DocuWorld — A 3-D User Interface to the Semantic Desktop

Katja Einsfeld[1], Achim Ebert[1], Stefan Agne[2], and Bertin Klein[2]

[1] Visualization Lab, Technical University Kaiserslautern
P.O. Box 30 49, 67653 Kaiserslautern, Germany
{k_einsfe,ebert}@informatik.uni-kl.de
[2] German Research Center for Artificial Intelligence (DFKI GmbH)
P.O. Box 20 80, 67608 Kaiserslautern, Germany
{stefan.agne,bertin.klein}@dfki.de

Abstract. With the aim of replacing the 20-year-old 2-D desktop metaphor with a more natural, intuitive to use, and immersive 3-D interface, we developed a prototype application that allows the dynamic exploration of document collections. Document metadata and relations between documents are visualized with the help of information visualization techniques. The general visualization and navigation metaphor called *Thought Wizard Metaphor* allows user- and context-sensitive adaption of visualization modes and visualization environments. The automatic document organization modes provide different insights in document semantics and search results. These possibilities are supplemented by the possibility to store documents at user-defined, hierarchically organized locations, which allows quick relocation using spacial memory.

1 Introduction

Conventional document or file system viewers like the Windows Explorer or KDE's Konqueror are quite unintuitive to computer newcomers who know nothing about file system architecture. They sometimes have problems understanding where they "are" in the file system and remembering where their documents are stored. But advanced computer users with huge document collections stored in different hierarchical folders also suffer from the disadvantages of conventional file system viewers: they have to use their cognitive energy remembering names and locations of documents, navigating though the file system hierarchy, opening and closing folders, and looking through alphabetically ordered document lists. There is no possibility to use semantic metadata in an adequate way.

Advances in computer graphics hardware and software offer new possibilities not only for computer games but also for daily-use applications. Information visualization techniques, together with the ability to generate realistic real-time, interactive applications can be leveraged in order to create a new generation of document explorers. An application environment that resembles real environments can be used more intuitively by persons without any prior computer knowledge. The users perceive more information about the documents, such as their size, location, and relation to other documents, visually using their natural capabilities to remember spatial layout and to navigate in 3-D environments, which moreover frees cognitive capacities by shifting part of the information-finding load to the visual system. It is hoped that this leads to a more efficient combination of human and computer capabilities: the computer has the ability to quickly search through documents, compute similarities, calculate and render document layouts, and provide other tools not available in paper document archives, while humans can visually perceive irregularities and intuitively interact with 3-D environments. Last but not least, a well-designed virtual-reality-like graphical document explorer is more fun to the average user than a conventional one and thus more motivating.

This work seeks to evaluate existing 3D file system and document collection approaches and, using these experiences together with new ideas and techniques, to develop and implement a stereoscopic and immersive system for efficiently organizing documents. Special emphasis will be on navigation in information spaces, which includes, for example, the visualization of relations between documents and the possible ways of somehow "moving" from one document to another—i.e., the realization of a seamless 3-D user interface to the semantic desktop.

2 State of the Art

2.1 The Definition of Information Visualization

Visualization is the mapping of abstract data to visually perceptible attributes like color, shape, size, location, texture, etc. with the intention of leveraging the visual system to recognize or view patterns, trends, irregularities, and connections that are not easily perceivable in text-only data collections. This can be done because evolution equipped humans with the ability to think, abstract, remember, understand, and organize visually. Thus, using these visual capabilities is completely natural, ergonomic, and efficient for human beings.

Scientific visualization focuses on the visualization of data obtained from experiments or simulation of natural or technical phenomena. This data is mainly of the numerical type and the values often have an intrinsic spatial representation. *Information visualization*, by contrast, is concerned with the visualization of *any* kind of data, from graphs to film to tables to web pages. Since information visualization plays a central role in this work, an overview will be given in the following subsections.

As Tufte put it, the goal of information visualization is the display of as much data as possible in an area as small as possible in a way such that the perception of information per time is maximized. The intention of a visualization might be the presentation of information, the "examinination" analysis of hypothetical assumptions, or the "exploration" analysis of yet-unknown datasets. The data to be visualized can have the form of tables, graphs, hierarchies, multimedia data, processes, hybrid data, or any other form. In order to accomplish with this variety of data, a huge diversity of information visualization methods has been developed.

2.2 Visualizations and Navigational Concepts for Document Information Spaces

There are many examples of visualizations of document information spaces in the literature, and many concrete applications presented on the Web. This section aims to give a short overview of various approaches.

The *information cube* introduced by Rekimoto et al. [1] can be used as shown in Fig. 1 to visualize a file system hierarchy. The nested-box metaphor is a natural way of representing containment. Rendering the boxes semitransparently allows the user to gain an overview of the structure. One problem with this approach is the difficulty in gaining a global overview of the structure, since boxes contained in more than three parent boxes or placed behind boxes of the same tree level are hard to observe. The authors also suggest using transparency to visualize groups like aggregated nodes in a 3-D graph by rendering semitransparent spheres containing the aggregated nodes.

Fig. 1. Rekimoto et al.'s InfoCube

Card et al. [2] present some metaphors for the visualization of collections of web pages. The web pages resulting from, for example, a search query or from a person's home page, are visualized as a so-called *WebBook*, which is labeled and allows page-turning and the ruffling of pages (both animated) as in real books. The authors claim that the familiar book metaphor with its various features allows intuitive use. WebBook can be zoomed in and out to different degrees, as with a magnifying glass. WebBook is integrated in an 3-D environment called *Web Forager* (see Fig. 2). In this 3-D space, web pages and WebBooks can be placed in different distinct locations like a desk, a focus plane, the air, or a bookcase.

Fig. 2. The Web Forager

Taking into account the cost structure of information workspaces and the locality of reference principle (which states that users tend to interact repeatedly with a small number of documents), Card et al. decided to create a *hierarchical workspace* with three main levels that can be used with increasing cost of time.

The first level is the focus plane, where one page or a book can be shown. The second level is called "immediate memory space". In this level the documents can be placed on the desk or at several distinct z-planes in the air. The third level is a bookcase, placed parallel to the z-axis in 3-D space, to store WebBooks. These books can be moved to other places in space and opened, but there is also the possibility for the user to move in an animated way to face the bookcase.

WebForager's hierarchical workspace is certainly a useful way of organizing and navigating documents in 3-D space. One drawback of this system, however, is that the user gets (apart from the possibility of one search query) no computer assistance in organizing the documents in space as well as in mental categories.

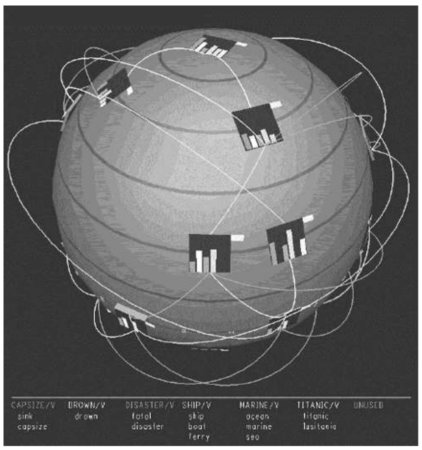

Fig. 3. The 3D NIRVE system

3D NIRVE, the 3-D information visualization presented by Sebrechts et al. [3] and shown in Fig. 3, helps the user to organize documents that result from a previous search query. Concepts consisting of one or more keywords are named and color-coded in the bottom of the window. Documents that contain all keywords of a concept are grouped together in a cluster and placed on the surface of a sphere. The more concepts are found in a document, the closer it will be to the north pole. In general, documents on the same latitude of the sphere all have the same number of concepts. The most important features of this system, with respect to user behaviors, are the grouping of documents into clusters and the color-coding of concepts.

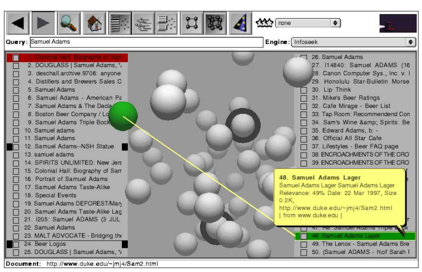

Fig. 4. The Lighthouse system

The *Lighthouse* system described by Leuski et al. [4] is used to visualize search results of web queries by integrating the traditional ranked list with a clustering visualization; the user thus benefits from the advantages of both approaches. Figure 4 shows how the ranked 2-D text lists of the results are connected to the visualization of the spheres representing the web pages. The cluster visualization places the document-spheres according to statistically determined similarities (similar documents are drawn close to each other). Since closely associated documents tend to be relevant to the same requests (cluster hypothesis), the cluster visualization gives the user more possibilities to find what she or he is looking for than a ranked list.

There are also some extra features to let the user judge the relevance of some of the documents and to use this information to calculate the relevance of unjudged documents. When in suggestion mode, this relevance value is visualized by shades of red and green, plus stars to mark the potentially most relevant documents.

A disadvantage of this system is that the double representation of documents forces the eye to often look to and fro, and maybe even to adapt the focus, when the sphere is positioned in the distance.

Fig. 5. The Task Gallery

Robertson et al. [5] developed *Task Gallery*, a 3-D window manager, that can be regarded as a simple conversion of the conventional 2-D windows and desktop metaphors to 3-D. The metaphor used for this approach is an art gallery—a linear hallway with a platform at the end. The current task is placed on the platform, and an ordered and an unordered stack help organizing the windows. The user carries a so-called "start palette" in his or her virtual hand. It looks like an artist's palette (in line with the gallery metaphor) and can be used exactly in the same way as the MS-Windows Start Menu to start applications or to open documents. Inactive tasks are represented by screenshots of their window configuration when they were on the platform, placed like artwork on the walls or even on the floor or ceiling of the gallery. By clicking on them, an animation is started to move the current task to its place on the wall and to move the selected task to the platform, making it the new active task.

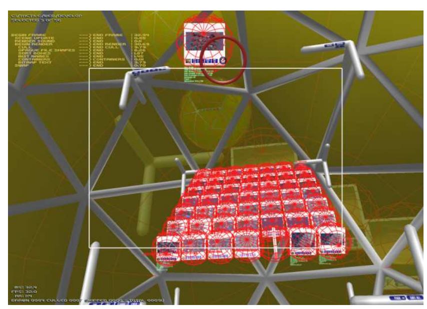

Fig. 6. Tactile 3-D

Figure 6 shows a screenshot of the *Tactile 3D* [6] system, a commercial 3-D user interface for the exploration and organization of documents; it is still in development. The file system tree structure is visualized in 3-D space using semitransparent spheres that represent folders and that contain documents and other folders. The attributes of documents are at least partly visualized: the file type defines the shape and the texture of the 3-D object representing the file. Special folders like drives also have special shapes. Hidden or read-only status is indicated by semitransparent material or by a red hue, respectively. The rotation speed of the objects informs the user about the last modification date of the object.

Objects within a container can be placed in a sorting box and thus be sorted by various sorting keys in the conventional way, forming 3-D configurations like a double helix, pyramid, or cylinder. The objects that are not in the sorting box can be organized by the user.

3 DocuWorld Design Decisions and Realization

3.1 The Data Visualized in DocuWorld

The data items to be visualized in the context of DocuWorld are documents with some metadata and some relations to each other. One problem in the development of DocuWorld was that not many concrete document metadata and document relations of special importance could be defined at the state of development of @VISOR, the project which subsumed this work. Therefore, some possibly useful document properties had to be defined.

Some possible attributes which might be considered as metadata of the documents, are as follows:

- (in some cases) **explicitly given data** such as the title, physical path, size (number of pages is usually more than file size), date of creation, date of last modification, date of last use, file format (e.g., PDF, HTML, plain text), document type (e.g., email, paper, report), semantic category, language, origin (author, institute), text body, and summary;
- **directly computable data** such as a thumbnail of the first page or a logo;
- **values computed by artificial intelligence** such as relevance (e.g., number of hits) according to search queries or relevance according to some example user-rated documents, important keywords, and word indexes; and
- **user defined data** such as notes, annotations, ratings, and comments.

When some of the data (for instance, number of pages) is not explicitly given, it can be computed in some cases. In the chosen document visualization, the categorical data "document category" is represented by the color of the book cover, the ordinal data "number of pages" is visualized as the thickness of the book, and a dynamic quantitative value is represented by the yellowness of the front page.

Possible types of relations could be organized in similar categories:

- **Explicitly given relations** are tree (e.g., file system tree) or graph structures.
- **Computed relations** can be, for example, similarities between documents according to search queries or other criteria.
- **User-defined relations** might be task-specific connections between documents (e.g., document A reminds the user of document B that might also be of interest to the current task).
- **Implicitly given relations** between documents are those that become clear when comparing several document attributes and recognizing similarities and differences between documents or groups of documents. These relations

are usually visualized implicitly through the visualization of document attributes.

In the case of DocuWorld, the example documents are about 200 PDF documents in the two categories "knowledge management" and "information visualization". Data for these documents includes the file name, number of pages, category, and text, including a Lucene [7] word index that allows powerful search, and is given though input methods. Other data, like relations to other documents, is artificially added to each document by dynamic initialization or defined by the user during execution time, as in the case of user-marked documents. In the final version of the application, the relations will be acquired by systems like Gnowsis [8, 9].

3.2 Global Design Decision: The Thought Wizard Metaphor

When considering the organization of documents in 3-D space, the *locomotion metaphor* "library", where the user moves through rows of bookcases, might seem a natural solution at first. There are, however, some valid arguments for the use of a *theater metaphor*, with which no movement is necessary:

- The library metaphor would work only if each document had a static position in 3-D space that could be remembered and relocated by the user with the help of his or her spatial memory. Useful information spaces should, however, allow the user to interact with the information and the visualization should react dynamically to the user's requests and needs; this is hard to realize with a static library metaphor. The concept of dynamically reorganizing document collections could be implemented with variations of the theater metaphor.
- According to the characterizing differences between locomotion and theater metaphors, locomotion metaphors are "egocentric" and theater metaphors "docucentric". In the context of a document collection application, the user's position is less important or even of no particular meaning. Therefore, the information load needed for understanding and suitably changing the user's position can be used for the more important analysis of document positions, content, and relations between documents.
- Theater metaphors also have the advantage that users with no experience in navigating in virtual 3-D environments and, in complex environments, also expert users are not in danger of getting lost in 3-D space.

The global visualization and navigation metaphor that was chosen with regard to these considerations is a variation of the theater metaphor and will henceforth be referred to as the *thought wizard metaphor*. When using this metaphor, the user (i.e., the viewpoint) is fixed and the visualization elements are dynamically moved around with respect to the current user need. Apart from restricting the viewpoint to a static point of reference, this allows the use of multiple creative dynamic visualizations: small scene elements like document representations can be moved around in the viewing volume and create various 3-D structures, decorations or global structure elements like walls can change their appearance, and all visualization elements can be moved out of the field of view, possibly revealing other elements behind them. The mental model that should develop in the user's mind when interacting with a application implementing the thought wizard metaphor might be close to that of being an old wise woman using magic tricks to get all the information she needs, directing multiple documents simultaneously to a specified 3-D location without effort, and having the power and cleverness to do all necessary knowledge work by comfortably sitting behind her desk and using gestures and magic words.

3.3 A Short Description of DocuWorld

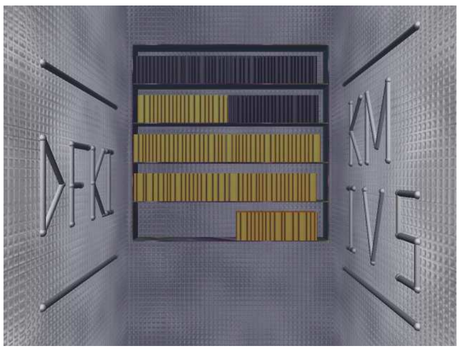

Fig. 7. Initial layout of DocuWorld

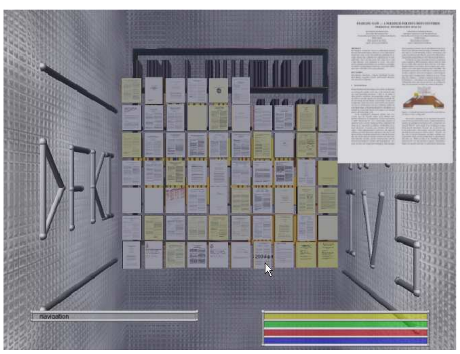

Fig. 8. Preselected documents

The DocuWorld application starts with a visualization of all documents in the state of being stored in some place in the background. As it can be seen in Fig. 7, the start layout of the final realization of DocuWorld displays all documents in the form of books standing at a bookcase in the back of the room. The user can preselect documents of this collection by typing a search query in the gray search query panel. The search panel is invoked by a gesture, which moves it to the bottom of the screen with the gray search field on the left side.

The preselected documents can be thought of as belonging to a higher semantic zoom level than those in the bookcase and should therefore be displayed with more detail. Figure 8 illustrates the visualization state of DocuWorld when some documents have been preselected. These documents are moved out of the bookcase, rotated so that their textured front page faces the user, and flown to their place in the start configuration of the plane layout (which will be described later). While the preselection search query has to be typed in the left search query panel, additional search queries typed in the color-coded right search query panels can be used for more detailed requests. The results of these search queries are visualized in various modes that allow different insights in the data collection.

Important documents will show an animated pulsing behavior and thus instantly catch the user's eye. There is also the possibility to mark documents with red crosses in order to define user- or task-specific interests and not to lose sight of them when they move to different places in different modes. To get a closer view of single documents, they can be moved to the front (the reading position). Although DocuWorld does not support these features, ruffling through pages and annotating the text could be allowed in reading mode. In addition to the reading mode, a preview image of the front page of the documents is shown in the upper right corner of the screen, when the mouse moves over documents (see Figure 8). To gain better insight in the 3-D structure of documents organized in 3-D space, the structure can be rotated by mouse gestures or rocked with an efficient predefined rotation.

3.4 Special DocuWorld Visualization Concepts

Structure on Demand The term "details on demand" is commonly used in the context of semantic zoom and interactive focus and context views. In these systems, the user starts with an overview visualization and demands more details of special visualization items. In trying to find a good compromise in the trade-off between the number of information items being displayed and the perceivability of the structure, the "details on demand" concept was transferred to the "structure on demand" concept. The idea of this concept is to find a context- and user-sensitive solution to the information structure trade-off by allowing the user to call up and remove certain structure elements depending on the current task and user specific preferences.

In the implementation of DocuWorld, the search panels can be called up and removed by a gesture, which moves them in or out in an animated way. The memory-line structure described in Sect. 3.7 and shown in Fig. 9 can also be called up on demand, which is quite useful as it is positioned in front of all other documents and thus party occludes them. Even the main DocuWorld structure, the walls, can be removed to reveal the area behind it in which 3-D structures for user-specific placement of documents are located, and called again when the walls are needed to perceive the 3-D structure of the main document visualization more easily. The bookcase and the main document visualization area can also be called up on demand.

Ghost Documents DocuWorld will allow artificially computed semantic structures as well as documents organized by users. This leads to the problem of how to treat user-placed documents: on the one hand, the mental model of the document representation is that of an object with an identity that can be located at only one place at a time; on the other hand, documents placed at user-specified locations are of special interest to the user and it might be helpful to see their position in a computed semantic structure as well, but displaying one document at different locations would contradict the idea of document identity.

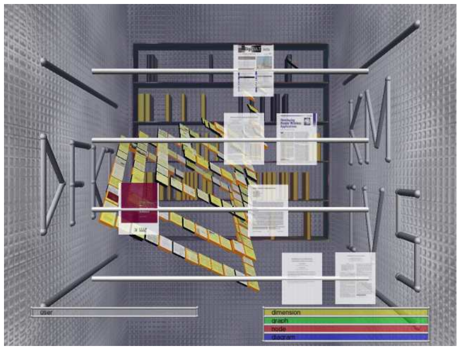

Fig. 9. Ghost documents at memory-line

The solution implemented in DocuWorld is that of *ghost documents*: "real" documents are located only in artificially organized areas. Whenever the user decides to store a document at a specific place, a "ghost" of this document is automatically created which is moved from the original document to the selected storing place, while the original document stays where it was. Figure 9 shows a computed document structure in the background and some ghost documents on the memory-line structure described in Sect. 3.7. Ghost documents, in contrast to real documents, have no 3-D shape, but only a flat, slightly transparent surface which makes them look immaterial and ghost-like. It is possible to visually indicate the document the ghost stands for with a special gesture that triggers the rendering of a semitransparent blue box around the corresponding real document.

Tactical Zoom In conventional 3-D visualizations with perspective projection, the operation of scaling and the operation of moving the viewpoint along the z-axis are isomorphic operations, which means that moving the viewpoint has the same result as a particular scaling operation. This is, however, only true for 2.5-D visualizations: as soon as stereo viewing is enabled, a scaled down object in the front can be distinguished from an object moved to the back due to binocular disparity, that is more or less obvious depending on the quality of the 3-D output device. Nevertheless, the (near) isomorphism between scaling and translating in the z-direction makes it possible to use one operation instead of the other or to combine both operations in order to create new creative visualization or interaction tools. Mine et al. [10], for instance, used automatic scaling for their *scaled world grab*: instead of moving the object to the user's hand or the user to the object, the world containing the object is scaled until the user can reach the object.

For the purposes of DocuWorld it was decided to experiment with z-axis translation and scaling in a different approach: abolishing the natural semantic connection between translation and scaling might allow use of the scale factor as an additional visualization dimension. The idea, called *tactical zoom*, is to use z-axis translation simply and intuitively to organize documents in 3-D space and to use the scale or zoom operation not as a substitute for z-axis translation, but for other, tactical purposes. In accordance with this idea, the physical size of documents would not be fixed, but rather opportunistically changed depending on the visualization situation, thus making more efficient use of 3-D space. If, for example, the documents in the bookcase need to have a particular size to all fit into the bookcase, and another document size would be more appropriate for organizing the documents in 3-D space, they could simply be resized. If, instead of rescaling, only a conventional translation along the z-axis had been used, this would have lead to a non-optimal use of space, as in the example of the bookcase standing in the middle of the scene or even objects intersecting the borders of the viewing frustum in the case of the organic structure moved closer to the viewer. Tactical zoom is also used when objects enter or leave the reading state: as it was decided to move documents for optimal readability to the focus plane perpendicular to the z-axis because there is no binocular disparity and non-rescaled documents are still too small to be readable at that distance, they are rescaled to cover the whole screen. Of course, this also leads to conflicting depth clues. Nevertheless, it is not important to the success of the application to understand at what distance the documents in PlaneMode are compared to the one in reading mode or the ones at the bookcase. It is only essential to perceive the spacial relationships and distances between documents in one mode.

3.5 Visualization of Search Results

One DocuWorld visualization mode for search results is PlaneMode (see Fig. 10). The preselected documents are organized with the front page facing the user in an xy-plane in front of the bookcase. The size of these documents is optimized depending on the number of preselected documents according to the tactical zoom concept. Semitransparent color-coded score bars in front of the documents display the relevance of the individual search queries for each document. The translation of every document along the positive z-axis according to its search result depends on the sum of the relevance values returned by the search queries, bringing the most relevant documents to the front and thus to the focus of the user.

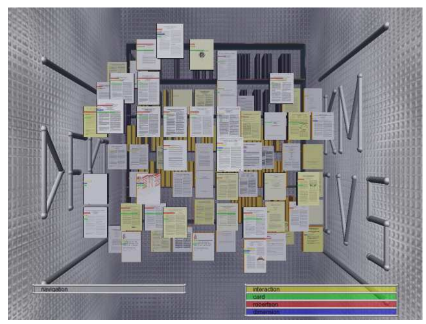

Fig. 10. Search results in PlaneMode

3.6 Document Relation Layout

The three dynamically generated relation types in the general form of "document A is related to document $X, Y, \ldots,$ and Z", which create a directed graph structure, where each document is related to 0 to 25 other documents, are visualized in three different ways—one relation visualization for each relation type—which should be seen as example relation visualizations that could be exchanged or supplemented with other relations.

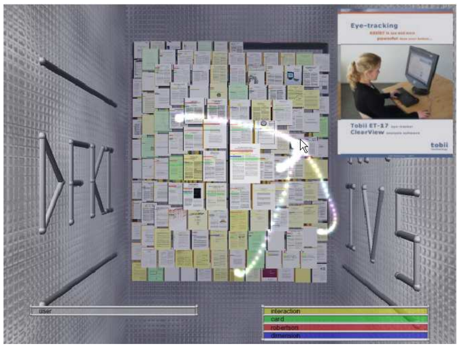

Fig. 11. Relation Types 1 and 2.

The first two relation types are visualized as shown in Fig. 11. Relation Type 1 is intuitively represented by connecting the document under the mouse pointer by curves to the documents to which it is related. These curves might create a mental model like that of *thought flashes* moving the user's attention from the current document to related documents. This mental model is suggested not only by the way the curves are rendered with shining particles, but also by the fact that the curves are animated, starting at the document under the mouse pointer and propagating to the related documents. The advantage of this way of representing relations is that documents are visually strongly connected by curves which moreover create an impressive 3-D structure. The disadvantage of literally connecting documents is that the thought flashes occlude the documents behind them, which makes it sometimes hard to understand which documents the curves are pointing to.

The visualization of Relation Type 2 tries to make up for this disadvantage: the documents which are related to the document under the mouse pointer are simply visualized with semitransparent green boxes around them. Thus, they can be perceived preattentively and do not occlude anything as the boxes are not much larger than the documents themselves. Also, the box representation can be used without complications in PlaneMode and in the bookcase. Compared to Relation 1, the connection of the document under the mouse pointer and the ones in green boxes are, however, not as intuitively clear as in the case of connecting thought flashes. One possible solution to this dilemma might be to combine the visualization of both Relation Types and their advantages in the visualization of a single relation type—i.e., to use thought flashes that end at documents in green boxes.

The visualization and use of Relation Type 3 is different from the previous ones. When enabled, the user has to trigger the display of the relation by clicking on a document. This will display the documents the selected document is related to in the form of ghost documents hovering before the main document space. The user now has the possibility to select one of these ghosts by clicking on it, which will trigger an animation that seems to collapse all of the ghost documents to the real document the selected ghost stands for, and to further visually emphasize this document by displaying a semitransparent red box around it in the moment of collapse. With this visualization technique for relations the user's attention is automatically moved from the selected document to all related documents and finally to the related document that seems to be most useful to him or her. From this document he or she might start the process again to move one step further to related documents.

3.7 Documents at User Defined Locations

The previously described document collection visualization modes are for the artificially computed organization of documents. To give the user the possibility to organize documents according to his or her own organization principles and to remind him or her of previous thoughts, it was decided to also support user-specific structures. Section 3.4 introduced ghost documents as a way of side-by-side use of user structures and computed structures without confusing double representations of documents.

DocuWorld provides two different areas for user-defined document placement: the memory-line in front of all other structure elements which can be seen in Fig. 9, and the area behind the environment walls where ghost documents can be placed on human statues as illustrated in Fig. 12, in which the wall structure is at an intermediate state of the animation of being lifted to an invisible place. The memory-line, which was derived from the idea of a clothes-line, allows the user to hang about twenty important or frequently accessed documents on it. In the area behind the walls can be placed more ghost documents that are less frequently accessed. Thus, DocuWorld provides a hierarchical workspace similar to Card's [2] WebForager presented in Sect. 2.2. The choice of human figures as a storing place was made because human figures, small changes in their appearance, and thus also objects stored on human bodies are easy to recognize and to remember—much easier than remembering the location of documents on a conventional bookcase. The current DocuWorld implementation uses only two human statues at fixed locations. This concept could, however, be extended to add (differently shaped) statues when the user requests it, to drag and drop them to any place behind the environment walls, and to move them to the center of the visualization area and rotate them.

Fig. 12. Statues as storing places

4 Evaluation and Conclusion

In this work existing document collection systems were evaluated and DocuWorld, a prototype for a new immersive 3-D system was developed.

The difficulties of moving from 2-D interfaces to 3-D interfaces were identified and solved: it was explained in detail how the layout of 3-D visualizations should be optimized to make good use of the available 3-D space and the Thought Wizard metaphor was developed as an efficient and flexible global visualization and navigation concept that dynamically moves the visualization items to the user instead of forcing the user to orient him- or herself in a 3-D environment. The need for two types of document areas, one for user-specific document management and one for automatically computed document organizations, was established, memory-line and human statues were developed as ergonomic user storing areas, and as a solution to the problem of double representations of documents, immaterial ghost documents were used.
Compared to the *Web Forager* approach presented in section 2.2 DocuWorld provides the user with a more natural environment, more flexible possibilities to store documents at user defined locations, and the visualization of document semantics through automatic document organization modes and visually indicated document relations — a feature not available in the *Web Forager* approach.
All these aspects, together with the visualization of the document metadata, search results, and the relations between documents, allow preattentive processing of the data, easy and intuitive use, immersion in the animated 3-D environment, and navigation from one document to the other by simply using the Thought Wizard tools or displaying document relations. Although still in a prototype state, DocuWorld is an example realization of a 3-D document collection system that was tested with research lab staff. During these tests the intuitive and animated visualization proved to be easy to use and more enjoyable than a conventional document explorer although there still are many restrictions in the provided functionality. Moreover, the visual representation of document semantics as the relevance to search queries and relations between documents allow the user to quickly perceive relevant documents and to navigate from one document along the displayed relation to a related document. These possibilities of representing document semantics can hardly be provided by traditional 2D text based interfaces, which indicates that a 3D user interface like DocuWorld is the appropriate way of implementing the Semantic Desktop.
To move the presented approach from prototype state to a usable application, additional research and implementation efforts will be necessary. Future work will, for example, have to focus on the replacement of keyboard gestures with more natural, easy-to-remember, and immersive (hand or body) gestures. The 3-D graphics could also be improved to create more convincing and impressive effects. Furthermore, the Gnowsis framework will be integrated in DocuWorld to acquire relations and other semantic metadata. Moreover, supplementary possibilities of information retrieval and knowledge acquisition as well as possibilities to interact with document collections like dragging documents through space or saving interesting configurations and document arrangements for later use will have to be added to create a powerful application that does not restrict the user.

References

1. Rekimoto, J., Green, M.: The information cube: Using transparency in 3d information visualization. In: Proceedings of the Third Annual Workshop on Information Technologies & Systems (WITS'93). (1993) 125–132
2. Card, S.K., Robertson, G.G., York, W.: The webbook and the web forager: an information workspace for the world-wide web. In: CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems, New York, NY, USA, ACM Press (1996) 111–ff.
3. Sebrechts, M.M., Cugini, J.V., Laskowski, S.J., Vasilakis, J., Miller, M.S.: Visualization of search results: a comparative evaluation of text, 2d, and 3d interfaces. In: SIGIR '99: Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval, New York, NY, USA, ACM Press (1999) 3–10
4. Leuski, A., Allan, J.: Lighthouse: Showing the way to relevant information. In: INFOVIS. (2000) 125–130
5. Robertson, G.G., van Dantzich, M., Robbins, D.C., Czerwinski, M., Hinckley, K., Risden, K., Thiel, D., Gorokhovsky, V.: The task gallery: a 3d window manager. In: CHI. (2000) 494–501
6. : (Tactile 3D User Interface: http://tactile3d.com/)
7. Lucene: (http://lucene.apache.org/)
8. Sauermann, L.: The gnowsis — using semantic web technologies to build a semantic desktop. Master's thesis, Technical University of Vienna (2003)
9. Sauermann, L.: The semantic desktop - a basis for personal knowledge management. (2005) 294–301
10. Mine, M.R., Brooks, Jr., F.P., Sequin, C.H.: Moving objects in space: Exploiting proprioception in virtual-environment interaction. Computer Graphics **31** (1997) 19–26