

DECHO - A Framework for the Digital Exploration of Cultural Heritage Objects

DANIEL G. ALIAGA^{*}, ELISA BERTINO^{*}, STEFANO VALTOLINA⁺

^{*} Purdue University, USA

⁺ University of Milan, Italy

We present a framework for the digital exploration of cultural heritage objects. Today computing and information technology is pervasive and ubiquitous and has boosted at unprecedented levels information diffusion and productivity. Such technology is today a ripe context for succinctly gathering knowledge by combining in innovative ways powerful visualization tactics, rapid access to a significant amount of relevant information, domain-specific knowledge, and rich and pervasive tools to sort, group, and slice the information and knowledge in different ways. To this end, we present a complete framework that is easy to use, does not require expensive/custom equipment, and has been designed for helping archaeology researchers and educators reconstruct and analyze the historical context of cultural heritage objects. Our main inspiration is that archaeology today would benefit significantly from having spur-of-the-moment access to information from a variety of heterogeneous data sources and being able to have multiple participants visually observe factual and visual data in an intuitive and natural setting. While we present a framework geared towards archaeology, in the long term we envision reusing it in a variety of fields.

Concretely, our framework includes data acquisition, data management, and data visualization components. The data acquisition component enables the fast, easy, and accurate addition of 3D object models and factual data, including narrations. The data management component includes a novel semantic database system that provides an intuitive view of the available contents in terms of an ontology, supports the addition of narrations, integrates data stored by other databases, and supports object retrieval, browsing and knowledge navigation. The data visualization component provides visual feedback which is a crucial part of an exploratory endeavor. It provides the ability to alter the appearance of archaeological objects, complete fragments of 3D object models, and several compelling forms of digital inspection and information visualization. All algorithms exploit knowledge from the database and from the obtained 3D models. Visuals can be applied on top of the physical object or on a 3D model shown in a traditional display, controllable via a webpage interface.

Categories and Subject Descriptors: I.3 [Computer Graphics], I.3.3 [Picture/Image Generation], I.3.7 [Three-dimensional Graphics and Realism], I.4.1 [Digitization and Image Capture], J.5 [Arts and Humanities], H.3 [Information Storage and Retrieval], H.5 [Information Interfaces and Presentation], H.5.4 [Hypertext/Hypermedia].

General Terms: cultural heritage, acquisition, modeling, narration, ontology.

Additional Key Words and Phrases: computer graphics, databases, semantic, ontology-based models.

1. INTRODUCTION

In this article, we introduce DECHO, a novel Framework for Digital Exploration of Cultural Heritage Objects. DECHO has been designed for helping archaeology researchers and educators reconstruct and analyze the historical context of cultural heritage (CH) objects. Archaeology relies on a significant amount of intuition, insight, and analysis which must be done in order to regenerate the past environment and to draw conclusions about a culture which, typically, is non-existent today. Our framework exploits computer and information technology that today provides innovative and powerful visualization tactics, rapid access to a significant amount of relevant information, domain-specific knowledge, and rich and pervasive methods to sort, group, and slice the information and knowledge in different ways (Figure 1). DECHO has been developed with the main goal of providing archaeologists with a semantically-rich information space, organized according to their conceptual view of the CH content and which can be explored and visualized and enriched in different ways. However DECHO should also be easy to use and should allow archaeologists to acquire contents without

having to use expensive equipment. Another important goal is to be able to use DECHO for the dissemination of CH information to the broad public and for educational purposes, as it is important that this information, possibly enriched by contributions by domain experts, be widely available to a broad community.

While we present a framework geared towards archaeology, in the long term we envision reusing DECHO in a variety of fields including art and sculpture, computer-aided design, engineering, and product life cycle management.

1.1 Background

Many research efforts have focused either on obtaining virtual 3D reconstructions or on creating information databases of CH objects and archaeological excavations. Virtual reconstruction methods focus on passive or active 3D acquisition methodologies. For example, methods that illuminate the object using a digitally-controlled light source (e.g., structured light [Scharstein and Szeliski 2003]) and laser-scanning methods [Levoy et al. 2000]) focus on assembling detailed geometric scans of large objects but require

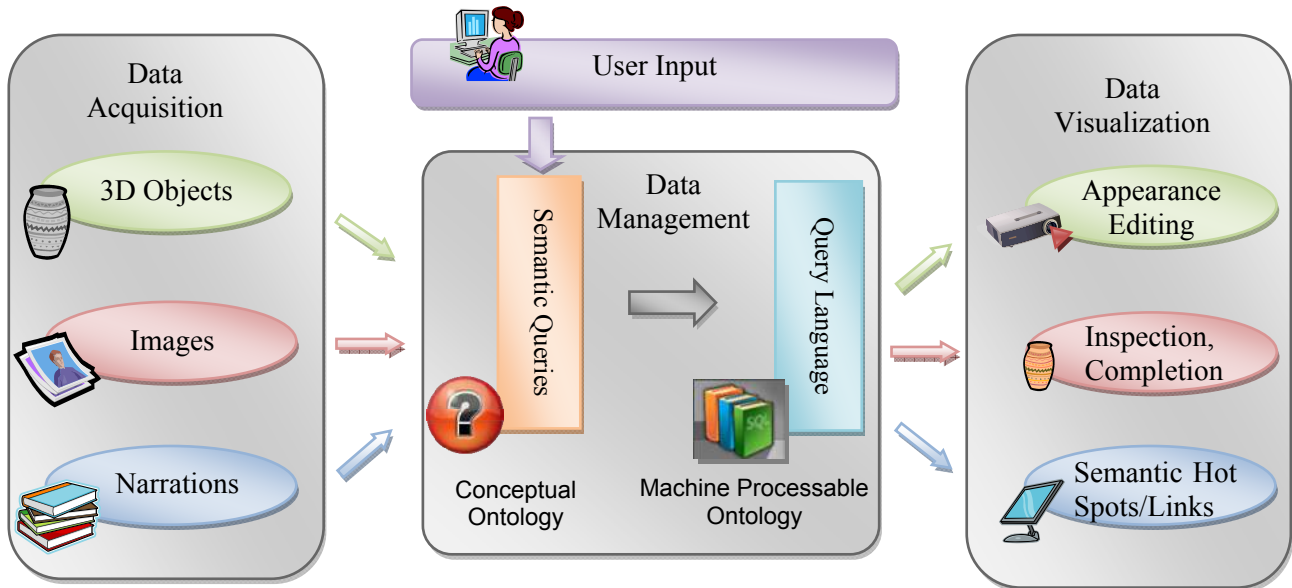


Figure 1. DECHO. We show an overview of the DECHO system, designed for helping archaeology researchers and educators reconstruct and analyze the historical context of cultural heritage objects.

high-technical expertise and sometimes many expert users. Image-based modeling methods have also acquired objects that support high-quality photometric operations. Photometric operations work directly on the intensity values of the images (e.g., relighting and detail enhancement) and do not typically compute accurate geometric models (e.g., [Gunawardane et al. 2009]). Many traditional computer vision methods focus on passive (i.e., unobtrusive) methods but lack in automation and robustness (e.g., [Pollefeys et al. 2004]).

With regards to information databases, a variety of digital libraries for archaeology have been developed and supported. Some of the large initiatives include the Digital Library Project at UC Berkeley [BerkeleyURL], the Alexandria Digital Library Project at UC Santa Barbara [AlexandriaURL], the Theban Mapping project at the American University at Cairo [ThebanURL], the Perseus Project at Tufts University [PerseusURL] and finally EUROPEANA [EuropeanaURL], a platform for collecting European CH and for enabling searching and exploring said heritages across several European cultural institutions. These large efforts focus on the gathering of the information and on passively providing a wide range of source materials to as large an audience as possible. Other efforts have focused, for example, on improving access to archaeology libraries by using multilingual glossaries and ontologies (e.g., [Monroy et al. 2010]), on integrating and handling highly heterogeneous data (e.g., [Ravindranathan et al. 2004]), and on using virtual reality techniques to present archaeology data (e.g., Vote et al. [2000]).

1.2 Overview

A key inspiration for our framework is that modern archaeology would benefit significantly from having spur-of-the-moment access to information from a variety of heterogeneous data sources and being able to have multiple participants visually observe factual and visual data in an intuitive and natural setting. The primary difference, as well as advantage, between contemporary archaeology and “19th century archaeology” is that the investigation of findings is not done in isolation but in relation to historical, social, economical, and geographical contexts. Although this is enabled by the variety of scientific techniques and analysis used, there are two fundamental problems encountered during archaeological research and pedagogy which we address in this work. First, archaeology must attempt to link a dynamic and unseen past filled with humans to a static present where only the partial physical and inanimate remains of that past behavior are left in the archaeological record. Second, archaeology must contend with the ethical dilemma of leaving the discovered artifacts (and sites) unaltered and of desiring to physically alter them so as to recreate the environment of the studied location.

To address such problems, DECHO integrates information and visualization abstraction. In particular, DECHO yields fast and intuitive access to historical information and analysis from heterogeneous sources and provides the ability to easily incorporate and virtually modify the virtual and physical appearance of CH objects for several forms of information visualization. Due to the use of many scientific techniques and to the inherently fragmented nature of

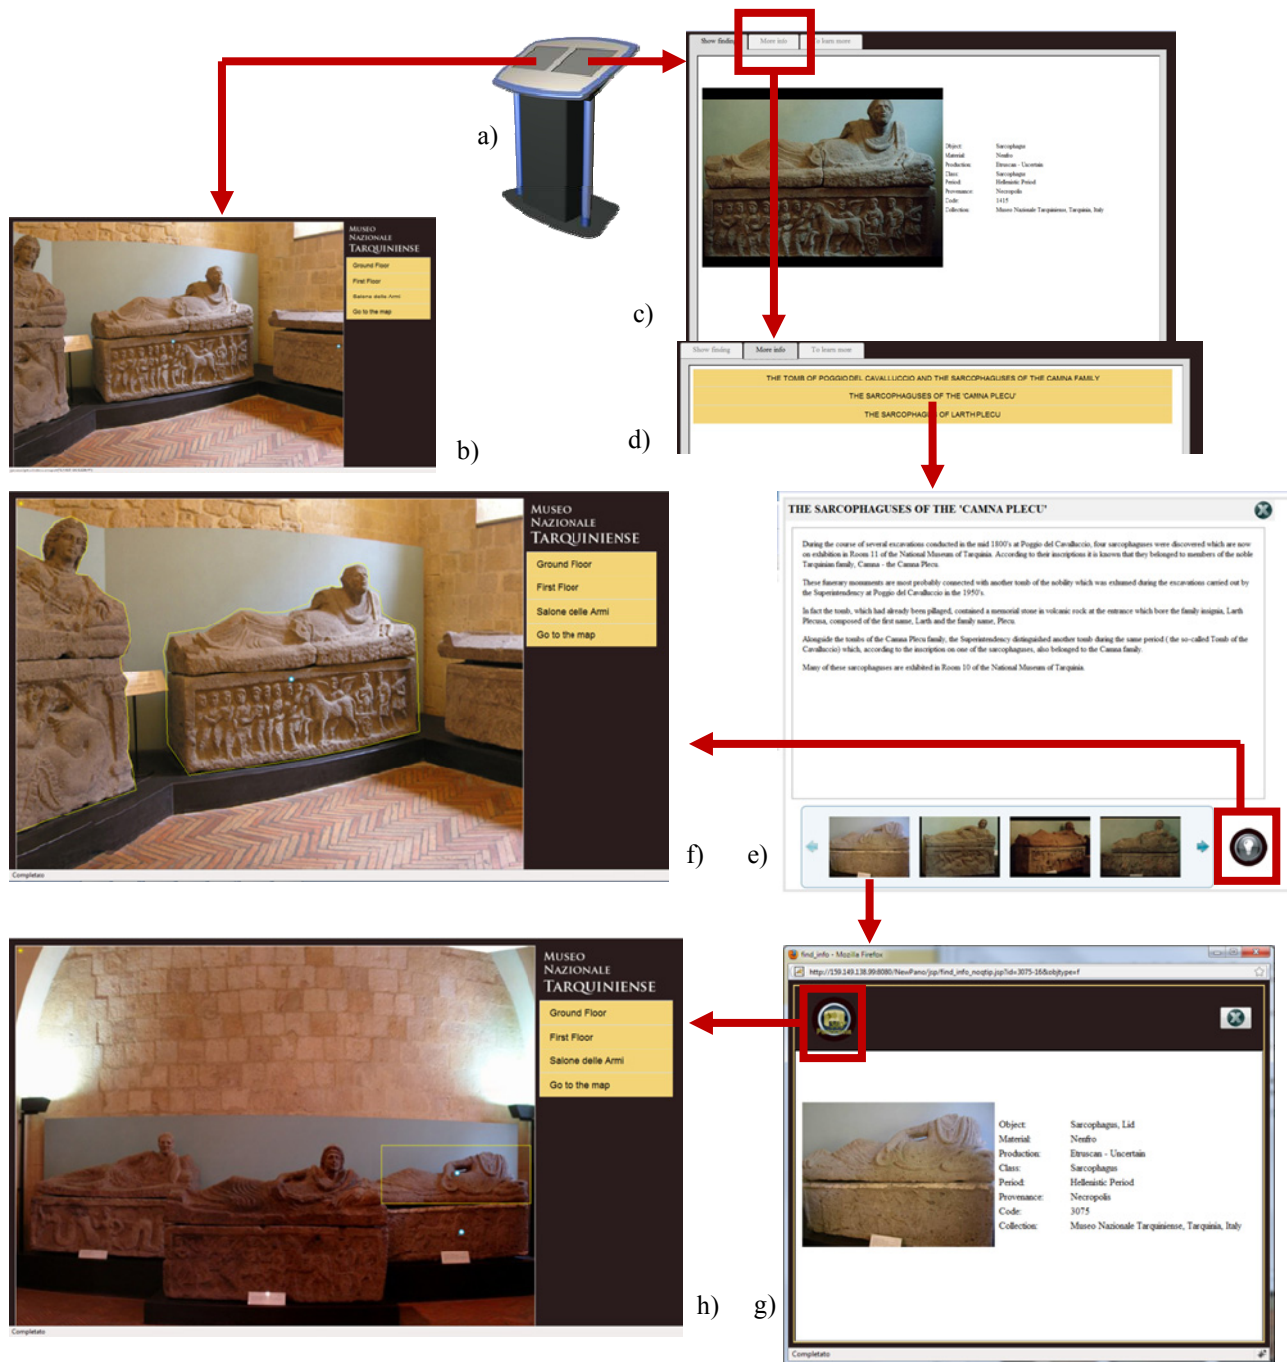


Figure 2. Example. We depict an early prototype and deployment of the DECHO system. Our two screen interactive kiosk is used for navigation (left screen) and detail enhancement (right screen) for personalized exploration.

the observations, information of interest to archaeological researchers is stored in a diverse set of archives and databases, often owned by different scientific and cultural institutions. We provide mechanisms for quick access and organization of such distributed information by using a domain-specific

conceptual mind map. Further, to enhance the virtual reconstruction of historical contexts, DECHO also provides the ability to digitize and virtually modify the virtual or physical appearance of CH objects, enabling an intuitive and natural setting for digital inspection, completing partially acquired objects, and appearance

alteration (e.g., restoring deteriorated objects). Altogether, our approach supports going from the factual information of the findings to virtual reconstructions of their historical, social, economical, geographical and anthropological contexts.

In more detail, our information-abstraction component organizes the database in a way that is more natural for archaeologists to use. It relies on an ontology (e.g., “a description of the concepts and relationships that can exist for an agent or community of agents” [Gruber 1993]) organized into two levels, and on a specialized tool, referred to as *ontology management system* (OMS), for managing it. The top level of the ontology presents a view that is suitable for non-computer experts while the bottom level is suitable for interacting with the computing infrastructure. The top level ontology exploits the concept of an “archeological mind map” for producing a representation of concepts and relationships suitable for archaeologists. The information core also supports the fast and flexible addition of narrations and the ability to perform object retrieval, browsing, and knowledge navigation.

The visualization-abstraction component encompasses solving the competing goals of automation, robustness, and high-quality 3D modeling while enabling compelling visualization tools. Adding images and 3D objects to a framework’s database should be fast, simple, accurate, and low-cost. A fast and simple approach is crucial to facilitate adding many objects to the database by a range of non-expert and expert users. An accurate acquisition is important so as to enable valid and useful scientific discovery. A low-cost method is also particularly important in practice because funding for expensive equipment is not always possible or such equipment is not available (e.g., during an excavation). We provide a *photogeometric capture system* (PCS) which satisfies the above design goals by combining the advantages of active geometric methods and of photometric-based methods. In addition, we use the same underlying system to provide an *appearance editing system* (AES) which enables visualization by altering the visual appearance of the actual historical object using only projected light. Our tool superimposes on the actual CH objects one of several forms of information visualization, surface analysis data, chemical analysis results, material compositions, tentative visual restorations, and choroplethic maps. Our method does not require the use of any sort of physical goggles, rather all visualizations can be observed by casual observers using the naked eye. Thus, this setting provides a very natural visualization where analysis data can be directly and spatially associated with the object.

Further, we provide several visually-based digital inspection tools and a method to produce completed 3D models from acquired object fragments.

1.3 Example Use

Our system has been applied to several historically-significant objects, including Etruscan objects from ancient Italy, and to data and analysis collected as part of the Tarchna project [Bagnasco et al. 2008, TarchnaURL]. An early prototype of DECHO based on interactive multimedia kiosks (Figure 2a) has been deployed in several European museums including the National Museum of Tarquinia (Italy), the Etruscan Necropolis in Tarquinia, the Louvre Museum, the Ny-Carlsberg Glyptotek Museum, the National Museum of Copenhagen, the National Museum of Warsaw, and the Archaeological Civic Museum of Milano. For users, DECHO enables a visitor to use the kiosk to create a custom tour of the museum being visited. For archaeologists, DECHO enables easily adding information and objects to its database. For example, Figure 2b presents the panorama of a room of the National Museum of Tarquinia, called “Hall of Sarcophaguses”, which appears in the left (context) screen of the kiosk. In Figures 2c-h, the visitor then interactively selects a sarcophagus for more detailed viewing in a second (detail) screen in the form of narrations (e.g., text describing the finding at one of several possible levels of detail), clickable hot spots, and other visualization means – more information on this example, and others, is in Section 6. This prototype offers the visitor a high degree of personalization. The information can be presented at different amounts of detail in order to encompass various possible levels of investigation. DECHO helps in overcoming the rigidity inherent to a real museum’s exhibition, by giving the user an opportunity to customize the visit (via our ontology management system) and the visualization (via our appearance editing system) according to their own interest or curiosity. Moreover, DECHO also provides the archaeologist with tools to more intuitively record the discovered information and interpretations (via our ontology management, including narrations, and via our capture system).

1.4 Contributions

Succinctly our main technical contributions include

- the adaptation of a fast, accurate, and self-calibrating acquisition method for use in obtaining 3D models of archaeological objects,
- the development of a novel two-level ontology system able on one side to support a representation of objects of interest and their properties which is natural for archaeologists,

and on the other side to support the integration of data from heterogeneous data sources, and

- visualization tools including the support of rendering enhancements, semantic hot spots, 3D object completion, and overall changing the appearance of archaeological objects while simultaneously providing fast and efficient access to a wealth of factual data.

1.5 Article Organization

The remainder of this article is organized as follows. First, we provide a summary of related work (Section 2). Then, we describe how to incorporate information into our framework, with a particular emphasis on object acquisition (Section 3). Afterwards, we focus on our information component (Section 4) and our visualization component (Section 5). Finally, we provide some additional results (Section 6) and conclude (Section 7).

2. RELATED WORK

No actual system or framework exists that can be directly compared to DECHO, as DECHO is the first framework to seamlessly integrate 3D object acquisition techniques, knowledge management, heterogeneous information integration, visualization and appearance editing techniques. However, DECHO selectively builds upon methods for object acquisition, the management of data related to the objects, and object visualization -- we describe major related work in these areas.

2.1 Data Acquisition

For 3D object acquisition, both geometric-based and photometric-based methods have been used in archaeology. Passive geometric-based methods have the advantage of being unobtrusive. However, they need to establish robust correspondences using only natural features (e.g., [Pollefeys et al. 2001, Pollefeys et al. 2004]). On the other hand, active geometric methods explicitly generate correspondences. One popular active method option is laser scanning (e.g., [Allen et al. 2003, Williams et al. 2003]). But, the process is expensive and time-consuming, does not necessarily capture color information, and does not necessarily produce smooth and accurate normals. For example, a Cyberware scanner [CwareURL] costs about \$80k and requires significant post-processing -- both of which prevent wide-dissemination and frequent-use because the archaeologist is burdened by this task or must plan ahead to send the object elsewhere to be scanned. Moreover, while some active systems specialize on acquiring non-diffuse objects

(e.g., [Park and Kak 2008]), most assume strictly diffuse objects.

Most structured-light methods (e.g., [Davis et al. 2002, Zhang et al. 2003b, Scharstein and Szeliski 2003]) use a priori calibration and projected light patterns to robustly establish correspondences and, similar to lasers, to reconstruct mostly-diffuse objects. While some self-calibrating structured-light systems have been proposed [Furukawa and Kawasaki 2005], the majority require a pre-calibrated setup. Self-calibration methods often rely on features and on either assumed scene or geometry constraints to estimate parameters [Hemayed 2003]. Further, convergence is difficult and sometimes not possible [Sturm 2002].

An alternative is a photometric-based approach which more directly benefits from the high-resolution of recent digital camera technology. Modern digital cameras easily sample 10 million pixels which correspond to approximately a 0.1mm resolution of a handheld object at arm's length. For example, using an un-calibrated photometric-based approach, three digital pictures captured from the same static camera but using a flash located at three different locations is sufficient to reconstruct fine surface details of the object (e.g., [Woodham et al. 1991, Raskar et al. 2004, Basri et al. 2007]). However, the overall macro-structure of the object cannot be obtained. In particular, it is difficult to surmount the fundamental ambiguity of the generalized bas-relief transformation [Belhumeur et al. 1999]. To improve upon the ambiguity, some previous methods rely on feature tracking [Lim et al. 2005], on a structure-from-motion refinement [Zhang et al. 2003a], or on accurate finding of contours [Hernandez et al. 2008].

The acquisition approach in DECHO builds upon our method that combines a photometric and geometric capture into a single photogeometric technique and has the benefits of both captures [Aliaga and Xu 2008]. The geometric processing uses structured light to produce robust details of the overall object structure. The photometric processing captures very fine details throughout the object's surface. The combination of both, through an optimization framework, enables arriving at a single object surface that incorporates both measurement sources. Hence, this method overcomes the low-frequency deformations of photometric methods and produces 3D reconstructions at camera resolution (as opposed to projector resolution which would be the case for a strictly geometric method). As opposed to other photogeometric methods (e.g., [Rushmeier and Bernardini 1999, Nehab et al. 2005]), ours is self-calibrating and multi-viewpoint. This makes the setup significantly more flexible, practical, and able to obtain more complete 3D models using

only uncalibrated and off-the-shelf digital projectors and cameras. We use the exact same equipment for both geometric and photometric acquisition without needing co-location of the hardware devices.

2.2 Data Management

Several systems have also been developed that are specialized for access, management, and enjoyment of cultural assets [Magnenat-Thalmann et al. 2004, Papagiannakis et al. 2002]. With respect to data management, state-of-the-art research consists of works in the area of heterogeneous knowledge system integration. In general, the problem of integrating heterogeneous databases has been widely investigated during the past 20 years and many approaches, proposals, and tools exist [Bukhres and Elmagarmid 1996, Rahm and Bernstein 2001]. Moreover, the interoperability problem has sparked vigorous discussions also in the digital library community [Fix et al. 2001]. In Cassel et al. [2010], the authors claim that the information must be available from wherever potential users are by means of multiple community oriented entry points to multiple sources.

Some research (e.g., [Bloehdorn et al. 2007, Mazzocchi et al. 2006]) has focused on approaches that combine several semantic technologies on top of a digital library, including ontology management, ontology learning and reasoning, keyword-type search, and text classification in order to allow for the flexible and versatile answering of user-provided questions. For example, the system presented in Mazzocchi et al. [2006] aims at enhancing interoperability among digital assets, vocabularies, metadata, and integration services that formulate queries in a uniform way. In order to tackle this problem, in Payette et al. [2002] the authors adopt a mediation mechanism for mapping, translating and aggregating distributed data into complex objects by means of web services. In Monroy et al. [2010] the authors propose to use an ontology and a multilingual glossary for enhancing the contents of a digital library by showing relationships among terms, definitions, and translations in context with the original sources and thus helping to better understand the contents of the collection.

A major shortcoming of these approaches is that they do not support high-level presentation of information using notions such as the conceptual ontology that we use in order to describe the different cultural aspects characterizing a given information domain. The novelty of our approach lies in the combination of different semantic technologies for retrieving data from a variety of knowledge sources in an integrated manner with well-defined semantics provided by the underlying two-level ontology. The first level is used to integrate museum databases and the second level is for

proving a knowledge representation of a given CH fitting the mental model of domain experts.

The first level is based on an important requirement for knowledge models in the context of the Cultural Heritage representation: the ability to abstract from the different storage strategies of various cultural archives. To address such requirements, several research projects (e.g., [Thomas et al. 2001, Wexler 2001, Lenzerini 2002, May 2005, Mena et al. 2001, Leone et al. 2005]) have led to the definition of ontological models that allow one to describe a given cultural domain and then to retrieve the associated context information from distributed data sources [Valtolina 2008]. Such research has investigated the use of ontology schemas for modeling implicit and hidden knowledge in order to integrate different databases owned by different museums. For example, in a solution called “virtual approach” [Valtolina 2008, Lenzerini 2002, May 2005] data residing at the sources are accessed during query execution, and are not replicated in the integrated system. This approach uses the knowledge base as a semantic access point to the information that can then be retrieved from databases federated by means of the ontology schema.

In the context of CH, a standard reference model has been proposed in recent years: the CIDOC Conceptual Reference Model [Crofts et al. 2006]. The CIDOC-CRM ontology is intended to promote a shared understanding of CH information by providing a common and extensible semantic framework that any CH information can be mapped to. The CIDOC-CRM ontology represents a suitable basis for a common language and for supporting conceptual modeling activities in each cultural context. However, a shortcoming of this model is that during use it requires instantiating the concepts and relationships by means of a specific domain language, in our case an archaeological language. This is to ensure that the knowledge base can effectively express the intrinsic characteristics of a specific cultural context, an operation that an abstract model is not able to support. In practice, the knowledge base to be defined must be the result of a process able to fit the conceptual reference model CIDOC-CRM within the selected cultural domain. Our approach builds upon the CIDOC-CRM to provide a novel two-level ontology controlled by an ontology management system and applied to archaeology.

2.3 Data Visualization

The notion of augmenting physical objects with virtual visualization content has been explored before. However, it often requires the use of computer displays, projection screens, or special viewing goggles (e.g., a head-mounted display). In both Virtual Reality

(VR) and Augmented Reality (AR), significant problems must be addressed such as ensuring the heavy head-mountable unit is securely fastened, having a sufficient vertical refresh rate, and, in the case of AR, having to acquire and/or track the object and/or camera [Azuma et al. 2001]. In contrast to these methods, our approach omits the need for any special viewing devices or replicas. Rather, only the naked-eye is needed and thus provides a significantly more intuitive method of interaction that exploits the natural cues of depth perception and physical inspection.

Building upon our previous work [Aliaga et al. 2008], DECHO includes the ability to change the appearance of archaeological objects placed on an appearance editing stage, to use digital inspection tools, and to have fast access to a wealth of factual information. Researchers have previously developed systems that augment replicas or custom-made “white” objects with projected content [Raskar et al. 2001]. Several radiometric compensation algorithms have also been proposed for calibrating projectors and cameras [Grossberg et al. 2004, Fujii et al. 05, Mitsunaga et al. 1999, Raskar et al. 2003, Wetzstein and Bimber 2007] often for the purpose of projecting movies on top of arbitrary surfaces. However, in archaeology we wish to work with the actual objects. Furthermore, producing sufficiently accurate replicas is time consuming and thus detracts from spur-of-the-moment visualizations and discoveries.

In DECHO, our inspection tools go one step beyond previous inspection methods that enhance the illustration of an object via careful lighting design and non-photorealistic rendering strategies (e.g., [Akers et al. 2003; Aliaga 2008; Barla et al. 2006; Bartesaghi et al. 2005; Lee et al. 2004; Rusinkiewicz et al. 2006 ; Winnemoller et al. 2005]). In addition to rendering enhancements, narrations and ontological information can further enrich the visualization by displaying connections to related information through semantic hot spots embedded in images and on the 3D objects [Mazzoleni et al. 2006]. The combination of narrations, ontology and semantic hot spots are key to supporting seamless navigation among networks of objects and information of interest, to providing semantic-rich user experiences and to fomenting the creation of links between seemingly unrelated data. Further, none of the existing rendering and visualization methods support our goals of changing the appearance of archaeological objects while simultaneously providing fast and efficient access to a wealth of factual data.

3. DATA ACQUISITION

DECHO includes infrastructure to add and manipulate data from a variety of existing data sources.

Archaeological methods include a variety of scientific techniques, such as sophisticated chemical analysis to detect, for example, the production site of a pottery fragment produced at a particular time in the past, as well as the movement of that pottery fragment through geographical space. To incorporate and collect data into our system we provide four general mechanisms. Together these mechanisms provide a comprehensive environment able to address all the data collection requirements by archaeologists

- (i) a database application supporting the inclusion of images, e.g., pictures, X-rays, drawings, maps;
- (ii) a database application for storing several types of meta-data such as from chemical tests;
- (iii) an object reader for 3D models of a priori digitized objects in one of several modeling formats; and
- (iv) a flexible and self-calibrating method to digitize a given object using only off-the-shelf hardware.

The addition of the aforementioned data occurs mainly in two distinct locations. On the excavation site data could be recorded about precise point of extraction of the finding and immediate observations. More advanced analyses (e.g., chemical tests to determine the constitutive material) are performed in a laboratory a posteriori. The acquisition of the 3D model of a finding can be performed in either of these two environments. Since network access is not guaranteed at an extraction site, the information can be cached on a portable computer and then synchronized with the main repository at a later point in time.

The fourth mechanism is motivated by the fact that a key requirement is supporting the easy addition of new 3D objects and fragments (e.g., in one of several possible 3D model formats, such as STL or PLY). In a full deployment, we want users to be able to directly focus on the content and on the research tasks. Such a method is described in more detail in the following subsections.

3.1 Photogeometric Capture System (PCS)

Our photogeometric acquisition method (Figure 3) uses one of three fundamental configurations with $C \geq 1$ uncalibrated digital cameras and $R \geq 1$ uncalibrated digital projectors to obtain object point samples $S = \{S_i\}$ where $i \in [1..N]$ and N is desired to be large. The configurations (one projector and camera, one projector and multiple cameras, or multiple projectors and cameras) are a generalization of structured-light systems and support multi-viewpoint reconstructions. They all provide self calibration and increase the



Figure 3. Photogeometric Capture System. *a) Our method captures photometric (top) and geometric (bottom) observations of an object. b) A novel viewpoint rendering of a 3D model is produced using our photogeometric method. c) A close-up of the model rendered with texture mapping. d) An even closer view of the model now rendered using wireframe and synthetic shading showing details beyond those possible using only geometric observations. Our approach uses the same hardware as standard structured-light but is fully self-calibrating and able to capture models at the resolution of the camera (in this example camera resolution is 10x greater than projector resolution). The average triangle edge length is 0.09035 or about 0.1 mm.*

capture resolution of almost any structured-light system.

For each projector, a sequence of geometric and photometric patterns is projected. Changing the number of patterns enables different time-quality trade-offs. In the limit, using zero geometric patterns yields a purely photometric-based capture while using zero photometric patterns produces a purely geometric-based capture.

The image data for geometric observations G_i is robustly captured using a structured-light based system. Nevertheless, with a structured-light system the appearance of specular highlights and the use of non-diffuse surface materials confounds the process of establishing image correspondences. Our later-described up-sampling process “fills-in” the missing samples in regions where correspondences were not established. The missing details are obtained from photometric processing. An optional and additional approach to further improve the correspondences is to use a more robust pixel classification method, such as our related method [Xu and Aliaga 2009].

The image data for the photometric observations consists of the color intensity of the samples S_i as lit by the projectors. For the first phase of un-calibrated photometric stereo, we assume a diffuse object (i.e., Lambertian) illuminated by each projector acting as a diffuse light source. For the second phase, after geometric self-calibration, the light sources are known and it facilitates one of several improved photometric reconstructions.

3.2 Photometric Modeling

Our method uses photometric modeling to first perform an uncalibrated photometric stereo reconstruction of the object and then again later during the process of multi-view photogeometric optimization. In the first stage, we assume a diffuse object-model and its results

of lighting directions and surface estimate are fed to the next stage of geometric modeling.

Uncalibrated photometric stereo recovers surface normals and lighting directions. The well-known illumination model for diffuse objects lit by distant and directional lights can be written as

$$NL^T = C \quad (1)$$

where N is a $k \times 3$ matrix of k outward-facing surface normals, L is a $l \times 3$ matrix of l light directions pointing towards the light, and C is a $k \times l$ matrix of the observed pixel intensities. Based on Woodham [1991] and Basri et al. [2007], we look closely at the case of one pixel/normal and three lights, namely $k=1$ and $l=3$. The problem of simultaneously solving for the normals N and the lights L is reduced to solving a large (sparse) linear squares optimization given at least six pixels and three intensities per pixel. A surface height field is then calculated from N up to a missing scale factor [Frankot and Chellappa 1988].

3.3 Geometric Modeling

For geometric modeling, we use the photometrically-estimated surface, approximate lighting directions, and tailored reprojection equations to obtain a self-calibrated reconstruction of a subset $S' = \{S'_i\}$ of the object point samples S_i . The size of S' affects the time and quality of the resulting reconstruction. Our approach estimates both the focal length and pose of the projectors, acting as virtual cameras. For all configurations, the projectors will be necessarily fully calibrated but the physical cameras are only optionally calibrated – their calibration is not needed for mapping images and normals onto the object. Geometric modeling seeks to minimize reprojection error expressed by the well-known nonlinear system of equations

$$\sum_j \sum_i \left(\frac{1}{h_{ijz}} \begin{bmatrix} h_{ijx} \\ h_{ijy} \end{bmatrix} - \begin{bmatrix} u_{ij} \\ v_{ij} \end{bmatrix} \right)^2 \text{ where } h_{ij} = F_j(R_j p_i + T_j) \quad (2)$$

where R_j , T_j , and F_j are the unknown 3x3 rotation matrix, 3D translation vector, and 3x3 perspective projection matrix.

In an initialization phase, a sparse and uniformly-distributed subset of object point samples of S' of size n are used to estimate the distance from each projector to the object's center as well as the global projector focal length f . Initial values for p_i and l_j come from the aforementioned photometric modeling. To bring the reprojection of the object points into rough alignment with the observed projections (u_{ij}, v_{ij}) , we optimize the following simplified nonlinear system of equations of only $n+1$ unknowns (f and z_j for $j \in [1, n]$)

$$\sum_j \sum_i \left(\frac{\hat{p}_{ijx} f}{\hat{p}_{ijz} + z_j} - u_{ij} \right)^2 + \left(\frac{\hat{p}_{ijy} f}{\hat{p}_{ijz} + z_j} - v_{ij} \right)^2 \quad (3)$$

where $M_j = [l_j \times w \quad (l_j \times w) \times l_j \quad -l_j]$, $w = [0 \ 1 \ 0]^T$, and $\hat{p}_{ij} = M_j p_i$.

In a next step, our method optimizes for a linear correction to each projector location and performs a global bundle adjustment. In particular, in Equation (2) F_j is replaced by a perspective projection matrix parameterized by f , T_j is replaced by $[0 \ 0 \ z_j]^T$ and R_j is replaced by $Q_j M_j$. Each matrix Q_j is computed using the following linear system of equations in the 8 unknowns of the matrix (i.e., $q_{33}=1$)

$$\sum_i \begin{bmatrix} q_{ij} - u_{ij}(q_{ij} + z_j)/f \\ q_{ij} - v_{ij}(q_{ij} + z_j)/f \end{bmatrix} \text{ where } \hat{q}_{ij} = Q_j M_j p_i \quad (4)$$

Using an iterative process we then include all the remaining object point samples of S' , optimize the projector poses and object points, and remove outliers. The optimization and image-space and world-space culling are repeated until convergence. At the end, M_j is updated to the computed projector pose matrix.

3.4 Photogeometric Optimization

In this final step, we combine the photometric data with the multi-view geometric data into a single linear optimization (Figure 4). Our approach enables a time-quality tradeoff whereby a variable amount of geometric modeling is performed and an approximation of the missing details is obtained from the faster photometric processing. Effectively, it enables the conversion of the traditional nonlinear modeling of multi-million point sample reconstruction into a fast and specialized nonlinear optimization of a small set of points followed by a linear up-sampling and linear multi-view optimization of all points.

The photogeometric optimization process alters the object points so as to best match both photometric and geometric measurements. To search for a best displacement of the current object point estimates, we attempt to

- (i) minimize reprojection error onto the projectors,
- (ii) keep similar relative position of object points, and
- (iii) reduce the difference between photometrically- and geometrically-computed normals.

The over-constrained linear equation to minimize that satisfies the aforementioned triple of properties is

$$e_t = (1 - \alpha)(1 - \beta)\kappa_g e_g + \alpha\kappa_p e_p + \beta\kappa_r e_r \rightarrow 0 \quad (5)$$

$$e_g = \sum_j \sum_i \begin{bmatrix} \hat{p}_{ijx} - \left(\frac{u_{ij} \hat{p}_{ijz}}{f} \right) \\ \hat{p}_{ijy} - \left(\frac{v_{ij} \hat{p}_{ijz}}{f} \right) \end{bmatrix} \quad (6)$$

$$e_r = \sum_i \delta_{ik} ((p_i - p_{ik}) - d_{ik}) \quad (7)$$

$$e_p = \sum_i \delta_{ik} (n_i \cdot (p_i - p_k)) \quad (8)$$

where d_{ik} is the initial distance between point p_i and p_k , δ_{ik} is 1 when p_k is considered a neighbor of p_i and 0 otherwise, and the unknowns are the 3D coordinates of each p_i . Equation (6) maintains points near their correct geometric location. Equation (7) keeps the relative distance between the object points similar to the original. Equation (8) encourages similarity between geometric and photometric normals. Equation (5) can be written as $Ax = b$ and solved as an over-constrained sparse linear least squares problem with empirically determined values used to control the tradeoff between geometric error, photometric error, and relative distance error (e.g., values for κ_g , κ_p , and κ_r).

4. DATA MANAGEMENT

The information core of DECHO is a semantic database system that provides a higher level view of the database contents in terms of an ontology, supports the addition of narrations, integrates data stored by other databases, and supports object retrieval, browsing, and knowledge navigation. The presented approach enables creating a comprehensive knowledge base for CH objects and is able to integrate heterogeneous data-sources, contextualize their information with the active participation of domain experts, and disseminate them using interactive tools.

The novelty of this approach lies in the combination of different semantic technologies for retrieving data from a variety of knowledge sources in an integrated manner with a well-defined semantics provided by the underlying two-level ontology. On the one hand, this ontology defines concepts and relationships characterizing the CH domain fitting the mental model of the cultural experts (represented by an appropriate knowledge base) and, on the other hand, it defines a

machine understandable knowledge representation for retrieving knowledge spread in different data sources.

4.1 Ontology Management System

The DECHO ontology management system (OMS) defines and maintains a two-level information abstraction providing fast and organized access to all relevant information. The top-level conceptual ontology provides a domain-specific view of the information. The definition of this ontology is based on the consideration that, in the cultural field, the presentation of the information has to be defined according to a specific context based on an organization of the knowledge that meets the domain experts' requirements and needs. Devedžić [2005] states that "there is no tool that will provide the terminology exactly the same with what the users expect it to be". Hence, limiting oneself to use a specific machine-processable ontology for CH is not sufficient for a good knowledge representation. Using only such an ontology domain experts are not able to fully exploit their capabilities of expression and communication.

In our approach, a mind map is used to represent the conceptual ontology. A mind map is a graphical diagram by means of which domain experts create a mental model of the knowledge domain independently from the technologies used for organizing and storing the information content itself. The main advantage of the mind map is it enables the domain experts to highlight the proper historical and anthropological meanings of cultural data under which prospective the data is to be disseminated.

Once this conceptual ontology is defined, it is necessary to translate it into a form understandable by the machine. This bottom level machine-processable ontology is expressed in a language that is processable by a computer system. Creating a new ontology from scratch is a time consuming task and, therefore, it is better to exploit a general ontology from which customized cultural ontologies can be derived. In our case, to develop such an ontology we leverage the CIDOC Conceptual Reference Model (CIDOC-CRM). This machine-processable ontology is then mapped onto the underlying data repositories, which in most cases are relational databases (DB).

CIDOC-CRM focuses on interoperability of information systems and databases and abstractly defines concepts and relationships. It hides the peculiarities of any domain-specific language or the intrinsic features of a well-determined cultural context, in our case archaeological heritage. The machine-processable ontology of DECHO is derived from the

CIDOC-CRM by selecting only the relevant classes and properties. This knowledge representation model is for communication with the machines. By contrast, the top-level conceptual ontology allows domain experts to express concepts and concept relationships in the domain language of the culture under study and thus it reflects human reasoning and domain-relevant concepts. Mappings between the two levels of the ontology have also been defined and are an important component of our system.

For a technical point of view, the DECHO ontology uses a machine-readable format such as RDF [RDFURL]. Therefore, the names of classes and properties are encoded using RDF labels. Moreover, the two mappings, the one between the mind map and the machine-processable ontology and the one between the ontology and the relational schema of each database integrated in DECHO, are encoded in RDF.

Technically speaking, by "mapping" we mean a set of equivalence statements between mind map concepts and database schema in the ontology. This mapping allows for defining transformation algorithms (implemented in JAVA) with which it is possible to translate a mind map paths into a semantic query (expressed in SeRQL – a RDF Query Language) and then the SeRQL query into SQL statements. The SQL statements enable accessing the integrated databases by means of Sesame [SesameURL] – an open source semantic Java framework. More information is in Sections 4.3 and 6.

Since the algorithms are written in JAVA and the semantic engine is based on Sesame the use of the DECHO system is allowed without a priori licenses.

Figure 5 shows a mapping between a portion of the archaeological mind map, a portion of the CIDOC ontology and a portion of a DB schema. The portion of the mind map surrounded by the red circle at the top of the figure highlights the relation connecting the concept of archaeological finding with its 3D reconstruction. The top-level mapping shows how this relation is translated onto proper CIDOC classes and properties. The "finding" concept is mapped onto the "E24-Physical Man-Made Thing" class and the "3D reconstruction" concept is mapped onto the "E36-Visual_Item" class. The link is translated using the property "P128 carries" that connects "E24-Physical Man-Made Thing" with "E73-Information Object", a parent class of the E36 class from which the property is inherited.

In addition to providing a two-level knowledge representation, OMS also supports the integration of information from different repositories. To integrate

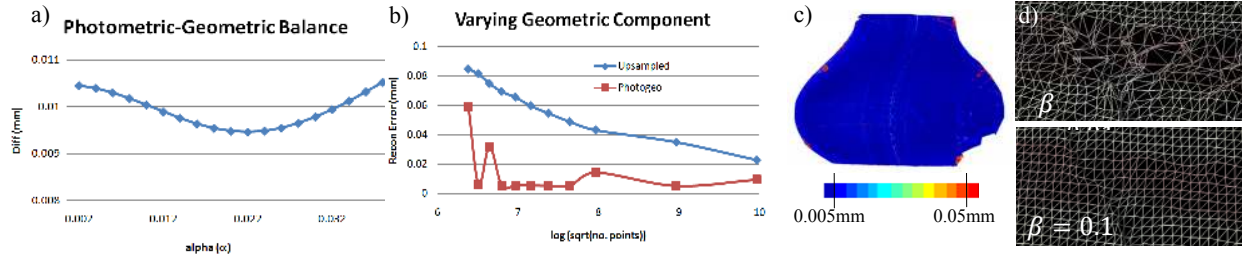


Figure 4. Photogeometric Optimization. a) Close-up of smallest reconstruction errors resulting from varying the photogeometric balance α . b) Improvement achieved by photogeometric optimization using different amounts of geometrically-computed points. c) Color-coded difference between a photogeometrically-computed surface and an ideal obtained by a standard structured light approach. d) View of the distribution of the reconstructed points of same reconstruction accuracy but using $\beta=0.01$ (top) and $\beta=0.1$ (bottom).

information from heterogeneous databases, semantic queries expressed against the machine-processable ontology are translated into queries expressed in the query languages of such databases. This translation is implemented by means of a second level of mapping between the CIDOC ontology and each DB schema. In the current version of DECHO, information is assumed to be stored in relational databases and thus the bottom-level ontology is mapped onto SQL.

For example, in Figure 5 the second level of mapping shows how the CIDOC classes and properties are mapped onto a given DB schema. The information is translated using the table introduced during the design of the DB schema in order to model many-to-many relations (table “Multimedia_link” in the figure). Since this type of table does not model any domain concepts, but it is only useful to define the structure of the relational schema, it does not make sense to try to map it on some ontology entity. For this reason, tables of this kind can be left out from the mapping, without incurring the risk of losing valuable information. The information describing these two mapping levels is represented and stored by the ontology itself. In order to be able to represent this information, the CIDOC-CRM reference model has been extended by adding two classes: DB_Class_Mapping and DB_Property_Mapping. Because these classes do not model any domain concept, they have been placed outside the original CIDOC-CRM class hierarchy. The two classes are endowed with a set of properties which refer the information related to the mapping between the ontology, the mind map and each DB integrated in the OMS.

4.2 Narrations

An important goal of DECHO is to support participatory and collaborative content contributions by archaeologists. To this end, DECHO allows domain experts to enrich alphanumeric content and multimedia

data, such as 3D objects, by *narrations*. A narration [Valtolina et al. 2007] is a communication medium used for disseminating knowledge throughout communities of different users from students to expert users such as scholars and researchers. Narrations allow users to tailor data according to their interests and to convey these tailored data to others. Narrations thus represent a vehicle for knowledge transfer and a tool for context-aware data tailoring.

In DECHO narrations are documents written by domain experts and are able to contextualize CH data in different ways by relating different types of objects. Therefore, domain experts can specify how portions of the data are to be materialized according to historical and anthropological perspectives. From an author’s stand point, a narration is composed of a title, a text and references to data from different databases.

A narration typically focuses on a given topic, such as a description of iconography and its formal origin, a class of monuments or paintings, the role of a given group of objects and so on. In this case the context expresses the network of relations that connect the narration text with the set of CH objects related to it. An example is the narration focusing on the theme of devotional practices or the music in ancient Etruria.

It is important to note that each DB, storing data about CH objects, can be updated, for example because the institution maintaining it acquires some new exhibits. Moreover a new member institution can contribute with a new database. In all these cases the set of CH objects collected as the result of the interpretation of a context in narration is different from the previous sets. Hence the context and the narration itself are time-varying and are able to dynamically evolve as the underlying DB evolves.

In order to integrate the narration concept in the final archaeological ontology, we identified and introduced a specialized ontology concept referred to as narration class. Narrations thus become a key component for

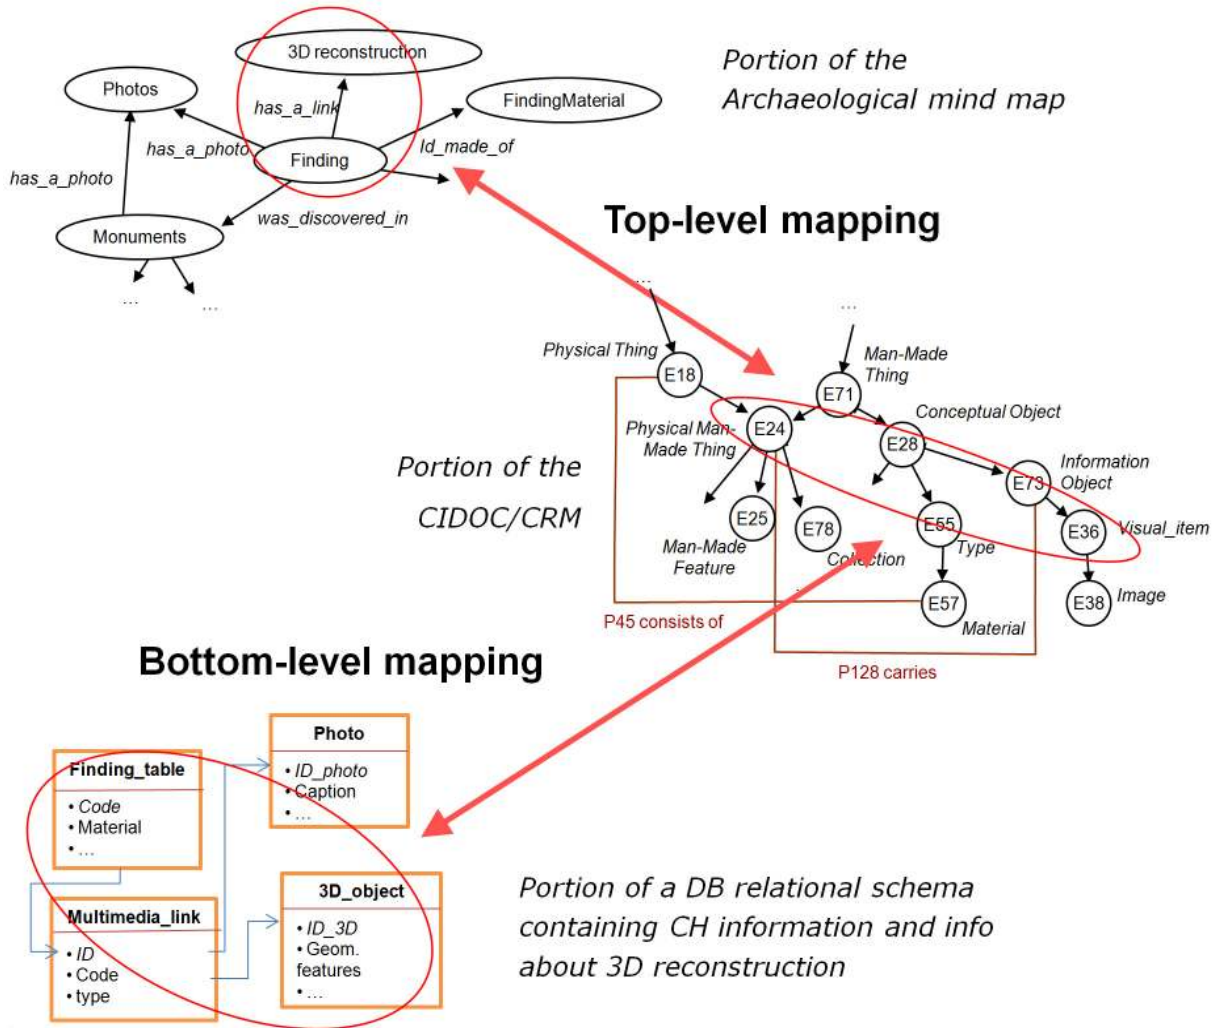


Figure 5. Ontology Management System. Using our ontology management system, we show how from a top-level mapping (conceptual mind map), we access a portion of the CIDOC/CRM standard and then access a bottom-level mapping (machine level map). This two layer abstraction enables fast and intuitive access to a heterogeneous set of data sources which all seemed to be organized in a domain-coherent fashion. The circled regions roughly correspond to the same concepts in the different mappings.

recording creative ideas and knowledge. They provide a medium to record knowledge not explicitly expressed in each database and to enhance the semantics of the collected data. But more importantly they provide a mechanism to facilitate knowledge sharing and to support rich semantic association among data. Such associations are key to enhancing the value of information during research.

In order to disseminate CH information presented in different DBs such as museum collections, chemical DB, 3D reconstructions repository along with the narrations specified by domain experts, the OMS exploits a set of web services, which implement the interface between the final applications and the

archaeological knowledge base. By exploiting these web services, the applications can access the narration repository and all information available in the DB. This interface is the only systems component visible to applications. Such separation simplifies the development of applications by making them independent from the other system's components, from the digital archives accessed, and from the available narrations.

4.3 Information Retrieval

Key DB operations that should be supported, in addition to archiving objects, analysis, and narrations, are *retrieving* and *matching* records of information.

Going beyond standard digital retrieval operations, DECHO exploits the ontology expressing the concepts relevant for the domain and uses it to integrate the available data sources, providing the archeologists with a uniform point of access to all information. A semantic mediator allows the user to formulate queries in terms of the domain's concepts rather than entities defined in the databases' logical schemas; e.g., "retrieve all objects related to the funeral ritual" or "retrieve all objects carrying the craftsman's signature". The semantic query, expressed by the archeologists through a form-based interface, is then mapped onto a semantic query and from this query onto a SQL query. For example, a query to retrieve all findings made of bronze is first translated onto the semantic query:

```

SELECT Finding
FROM {Finding} rdf:type {owl:E22.Man-
Made_Object}, owl:p45f.consists_of
{Material}; {Material} rdf:type
{owl:E57.Material};{Material}
rdfs:comment {MaterialName}
WHERE label(MaterialName) like
"bronze" USING namespace owl =
<"http://www.w3.org/2002/07/owl#">

```

And then onto the SQL query:

```

SELECT FindingID
FROM Material JOIN Findings
ON Material.Material =
Findings.MaterialID
WHERE Material.Material="bronze"

```

In addition to such queries, specialized operations are supported to go beyond current capabilities provided by conventional digital libraries or semantic database systems. One such operation, of particular interest in archaeology, is matching fragments of objects. Several fragments can be discovered nearby and joining them into a single object record is important. A purely geometric approach would analyze the shape and attempt to piece together the whole object. However, sometimes not all fragments are available. This makes reconstructing the entire object difficult and makes even more difficult cataloging this object. As a consequence, many object fragments in archaeological findings end up ignored. Having a system that immediately enables finding similarities and relationships between findings and previous objects, based on several dimensions of similarity beyond geometric shape, yields significant advantages to archaeological research.

Under this perspective the narration concept plays an important role in the reconstruction of a context related to a fragment of a finding. As we discussed in Section 4.2 a narration provides a mechanism for connecting data from different databases according to a given anthropological, historical, or artistic point of view. These points of view represent the new dimensions that together with geometrical features can be used to discover similar fragments usable to lead towards a virtual reconstruction of an entire object or site.

5. DATA VISUALIZATION

Providing visual feedback is a crucial part of an exploratory endeavor. Archaeology requires insight on the part of the researcher to establish relationships between the fragments of factual information and to imagine the culture and influences that affected the findings. DECHO provides the ability to alter the appearance of archaeological objects by precisely controlling the position, color, and intensity of projected light. This display mechanism permits various compelling forms of information visualization. The algorithms exploit knowledge from the database and from the obtained 3D model. All of the methods can be applied on top of the physical object or on a 3D model shown in a traditional display.

5.1 Appearance Editing System (AES)

Our AES enables visually exploring and comparing multiple presentation states of colored and arbitrarily-shaped objects (Figure 6). On a single object we can produce an original as-made appearance that suggests or completes missing designs or forms, a superimposition of the results of analysis, or an accentuation of selected details of the object's surface. A key advantage of our display method is that it circumvents the potential for ethical dilemmas caused by physical alteration. Since our appearance edits do not require the physical alteration of the object and are generated with adjustable digital light sources, they are produced immediately, are completely reversible, and provide an ability not usually possible when studying historical artifacts. Further, since only the naked-eye is needed (i.e., no special viewing devices or goggles), AES immediately enables multiple participants to simultaneously experience new physical appearances and visualizations in an intuitive setting, thus leading to improved collaboration.

To accomplish an appearance edit we start by digitizing the object (or using its already digitized model from Section 3) and capturing a digital image of the object. A desired new appearance is then generated (e.g., by the computer, by the participant, or as a result of an analysis) and a visual compensation algorithm

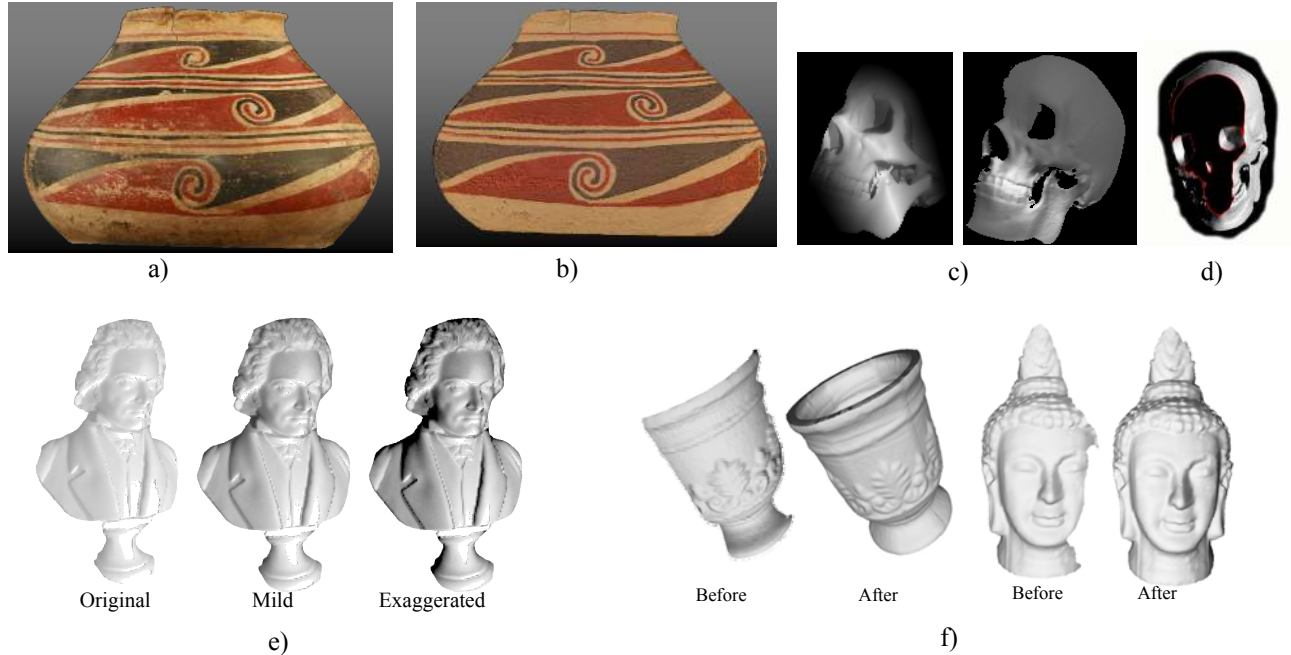


Figure 6. Visualization. We show an example result of our appearance editing system: a) an original Mexican vessel from 1200-1400AD and b) the same object after appearance editing. Both appearances are visible to the naked eye. We also show example visualizations produced via our interactive digital inspection tools, including c) depth-based virtual illumination, d) virtual slicing, and e) shading exaggeration. Finally, in f) we show results using our described object completion algorithm on two examples.

calculates the position, color, and intensity of light needed to project onto the object in order to produce the new appearance. The AES uses a geometrically and radiometrically calibrated projection system – it could be the same used by PCS in Section 3. An optimal compensation image is then projected onto the object from several digital projectors. We define an optimal compensation image as one that achieves a desired balance between supporting the maximum light intensity and maximum color variation while still producing a smooth visual compensation. Further, in order to limit damage due to prolonged exposure to light, the system constrains the maximum amount of light intensity to be incident per unit surface area of the object below a fixed constant E_{bound} . The amount of light that the object receives is particularly important when working with light-sensitive objects, as is the case with important historical artifacts. In general, the user only needs to define E_{bound} and our system will balance out the projectors' light contributions to the visual compensation in such a way to achieve a visually appealing result. Altogether, given a new appearance, the compensation image is computed and projected within a few seconds or less.

5.2 Inspection

In many application scenarios, the archaeologist might have a large number of object fragments to inspect. For

example, an archaeologist (or artist or historian) might wish to digitally magnify surface details, to relight the objects from numerous orientations, and to create synthetic illustrations on the fly to help with the analysis. In an educational setting, the archaeologist (or anthropologist) might explain anatomy to young students using a handheld mock-up of a skull. Being able to virtually slice the object on the fly, even approximately, would add powerful interactive visuals to the didactic task. For the inspection of known objects, a collection of shape profiles of a vessel could be preliminarily inspected and used to prime a search through an object database.

Our framework includes digital tools for immediately performing digital inspection of captured objects by exaggerating important surface details and by generating realistic and non-photorealistic illustrations (Figure 6c-e). In particular, our system provides the following visualization methods:

- exaggerated shading - brings out surface detail by exaggerating the shading of the surface,
- depth-based shading - in addition to shading varying according to surface normals, surface height can be used to apply different shading strategies to different depths of the object; we provide two depth-based shading methods (i.e.,

“depth lights” and “depth-based detail modulation”),

- object slicing - our system also supports the slicing of objects at selected distances from the reference plane and the rendering of multiple iso-distance curves; using the reference plane, we can cull object points that are “behind” the reference plane; by performing this task interactively, we in fact obtain a live rendering of the object’s contour intersecting the reference plane, and
- restoration - the restoration method provides an interactive application whereby a user can restore the appearance of the observed object using an iterative algorithm to compute the image necessary to give the object the illusion of being restored; the algorithm uses an energy minimization method to enforce a set of criteria over the surface of the object and provides an interactive tool to the user which can then further modify the restoration in a few minutes.

5.3 Completion

Our completion method takes advantage of symmetries present in the geometric shape of many natural and man-made objects in order to yield plausible complete models of observed archaeological objects (Figure 6f). Our tool quickly and efficiently classifies the geometric shape into one of a small set of symmetry cases, and then reconstructs more than what is visible in the provided set of input images. In fact, an entire object can be reconstructed from a single camera viewpoint. Using an optimization-based technique we discover and exploit the symmetric structure of the object and using blending and texture-mapping we transfer the color and texture information to the unobserved parts of the object. Unlike previous approaches, our model completion does not require an object database [Pauly et al. 2005] nor assume small holes [Sharf et al. 2004], but it does produce colored models and supports at least three related families of symmetric objects. The end result is the ability to infer information about the non-visible parts of observed objects. The symmetry of an object can be exploited to aid in completing self-shadowed or poorly-sampled areas within the field of view of the images and to complete parts that are back-facing or outside the field of view.

Given an acquired object fragment, our method maps the fragment’s symmetric properties to one of our three supported types of symmetry which we find to be prevalent in many natural and man-made objects: bilateral symmetry, rotational symmetry, and surface-of-revolution symmetry. Using discovered symmetry, our algorithm replicates the necessary geometric fragment to instantiate a complete 3D model. This step involves a custom-made zippering process which does

not require an alignment process such as in iterative closest point methods (e.g., [Besl et al. 1992, Rusinkiewicz and Levoy 2001]) and handles geometry which is not purely symmetric or suffers from global deformations. In addition, under user control, we assume the object to be either solid or hollow and either close the mesh or add an inset to create a plausible interior surface for the object. We have applied our method to capture individual models ranging up to 900,000 triangles and multi-object scenes ranging up to 4.1 million triangles, and the resulting models have an average sampling resolution of about 0.25mm.

The completed object can then be used for digital inspection. Although it cannot be known if the completed object truly matches the physical object from which the fragment came, it does provide a plausible visualization that helps with reconstructing the historical context. The AES system can alter the appearance of the fragment itself, but at present cannot produce a view-independent reconstruction of the completed object. By using a background surface near the object (e.g., a wall or poster-board), the completed object can be visually produced but will only appear correct from the vantage point of the camera used for acquisition.

6. USING DECHO FOR THE WIDE DISSEMINATION OF CULTURAL HERITAGE

DECHO supports a semantically-rich and flexible visualization of CH contents. Hence, it also lends itself to support applications aiming at the dissemination of these contents to the general public, such as the research performed by the archaeologists interested in the same objects. Using the DECHO techniques, we created a web site containing panoramic images able to create a unified and seamless network of knowledge promoting a contextualized access to the Etruscan heritage. This website offers a new kind of virtual visit able to cross through several museum sites composing a semantically interconnected network of panoramas according to the thematic tour chosen by the viewer. The application is accessible both through the Web and by multimedia installations placed inside museums and near objects whose appearance is to be altered using the AES system. The panorama web framework is developed using pages dynamically generated according to the tour chosen by the viewer. Combining JSP (Java Server Page) and JavaScript, our solution generates calls to a set of web services that accordingly generate a semantic query based on the content of the selected area. The query is then forwarded to Sesame, an open source Java framework for storing, querying and reasoning with an ontology; this query is expressed in SeRQL, a RDF/RDFS query language combining traditional features of classic query languages with the

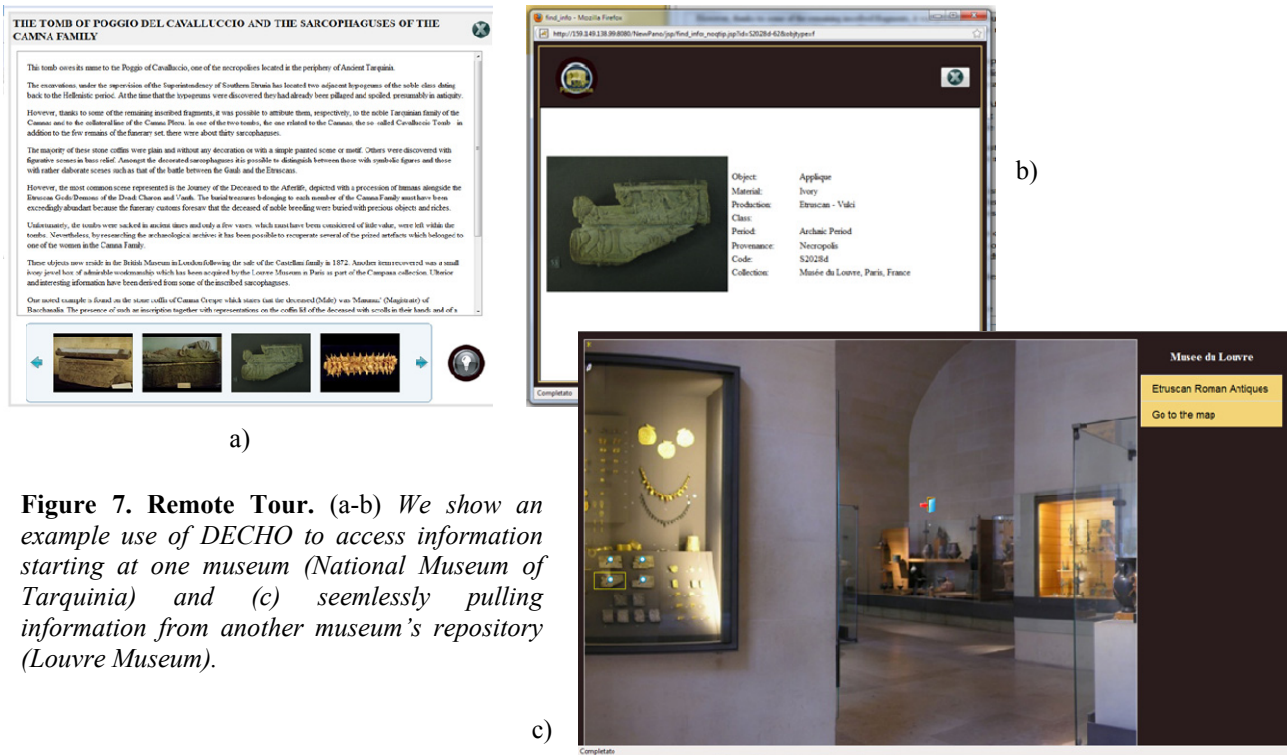


Figure 7. Remote Tour. (a-b) We show an example use of DECHO to access information starting at one museum (National Museum of Tarquinia) and (c) seamlessly pulling information from another museum's repository (Louvre Museum).

possibility of going through a semantic net. Sesame accesses the ontology representing the knowledge domain in order to enrich the panorama with mutable and flexible information according to the number and types of the databases from which to retrieve data.

Figures 2, 7, and 8 show example virtual tours. As mentioned in the introduction, Figure 2 explores the “Hall of Sarcophaguses” in the National Museum of Tarquinia. Using the kiosk (Figure 2a) and the panorama (Figure 2b), the user explores a sarcophagus (Figure 2c). By selecting the tab (red box in Figure 2c), the visitor can read relevant narrations (shown in Figure 2d); e.g., more information about the sarcophagus itself (the sarcophagus of Larth Plecu) or about an aspect of the Etruscan culture related to this object (the sarcophaguses of the Camna Plecu family of which Larth Plecu is a member, or the tomb of Poggio Cavalluccio in which the sarcophaguses of the Camna family were discovered). In Figure 2e the narration titled “The sarcophaguses of the Camna Plecu” is shown. Below the text of the narration, the DECHO system shows all objects that are related to this narration, i.e., all sarcophaguses of the Camna family. Clicking on a light bulb shaped icon, in the panorama, all sarcophaguses related to the narration and belonging to the Camna family are highlighted (Figure 2f). In this case the first two sarcophaguses belong to the Camna family and the last one does not. It is worth noting that in Figure 2e, the thumbnails of the sarcophaguses related to the narration are extracted

from the databases of several different museums and not only from the National Museum of Tarquinia's repository. By further clicking on a thumbnail, a next view appears reporting detailed information about the selected object (Figure 2g). If the panorama containing this object is present in the DECHO system an icon appears on the top left corner of this window (red box in Figure 2g). By clicking on this icon, it is possible to jump into a panorama containing this object and repeating the process (Figure 2h).

Figure 7 presents another scenario in which the visitor selects the narration titled “The tomb of Poggio Cavalluccio and the sarcophaguses of the Camna family”. Below the text of the narration, the DECHO system retrieves all objects that are in some way related to this narration, e.g., all objects discovered in the tomb of Poggio Callinari. In Figure 7a, the visitor has clicked on the thumbnail representing the ivory appliqué and a new window appears reporting further information about it (Figure 7b). This object is not in the National Museum of Tarquinia but instead in the Louvre Museum. By clicking on the proper icon, the user has the possibility to visit the room of the Louvre museum containing this object (Figure 7c).

Finally, Figure 8 showcases another example in the National Museum of Tarquinia. A user browsing a panorama clicks on a “hot-spot” representing a musical instrument: the lute (Figure 8a). As a consequence, a window appears in which the user can view a list of narrations related to the lute (Figure 8b). By selecting

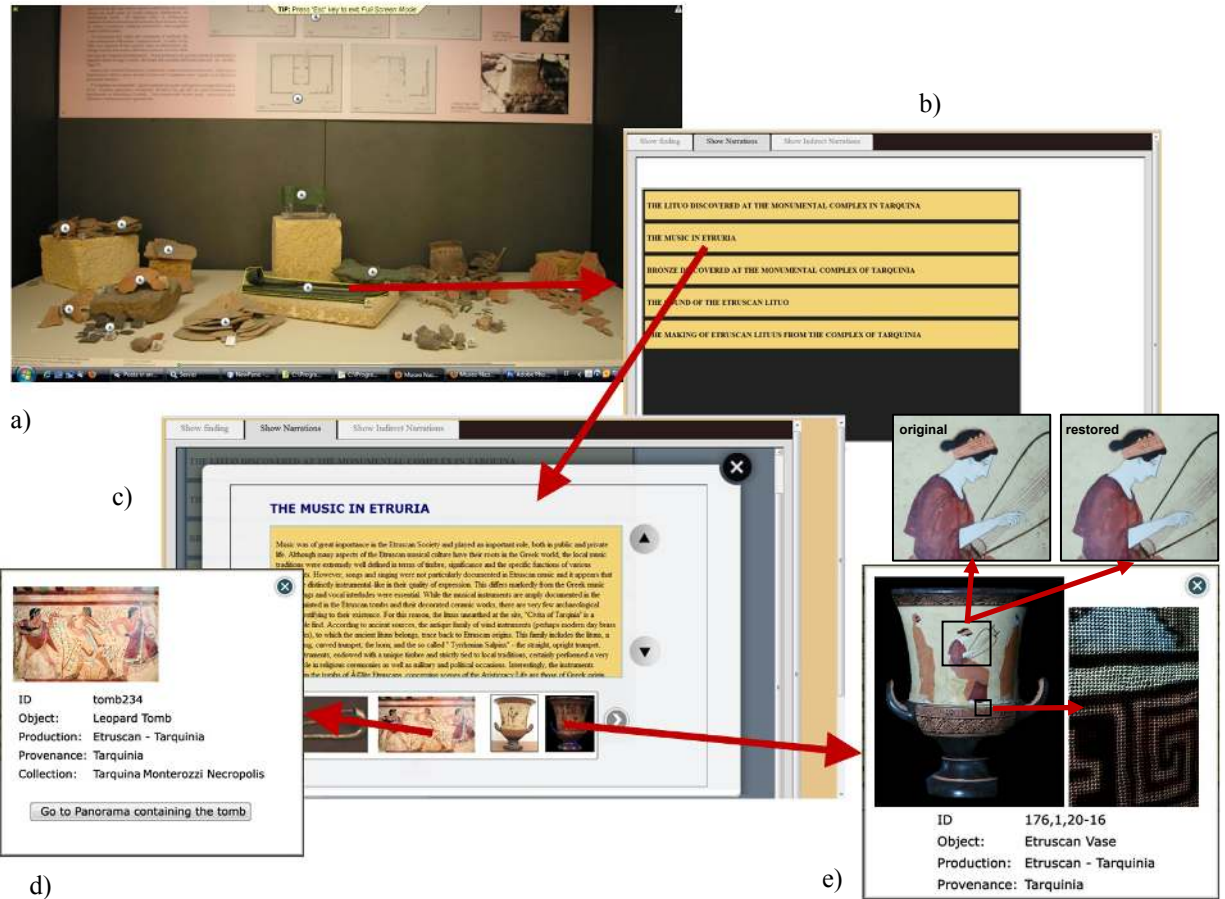


Figure 8. Personalized Exploration. We show an example result of using a subset of DECHO for cultural heritage dissemination: a) the showcase in the National Museum of Tarquinia, b) related narration, c) functions and roles of the artifact in Etruscan civilization, d) information about a related fresco, and e) 3D reconstruction and plausible restored appearance.

the narration titled “The music in Etruria” a new window pops up containing text describing the social role of music and its features and the roles and functions of musical instruments in the Etruscan age (Figure 8c). The window contains a set of pictures of related artifacts displayed at the bottom of the window (bottom row of Figure 8c): the first image shows the lute; the second image displays a fresco from the Leopard Tomb, displaying musicians playing musical instruments; the third image shows a vase on which a musician is represented; and finally the last image displays a 3D reconstruction of an Etruscan vase. Clicking on each image the user can view additional information. Figure 8d shows information about a related fresco and Figure 8e shows a 3D reconstruction, wireframe rendering, and plausible restoration of a related vase. Both images also contain information such as production and provenance. Moreover, this artifact is also related to a hot-spot in another panorama; thus by pressing the button “go to panorama

containing the tomb” it is possible to reach a panorama similar to that in Figure 8a and further explore other objects. Finally, all the objects and data are retrieved from different DBs, and allows one to see and explore artifacts (instrument, frescos, and vase) despite them being physically located at different, geographically dispersed institutions and stored in various data repositories.

7. CONCLUSIONS AND FUTURE WORK

In this paper we have discussed the key components and techniques of DECHO – a comprehensive framework supporting the semantic enrichment and exploration of CH objects. DECHO is aimed at a wide range of users from archaeologists to students and the general public. DECHO employs and extends many state of the art techniques from several fields including visualization, 3D acquisition, ontology management and database integration.

DECHO and its techniques can be extended in many different directions. The notion of the two-level ontology is quite novel and can be extended with several features. First, the ontology system can be enhanced to support the addition and modification of the mind-maps by the end-users. Archaeologists should be able to add/remove links, nodes, and narrations. Further, because many research activities are collaborative, different versions and views of a same mind-map must be supported. As such, issues related to versioning and view mechanisms for the two-level ontology must be investigated. Access control is also important since certain portions of the ontology, especially concepts, links or narrations added by a specific user, may not be visible to all users of the systems, in that they represent working hypotheses that need to be tested before being shared. Hence, mechanisms for selective access to the two-level ontology and the corresponding data have to be investigated. An access control model, based on the notion of group-based data sharing, could be applied to the two-level ontology. Access control models specific to ontologies have not been investigated so far. Finally techniques for comparing ontologies will also be devised to allow archaeologists to compare their conceptual representations of certain sets of objects. All these issues should be addressed by keeping usability as a key requirement.

Another research direction is related to the use of the archaeological ontology for integrating data sources other than relational databases. In particular, a relevant direction is to use the ontology in order to support a semantic orchestration of web-services. In this scenario the data sources are accessed using specific web-services and the OMS is used for coordinating them according to the conceptual ontology.

References

AKERS D., LOSASSO F., KLINGNER J., AGRAWALA M., 2003. Conveying Shape and Features with Image-Based Relighting, *Proceedings of IEEE Visualization*, 46-53.

ALEXANDRIAURL, 2009. Alexandria Digital Library, University of California Santa Barbara, <http://alexandria.ucsb.edu>, accessed Nov. 1, 2010.

ALIAGA D., 2008. Digital Inspection: An Interactive Stage for Viewing Surface Details, *Proceedings of ACM Symposium on Interactive 3D Graphics (I3D)*, 53-60.

ALIAGA D., XU Y., 2008. Photogeometric Structured Light: A Self-Calibrating and Multi-Viewpoint Framework for Accurate 3D Modeling, *Proceedings of IEEE Computer Vision and Pattern Recognition*, 8 pages.

ALIAGA D., LAW A., YEUNG Y. H., 2008. A Virtual Restoration Stage for Real-World Objects, *ACM Transactions on Graphics*, 27(5), 10 pages.

ALLEN P., STAMOS I., TROCCOLI A., SMITH B., LEORDEANU M., AND HSU Y. C., 2003. 3D Modeling of Historic Sites using Range and Image Data”, *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 145-150.

AZUMA R., BAILLOT Y., BEHRINGER R., FEINER S., JULIER S., AND MACINTYRE B., 2001. Recent Advances in Augmented Reality, *IEEE Computer Graphics and Applications*, 21(6), 34-47.

BAGNASCO GIANNI G. (EDITOR), ANDREADIS G., AZZAOU M., BERTOLOTTO M., BONGHI JOVINO M., CATALDI M., CHRISTIANSEN J., DOBROWOLSKI W., ELGHAZEL H., FRANZINI E., FRANZONI T., GEROLI M., HACID M.S., HAUS G., KATSAROS D., KOINIS A., LEGRAND J., LEVIALDI S., MAC AOIHD E., MCCAULEY J., MORETTI SGUBINI A.M., MOSCATI P., MUSSIO P., NIKOU I., PALLIS G., PEREGO L.G., RIDI C., STAMOS K., VAKALI A., VALTOLINA S., 2008. Bridging Archaeological and Information Technology Culture for community accessibility, [a cura di] G. Bagnasco Gianni. - [s.l.]: L'Erma di Bretschneider - (Tarchna supplemento; 2).

BARLA P., THOLLOT J., MARKOSIAN L., 2006. X-Toon: An Extended Toon Shader, *Proceedings of International Symposium on Non-Photorealistic Animation and Rendering*, 127-132.

BARTESAGHI A., SAPIRO G., MALZBENDER T., GELB D., 2005. Three-dimensional shape rendering from multiple images, *Graphical Models*, 67(4), 332-246.

BASRI R., JACOBS D., KEMELMACHER I., 2007. Photometric Stereo with General, Unknown Lighting”, *International Journal of Computer Vision*, 72(3), 239-257.

BELHUMEUR P., KRIEGMAN D., YUILLE A., 1999. The Bas-Relief Ambiguity, *International Journal of Computer Vision*, 35(1), 33-44.

BERKELEYURL, 2009. Berkeley Digital Library SunSITE, University of California Berkeley, <http://sunsite.berkeley.edu>, accessed Nov. 1, 2010.

BESL P., AND MCKAY N., 1992. A Method for Registration of 3-D Shapes, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2), 239-256.

BLOEHDORN, S., CIMIANO, P., DUKE, A., HAASE, P., HEIZMANN, J., THURLOW, I., VÖLKER, J., 2007. Ontology-Based Question Answering for Digital Libraries. In *Research and Advanced Technology for Digital Libraries*. Springer Berlin / Heidelberg, 4675, 14-25.

BUKHRES O., AND ELMAGARMID A., 1996, Object-Oriented Multidatabase Systems: A Solution for Advanced Applications, Prentice-Hall.

CASSEL, L.N., FOX, E. A., FURUTA, R., DELCAMBRE, L.M.L., 2010. Many-to-many information connection connections in a distributed digital library portal. In *Proceedings of the 10th Annual Joint Conference*

- on *Digital Libraries* (JCDL '10). ACM, New York, NY, USA, 369-370.
- CROFTS, N., DOERR, M., GILL, T., STEAD, S., STIFF, M., 2006. Current official version of the CIDOC CRM version 4.2.1, last updated October 2006, http://cidoc.ics.forth.gr/docs/cidoc_crm_version_4.2.1.pdf
- CWAREURL, 2009. Cyberware, www.cyberware.com, accessed Nov. 1, 2010.
- DAVIS J., NEHAB D., RAMAMOORTHY R., RUSINKIEWICZ S., 2002. Spacetime Stereo: A Unifying Framework for Depth from Triangulation, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(2), 296-302.
- DEVEDŽIĆ, V., 2005. Research community knowledge portals, *International Journal of Knowledge and Learning*, 1(1-2), 96-112.
- EUROPEANAURL, 2009. <http://www.europeana.eu>, accessed Nov. 1, 2010.
- FOX, E., AND MARCHIONINI, G., 2001. Digital libraries: introduction. *Communications of the ACM*, 44(5), 30-32.
- FRANKOT R., CHELLAPPA R., 1988. A Method for Enforcing Integrability in Shape from Shading Algorithms, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4), 439-451.
- FUJII, K., GROSSBERG, M. D., and NAYAR, S. K., 2005. A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments. *Proceedings of IEEE Computer Vision and Pattern Recognition*, 1, 814-821.
- FURUKAWA R. AND KAWASAKI H., 2005. Uncalibrated multiple image stereo system with arbitrarily movable camera and projector for wide range scanning, *Proceedings of 3D Digital Imaging and Modeling*, 302-309.
- GROSSBERG, M. D., PERI, H., NAYAR, S. K., and BELHUMEUR, P. N., 2004. Making One Object Look Like Another: Controlling Appearance using a Projector-Camera System. *Proceedings of IEEE Computer Vision and Pattern Recognition*, 1, 452-459.
- GRUBER, T.R., 1993. A translation approach to portable ontologies, *Knowledge Acquisition*, 5(2), 199-220.
- GUNAWARDANE P., WANG O., SCHER S., RICKARD I., DAVIS J., AND MALZBENDER T., 2009. Optimized Image Sampling for View and Light Interpolation, *Proceedings of International Symposium on Virtual Reality, Archeology and Cultural Heritage (VAST)*.
- HEMAYED E., 2003. A Survey of Camera Self-Calibration, *Proceedings of IEEE Conference on Advanced Video and Signal Based Surveillance*, 351-357.
- HERNANDEZ C., VOGIATZIS G., CIPOLLA R., 2008. Multi-view Photometric Stereo, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3), 548-554.
- LEE C., HAO X., VARSHNEY A., 2004. Light Collages: Lighting Design for Effective Visualization, *Proceedings of IEEE Visualization*, 281-288.
- LENZERINI, M., 2002. Data integration: a theoretical perspective, *Proceedings of ACM International Conference on Management of Data (SIGMOD)*, 233-246.
- LEONE, N., EITER, T., FABER, W., FINK, M., GOTTLÖB, G., AND GRECO, G., 2005. Boosting information integration: The Infomix system, *Proceedings of Italian Symposium on Advanced Database Systems (SEBD)*, 55-66.
- LEVOY ET AL. 2000. The Digital Michelangelo Project: 3D Scanning of Large Statues, *Proceedings of ACM SIGGRAPH*, 131-144.
- LIM J., HO J., YANG M., KRIEGSMAN D., 2005. Passive Photometric Stereo from Motion, *Proceedings of IEEE Conference on Computer Vision*, 1635-1642.
- MAGNENAT-THALMANN N., PAPAGIANNAKIS G., FONI A., AREVALO-POIZAT M., CADI-YAZLI N., 2004. Simulating Life in Ancient Sites Using Mixed Reality Technology, *Proceedings of Congreso Español de Informática Gráfica (CEIG)*, Seville, Spain, May 2004.
- MAY, W., 2005. Logic-based XML data integration: a semi-materializing approach, *Journal of Applied Logic*, 3(2), 271-307.
- MAZZOCCHI, S., GARLAND, S., LEE, R., 2006. Simile: Practical metadata for the semantic web. *XML.com*.
- MAZZOLENI P., VALTOLINA S., FRANZONI S., MUSSIO P., BERTINO E., 2006. Towards a contextualized access to the cultural heritage world using 360 panoramic images, *Proceedings of the Eighteenth International Conference on Software Engineering & Knowledge Engineering (SEKE)*, San Francisco, CA, July 5-7.
- MENA, E., ILLARRAMENDI, A., 2001. *Ontology-Based Query Processing for Global Information Systems*, Kluwer Academic Publishers, ISBN 0-7923-7375-8, 215.
- MITSUNAGA, T., NAYAR, S. K., 1999. Radiometric Self Calibration. *Proceedings of IEEE Computer Vision and Pattern Recognition*, 1, 374-380.
- MONROY, C., FURUTA, R., AND CASTRO, F., 2010. Using an ontology and a multilingual glossary for enhancing the nautical archaeology digital library. *In Proceedings of the 10th Annual Joint Conference on Digital Libraries (JCDL)*. ACM, New York, NY, USA, 259-262.
- NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHY R., 2005. Efficiently Combining Positions and Normals for Precise 3D Geometry, *ACM Transactions on Graphics*, 24(3), 536-543.
- PAPAGIANNAKIS G., PONDER M., MOLET T., KSHIRSAGAR S., CORDIER F., MAGNENAT-THALMANN N., AND THALMANN D., 2002. LIFEPLUS: Revival of Life in Ancient Pompeii, *Proceedings of Virtual Systems and Multimedia (VSMM)*, 25-27.
- PARK J., AND KAK A., 2008. 3D Modeling of Optically Challenging Objects, *IEEE Transactions on*

- Visualization and Computer Graphics*, 14(2), 246-262.
- PAULY M., MITRA N., GIESEN J., GROSS M., GUIBAS L., 2005. Example-based 3D Scan Completion, *Proceedings of Symposium on Geometry Processing*, 23-32.
- PAYETTE, S., STAPLES, T., 2002. The Mellon Fedora Project: Digital library architecture meets XML and Web services. In: Calzarossa, M.C., Tucci, S. (eds.) Performance, LNCS, Springer, Heidelberg, 2459, 406-421.
- PERSEUSURL, 2009. Perseus Digital Library, Tufts University, <http://www.perseus.tufts.edu/hopper>, accessed Nov. 1, 2010.
- POLLEFEYS M., VAN GOOL L., VERGAUWEN M., CORNELIS K., VERBIEST F., TOPS J., 2001. Image-based 3D Acquisition of Archaeological Heritage and Applications, *Proceedings of Virtual Reality, Archaeology, and Cultural Heritage*, 255-262.
- POLLEFEYS M., VAN GOOL L., VERGAUWEN M., VERBIEST F., CORNELIS K., TOPS J., AND KOCH R., 2004. Visual modeling with a hand-held camera, *International Journal of Computer Vision*, 59(3), 207-232.
- RAHM E., AND BERNSTEIN P.A., 2001, A Survey of Approaches to Automatic Schema Mapping, *VLDB Journal*, 10(4), 334-350.
- RAVINDRANATHAN U., SHEN R., GONÇALVES M. A., FAN W., FOX E., AND FLANAGAN J., 2004. ETANA-DL: managing complex information applications - an archaeology digital library, *Proceedings of the Fourth ACM/IEEE Joint Conference on Digital Libraries*, 7-11.
- RASKAR, R., WELCH, G., LOW, K. L., AND BANDYOPADHYAY, D., 2001. Shader Lamps: Animating Real Objects With Image-Based Illumination. *Proceedings of Eurographics Workshop on Rendering Techniques*, 89-102.
- RASKAR R., VANBAAR J., BEARDSLEY P., WILLWACHER T., RAO S., FORLINES C., 2003. iLamps: Geometrically Aware and Self-Configuring Projectors, *Proceedings of ACM SIGGRAPH*, 10 pages.
- RASKAR R., TAN K., FERIS R., YU J., TURK M., 2004. Non-photorealistic Camera: Depth Edge Detection and Stylized Rendering using Multi-flash Imaging, *Proceedings of ACM SIGGRAPH*, 10 pages.
- RDFURL. 2009, http://www.w3schools.com/RDF/rdf_intro.asp, accessed Nov. 1, 2010.
- RUSHMEIER H., AND BERNARDINI F., 1999. Computing Consistent Normals & Colors from Photometric Data, *Proceedings of 3D Digital Imaging and Modeling*, 99-108.
- RUSINKIEWICZ S., LEVOY M., 2001. Efficient Variants of the ICP Algorithm, *Proceedings of 3D Digital Imaging and Modeling*, 145-152.
- RUSINKIEWICZ S., BURNS M., DECARLO D., 2006. Exaggerated Shading for Depicting Shape and Detail, *Proceedings of ACM SIGGRAPH*, 1199-1205.
- SCHARSTEIN, D. AND SZELISKI, R., 2003. High-Accuracy Stereo Depth Maps Using Structured Light. *Proceedings of IEEE Computer Vision and Pattern Recognition*, 195-202.
- SESAMEURL, 2010. User Guide for Sesame. Updated for Sesame release 1.2.3., <http://www.openrdf.org/doc/sesame/users/index.html>, accessed on Nov. 1, 2010.
- STURM P. 2002. Critical motion sequences for the self-calibration of cameras and stereo systems with variable focal length, *Image and Vision Computing*, 20(5-6), 415-426.
- TARCHNAURL, 2009. http://homes.dico.unimi.it/~valtolin/tarchna_brochureEnglish.pdf, accessed Nov. 1, 2010.
- THEBBANURL, 2009. Thebban Mapping Project, American University in Cairo, <http://www.thebanmappingproject.com>, accessed Nov. 1, 2010.
- THOMAS, J.C., KELLOGG, W.A., ERICKSON, T., 2001. The knowledge management puzzle: human and social factors in knowledge management, *IBM Systems Journal*, 40(4), 863-884.
- VALTOLINA S., MUSSIO P., MAZZOLENI P., FRANZONI S., BAGNASCO G., GEROLI M., RIDI C., 2007. Media for Knowledge Creation and Dissemination: Semantic Model and Narrations for a New Accessibility to Cultural Heritage, *Proceeding of the 6th Creativity & Cognition Conference*, 107-116.
- VALTOLINA S., 2008. Design of Knowledge Driven Interfaces in Cultural Contexts, *International Journal on Semantic Computing - Special Issue on Human Centric Communications*, 5(3), 525-553.
- VOTE E., ACEVEDO D., JOUKOWSKY M., AND LAIDLAW D., 2000. ARCHAVE - A Three Dimensional GIS for a CAVE Environment, *Proceedings of CAA: Computing Archaeology for Understanding the Past*.
- WETZSTEIN, G., AND BIMBER, O. 2007. Radiometric Compensation through Inverse Light Transport. *Proceedings of Pacific Graphics*, 391-399.
- WEXLER, M.N. 2001. The who, what and why of knowledge mapping, *Journal of Knowledge Management*, 5(3), 249-263.
- WILLIAMS N., HANTAK C., LOW K.L., THOMAS J., KELLER K., NYLAND L., LUEBKE D., AND LASTRA A., 2003. Monticello Through the Window, *Proceedings of the 4th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*.
- WINNEMOLLER H., MOHAN A., TUMBLIN J., GOOCH B., 2005. Light Waving: Estimating Light Positions from Photographs Alone, *Computer Graphics Forum (Eurographics)*, 24(3), 433-438.
- WOODHAM R., IWAHORI Y., BARMAN R., 1991. Photometric Stereo: Lambertian Reflectance and Light Sources with Unknown Direction and Strength, *University of British Columbia*, TR-91-18.

- XU Y., AND ALIAGA D., 2009. An Adaptive Correspondence Algorithms for Modeling Scenes with Strong Inter-reflections, *IEEE Transactions on Visualization and Computer Graphics*, 15(3), 465-480.
- ZHANG L., CURLESS B., HERTZMANN A., SEITZ S., 2003a. Shape and Motion under Varying Illumination: Unifying Structure from Motion, Photometric Stereo, and Multi-view Stereo, *Proceedings of IEEE Computer Vision and Pattern Recognition*, 618-625.
- ZHANG L., CURLESS B., SEITZ S., 2003b. Spacetime Stereo: Shape Recovery for Dynamic Scenes, *Proceedings of IEEE Computer Vision and Pattern Recognition*, 367-374.