enormous number of acquisitions, in a way that general trends with respect to exploration behavior, preferences and attractivity of features emerged clearly.

7. CONCLUSIONS

In this paper, we presented visual and quantitative methods, as well as analysis results, to study user activity during virtual explorations of a collection of protostoric Mediterranean sculptures. The interactive stations were used by tens of thousands of visitors during 24 months at the National Archaeological Museum of Cagliari, 10 months at the Civic Archaeological Museum of Cabras, six months at the National Prehistoric Ethnographic Museum Luigi Pigorini in Rome, and six months at the Civic Archaeological Museum in Milan. This user study focused on discerning the main visitor behaviors during 3D explorations by employing tools for deriving focus measures on surfaces and for clustering and knowledge discovery from high-dimensional data. Results confirmed that the system permits casual users to inspect and appreciate the artworks at different scales: from general views to very close details that would be otherwise difficult to observe on the real sculptures. This usage analysis can provide useful insights into the development of constrained 3D interfaces for virtual exploration of scenes and models, and for creating precomputed paths with additional informative contents [Balsa Rodriguez et al. 2015]. We believe that the derivation of focus maps (as well as of tools for spatial visualization and analysis of recorded view positions) can allow technology assessment and sociological analysis. Furthermore, the analysis of the exploration workspace can provide useful information for planning future temporary and permanent exhibition layouts. Specifically, gathered data on virtual exploration can be used as input for deriving the optimal positions for sculptures in a way to provide the maximum accessibility to the details considered most interesting. We plan to further explore optimal positioning in future work.

Acknowledgments. This work was partially supported by *Soprintendenza Archeologia della Sardegna* under project MONTEPRAMA4. We also acknowledge the contribution of Sardinian Regional Authorities under projects VIGEC and TDM. The project also had support from King Abdullah University of Science and Technology. We also thank Virginia A. Unkefer, from Research Publication Services at King Abdullah University of Science and Technology, for carefully proof-reading and revising the manuscript.

REFERENCES

- Marco Agus, Fabio Marton, Fabio Bettio, and Enrico Gobbetti. 2016. Interactive 3D exploration of a virtual sculpture collection: an analysis of user behavior in museum setting. In *The 13th Eurographics Workshop on Graphics and Cultural Heritage*.
- 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 Bettio, Fabio Marton, and Enrico Gobbetti. 2016. Digital Mont'e Prama: Exploring large collections of detailed 3D models of sculptures. *ACM Journal on Computing and Cultural Heritage* (2016). To appear.

- 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, Enrico Gobbetti, Emilio Merella, and Ruggero Pintus. 2013. Improving the digitization of shape and color of 3D artworks in a cluttered environment. In *Proc. Digital Heritage*. 23–30.
- 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.
- Silvia Biasotti, Andrea Cerri, Bianca Falcidieno, and Michela Spagnuolo. 2015. 3D Artifacts Similarity Based on the Concurrent Evaluation of Heterogeneous Properties. *J. Comput. Cult. Herit.* 8, 4, Article 19 (Aug. 2015), 19 pages.
- Tamy Boubekeur. 2014. ShellCam: Interactive geometry-aware virtual camera control. In *Proc. ICIP*. 4003–4007.
- Umberto Castellani, Marco Cristani, Simone Fantoni, and Vittorio Murino. 2008. Sparse points matching by combining 3D mesh saliency with statistical descriptors. In *Computer Graphics Forum*, Vol. 27. Wiley Online Library, 643–652.
- Paolo Cignoni, Fabio Ganovelli, Enrico Gobbetti, Fabio Marton, Federico Ponchio, and Roberto Scopigno. 2004. Adaptive tetra-puzzles: efficient out-of-core construction and visualization of gigantic multiresolution polygonal models. *ACM TOG* 23, 3 (2004), 796–803.
- Dane Coffey, Nicholas Malbraaten, Trung Bao Le, Iman Borazjani, Fotis Sotiropoulos, Arthur G. Erdman, and Daniel F. Keefe. 2012. Interactive Slice WIM: Navigating and Interrogating Volume Data Sets Using a Multisurface, Multitouch VR Interface. *IEEE TVCG* 18, 10 (2012), 1614–1626.
- Livio De Luca. 2014. Methods, formalisms and tools for the semantic-based surveying and representation of architectural heritage. *Applied Geomatics* 6, 2 (2014), 115–139.
- Paul Debevec. 2005. Making "The Parthenon". In *Proc. VAST*.
- Vincenzo Deufemia, Luca Paolino, Genoveffa Tortora, A Traverso, Viviana Mascardi, Massimo Ancona, Maurizio Martelli, Nicoletta Bianchi, and Henry De Lumley. 2012. Investigative analysis across documents and drawings: visual analytics for archaeologists. In *Proceedings of the international working conference on advanced visual interfaces*. ACM, 539–546.
- Marco Di Benedetto, Fabio Ganovelli, Marcos Balsa Rodriguez, Alberto Jaspe Villanueva, Roberto Scopigno, and Enrico Gobbetti. 2014. ExploreMaps: Efficient Construction and Ubiquitous Exploration of Panoramic View Graphs of Complex 3D Environments. *Computer Graphics Forum* 33, 2 (2014), 459–468. <http://vic.crs4.it/vic/cgi-bin/bib-page.cgi?id='DiBenedetto:2014:EEC'> Proc. Eurographics 2014.
- Helin Dutagaci, Chun Pan Cheung, and Afzal Godil. 2012. Evaluation of 3D interest point detection techniques via human-generated ground truth. *The Visual Computer* 28, 9 (2012), 901–917.
- 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.
- H. J. Falk and L. D. Dierking. 2000. *Learning from Museums: Visitor Experience and the Making of Meaning*. Rowman & Littlefield.
- N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. 2005. Robust Global Registration. In *Symposium on Geometry Processing*. 197–206.
- Mark Harrower and Cynthia A Brewer. 2013. ColorBrewer. org: an online tool for selecting colour schemes for maps. *The Cartographic Journal* (2013).
- Jiaxi Hu and Jing Hua. 2009. Salient spectral geometric features for shape matching and retrieval. *The visual computer* 25, 5-7 (2009), 667–675.
- Anil K Jain. 2010. Data clustering: 50 years beyond K-means. *Pattern recognition letters* 31, 8 (2010), 651–666.
- Piotr Jankowski, Natalia Andrienko, Gennady Andrienko, and Slava Kisilevich. 2010. Discovering landmark preferences and movement patterns from photo postings. *Transactions in GIS* 14, 6 (2010), 833–852.
- Takeshi Kurashima, Tomoharu Iwata, Go Irie, and Ko Fujimura. 2010. Travel route recommendation using geotags in photo sharing sites. In *Proceedings of the 19th ACM international conference on Information and knowledge management*. ACM, 579–588.
- Chang Ha Lee, Amitabh Varshney, and David W Jacobs. 2005. Mesh saliency. In *ACM transactions on graphics (TOG)*, Vol. 24. ACM, 659–666.
- Max Limper, Arjan Kuijper, and Dieter W Fellner. 2016. Mesh Saliency Analysis via Local Curvature Entropy. (2016).
- Shusen Liu, Dan Maljovec, Bei Wang, Peer-Timo Bremer, and Valerio Pascucci. 2015. Visualizing high-dimensional data: Advances in the past decade. In *Proc. Eurographics Conf. Visualization*. 127–147.
- ACM Journal on Computing and Cultural Heritage, Vol. 0, No. 0, Article 0, Publication date: January 2017.

- Luigi Malomo, Paolo Cignoni, and Roberto Scopigno. 2016. Generalized Trackball for Surfing Over Surfaces. In *STAG: Smart Tools and Apps for Graphics*. Eurographics.
- Fabio Marton, Marco Agus, Enrico Gobbetti, Giovanni Pintore, and Marcos Balsa Rodriguez. 2012. Natural exploration of 3D massive models on large-scale light field displays using the FOX proximal navigation technique. *Computers & Graphics* 36, 8 (December 2012), 893–903.
- 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* 7, 2 (2014), Article 12.
- Robert E Patterson, Leslie M Blaha, Georges G Grinstein, Kristen K Liggett, David E Kaveney, Kathleen C Sheldon, Paul R Havig, and Jason A Moore. 2014. A human cognition framework for information visualization. *Computers & Graphics* 42 (2014), 42–58.
- Marco Potenziani, Marco Callieri, Matteo Dellepiane, Massimiliano Corsini, Federico Ponchio, and Roberto Scopigno. 2015. 3DHOP: 3D Heritage Online Presenter. *Computers & Graphics* 52 (2015), 129–141.
- Kun Qian, Jizhou Sun, Hui Chen, and Jiawan Zhang. 2015. Visual analysis method for cultural heritage site risk assessment. *Journal of Visualization* (2015), 1–12.
- Ivan Sipiran and Benjamin Bustos. 2011. Harris 3D: a robust extension of the Harris operator for interest point detection on 3D meshes. *The Visual Computer* 27, 11 (2011), 963–976.
- Ran Song, Yonghuai Liu, Ralph R Martin, and Paul L Rosin. 2014. Mesh saliency via spectral processing. *ACM Transactions on Graphics (TOG)* 33, 1 (2014), 6.
- Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. 2009. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. In *Computer graphics forum*, Vol. 28. Wiley Online Library, 1383–1392.
- Alessandro Torrisi, Giovanni Signorello, Giovanni Gallo, M De Salvo, and Giovanni Maria Farinella. 2015. Mining Social Images to Analyze Routing Preferences in Tourist Areas. (2015).
- 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.
- Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier, and Christophe Schlick. 2010. Radiance scaling for versatile surface enhancement. In *Proceedings of the 2010 ACM SIGGRAPH symposium on Interactive 3D Graphics and Games*. ACM, 143–150.
- N. Walter, O. Aubreton, and O. Laligant. 2008. Salient point characterization for low resolution meshes. In *2008 15th IEEE International Conference on Image Processing*. 1512–1515.
- Malte Weiss, Simon Voelker, Christine Sutter, and Jan Borchers. 2010. BendDesk: dragging across the curve. In *Proc. ACM ITS*. 1–10.
- Raphael Wimmer, Fabian Hennecke, Florian Schulz, Sebastian Boring, Andreas Butz, and Heinrich Hussmann. 2010. Curve: revisiting the digital desk. In *Proc. NordiCHI*. 561–570.
- Jiawan Zhang, Kai Kang, Dajian Liu, Ye Yuan, and E Yanli. 2013. Vis4heritage: visual analytics approach on grotto wall painting degradations. *Visualization and Computer Graphics, IEEE Transactions on* 19, 12 (2013), 1982–1991.