A platform for virtual museums
with personalized content

**B. Bonis · J. Stamos · S. Vosinakis · I. Andreou ·
T. Panayiotopoulos**

Abstract Presentation of content is an important aspect of today's virtual reality applications, especially in domains such as virtual museums. The large amount and variety of exhibits in such applications raise a need for adaptation and personalization of the environment. This paper presents a content personalization platform for Virtual Museums, which is based on a semantic description of content and on information implicitly collected about the users through their interactions with the museum. The proposed platform uses stereotypes to initialize user models, adapts user profiles dynamically and clusters users into similar interest groups. A science fiction museum has been set up as a case study for this platform and an evaluation has been carried out.

Keywords Virtual reality · Virtual museums · User modeling · Personalization

1 Introduction

The advances in 3D graphics and in specialized Virtual Reality (VR) hardware during the last decade have enabled the development of novel interactive and immersive applications that emphasize on the realistic representation of content and present a satisfying and intuitive experience to the users. An application area that is based on content presentation and can benefit from this emerging technology is Virtual Museums, i.e. computer generated environments that present exhibit collections from real or fictional museums and aim to educate and entertain users by offering them an experience similar to an actual museum visit. The notion of virtual museums has been introduced by Tsichritzis and Gibbs [36] as a

B. Bonis · J. Stamos · T. Panayiotopoulos
Department of Informatics, University of Piraeus, 80 Karaoli & Dimitriou str. 185 34, Piraeus, Greece

S. Vosinakis (✉)
Department of Product and Systems Design Engineering, University of the Aegean,
Hermoupolis, Syros 85100, Greece
e-mail: spyrosv@aegean.gr

I. Andreou
Department of Technology Education and Digital Systems, University of Piraeus,
80 Karaoli & Dimitriou str. 185 34, Piraeus, Greece

means to overcome the limitations of real museums and to enhance user experience. A synthetic collection of artifacts, which incorporates multimedia and virtual reality technologies, alleviates the problem of storing, preserving and protecting the real artifacts and allows synthetic museums to present an unlimited number of exhibits, to which users have access at any time and from any place. Furthermore, it may introduce new forms of presentation and interactivity that stand beyond the passive viewing of the artifacts and reading the description text, which is typically the case with traditional museum visits. Digital artifacts may be presented using a combination of various forms of media, such as 3D representations and rich hypermedia annotations and may also let the user interact with them in many intuitive and creative ways in order to learn and entertain themselves. Examples of rich interactive experiences could be the utilization, assembly and disassembly of mechanical artifacts in a science museum, or the inclusion of educational games that are thematically related to the museum collection.

A number of Virtual Museum applications have been presented, which run standalone [20] or over the Internet [6, 27, 35]. They serve either as complementary information source to existing museums or as individual approaches for the presentation of artifact collections [6]. There is a great diversity in terms of visualization and user interactivity in the available approaches, as a multitude of different technologies have been employed. As per the presentation of artifacts, the approaches include simple images, panoramic views, video and hypermedia presentations, and detailed 3D models. In terms of user interface, a variety of systems have been presented, ranging from a sequence of 2D pages containing the artifacts to immersive 3D environments in which users can navigate, explore the virtual space and get haptic feedback from the exhibits [16], using specialized hardware. The interaction modalities vary from simply viewing the artifacts to letting users have rich educative interactions with them [29]. Concerning implementation technology, virtual museums have been developed either as standalone multimedia applications or as web-based environments using various standards, such as Flash, Quicktime VR, VRML (Virtual Reality Modeling Language) and X3D. Finally, virtual museums can be single-user or multi-user environments, and in some cases they are presented as part of massive multi-user online worlds [7, 37].

It can be noted that most of these applications emphasize on the content and aim to the deeper understanding of entities and concepts through user navigation and interaction in 3D. However, the majority of them contain static collections arranged in predefined positions, and the design of the virtual space and its contents is based entirely on the author's point of view. As a result, the user's role is limited to a passive observer and the presentation of a large amount of artifacts may fall into the obstacle of navigational difficulties in 3D environments compared to reality, e.g. due to the lack of visual cues of distance, movement, direction and binocular vision [32, 38]. This may eventually lead to the reduction of user interest and to the disability to explore and search for the desired content [24]. Additionally, the complexity of designing static exhibitions increases as the amount of exhibits grows.

The problem of presenting and categorizing large quantities of content has been effectively addressed in Web [11, 23, 25, 34] and multimedia applications [30]. In these cases, user modeling techniques have been employed in order to personalize the content presentation according to the users' own interests. The authors claim that virtual environments could also benefit from user modeling and adaptation methodologies, in order to make assumptions about user interests and intentions concerning the application, and to construct the virtual space accordingly. Such a personalized space is expected to reduce the navigational burden and still retain the metaphor of being immersed in a 3D environment.

This paper presents a novel approach towards adaptive virtual museums, in which artifacts are dynamically distributed among rooms and users may experience personalized presentations based on assumptions drawn implicitly by their previous interactions with the system and explicitly by their own feedback. Additionally, the authors support the claim that users' experience in a Virtual Museum is enhanced through communication and collaboration with other users in a shared environment [17]. In this context, the proposed approach supports dynamic clustering of user groups based on the degree of similarity of user profiles, which may lead to the formation of e-societies with common interests. The authors have developed a platform for the creation of web-based multi-user adaptive 3D virtual museums that incorporates the above-mentioned characteristics and facilitates the design, maintenance and enrichment of large museum collections. A science fiction virtual museum has been set up as a case study for this platform and a user evaluation has been conducted. Evaluation results indicate a significant degree of user satisfaction from the system in general and a positive impact of the user modeling process on content selection.

The rest of the paper is structured as follows: Section 2 presents the related work towards adaptive presentation of content in virtual museums; Section 3 presents in detail the design, architecture and implementation of the proposed platform for adaptive virtual museums, emphasizing on the data representations and processes that drive the content adaptation and on the functionality provided to the designer and the visitors; Section 4 presents the set up of a science fiction virtual museum as a case study using the proposed platform and Section 5 presents an evaluation of the implemented case study. Finally, Section 6 draws some conclusions and presents the future work.

2 Virtual museums and adaptivity

The traditional roles of virtual museums and exhibitions are to provide a public space for the exhibition of artifacts and to serve as a centre of knowledge specialized in a thematic area [10]. The typical goals of virtual museums visitors could then be identified as: the exploration of the foyers, the browsing of the exhibits they are interested in, the interaction with them in order to learn by their experience, and the exchange of opinions with other visitors about the content and the related knowledge area. Therefore, virtual museums should facilitate navigation, exploration and communication by using intuitive user interfaces and navigation aids. The case of traditional 2D interfaces limits the user experience to simple page viewing and sequential browsing and leaves no room for any immersion. On the other hand, a 3D representation of the museum environment places the artifacts in a natural-looking setting and may offer a much more realistic and entertaining experience [19], provided that the environment will be enhanced with a number of navigational queues, such as mini-maps, landmarks and teleportation, in order to tackle the well-known navigational problems of virtual environments [1, 8].

Museum artifacts are usually distributed in foyers and display cases following some inherent categorization proposed by the museum curator. However, visitors have varied interests; some may be specifically interested in a subset of the categories, e.g. a student doing research in a specific historical period, and some may have broader interests, or may just wish to wander around, until they find something that captures their attention. Therefore, the distribution of artifacts in predetermined placeholders in real museums may not match all visitors' expectations, as it is practically impossible to present different subsets of the collection to individual visitors. Virtual museums, on the other hand, have no such limitations, as the content can be dynamically distributed and rearranged resulting to

an infinite number of varied presentations of the same museum collection. Furthermore, it is significantly hard to extend or alter a real museum collection, at least on a regular basis, whilst virtual museums can offer the ability to perform instant changes on the collection presentation and to expand the museum space infinitely. As a result, virtual museums may enhance traditional museum visits by offering the ability to adapt, expand and personalize the artifact collections.

According to Kobsa et al. [18] the personalization process is divided in three major tasks: the acquisition of information about users' characteristics and behavior, the representation of that data in a formal system that allows the drawing of assumptions about user preferences, and the production of personalized content. A number of successful adaptive systems have been developed using web-based and hypermedia technologies that provide personalized content to users. However, in the case of 3D environments, the tasks of recoding user behavior and dynamically producing the content introduces a significant amount of complexity, due to the fact that the degrees of freedom in navigation and interaction are significantly increased, compared to navigating and interacting with page-based content. A first approach towards personalized 3D exhibition environments is JubilEasy [33], a system that generates virtual visits in the city of Rome. The adaptation process is using as input a set of VRML models and images, the city topology and information received explicitly by the user, and derives a plan that is being used to generate the virtual visit.

Chittaro and Ranon [4] propose an approach for the development of adaptive 3D Web sites. They present the AWE3D architecture for the generation of dynamic VRML worlds that adapt to user preferences. Their proposed data acquisition process is based on sensors that record proximity and visibility to infer whether a user has seen an element of the environment or how much time she has spent in a place. User data are submitted to the server, recorded in a database, and processed by a rule-based system that updates the user model. The personalized environment is generated using a world generator process that selects elements from a VRML content database and generates the personalized environment. AWE3D has been used for the development of a personalized 3D e-commerce site and for automatically generated tours in Virtual Museums [5]. Another approach towards personalization in virtual environments is the intelligent virtual environment presented in [9], which adapts its structure and presentation according to the visitor preferences by inserting and deleting content. It uses a rule-based system with certainty factors to draw inferences and to update the user model accordingly. Furthermore, it contains an automated content categorization system based on machine learning techniques that aims to assist the spatial distribution of content in the environment. The system has been used in a prototype distant learning environment.

Celentano and Pittarello [3] propose an approach to facilitate adaptive interaction with the virtual environment, which is based on the following: a structured design of the 3D interaction space, the distinction between a basic virtual world layer and an interaction layer, and the recording of the environment's usage by the user in order to find interaction patterns. The aim is to facilitate the system's usage by monitoring user behavior and predicting future needs for interaction purposes. When the system recognizes the initial state of an interaction pattern, it executes the final state without letting the user engage in the intermediate ones.

Our proposed platform, PeVEP (Personalized Virtual Exhibition Platform), consists of a set of generic tools for the design and deployment of personalized virtual museum applications, based on an architecture that separates the content from the exhibition rooms and allows the dynamic generation of virtual exhibitions. The system functionality is

supported by a semantic graph, defined by the designer, which contains an ontological description of the content and drives the user modeling process. Therefore, the environment is context-independent, adaptive and provides a rich user experience.

3 The PeVEP platform for content personalization

The PeVEP platform for designing and implementing virtual exhibitions with content personalization is based on four methods that enhance a static 3D environment with dynamic characteristics: user model generation, content presentation, user model update and user clustering. The retrieval approach can be classified as Content Based [2]. In the next paragraph a full user interaction session with an application based on this platform is presented.

   When a new user enters the environment for the first time, a user model is assigned to her based on stereotypes [31] that correspond to her selection of an avatar, i.e. her graphical representation in the 3D environment. While the user is browsing the environment, her navigation and interaction with content are monitored and the recorded behavior is utilized to make assumptions about her interests and preferences, which are then incorporated into the user's profile. At any time, the user can ask to be transported to a personalized environment, which reflects her assumed preferences and recommends new content that might be of her interest. The user can also join communities with similar preferences, visit other personalized environments, and exchange opinions about the content. User interest groups are proposed by the environment through an automated clustering process.

   From the designer's point of view, the platform can be employed to construct new dynamic virtual museums without having to define explicit rules for content personalization and adaptation. The designer has to provide the 3D content, i.e. the rooms and objects of the environment, the semantic graph, i.e. an ontological description of the content, and the user stereotypes that contain templates of estimated initial user preferences concerning the content. A presentation process then creates the exhibition rooms and distributes the exhibits dynamically based on the above data. The personalized environments depend on the interaction history of the respective users. Furthermore, exhibitions generated using the proposed framework can be easily adapted or enhanced by altering or inserting new 3D content and making appropriate changes in the semantic graph and/or the user stereotypes.

   PeVEP is based on a client-server architecture (Fig. 1) and has been implemented using Web technologies, i.e. VRML, Java, Java3D, and TCP sockets. The users interact with the environment on the client side and user modeling takes place on the server side. This schema follows the paradigm of decentralized user modeling architectures [12], and allows clustering of users into groups with similar preferences. This paradigm is necessary to support multi-user environments and for the immediate adaptation and expansion of exhibition data by the designer. The rest of this section describes the architectural components in detail.

3.1 The semantic graph

A vast number of applications that utilize user modeling methodologies try to address the user's need for quick and efficient access to a subset of information that meets her interests and preferences, without having to search through a larger set of objects. A widely used term in the literature for describing these applications is recommender systems [2, 15, 25, 28, 30, 31, 34]. A distinctive characteristic of these systems, compared to information
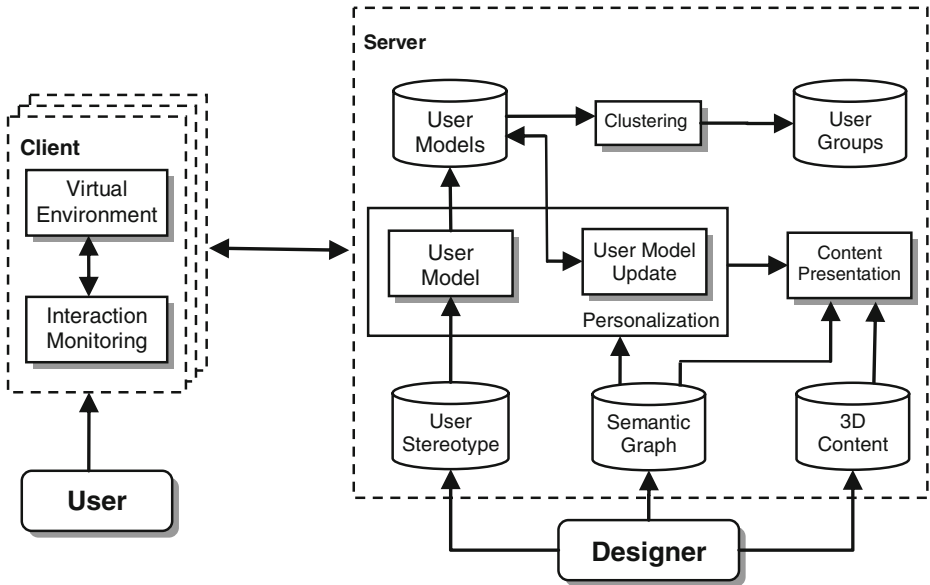
Fig. 1 The architecture of the PeVEP platform

retrieval and filtering systems and search engines, is the output of individualized information based on a priori knowledge about the content and assumptions about user preferences.

A thematically uniform set of objects can be grouped together and categorized based on a number of criteria; relations can be determined between objects and categories or between categories themselves, e.g. relations of affinity and inheritance. For example, in an art exhibition, exhibits can be grouped with respect to their creators, the epoch or the style. These categories can be generalized into broader categories or specialized into subcategories. A categorical hierarchy of this type forms a tree with nodes being the categories and edges being the relations between them. The entirety of the categories can thus be represented as a forest (a set of distinct trees). Parent nodes in each tree imply categories with broader meaning than their children, and the respective relations can be viewed as inheritance relations. Nodes of different trees can be connected implying a semantic union relation. Because of the union relations this categorization scheme is used as a graph instead of a forest in the context of determining related categories. Let level 1 be the level of the most specialized categorization nodes for each tree and the parent of each node at level N be placed at level $N+1$. The actual objects are attached, as nodes, to the categorization trees using connection(s) with one or more level 1 category nodes via an instance relation. Thus, the resulting structured hierarchical semantic taxonomy forms a directional graph, the Semantic Graph (SG), in which nodes (distributed into levels) stand for objects and concepts, edges represent the relations between them [13, 15, 22] and the levels represent the degree of generalization. In this taxonomy, the actual objects are considered as level 0 nodes, although they do not necessarily belong to only one tree.

Figure 2 presents a sample part of a semantic graph that is used for categorizing movies-related content. The dashed line is a union relation that connects two categories belonging to different trees. In the example the node Genre is divided into two subcategories, the nodes Science Fiction and Comedy. The Extra Terrestrial appears in a science fiction movie
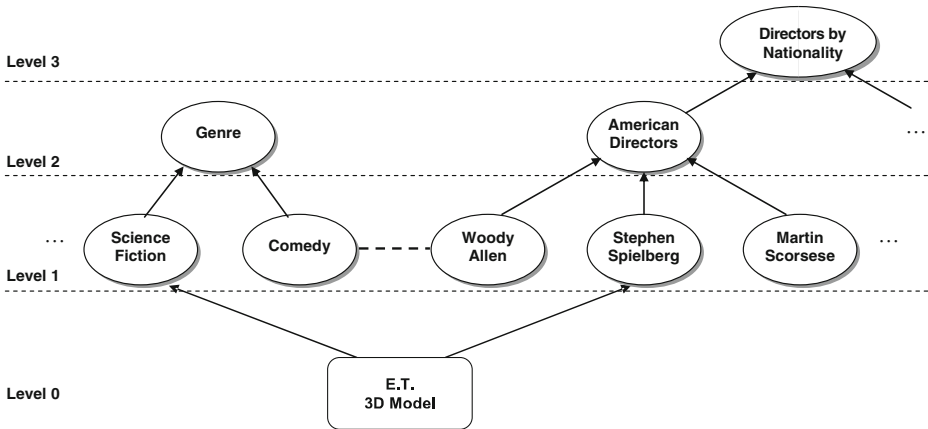
Fig. 2 A part of a semantic graph for categorizing movies. Level 0 is the objects level

directed by Stephen Spielberg. Thus, its 3D representation can be connected to the nodes 'Science Fiction' and 'Stephen Spielberg'. The union relation between 'Woody Allen' and 'Comedy' nodes implies an affinity association between the famous director and the comedy genre.

Let A be the node at the lower end of a categorization edge and B the node at the upper end. All edges (which represent the aforementioned relations) have a numerical weight in the range of (0, 1], which is the degree of membership of the object or concept that node A represents to the set that the node B portrays. The use of the degree of membership is analogous to the respective term from the Fuzzy Set Theory [22]. For example, Stephen Spielberg can be said to be a well known American Director, so the degree of membership of this director in the 'American Directors' category is significant. In the case of union edges, the weights represent the degree of association between the connected nodes. The choice of weights can be made by using expert knowledge or by using machine learning techniques. The degrees of membership and the degrees of association are used during the execution of the personalization-recommendation algorithm.

In the PeVEP platform, the SG is the core element, because it drives most computational procedures:

- modeling user preferences in the environment,
- providing recommendations to users,
- clustering (and searching for) users with similar preferences, and
- dynamically distributing the content into separate rooms and determining the connections between them.

The authors argue that the SG facilitates the access to the exhibition's content repository, by reflecting a natural content interpretation, effectively serving the users' informational needs and preferences.

3.2 Dynamic content generation and rendering

The hierarchical structure of the SG is used for the creation of the virtual exhibition's spatial structure. Each categorization node can be represented by a set of interconnected rooms. Rooms that represent conceptually relevant nodes (as portrayed in the graph) are also

connected via doors. This approach provides a multidimensional navigation paradigm, in which rooms are connected based on their semantic similarity. A user can navigate inside the exhibition environment and browse the content by following paths that correspond to any of the concepts that can characterize it. E.g. a room with objects from Steven Spielberg movies is connected with the room of American directors (at a higher level). By dynamically structuring the environment, the effort of designing and creating every room in the exhibition is reduced to a minimum.

The designer provides a number of template rooms that are used by the client application to construct the actual exhibition rooms. The system requires a default template room and, optionally, any number of thematic rooms related to existing categories of the semantic graph. Each template room is a 3D model of a section of the museum environment, in which two types of objects are inserted: doors and exhibit containers. When the client application has to construct a thematic room, it requests from the server to search for a predefined template that matches the respective category. In the case that no such template exists, the default room is used. Then, the system dynamically links the room with related ones using doors, each labeled by the name of the room it leads to. Finally, the exhibits are dynamically placed in the containers whose positions are defined by the designer. A default exhibition is constructed by creating an entrance room that is connected to all the general (top-level) category rooms of the semantic graph.

The collection of objects that will populate each room is generated by traversing the graph downwards, from the respective categorization node to the set of object nodes. The number of objects can be significantly large, especially when browsing the content of higher-level nodes. In that case, objects are distributed into a set of interconnected rooms, each containing a subset of the total exhibits of the same node. In a real exhibition, thematically related exhibits are, ideally, placed in adjacent rooms. Similarly, the doors in the virtual exhibition connect rooms that belong to the same categorization node, the parent node, or sibling categorization nodes i.e. nodes with the same parent node in the SG. The metaphor of a navigation panel provides access to rooms that represent lowest levels of categorization i.e. these rooms represent specialized categorization nodes that are connected via a generalization edge to the current room's node. This approach is used in order to avoid overwhelming the user with a potentially vast plethora of doors. A navigation panel can also be viewed as a dynamic door, since the room that the user will be lead depends on her choice.

Figure 3 sketches a room in the virtual museum, which could have been generated for a part of the graph in Fig. 2. This room is the second thematic room that corresponds to the "American Directors" categorization node, which is a level 2 categorization node. The room contains exhibits that are related (indirectly) with the room's node, through that node's children, which correspond to American directors. The user can navigate sequentially to the rest of the rooms that belong to this categorization node, through doors that exist in this room, i.e. there are two doors in this room that lead to the next and previous rooms for this node, since more than one room was generated for the "American Directors" node.

In order to facilitate horizontal navigation, each node should be linked to the previous and next sibling nodes (if any). This is implemented using doors in the corresponding rooms, i.e. each room of a node contains two doors that lead to the first room of the previous and next sibling nodes. In the diagram, the previous and next sibling rooms are the "French Directors #1" and "British Directors #1" rooms. There is also a door that leads to a room of the parent node, the "Directors" node. Finally, as the "American Directors" node is not a lowest level categorization node, a navigation panel exists, to provide options for moving to more specialized rooms, i.e. rooms of specific American directors.
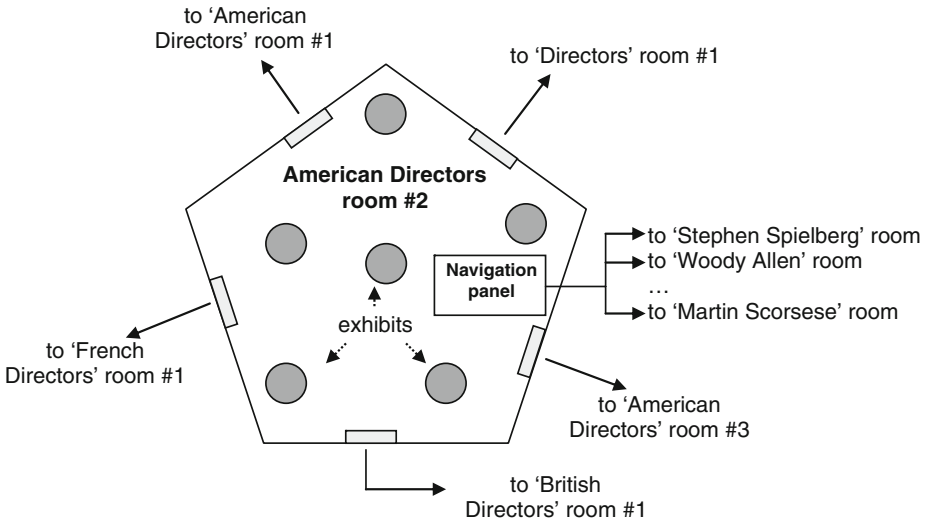
Fig. 3 Schema of a thematic room, generated from a semantic graph that categorizes movies

This arrangement allows every room, regardless of generalization (level) to contain exhibits. However, the rooms of the lowest level nodes provide the greatest degree of specialization. While navigating, the user moves between rooms of the same generalization level, as she would move inside a real exhibition. She uses the parent door and the navigation panel in order to move between different levels of generalization (as they are represented by levels in the graph) and browse rooms using different subtrees of the graph.

3.3 Personalization

All available objects in the PeVEP platform are categorized by the SG in a semantic taxonomy. When and if a user is interested in a certain object, it can be assumed that she is also interested in one or more categories to which the object belongs (she likes the movie series Star Wars because she likes sci-fi movies). If this belief is reinforced during the interaction with the system by showing a tendency to interact with similar objects, a recommendation set with members originating from this particular category would probably be a preferable choice. The user model is an instance of the SG, with all nodes being given a numerical value (positive or negative), which represents the degree of interest that the user is assumed to have for the term or object that the node portrays. The process of calculating these degrees of interest is analytically explained in the next paragraphs.

3.3.1 User stereotypes

As stated in [32], the explicit creation of a user profile may annoy users that are unwilling to state their interests and to provide information about themselves and thus lead to user models that do not actually reflect the user preferences. Therefore, an indirect method for creating an initial user profile has been utilized, in which a set of stereotypes [31] is used to initialize the model of a new user. At registration time, the user selects an avatar from a provided library. Each avatar is related with assumptions about a user, which are lexical values of properties defined by the designer. A possible set of properties can be age, sex,

education or anything that is considered to characterize the users and the respective values can be young, old, female, male, high, low etc. The stereotypes are rules that relate each value of each property with estimated degrees of interest for a set of nodes in the SG.

The degree of interest in the categorization nodes is calculated as follows. Initially all the categorization nodes have a zero value. For each value of each property, the degree of interest declared in a stereotype (if any) is added to the respective categorization node. This initial user profile is used for the formation of a recommendation set (which will be explained in the following paragraphs) prior to the user interaction with the 3D environment. This approach deals with the new user problem [2], with a reduced accuracy nevertheless. This is an initial estimation, however, and as the user interacts with the system, information is accumulated in the user model, updating it and thereby increasing its accuracy.

The aforementioned methodology is based on researches [14, 26] that state the possible relation between the choice of an avatar by the user and the users' intrinsic characteristics and personality trades. Expert knowledge and/or user segmentation researches can be used to create the set of stereotypes.

3.3.2 Data acquisition and user profile updating

The algorithm that updates the user profile based on her interactions uses a mechanism that propagates from the lowest levels of the semantic graph towards the upper levels, i.e. an upward propagation mechanism. Initially, it collects the summing degree of interest for each object node of the SG that the user has interacted with. This is the result of a monitoring process that takes place on the client side. A variety of acquisition methods is used, such as measuring the time spent by the user observing an object, taking into account the type of interaction, and providing a rating system that lets users express their preferences. The framework supports various types of acquisition methods and interpretations. At the end of this step, all object nodes have a degree of interest ranging from negative values (dislike for the particular object) to positive values (positive interest).

In order to include the union connections in the propagation mechanism, these connections are implemented as follows: each union relation in the SG is represented with a dummy node (union node), which is placed at one level higher than the maximum level of the categorization nodes it connects. These nodes are connected to the union node with dummy relations that have weights equal to w, where w^2 equals to the degree of association of the union relation. For the rest of the paper we refer to both categorization and union nodes with the term *semantic nodes*.

Let SN be a semantic node and $DI(SN)_t$ the degree of interest in node SN at time instance t. Let $CSN=\{CSN_1, CSN_2, \ldots, CSN_N\}$ be the set of children nodes of node SN, with $DI(CSN_i)_t$, $i \in [1,N]$ being the respective degrees of interest. Let also W_i be the weight of the edge that connects CSN_i to SN. Then, the degree of interest at time instance $t+1$ for node SN is calculated by:

$$DI(SN)_{t+1} = \sum_{i=1}^{N} \left(DI(CSN_i)_t \cdot W_i \right) + DI(SN)_t$$

Calculations begin at level 0 (the object nodes level) and they proceed level by level until the maximum level of the SG. Object nodes maintain the values from the initial step. The resulting degrees of interest comprise the updated user profile. This process is depicted in Fig. 4.
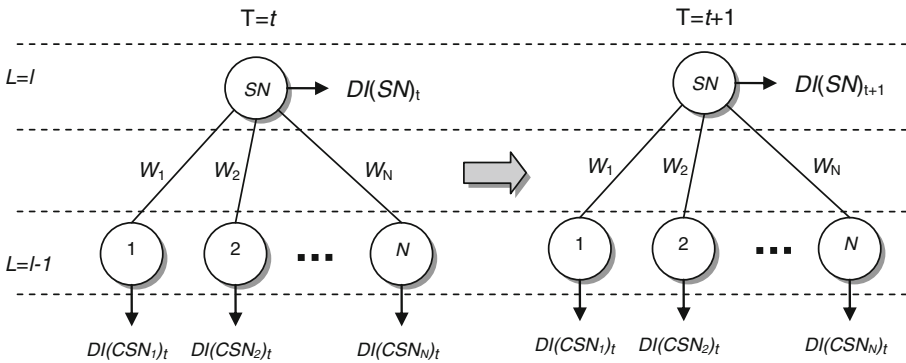
Fig. 4 Updating the user profile

3.3.3 Adaptation of the environment

For the creation of a personalized room, a set of objects are chosen to be recommended to the user by populating the room with them. This task is realized using a downward propagation mechanism based on the user model. During this process, an estimated degree of interest is assigned to the object nodes. The choice of objects that will be part of the user's personalized room depends on the ratings of the respective object nodes. The M top rated nodes of level zero (the level with the object nodes), where M is the capacity of the room, are chosen. Variations of the proposed exhibits can be produced by employing methods, such as to insert a few random elements to increase the diversity, or to avoid selecting objects that the user has already interacted with.

Let us describe the downward propagation mechanism. The following computations begin from the top level in the graph and reach level 1 (lowest concept nodes level). Let SN be a semantic node and $DI(SN)_t$ the degree of interest in node SN at time instance t. Let $PSN = \{PSN_1, PSN_2, \ldots, PSN_N\}$ be the set of parent nodes of node SN, with DI $(PSN_i)_t$, $i \in [1,N]$ being the respective degrees of interest. Note that the values of the nodes in the upper levels have already been calculated for time instance t. Let also W_i be the weight of the edge that connects PSN_i to SN. Then, the degree of interest at time instance t for node SN is calculated by:

$$\mathrm{DI(SN)}_t = \sum_{i=1}^{N} \left(\mathrm{DI(PSN}_i)_t \cdot W_i \right) + \mathrm{DI(SN)}_t$$

After the new degrees of interest have been determined for all the concept nodes, we normalize all the degrees of interest for these nodes, by dividing their values with the maximum calculated degree of interest. This normalization allows for newly acquired changes of interest to be able to significantly modify the interest values, even after a graph has been active for a long time. Subsequently, the degrees of interest for object (level 0) nodes are calculated by:

$$\mathrm{DI(SN)}_t = \sum_{i=1}^{N} \left(\mathrm{DI(PSN}_i)_t \cdot W_i \right)$$

The last equation certifies that the contents of the personalized rooms depend only on the interest values for concepts. Thus different exhibits may be selected with each visit or update, as long as they relate closely to the concepts the user is interested in.

Besides the selection of the recommendation set, a personal room should also have connections to other rooms of the virtual environment, allowing the user to further explore the exhibition. The set of rooms that are connected to the personalized room should also reflect the user's preferences. As mentioned earlier, every categorization node is represented as a set of rooms. So, to connect the personal room with the rest of the environment, L doors, where L is the door capacity of the personal room, are dynamically created. These doors are connected to a single room of each room set from the L top-rated level 1 categorization nodes of the user model.

3.4 Clustering and formation of user communities

In Internet based applications, added grouping capabilities can promote the formation of e-communities, thus increasing the sense of immersion in the virtual environment, enhancing communication opportunities, and satisfying the need for social interaction and awareness. To create a group of users with similar interests and preferences in the proposed framework, the user models must be compared. By assigning a unique index to every node of the SG, a user model can be taken as a numerical vector with each item being the degree of interest in the corresponding node. The dissimilarity between a pair of vectors is computed using Euclidean distance.

Using the above distance metric, the users of the system are clustered using the k-means algorithm [21] during an offline clustering process. The initial number of random centroids for the algorithm is determined by the number of current users. We chose to produce clusters of users that contain an average of ten users, in order to provide a fast and small list of users. An additional search option locates a given number of users with similar interests upon request.

The user can choose a personalized room from a subset of recommended rooms, the owners of which have been linked with her by the system, based on their interests. By using this approach, a personalized room can become a public space, in which different users can meet and communicate. The same holds for the rest of the exhibition. Considering the above, an environment created using the PeVEP platform can be viewed as a public exhibition space enriched with an arbitrary number of personal rooms, all being easily accessible.

3.5 The design process

The designer's primary task is to create and/or choose the rooms and the objects, which are models created in a 3D modeling tool. After all the content is collected, the designer has to create a semantic graph that represents a hierarchical categorization of the content, using the provided authoring tool. A screenshot of this tool is provided in Fig. 5.

At run time, the system will choose the room, the exhibits and the links with other rooms, depending on the client's request and available data. The actual position of the objects and doors is defined by the designer using a special kind of object called a placeholder. Special conventions are used inside the room template models to denote these placeholders. The graphical properties of these placeholders (translation, orientation and bounding boxes) are defined through the use of 3D primitives (such as the cones in Fig. 6) that are invisible during the rendering by the application client. Another task for the
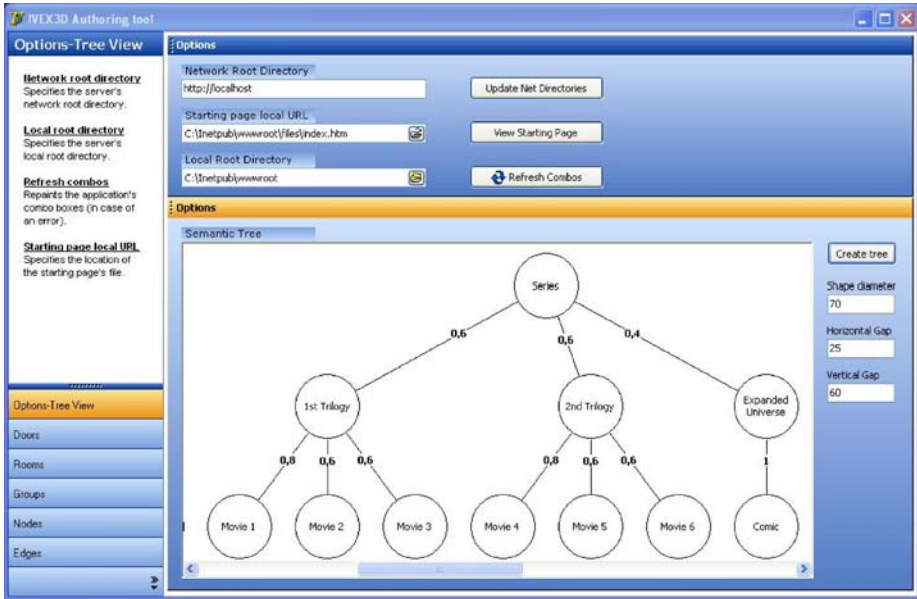
Fig. 5 A screenshot of the authoring tool

designer is to construct a set of avatars, which includes designing/selecting the respective 3D models and defining the avatar characteristics. The final task is to compose the stereotype rules that connect these characteristics to assumed user preferences i.e. to nodes in the SG.

Fig. 6 Construction of an exhibition room. The cones represent exhibit placeholders

As the generated degrees of interest rely on the organization of the SG and the weights of its edges, the designer of the exhibition is largely responsible for the successful arrangement of the objects in personalized rooms. During our experimentation with the system, we concluded that it is best to avoid graphs with high levels, as they tend to become unmanageable. On the contrary, including many low level trees in the graph helps to better semantically describe the concepts and objects and fend off weight changes that affect a large percentage of the graph.

4 Case study: a science fiction museum

In order to assess the PeVEP platform, a Virtual Museum concerning science-fiction movies, has been developed. The users can navigate in thematically different rooms (e.g. containing characters, vehicles etc.), enter a personalized room with various exhibits that the user might enjoy, according to her interests, or enter online personal rooms of users with similar interests. The exhibition's content (3D models) has been created using external modeling tools and imported to the environment, while the authoring tool presented in Section 3.5 facilitated the construction, manipulation and maintenance of the SG and the system's database management.

As the users navigate inside the virtual museum, they look at exhibits that fall inside their field of view. Furthermore, the users have the option of manipulating an exhibit, i.e. rotating it in order to gain a full perspective. Presentation of additional information concerning the exhibit is achieved through description pages in a provided panel. The users are also able to provide feedback about an exhibit by rating it and entering their personal message/comment. At the same time, users can view comments written by others, chat with them, observe their actions in the virtual space (via their avatars) and enter their personalized room.

The initial assumptions concerning user preferences are based on stereotypes related to the avatar gallery, and consequent interaction in the exhibition's space provides the appropriate degrees of interest for the profile update. Three types of interaction are monitored (ordered by the degree of influence upon the rating of an object, ascending):

- *viewing time*: the time spent looking at an exhibit. It is considered that the more time the user looks at an exhibit, the more interested she is in it.
- *manipulation of an exhibit*: interaction with an exhibit has a stronger impact than viewing time, as the user expresses her preferences more explicitly.
- *commenting and rating exhibits*: users can write comments that others can read and rate the exhibits accordingly. There are three available rating values: positive, negative and indifferent, which can reveal the user's opinion directly.

User interactions produce degrees of interest that are used by the server to assembly the personalized room and to distribute its contents dynamically upon user request. Users have the added ability to be instantly transported to any exhibition room (including the personalized room), in order to access rapidly a desired category/exhibit room.

The museum's content reached a total of 115 exhibits. Because of the intrinsic limitations of virtual environments that run over the Internet, the files have been processed by a polygon reducing tool. The content's categorization that produced the semantic graph was inspired by reviewing various sources regarding science-fiction movies. Another important issue was the creation of the museum's rooms and their structure (position of the placeholders etc). The content's aesthetics (futuristic style, space theme etc.) was taken into

consideration towards a satisfactory exposure for the exhibits: five room templates were designed, three of them for generic use and two of them for specific categories. Finally, to account for different interests and preferences among the users of the target (user) group, a small number of avatars were chosen, along with their characteristic attributes. An appropriate set of stereotypes was also formed, in order to associate these attributes to initial user interests, aiming to provide a smooth new user experience. Figure 7 presents a screenshot of the virtual museum application.

5 User evaluation

Evaluating personalized Virtual Environments is still an uncharted area. Systems that include a wide range of user interaction capabilities are difficult to evaluate in terms of efficacy. In most cases, such evaluations can only provide analysis on subjective observations produced by the usage of those systems; thus the accuracy of the results depends on the users' dedication and motivation during the evaluation and the produced results are affected by individual preferences. Despite these inherent flaws, a user evaluation of the case study has been considered necessary, in order to provide insight on the effectiveness of the PeVEP platform.

During this evaluation, we calculate the degree of satisfaction of the users from the recommendations and monitor the evolution of this value over time. Additionally, we produce an overall rating of the environment and summarize our observations about the usage of the system in general.

A group of 25 users participated in the evaluation process, all of which were undergraduate students of an Informatics university department. The users had to interact with the system and then fill out a questionnaire that was given to them. Out of the 25
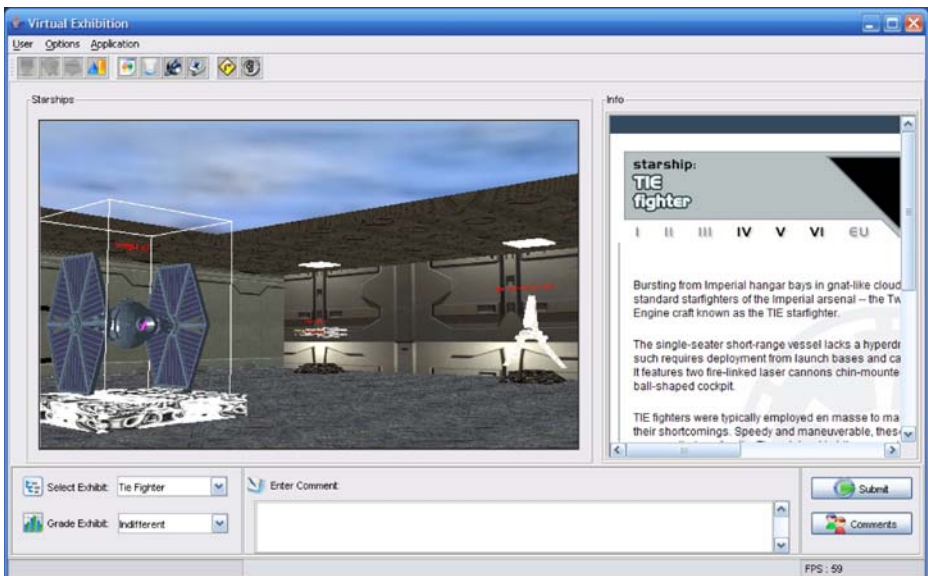

Fig. 7 A screenshot of the science fiction museum

participants, 17 stated that they had a fairly good experience in such 3D environments, while eight of them had almost never used a 3D application.

In the first step of the experiment process, the user had to choose an avatar from the provided library, register into the system and state her opinion about the degree of representation by the avatar. After logging into the system, the user had to complete the main process of the evaluation, which included navigating in the personalized room, interacting with the recommended objects and finally navigating through the rest of the museum for a short period of time, interacting with other objects of interest. Every object in the personalized room had to be marked in the questionnaire as known or unknown to the user and to be rated in a scale of 0 (total dissatisfaction) to 5 (total satisfaction). This evaluation process was repeated three times, while the system was adapting her user model, based on the user interaction with objects. At the end of the evaluation, the participant was asked to rate her experience with the system, and write down comments. A discussion with each user followed their evaluation and any written comments were analyzed.

The analysis of the questionnaire provided results that are displayed in Fig. 8; Fig. 8a depicts the distribution of overall object ratings, while Fig. 8b displays the evolution of the object ratings average in each evaluation cycle. The overall rating average was formulated at 3.1. Although these are preliminary results, it can be inferred that the contents of the personalized room have matched user preferences to a satisfying degree. Additionally, user satisfaction by the object recommendations has improved from the first round of the evaluation to the last by an average factor of 11.65%. This improvement, although small, was considered satisfactory by the authors, as the duration of the interaction periods was small and object recommendations are implicitly influenced by the initial collection of objects contained in the personalized room, during early user interactions with these objects. We have also measured the total user satisfaction from the system, which produced an average of 73%.

The discussions with the users and their comments produced the following observations. First, the quality of the rendered graphics played an important role in object rating and therefore affected total user experience, i.e. visually appealing, high-resolution models appeared to be more interesting than low-resolution models with poor quality textures. Additionally, the quality of navigation (smooth movement, collision detection etc.) and the interaction capabilities of the objects significantly influenced user opinion. Several users stated that if the aforementioned factors were improved, the total system rating would probably be higher. Overall, the experience gained during the conducted user evaluation,
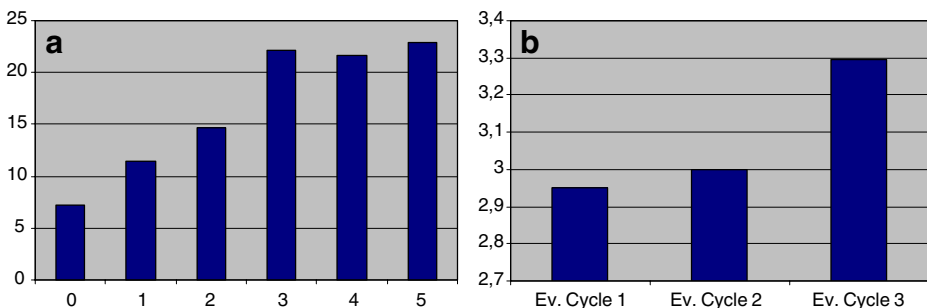

Fig. 8 User evaluation results. **a** Distribution of overall object ratings; **b** evolution of the average user recommendation ratings

indicates that intrinsic characteristics of dynamic 3D environments have an important impact on user satisfaction about the set of the recommended objects.

6 Conclusions and future work

This paper presented a user-oriented platform for designing and executing virtual exhibitions. Implicit generation and adjustment of user profiles allows applications that have been implemented on this platform to dynamically adapt content presentation to user interests and preferences, based on user stereotypes and prior interaction of users with exhibits. A semantic graph is used as a basis for both content categorization and user modeling. This semantic representation of content enhances presentation capabilities and simplifies the alteration and/or extension of existing environments. Additionally, it is possible to detect similarities among user models, leading to the formulation of user interest communities. A case study i.e. a science fiction virtual museum has been implemented, in order to gain insight about the effectiveness of this approach. A user evaluation of this case study produced favorable feedback towards the application of this framework in virtual museums and virtual content presentation systems in general.

In the future, the authors are planning to fine grain the PeVEP platform, through the development of more complex case studies, in order to better evaluate the user modeling process. Another objective for the research team would be to turn PeVEP into an open-ended platform for virtual reality content presentations thus promoting researcher, developer and designer cooperation in the field. Finally, an ongoing extension of this work is to enrich the platform by allowing users to customize their personal space, thus providing stronger feedback on the quality of their user model.

References

1. Benyon D, Höök K (1997) Navigation in information spaces: supporting the individual. In: Howard S, Hammond J, Lindegaard G (eds) Human-computer interaction: INTERACT'97, pp 39–46
2. Burke R (2002) Hybrid recommender systems: survey and experiments. User Model User-adapt Interact 12:331–370. doi:10.1023/A:1021240730564
3. Celentano A, Pittarello F (2004) Observing and adapting user behavior in navigational 3D interfaces. Proc. of AVI 2004, pp 275–282
4. Chittaro L, Ranon R (2002) Dynamic generation of personalized VRML content: a general approach and its application to 3D E-commerce, Proc. of the seventh international conference on 3D Web technology, pp 145–154
5. Chittaro L, Ranon R, Ieronutti L (2003) Guiding visitors of Web3D Worlds through automatically generated tours, Proc. of Web3D 2003: 8th International Conference on 3D Web Technology, pp 27–38
6. Ciabatti E, Cicnoni P, Montani C et al (1998) Towards a distributed 3D virtual museum. Proc. of Working Conference on Advanced Visual Interfaces, pp 264–266
7. Corbit M (2002) Building virtual worlds for informal science learning (SciCentr and SciFair) in the Active Worlds Educational Universe (AWEDU). Presence 11(1):55–67. doi:10.1162/105474602317343659
8. Darken RP, Sibert JL (1993) A toolset for navigation in virtual environments. In Proc. of UIST'93, pp 157–165
9. dos Santos C, Osorio F (2004) An intelligent and adaptive virtual environment and its application in distance learning. Proc. of the working conference on Advanced visual interfaces, pp 362–365
10. Economou D (2004) Virtual reality and museums: the development of a collaborative virtual environment for learning about ancient Egypt. Museology—International Scientific Electronic Journal 14(2)
11. Eirinaki M, Vazirgiannis M (2003) Web mining for web personalization. ACM Trans Internet Technol 3 (1):1–27. doi:10.1145/643477.643478

12. Fink J, Kobsa A (2000) A review and analysis of commercial user modeling servers for personalization on the world wide web. User Model User-adapt Interact 10:209–249. doi:10.1023/A:1026597308943
13. Itzkovich I, Hawkes LW (1994) Fuzzy extension of inheritance hierarchies. Fuzzy Sets Syst 62:143–153. doi:10.1016/0165-0114(94)90054-X
14. Kang H, Yang H (2004) The effect of anonymity on the usage of avatar: comparison of internet relay chat and instant messenger. Proc. of the Tenth Americas Conference on Information Systems, pp 2734–2743
15. Kim W, Kerschberg L, Scime A (2002) Learning for automatic personalization in a semantic taxonomy-based meta-search agent. Electron Commerce Res Appl 1:150–173. doi:10.1016/S1567-4223(02)00011-X
16. Kim YS, Kesavadas T, Paley S (2006) The virtual site museum: a multi-purpose, authoritative, and functional virtual heritage resource. Presence 15(3):245–261. doi:10.1162/pres.15.3.245
17. Ko D, Sumi Y, Choi Y et al (1999) Personalized Virtual Exhibition Tour (PVET): an experiment for Internet collaboration. Proc Syst Man Cybern IEEE SMC 99:25–29
18. Kobsa A, Koenemann J, Pohl W (2001) Personalized hypermedia presentation techniques for improving online customer relationships. Knowl Eng Rev 16(2):111–155. doi:10.1017/S0269888901000108
19. Lepouras G, Vassilakis C (2004) Virtual museums for all: employing game technology for edutainment. Virtual Real 8(2):96–106. doi:10.1007/s10055-004-0141-1
20. Lepouras G, Charitos D, Vassilakis C et al (2001) Building a VR-museum in a museum. Proc. of VRIC Virtual Reality International Conference
21. MacQueen JB (1967) Some methods of classification and analysis of multivariate observations, Proc. of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, pp 281–297
22. Minsky M (1975) A framework for representing knowledge. The Psychology of Computer Vision. McGraw-Hill, New York, pp 211–277
23. Mobasher B, Cooley R, Srivastava J (2000) Automatic personalization based on web usage mining. Commun ACM 43(8):142–151. doi:10.1145/345124.345169
24. Modjeska D, Waterworth J (2000) Effects of desktop 3D world design on user navigation and search performance, Proc. of IEEE Information Visualization, pp 215–220
25. Mukherjee R, Sajja E, Sen S (2003) A movie recommendation system—an application of voting theory in user modeling. User Model User-adapt Interact 13(1–2):5–33. doi:10.1023/A:1024022819690
26. Nowak KL, Rauh C (2005) The influence of the avatar on online perceptions of anthropomorphism, androgyny, credibility, homophily, and attraction. J Comput Mediat Commun 11(1):153–178. doi:10.1111/j.1083-6101.2006.tb00308.x
27. Patel M, White M, Walczak K et al (2003) Digitization to presentation—Building virtual museum exhibitions. Proc. of Vision Video and Graphics, pp 1–8
28. Perugini S, Goncalves MA, Fox EA (2004) Recommender systems research: a connection-centric survey. J Intell Inf Syst 23(2):107–143. doi:10.1023/B:JIIS.0000039532.05533.99
29. Petridis P, White M, Mourkousis N et al (2005). Exploring and interacting with virtual museums, Proc. Of Computer Applications and Quantitative Methods in Archaeology (CAA)
30. Raskutti B, Beitz A, Ward B (1997) A feature-based approach to recommending selections based on past preferences. User Model User-adapt Interact 7:179–218. doi:10.1023/A:1008291330418
31. Rich E (1979) User modeling via stereotypes. Cogn Sci 3:329–354
32. Rongen PHH, Schroder J, Dignum FPM et al (2005) Multi agent approach to interest profiling of users. Proc. of Multi-Agen Systems and Applications IV, pp 326–355
33. Schaerf M, Tessicini A (1999) JubilEasy: build a personalized 3D visit of Rome. In Proc. of VRML 99—Fourth Symposium on the VRML, pp 105–112
34. Schafer JB, Konstan JA, Riedl J (2001) E-commerce recommendation applications. Data Min Knowl Discov 5:115–153. doi:10.1023/A:1009804230409
35. Su CJ, Yen B, Zhang X (1998) An internet based virtual exhibition system: conceptual design and infrastructure. Comp Ind Eng 35(3–4):615–618
36. Tsichritzis D, Gibbs S (1991) Virtual museums and virtual realities. Proc. of International Conference on Hypermedia and Interactivity in Museums, pp 17–25
37. Urban R, Marty P, Twidale M (2007) A second life for your museum: 3D multi-user virtual environments and museums. In: Trant J, Bearman D (eds) Proc. of Museums and the Web 2007
38. Vinson NG (1999) Design guidelines for landmarks to support navigation in virtual environments. Proc CHI 99:278–285

Bill Bonis graduated at the Dept. of Informatics of the University of Piraeus, Greece, in 2004. He received a MSc in "Intelligent Human Computer Interaction" from the Dept. of Informatics of the University of Piraeus, in 2007. He is now a PhD candidate in Computer Science at the Dept. of Informatics of the University of Piraeus. His research interests include: Virtual Reality, Intelligent Virtual Agents, Adaptation and Personalization in 3D Environments, Personalization and User Modeling.

John Stamos graduated at the Dept. of Informatics of the University of Piraeus, Greece, in 2004. In 2007, he received an MSc in "Intelligent Human Computer Interaction" from the Dept. of Informatics of the University of Piraeus. His research interests include: Virtual Reality, Intelligent Virtual Environments, Adaptation and Personalization in 3D Environments, User Modeling, Human-Computer Interaction, Computer graphics.

Spyros Vosinakis is a Lecturer at the Department of Product and Systems Design Engineering of the University of the Aegean, Greece with subject "Virtual Reality". He holds a PhD in Informatics at the University of Piraeus, Greece with subject: "Intelligent Agents in Virtual Environments". In 2000, he received an MSc in "Computer Graphics and Virtual Environments" at the Department of Computer Science, University of Hull, UK. In 1998 he graduated at the Dept. of Informatics of the University of Piraeus, Greece. His research interests include: Virtual Reality, Intelligent Virtual Agents, Collaborative Virtual Environments, Adaptation and Personalization in 3D Environments.

Ioannis Andreou holds a PhD in Digital Systems and Technology Education from the University of Piraeus, Greece, specializing in Information Retrieval Techniques in Multimedia Systems and a Diploma in Computer Science from the University of Piraeus. His main research interests include (Content Based) Multimedia Retrieval, Computer Vision and Computer Graphics.

Themis Panayiotopoulos received his Diploma in Electrical Engineering and the Ph.D. degree from the National Technical University of Athens in 1984 and 1989 respectively. He has held a position of Assistant Professor with the Department of Mathematics at the University of the Aegean, Greece as well as a position of Research associate at the Institute of Informatics and Telecommunications, at the National Research Center Demokritos. He is a full professor at the Department of Informatics of the University of Piraeus, Greece. His main technical and research interests lie in the areas of Artificial Intelligence, Virual Reality and Web Technologies.