Revisiting 3D Information Landscapes for the Display of Art Historical Web Content

Doron Goldfarb, Max Arends, Josef Froschauer, Dieter Merkl
Institute of Software Technology and Interactive Systems
Vienna University of Technology
Favoritenstrasse 9-11, A-1040 Vienna, Austria
vsem@ec.tuwien.ac.at

ABSTRACT

As of today, a lot of different approaches have been dealing with the presentation of art history resources on the Web. While the majority of these focused on traditional 2D methods of display, some did introduce the application of 3D visualization metaphors. Such environments were, however, usually tailored to a specific collection or topic, such as a featured artist or epoch. Meanwhile, the increasing availability of valuable metadata resources has opened up the perspective for the automatic creation of such 3D environments by integrating semantic data sources through the Web. In this work we therefore present a prototype that automatically constructs a 3D environment from semantic art history related Web resources, offering users the opportunity to explore art history following the visualized structure of relations between historical actors of the field. Traversing this historical social network enables users to encounter previously unknown artists and their work in a serendipitous way.

Categories and Subject Descriptors

J.5 [**Arts and Humanities**]: Fine Arts; H.5.2 [**User Interfaces**]: Graphical user interfaces (GUI)

General Terms

Information Visualization, Human Computer Interfaces

Keywords

Web3D, Art History, Visualization

1. INTRODUCTION

The increased interest of Cultural Heritage (CH) institutions like galleries, libraries, archives and museums (GLAMs) in targeting new audiences through their Web presences has resulted in the significantly increased availability of related online resources. National and International projects like the Prometheus[1] image archive or the Europeana[2] initiative seek to combine such resources within dedicated platforms. At the same time, initiatives like Flickr[3] and Wikimedia[4] Commons strive to collect and make available reproductions of artworks that have been put into the Public Domain. Crowd-Sourced semantic knowledge bases like the Wikipedia derivate DBpedia[5] or Freebase[6] enrich the available information to extend beyond the professionally acquired knowledge from institutional data-sources.

Web based presentations of these resources usually rely on human authored content that is put into its respective context by experts or enthusiasts, inviting users to explore art history following various narratives. On the other hand, an increasing number of institutions offer Web based query front-ends to their collection databases. Given the sheer size of the combined content that is for example present in the Europeana portal - about 15 Million items - effective search paradigms have to be applied in order to provide overview and effective navigation within the vast information space.

Some platforms offer advanced search techniques such as faceted search and query refinement in order to allow users to narrow down the results returned from initial queries that often contain just single terms. These methods provide helpful means of acquiring overview on the contents of huge information spaces, especially for people who are not acquainted with the field. Such methods of search are an example of approaches that can be summarized under the term "exploratory search" [15], reflecting a view on the issue of searching large information spaces that focuses on searchers that do not exactly know what they are looking for, i.e. searchers that have only little knowledge of the domain in question, especially regarding knowledge of the structure of the information space and the used terminology.

Targeting users with such constraints poses a challenge, and especially with respect to the domain of Cultural Heritage, it might be of general interest to provide user interfaces that address such "casual" visitors, enabling them to explore unknown information terrains, offering them opportunities to broaden their knowledge during the search process itself.

Short presentation, ACE'2011 - Lisbon, Portugal

[1] http://prometheus-bildarchiv.de/en/index
[2] http://www.europeana.eu/portal/
[3] http://www.flickr.com/commons
[4] http://commons.wikimedia.org
[5] http://dbpedia.org
[6] http://www.freebase.com

We believe that methods from Information Visualization can effectively be employed to provide environments that are dedicated to facilitating information about the structure of the visited information space while, at the same time, offering access to the actual content, thus combining overview and detail within one and the same environment.

2D Information Visualization has been applied to the field of Cultural Heritage. Map displays are an increasingly applied method to convey geographical context information, while graph drawings are used to display relationships between important entities. Such additional structural information can be a valuable companion for exploring unknown collections. As long as just the information structure is displayed without the actual content, such 2D visualizations are sufficient for that purpose. Trying to combine structure and content within one and the same 2D visualization can, however, quickly lead to cluttered display space, effectively reducing the number of simultaneously displayable items to a rather handful.

As a solution to overloaded visualizations, the simultaneous presentation of different views on one information space has been extensively studied within the field of Exploratory Data Analysis. Multiple linked views are an effective way of providing insight when dealing with highly dimensional information spaces. We do, however, believe that such findings from the field of data analysis cannot be directly transferred to the field of Cultural Heritage. The main reason for our viewpoint is that while Exploratory Data Analysis mainly deals with abstract, quantifiable information, information from the field of Cultural Heritage - although of course offering many aspects that can indeed be quantified and abstracted - is tightly coupled to the involved people, their histories and, of course, the respective artifacts. Visualizing such complex coherencies across multiple views might provide additional insight, but, at the same time, drive attention away from the actual content.

3D visualization methods can be used in order to address the mentioned problems, displaying structure and content within one view while reducing the problem of clutter. Moreover, the immersive nature of 3D virtual environments can be utilized for drawing the user's attention to the content.

In a showcase addressing the mentioned issues, we therefore provide a 3D environment that allows casual Internet users to explore art history. The environment is dynamically created from underlying art historical data sources, evolving around an entry-point chosen by the user, in this case, the name of an artist. We perceive art history as a complex network of interactions between people. Artists supported by patrons of mundane or clerical origin, having students or assistants, themselves possibly influenced by other artists, etc.. Thus, one view on the structure of the information space "art history" can be perceived as a social network graph, with historical actors as nodes and their mutual relationships as edges.

The graph is visualized to the user, providing him/her with a narrative flow through time, following the relationships of involved people. Each node of the graph is represented through an image depicting the respective person. The person-nodes serve as entry points for viewing related artworks. We draw the graph in a 3D Information Landscape, providing an immersive environment where users can follow the structure of the information space, use it as navigation aid while at the same time being able to focus on the displayed content. The resulting scene can be accessed via Web browsers by using a specific plug-in[7]. We do believe, however, that the current effort put into establishing a Browser based Web3D[8] standard like WebGL[9] will eventually make plugins for Web based 3D content obsolete.

The paper is structured as follows. The second section provides related work and a comparison to our approach. The third section introduces the used data sources and how they are integrated. The fourth section describes our approach to visualizing art history. This is followed by a description of user interaction and navigation in section five. The sixth section offers a discussion of the current state of the project and strategies for evaluation. The seventh section presents our conclusion.

2. RELATED WORK

The application of Web3D methods to the field of Cultural Heritage is not new. In this section we therefore focus on some important examples that are relevant to our approach. For a general overview of the field and a discussion about the implications of virtuality on the presentation of real world artifacts see for example [13] or [7]. An interesting overview of the application of visualization to the field of art history can be found in [8].

A general distinction can be made between approaches aimed at representing a replica of the real world, simulating the environment of a real or imaginary museum building, or approaches that rather focus on the construction of abstract, data-driven environments. Another distinction can be made by whether the presented content or its layout is mostly human-authored, human-authored assisted by semiautomatic means, or derived from some form of data repository and then laid out automatically. In the following two subsections we give examples for each of these six categories and locate our approach in that respect.

2.1 Replicas of "real" museums

Although only accessible through dedicated client software and thus less representing a Web3D approach, virtual museum presentations in Second Life, as for example discussed in [6], are an important example for human authored approaches that rather stick to the "real world" visualization metaphor - but extending the possibilities of the real physical world by user's ability to fly around, or by the construction of "impossible" architectures.

The Minerva System [1] is an example for a semiautomatically authored real world replica museum. The system allows curators to cluster artifacts according to selected categories and assists the assignment of the items to thematically coherent exhibition rooms that are subsequently automati-

[7]http://www.unity3d.com
[8]http://www.web3d.org/realtime-3d/
[9]http://www.khronos.org/webgl/

cally laid out in the virtual space. The resulting virtual museum has the visual appearance of traditional exhibition premises and is displayed using VRML[10] technology.

Another system proposed in [4] focuses on the direct translation of the results returned from queries against a database into a virtual museum architecture, thus providing an example for an automatically generated virtual museum replica.

Of course, there are many more examples for virtual museums that aim at reproducing the "real" world within a virtual environment, as for example virtual city-tours or walk-throughs through reconstructed archaeologic sites.

2.2 Synthetic data environments

Looking at the other category of virtual museums - approaches that do not try to simulate reality, but rather use the extended capabilities of synthetic 3D environments to invent a "new" reality or display purely abstract data environments - the scope has to be broadened up in order to include systems that are not immediately perceived as being a museum, but nevertheless aim at facilitating cultural or historical knowledge through 3D visualization. Nevertheless, such approaches can also be further classified by the amount of human activity involved in the authoring process.

A good example for an abstract environment based on human authoring is the information landscape approach presented in [11], offering an environment for learning about various aspects of a graphic art collection.

SemaSpace [9] is an example for a synthetic 3D visualization environment dedicated to the display of knowledge networks. It has been used to create a thesaurus on media art based on the archives of the annual Ars Electronica[11] festival. Besides offering tools to explore and analyze large network graphs, it can be also used to create such graphs, thus serving both as viewer and as editor tool. Therefore, it can be seen both as a human-authored as well as an automatic environment, depending on the usage scenario. The authors state an interesting metaphor for this kind of visualization system, the so-called *memory theater*, as also mentioned in [8]. This term dates back to the times of the Renaissance, when the then dramatically growing amounts of produced knowledge demanded new ways of knowledge organization.The idea of the memory theater can be seen as one of the first approaches to the organization and visualization of abstract concepts.

2.3 Comparison

We see our approach both as presentation environment, as well as a search interface for the exploration of art history. We believe that abstract information visualization can be used as an aid for both purposes. Thus, in that respect, our approach falls under the "synthetic environment" category.

In contrast to the approach from [11], we focus on the automatic creation of the environment from heterogeneous Web-based data sources. This is necessary for providing the "exploratory search" aspect of our system. Compared to [9], we want to offer a synthetic environment that nevertheless still contains references to some "real world" counterpart, i.e. a ground surface, a sky above, in other words, simple indicators of known space, in order to simplify navigation and to foster immersion by providing a virtual landscape. Thus, we adopt the notion of information landscapes for Cultural Heritage as stated in [11], although we choose a different approach to constructing such a landscape.

The concept of information landscapes is quite old and dates back to the early nineties of the previous century. The metaphor has been proposed for visualizing hierarchical document collections [2] or digital libraries [17]. With respect to the previously mentioned issues regarding the simultaneous display of structure and content, we nevertheless believe that this approach is well suited to the domain of art history. Moreover, the recent advances in Web3D plugins and upcoming standards like WebGL, in combination with the current wave of affordable consumer-level 3D display devices present on the market, do, however, encourage us to revisit this approach, because we believe that the mentioned technological advances are close to reaching a point where such approaches might eventually be ripe for mass deployment.

3. DATA SOURCES

As mentioned earlier, we seek to visualize art history as a process between the involved people. Therefore, we need data sources that provide information about and reproductions of artworks, and sources that provide us with biographical information about art historical actors.

Regarding publicly available image collections, the Web Gallery of Art (WGA)[12] is a well known representative. This data collection offers information on more than 24.000 artworks and information on more than 2.900 artists covering a timespan between approximately 1000–1900. The provided metadata is, however, limited to information on artist name and birth-and-death dates and places, artwork title, its current location and rough content classification such as type of artwork (painting, sculpture, etc.) or general theme of the artwork (religious, still-life, genre, etc.).

Well established databases that are specifically tailored to serve as contextual information sources for Cultural Heritage related applications are the Getty Thesauri. They represent controlled vocabularies that are providing *terms, names, and other information about people, places, things, and concepts relating to art, architecture, and material culture*[13]. Especially the Getty Union List of Artist Names (ULAN) offers state-of-the-art biographical information regarding artists, patrons and other art historical actors, as well as their mutual relations between each other. Linking the WGA's image collection with the ULAN data thus results in a combined knowledge base of artworks, their creators, and their underlying social network that we seek to visualize.

Regarding the data storage method and platform we chose to follow the approach proposed in [16], as it already provides a framework for an RDF[14] based integration of the

[10] http://www.web3d.org/x3d/specifications/vrml/
[11] http://www.aec.at/festival/
[12] http://www.wga.hu
[13] http://www.getty.edu/research/tools/vocabularies/index.html
[14] http://www.w3.org/RDF/

Getty Thesauri with visual resources that are described using the VRA Core[15] standard. Moreover, we use their open-sourced platform ClioPatria[16] as data storage platform, as it provides a standard SPARQL[17] endpoint and various API's for data retrieval and presentation. As a consequence, we created a VRA Core conforming RDF representation of the WGA data and linked with the respective RDF representations of the Getty Thesauri.

In addition to the mentioned data sources, we also include information from crowd-sourced platforms like DBpedia and Freebase. Currently, we use information from DBpedia to gather portraits of artists, patrons and other important historical figures. The portraits - where available - are displayed as representatives of the respective nodes in the graph.

Moreover, the Genealogy of Influence[18] project on Freebase offers additional relational information between historic figures in form of "influenced/influenced_by" relations. In a comparison of these relationships with the social relations contained in the Getty ULAN, we have found out that the relations from Freebase tend to connect actors across much larger timespans. This is due to the fact that the relationships from the Getty ULAN mainly represent personal connections between people, such as parent/child, teacher/student or artist/patron relations that only exist during the lifespan of the respective persons. Therefore, the integration of the relationships from Freebase enable the connection of such "local" neighborhoods across multiple centuries, leading to a more complete coverage of art history. Figure 1 shows how the local neighborhoods of the artists Raphael and Pierre-Auguste Renoir are connected through a Freebase "influenced" relation.

4. VISUALIZATION

Since we focus on the (historical) social network aspect of art history, we choose a graph based visualization method for display. As we are dealing with historical processes, we are interested in displaying them in a chronological order.

Layered graph layout algorithms as initially proposed in [14] are a useful method for chronological graph visualizations, as each layer can represent a certain point in time, containing all nodes of the graph that are tied to that certain date. The order of the nodes within each layer is determined with respect to the minimization of edge crossings.

In our initial prototype, we therefore first tried to directly use the well-known 'dot' application from the Graphviz[19] software package on the server-side, returning a readily calculated layout for each client request. This, however, proved to be inefficient both from performance considerations regarding server-load as well as being somewhat inflexible by allowing less control over the layout from the client perspective. Thus, we have implemented an adaptation of the 'dot' algorithm as described in [5], where the final stage of the algorithm has been replaced by an approach proposed in

[15]http://www.vraweb.org/projects/vracore4/
[16]http://e-culture.multimedian.nl/software/ClioPatria.shtml
[17]http://www.w3.org/TR/rdf-sparql-query/
[18]http://www.mike-love.net/genealogy/
[19]http://www.graphviz.org

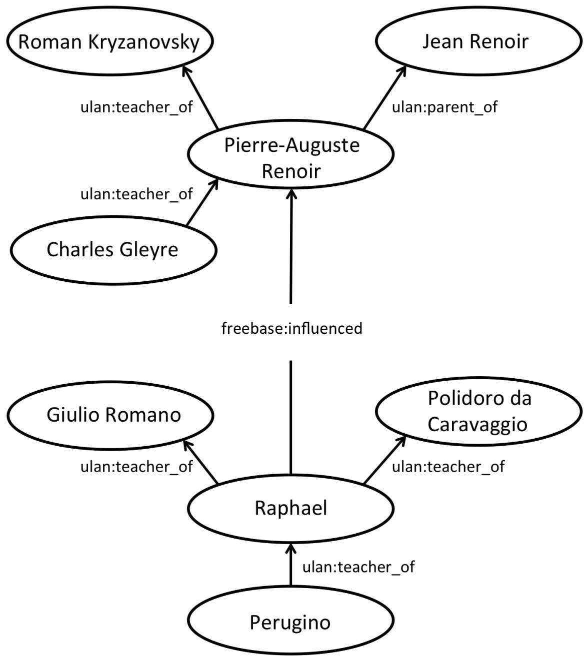

Figure 1: Extending local neighborhoods

[3]. Our adapted algorithm is part of the client which potentially allows for more user control over the layout process.

Another modification of the basic graph layout is introduced in the stage of constructing dummy vertices for edges spanning multiple layers. Here we introduce a concept derived from so called flow map layouts, as for example proposed in [10]. Flow maps are usually used for the visualization of geographical flows like migration, transfer of goods, etc.. We use the concept of flow maps in a chronological sense instead. We consider outgoing relations from a person, i.e. relations where the counterpart usually lies in the future (parent_of, teacher_of, etc.), as a flow in time, representing the amount of influence that a historic person has had on following generations. The amount of influence is represented by bundling the outgoing edges of a node together, resulting in streams whose width corresponds to the number of outgoing relations. We refrain from straightening long edges, as we believe that their meandering appearance reflects a beautiful metaphor of rivers flowing through the landscape. Figure 2 shows a snapshot of the graph layout with the resulting bundles of outgoing edges. Time is flowing from left to right and the edge bundles get thinner and thinner as the contained edges reach their endpoints.

We transform the graph layout into a landscape by promoting nodes of high degree, i.e. persons that have a large number of in- and outgoing relations. Thus, nodes are assigned with a height value corresponding to their degree. The coloring of the edges reflects the different types of relationships. In order not to clutter the display space, all the edges are drawn with high transparency. Only edges ending or starting at the currently visited node are drawn with full opacity. This enables the visitor to observe relevant connections for the current position without destroying the overall

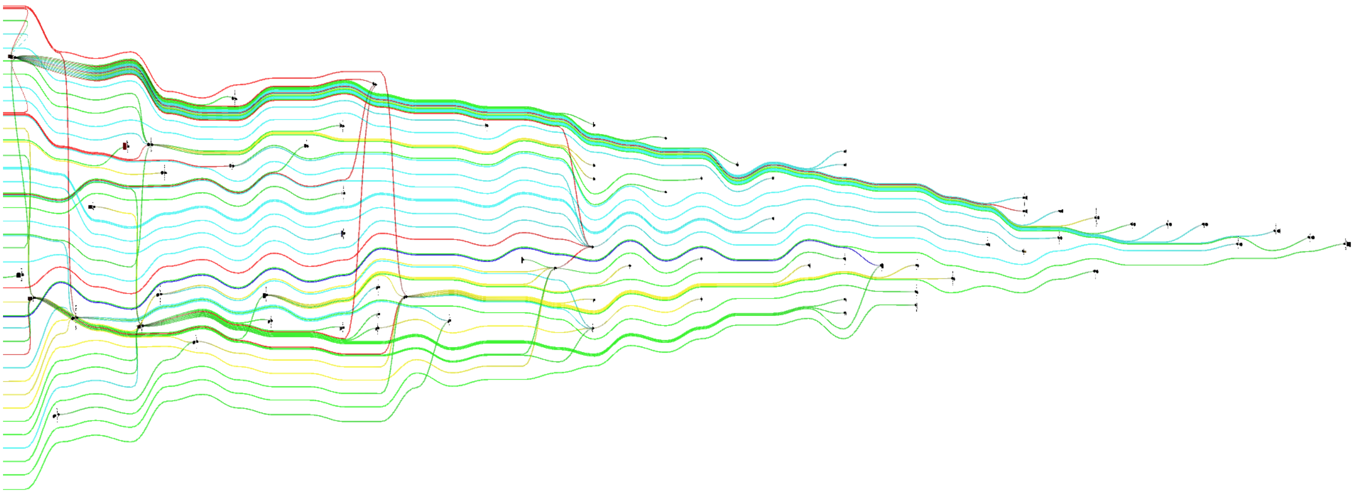

Figure 2: A part of the graph layout showing the different flow widths

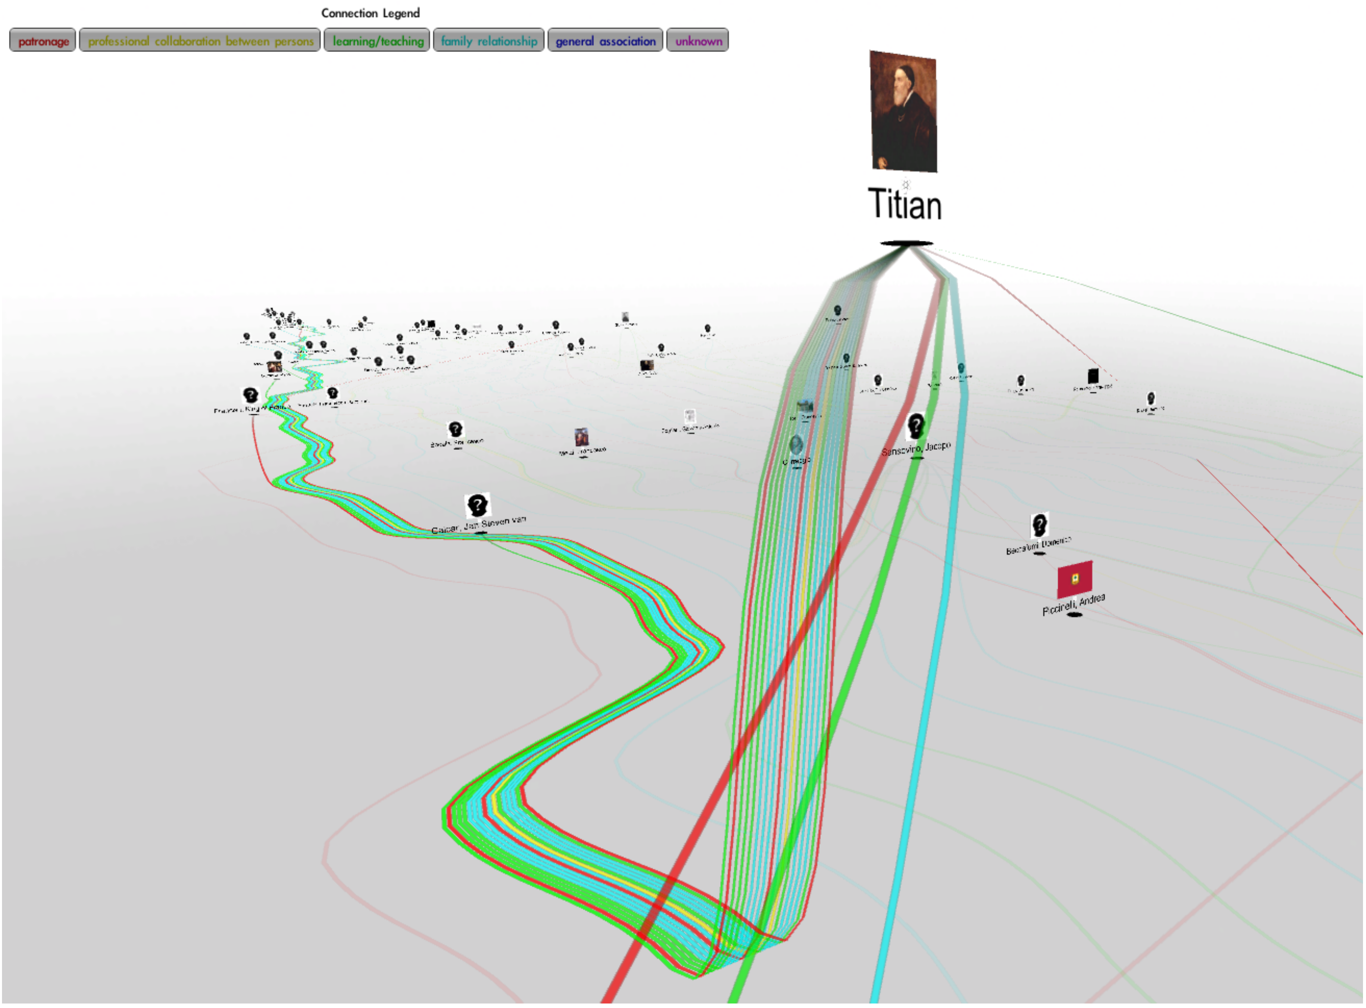

Figure 3: In the neighborhood of Titian, looking into the future

notion of landscape. In order to provide a sense of orientation within the highly abstract environment, the ”ground” is shaded in light grey, while the ”sky” remains white in order not to interfere with the content.

The names of the respective persons are placed above the nodes. If available, a representative portrait is downloaded from the corresponding DBpedia resource and shown above the name. In the case that no portrait can be found, a default image showing the silhouette of a person’s head with a question mark on it is displayed instead. Figure 3 shows the graph from the user’s perspective. The node representing the painter Titian is currently activated, his relationships to other people highlighted in the landscape. Nodes representing persons with many relationships are visible from far away, thus potentially drawing attention.

5. USER INTERACTION

A session is initiated by entering the name of an artist into a text-box. An autocomplete feature offered by the ClioPatria server assists in finding persons of interest. After the choice has been made, the neighborhood of the selected person is recursively traversed along his or her connections. The recursion currently stops after four steps. After the loading sequence, the graph layout sequence is initiated. After the layout phase, the environment is ready to be visited, portraits of people are loaded in parallel.

In the beginning, users are positioned in front of the initially chosen person. By turning around, they can get acquainted with the scene, observing incoming and outgoing relations to and from the currently activated person. As mentioned before, edges entering or leaving the currently visited node are highlighted by setting their opacity to a high level. This avoids visual interference with the remaining edges of the scene, which are drawn with higher transparency.

Currently, users have to use the mouse to fly through the landscape. A camera zoom is controlled by the mouse scroll wheel. The focal length of the zoom also determines the speed of flight. The idea behind the zoom functionality is to provide users with the ability to better see what lies ahead (or behind) and to reach positions of interest more quickly. The ability to zoom does also support the landscape metaphor by putting the visitor into the role of a hiker through the cultural landscape, stopping at points of interest, using binoculars to search for other interesting spots. High focal lengths do also result in a flat perspective, effectively re-transforming the 3D landscape into a 2D canvas. Figure 4 shows Titian and his successors at high zoom level. The high focal length reduces differences in size and thus makes it possible to compare distant nodes in a more side-by-side fashion. In order to achieve a zoom functionality that enlarges distant images in the scene while keeping the display size of the currently viewed image constant, we employ a well-known method from cinematography, the so-called dolly-zoom as described in [12].

As a navigational aid, an overview map can be activated by key-press. Initially, we tried to provide overview by showing a map containing the graph from the bird’s eye perspective. This turned out to be unfeasible due to screen resolution issues. Moreover, a second view of the graph itself could

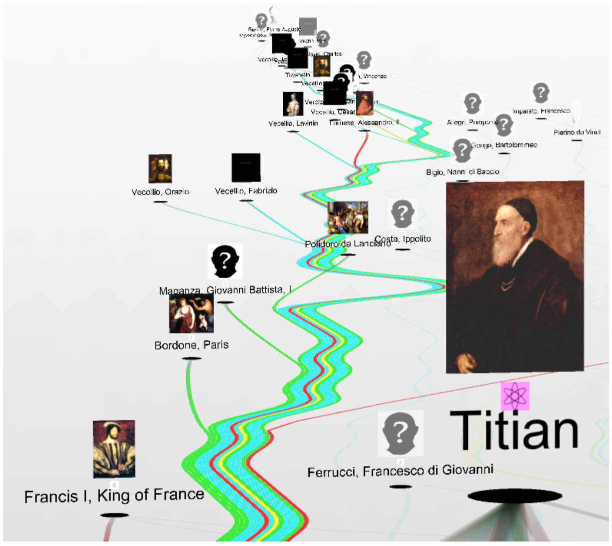

Figure 4: Zooming on Titian and his successors

lead to confusion. Therefore, we decided to provide an abstraction of the bird’s eye perspective by only showing the names of the persons along the current year. Since the simultaneous display of all node names currently displayed in the landscape would quickly lead to a cluttered map, we only show names that are connected to the currently visited node. The name of the currently visited person is highlighted in black, and the colors of the names of all connected persons reflect the type of relationship. Figure 5 shows the resulting

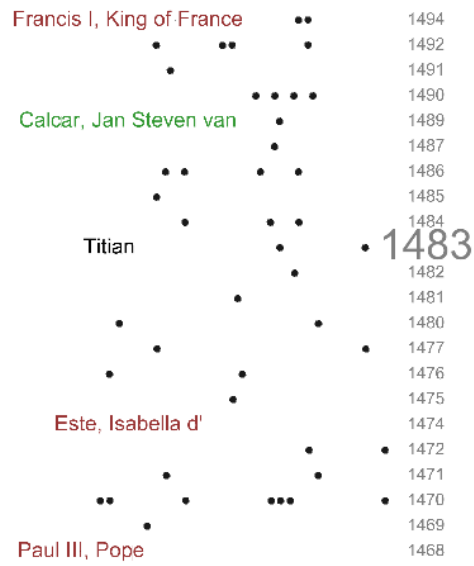

Figure 5: Part of the overview map around Titian

overview map. The overview map can be used for ”teleporting” to points of interest. As all connected persons are highlighted together with the currently visited actor, they can be clicked and the user’s position is translated to the respective position. The nodes representing artists also serve as entry points for the display of respective artworks. Repeated Clicking on a node cycles through the available artworks. Additional information is displayed when the user hovers the mouse over an image. Figure 6 shows the display of biographical information about Titian. The number of available artworks is displayed along the title of the currently displayed image.

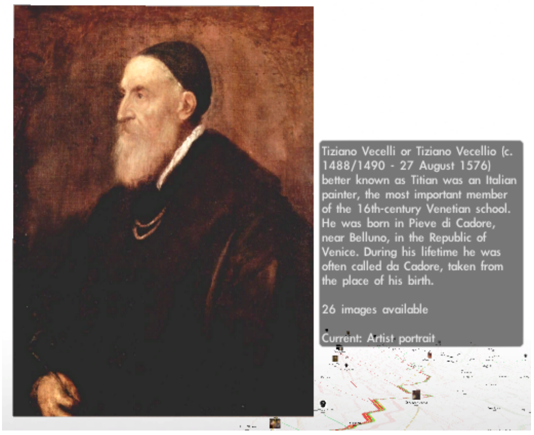

Figure 6: Displaying additional information

In order to support exploratory search behavior, we also provide means to narrow down the displayed set of persons and relations. Users can choose what types of actors should be displayed, e.g. only artists. Moreover, the types of relations to be displayed can be selected as well. If a user is only interested in, for example, displaying teacher/student relationships, he or she can disable the display of other relation types directly in the connection legend (see Figure 3).

By moving along the relations between art historical actors, users will most likely encounter historic figures that were not known to them before. We expect the effect of finding new content that is related to known content by social relations of the respective creators to be an effective method to provide a serendipitous browsing and learning experience.

6. EVALUATION

The current state of the environment is in its third revision. The initial approach was to provide a more natural

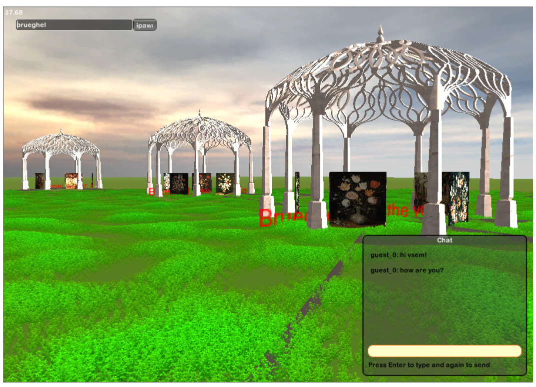

Figure 7: A screenshot of the first prototype

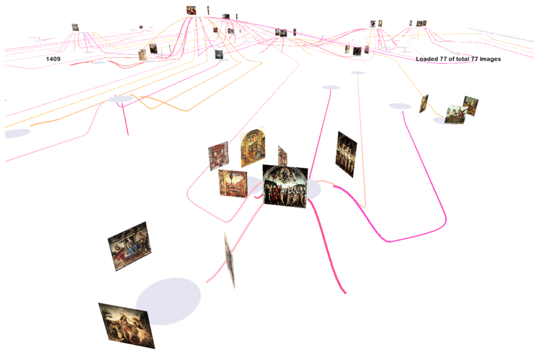

Figure 8: A screenshot of the second prototype

environment where artists were represented by pavilions, in which the respective artworks were shown. This visualization metaphor, shown in Figure 7, was, however, not considered as being appropriate during internal pre-evaluations. The main criticism was that the environment drew away too much attention from the actual content. Therefore, the whole scene was reduced to its current, graph based form, using the pre-calculated layout from the 'dot' application (see Figure 8). In a following internal pre-evaluation, potential users complained about the confusion caused by the placement of edges as determined by the crossing minimization, as well as about the somewhat artificial appearance of the straightened, parallel edges. This, and the previously mentioned considerations regarding server-load and control over the layout process, led to the implementation of our own layout algorithm including the introduction of the edge bundling functionality. Moreover, in order to reduce the potentially confusing variety of colors representing all the different person-person relationships, we reduced the number of available relationship-categories by creating superclasses for common types like family-related, professional, etc..

Additional feedback was collected for the navigation aids. The initial, static version of the overview map was rejected because it was not usable due to insufficient screen resolution. Then, we tried to provide a "sliding window" like map overview. The additional movement of the graph's bird's eye perspective together with the actual visualization was, however, considered as confusing. This lead to the current state of the overview map showing names and dates only.

We consider the current state to be ready for a more formal evaluation and expect to refine the approach through additional revisions. We are planning to evaluate the system twofold. A first field-test should be performed by offering an online version of the system and collecting responses by the users in form of a questionnaire. This should provide us with further feedback on the usability of the system in order to identify the main problems encountered. A more controlled lab evaluation should then try to assess the advantages of the system compared to traditional methods of presentation. We plan to create a Wikipedia like hypertext version of the same content that is used for our environment and compare the users experiences in form of an experimental and a control group.

7. CONCLUSIONS

In this paper, we have described our approach to dynamically present online art historical Web resources in form of 3D information landscapes. We have introduced our motivation to perceive art history as social network of relations between historical actors and described how we have gathered respective data sources for that purpose. Then we have outlined a method how to transform the network into a 3D representation on the basis of a layered graph layout algorithm, followed by a discussion of the applied visualization metaphors and user interaction principles. The proposed method enables non-expert users to intuitively follow historical interpersonal relationships in order to get a better understanding of the evolution of art throughout time. Finally, we have presented insight gained from initial design iterations with intermediate pre-evaluations and our plans for the first larger-scale evaluation. Our next steps will contain a Web-based usability study and an evaluation of the effectiveness of information transfer compared to classic text-based methods.

8. ACKNOWLEDGMENTS

This work was funded by the FWF Austrian Science Fund, Project No. L602, *The Virtual 3D Social Experience Museum.*

9. REFERENCES

[1] F. Amigoni and V. Schiaffonati. The minerva system: A step toward automatically created virtual museums. *Applied Artificial Intelligence*, 23(3):204–232, 2009.

[2] K. Andrews, M. Pichler, and P. Wolf. Towards rich information landscapes for visualizing structured web spaces. In *Proceedings of the 1996 IEEE Symposium on Information Visualization (INFOVIS '96)*, pages 62–63, 121. IEEE Computer Society, 1996.

[3] U. Brandes and B. Koepf. Fast and simple horizontal coordinate assignment. In *Proceedings of the 9th International Symposium on Graph Drawing (GD '01) (LNCS 2265)*, pages 31–44. Springer, 2002.

[4] E. Ciabatti, P. Cignoni, C. Montani, and R. Scopigno. Towards a distributed 3d virtual museum. In *International Working Conference on Advanced Visual Interfaces (AVI '98)*, pages 264–266. ACM, 1998.

[5] E. R. Gansner, E. Koutsofios, S. C. North, and K.-P. Vo. A technique for drawing directed graphs. *IEEE Transactions on Software Engineering*, 19(3):214–230, 1993.

[6] S. Hazan. Musing the metaverse. In *Heritage in the Digital Era.* Multi-Science Publishing, 2010.

[7] E. Huhtamo. On the origins of the virtual museum. In *Virtual Museums and Public Understanding of Science and Culture.* Nobel Foundation, 2002.

[8] K. Kwastek. Visualising art history. In *Proceedings of the 10th Conference on Computers and the History of Art (CHART '03).* CHART, 2003.

[9] D. Offenhuber and G. Dirmoser. Semaspace - semantic networks as memory theatre. *http://residence.aec.at/didi/FLweb/semaspace.pdf*, 2006.

[10] D. Phan, L. Xiao, R. Yeh, P. Hanrahan, and T. Winograd. Flow map layout. In *Proceedings on Information Visualization (INFOVIS '05)*, pages 219–224. IEEE, 2005.

[11] E. Ruffaldi, C. Evangelista, V. Neri, M. Carrozzino, and M. Bergamasco. Design of information landscapes for cultural heritage content. In *Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts (DIMEA '08)*, pages 113–119. ACM, 2008.

[12] K. Singh, C. Grimm, and N. Sudarsanam. The IBar: a perspective-based camera widget. In *Proceedings of the 17th annual ACM symposium on User interface software and technology(UIST '04)*, pages 95–98. ACM, 2004.

[13] S. Sylaiou, F. Liarokapis, K. Kotsakis, and P. Patias. Virtual museums, a survey and some issues for consideration. *Journal of Cultural Heritage*, 10:520–528, 2009.

[14] M. Toda, K. Sugiyama, and S. Tagawa. Methods for visual understanding of hierarchical system structures. *IEEE Transactions on Systems, Man, and Cybernetics*, 11:109–125, 1981.

[15] R. White and R. A. Roth. *Exploratory Search: Beyond the Query-Response Paradigm.* Morgan & Claypool, San Rafael, CA, USA, 2009.

[16] J. Wielemaker, M. Hildebrand, J. van Ossenbruggen, and G. Schreiber. Thesaurus-based search in large heterogeneous collections. In *Proceedings of the 2008 International Semantic Web Conference (ISWC '08), LNCS 5318*, pages 695–708. Springer Verlag Berlin Heidelberg, 2008.

[17] J. A. Wise, J. J. Thomas, K. Pennock, D. Lantrip, M. Pottier, A. Schur, and V. Crow. Visualizing the non-visual: spatial analysis and interaction with information from text documents. In *Proceedings on Information Visualization (INFOVIS '95)*, pages 51–58. IEEE, 1995.