A Cup of Tea & a Piece of Cake
Integration of Virtual Information Workspaces
Inspired by the Way We May Think

Thomas Pederson, March 24 1997
top@medialab.ericsson.se

Abstract

This thesis concerns the design of an information navigation and retrieval system where some navigation mechanisms have inherited functionality from concepts found in cognitive science. These mechanisms are intended to improve navigation and information retrieval in large information spaces like the World Wide Web, as well as enhancing the possibility for creative thinking.

A model describing the nature of different information spaces and how they could be integrated is developed, and a proposal for a 3D user interface based on the model is introduced.

Keywords

Cognitive Tools, Creativity, Information Retrieval, Information Visualization, Information Workspaces, Knowledge Work, Semantic Networks, Virtual Reality, World Wide Web.

Foreword

Our society is in many respects technology-driven. Ever since the industrial revolution, people's way of living has been constantly revised due to technology development. When new technologies emerge, we often try to incorporate them into the existing society architecture as smoothly as possible. One important reason for this is of course that we don't want to re-learn how to live our lives every ten minutes. We need to feel that our earlier experiences in life are useful.

However, sometimes this smooth integration (using old metaphors etc.) leads the development of new technology into an unnecessary narrow non-innovative path, transforming the very promising technology into just a slight improvement of already known technologies. In these cases, it can be fruitful trying to see the new technology from a different perspective. Not to, as usual, stand on top of the mountain called 'Technology Frontline' looking towards the future horizon, but instead climb another mountain, say the mountain called 'Human Thinking', and from *there* study the top of 'Technology Frontline' and the future horizon simultaneously.

This is the viewpoint of this thesis.

Contents

Falling Down

Dalbello & Stefano Dalbello (1995)

like the times that i swore
like the wine that i poured
like a blinding fact
you can't ignore

like the motion
that delivers you forward
'til you can't go back
not anymore

can't dive into an ocean
and not leave the ground
can't trade in your devotion
for what you haven't found
can't go beyond the limit without falling down
can't go beyond the limit without falling down

and it's not what i say
it's alot what i do
'cause you can't go back
not on you word
it's the moment
when you're riding on fear
that you can't look back
you live and learn

can't start at the beginning
when you're in-between
can't know if where you're going's
where you haven't been

can't give away the centre
without spinnin' 'round
can't go beyond the limit
without falling down

like the signs that i feared
and the lines in the mirror
*like a haunting ghost*
that will return

like the moment
when its suddenly all clear
what i needed most
to live and learn

what a lovely day it is for falling down
(mmm..what a lovely day it is for f-f-allin')
what a lovely day it is for falling down
(mmm..what a lovely day it is for f-f-allin')
can't go beyond the limit without falling down
can't go beyond the limit without falling down

can't ruise below the surface
what's not above the skin
can't lose the war on purpose
and not breathe to win

you can't throw away the minutes
you can't go beyond the limit
you can't tell the tale to live it ......
.....without fallin down

and it's not what i say
it's alot what i do
what a lovely day for falling down
what a lovely day it is for falling down
what a lovely day it is for falling down
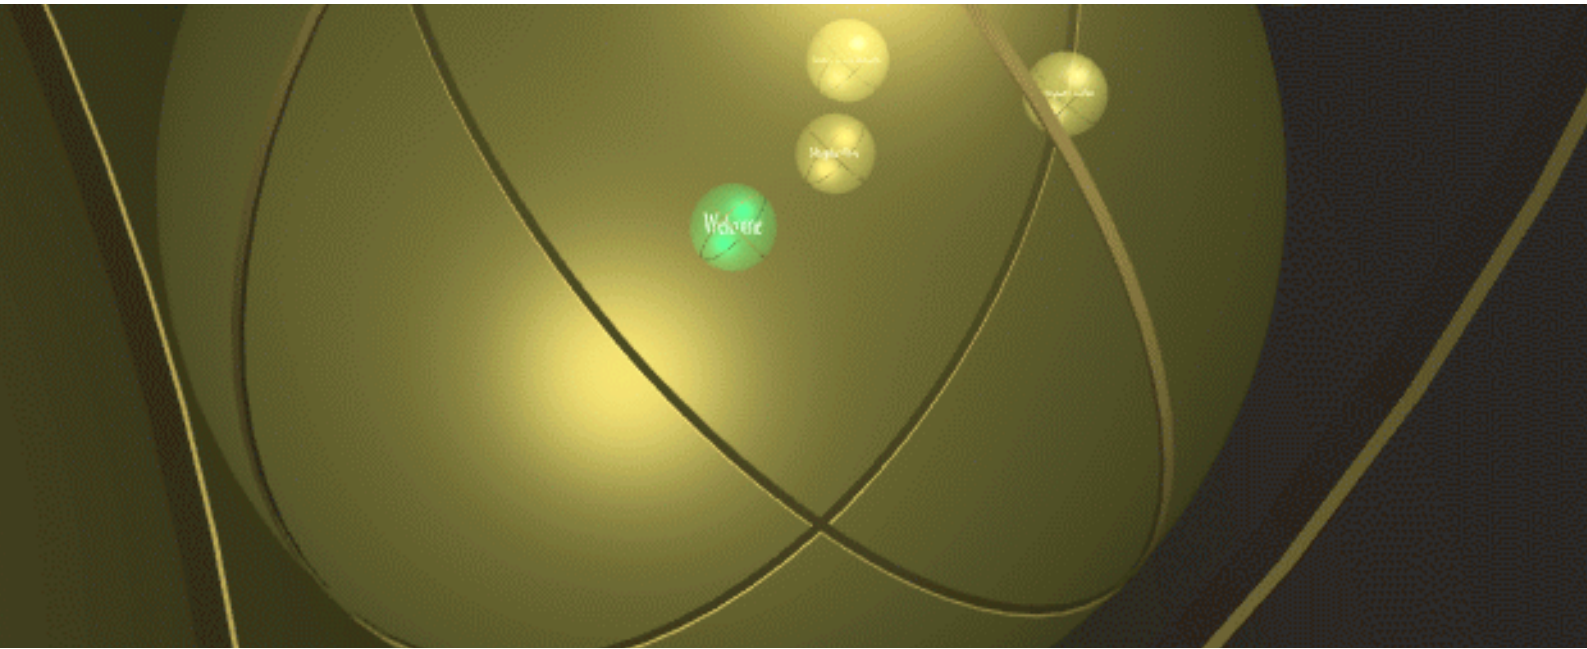
1. Introduction

What It's All About

Information work has become a common occupation in the modern society. The computer has been seen as an important tool when it comes to retrieving and publishing information.

However, as more and more sophisticated artifacts are developed aimed for this purpose, they begin to concern processes that until now have been entirely devoted to humans; processes based on how we think. When we create such artifacts, cognitive tools that operate close to our own minds, it is important to let the human way of thinking have influence on the design. Otherwise, their use will be limited.

This thesis concerns the design of an information navigation and retrieval system where some navigation mechanisms have inherited functionality from concepts found in cognitive science. These mechanisms are intended to improve navigation and information retrieval in large information spaces like the World Wide Web, as well as enhancing the possibility for creative thinking.

The system can be viewed as an enhanced web browser visualizing the web pages in a 3D space, but also making transfer between the shared information space (the WWW) and the private information space (the user's personal computer environment) easier to perform in both directions.

Thesis Sections

The thesis is split into two major parts, one part focusing on theory around information work and visualization, one part explaining practical issues discovered during the project.

Sections 2, 3 and 4 — theory

Experience has shown that finding information in large information spaces, is difficult. The system proposed in this essay tries to overcome some of the problems occurring when interacting with large information spaces by introducing mechanisms based on the nature of human thinking. Therefore, some important concepts present in the process of thought are introduced in section 2, "Background". Since the system is also concerned with information and information transfer in general, as well as the new world-wide media called Internet, these subjects are also discussed briefly in this section.

The proposed system relies heavily on automatically generated information spaces that are based on similarity between information objects created by the users of the system. A model describing the nature of different information spaces and how they could be integrated, is developed in section 3, "Virtual Information Workspaces".

Finally, a proposal for a 3D user interface based on the previous sections is presented in section 4, "Designing a User Interface for Integrated VIWs".

Section 5 and 6 — practice section

A 3D direct-manipulative information tool, GMD's RelevanceSphere, that has inspired the user interface design in section 4 is informally evaluated in section 5, "Implementation". The tool is modified slightly to suit a data set originating from a Music On Demand system developed at Ericsson Media Lab.

In section 6, "About This Project", turns of focus and project goals are discussed in the light of discovered problems.

Section 7, 8 and 9 — summing it up

Conclusions, claims and future visions are proposed in section 7, "Conclusions, Claims and Visions", and the people that have contributed to the project are presented in section 8, "Acknowledgements". Section 9 is a list of references to related litterature.

2. Background

Basic Building Blocks
This section is concerned with some concepts found in cognitive science and informatics. The internet phenomena is also discussed briefly. The purpose is to give the reader some understanding about basic building blocks used in the forthcoming sections.

> *"I had, indeed, plenty to think about. There was general confusion in my mind — things wouldn't fit together. And, although certain impressions stood out, there were too many of them for any one to stand out and give meaning to the whole picture. Things seemed to flash by and vanish at random, and I myself felt reluctant to concentrate on any one thing in particular or try to bring order into the chaos."*
> Fyodor Dostoevsky (1874) [22]

2.1. Modelling the Thinking Process

Abstraction, Association and Context
This essay is concerned with some of the mechanisms present in our thinking process at a mid-abstraction level; avoiding the details of how our brain works at the neuron level, and also avoiding the more philosophical/psychological questions why we want to perform current tasks. This mid-level can be referred to as the cognitive level. Two mechanisms that seem to be important for creative thinking, but also used in seemingly non-creative thinking, are *association* and *abstraction*. Additionally, an important meta-concept is something we can denote *context space*, which affects and is affected by both of the earlier mentioned mechanisms.

Association
Association denotes how new concepts come to mind in a sequence based on previous active (conscious) concepts. A simple but illustrative representation of the nature of association sequence is the propositional network, a model originating from cognitive psychology. In this model, the current active concept 'primes' related concepts according to a process called spreading activation. If the priming value of a certain concept in the network reaches a certain level, it becomes accessible to conscious attention. The view of the association process is often simplified so that only the concept with the highest prime value is considered in each certain point of time, resulting in a single 'trail of association'[10].

Abstraction
Abstraction is our capability of looking on things at a distance, to group similar concepts into sets and to treat these sets as single entities, more or less aware of the simplification done during the generalization phase (often based on induction) along the way. The reverse process, to apply general, abstract knowledge on a specific concrete task/concept, is of course equally important. A common representation scheme which relies heavily on abstraction mechanisms as well as the associ-

ative properties of the propositional network is the so-called semantic network, originally developed by psychologists and used heavily by researchers in the classic Artificial Intelligence (AI) discipline.

Context space
Context space is an imaginary space containing all more or less primed concepts at a certain point of time. Concepts outside this space are no candidates to become the next concept in the trail of association. The context space is the result of the unconscious spreading activation processes based on the relational weights between concepts. These weights are based on long term (earlier experiences in situations similar to the current) or short term (perception of present environmental conditions) experiences.

> *"A third criticism of task analysis has been that many methods tend to view the user solely as a goal directed or top down problem solver. However when considering IVAs[1] it seems more appropriate to view the user as an opportunistic planner (Green 1989) who uses information from the external environment to mediate cognition."*
> Lisa Tweedie (1995) [81]

New Tasks Need New Models
Within the field of Human-Computer Interaction (HCI), quite a few models of human goal building, during performance of certain tasks, have been developed and used. Some of them have been heavily influenced by AI research. These 'rigid' models are well-suited for studying humans performing traditional routine-based tasks with computers but do seldom comply to the situations where the task is more of a heuristic nature, or rather, if the user *chooses* to attack the task at hand in a heuristic fashion. Today's user interfaces (UIs) do in many cases allow the user to solve task subgoals in only one way (using a software designer's predefined tools in a predefined sequence), which makes the heuristic approach hard to perform. Thus, if we continue to make rigid UIs we can also continue to use these AI-inspired HCI models since the heuristic approach isn't an option![2]

A positive indication from the software industry, regarding the UI flexibility needed for heuristic tasks, is the rising interest for smaller applications like Sun JavaBeans or Microsoft ActiveX components. This fragmentation of software will hopefully make it easier for users to design their own virtual workspace[3], just like they choose tools in the real world. I would like to point out that I don't mean that user's should be forced to

---

1. "Interactive Visualization Artifacts (IVAs) are complex applications that allow users to manipulate, encode and organise data graphically." [81].
2. This argumentation is not entirely true since comprehensive HCI models have to cope with the fact that we in many situations choose heuristic approaches within the task but outside the interaction with the computer.
3. A virtual workspace is a mix of digitally stored information objects and tools that the user can use to manipulate these objects. The desktop environment present in many personal computers today, containing file folders and word processor software, is an example of a virtual workspace.

become their own software designers. What I mean is that they will be able to choose freely among a lot of different smaller tools and that they will be able to use these tools in whatever environment they like.

I believe that computers will increasingly be used as a tool in the creative process, irrespective of task domain[60]. This will make it necessary to develop HCI models that take into account the creative process. In any case, applications designed to be used in tasks that are most probably attacked using heuristic approaches (such as most kinds of knowledge work) should have the human way of thinking/problem solving as a basic role model of design. In this way, the system will in a cost-effective manner minimize the conceptual bridges[59] needed in the interaction between man and machine. Examples of HCI theories aimed for heuristic tasks are the Flower & Hayes model for the writing process[26] and Bødker's activity approach[9].

2.2. Communication

Before we study how information is exchanged using computers, a more general presentation of the concept of communication and information transfer might be in its place.

Communication is Important
Communication can be viewed as social interaction between humans. It is the schema for how thoughts and knowledge is shared. There are many reasons why we communicate. Among them are that we want to gain understanding of other persons views in certain areas of interests and to learn new things. In theory, if the participants in a communication act all have the same background and discuss a to all parties well-known area, little or nothing will result (if we choose to over-look the social aspects). In practice, all humans do of course have different backgrounds, so communication is in this respect always rewarding. It is in any case obvious that we gain a lot by communicating with people from different backgrounds. Internet is a new media through which this kind of communication is made more easily, since this media is geographically independent.

Direct and indirect communication
Communication can either be direct or indirect. Direct communication denotes the situation where information is shared in real time, i.e. speech, physical interaction etc. Indirect communication is when the thoughts are transformed into some kind of time-persistent media at one point in time and shared/interpreted by others at a later stage, like in literature, movies etc. In this paper we will focus on indirect communication.

2.2.1. Information and Information Transfer

The Difference Between Information and Data
Information is quite a complex word with different meanings in different contexts [43]. To clarify how the word is to be interpreted in this essay, I will use a phenomenological definition: Consider the number 30. If this number is presented in a more or less empty context, like in the preceding sentence, the number con-

tains little or no information. It does contain some information because it is presented in this thesis and below a certain header that can give some clues of in what purpose the number is presented. Also, to a mathematician, the number belongs to the set of natural numbers and can be manipulated using certain algebraic rules. However, it is very hard to figure out what the number is supposed to describe right here and right now, and it does therefore contain little or no information.

By saying that the above mentioned number actually is the number of pages of this document, it suddenly contains a lot more information. Generally speaking, we have connected the information-poor number to a semantic attribute which has given it a meaning. Before this fabulous connection, the number could be described as a data item, afterwards it has become a piece of information.

Information at one abstraction level might be interpreted as data on another level, all depending on context. This make things more complex, but as long as one knows at which level one is, no harm is done. The above definition will in any case do for the purpose of this essay.

> *"Data is... probably the only thing humans have in common with computers"*
> D. Benyon (1992) [6]

Computers and Information
Computers are brilliant data-processors. To make their power of handling data useful to humans, we apply certain semantics to the data by presenting the data in software environments providing context clues how to interpret and understand the data, making it easier for us to solve our tasks at hand. It should be the software developer's main goal to design these environments so that the underlying semantics provided by the user interface match the user's cognitive model of the task.

The science
Computer science is the science of processing data using computers. Information science traditionally address the process of sharing information; communication. If we combine the two, we end up in a scientific discipline called informatics [41]. Most of the discussion in this essay is related to the domain of informatics.

Information Transfer
Information can be viewed as a media through which cognitive structures (knowledge) is transferred between humans in the communication process. As illustrated in figure 1, this process can be performed at different levels of abstraction.

Since very few human beings are equipped with telepathic skills, communication is usually done through one of the lower abstraction levels i.e. using speech in free air. The original individual cognitive structures are transformed to one of the lower-level forms understandable by other humans. During this transformation, certain loss of information quality is present —sometimes it is very hard to express a feeling or an intention (cognitive structures) in for instance natural language. When

using computers in the communication process, it is the computer system designer's task to design the system in such a way that this kind of information quality loss is minimal. The new-born scientific area called Information Visualization has a potential to make it possible to share information on a higher abstraction level than for instance natural language, just like a picture can say more than a thousand words.

Many more agents with different backgrounds
Information transfer between humans using computers as media is affected by semantic and articulatory distances [59] just like any other human-computer interaction. Information is often regarded as being 'built-in' into the computer system (by a software designer) which however is a too rough simplification when discussing information transfer. The added complexity lies in the semantic and articulatory distances between the author's cognitive model of the information to be transferred and the computer systems ability to support that model. This is because software designers (or groups of them, working on the same project) can be regarded as one single agent whereas the often much larger group of authors involved in a general computer-based information transfer system must be viewed as individual agents or groups of agents having different cognitive models of the information, different motives for supplying the information and so forth. Also, the user group (information retrievers) is much more heterogene than the classic 'user' groups for whom for instance word processor software is designed.

I would like to clarify that the above discussion addresses the *communication problem* of having many information providers and retrievers on the same information sharing system, rather than the problem of designing a system supporting it, although that particular task is of course also affected.

2.2.2. Internet

The New Media
Tim Berners-Lee's invention of the World Wide Web and the implementation of the first public web browser Mosaic in the early 1990's has caused the amount of Internet users to grow enormously. Soon, Internet will be as common as the telephone or the TV in the Western world. Therefore, Internet will be one of the most important media there are.

For a media to be successful it has to come up with new features which the 'old' ones doesn't provide, and surely, Internet can. Among the features available today are the possibilities to send low-cost fast-delivered letters all over the world, to publish information to the whole Internet-community as easy or even easier than in the local press, and to get in touch with people who share common interests regardless of geographical location.

As can be noted from the above examples, one significant feature of this new media is the absence of geographical boundaries. Another important feature is the
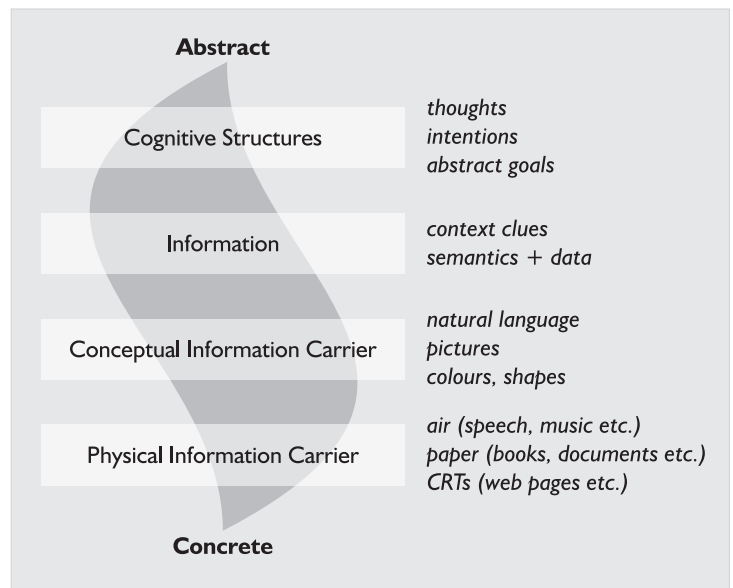

Figure 1: Different levels of abstraction in information transfer.

generality. Internet can (or will) be used to share information in many different forms: text, pictures, speech, music and movies. In theory, it's hard to find 'old' media that Internet cannot cope with.

The limitation is related to computing power and network bandwidth but nothing indicates that the enormous development speed in these areas are losing strength. And since Internet has become an important economical resource, even more effort will be put into research and development of hardware and software to support the demand.

3. Virtual Information Work-spaces

Getting a Grip

Where is information located? In what context are information objects found? Who decides? Who designs? This section is dedicated to find answers to these questions. Different virtual information spaces are defined and their content and purpose are discussed. The goal is to get a clearer view of the conceptual similarities and differences between the spaces so that an integrated system dealing with all of them can be designed more easily (in the next section, "Designing a User Interface for Integrated VIWs").

Although discussed in general terms, the spaces are highly influenced by the common current information work situation, consisting of a personal computer connected to the WWW and a user wanting to do some really creative work.

In the end of the section, weaknesses in today's knowledge work situation are discussed.

The General Idea

Information work is cyclic

Traditional information retrieval models doesn't take into consideration all the different processes involved in the knowledge-worker's task [66][20]. The information retriever is considered to retrieve the desired information and then use it *outside* the IR-system. Whatever happens after the retrieval is out of scope of the model. This model has its historical reasons, but the way information is retrieved today and the way it is worked with (or will be) makes the model inappropriate when studying knowledge-workers.

Nowadays, many knowledge-workers use their own personal computers in the retrieval process, continues to work with the retrieved information within more or less the same environment, and when the task or a specific part of the task is finished, usually intends to re-enter the new self-built piece of information into the public information store so that result can be shared with others. And then it starts all over again. The idea of information workspaces tries to incorporate this cyclic way of processing information.

Abstract spaces

Figure 2 is a rough sketch of the information work process viewed as a chain of cycles. The spaces shown should, unlike the spaces discussed earlier in this section, be viewed as abstract situation-based rooms, existing in our minds, and not as physical or virtual visualizations. All process cycles depend on various tools to make the processes feasible. In a highly computerized environment, such as the one we are facing, these tools should, to the greatest extent possible be provided within the same computerized environment to avoid unnecessary cognitive overhead similar to the overhead present in the inter-world leaps discussed on page 9.

In the following discussion, the word 'space' should be interpreted more concrete, with physical attributes. It
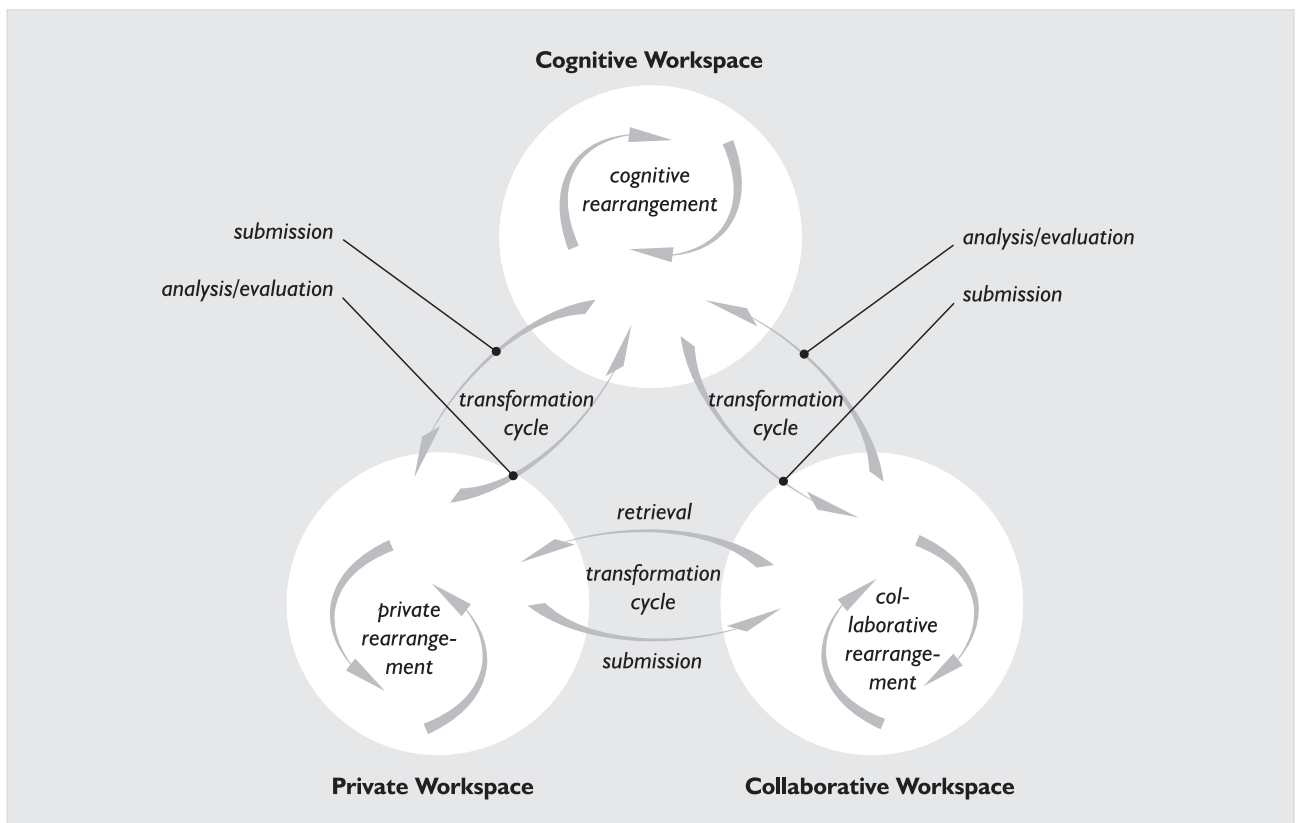

Figure 2: The information work process viewed as a chain of cycles. The three spaces are abstract and do not have to correspond to any physical equivalences.

could, however, be good to keep the abstract spaces shown in figure 2 in mind.

Virtual information tools are powerful
It is hard or even impossible to create as good information tools in the real world as in the virtual world. To illustrate this, try replacing all occurrences of the word 'information' with 'cognitive food' in this paper: If you have the physical version that is printed on the media made of trees it will be a time consuming task giving quite a sloppy result. If you instead have the 'virtual' version in your word processor, it will take you no-time and the result will be fully acceptable. (Some people would of course remark on the strange choice of words in certain parts of the document but that is another story.)

The information we need is increasingly stored in digital form regardless of physical location —maybe it's on your local hard disk, maybe it's somewhere else, but it's almost certainly digital. Conclusion: information that is to be used more than being just perceived (this, I would argue, is viable for most kinds of information use, except maybe when interpreting certain kinds of art) gains quality by just *being* in virtual form, providing appropriate virtual information tools are present. These tools are embedded in the virtual information workspace (VIW) concept.

Avoiding the inter-world leaps
Instead of transforming the information from the virtual store to the real world, performing the information work, and re-transform the result into the virtual world, why not put the entire workspace itself into the virtual world? By doing this, we can benefit from the power of computers as data-processors without having to do the quality-degrading transformations between the two worlds. This is essentially the idea of VIWs.

Summary
A VIW is a space containing intrinsically related information objects and tools making it possible to alter these objects and relations in various ways (changing location, size, content etc.).

3.1. A Simple VIW Model

VIWs (further on called spaces) can be divided into two subgroups, user-designed and system-designed spaces. As the names suggests, the former denotes spaces that are more or less hand-crafted by users using direct manipulation or more sophisticated tools, the latter kind of spaces are more or less automatically generated by computer systems. With space design I mean the design of *relations* between information objects in the space, not necessarily the design of the objects themselves. This section discusses the nature of different spaces and their different tasks in the knowledge-worker's Overall working process.

3.1.1. User-Designed VIWs

User-designed spaces can in turn be divided into two subspaces, private and public. The private space is accessible by the designer/owner only, whereas the pub-
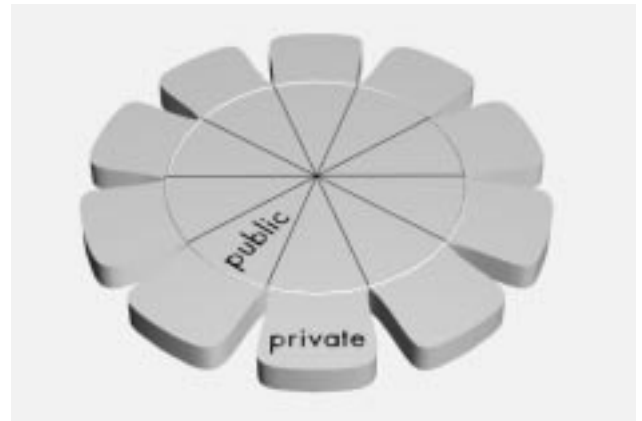

Figure 3: User-Designed Virtual Information Workspaces.

lic space is accessible by all other users in the shared space. The shared space is simply the union of all public parts of the users' self-designed spaces. Figure 3 illustrates the user-designed spaces.

Private User-Designed VIWs — Mental Support
An easy-accessed external storage for information objects involved in information based tasks is an invaluable tool in solving the task by helping the knowledge-worker to build an internal cognitive representation, unveiling relations between the objects, and their intrinsic relative importance.

When building cognitive models, which we do all the time but most evidently when we learn things or when we create new things, it is important to have immediate access to all relevant information pieces, hereafter called objects. If the number of objects is too small, the model will be built on too few premises eventually leading to a distorted and unnecessary narrow view. If the number of objects is too large, we have difficulties in creating a clear structure, a complete sound impression of the relevant objects, and therefore of the task itself. A phenomena generally called information overload[84]. If the retrieval of relevant objects involves a high cognitive overhead (in other words, the objects are hard to retrieve and/or to understand) it will affect the model-building process by causing continuous task breakdown situations [9].

(It is my belief that) private spaces can extend the amount of objects accessible during the model-building process [60]. This space can be viewed as an external self-designed 'cache'-memory between larger information sources (like Internet or the world in general[1]) and the knowledge-worker, enhanced with tools to support the information transformations involved in the specific task.

Today's personal computers, loaded with virtual information tools like word processor software and drawing applications, serve as a limited private user-designed VIW. However, this virtual environment does seldom provide sufficient support for organizing the different information objects to show relations between

---

1. The private workspace-application Merz[47][48][49] has in fact borrowed it's name from the German artist Kurt Schwitters who incorporated bits and pieces found in the surrounding world in his art.

them (often they can be placed freely on the 'desktop' but there can only be one such 'semantic map'). Also, most information tools do not operate within the desktop. Instead, the information object is transferred into a different application area, enhancing cognitive overhead. The component technology mentioned in section 2 is a positive trend, potentially moving the tools closer to the information space (the desktop).

Public User-Designed VIWs — Publishing Area
Most work would be of little use if just kept to oneself. Public spaces serve as publishing forum for work done in private spaces. It is closely coupled to the private space, making it easy to move information structures from one to the other. The designer of a public space can decide whether to allow other users to alter the space or if it should be fixed, i.e. only viewable. A World Wide Web (WWW) site is a good representative for a fixed public space, though — until recently[1] — the workspace tools for altering WWW sites is not available within the space itself.

Shared User-Designed VIWs — Platform for Collaboration
The shared space is the union of all user-designed public spaces. It is an environment where all users have a nearly identical view of all the public spaces designed by all other users. If one person alters some part of the space, using one of the tools provided, the result is immediately evident to the other users. The users themselves can be either visible, i.e. represented by so called avatars, or invisible. The main goal for a shared space is to support collaborative information work. Public spaces, spanning up the shared space, alterable by other users than the original designer vouch for direct collaborative work. The WWW is a good representative for a (in most cases non-direct collaborative) shared space[2].

3.1.2. System-Designed VIWs

Shared System-Designed VIWs — Automated Relations
Based on shared user-designed spaces, algorithms can generate a space where automatically generated relations between objects in different user-designed spaces appear. Heuristics for this automated process involve some kind of content analysis of objects and general, non-user-specific, knowledge databases. The resulting automated space highly depends on the information quality of single objects in the user-designed space. Information quality is in this context equal to how easy the algorithms can extract a suitable corpus for the object at hand. A corpus is a system-internal representa-

tion of a given object, used for measuring degree of similarity between objects.

Meta information
The use of so-called meta information tags (designer-added standardized labels describing objects or parts of objects) enhance information quality, making the automated space generation easier. However, this meta-tagging puts a burden on the designer. There is a limit to how large the set of standardized tags can be — at a certain point, space designer will lose track of available tags potentially leading to the decision not to use them at all. Also, it is easier for algorithms to cope with few well-defined tags, than the opposite which at an extreme leads to the 'tag set' called natural language. Thus, since the amount of available tags has to be restricted, non-meta based content-analysis is in any case useful as complementary mechanism.

Today and Tomorrow
On the web today, shared system-designed information spaces are commonly represented by results returned from Internet search engines. These 'spaces' are primitive both regarding the quality of the contained information objects, and the way semantic relations between them are visualized. However, work is going on to improve the search engines from at least two fronts. Front 1: the change in database modelling architecture from being boolean-based 'exact match' to more fuzzy-logic orientated, probabilistic 'best match' driven. Front 2: Visualizing the results using two or three-dimensional visualization techniques based on spatial measurements. Frontier projects worth mentioning are WebSOM[37], SemioMap[71], Knowledge Browser[20] and Visual SiteMap[51].

As an example, WebSOM uses Kohonen's Artificial Neural Network mechanism called Self-Organizable Map to generate a 2D space where similar objects and object clusters automatically are positioned near each other. The user is allowed to 'drill-down' in a third dimension to inspect desired locations more closely.

Private System-Designed VIWs — User-Based Automated Relations
The public system-designed space, discussed in the previous paragraphs, can in turn be used as base for a more personal view of the entire shared space, where knowledge about a specific user and task implicitly is allowed to affect the appearance of the space. This knowledge can for instance be based on long and short term history of the user's actions as well as the task at hand.

User agents
The programs performing this alteration of information space, doing their work in the background, are often called user agents. Each agent uses a specific part of the total knowledge base that model user activities, to perform it's specific task. It can involve highlighting potentially important objects, excluding unimportant objects from the space or for instance grouping objects into clusters according to some criteria. The alteration is presumed to guide the user by putting interesting objects close to the current view of the space and uninteresting

---

1. Some recent WWW-browsers like Netscape's Navigator Gold 3.0 and W3C's Arena allow users to alter web pages within the browser environment. Arena even allow users to edit webpages located on foreign sites. The result is however still saved locally.
2. However, the transformation that has to be performed when moving an information structure from the private space (constituted by the office, the desktop, the personal computer etc.) to the public (and shared) WWW space is sometimes quite complex.
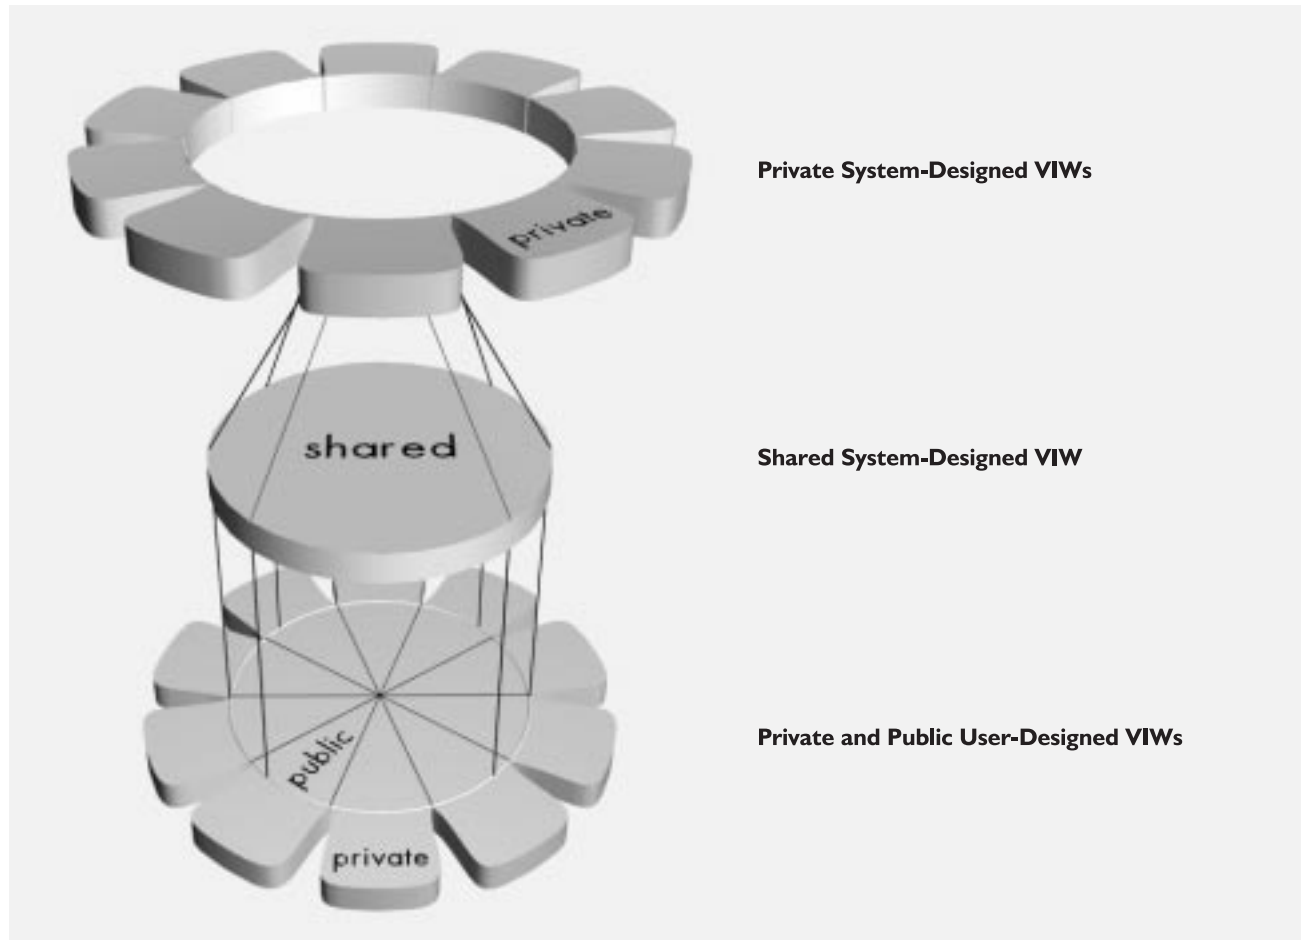
Figure 4: User-Designed and System-Designed Virtual Information Workspaces.

objects further away. This kind of space distortion goes under the name of fisheye-view[25][24][65] and has become an important mechanism when visualizing large information sets within the area of Information Visualization. Note: fisheye-views can of course be generated by other computer mechanisms than user-agents.

It is important to differ this kind of implicit space alteration from the explicit that is performed by the user directly. If we choose to view the entire knowledge-work system from a cognitive perspective, these agents can be seen as unconscious cognitive processes that prime certain objects of presumed interest. The explicit alteration performed by the user is done consciously, i.e. in the foreground.

3.1.3. Completing the Model

Figure 3 shows the two system-designed spaces added to the user-designed space model shown earlier. This figure will serve as a framework for forthcoming discussions about space integration.

The private user-designed spaces is the set of unconnected leaves surrounding the bottom disc. Each leaf corresponds to a specific user's private space. The public user-designed spaces are pieces of the cake called shared user-designed space, illustrated by the bottom disc. This space is projected onto another disc using information analysis algorithms generating the shared system-

designed space. In turn, this space can be transformed into the private system-designed space by user agents. Like the private user-designed space, this space is not shared and therefore placed outside the shared area, above the specific user's self-designed private space.

3.2. Objectivity and Subjectivity

"We suspect that the need to differentiate between objectivity and subjectivity in VR systems will increase as the number of truly co-operative VR application grows."
Benford et al. (1995) [5]

Shared VIWs give the user a kind of objective view of information since the design is done in a collaborative fashion. The view can be defined implicitly (system-designed) or explicitly by a group of users (shared user-designed space constituted of a set of public user-designed spaces).

Private VIWs on the other hand, allow the users to illustrate her own partial, personal (subjective) view of a shared information space.

3.3. Space Integration

All information spaces discussed above exist today or will exist in the near future. However, their shapes and interaction schemes differ widely, making inter-space work hard to perform. Wouldn't it be nice if we had a system that could offer access to all spaces through one single user interface using as far as possible identical interaction schemes regardless of space? Such a system would not only make the inter-space work easier for knowledge workers, but also most probably reveal new interesting functionalities never considered before.

Integrating software environments is a tough task since the environments tend to constantly diverge to become better at their respective task (which is a good reason). However, although a full integration might be impossible, it is certainly an ideal worth striving for. If we could at least come half-way to an integrated information work environment, a lot of boring irrelevant tasks we are concerned with today would become unnecessary.

3.4. Weaknesses in Today's Knowledge Work Situation

In a sense, we already have access to most spaces through a single user interface, namely our virtual PC desktop. We can access information tools (word processors etc.), private spaces (the desktop, file containers etc.), shared spaces (through different kind of browsers we can access the WWW, Internet and Intranets in general). Primitive shared system-designed spaces we get by using Internet search engines within browsers. Software for the generation of private system-designed spaces is also available in the form of user agents returning URLs of potential interest. All this on our own single PC.

However, from the view of the modern knowledge worker, whose task involves the study of many different spaces and information transfer between them, this arrangement is far from optimal. A few weaknesses and potential enhancements are pointed out below (in no particular order).

Incoherence
Tools performing the same function in different spaces have different shapes and different interaction schemes. This is a general HCI problem that is showing up continuously. A great deal of work has been done on this though. First, incoherence was attacked on a single-application level, so that tools within applications were uniformly designed and didn't mean different things in different application modes. If the letter 'S' means 'Start' in one mode, It shouldn't mean 'Stop' in another. Then, exemplified by the strict developing guidelines delivered with the Apple Macintosh in the early eighties, the incoherence matter was brought up in the light on an inter-application level.

The question is, when are we going to consider the tool incoherence that exists between information spaces, i.e. between PC desktop tools and web tools?

Large Conceptual Distances Between Spaces
The nature of the different spaces are based on different metaphors. Every space uses one single or a combination of metaphors, inherited from old environments (like the PC desktop metaphor) or created more or less without reference (like the WWW). No matter what metaphor that is chosen, it is argued to be good when working in that particular space. However, little or no consideration is attached to the fact that the space is used in combination with other spaces in the real-world knowledge work situation. Within the desktop environment of our PCs, the different applications use a standard set of buttons and standard interaction schemes when we want to perform standard tasks, but why does not the WWW operate in the same way as our virtual desktops or why is not my desktop a WWW-browser?

The difference between the spaces can be explained by the fact that the spaces were born at different times, built on different technologies and for different purposes. However, when talking about knowledge work, I don't think that the last argument is very strong since they are all essentially VIW user interfaces. And the technology-part (networks vs. local hardware) is something that can be disguised from the user. The WWW interface is actually a quite impressive disguise of the complex technology it relies on, and that is probably one important reason for it's success.

Space-Dependent Information Tools using Space-Dependent Object Syntax
Why is the WWW version of this document a HTML file and my local word processor file a word processor-kind of file? Or rather, why should I have to bother? They do both contain almost exactly the same information, still I need to consider both file types because I want the essay to be accessible in both my personal space (the virtual desktop in my PC) and the shared space (WWW). The conversion is neither user-transparent nor automatic.

Wouldn't it be great if I could bring my favourite information tool, say my word processor, into any space I like and to alter information objects on location? The same tool should work in my private space *and* my self-designed public space *and* in the transformed public (the system-designed) spaces.

Little Use of Spatiality
Humans come with built-in unconscious talents of which measuring things spatially is one. Every unconscious talent that can be used in HCI is fruitful, because it releases more conscious cognitive power useful for other task related things. The use of spatiality is not used very often in today's user interfaces, even though it might be appropriate in many cases. The movement from text-based interaction to graphically based interaction has made this concept more used though.

Showing similarity between objects using spatial positions (objects closely related are spatially close to each other and vice versa) is a good way of making use of spatiality. This is also one important source of inspiration for the design of the user interface introduced in the next section.

Little Support for Creative Thinking
Knowledge work involves many different kinds of information processing like creating, organizing, publishing, relating, retrieving, sharing, transferring, transforming and viewing information. Many of these processes could more or less be supported by computer systems.

Supporting perspective shifts
A tempting process to support is the creative process. An important part of creative thinking (and learning) is to try to see objects from different perspectives, in different contexts. Since information objects tend to be digital and located within a shared space, supporting these perspective-shifts could be done by allowing users to switch between similar objects located in different user- or system-designed spaces, i.e. to change viewing context of a specific object.

This kind of mechanism plays an important role in the system designed in the next section.

Primitive Temporary System-Designed Spaces
The procedure of finding related information in the shared space called WWW is either relying on user-designed collections of (links to) related objects or highly temporary system-designed spaces created by search engines. The interaction paradigm used together with most search engines (text-based query) differ from the way we otherwise navigate in the space (mouse point and click).

Why is not search engine queries by default set to look for objects similar to the object the user is currently studying? This would make it possible to create a non-temporary system-designed space based on similarity between objects, possible to navigate through using the same navigation paradigm as we use navigating user-designed spaces; mouse point and click. Of course, text-based queries would still be a good complement if we want to do long-distance jumps.

This kind of 'similarity-browsing' is a key feature of the system designed in the next section.

4. Designing a User Interface for Integrated VIWs

Disclaimer

Based on the VIW-model and on the criticism of today's knowledge work situation discussed in section 3, a user interface proposal is presented. The ideas are not well thought-out so this section should just be interpreted for what it is, a proposal based on ideas. Details have to be worked out before a UI like the one proposed is actually implemented.

The UI focuses on navigation, information retrieval and space integration. Agent technologies and space-independent information tools are not considered.

The section starts with an introduction to the space principles of the UI and ends with a scenario showing the system in action.

4.1. Space Principles

The Information Space

The information space is materially finite (it contains a finite, potentially very large, set of information objects) but spatially infinite. It *cannot* be described as a simple three-dimensional space where every physical location is unique and can be entered only by first traversing one member of a fixed set surrounding locations. Instead, the number of surrounding locations is dynamic, defined by non-physical contextual attributes. These attributes can be of a temporal short-term nature, like a history of previous visited information objects, or more long-term related, defined by single users or user groups designing their own surrounding locations to reflect their particular view.

Combining three dimensions with n dimensions

The main reason why the space becomes so complex is that it inherits its characteristics not only from the concrete three-dimensional world around us, but also from the nature of information, the nature of knowledge and the way the human mind processes the second last concept to obtain the latter. Note that the technology making it all possible *not* necessarily has to affect the UI design.

The space can be considered to be three-dimensional when viewed from a certain location and n-dimensional when it is explored by navigation or viewed (imaginary) as whole. Figure 5 illustrates the different sources affecting the design of the information space, along with their contributions.

Atomic and non-atomic objects

An objects space, in this UI visualized as a sphere, is like a webpage. However, there are no restrictions for contents of information objects. They might for instance, sooner or later in the user browsing chain, contain themselves. An object containing only directly interpretable information, for instance text without references to other information objects, are said to be atomic. All other objects are non-atomic. Put more formally, using context-free grammar:

$$\langle object \rangle \rightarrow \langle atomic\ object \rangle \langle object \rangle$$
$$\langle object \rangle \rightarrow \langle atomic\ object \rangle$$

Navigation

Exploration of the information space is done by navigation in three dimensions. Support to explore the higher-dimensioned global space is partially realized implicitly, by navigation *through* objects (spheres) opening up new spaces, partially explicitly by selecting a different *perspective* from a dynamic set of possible perspectives. The latter operation is performed by jumping between different kind of spaces and is not incorporated in the three-dimensional space metaphor.

The object *itself* is the link between knowledge you have (information you understand[1]) and knowledge you don't have (information never encountered before).

Information object positions

The links are themselves information objects (spheres) positioned somewhere within the current space. Physical positions of information objects are determined by the space designer, a human (user-designed spaces) or the system (system-designed spaces). Alteration of object positions during browsing can be performed by the user. Often these alterations are only temporal,

---

1. In a sense, you are presumed to understand the information object since you have found it in a self-created context using navigation.
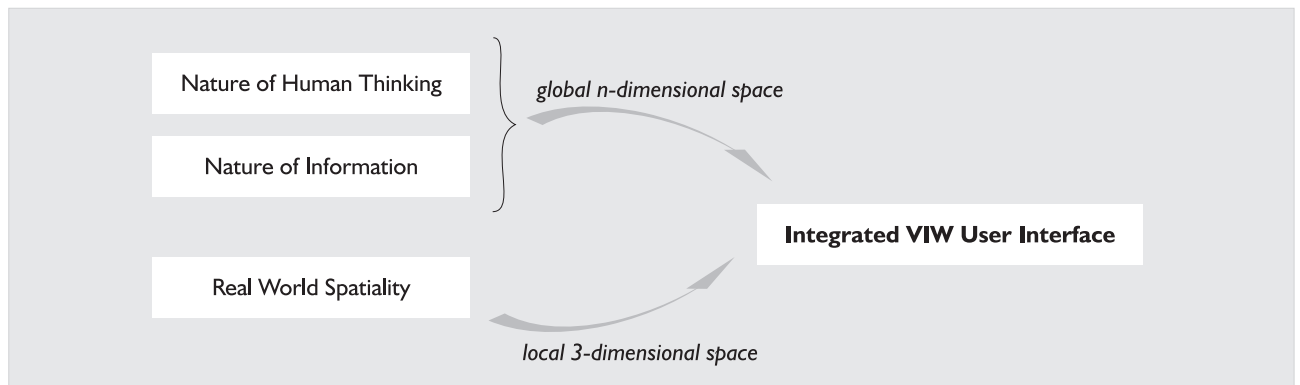

Figure 5: Different sources affecting the design of the information space.

helping the user in his search for interesting information. However, the user can also make static changes to object positions, for instance within his private space.

The Link to Human Thinking
Although being a simplified model of human thought processes, the propositional model (see "Modelling the Thinking Process" on page 5) could serve as a connection surface between Cognitive Science, Information Retrieval and Information Visualization.

Navigating in a virtual information space can be viewed as a high-level mind process where the different spatially visualized information objects are more or less primed concepts within a propositional network. It is the systems task to prime the right information objects (Information Retrieval) and create the context space (Information Visualization) so that the user's current information-need is satisfied. Human thought processes like abstraction and association have their corresponding computer mechanisms in automatic clustering of information objects and links between these clusters.

4.2. User Interface Scenario

4.2.1. Background

The Vannevar Bush Scenario
I use Vannevar Bush's Memex scenario presented in the classical essay "As We May Think" [10] as a base for showing the interaction between user and system.

> "The owner of the memex, let us say, is interested in the origin and properties of the bow and arrow. Specifically he is studying why the short Turkish bow was apparently superior to the English long bow in the skirmishes of the Crusades. He has dozens of possibly pertinent books and articles in his memex.
>
> First he runs through an encyclopedia, finds an interesting but sketchy article, leaves it projected. Next, in a history, he finds another pertinent item, and ties the two together. Thus he goes, building a trail of many items. Occasionally he inserts a comment of his own, either linking it into the main trail or joining it by a side trail to a particular item. When it becomes evident that the elastic properties of available materials had a great deal to do with the bow, he branches off on a side trail which takes him through textbooks on elasticity and tables of physical constants. He inserts a page of longhand analysis of his own. Thus he builds a trail of his interest through the maze of materials available to him."
>
> Vannevar Bush (1945) [10]

4.2.2. Scenario

The scenario illustrations are snapshots from an animation (available on the WWW) designed by Dinka Pignon. Because of time limits, this animation covers only the first half of the scenario. The latter parts of the scenario is therefore not illustrated.

Note: Before reading this scenario, a quick look on the presentation of the tool called RelevanceSphere on page 20 is recommended.

The windows
The user interface consists of a few open windows containing different information workspaces, a collection of iconified workspaces and a toolbar invoked by a click on the right interaction device button. In the centre of the shared space, all information objects (user-designed spaces) in the space are located.
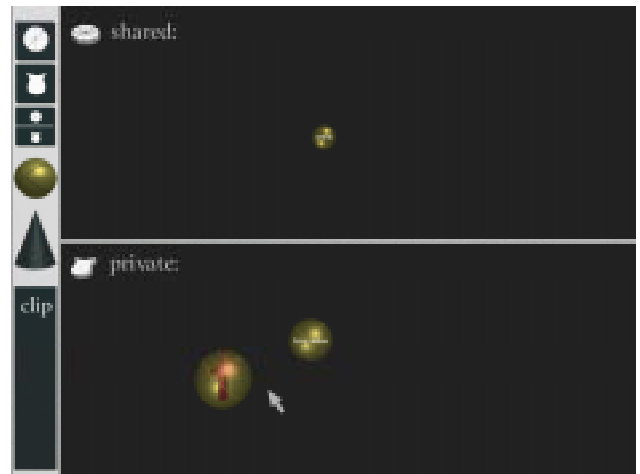

Creating the initial space
The user, let's say John Doe, drags an empty Relevance-Sphere from the toolbar to the shared space window. The sphere centres itself around the collection of information objects.
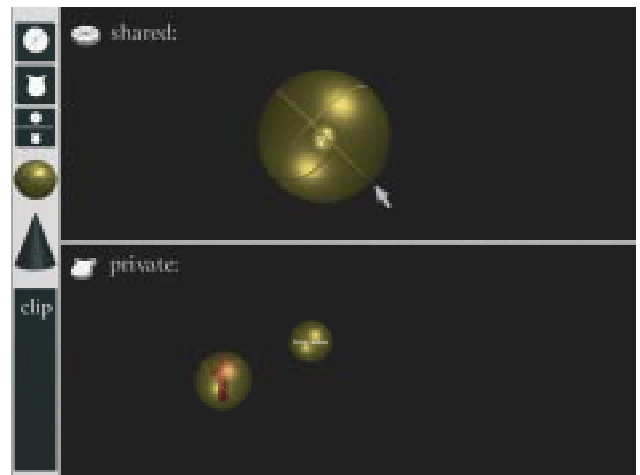

John drags two term spheres from the toolbar into the shared space. They are automatically attached to the surface of the RelevanceSphere. He moves them onto desired positions and assigns attributes by clicking on them and typing on his keyboard. One of the terms is assigned the name "Turkish bow" and the other "English long bow". As soon as he enters the attributes, some

of the objects located in the centre of the Relevance-Sphere (and the shared space) move closer to the term spheres. These are the information objects (user-designed spaces) that the system finds to be most relevant to the terms John has specified.
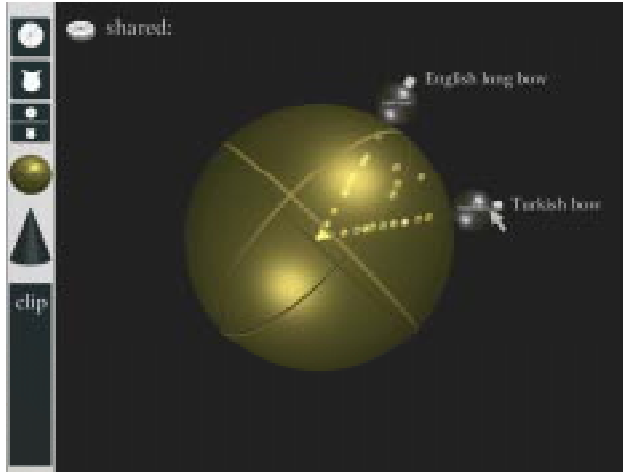

The first glance

John grabs the handle of his 3D-interaction device and starts to explore the shared information space. At a distance, each information object looks like a sphere containing a conventionalized 2D or 3D physical object symbolizing the real content of the sphere. These conventionalized objects are put there by the designers of the information objects themselves. Possible symbols would be for instance a book cover if it contains the texts of a book, a record cover if the object contains music from a certain record, or if the object constitutes the 'homepage' of some person it might be a picture showing that persons face.
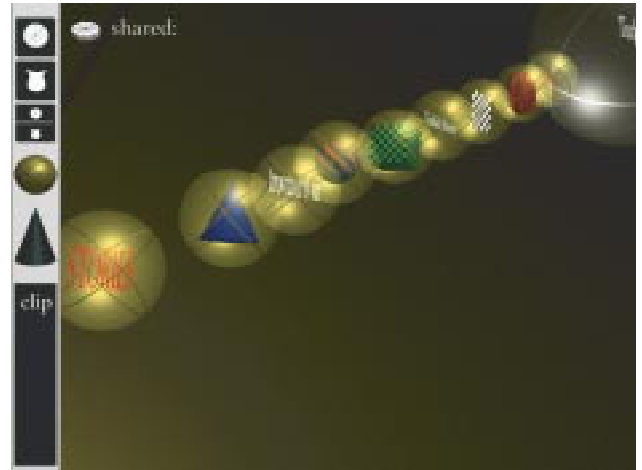

Moving closer

John inspects the objects lying close to the term spheres first. He concentrates on the objects between the term spheres since he wants to find information objects that concern both the Turkish bow and the English long bow. As he moves closer to information objects, the symbols for the objects fade away and the real information space contained by the object spheres are shown.
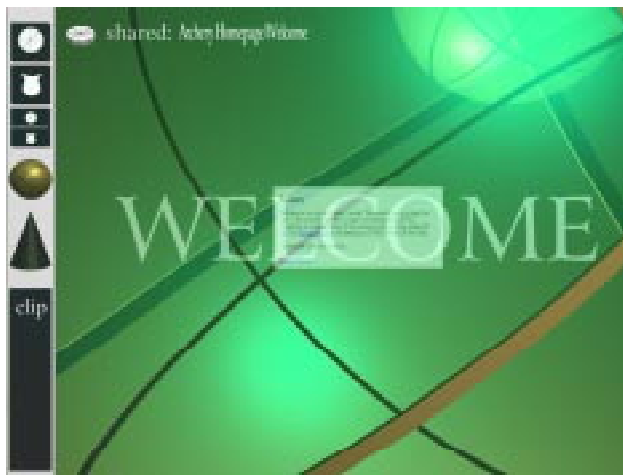
At this close distance from objects, the two term spheres John defined earlier show up on the surface of each single information object as well. If he wants to, he can increase their attraction so that information objects related to the two bow types are attracted to the terms as usual.
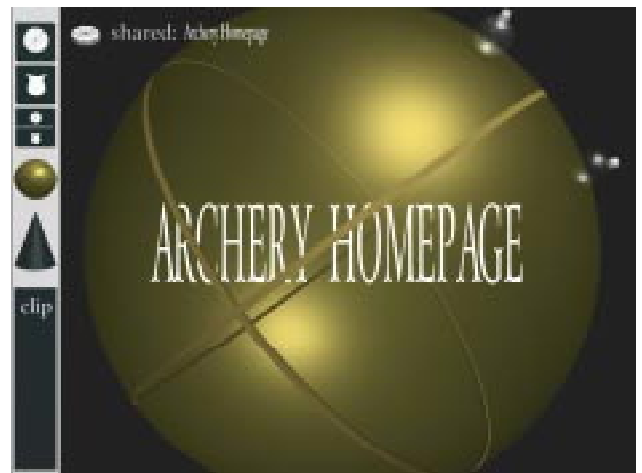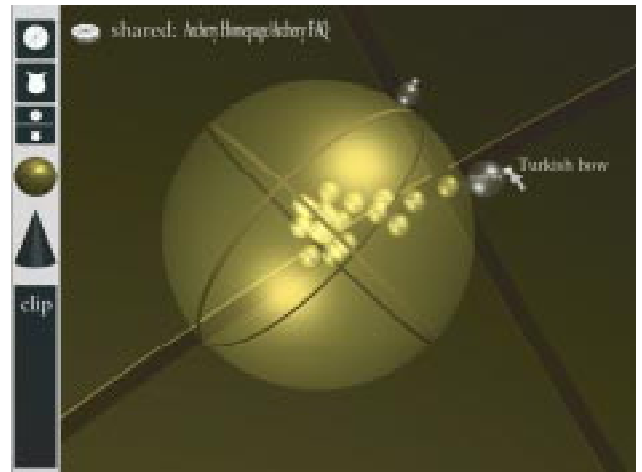
In many of the spheres, the space designers have put a 'welcome' sphere that explains the content of the current user-designed space.


After inspecting the 'welcome' sphere of the information object "Archery Homepage", John decides to take a closer look on the information object called "Archery FAQ". He manipulates the space using the term spheres on the information object surface,

and dives into the information object that has been most attracted to the term "Turkish Bow", which is "Asian & Turkish bow construction FAQ". He repeats this space manipulation process once again and inspects the information object named "The Composite Turkish Bow".

Changing object granularity
This information object is atomic, i.e. it doesn't contain any other objects but a single piece of information: a text document which reacts very strong on the attraction of the term "Turkish bow".

To find out which sections of the document that makes this relation so strong, he activates a granularity tool from the toolbar and increases the granularity level of the document by pulling a slider within the granularity tool. The effect is direct: sections covering the bow are pulled closer to the term while non-relevant sections get positions closer to the centre of the object sphere.
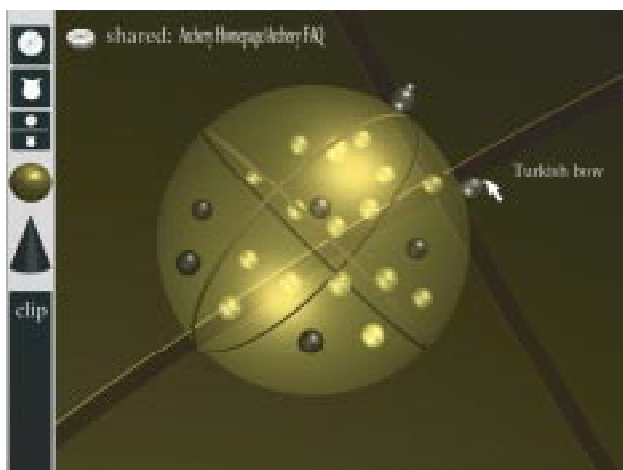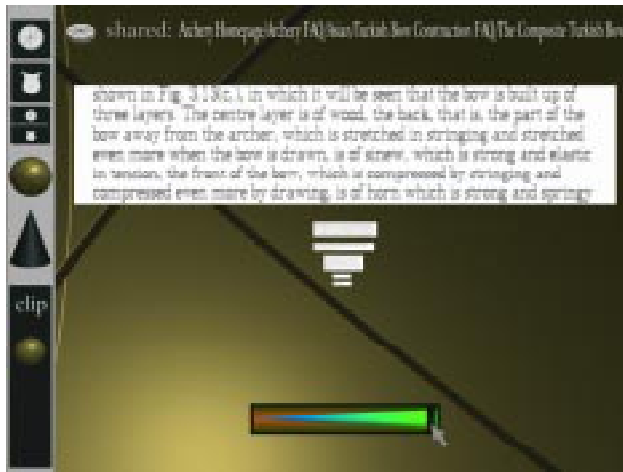
Making a private copy

When the desired granularity factor is reached, he inspects the attracted part more closely and finds detailed descriptions of the Turkish bow. What makes it special is apparently that it is constructed out of wood, sinew and horn in three layers. The sinew was cut in 50 mm pieces. "Strange" he thinks and decides to make a copy of this information object and drags the copy into his private space.

The terms and the granularity factor remains unchanged in the copy so that the interesting part of the object will be easy to find whenever he will inspect this private object again. Also, the relation between the object and it's parent object "Archery Homepage" is preserved so that if he in the future decides to zoom out from his private object copy, he will eventually end up inspecting the parent object, located in the shared space, and see in what context the object has been placed by the original space designer.

Switching between spaces

He is curious of what kind of person or organization has created this space. Maybe it contains information on the English long bow too? John switches to user-designed space mode in the shared space window. Since he is already positioned within a user-designed space ("Archery Homepage") this has no noticeable effect on the space currently in view. However, the system-designed space John started his journey within, formed by his terms "Turkish bow" and "English long bow", is exchanged with another system-designed space containing a collection of user-designed objects that in turn all contain "Archery Homepage". The system names this space "Archery Homepage Parent Space".

When John zooms out of "Archery Homepage", he could in theory land in any one of the objects contained in this system-designed parent space. By default, the system chooses a space designed by the same designer as the child space. On the other hand, if John decided to zoom out all the way, passing through the parent of "Archery Homepage", to the "Archery Homepage Parent
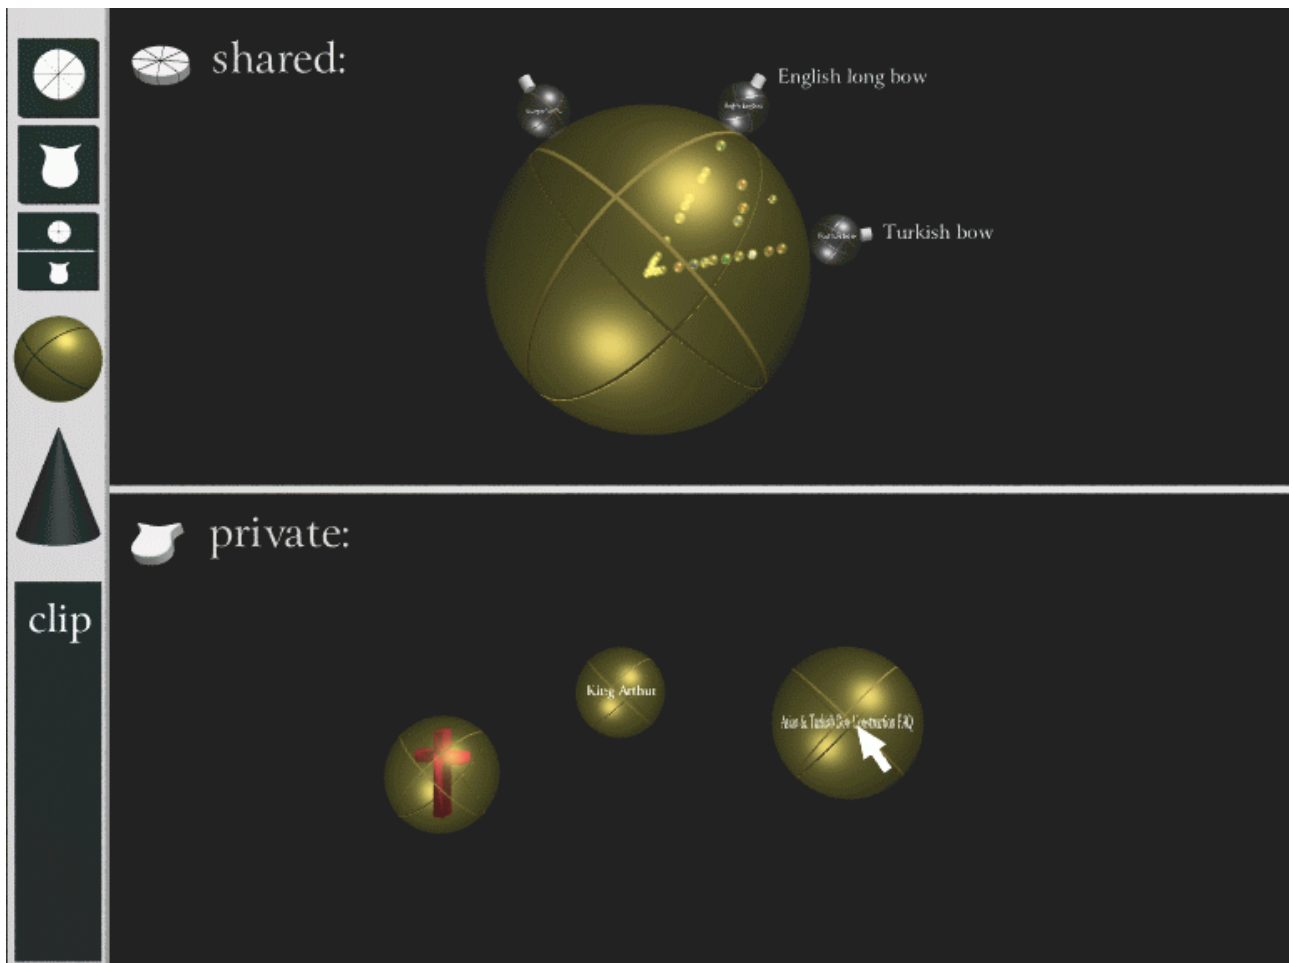
Space", he would have the opportunity of inspecting alternative user-designed spaces containing "Archery Homepage", e.g. to view this object in different contexts.

Another option for John would be to let the system generate a similarity based space. This space would contain objects similar to the one he is currently watching having the current object as a term on the space surface so that the objects that the system finds most similar to the current would be located close to the term (as usual).

Success!
John zooms out from "Archery Homepage" and lands in the space "University Archery Club: Sagittarius Twente", a space presenting an archery club located in the Netherlands. From here, he dives through the object space "Other Archery Web-Sites" and lands in a space named "Archery History sites" containing a set of objects that are rated regarding their contents by the space designer. John dives into the object called "Mary Rose Homeport" which happens to be the highest rated object. This space is a virtual museum for the British 16th century ship Mary Rose rediscovered in the 1970s.

John enters the object "The Virtual Museum Tour" and gets so interested in the ship that he forgets his original mission to find information about the English long bow. Eventually, the information text contains a link (object) called "Firepower" which makes him recall his interests in archery. John enters the firepower space and finds a description of Mary Rose's armoury. Well, ain't John a lucky guy! The text states that the ship crew assisted their iron guns with long bows made of jew in apparently one piece. John concludes that the Turkish bow had to be superior because of its composite construction!

Finishing up
John copies the object "Firepower" to his private space, and invokes the word processor tool within the private space. He writes a short story about his discovery and places it right beside the other two objects "Asian & Turkish bow construction FAQ" and "Firepower". He then places all three objects in an empty space (sphere) which he denotes "Turkish bow vs. English long bow".

For private use only
More than satisfied with his achievement, John goes into the kitchen and makes himself a cup of tea. During the break he wonders: "Why in the world did the Turks cut up the sinew in 50 mm pieces?". To find out, he takes his cup of tea with him and uses the system as described above. He finally lands in a space containing tables of physical constants. "Can't figure out a thing from these tables. These Turks must have been really smart..." John, who has never been to Turkey, takes another sip from his tea and thinks "Wonder if they make any tea in Turkey? If they make as good tea as they make bows, it has to taste lovely." He decides to ask his local tea dealer next time he goes shopping.

If he wanted, he could publish his new object by moving or copying it to one of his public spaces. However, although the conclusion he drew, the connection he made and the short story he wrote were valuable to himself (and will be in the future according to the Bush sce-

nario), he decides not to publish it. "It's probably of little use for anybody else."

Comments
All user-designed spaces mentioned in the scenario exists in the form of webpages sometimes linked to each other within the WWW at the time this is written. The zoom-out from "Archery Homepage" to its parent space "University Archery Club: Sagittarius Twente" is performed in current WWW browsers by clicking on a "Home"- button on the first page. However, there is currently no support in the WWW for general backtracking to pages that have links to a current page.

5. Implementation

Programming
This section describes the practical part of the project which consisted of applying a 3D direct manipulative information tool onto a user data set originating from a Music On Demand system.

5.1. Background

5.1.1. The RelevanceSphere

LyberWorld is a C++-library built on top of Open Inventor. LyberWorld is developed by GMD-IPSI in Germany. Different 3D direct manipulative tools are implemented in the toolkit and one of them is the Relevance-Sphere.

A Brief Description[1]
This tool is a transparent sphere containing a set of information objects visualized for instance as cubes (se figure 6). On the sphere surface are a set of smaller sized spheres located. These spheres are called terms and can attract the information objects within the RelevanceSphere. The size of a term visualizes it's attraction strength relative a set of information objects and can be altered by the

---

1. For a more fair and complete description on the nature of the RelevanceSphere, please read the authors' original paper on LyberWorld[35].
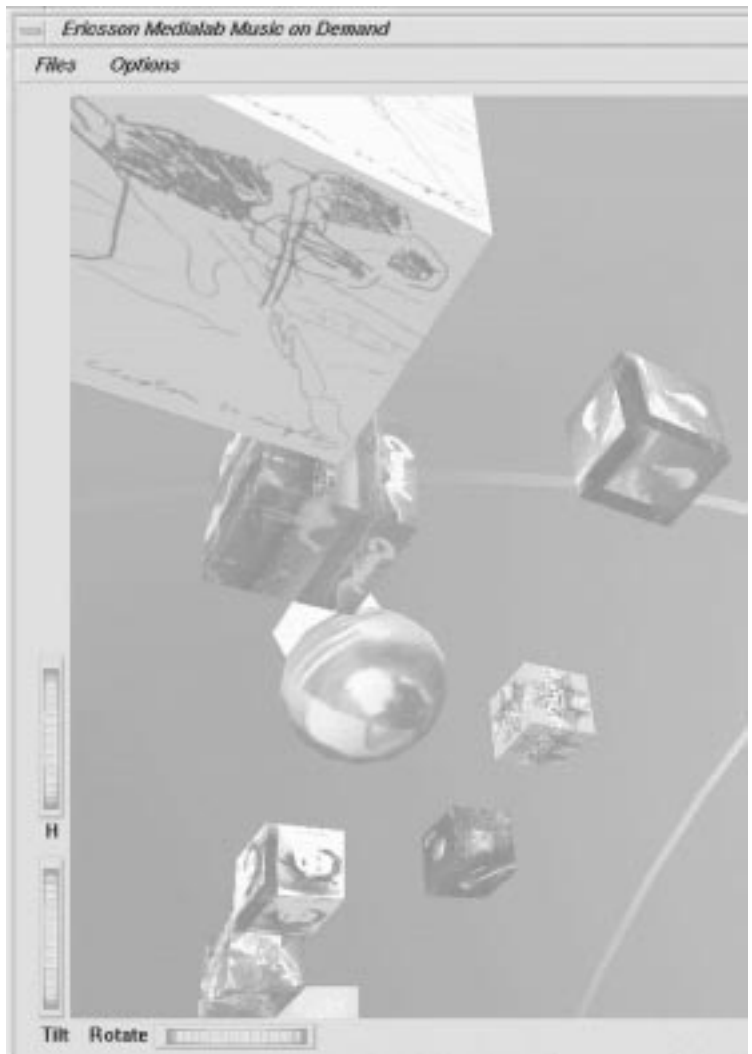

Figure 7: Browsing within the GMD RelevanceSphere. The cubes are records from the Ericsson Media Lab MOD system. Among the cubes is a user profile term sphere.
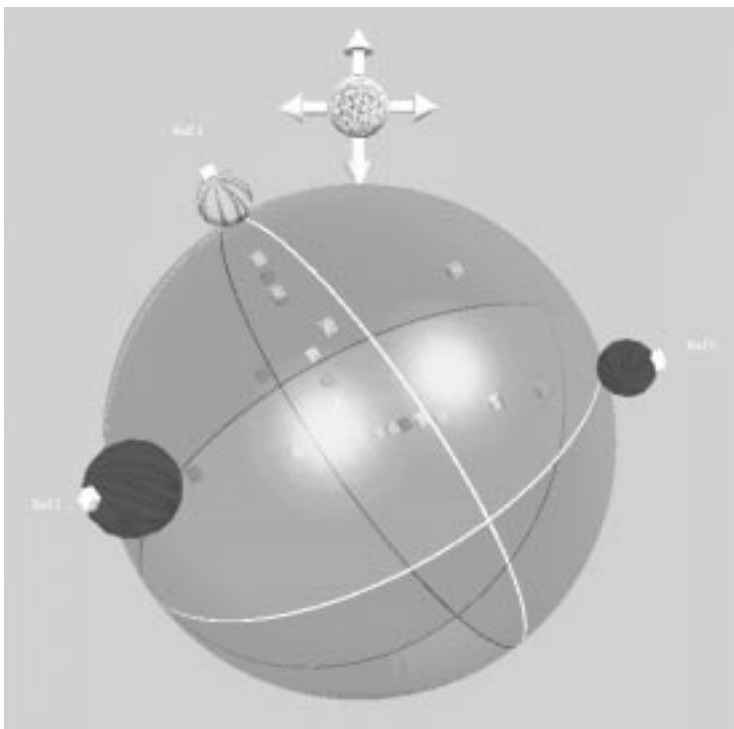
user by direct manipulation. The terms are freely movable around the sphere surface.

An important feature of the RelevanceSphere is its spatiality. Relations between information objects are shown by the spatial distances between them. Objects closely related to a term are close to that term in the space.

When a term is moved, all information objects related to the term react upon the movement based on the relation strength between the term and the object. The thing is that this happens in real time. You don't have to sit and wait (unless you use a slow machine of course).

In-space located dynamic queries
The above mentioned direct feedback mechanism, sometimes called 'Dynamic Queries' (DQ), is hard to describe in words but seems to have become quite a hot subject within Information Visualization and Information Retrieval. See for instance IVEE Spotfire[1][2] that makes heavy use of this kind of direct visual feedback. The RelevanceSphere's DQ mechanism differs from the one present in Spotfire in one important aspect: RelevanceSphere queries are 'asked' inside the visualization
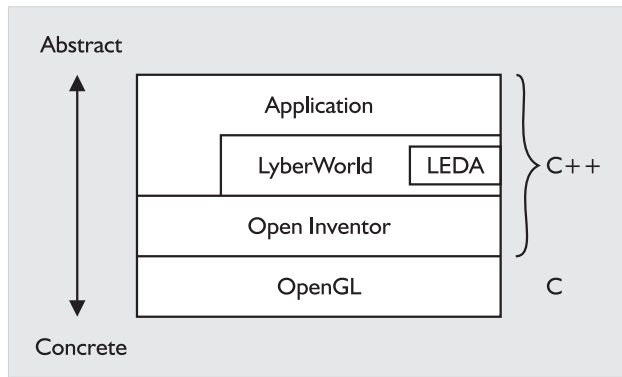

Figure 6: The GMD RelevanceSphere.

Figure 8: Developing Environment API layers.

environment whereas Spotfire DQs are asked using interaction tools located outside the visualization environment. Spotfire's flexible set of DQs (defined by the data set) makes this design necessary. However, if the data set is well-defined and if the amount of possible queries (or rather the amount of possible queries mapped to visual tools) are kept to a low number, they are possible to integrate in the information space, like it is done in the RelevanceSphere case.

5.1.2. Applying the RelevanceSphere on User Data

As explained in section 6, a full implementation of the intended system was not possible. Instead I just altered the RelevanceSphere so that it visualizes relations in an already existing data set, in a suitable way. The data originates from the Music on Demand (MOD) system developed at Ericsson Media Lab and consists of information about how many times a song has been listened to by a specific user. The songs are organized in record-clusters located on a large hard disk and each user's log-file is updated every time a song is played. (Users have been informed that this kind of statistic logging is done.)

Each user is dedicated a term on the Relevance-Sphere surface. A picture of the user is texture-mapped onto the term sphere. Each record is visualized as an information object (a box texture-mapped with the record cover on all sides) within the sphere. If a user has played a record many times, the corresponding record box is located close to the specific user term sphere.

5.2. Programming

5.2.1. API Layers

The application programming interface (API) layers for the developing environment used in this project is illustrated in figure 3. The main application was implemented using C++ and the LyberWorld toolkit.

5.2.2. Modifications to LyberWorld

The LyberWorld distribution I downloaded included the two tools RelevanceSphere and NavigationCone along with small example programs allowing the user to interact with the two tools respectively. My application is

```
#Generic sphere tool ascii format V1.0
NumberOfReferences 9
NumberOfInformation 50
    References {
            id hanna
            type 0
            texture/music1/music/.users/hanna.gif
            attraction 1.0

            id pelle
            type 0
            texture/music1/music/.users/pelle.gif
            attraction 1.0


            .
            .
            .

            id svenne
            type 0
            texture/music1/music/.users/svenne.gif
            attraction 1.0
    }
    Information {
            id 36 Svenska Klassiker 1980-1990
            type 1
            texture/home/www/demo/mod/samling/
                svenska_klassiker_1980_1990/
                svenska_klassiker.jpg
            relevance 1 19
            relevance 4 42
            relevance 8 16

            id Debut
            type 1
            texture/home/www/demo/mod/bjork/debut/
                bjork.jpg
            relevance 0 11
            relevance 4 3
            relevance 5 82
            relevance 7 25
            relevance 8 6


            .
            .
            .

            id Sondag i sangen
            type 1
            texture/home/www/demo/mod/
                bo_kaspers_orkester/
                sondag_i_sangen/
                bo_kaspers_orkester.jpg
            relevance 1 20
            relevance 5 2
            relevance 8 37
    }
#EOF
```

Figure 9: The data file used as input to the RelevanceSphere.

based on the example program using the Relevance-Sphere.

The Data File Parser
The data file parsing section of the application was modified so that information about record picture files and user picture files could be imported via the input data file. These were used for texture-mapping information object boxes and user term spheres with the record cover and user picture respectively.

An example input file is shown in figure 3. It is divided in two sections, one describing the terms (called 'References' in the file) and one section describing the information objects (called 'Information'). The terms are implicitly assigned numbers starting with 0 which are referred to in the Information section. The 'relevance' field in the Information section connects current infor-
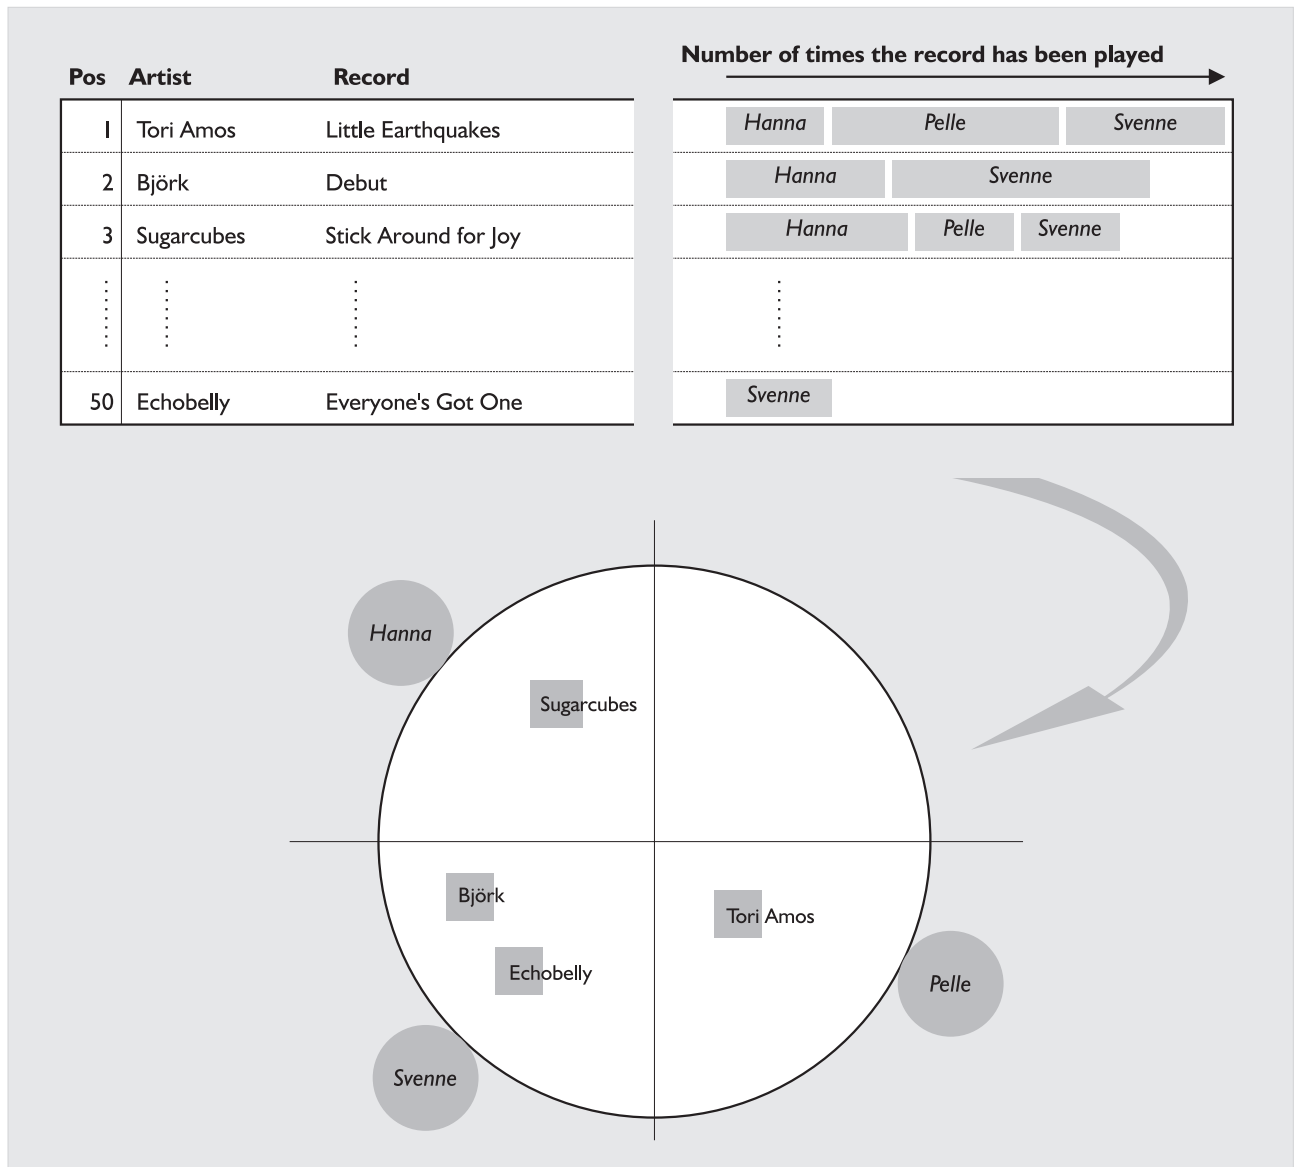
Figure 10: The MOD user data's way to the RelevanceSphere.

mation object to term no. x with strength y, where x is the first number and y is the second (see figure 3). The fields I have added to the file 'language' are the texture fields, denoting paths to picture files.

The Information Object Class
Most programming effort was put into a quite extensive alteration of the information object class. The original information objects were all defined as one single curve set instead of separate simple cubes (at least that's my interpretation of it) in the OI object world called Scene-Graph. I guess the reason for this was efficiency.

However, although I tried very hard, I couldn't map textures in a suitable way onto these curve sets. So I changed the curve set to single cubes belonging to each information object instead, which made the texture mapping a piece of cake. The backside was that I had to rewrite all methods concerned with these objects in the information object class (methods for updating object positions, methods for highlighting objects etc.).

Other Changes
The term class (actually called reference object class in the code) had to be altered to texture-map user pictures onto the term spheres.

Member variables had to be created in the relations class (handles data used when information object positions should be updated caused by term movements by calling the LEDA library), to store references to information object texture files.

The information object id's (in my case record names) was not passed on from the main program to the RelevanceSphere originally. Since many record covers don't show the record name and/or musician(s), I implemented a mechanism for showing the name of the record when clicking on it. To make this possible I had to change some methods so that the record name could be printed out when necessary.

I changed the viewer from a static positioned type to a viewer that allowed movement i z-direction. This was to allow the user to inspect the data set more closely (to see which records that was positioned where). Also I

wanted to see how navigation within a space filled with information objects feels. (My original intention was to implement a system based on this kind of browsing.)

5.2.3. Generating the Input File

To generate the input file for the RelevanceSphere, I wrote a Perl-script (based on a top-ten script written by Thomas Stephanson at Media Lab) that constructs the file based on relevant data from the MOD system. Selecting the set of records to visualize can be done in many ways. I chose simply to filter out the top 50 most played records. The number of times a user has played a specific record decides how high default attraction that user term will have on the specific information object, and is written as it is to the input file. An not to unfair assumption is made that a user likes records he or she plays often. LyberWorld normalizes the values at a later stage to get all information within the sphere surface.

The process of generating the RelevanceSphere based on MOD user data via the Perl-script is illustrated in figure 10.

5.3. Motives

... for applying the 3D-tool on the MOD system
Why construct a MOD user interface based on 3D browsing and spatial relations? Wouldn't some kind of ordinary text list do? I can't say that the chosen data set is optimal for the RelevanceSphere tool but this kind of interactive visualization has a potential to make at least two tasks easier for the MOD-user:

- Getting in touch with people having similar taste in music. By studying other peoples music tastes, you can find musical friends.
- Finding new interesting music by association. When you find a user (-sphere) that attracts many of the records you are listening to yourself, you will almost certainly find records in his or her record set that you have never seen (or heard) before. Since the other records are of the kind that you listen to yourself, you will probably like the new ones too.

... for using 3D instead of 2D (in general)
In general, the visualization of spatial relations between objects using a 3D space instead of 2D has at least two motives:

- It can be proven mathematically that the higher degree of freedom introduced when using 3D instead of 2D reduces the risk of 'unclear interpretations'. That is, a situation where relations between information objects and terms can be interpreted in more than one way. This phenomena is discussed in detail in the LyberWorld paper[35]. In fact, in the paper, the RelevanceSphere is introduced as a 3D-enhancement of a 2D tool within another visualization environment, VIBE[5], just because of this reduced risk for redundancy.
- 3D space is superior vs. 2D space when it comes to

visualization of large information sets because it can present more information objects within the same limited screen space compared to 2D. This comes from our human power to remember spatial locations of objects when we navigate in 3D space, and that we easily can remember how object looks like even if they are partially obscured by other objects.

... for not using 3D (in general)
All good things have seamy sides. 3D visualization has at least three:

- Object positions in the z-plane (depth) can sometimes be hard to tell. This is overcome by rendering techniques using for instance a light source located near the viewing point so that objects further away are darker. Also, by allowing the user to easily move the viewpoint from one side to another (like owls do to get perspective), object positions in the z-plane are easier to see.
- It can be tricky to navigate in 3D space. Using mouse and a combination of buttons requires practice. 3D-mice, gloves and gesture-navigation are all enhanced 3D navigation and interaction tools making life easier in the virtual worlds and will probably become more common in the future.
- 3D graphics are expensive, both in monetary units and in computer power. However, this argument is definitely beginning to loose its strength since cheap 3D-accelerated graphic cards for low-end computers are starting to show up on the market.

5.4. Usage Reflections

No real user test has been performed, but many members of Media Lab couldn't keep their fingers away, so I got some feedback anyway:

Many members thought that it was fun just seeing what the other users prefers to listen to, and to find common music interests. Funniness is not a bad property for a visualization tool!

To actually visualize the records, instead of for instance show them as a text list, gives a totally different kind of impression. It invites to exploration of the space; to run into records 'by accident', just like web browsing does compared to more traditional text-query based information retrieval systems.

In a sense, music records (or rather: music record covers) are easy to visualize since they differ very much in visual design and are therefore easy to remember and easy to differ from each other. Thus, many users could very fast sort out which of the records were unknown and which of them the user has played her/himself before.

5.5. Extensions

For the time being, the data file is generated by pushing a certain 'View 3D-relations'-button within the MOD user interface, causing a data file based on the current state of user log files to be generated. If the user runs a

SGI machine, the file can be imported into my Relevance-eSphere application and inspected.

Integrating the Application with the MOD system

Jump-starting the application
The application could be tighter coupled to the MOD user interface by being automatically started when pushing the button for the 3D-viewer, avoiding the annoying step of importing the file into the application manually.

More users
The porting of LyberWorld to the Windows NT platform is reported to be finished. If the application should be compiled on this platform, many more users could benefit from it.

Running the application within MOD
TGS Software has developed an ActiveX-based solution making it possible to run OI applications within Microsoft Explorer and are currently working on a Netscape Navigator plug-in version. Licensing the latter software, and licensing LyberWorld for Windows NT (see above) would make it possible to run my application within the Virtual Information Workspace it belongs, namely the MOD environment (based on Netscape Navigator).

Extending the Application Functionality
The space filled with records could be used in many more ways than it is for the time being. Various functionality present in the existing MOD system could be integrated in the 3D-view:

A list of tracks could be found inside every record-cube. When clicked at, the desired track is played. Music from a record could be played automatically if the user comes close enough to the specific record box, fading in volume if the user turns away. Homepages belonging to the different users could be linked to their respective term sphere, accessible on mouse click.

In principle, the entire MOD user interface could be transformed into a 3D UI. Some functionality would benefit more than other, so what is found to be better done in 2D could be kept 2D even within the new environment. A transformation of a 2D user interface like the current MOD one, into 3D spatially based environment, would be a great challenge opening for many questions how traditional 2D interaction schemes could best be visualized and implemented in this more 'space-free' environment.

6. About This Project

Project Development
This section presents what I have done, what I haven't done, turns of project focus and time schedule.

6.1. Background

6.1.1. The Idea

The idea for this project originates from papers written during two specializing courses taken at the Dept. of Computing Science at Umeå University autumn term 1995 and spring term 1996 respectively. During the course 'Human Computer Interaction Advanced Course' I wrote an essay about the creative process and the need for considering this process when modelling Human-Computer Interaction (HCI) [60]. During the course 'New Media' I had creativity as main design parameter when developing a proposal for a hypermedia visualization system [61].

The current master thesis was intended to result in a prototype implementation of the system presented in the above mentioned hypermedia essay. However, it turned out to be a too optimistic plan…

6.2. Project Development

6.2.1. Theory

Studying Information Visualization
The hypermedia essay mentioned earlier described very briefly an imaginary creativity-supportive hypermedia system based on human thinking processes like abstraction and association. To find out possible visualization alternatives, I started the project by reading articles mainly concerning general information visualization and specifically visualization of large data sets. Approximately half of the total 20 project weeks have been used to study what people have done and currently are doing within this scientific area. Other areas which I have taken a glance at during this period involve Information Theory, Information Retrieval, Knowledge Work and Computer Supported Cooperative Work.

Conceptually Designing the User Interface — Not
Initially I believed that this part of the project would be a small one. However, soon I realized that considerable effort was needed to find solutions to problems not concerned in my hypermedia essay.

Designing a Model
I decided to construct a model that described the nature of the underlying information sets that was to be interacted with through the user interface. By doing this I would hopefully gain understanding about which functionality the user interface should possess and what kind of metaphors to use, if any. This work resulted in the VIW model discussed in section 3, visualized in figure 3 on page 9.

6.2.2. Practice

Finding the Right Developing Environment
Along with the reading, I searched for a suitable developing environment for my system. The two main tools to be found was a good visualization environment and a suitable database engine.

Visualization
For various reasons, Silicon Graphics' C++ 3D-visualization toolkit Open Inventor[74] (OI) was chosen. Another hot candidate was the –at the time brand new– VRML2.0[64]. It was rejected because my prototype needed a higher degree of flexibility which VRML2.0 couldn't support[1].

One of the most interesting articles about information visualization was Hemmje, Kunkel, & Willet's paper about LyberWorld[35]. Among other things, they presented a 3D direct-manipulation tool called RelevanceSphere that enclose sets of information objects in way very similar to how I would like to do it in my own system. Here was a toolkit that would enable me to develop the system at a high abstraction level but still – since it was based on OI– with possibilities to control 3D-objects and rendering in detail. After some email, fax and mail-correspondence with Mr. Matthias Hemmje and other people at GMD-IPSI, a right-to-use license and a two-week installation support contract was signed for LyberWorld.

Database engine
The most interesting database engine I ran into was INQUERY developed at the Information Retrieval Laboratory of University of Massachusetts, based on a probabilistic representation model called inference network. The reason of interest was that this kind of representation would allow atomic objects content as well as object cluster content to be represented internally in a fully uniform manner, namely as inference networks only differing in size. Coincidentally, the LyberWorld team (see below) happened to use exactly this database engine together with their user interface! After quite a struggle with licensing administration[2], a temporary license of INQUERY was achieved.

Technical complications
After deciding to use OI, a suitable hardware platform had to be found. At first, a Windows NT-based machine was considered since it was available at the lab and OI was available on that platform too. This alternative was rejected due to the high price of OI for Windows NT. Instead, Open Inventor was licensed and installed on a for the time being not used SGI Indy-machine. This machine had to be enhanced with a larger hard disk though. After a few days it was decided that the Indy was too weak to serve as a developing machine.

By that time, an SGI Indigo2 Impact (more power) machine was free to be used so OI was installed on that machine instead. Since SGI's C++-compiler wasn't licensed for this machine I tried to use GNU C++ but for some reason, they were not on talking terms. A new C++-compiler was licensed and installed. This state was reached after 7 weeks.

LyberWorld Installation
At this point I read the article describing LyberWorld and contacted Mr. Hemmje at GMD-IPSI to find out the possibilities of using it in my project.

The first distribution was downloaded from the net after 3 weeks from the point when I contacted GMD for the first time. The installation procedure was not easy (Mr. Hemmje had informed me that LyberWorld was still a research prototype so this came as no surprise), but with good help from Mr. Arnd Steinmetz and Mr. Joachim Nelz at GMD I finally got it up and running. Along with the traditional system-dependent stuff like the altering of 'makefiles' etc., I had to download the shareware optimized computational algorithm C++-library called LEDA, developed by LEDA Software GmbH and used by LyberWorld to compute relative positions of objects.

Finally, with 7 weeks left, I had a fully functioning developing environment.

INQUERY
The INQUERY-license took another 4 weeks, so I had to abandon the use of it. I started to realize that I wouldn't manage to produce the kind of prototype I wanted from the start anyway:

If you don't succeed, redefine success.
(Unknown.)

6.2.3. Revising the Goals

After consultation with my project supervisors I decided to change the goals of the project. Instead of actually implementing a prototype of the hypermedia system I would present a design proposal for it. The implementation part would consist of some application that made use of the LyberWorld toolkit. In large, this is also the final outcome of the project.

---

1. I would like to thank Mr. Sougata Mukherjea, user of both systems and co-author of [55][56][57][58], for advice in the matter.
2. I would like to thank Mr. Lars-Erik Janlert and Mr. Per Lindström at the Dept. of Computing Science, Umeå University, for supporting me on this.

7. Conclusions, Claims and Visions

Today's knowledge worker has access to an enormous amount of information stored digitally. Visualizing large sets of information makes it possible for humans to handle more information since it relieves the burden of keeping every single piece of information in our minds. 3D visualization is a good alternative to its 2D counterpart since it is able to present larger sets of information pieces.

Still, the total amount of information available can never be visualized on one screen, and 99.9% of it is in any case irrelevant for the user in his/her task. Therefore, all help we can get from the system to find the right information is appreciated. Since knowledge work essentially depends on our own thinking process, a good source of inspiration for system mechanisms supporting information retrieval could be cognitive science.

The future knowledge worker needs systems that make it easy to transfer pieces of information between private and shared information spaces. A system based on a model of the most important digital information spaces available, is a move in the right direction towards such an integrated environment. Such a system should, in line with Vannevar Bush's thoughts, also provide object and context 'freezing' mechanisms so that certain views within the information space could be stored, combined and re-retrieved at a later stage.

Humans have unconscious talents not fully utilized in user interfaces. To find these talents and to utilize them is an important task for software designers in general and knowledge work system designers particularly. To measure things by their spatial positions is one talent that could be utilized when visualizing information sets, by showing relationships between pieces of information spatially. Direct manipulation and the use of Dynamic Queries gives dynamic direct feedback of user actions. These are powerful user interface concepts and should be considered whenever possible.

Encapsulating information is necessary in every communication process. However, such packaging removes the particular piece of information from its original context. Systems supporting communication should as far as possible support interpretation of information pieces in relevant contexts. To be able to study a particular piece of information in different (but relevant) contexts could enhance ease of learning but also be a mechanism enhancing creative thinking.

To base knowledge work systems on a metaphor partly consisting of information pieces encapsulated in disconnected 'files' of fixed granularity is not optimal. Instead a more flexible information package metaphor that allow relations between information objects and parts of information objects to be visualized and used when retrieving and submitting information, should be considered. Such a variable granularity-based encapsulation would be even more feasible if the system could support the user with context suggestions.

8. Acknowledgements

9. Bibliography / References

Many of the papers listed here are not referenced to from within the essay, but have influenced the work nevertheless. Most of the non-referred papers concern Information Visualization in some form.

Note: Many URLs are links to the at the time latest results of research done by their respective authors, often containing links to older published papers. All URLs were fully functional October 6 1996.

1. Ahlberg, C., et al. (1996) "Spotfire", IVEE Development AB, Sweden.
   `http://www.ivee.com/`

2. Ahlberg, C., & Shneiderman, B. (1994) "Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays" in *Proceedings CHI'94: Human Factors in Comp. Systems.*
   `http://www.cs.chalmers.se/SSKKII/tech-reports/SSKKII-94.02.ps`

3. Bederson B., & Hollan, J., & Perlin, K., & Meyer, J., & Bacon, D., & Furnas, G. (1996) "Pad++: Zoomable Graphical Interfaces", University of New Mexico and NYU Media Research Laboratory.
   `http://www.cs.unm.edu/pad%2b%2b/begin.html`

4. Belkin, N. J, & Croft, W. B. (1992) "Information Filtering and Information Retrieval: Two Sides of the Same Coin?" in *Communications of the ACM*, December 1992/Vol.35, NO.12.

5. Benford, S., & Snowdon, D., & Greenhalgh, C., Ingram, R., & Knox, I., & Brown, C. (1995) "VR-VIBE: A Virtual Environment for Co-operative Information Retrieval", Eurographics'95, 30th August - 1st September, Maastricht, The Netherlands, pp 349-360.

6. Benyon, D. (1992) "Task analysis and systems design: the discipline of data" in *Interacting with Computers* **4** (1), pp246-259.

7. Bier, E., & Fischkin, K., & Pier, K., & Stone, M. (1995) "The Magic Lense Interface Project", Xerox PARC.
   `http://www.parc.xerox.com/istl/projects/MagicLenses/`

8. Blair, D. C. (1990) "Language and Representation in Information Retrieval", Elsevier, Oxford, U.K.

9. Bødker, S. (1989) "A Human Activity Approach to User Interfaces" in *Human-Computer Interaction*, 4:171-195.

10. Bush, V. (1945) "As We May Think" in the July 1945 issue of *The Atlantic Monthly.*
    `http://www.isg.sfu.ca/~duchier/misc/vbush/`

11. Callan, J. (1996) "Document Filtering With Inference Networks." in *Proceedings of the Nineteenth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Zurich, Switzerland. (pp. 262-269).
    `http://ciir.cs.umass.edu/info/psfiles/irpubs/callansigir96a.ps`

12. Callan, J.P., & Croft, W.B., & Harding, S.M. (1992) "The INQUERY retrieval system." in *Proceedings of the Third International Conference on Database and Expert Systems Applications*, (pp. 78-83). Valencia, Spain: Springer-Verlag.
    `http://ciir.cs.umass.edu/info/psfiles/irpubs/callancroftdexa92.ps`

13. Candy. L., & Edmonds, E. (eds.) (1996) "Creativity and Cognition 1996", Proceedings, LUTCHI Research Centre, Loughborough University, UK.
    `http://lucy.lboro.ac.uk/ccog2.html`

14. Card, S. K. (1996) "Visualizing Retrieved Information: A Survey", Special Report, Computer Graphics and Visualization in the Global Information Infrastructure, CG&A, Vol. 16, No. 2, March 1996.
    `http://www.computer.org/pubs/cg&a/report/g20063.htm`

15. Century, M. (1996) "Notes on Merz"
    `http://mcentury.citi.doc.ca/Notes_on_Merz.html`

16. Chalmers, M. (1992), German translation by Schlett, M. (1995) "BEAD : EXPLORATIONS IN INFORMATION VISUALIZATION" (German language).
    `http://www-cui.darmstadt.gmd.de/visit/Activities/Courses/VirtuelleRealitaet.SS.95/Reports/BEAD/#4`

17. Chuah, M.C., & Roth, S.F., & Mattis, J., & Kolojejchick, J. (1995) "SDM: Malleable Information Graphics" in *Proceedings of Information Visualization*, Atlanta, GA, October 1995, pp. 36-42.
    `http://www.cs.cmu.edu/Web/Groups/sage/SDM-InfoViz95/SDM-InfoViz95.html`

18. Chuah, M.C., & Roth, S.F., & Mattis, J., & Kolojejchick, J. (1995) "SDM: Selective Dynamic Manipulation of Visualizations" in *Proceedings of the ACM Symposium on User Interface Software and Technology*, Pittsburgh, PA, November 1995, pp. 61-70.
    `http://www.cs.cmu.edu/Web/Groups/sage/SDM-UIST95/SDM-UIST95.html`

19. Colby, G., & Scholl, L. (1991) "Transparency and Blur as Selective Cues for Complex Visual Information" in *Proceedings of ISOE*, Volume 1460, 114-124.

20. Daessler, R., & Otto, A. (1996) "3D-Navigation in Virtual Information Spaces: From Text-based Retrieval to Cognitive User Interaction".
    `http://www.gfz-potsdam.de/drz/geolis/knowledge/knobro.ps`

21. Daessler, R. (1995) "Visualization of Abstract Information", Minnesota Supercomputer Institute.
    `http://banzai.msi.umn.edu/~roelf/tut.html`

22. Dostoevsky, F. (1874) "The Adolescent", english translation by Andrew R. MacAndrew in 1971, W. W. Norton & Company, Ltd, London WC1A 1PU.

23. England, D. (1996) "Critique of Approaches to Interaction and Visualisation", GMD - German National Research Center for Information Technology, Institute for Applied Information Technology, Sankt Augustin, Germany.
    `http://orgwis.gmd.de/~de/Research/comic_crit.html`

24. Fairchild, K. M. (1996) "Information Visualization", National University of Singapore.
    `http://panda.iss.nus.sg:8000/kids/fair/webdocs/viz/viz-1.html`

25. Fairchild, K. M., & et al. (1988) "SemNet: Three-Dimensional Graphic Representations of Large Knowledge Bases", pp200-233 in *Cognitive Science*

And its Applications for Human-Computer Interaction, Raymonde Guindon (ed.), ISBN 0-89859-884-2, Lawrence Erlbaum Associates Inc., USA.
`http://panda.iss.nus.sg/kids/fair/webdocs/semnet/`
`        semnet-1.html`

26. Flower, L., & Hayes, J. R. (1980) "A Cognitive Process of Theory of Writing", pp365-387 in College Composition and Communication **32**.

27. Fox, D. (1995) "Tabula Rasa: A Pad-style Multi-Scale User Interface System", New York University.
`http://galt.cs.nyu.edu/students/fox/tab.html`

28. Fox, D. (1994?) "Graphical Elements for Information Browsing Systems", Survey paper submitted towards fulfillment of Ph. D. requirements, New York University's Department of Computer Science.
`http://galt.cs.nyu.edu/students/fox/area.html`

29. Fox, D. (1994?) "The Design of Multi-Scale User Interface Systems" (draft), New York University's Department of Computer Science.
`http://found.cs.nyu.edu/fox/multiscale.html`

30. Francis, P., & Kambayashi, T., & Sato, S., & Shimizu, S. (1996) "Ingrid: A Self-Configuring Information Navigation Infrastructure", NTT Software Labs, Japan.
`http://www.ingrid.org/francis/www4/Overview.html`

31. Furnas, G. W., & Zacks, J. (1994) "Multitrees: Enriching and Reusing Hierarchical Structure", CHI'94 Proceedings.
`http://matia.stanford.edu/~zacks/multitrees/`
`        multitrees.ps`

32. Furnas, G. W., & Bederson, B. B. (1994) "Space-Scale Diagrams: Understanding Multiscale Interfaces" in *ACM SIGCHI '95.*
`http://community.bellcore.com/gwf/CHI95.html.paper/`
`        gwf_bdy_inline.html`

33. Goldstein, J., & Roth, S.F., & Kolojejchick, J., & Mattis, J. (1994) "A Framework for Knowledge-Based, Interactive Data Exploration" in *Journal of Visual Languages and Computing*, 5, December 1994, pp. 339-363.
`http://www.cs.cmu.edu/Web/Groups/sage/JVLC/JVLC.html`

34. Goldstein, J., & Roth, S.F. (1994) "Using Aggregation and Dynamic Queries for Exploring Large Data Sets" in *Proceedings of the Conference on Human Factors in Computing Systems (SIGCHI '94)*, Boston, MA, April 1994, pp. 23-29.
`http://www.cs.cmu.edu/Web/Groups/sage/`
`        IntegratingTechniques/`
`        IntegratingTechniques.html`

35. Hemmje, M., & Kunkel, C., & Willet, A. (1994) "LyberWorld - A Visualization User Interface Supporting Fulltext Retrieval" in *Proceedings of ACM SIGIR '94*, July 3-6, Dublin.
`ftp://ftp.darmstadt.gmd.de/pub/VISIT/papers/hemmje/`
`        SIGIR94.ps`

36. Hirata, K., & Hara, Y., & Shibata, N., & Hirabayashi, F. (1993) "Media-based Navigation for Hypermedia Systems" in *Proceedings of Hypertext '93 Conference*, pp159-173, November 1993.

37. Honkela, T., & Kaski, S., & Lagus, K., & Kohonen, T. (1996) "WEBSOM - Self-Organizing Map for Internet Exploration", Helsinki Univ. of Technology.
`http://websom.hut.fi/websom/`

38. Hutchins, W. J. (1978) "The concept of 'aboutness' in subject indexing", Aslib Proceedings, 30(5) 172-181.

39. Ingram, R., & Benford, S. (1995) "Improving the Legibility of Virtual Environments", in *Proceedings of the 2nd Eurographics Workshop on Virtual Environments*, Monte Carlo, 31 Jan - 1st Feb 1995.
`http://www.crg.cs.nott.ac.uk/crg/Research/leads/`
`        euro.html`

40. Ingram, R., & Benford, S. (1995) "Legibility Enhancement for Information Visualisation" in *Proceedings of Visualization'95*, Atlanta, Georgia, November 1995.
`http://www.crg.cs.nott.ac.uk/crg/Research/leads/`
`        viz95.html`

41. Ingwersen, P. (1995) "Information and Information Science" in *Encyclopedia of Library and Information Science,* pp137-174, Marcel Dekker Inc., 270 Madison Avenue, New York, New York 10016. ISBN 0-8247-2056-3.

42. Ingwersen, P. (1992) "Information Retrieval Interaction", Taylor Graham, London.

43. Janlert, L-E. (1995) "Hemma i cyberspace", Scandinavian University Press, ISBN 91 88584 06 2.

44. Johnson, B., & Shneiderman, B. (1993) "Treemaps: a space-filling approach to the visualization of hierarchical information structures" in *Sparks of Innovation in Human-Computer Interaction*, Shneiderman, B. (ed.), Ablex Publ., Norwood, NJ (1993).
`ftp://ftp.cs.umd.edu/pub/papers/papers/2657/`
`        2657.ps.Z`

45. Judge, A. J. N (1994?) "ENVISAGING THE ART OF NAVIGATING CONCEPTUAL COMPLEXITY", Union of International Associations, Brussels, Belgium.
`http://www.uia.org/uiadocs/artnavig.htm`

46. Koike, H., & Yoshihara, H. "Fractal Approaches for Visualizing Huge Hierarchies" in *Proceedings of the 1993 IEEE Symposium on Visual Languages*, pp.55-60, IEEE/CS, 1993.
`http://roma.vogue.is.uec.ac.jp/~koike/papers/vl93/`
`        vl93.html`

47. Lenman, S., & See, H. (1996) "Personalized Information Spaces"
`http://merzbau.citi.doc.ca/merz/docs/w3c4/`
`        personalspaces.html`

48. Lenman, S., & See, H., & Century, M. (1996) "Merz: Constructing Personal and Shared Spaces on the World Wide Web", Centre d'innovation en technologies de l'information, Canada.
`http://merzbau.citi.doc.ca/merz/docs/w3c5/`
`        Overview.html`

49. Lenman, S., & See, H., & Century, M. (1996) "Merz, Hypertext '96 Poster", CITI, 1575 boul. Chomedey, Laval, Québec, Canada H7V2X2.
`http://merzbau.citi.doc.ca/merz/docs/ht96Poster/`
`        poster.html`

50. Lieberman, H. (1994) "Powers of Ten Thousand: Navigating in Large Information Spaces" in *Proceedings of UIST'94 Symposium on User Interface Software and Technology*, Marina del Rey, November 1994. New York: ACM.

51. Lin, X. (1996) "Visual SiteMap —Map of World Wide Web Sites", University of Kentucky.
`http://lislin.gws.uky.edu/Sitemap/Sitemap.html`

52. Long J. (1995) "Intelligent Archive —Integrated Data Access and Organization for Scientists", Lawrence Livermore National Laboratory (LLNL).
`http://www.llnl.gov/ia/ia_flier.html`

53. Lucas, P., & Roth, S. F., & Gomberg, C. C. "Visage: Dynamic Information Exploration" in *Conference Companion of the Conference on Human Factors in Computing Systems (SIGCHI '96)*, Vancouver, BC, Canada, April 1996.

54. Mackinlay, J. D., & Card, S., & Robertson, G. (1991) "Perspective Wall: Detail and Context Smoothly Integrated" in *Proceedings of the ACM SIGCHI '91 Conference on Human Factors in Computing Systems*, pages 173-179, New Orleans, La, April 1991.

55. Mukherjea, S., & Foley, J. D., & Hudson, S. (1995) "Visualizing Complex Hypermedia Networks through Multiple Hierarchical Views" in *CHI '95 Proceedings?*.
`http://www.cc.gatech.edu/gvu/people/Phd/sougata/chi95/sm_bdy.html`

56. Mukherjea, S., & Foley, J. D. (1995) "Visualizing the World-Wide Web with the Navigational View Builder" in *WWW '95 Darmstadt Proceedings*.
`http://www.igd.fhg.de/www/www95/proceedings/papers/44/mukh/mukh.html`

57. Mukherjea, S., & Foley, J. D., & Hudson, S. (1994) "Interactive Clustering for Navigating in Hypermedia Systems" in *Proceedings of the ACM European Conference of Hypermedia Technology*, pages 136-144, Edinburgh, Scotland, September 1994.

58. Mukherjea, S., & Foley, J. (1994) "Navigational View Builder: A Tool for Building Navigational Views of Information Spaces" in *ACM SIGCHI '94 Conference Companion*, pages 289-290, Boston, Ma, April 1994.

59. Norman, D. A. (1998) "Cognitive Engineering" pp31-61 in *User Centered System Design*, Norman, D. A., & Draper, S. W., (eds.), Lawrence Erlbaum.

60. Pederson, T. (1996) "Det ensamma trädet får drängen sköta – om kreativitet, MDI och de nya datorbaserade arbetsuppgifterna" (Swedish language), unpublished, Dept. of Computing Science, Umeà University, Sweden.

61. Pederson, T. (1996) "Nyfiken i en rymd – hypermediesystem baserat pà associationsplan i 3D - ett exjobbsförslag" (Swedish Language), Dept. of Computing Science, Umeà University, Sweden.
`http://www.cs.umu.se/tdb/kurser/TDBD07/VT-96/laborationer/GRP1/nm_upp.html`
`http://www.cs.umu.se/~dvltpn/nm_upp.html`

62. Pitkow, J., & Bharat, K.(1994) "WEBVIZ: A Tool for World-Wide Web Access Log Visualization" in *Proceedings of the First International World-Wide Web Conference*, Geneva, Switzerland, May 1994.

63. Potter, W. D., & Trueblood, R. P. (1988) "Traditional, Semantic, and Hyper-Semantic Approaches to Data Modeling", pp53-63 in *Computer* (ISSN 0018-9162), June 1988, Computer Society of the IEEE, USA.

64. rikk@best.com, & cmarrin@sgi.com, & gavin@acm.org (1996) "VRML 2.0 Cover Page".
`http://vrml.sgi.com/moving-worlds/`

65. Risch, J. S., & May, R. A., & Downson, S. T., & Thomas, J. J. (1996) "A Virtual Environment for Multimedia Intelligent Data Analysis" in *IEEE Computer Graphics and Applications*, November 1996, pp33-41.

66. Robertson, G. G., & Card, S. K., & Mackinlay, J. D. (1993) "Information Visualization using 3D Interactive Animation" in *Communications of the ACM*, April 1993/Vol.36, No.4.

67. Robertson, G. G., & Mackinlay, J. D., & Card, S. (1992) "Cone Trees: Animated 3D Visualizations of Hierarchical Information" in *Proceedings of the ACM SIGCHI '91 Conference on Human Factors in Computing Systems*, pages 189-194, New Orleans, La, April 1991.

68. Sarkar, M., & Brown, M. H. (1992) "Graphical Fisheye Views of Graphs" in *Proceedings of the ACM SIGCHI '92 Monterey*, May 3-7 1992, pp83-91.

69. Schaffer, D., & Greenberg, S. (1993) "Sifting through hierarchical information" in *Proceedings of posters and short papers ACM INTERCHI '93 Conference on Human Factors in Computing Systems*, Amsterdam, Holland.

70. Schaffer, D., Zuo, Z., Greenberg, S., Bartram, L., Dill, J., Dubs, S. and Roseman, M. (1996). "Navigating Hierarchically Clustered Networks Through Fisheye and Full-Zoom Methods" in *ACM TOCHI Transactions on Computer-Human Interaction*, 3(2), p162-188, June.
`http://www.cpsc.ucalgary.ca/projects/grouplab/papers/FisheyeTochi/FisheyeTOCHI96.html[Html]`
`http://www.cpsc.ucalgary.ca/projects/grouplab/papers/vis.tois-full.ps[Postscript, 2.3mb]`

71. Semio Corporation (1996) "SemioMap -- The Discovery Search Tool".
`http://www.semio.com/`

72. Sharma, R. (1995) "Open Inventor: A toolkit for 3D graphics", Discovery Publishing Group.
`http://www.unx.com/DD/advisor/docs/dec95/dec95.rajesh.shtml`

73. Silicon Graphics Inc. "MineSet Visual Data Mining Tools".
`http://www.sgi.com/Products/software/MineSet/products/vtools.html`

74. Silicon Graphics Inc. "Open Inventor Homepage".
`http://www.sgi.com/Technology/Inventor.html`

75. Silvers, R. (1995) "Livemap: A System for Viewing Multiple Transparent and Time-Varying Planes in Three Dimensional Space" in *Conference Companion, CHI'95 Conference on Human Factors in Computing Systems*, Denver, May 1995. New York: ACM.

76. Stone, M., & Fishkin, K, & Bier, E. A. (1994) "The Movable Filter as a User Interface Tool", Xerox PARC, 3333 Coyote Hill Rd., Palo Alto CA 94304.
`http://www.parc.xerox.com/istl/projects/MagicLenses/94CHIFilters.html`

77. Travers, M. (1989) "A Visual Representation for Knowledge Structures", in *Hypertext '89 Proceedings*, pp147-158.

78. Tufte, E. R. (1990) "Envisioning Information", Graphics Press, Cheshire, CT.

79. Tufte, E. R. (1983) "The Visual Display of Quantitative Information", Graphics Press, Cheshire, CT.

80. Tweedie, L. (1995) "Externalising Abstract Mathematical Models" to appear in the *Proceedings of CHI'96*, Vancouver, Canada, ACM Press.
```
http://www.ee.ic.ac.uk/research/information/www/
     LisaDir/CHI96/lt1txt.html
```

81. Tweedie, L. (1995) "Interactive Visualisation Artifacts: how can abstractions inform design?" in *People and Computers X, Proceedings of the HCI'95 Conference*, Huddersfield, M.A.R. Kirby, A.J. Dix and J.E.Finlay pp247- 265.
```
http://www.ee.ic.ac.uk/research/information/www/
     LisaDir/DIVA/DIVA.html
```

82. Waterworth, J. A. (1996) "Personal Spaces: 3D Spatial Worlds for Information Exploration, Organisation and Communication" in *Proceedings of BCS Conference on 3D and Multimedia on the Internet, WWW and Networks*, 17-18 April 1996, Bradford, UK.

83. Weiss, R., & Velez, B., & Sheldon, M. A., & Nemprempre, C., & Szilagyi, P., & Duda, A., & Gifford, D. K. (1996) "HyPursuit: A Hierarchical Network Search Engine that Exploits Content-Link Hypertext Clustering" in *Proceedings of the Seventh ACM Conference on Hypertext*, Washington, DC, March 1996.
```
http://www.psrg.lcs.mit.edu/ftpdir/papers/
     hypertext96.ps
http://www.psrg.lcs.mit.edu:80/Projects/CRS/
     HyPursuit/
```

84. Wilson, P. (1988?) "Some Consequenses of Information Overload and Rapid Conceptual Change".

85. Wood, A., & Drew, N., & Beale, R., & Hendley, B. (1995) "HyperSpace: Web Browsing with Visualisation", University of Birmingham, United Kingdom.
```
http://www.igd.fhg.de/www/www95/proceedings/posters/
     35/index.html
```