

Adaptive Virtual Exhibitions

Bill Bonis*, Spyros Vosinakis**, Ioannis Andreou* and Themis Panayiotopoulos*

**Deptt. of Informatics, University of Piraeus, Piraeus, Greece
E-mail: bonisb@unipi.gr, iandreou@unipi.gr, themisp@unipi.gr*

***Deptt. of Product and System Design Engineering
University of the Aegean Hermoupolis, Syros, Greece
E-mail: spyrosv@unipi.gr*

ABSTRACT

The paper presents a generic framework for designing and implementing multiuser virtual exhibitions that adapt visitors' preferences and goals and foster the emergence of communities with common interests. A user-oriented platform for designing and executing virtual exhibitions has been presented in which implicit generation and adjustment of user profile allows exhibition to dynamically adapt content presentation to users' interest and preferences.

Keywords: Virtual exhibitions, adaption techniques, semantic graph, taxonomy

1. INTRODUCTION

Virtual exhibitions are single- or multi-user realistic three dimensional (3D) representations of artifact collections, in which visitors navigate, observe the exhibits, learn related information presented in various media, and in some cases interact with them. Given that virtual exhibitions have no physical restrictions, they may host a limitless number of exhibits in a very large presentation space. On the other hand, studies have shown that visitors have various goals concerning their experience within an exhibition space, they follow different navigation strategies in browsing the exhibits, and they may be interested only in parts of the collection. These differentiations among online users raise a need for adaptation and personalisation of the 3D environment and its content.

Researchers in adaptive hypermedia distinguish between adaptation of content, adaptation of presentation and modality, and adaptation of structure. In the case of virtual exhibitions the content can be adapted by dynamically altering the artifacts presented in the exhibition spaces based on given criteria, e.g. user interests. The presentation and modality of the exhibition collection can be adapted by selecting between different media that may accompany the artifacts and by adjusting the presentation text to the visitors' preferences and goals. Finally, the adaptation of structure of 3D environments is not as trivial as in the case of web-based and hypermedia

presentations, where hyperlinks can be added or removed and navigation maps can be dynamically generated. In the case of virtual exhibitions the environment is assisting visitors' navigation towards places of interest rather than restructuring the space. This assistance can be based on animated agents, drawn paths, etc.

The paper presents a generic framework for designing and implementing multi-user virtual exhibitions that adapt to the visitors' preferences and goals and foster the emergence of communities with common interests. In the proposed architecture the exhibition content is dynamically arranged in thematic rooms based on semantic proximity, and visitors can optionally enter personalised spaces containing exhibits that may possibly interest them. The recommendation of artifacts takes place according to information implicitly collected about the users through their interactions with the museum. The proposed framework uses stereotypes to initialise user models, adapts user profiles dynamically and clusters users into similar interest groups, thus allowing virtual communities to emerge. The PeVEP platform has been implemented as a prototype of the proposed framework. It offers tools for the interactive design and maintenance of virtual exhibitions containing 2D or 3D artifacts with related information pages.

2. VIRTUAL EXHIBITIONS

The main objectives of exhibitions are to provide a public space for the presentation of artifacts and

to serve as a centre of knowledge specialised in a thematic area¹. The 2D interfaces (e.g., museum websites) limit the user experience to simple page viewing and sequential browsing and leave no room for any immersive experience. On the other hand, a 3D representation of an exhibition environment places the artifacts in a natural-looking setting and may offer a much more realistic and entertaining experience². The notion of virtual museums and exhibitions has been introduced by Tsichritzis & Gibbs³ as a means to overcome the limitations of the physical space and to provide a vivid experience to remote visitors. A synthetic collection of artifacts, which incorporates multimedia and virtual reality technologies, alleviates the problem of storage, preserving and protecting the real artifacts and allows virtual spaces to contain a limitless 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 accompanying legend, which is typically the case with traditional museum and exhibition 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 utilisation, assembly and disassembly of mechanical artifacts in a science museum, or the inclusion of educational games that are thematically related to the museum collection.

Various virtual exhibitions have been implemented as commercial or research projects during the last decade, taking advantage of the tremendous increase in the efficiency of graphics rendering and processing of modern PCs and of the availability of 3D scanning hardware. In some cases these applications serve as complementary information source to existing exhibitions and they are usually found in the web pages of museums or galleries and contain artifacts that belong to their collection. On the other hand, there are virtual exhibitions that host collections of artifacts that may be abstract, imaginary, restorations of damaged objects, hypothetical models of real artifacts that no longer exist, etc.⁴

There is a great diversity in terms of visualisation and user interactivity in the available approaches, as a multitude of different technologies has 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 exhibits⁵, using specialised hardware. The interaction modalities vary from simply viewing the artifacts to letting users have rich educative interactions with them⁶.

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. They may run remotely in the users' home computers or in public installations inside existing exhibition spaces. Virtual exhibitions may also be presented in multi-user environments, where users are represented with avatars and may communicate using text or voice chat. In this case, users can have a collaborative visiting experience by meeting people with similar interests, commenting on the exhibits, exchanging information about the area of interest, making new friends, etc. Virtual museums and galleries are also being built inside persistent virtual worlds, such as the popular second life platform⁷, thus making the exhibition accessible to a large, already established, user community, with no need to install any additional software or plug-in. An overview of virtual museum technologies is presented by Sylaiou⁸, *et al.*

3. ADAPTIVITY OF EXHIBITION SPACES

Distribution of artifacts in predetermined placeholders in physical exhibitions may not match all visitors' expectations, as it is practically impossible to present different subsets of the collection to individual visitors. Virtual exhibitions, 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 collection. Furthermore, it is significantly hard to extend or alter a physical collection, at least on a regular basis, whilst virtual exhibitions can offer the ability to perform instant changes on the collection presentation and to expand the museum space infinitely. As a result, virtual exhibitions may enhance traditional museum visits by offering the ability to adapt, expand, and personalise the artifact collections.

The presentation and categorisation of large collections has been effectively addressed in web⁹⁻¹² and multimedia applications¹³. User modeling techniques have been employed in order to personalise the content presentation according to the users' own interests. Virtual exhibitions 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 personalised space is expected to reduce the

navigational burden and still retain the metaphor of being immersed in a 3D environment.

3.1 Visitor Behaviour and Goals

Visitors of exhibitions may 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. Parker & Ballantyne¹⁴ studied the motivational factors of physical exhibition visitors using questionnaires in three different exhibition sites. The results show that exhibition visitors have various goals, such as learning and discovery, passive enjoyment, restoration, social interaction and self fulfillment. Furthermore, findings indicate a dependency between the visiting experience and the theme of the exhibition. For example, most aquarium visitors perceived their experience as entertaining, whilst most museum visitors appreciated the educational aspects of their visit.

Concerning their behaviour inside the physical space, Veron & Levasseur¹⁵ identified four different visiting styles using metaphors from animal motion behaviours:

- The ant visitor, who spends a long time to observe all exhibits and moves close to walls and exhibits avoiding empty space,
- The fish visitor, who walks mostly through empty space making just a few stops and sees most of the exhibits but for short time,
- The grasshopper visitor, who sees only exhibits he/she is interested in. He/she walks through empty space and stays for long time only in front of selected exhibits, and
- The butterfly visitor, who changes frequently the direction of visit, usually avoiding empty space. He/she sees almost all exhibits, but times vary between exhibits.

Umiker-Sebeok¹⁶ presented four major visitor types with respect to their interpretation of the exhibition space:

- The Pragmatic interprets the exhibition as a classroom or workshop and is interested in 'useful' information,
- The Critical visitor interprets the exhibition as a museum and is interested in the aesthetics of displays, the structure of the collection and the classification of exhibits,
- The Utopian visitor interprets the exhibition as an encounter session and his/her main goal is the social interaction, and

- The Diversionary visitor sees the exhibition as an amusement park and his/her goal is to have fun during the visit.

Finally, if physical or virtual exhibits are considered as information placeholders, one may interpret the visiting experience as browsing an information space. Dourish and Chalmers¹⁷ present three modes of navigation in information spaces: spatial, semantic and social. Spatial navigation is based on the arrangement of artifacts in their presentation space, semantic navigation is the browsing of artifacts based on semantic relevance, and social navigation is driven by the actions and preferences of other users in the information space.

Following from the study of the categorisations presented above, one may notice the great diversity between the needs and behavior of various exhibition visitors. Some visitors are focusing on the learning aspect of their visiting experience, may be interested in specific categories of exhibits, and may wish to discover all relevant information. Other visitors are more interested in the aesthetics of the experience and perceive their exhibition visit as a means of restoration, and finally some users are primarily interested in the social aspects of their visit and may possibly need to discuss their experience with others, meet other people with similar interests, etc. This diversity of interests could be supported in multi-user virtual exhibitions by offering adaptive content that satisfies the semantic browsing and social interaction needs of users based on their own preferences.

3.2 Adaptation Techniques

According to Kobsa¹⁸, *et al.* the personalisation 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 personalised content. A number of successful adaptive systems have been developed using web-based and hypermedia technologies that provide personalised 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.

Virtual environments are a special kind of medium, not widely used until a few years ago, which has major differences from the well known hypermedia metaphor used in web and multimedia applications. The elements are not arranged within a page (which can be easily explored using the

scroll bar and hyperlinks), but within a shared 3D space. Given that the visibility is limited (e.g., one cannot see behind walls), a lot of traveling is required from the users to explore the environment and search for information artifacts. Furthermore, virtual worlds cannot be 'hyperlinked' to each other due to a lack of common platform; each world has its own user base and content and sharing is not possible. The only equivalent of hyperlinking in 3D environments is the ability to "teleport" the user to distant places. However, similarly to the hyperlink metaphor, extensive use of this kind of navigation may distort the users' sense of presence and cause disorientation¹⁹.

The collection of user interests can be based either on user supplied information, e.g., by explicitly asking the user before entering the environment, or on observation of user behavior that may form assumptions about his/her likes and dislikes. The second case is not as straightforward as in hypermedia applications, because users are not explicitly clicking on hyperlinks and media elements; they are simply navigating inside the 3D space and observing areas of the environment. Therefore, the system has to adopt strategies, such as the user's field of view calculation of total viewing times to infer the elements of the environment that are possibly capturing the users' interest. Chittaro & Ieronutti²⁰ have proposed a tool for recording user behavior in virtual environments. It can collect information about the places the user has visited and the time spent in each position, as well as about the parts of the environment that he has seen during his visit.

Adaptation in hypermedia systems is generally applied in the following categories:

- (a) Adaptation of content: Personalisation of content based on users, usage and environment,
- (b) Adaptation of presentation: Changing the presentation and media format and the interaction elements, and
- (c) Adaptation of structure: Personalising the presentation of links.

In the case of 3D environments, the adaptation of content and structure is not as simple as the process of generating a dynamic web page. The form and topology of the interior and exterior spaces cannot be easily adjusted to custom content size. One approach is to use placeholders in predefined places of the virtual exhibition and to change the exhibits based on some criteria, however, this approach can only be applied for a static number of exhibits. This method has been used in the AWE3D environment where the contents of an on-line store are dynamically arranged based on visitor preferences²¹. Another approach for structure adaptation is the use of a

'rooms and corridors' metaphor and the dynamic arrangement of the rooms based on some criteria. A prototype virtual environment for distance learning²² has been implemented using this metaphor, where each thematic area is represented as a different room and visitors walk through a corridor with the rooms arranged based on their interests.

Probably the most widely used adaptation technique in 3D environments is the adaptive user navigation. Systems adopting this technique are either generating automated tours based on user interests or contain embodied agents as virtual guides with which users can interact using dialogs or natural language. A first approach towards personalised 3D exhibition environments is JubilEasy²³, a system that generates virtual visits in the city of Rome based on information received explicitly by the user. Chittaro²⁴, *et al.* present a Virtual Museum of Computer Science in which an embodied agent offers guided tours to visitors using natural language to present the exhibits. Finally, an agent that presents a virtual exhibition in second life is being presented by Oberlander²⁵. The agent can adjust the language and style of presentation based on user preferences.

3.3 Related Approaches

Chittaro & Ranon²¹ propose an approach for the development of adaptive 3D websites. 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 personalised environment is generated using a world generator process that selects elements from a VRML content database and generates the personalised environment. AWE3D has been used for the development of a personalised 3D e-commerce site and for automatically generated tours in virtual museums²⁶. Another approach towards personalisation in virtual environments is the intelligent virtual environment presented²², 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 categorisation 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 & Pittarello²⁷ propose an approach to facilitate adaptive interaction with the virtual

environment, which is based on 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 behaviour and predicting future needs for interaction purposes. When the system recognises the initial state of an interaction pattern, it executes the final state without letting the user engage in the intermediate ones.

Lepouras & Vasilakis²⁸ present an overview of adaptive virtual museums on the web and propose the architecture of an adaptive virtual environment creation system. In their proposed approach, designers provide templates of rooms, objects and interaction techniques and an environment generator is generating personalised museums based on user models. A usage data recorder is monitoring user motion in the environment and is updating the user profile accordingly.

4. MODELING ADAPTIVE VIRTUAL EXHIBITIONS

This work presents a framework for virtual exhibitions, which adapt to user preferences and experience. Exhibits are selected and presented in the context of a hypothetical user's interests. An abstract model represents the interests of each user. This model must be initialised upon entering the exhibition environment and must be subsequently updated to reflect to its owner. In order to satisfy these requirements the user model encapsulates user characteristics as they are deducted by her initial preferences and her recorded actions inside the presentation space itself. Overall, the proposed framework contemplates on generating and updating the user model, presenting content and grouping of users based on similar characteristics. The retrieval approach can be classified as content-based²⁹.

Ontologies are used in a wide variety of applications, as they can effectively describe classes, objects and relations among them in a high level in a representation usable by humans and algorithms. Gruber³⁰ defines Ontology as an explicit specification of a conceptualisation. Ontology-based profiles have been used in web-search³¹, in which case a hierarchical taxonomy of human interests is adapted over time to improve retrieval precision. Knowledge management systems also employ ontologies in user modeling³² while others³³ attempt to incorporate abstract user concepts and properties into ontology-based user profiles. In our framework, the exhibitions as well as the users are modeled as ontologies, with the latter's structure relying on the former's. On the most abstract level, the design of virtual exhibition

is addressed by organising the exhibits domain into categories (classes) and specifying relations of generalisation/specialisation and affinity among categories. The basis of the exhibition design is therefore a hierarchical taxonomy of the content categories. Utilising the degree of a user's interest in categories it is possible to deduce her degree of interest in the actual contents. Furthermore, these categories (classes) and relations among them participate in both exhibition and user models.

User models are initialised through a process that assigns degrees of interest to the categories of the exhibition's content taxonomy. The initialisation method should be both user and context dependent³⁴ and preferable implicit thus facilitating users without posing any extra burden. The degree of interest in each of those categories are the dynamic components of the user model. However, the importance of adaptivity in virtual exhibitions of large scale can be as important as in hypermedia systems³⁵. While the user is browsing the environment, her navigation and interaction with content are monitored and the recorded behavior is utilised to make assumptions about her interests and preferences, which are then incorporated into the user's model. Combining the contents of the user ontology and the exhibition ontology it is possible to produce a personalised environment, which reflects the user's assumed preferences and to recommend new content that might be of her interest. Overall, the exhibition framework can be implemented using only transparent processes for personalisation³⁶. Additionally, user groups can be formed by comparing user models, thus introducing social awareness and interaction in the virtual environment. Figure 1 presents an overview of the Adaptive Virtual Exhibitions Framework. To elaborate on this framework, a virtual exhibition platform has been implemented; an analytical description of this platform follows.

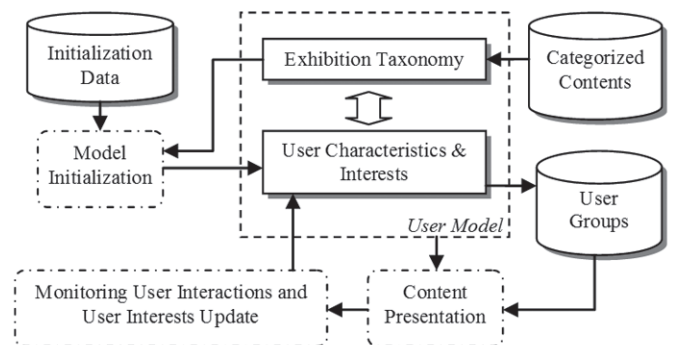


Figure 1. An adaptive virtual exhibitions framework.

5. PeVEP PLATFORM FOR CONTENT PERSONALISATION

PeVEP (Personalised Virtual Exhibition Platform) is a platform for designing and implementing virtual

exhibitions with content personalisation. When a new user enters the PeVEP environment for the first time, a user model is assigned to her based on a selection of an avatar. Her user model is consequently updated according to her monitored behavior as she is browsing the environment. At any time, the user can ask to be transported to a personalised 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 personalised 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 personalisation 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 personalised 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.

5.1 The Semantic Graph

A vast number of applications that utilise 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 systems³⁷⁻⁴³. A distinctive characteristic of these systems, compared to information retrieval and filtering systems and search engines, is the output of individualised 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 categorised 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 generalised into broader categories or specialised 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. Let level 1 be the level of the most specialised categorisation 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 categorisation trees using connection(s) with one or more level 1 category nodes via an instance relation. Thus, the resulting ontology formulates a directional graph, the Semantic Graph (SG), and its core is a hierarchical semantic taxonomy. In the semantic graph the nodes (distributed into levels) stand for objects and concepts, edges represent the relations between them^{42, 44, 45} and the levels represent the degree of generalisation. In this taxonomy, the actual objects are considered as level 0 nodes and they do not necessarily belong to only one tree.

Figure 2 presents a sample part of a semantic graph that is used for categorising movies-related content. 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 directed by Stephen Spielberg. Thus, its 3D representation can be connected to the nodes 'Science Fiction' and 'Stephen Spielberg'.

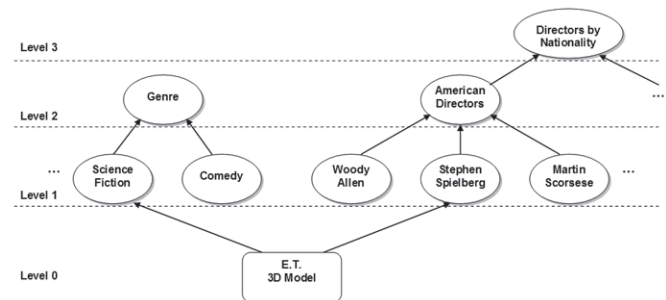


Figure 2. A part of a semantic graph for categorising movies. Level 0 is the objects level.

Let A be the node at the lower end of a categorisation 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 Theory⁴⁴. 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. 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 personalisation-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,
- Grouping (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.

5.2 Dynamic Content Generation and Rendering

The hierarchical structure of the SG is used for the creation of the virtual exhibition's spatial structure. Each categorisation node can be represented by a set of interconnected rooms. Rooms that represent conceptually relevant nodes (as portrayed in the graph) are 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 characterise it. For example, 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. A default template room is required 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 application has to construct a thematic room, it searches for a predefined template that matches the respective category. In the case that there are no such template exists, the default room is used. Each room's name is derived by the corresponding node it represents. 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 exhibit 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 categorisation 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 categorisation node, the parent node, or sibling categorisation 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 categorisation i.e. these rooms represent specialised categorisation nodes that are connected via a generalisation 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 led to 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' categorisation node, which is a level 2 categorisation 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 categorisation 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.

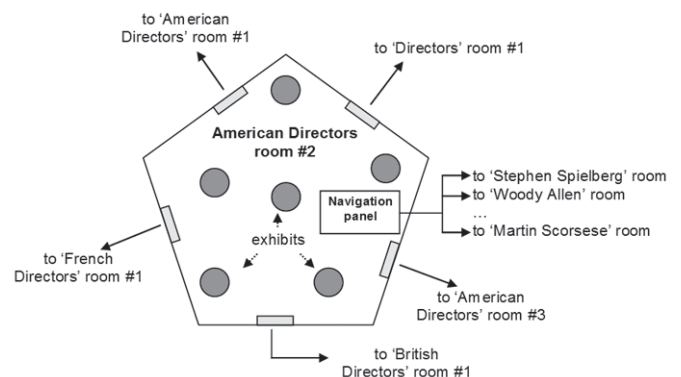


Figure 3. Schema of a thematic room, generated from a semantic graph that categorises movies.

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 categorisation node, a navigation panel exists, to provide options for moving to more specialised rooms, i.e., rooms of specific American directors.

This arrangement allows every room, regardless of generalisation (level) to contain exhibits. However, the rooms of the lowest level nodes provide the greatest degree of specialisation. While navigating, the user moves between rooms of the same generalisation 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 generalisation (as they are represented by levels in the graph) and browse rooms using different subtrees of the graph.

5.3 Personalisation

All available objects in the PeVEP platform are categorised 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 (e.g., 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.

5.3.1 User Stereotypes

As stated by Rongen⁴⁶, *et al.*, 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 utilised, in which a set of stereotypes^{41, 47} is used to initialise 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 characterise 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 categorisation nodes is calculated as follows. Initially all the categorisation 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 categorisation 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 problem⁴⁰, 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 researches⁴⁸⁻⁴⁹ that state the possible relation between the choice of an avatar by the user and the users' intrinsic characteristics and personality traits. Expert knowledge and/or user segmentation researches can be used to create the set of stereotypes.

5.3.2 Data Acquisition

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 virtual exhibition's application. 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, providing a rating system that lets users express their preferences, etc. 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).

Objects in the virtual environment can provide a multitude of interaction modalities which may vary between the available set of objects. Because some objects may be considered to be of significant importance compared to others, the designer can assign an initial degree of interest for the objects of his choice. Exhibits with higher initial degree will

have a comparative advantage, especially for new users' profiles, given the fact that the system will not have sufficient information about their interests. The values for the initial degrees can be arbitrary, although there should not be significant differences between objects in order to avoid producing over-biased recommendations. The default value is 1.

The total degree of interest for each object is calculated

$$\text{by: } TDIn = InD \cdot \left(\sum_{n=1}^N (\text{Inf } F \cdot SR)_{\text{modality}_n} \right) + \sum_{k=1}^K OutR$$

where, $TDIn$ is the total degree of interest for a specific object, InD is the initial degree of interest, $InfF$ is the influence factor of the objects, n is interaction modality, SR (special rating) is the degree of interest for the n interaction modality and $OutR$ (outside rating) is any rating provided by outside factors (if any) like user-rating. The influence factor determines the extend in which a particular interaction modality effects the total degree of interest. The default value is 1. The special rating varies between interaction modalities. If for example, an interaction takes a certain amount of time to be completed, the special rating can be the time interval, or if it is a goal oriented interaction, the special rating can be 0 if the goal was not achieved or a non zero value otherwise.

5.3.3 User Profile Updating

The first step in the user profile updating process is to normalise the objects' degrees of interest from the previous monitoring process by dividing with the maximum degree.

Let SN be a semantic node and $DI(SN)_t$ the degree of interest for node SN at time instance t . Let $CSN = \{CSN_1, CSN_2, \dots, CSN_N\}$ 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:

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

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

After the completion of the calculations, the new degrees of interest are normalised by dividing with the maximum degree. Because of the additive

nature of the previous equation, without normalisation each profile update would gradually increase the profile's degrees of interest and the effect of the objects' interaction data to the profile would in time decrease. This normalisation 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.

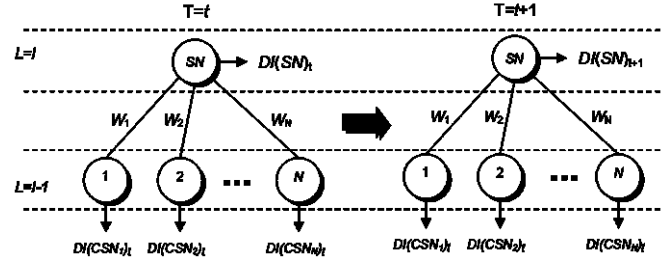


Figure 4. Updating the user profile

The algorithm's next step uses two additional set of data, the maximum recorded degree of interest for all nodes in each user profile and the previous degrees of interest i.e., the previous profile. First, the new user profile is compared node by node with the corresponding maximum degrees of interest and nodes that have decreased degrees by 50 % or more are flagged to be downgraded. The rest of the nodes are compared with the previous profile. If a node exhibits a decrease by 70 % or more, a flag called pre-reduction indicator is checked and if true the node is also flagged to be downgraded. If the indicator was false then it is given the value true. The degrees of interest for all nodes flagged with the downgrade indicator are reduced by half. Both indicators are initialised to false before the process starts. A less than 70 % decreased degree of interest is not considered an adequate indication of indifference because a user can have short term lack of interest periods and at the same time she can maintain a long term interest. The pre-reduction indicator introduces a delay for one calculation circle for any needed downgrades.

The previous process enables the system to react to the users' interest dynamics and to base its recommendation decisions not only in the current information collected in previous steps, but also in the immediate short term interests' history.

5.3.4 Adaptation of Environment

For the creation of a personalised room, a set of objects are chosen to be recommended to the user by populating the room with them. This task is realised using a downward propagation mechanism based on the user model. This mechanism is computationally equivalent to the upward propagation mechanism described before, having only the opposite direction.

During the downward propagation procedure, an estimated degree of interest is assigned to the object nodes. The choice of objects that will be part of the user's personalised 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 $DI(SN)_t$ is the degree of interest for node SN at time instance t . Let $PSN = \{PSN_1, PSN_2, \dots, PSN_N\}$ be the set of parent nodes of node SN , with $DI(PSN)_i$, $i \in [1, N]$ 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 $t+1$ for node SN is calculated by:

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

Subsequently, the degrees of interest for object (level 0) nodes are calculated by:

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

Last equation certifies that the contents of personalised 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 of user interest.

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 personalised room should also reflect the user's preferences. As mentioned earlier, every categorisation 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 categorisation nodes of the user model.

5.4 Grouping 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. The main goal of the grouping subsystem in the proposed framework is to create interest groups, each member of which would have similar objects proposed to her by the personalisation algorithm. The user model, i.e. an arbitrary set of nodes from the SG with a specific degree of interest for each node, and the set of proposed objects for each user provide the needed data for the grouping process. Since all registered users have a working profile and the personalisation algorithm can produce a recommendation when needed, all data are readily available. All degrees of interest, object and user model alike will be normalised before being used by the grouping subsystem, thus maintaining a level ground for the comparisons between user data.

To calculate the interest proximity or dissimilarity between two users we need first to define a distance metric. A metric like the Euclidean distance used on indexed user models, therefore viewed as vectors, provides a relatively simplistic approach because,

- (a) Measurements represent geometric distances in a two-dimensional space deprived of knowledge insight, and
- (b) Comparison data have high informational content that should be taken into consideration, content that is stripped from the data if viewed as simple vectors.

As mentioned before, there are two different comparison data sets, the recommended objects and the user model. First let us define the distance between two sets of recommended objects. The metric should be affected by the ordering of objects based on their degrees of interest since dissimilarity between highly rated objects is more important compared to dissimilarity between objects lower in the interest ordering. Let k be the number of the first k objects in each object ordering used for the distance measurement. Different exhibition characteristics e.g. the room's objects capacity, can be taken into account for the choice of k . Let n be a specific position in both object orderings. If the same object resides in n then the distance will not be increased. If there are different objects in n then the distance will be increased by $k \cdot (n-1)/k$, a fraction that decreases by $1/k$ th for each position further down the ordering. Thus, the object distance metric is:

$$dist_{obj} = \sum_{i=1}^n x_i, \text{ where } 1 \leq n \leq k \text{ and}$$

$$x_i = \left\{ \begin{array}{l} 0, \text{ if identical objects in } n \\ \frac{k - (n-1)}{k}, \text{ if different objects in } n \end{array} \right\}$$

This particular distance metric has a pair wise comparative value and does not reveal the degree in which two user models or recommendation sets are similar.

By calculating the percentage of objects in the set of k objects that are contained in both recommendation sets and have a degree of interest greater than a certain limit, a quality metric is stemmed and for small values two recommendation sets cannot be considered similar enough. The minimum percentage and the minimum degree of interest can be adjusted according specific exhibition requirements, e.g., room object capacity. A higher percentage and a higher degree of interest formulate a more demanding quality metric.

If the main goal of the grouping process was only the formation of interest groups or more accurately groups with similar recommendation sets, the previous distance metric would be sufficient. Because there is a need for formation of groups with users having similar user models, an added user model distance metric should be introduced. The reason for that resides in the fact that even though two similar user models will produce similar recommendation sets, two similar recommendation sets are not necessarily originated by similar user models. This means that the object distance metric does not portray an accurate grouping based on user interests. Let l be the number of nodes in the SG that will be used in the pair wise user model comparison. Let m be a specific position in both node orderings based on the respective degrees of interest.

Let $dif_{AB} = |di_A - di_B|$, where di_x is the degree of interest for the node X and nodes A and B , the same nodes in the SG for each user model. The user model distance metric is defined :

$$[dist]_{um} = \sum_{i=1}^l (t=1)^{\uparrow} m \equiv y_{\downarrow} i, \text{ where } 1 \leq m \leq l \text{ and } y_{\downarrow} i$$

$$= \left\{ \begin{array}{l} 0, \text{ if nodes } A \text{ and } B \text{ both in } m \text{ and } [dif]_{\downarrow} A^{\uparrow} \end{array} \right.$$

The total distance is defined by:

$$dist_{tot} = dist_{obj} + dist_{um}$$

There is no need for a dedicated process for the removal of outliers in the data due to the usage of the quality metric mentioned earlier. A user model that does not satisfy that requirement will not be included in a group.

The grouping process itself needs a specific grouping algorithm that will use the aforementioned distance metric and create the actual groups. There are some limitations posed to the candidate algorithm.

- (1) Number of groups is not known beforehand, and
- (2) Number of users in each group is also not known beforehand.

Even though there might be a substantial number of registered users in the system, the volume of the comparison data is not very big thus posing no serious calculating restrictions. Hierarchical agglomerative clustering algorithms⁵⁰ can find arbitrary number of clusters with arbitrary number of members each in a multidimensional dataset, thus being ideal for this application. The process starts with each user data set being a single cluster and terminates when every new merging violates the aforementioned quality percentage criterion.

The system supports two different techniques for choosing the distance between clusters during the merging process, complete link and average link technique⁵⁰. The single link technique is not used because it suffers from a chain effect (clusters straggly and elongated)^{50, 51} that eventually will leave user data not being included in any cluster due to the quality criterion, data that otherwise under certain conditions could form an individual cluster.

The user can choose a personalised 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 personalised 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.

5.5 Designing Process-Maintenance of Exhibition

The main purpose of PeVEP platform is to provide users with an enhanced experience through the virtual exhibition and with its central entity the content

itself. Thus, the designer's primary task is to collect the 3D content, i.e., the exhibits, the set of which will define the exhibition's theme. After all content is collected, the designer has to create a semantic graph that represents a hierarchical categorisation of the content, using the provided authoring tool. A screenshot of this tool is provided in Fig. 5.

through the environment 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

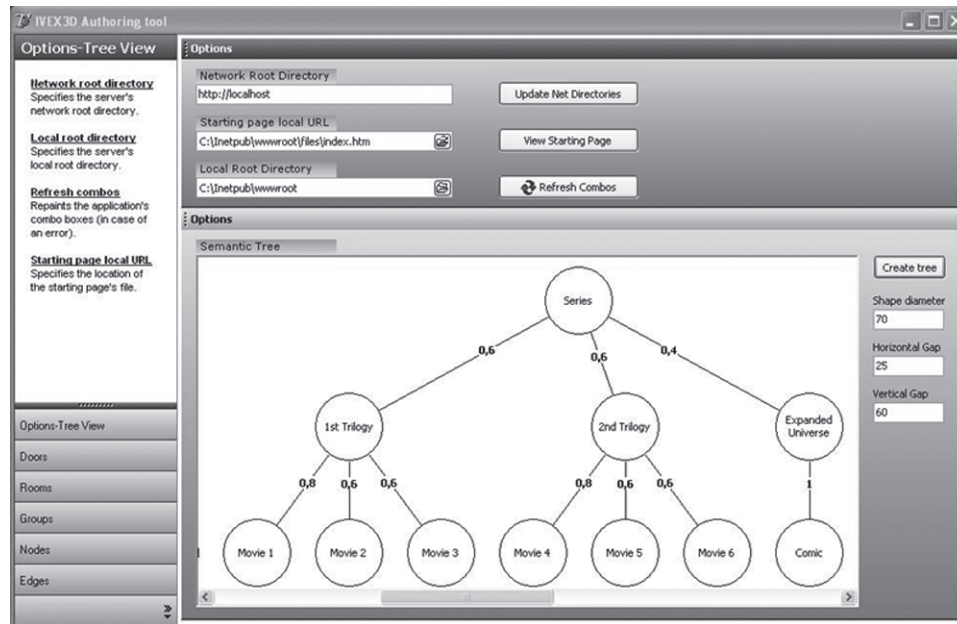


Figure 5. IVEX 3D authoring tool.

This process can be viewed as a series of distinct steps. The designer first has to formulate a natural categorisation of the content based on domain specific criteria and the available content. Next, the categorisation will be structured into a hierarchy. After all categories have been identified and linked where necessary, the content will be correlated with the corresponding nodes/categories of the lowest level in the newly formulated SG. Since the SG not only facilitates the content's categorisation, but also the user modeling process, it should also reflect the users' possible interests. This would, at some cases, lead to the deletion of higher level nodes which group categories into more generalised ones that do not provide any further information. An added reason to avoid over generalised nodes is the fact that all nodes will eventually produce a room (or set of rooms) in the exhibition and these particular rooms will probably not illustrate a clear thematic cohesion, as viewed from the users' standpoint. As soon as the structuring of the SG is finalised, the degrees of membership between linked nodes should be chosen. This last step concludes the process of designing the SG.

The designer's next task is to create a set of rooms for the exhibition using a 3D modeling software tool of his choice. At run-time, the system will choose the room, the exhibits and the links with other rooms, depending on the user's navigation

defined through the use of 3D primitives (such as the cones in Fig. 6) that are invisible during the rendering by the application (Fig. 7). Another task for the designer is to construct a set of avatars, which includes designing/selecting the respective 3D models, defining the avatar characteristics and to compose the stereotype rules that connect these characteristics to assumed user preferences, i.e., to nodes in the SG.




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

The final task is to adjust the various system specific variables like the interaction modalities degrees and clustering variables. All data structures, files and system variables can be altered, refined




Figure 7. An exhibition room at (a) design time and at (b) run time. The cones represent exhibit placeholders.

or enriched while the exhibition is online. Special consideration should be taken if any changes in the SG are deemed necessary. Because the user's profile is stored in an instance of the SG, the deletion of nodes will probably, under certain circumstances, lead to a significant diversification of the recommended content. Furthermore, new node additions in the SG, although they will not alter the user's profile, they will exhibit a handicap compared to existing nodes because of the latter's head start. As more usage data are gathered, any differences will be eventually assimilated.

6. CASE STUDIES

To assess the PeVEP platform, two virtual exhibitions have been developed, a Virtual Museum about science fiction movies and one about the space exploration program. In both museums the users can navigate in thematically different rooms, enter a personalised room with various exhibits that the user might enjoy, according to her interests, or enter online personal rooms of users with similar interests. The exhibitions' 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 each SG and the systems' database management.

As the users navigate inside virtual museums, 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 personalised room.

The initial assumptions concerning user preferences are based on stereotypes related to the avatar gallery, and consequent interaction in the exhibitions' 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):

- (a) 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.
- (b) Manipulation of an exhibit: Interaction with an exhibit has a stronger impact than viewing time, as the user expresses her preferences more explicitly.
- (c) 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 personalised 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 personalised room), in order to access rapidly a desired category/exhibit room.

The Science Fiction museum's content reached a total of 115 exhibits, whereas the Space Exploration museum's content reached a total of 104 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 categorisation that produced the semantic graph

in either case was inspired by reviewing various sources regarding science-fiction movies for the Science Fiction museum and on-line resources for the Space Exploration museum, like NASA's site and Wikipedia. Another important issue was the creation of the museums' rooms and their structure (position of the placeholders, etc.). For the Science Fiction museum five room templates were designed, three of them for generic use and two of them for specific categories and for the Space Exploration museum four room templates were designed, three for generic use and one 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 for both museums, 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 8 presents a screenshot of the Science Fiction museum application and Fig. 9 presents a screenshot of the Space Exploration museum.

The applications were based on client/server architecture and have 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 decentralised user modeling architectures⁵².

7. CONCLUSIONS

This paper presented an overview of virtual exhibition technologies and the related work in content presentation adaptivity and personalisation. A user-oriented platform for designing and executing virtual exhibitions has been presented, in which implicit generation and adjustment of user profiles allows exhibitions to dynamically adapt content presentation to user interests and preferences. A semantic taxonomy of content simplifies the alteration and/or extension of existing environments and enhances presentation capabilities. Additionally, it is possible to detect similarities among user models, leading to the formulation of user interest communities. Two case studies, a science fiction and a space exploration museum, have been implemented, in order to gain insight about the effectiveness of this approach.

REFERENCES

1. Economou, D. Virtual reality and museums: The development of a collaborative virtual environment for learning about ancient Egypt. *Museology - Inter. Sci. Electr. J.*, 2004, **14**(2).
2. Lepouras, G. & Vassilakis, C. Virtual museums for all: Employing game technology for edutainment. *Virtual Reality*, 2004, **8**(2), 96-106.
3. Tschritzis, D. & Gibbs, S. Virtual museums and virtual realities. *In Proceedings of International Conference on Hypermedia and Interactivity in Museums*, 1991, pp. 17-25.
4. Ciabatti, E., et al. Towards a distributed 3D virtual museum. *In Proceedings of Working Conference on Advanced Visual Interfaces*, 1998, pp. 264-66.
5. Kim, Y.S.; Kesavadas, T. & Paley, S. The virtual site museum: A multi-purpose, authoritative, and functional virtual heritage resource. *Presence*, 2006, **15**(3), 245-61.
6. Petridis, P., et al. Exploring and interacting with virtual museums. *In Proceedings of Computer Applications and Quantitative Methods in Archaeology (CAA)*, 2005.
7. Marty, P.F.; Twidale, M. & Urban, R.A. Second life for your museum: 3D multi-user virtual environments and museums. *In Proceedings of Museums and the Web*, 2007.
8. Sylaiou, S., et al. Virtual museums: A survey and some issues for consideration. *J. Cultural Heritage*, 2009, **10**(4), 520-28.
9. Mukherjee, R.; Sajja, E. & Sen, S. A movie recommendation system: An application of voting theory in user modeling. *User Model. User-Adapted Interaction*, 2003, **13**(1-2), 5-33.



Figure 8. Science Fiction museum.

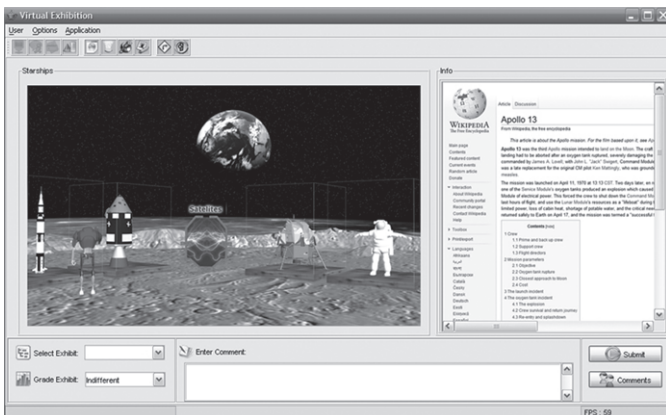


Figure 9. Space Exploration museum.

10. Mobasher, B.; Cooley, R. & Srivastava, J. Automatic personalisation based on web usage mining. *Communications of ACM*, 2000, **43**(8), 142-51.
11. Eirinaki, M. & Vazirgiannis, M. Web mining for web personalisation. *ACM Trans. on Internet Technol.*, 2003, **3**(1), 1-27.
12. Schafer, J.B.; Konstan, J.A. & Riedl, J. E-commerce recommendation applications. *Data Mining and Knowl. Discov.*, 2001, **5**, 115-53.
13. Raskutti, B.; Beitz, A. & Ward, B. A feature-based approach to recommending selections based on past preferences. *User Model. User-Adapted Interaction*, 1997, **7**, 179-218.
14. Packer, J. & Ballantyne, R. Motivational factors and the visitor experience: A comparison of three sites. *Curator: The Museum J.*, 2002, **45**(3), 183-98.
15. Veron, E. & Levasseur, M. Ethnographie de l'Exposition. Bibliothèque publique d'Information, Centre Georges Pompidou, Paris, 1983.
16. Umiker-Sebeok, J. Behavior in a museum: A semio-cognitive approach to museum consumption experiences. *Signifying Behavior. J. Res. Sem., Comm. Theo, and Cognitive Sci.*, 1994, **1**(1), 52-100.
17. Dourish, P. & Chalmers M. Running out of space: Models of information navigation. *In Proceeding of HCI 94*, 1994.
18. Kobsa, A.; Koenemann, J. & Pohl W. Personalised hypermedia presentation techniques for improving online customer relationships. *The Know. Eng. Rev.*, 2001, **16**(2), 111-55.
19. Bowman, D.; Koller, D. & Hodges, L.F. Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. *In Proceeding of Virtual Reality Annual International Symposium (VRAIS)*, 1997.
20. Chittaro, L. & Ieronutti, L. A visual tool for tracing users' behavior in virtual environments. *In Proceedings of the Working Conference on Advanced Visual interfaces AVI 2004*, 2004. pp. 40-47.
21. Chittaro, L. & Ranon, R. Dynamic generation of personalised VRML content: A general approach and its application to 3D e-Commerce. *In Proceedings of 7th International Conference on 3D Web Technology*, 2002, pp. 145-54.
22. dos Santos, C. & Osorio, F. An intelligent and adaptive virtual environment and its application in distance learning. *In Proceedings of Working Conference on Advanced Visual Interfaces*, 2004, pp. 362-65.
23. Schaerf, M. & Tessicini, A. JubilEasy: Build a personalised 3D visit of Rome. *In Proceedings of VRML 1999: 4th Symposium on the VRML*, 1999, pp. 105-12.
24. Chittaro, L.; Ieronutti, L. & Ranon, R. Navigating 3D virtual environments by following embodied agents: A proposal and its informal evaluation on a virtual museum application. *Psychology Journal*, 2004, **2**(1), 24-42.
25. Oberlander, J., et al. Building an adaptive museum gallery in second life. *In Proceedings of Museums and the Web 2008*.
26. Chittaro, L.; Ranon, R. & Ieronutti, L. Guiding visitors of web3D worlds through automatically generated tours. *In Proceedings of Web3D 2003: 8th International Conference on 3D Web Technology*, 2003. pp. 27-38.
27. Celentano, A. & Pittarello, F. Observing and adapting user behavior in navigational 3D interfaces. *In Proceedings of AVI 2004*, 2004. pp. 275-82.
28. Lepouras, G. & Vassilakis C. Adaptive virtual reality museums on the web. *In Adaptable and Adaptive Hypermedia Systems*, edited by S.Y. Chen & G.D. Magoulas, 2005.
29. Burke, R. Hybrid recommender systems: Survey and experiments. *User Model. User-Adapted Interaction*, 2002, **12**, 331-70.
30. Gruber, T.R. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 1993, **5**(2), 199-220.
31. Trajkova, J. & Gauch, S. Improving ontology-based user profiles. *In Proceedings of RIAO*, 2004, pp. 380-89.
32. Razmerita, L.; Angehrn, A. & Maedche, A. Ontology based user modeling for Knowledge Management Systems. *In Proceedings of User Modeling Conference*, 2003, pp. 213-17.
33. Golemati M., et al. Creating an ontology for the user profile: Method and applications. *In Proceedings of 1st IEEE International Conference on Research Challenges in Information Science (RCIS)*, 2007.
34. Nguyen, Q.N. & Ricci, F. User preferences initialisation and integration in critique-based mobile recommender systems. *In Proceedings of Workshop on Artificial Intelligence in Mobile Systems*, 2004.
35. Brusilovsky, P. Adaptive Hypermedia. *User Model. User-adapted Interaction*, 2001, **11**(1-2), 87-110.
36. Kühme, T. A user-centred approach to adaptive interfaces. *Knowledge-based Sys.*, 1993, **6**(4), 239-48.
37. Mukherjee, R.; Sajja E. & Sen S. A movie recommendation system: An application of voting theory in user modeling. *User Model. User-Adapted Interaction*, 2003, **13**(1-2), 5-33.
38. Schafer, J.B.; Konstan J.A. & Riedl J. E-commerce recommendation applications. *Data Mining Knowl. Discov.*, 2001, **5**, 115-53.
39. Raskutti, B.; Beitz, A. & Ward, B. A feature-based approach to recommending selections based on past preferences. *User Model. User-Adapted Interaction*, 1997, **7**, 179-218.

40. Burke, R. Hybrid recommender systems: Survey and experiments. *User Model. User-Adapted Interaction*, 2002, **12**, 331-70.
41. Rich, E. User modeling via stereotypes. *Cognitive Science*, 1979, **3**, 329-54.
42. Kim, W.; Kerschberg, L. & Scime, A. Learning for automatic personalisation in a semantic taxonomy-based meta-search agent. *Electr. Commerce Res. Appl.*, 2002, **1**, 150-73.
43. Perugini, S.; Goncalves, M.A. & Fox, E.A. Recommender systems research: A connection-centric survey. *J. Intelli. Inf. Sys.*, 2004, **23**(2), 107-43.
44. Minsky, M. A framework for representing knowledge. *The Psychology of Computer Vision*, McGraw-Hill, 1975, pp. 211-77.
45. Itzkovich, I. & Hawkes, L.W. Fuzzy extension of inheritance hierarchies. *Fuzzy Sets and Sys.*, 1994, **62**, 143-53.
46. Rongen, P.H.H., *et al.* Multi agent approach to interest profiling of users. *In Proceedings of Multi-Agent Systems and Applications IV*, 2005, pp. 326-55.
47. Lock, Z. Performance and flexibility of stereotype-based user models. University of York, 2005, PhD Thesis.
48. Kang, H. & Yang, H. The effect of anonymity on the usage of avatar: Comparison of internet relay chat and instant messenger. *In Proceedings of 10th Americas Conference on Information Systems*, 2004, pp. 2734-43.
49. Nowak, K.L. & Rauh C. The influence of the avatar on online perceptions of anthropomorphism, androgyny, credibility, homophily, and attraction. *J. Comp. Mediated Comm.*, 2005, **11**(1), 153-78.
50. Jain, A.K.; Murty, M.N. & Flynn, P.J. Data clustering: A review. *ACM Comp. Surveys*, 1999, **31**(3), 264-323.
51. Dunham, M.H. Data mining and advanced topics. Prentice Hall, 2004.
52. Fink, J. & Kobsa, A. A review and analysis of commercial user modeling servers for personalisation on the world wide web. *User Model. User-Adapted Interaction*, 2000, **10**, 209-49.