Dynamic Timelines

Visualizing Historical Information in Three Dimensions

by Robin Lee Kullberg

B.S. Environmental Science
Massachusetts Institute of Technology
Cambridge, Massachusetts
1991

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, in partial fulfillment of the requirements for the degree of

Master of Science in Media Arts and Sciences at the
Massachusetts Institute of Technology
September 1995

Signature of Author ______________________
Program in Media Arts and Sciences
August 11, 1995

Certified by ______________________
William J. Mitchell
Dean of School of Architecture and Planning
Massachusetts Institute of Technology
Thesis Advisor

Accepted by ______________________
Stephen A. Benton
Chairperson
Departmental Committee on Graduate Students
Program in Media Arts and Sciences

Dynamic Timelines
Visualizing Historical Information in Three Dimensions

by Robin L. Kullberg

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, on August 11, 1995, in partial fulfillment of the requirements for the degree of
Master of Science in Media Arts and Sciences

Abstract

This thesis considers the form and function of the visual communication of historical information in computer-based media. By applying new visual techniques derived from traditional graphic design and cinema, such as infinite zoom, translucency, and animation, the traditional timeline is transformed into a dynamic, three-dimensional framework for the interactive presentation of historical information.

I argue that current static and non-interactive presentation limit the ability of the designer to visualize complex historical information. Dynamic, interactive design solutions address the communicative goals of allowing seamless micro and macro readings of information at several levels of detail and from multiple points of view.

Experimental software for visualizing the history of photography was created to examine and compare various visualization techniques. Selected examples from the system illustrate advantages and disadvantages of various methods of dynamic visualization and interaction.

Thesis Advisor: William J. Mitchell
Title: Dean, School of Architecture and Planning
Sponsors: NYNEX and the MIT Media Laboratory News in the Future Consortium

Dynamic Timelines
Visualizing Historical Information in Three Dimensions

by Robin L. Kullberg

The following people served as readers for this thesis:

Reader

Glorianna Davenport
Associate Professor
MIT Program in Media Arts and Sciences

Reader

Ken Haase
Assistant Professor
MIT Program in Media Arts and Sciences

Reader

Ron MacNeil
Principal Research Associate
MIT Media Laboratory

acknowledgments

Thanks to my colleagues and friends at the Visible Language Workshop and in the MIT Media Lab. Thanks especially to Nancy for her patience and sense of humor in trying times, to Mike and Dave for distraction, to Yin Yin for teaching me so much, to Dave for always being there, and to Suguru for cracking the whip.

This research has taken place in the Visible Language Workshop of the MIT Media Laboratory. The Visible Language Workshop was founded by Professor Muriel Cooper. Her profound vision of the future of design at the computer interface has guided this research.

Thank you Ishantha, Lisa, Earl, Louie, and XY for being such wonderful colleagues. Love and thanks to Darcy, Ryan, Sekhar, Bruce, Mike, Mom, Dad, and my two wonderful sisters, for helping me through and for putting up with me in times of stress.

Special thanks to Bill for advising me on such short notice, and to my readers for their advice and inspiration. I would also like to express my gratitude to Justine Cassell and Leila Kinney for their time and input, and to Andrew Eskind of the George Eastman House for the wonderful database.

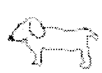

this thesis is dedicated to Suguru Ishizaki. Suguru, you have made these past two years invaluable. the words do not exist with which i could adequately thank you for your patience, your insight, your inspiration, and your friendship.

Table of Contents

1 Introduction

This thesis considers the form and function of the visual communication of historical information in computer-based media. By applying new visual techniques derived from traditional graphic design and cinema, such as infinite zoom, translucency, and animation, the traditional timeline is transformed into a dyna-mic, three-dimensional framework for the interactive presentation of historical information.

Experimental software for prototyping three-dimensional visualizations of the history of photography was created to examine and compare various visualization methods. Content for the visualizations consists of photographs from the collection of the George Eastman House in Rochester, New York. The final software implementation presents a database of 265 annotated images considered to be representative of an overview of the history of photography from approximately 1830 through 1950. This information is shown in its historical context of concurrent important events from the history of art, the history of technology, and political history.

The image database includes textual annotations which identify each image by title, author (including nationality and activity dates), date of creation, place of creation, photographic process, size, and keywords. At run time, the visualization software reads text annotations of the database of images and uses them to construct a knowledge representation, which will be described in greater detail below. Supplementary information intended to provide historical context is also read and added to the knowledge representation at run time. This contextual information consists primarily of major events from cultural, technological, and political history gathered from art history textbooks.

1.1 Problem

A timeline is an atlas of history, a map of events in time. We use timelines for some of the same reasons we use geographical maps: to locate an event in time, as we would locate a city on a geographical map; to see the time elapsed between events, as we would see the distance between two cities; to get an overview while being able to focus on detail in its correct context, as we would view a city in the larger context of its state while being able to discern information particular to the city. When examining events in time, we are not only concerned with finding the what, when, where. We also look for causal relationships. We look at other events and the historical context, and try to understand why and how.

This thesis argues that the current static and/or non-interactive media (e.g. print,film,hypertext) limit the ability of the designer to communicate complex historical information that often requires multiple points of view and several levels of detail. By re-inventing the static, two-dimensional timeline as a dynamic, three dimensional time*space*, the designer can facilitate the user's ability to access, browse, and understand historical information. This thesis addresses the following three problems in the visual communication of historical information:

1. The difficulty of presenting multi-dimensional information as a two-dimensional visual structure. This is especially problematic in a rigid two-dimensional timeline in which time is one dimension. Typographical elements and images representing events are arranged along the time axis. These graphical elements often invade the time dimension and result in crowding and confusion which obscures the visual structure of the information. Figure 1 shows part of a timeline of the history of art included in a well-known art history textbook [Janson 1991]. The labeling across the top of the diagram clearly defines the time dimension, and dark horizontal bars place artistic movements along the time dimension. The problem arises with the typographic presentation of overlapping artistic movements and the added dimensions of artists and their artwork. To retain legibility of the typographic elements, the artis-

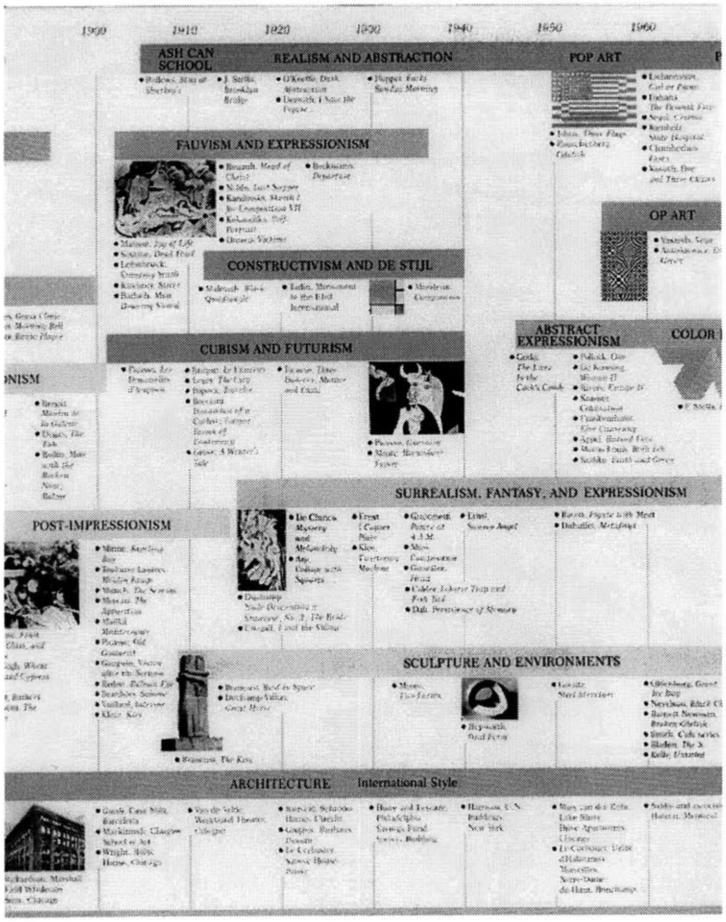

Figure 1. (left) Timeline of the history of art illustrating the problem of designing multi-dimensional information as a two-dimensional visual structure; from Janson's History of Art, *1991.*

Figure 2. (below) Timeline of the history of art illustrating the problem of presenting both overview and detail, as well as providing historical context; from Hartt's Art: A History of Painting, Sculpture, and Architecture, *1989.*

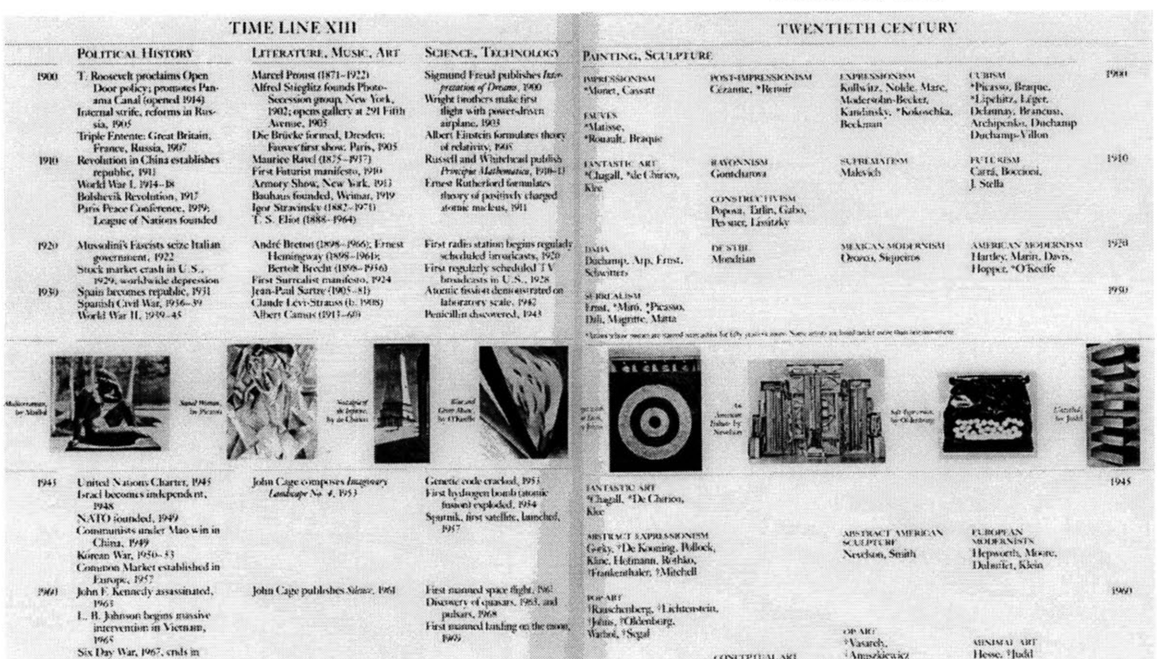

tic movements and their related artists and works extend downward along a meaningless vertical axis. The effect is clutter and an obscuring of relationships between the time dimension, the movements, and their related artists.

2. The difficulty of providing an overview of historical information while at the same time allowing access to different levels of detail. In static media, this problem is often addressed by providing the reader with several different timelines at different granularities, leaving the reader the difficult task of assimilating the events from the different time scales into a cohesive historical understanding. Another static solution is often to create a timeline with a high level of detail which would expand to fill several pages, forcing the reader to flip back and forth between views. Figure 2 is a timeline of twentieth century art [Hartt 1989] which demonstrates one aspect of the problem of overview vs. detail. In this design, all events are expressed in the same manner typographically, making it difficult to form the big picture of twentieth century art. Landmark events which are commonly referenced in art history texts, such as the advent of World War II or the painting of *Les Demoiselles D'Avignon*, are obscured or missing.

3. The difficulty of communicating the context in which historical events occur. Current trends in the study of history emphasize contextualization of events within their broader historical context [Tarnas 1991]. Some two-dimensional timelines attempt to address this problem by including contextual events in one large timeline, separating them spatially into columns or rows according to categories such as "politics" and "arts." This spatial separation makes it difficult to see relationships between events. Figure 2 also demonstrates this problem. Because events are categorized and spatially separated, important causal relationships, such as the relationship between the Bolshevik Revolution and the Constructivist movement, are not indicated.

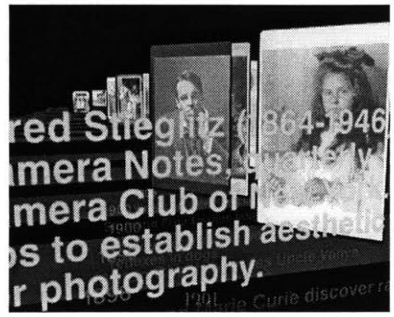

Figure 3. Scenes from photographic history visualization.

1.2 Approach

By re-inventing the traditional static timeline in a dynamic computational environment, the communication of complex, multi-dimensional historical information can be improved. Consider this scenario:

> *A user* sits down at her computer and connects to a remote database of information on the history of photography. This database is extensive, perhaps linking to other databases with more detailed information on specific topics. Upon first view, the student is presented with an overview of the information represented in the database. This overview spreads the available information out before her as a three-dimensional landscape of information through which she can travel to find the information she needs (Figure 3). The user "flies" over this overview until she comes to an area of interest. As she moves closer to the information, more and more detailed information appears about the lives and works of particular artists. Contextual information from political and social history hovers in the background, moving as she moves in order to continuously reflect the dates that she is examining.*

In this scenario, the task of the designer of this three-dimensional interactive information space surpasses the lessons learned in traditional graphic design. Now designing in time as well as space, the problem becomes multidisciplinary, suggesting an approach which combines the language of tra-

**The term "user" is intended to denote a viewer who is an active participant in the dynamic, interactive presentation of the information. The user can browse the database, or conduct a search on a particular topic of interest.*

ditional graphic design with the language of a dynamic visual form such as cinema. Additionally, the design of historical information should be informed by knowledge from the field of cognitive psychology about the way we perceive time.

The primary focus of investigation is the use of new interactive and dynamic visualization techniques derived from traditional graphic design and cinema that are unique to computers, such as infinite zoom, translucency, and animation.

Real-time computer animation can be used to create smooth visual transitions when moving in a three-dimensional graphical environment. Visually continuous transitions help the user to maintain her orientation in complex information spaces by constantly providing a visual context. This research uses a camera metaphor to describe how the user's view of the three-dimensional space is captured and animated on the computer screen. Zooming, panning, and tilting of the camera are cinematic techniques which can be applied to a virtual space to structure and inform animated transitions. As in cinema, the visual language of camera framing in time is inextricably linked to the semantic interpretation of the information being visualized.

The infinite zoom is similar to dollying a camera which has infinite depth of field. As the camera moves closer and closer to an object, it remains in focus. This technique was originated in the film *Powers of Ten* by The Office of Charles and Ray Eames, in which the viewer is taken on a whirlwind tour of the physical universe from the scale of galaxies to that of quarks [Eames 1978]. Applied to a virtual space, the infinite zoom offers the ability to present increasingly detailed levels of information through animation. Using the technique of zooming into multiscaled hierarchical information, the finite space of the computer screen becomes an infinite space with infinite resolution. In a dynamic timeline with several levels of detail, the infinite zoom offers great possibilities for moving between levels while retaining context. This thesis includes the implementation of the infinite zoom and touches upon questions raised regarding maintaining orientation when zooming into areas of detailed information.

Transparency is a technique that can be used to draw the user's attention to important areas of information by taking advantage of human visual perception. Information can be emphasized by using techniques which simulate depth. This effect is achieved by heightening the contrast between the information being emphasized and "background" information, either by blurring the background or by using transparency to fade the background [Colby 1992].

In a three-dimensional visualization of a timeline which contains perhaps thousands of graphical objects such as text and images, continuous gradual transitions between transparency and opacity and can be used to visually distinguish relevant information while retaining context. For example, a user might want to search a timeline for details of a specific event, such as a particular battle during World War II. Transparency could be used to focus attention on the specific details of the battle of interest, while the general context of World War II, though de-emphasized, would still be visible in the background. Smooth visual transitions between transparency and opacity provide the access to detailed information while simultaneously placing the battle in the correct historical context.

In the process of narrowing down the scope of this research, decisions were made regarding the nature of the information being presented. First of all, this research makes the simplifying assumption that the historical information presented is static. However, historical information is not static. Not only do databases change over time, but history itself is rewritten. The second assumption is that of objectivity. In the historical database used in experimentation, each photograph is treated equally. There is no indication of photographic movements, lines of influence, or exceptional historical significance. However, it is important to realize that any historical database is subjective by nature, if only by the exclusivity of the information it contains, and any visual presentation is equally biased. This research recognizes the highly subjective nature of the task of visualization of historical events, but attempts to minimize the editorial subjectivity of the designer. No event is presented as being more or less important than any other event, and no opinions as to issues of influence are expressed graphically.

1.3 Research Methodology

It is important to recognize the distinction between the research process and the design process. The design process is an important part of the research process; however, their goals are quite different. The result of the design process is a completed design, whereas the result of the research process does not take the form of any one design, but rather a range of guidelines and arguments that can be referred to by designers tackling similar problems in the future.

The research process begins with the identification of a set of design problems in the visual communication of information. The design problems being researched are general rather than specific, addressing broad issues in communication design. Next, the design process begins. A range of design approaches are identified, and several experimental design solutions are planned. These experimental design solutions are visualized using testbed software which was implemented for this purpose. The resulting prototype solutions are compared and evaluated as to how well they succeed in solving the identified design problems. This stage of evaluation takes the form of qualitative discussion of the areas in which the experimental design solution succeeds in fulfilling communicative goals, and the areas in which it fails. In the process of prototyping and evaluating the experimental design solutions, the designer and her critics often discover new design problems which also must be solved in order to arrive at a successful design solution. Thus, the design process is a naturally recursive process in which the continuous identification of problems plays a vital role.

In design research, the most important aspect of the design process is discussion of *why* and *how* different aspects of a design succeeded or failed, including the identification of related subproblems. The result of this discussion is a range of arguments and relevant issues which can help guide future designers. In this thesis, design prototypes were evaluated primarily by myself, my advisor and readers, my colleagues in the MIT Media Laboratory and the design community, and by a professor of the history of art.

1.4 Software

Testbed visualization software is composed of several modules coded in C++ which utilize the Silicon Graphics Performer graphics library. The Performer graphics library is a C library designed to render a three-dimensional virtual world in real-time. My testbed software allows quick prototyping by providing high level C++ code for:

1. creating and managing dynamic, three-dimensional graphical objects
2. reading and representing a database of historical information
3. dynamic manipulation of a virtual camera using cinematic language

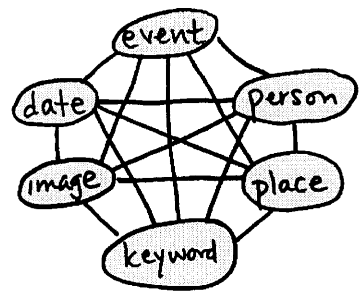

Figure 4. Data elements used in object-oriented knowledge representation of database. Any data lement can be semantically linked to any other data element.

An annotated image database is represented using a relational network data elements. Each data element can be identified as one of the following: an event object, a date object, a place object, a person object, an image object, or a keyword object. Any data element can be semantically linked to any other data element (Figure 4).

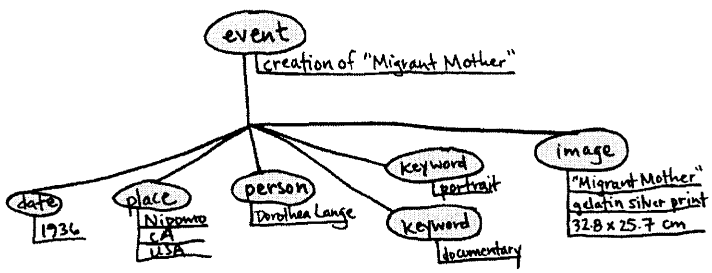

An event, such as the creation of a photograph, has a list of links from it to at least a few of these objects, depending upon the information available in the historical image database. For example, the event of the creation by Dorothea Lange of her photograph entitled *Migrant Mother* (Figure 5) is an event object which contains a list of links to the following: a person object (Dorothea Lange, the photographer), a date object (1936), a place object (Nipomo, California, USA), an image object (*Migrant Mother*, also containing information about the process used and the size of the image), and several keyword objects (portrait, documentary, etc). The various objects to which the event object is linked also contain a list of links, one of which points back to the event object. Thus, if one wanted to find all the portraits, one would follow the list of links out of the keyword object called "portrait" to the related event objects (Figure 6).

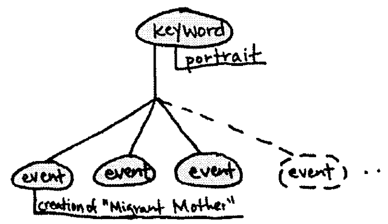

Figure 6. To find the events related to portraiture, one follows the list of links coming from the keyword object named "portrait."

1.5 Summary of Results

In a process of design experimentation using a historical database of images, this research has identified design issues and solutions which are particularly relevant to the problem of visualizing historical information in three dimensions. Results include how to provide and maintain a single, stable visual structure whose framework can allow the communication of overview, levels of detail, and historical context in response to user interaction.

2 Background

In order to approach the previously identified problems in designing a dynamic timeline, this research borrows experimental findings, techniques, and vocabulary from four different fields: information visualization, cinema, new graphical techniques, and temporal cognition.

The purpose of this chapter is twofold. First, it aims to introduce the reader to the background in which this research is rooted. Second, it establishes a common vocabulary upon which later discussion of design experiments and results is based.

Before delving into the four different fields which make up the background of this research, I would like to place this research in the current context of the increased interest in virtual museums. One possible application of this research is the visualization of a museum's collection.

The virtual museum has many advantages: you can visit a virtual museum without leaving your home; you can avoid the discomfort of a crowded gallery; you can view the entire collection of the museum rather than being limited to those works which are being currently displayed in the museum due to lack of space; you can construct a personalized path through the works that are of interest to you [Mitchell 1995].

In the dynamic environment of the virtual museum, there appear to be few limits on the richness and variety of information which can be presented. There are limits, however, to the user's perception of this information. Regardless how rich with information a space might be, the information is worthless unless it can be communicated.

2.1 Information Visualization

Information visualization helps to communicate information by displaying the structural relationships that would be difficult to discern using decontextualized information [Card 1991]. It has been shown that by visualizing information and by allowing dynamic user interactivity, complex information spaces can be more easily explored and understood [Ahlberg 1994]. Tufte stresses that in order to navigate through a complex information space, the user needs some kind of overview, or macroview onto the information, by which he can orient himself [Tufte 1990].

The Visible Language Workshop of the MIT Media Laboratory has produced groundbreaking design research regarding dynamic design and three-dimensional information visualization. This work draws heavily from the lifelong work of the late Professor Muriel Cooper, a pioneer of design for computational media [Cooper 1989]; Grace Colby and Laura Scholl's research into the use of transparency in dynamic information design (Figure 7)[Colby 1992]; Muriel Cooper and David Small's *Information Landscapes*, an interactive visualization of an abstract information space [Abrams 1994]; Suguru Ishizaki and David Small's *Typographic Space*, an early exploration into the use of typography in a three-dimensional virtual space [Small 1994]; Ishantha Lokuge and Suguru Ishizaki's *GeoSpace* project, which takes advantage of size and transparency in visualizing personalized geographical information [Lokuge 1995]; Lisa Strausfeld's *Financial Viewpoints*, which explores using point-of-view in visualizing abstract financial data in three dimensions [Strausfeld 1995]; Earl Rennison's *Galaxy of News*, a browser which uses animation and scale to allow the user to access a news database [Rennison 1995]; and unpublished work by my colleagues Yin Yin Wong and Suguru Ishizaki exploring issues of designing in space and time [Wong 1995; Ishizaki 1996].

Other related work includes the research by Steven Drucker regarding intelligent camera control in virtual three-dimensional environments such as virtual museums [Drucker 1994].

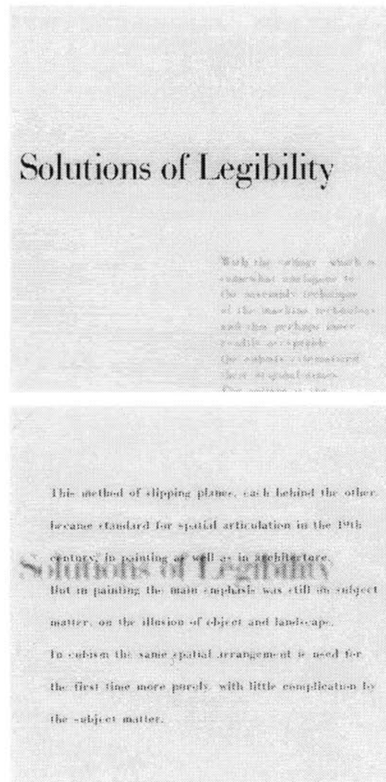

Figure 7. Two frames from Colby's interactive demonstration of the use of transparency in navigating through layers of graphical information

2.2 Cinema

The investigations of this thesis are also informed by the art of cinema. At the most basic level, both cinema and interactive virtual spaces are the dynamic presentation of a three-dimensional space on a two-dimensional screen.

Design discussion in this thesis uses many terms from the vocabulary of cinema. The metaphor of a virtual camera is used to describe the means by which the three dimensional information space is displayed on the computer screen. Later design discussion freely mixes terms from traditional graphic design with cinematic terms to describe the visual qualities of the images captured by the virtual camera. Framing, distance, angle, depth of field, focus, and point-of-view are a few of the cinematic terms which will be used to describe the quality of the "shot." Cinematic terms are also used to describe camera movement.

The language of cinema includes shot and editing techniques which enable the viewer to construct an understanding of spatial and temporal relationships by presenting a series of shots of non-abstract images which are often discontinuous in space and time [Bordwell 1990].

Important questions posed in this design exploration include how and if the language and techniques of cinema can be applied to a virtual environment containing abstract information. Two particular areas of investigation are (1) editing based on the graphic, temporal, and spatial relationships between shots; and (2) cinematographic qualities of the shot based on framing, distance, and angle.

2.3 New Graphical Techniques

With the advent of the computer and rapidly increasing processing speeds, many new graphical techniques have become available. This thesis investigates the use of the new techniques of animation, infinite zoom, and transparency in articulating complex historical information. Many of these new techniques have been explored in the research of the Visible Language Workshop described above.

Infinite zoom in a virtual space has been demonstrated in the Visible Language Workshop by Sabiston and later by Cooper, Small, and Ishizaki. Outside the Visible Language Workshop, Bederson and Hollan demonstrated a mutliscale hierarchical sketchpad in which moving from one area of information to another consists of zooming out to get an overview of the information, then zooming in to the new area of interest [Bederson 1993].

2.4 Temporal Cognition

We all use the conventions of calendars and clocks as external representations of time. We also draw upon internal, mental representations of time to construct chronologies, understand the concepts of past and future, and pick out patterns.

Mental imagery is one theory of temporal representation which has been widely described. It is evident from our external representations of time that time and space are closely related, and that time is often depicted spatially in both graphical representations and in verbal metaphor [Friedman 1990; Lakoff 1980]. Experimentation has shown that subjects use processes resembling visual perception to represent and manipulate temporal information. This idea is a core concept of my visualization, in which time is mapped to one dimension of a three-dimensional space.

A second theory which has been very helpful in describing other kinds of cognitive knowledge is that of the semantic network. Cognitive psychologists now believe that is impossible to explain time cognition without taking into consideration the meaning (structure and content) of events as they are related to internalized, prototypical events. The event is now considered the core concept in mental representations of time [Macar 1992]. Accordingly, my visualization is event-based, concentrating on discrete events rather than a more continuous view of history. This event-based approach is a core concept of my knowledge representation, which was described earlier in Section 1.4 *Software*. For example, each "photograph" from the database is represented by my software as the discrete event of the creation of that photograph. Other events from the history of photography, such as the invention of the color photograph, are also represented as discrete events in time. Spans of time are represented as a duration which is associated with each discrete event.

3 Design Experiments

This chapter describes the design experiments completed and discussion of their results. It is organized into three sections, each of which addresses one of the three design problems described in the introduction:

1. *the difficulty of presenting multi-dimensional information*
2. *the difficulty of providing an overview of historical information while at the same time allowing access to different levels of detail*
3. *the difficulty of communicating the context in which historical events occur*

In the process of approaching these design problems, each problem is broken down into subproblems; each section is organized into subsections, each of which addresses one of the subproblems in achieving the section's main communicative goal. It is important to note that although this document separates the three design problems into different sections, the author considered all three problems at once.

Visual techniques used in the experiments described below include three dimensional layout, color, transparency, framing, camera movement, and smooth visual transition through animation and infinite zoom.

3.1 Structure, Stability, and Motion

In this research, design experimentation addressing the problem of how to visually present multi-dimensional historical information in a three dimensional space began with an investigation of how to visually structure the abstract information space. The problem of providing visual structure includes the subproblems of how to visualize dense clusters of information, and how to maintain a feeling of stability when moving through the graphical information space.

It is important to recognize the difference between the visual structure of a virtual space and the structure of the information being visualized. The information being visualized has its own intrinsic multi-dimensional semantic structure. This structure is inherent in the given information and cannot be altered. The visual structure is determined by the designer. When the information is given graphical form in order to communicate the information, the graphical forms are often organized into a visual structure based on the information's intrinsic structure. A well designed visual structure can help the user understand the intrinsic structure of the information and the relationships among different dimensions of the information.

This research deliberately avoids recreating a physical three-dimensional architectural space, such as a building, gallery, or museum, in a virtual three-dimensional space. When a designer brings an architectural space into a virtual space, she brings with it its rich legacy of cultural and historical connotations. As a result, architectural metaphor often imposes meanings which are not intended by the designer. Rather than apply architectural metaphor, this research uses forms such as typography, two-dimensional images, and rectangles to visually structure an abstract information space. This abstract information space is sometimes referred to herein as virtual space, or, borrowing terminology from cinema, the scene.

3.1.1 Providing Visual Structure

This section describes experimentation regarding the problem of how to provide a three dimensional visual structure for a historical database of multi-dimensional information. Four experimental design solutions, which I will call prototypes, will be described in this section. The first two prototypes represent early attempts to provide a visual structure. The third prototype, in which information is arranged primarily by time and by person (photographer), is fairly successful in providing a visual framework. Small revisions in the third prototype lead to a fourth prototype, whose visual structure is retained for all later design experiments.

One solution was to design a space in which visual structure was constantly changing in relation to the information needs of the user. In this option, visualized historical events would dynamically reorganize in response to the user's interaction. This option was rejected from the beginning based upon my previous experiences navigating in abstract three-dimensional information spaces. I have found that it is very easy to become disoriented, especially in areas in which the visual information is less structured. A three-dimensional abstract information space lacks many of the visual cues with which we orient ourselves in the natural world. For example, in a typographic space devoid of recognizable physical objects, size of a typographic element can only be measured in relation to other typographic elements. As a result, it is easy to become lost when zooming in. Disorientation can also occur as a result of the lack of horizon or bounding walls, or from navigating in a space with no predefined "up" or "down." Because my goal was to optimize the communication of my database of information for educational purposes, I felt strongly in favor of providing a stable structure in which the user could easily navigate to explore the different dimensions of the historical information.

Early structural experiments included organizing the virtual space into a cartesian coordinate system in which each of the three dimensions represents a variable (e.g. time). Using this organization structure, the events contained in the database could then be placed in the three-dimensional space in the form of a graphical object, such as a digital image or a typographic element. Because this research is an investigation into timelines, the variable of time is always assigned to one dimension. The remaining two axes in the cartesian coordinate system can be assigned to other variables which can describe a historical event, such as the the geographic location of the event, the person related to the event, or by any of several categories into which the event falls.

In the three-dimensional virtual space described, the axes are defined as shown in Figure 8. When imagining the

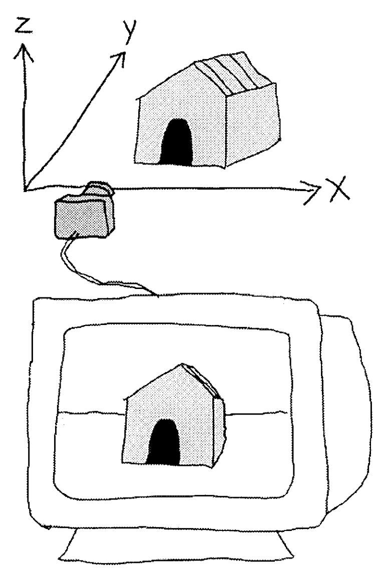

Figure 8. Definition of the axes and visualization of the camera metaphor.

space, it is helpful to think of the X-Y plane as the ground and the Z-axis as altitude. The X-axis runs across the screen from left to right, while the Y-axis is pointing away from the user into the plane of the screen. This definition of axes is common to flight simulator software.

First and Second Prototypes

Figures 9 and 10 show the first design prototype, in which time is assigned to the Y-axis. Events are arranged by their date, from oldest to most recent. In the case of a year which the database contains more than one event, additional events are spaced along the X-axis. The presentation of the events is sequential. As you move the camera forward, you are moving forward in time. Events which are farther away turn from transparent to opaque as they come within a certain range of the camera. I chose this as a beginning point because our metaphors for time include the spatial metaphor that "moving ahead" means moving to a later time, while "looking back" means looking to an earlier time.

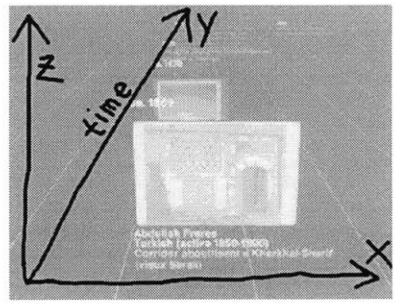

Figure 9. First prototype with definition of axes.

In the second prototype, shown in Figure 11, time is also assigned to the Y-axis; however, in this variation, the events are arranged from most recent to oldest. Photographs are arranged along the X-axis alphabetically by the name of the photograher. The presentation of the the images is sequential, as you move from left to right. The depth of the images in the Y direction gives an indication of the relative dates of the images—images which are farther away from the camera appear smaller, and thus indicate that they are older.

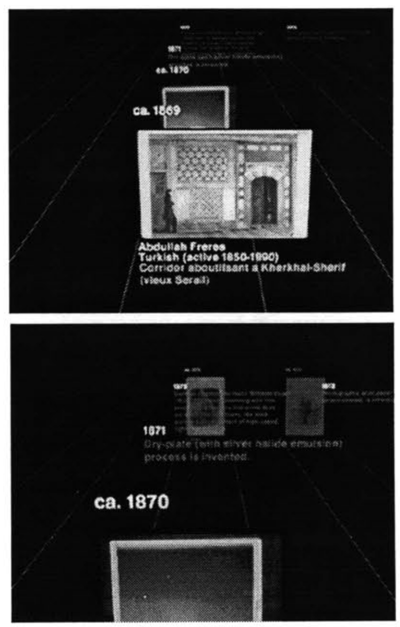

Figure 10. First prototype: these two images represent moving forward in time.

The visual structures of the first two prototypes are not completely successful. The first prototype succeeds in providing a way to quickly and intuitively browse forward and backward in time by moving the camera forward and backward. Delays resulting from drawing a great number of images on the screen at one time indicate periods of great activity in terms of the history of photography. This type of presentation enables the user to get a feeling for the overall content of the database in terms of the types of images and the periods of time it covers. However, because images are viewed sequentially over time, this organization fails to give a visual overview of the entire database. Furthermore, its visual structure is very limited—I found that I wanted to be able to see more of the intrinsic structure of this multi-dimensional information, such as geographic distribution.

The second prototype has many of the same qualities as the first prototype. Again, I found that while the sequential presentation of the images had an intuitive quality and was entertaining, it lacked a rich visual overview which could help the user understand the relationships among different dimensions of the information. For example, although images by the same photographer are juxtaposed sequentially along the X-axis, there is no other visual indication that these photographs were created by the same photographer. Furthermore, it is nearly impossible to see relationships between photographers based on geographic or temporal proximity.

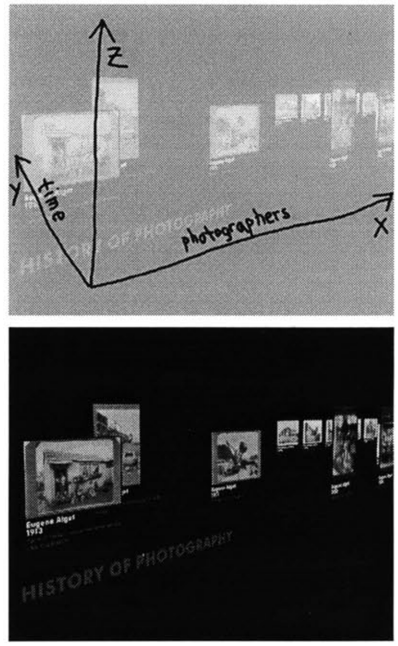

Figure 11. Second prototype: Y-axis shows time while the X-axis lists photographers.

Third Prototype

A third prototype (Figure 12) switches the time axis to the X-axis and arranges the photographers along the Y-axis alphabetically. Each photographer has his or her own timeline, indicated by a colored, partially transparent bar, within the larger timeline. Each photographer bar spans from left to right along the time axis, indicating the photographer's dates of birth and death within the larger timeline, and is placed at an angle. Events related to each photographer are presented graphically in the form of text and/or images that "hang" from that photographer's bar like clothes from a rack.

The word "hang" may seem inappropriate in an abstract space devoid of gravity; however, the arrangement of graphical objects in this prototype results in a feeling of "up" and "down," which is very helpful in orientation. One result of the first two experimental prototypes is that the way in which two dimensional images are placed in an abstract three-dimensional space defined a distinct feeling of "up/down" and "front/back." Placing images "behind" each other (Figure 13) gives the feeling that the images are "standing up" and we are looking at them as if through a camera placed on a tripod. Placing images side by side gives the feeling that the images are "laying down" and we are looking at them through a camera which is pointed down at the ground. Physical metaphor based on the body is pervasive and extremely useful in visually

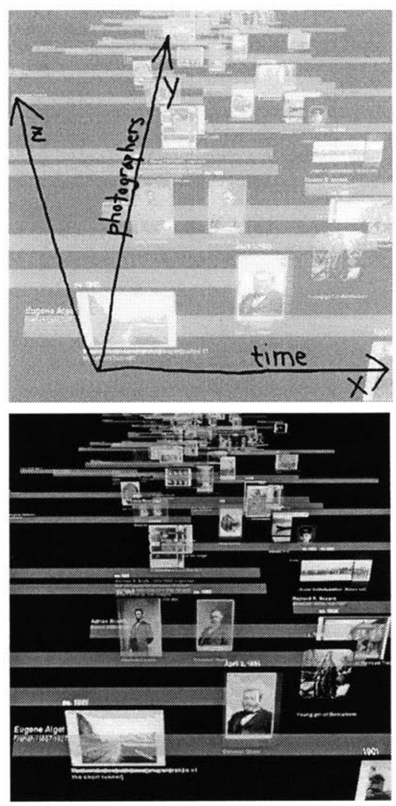

Figure 12. Third prototype: alphabetical listing of photographers along the Y-axis, with the X-axis as the temporal axis.

structuring and describing a virtual space [Strausfeld and Rennison 1995].

This prototype is more successful in providing a rich visual structure which communicates the internal structure of the information. First, because events related to the same photographer are suspended from the same bar, it is much easier to connect events related to the same person. Second, it is easier to make connections among contemporaries by scanning the Y-axis. And third, this prototype enables a visual overview of the information. It is now possible to see a contour of time which is carved out by the lifespans of the photographer represented in the database. Patterns in the information in the form of clusters of images and text are also made visible. A grid along the X-axis indicates units of ten years, and is helpful in providing a stable point of reference in terms of scale when zooming in and out.

This prototype is weak in several areas. It is still difficult to make distinctions between photographers based on their geographic location and their particular genre. Also, the hanging images are often altered in color by the overlap of the slightly transparent bars of each photographer. The bars use transparency as a technique to give structure and at the same time avoid obscuring information; however, because the clarity of the images is so important to the visualization of this this subject matter, the overlap becomes a problem (Figure 14). In approaching this problem, I experimented with limiting the camera to low angle framing of the images (Figure 15). From low angles the bars appear behind the images rather than in front. This solution has two drawbacks. First, the visual structure is designed such that all overviews are the result of high angle framing, i.e. camera placement *above* the bars looking down (the issue of providing overview will be discussed in detail in the following section). Moving from high angle to low angle through the bars was found to be unintuitive and disorienting. Second, low angle framing felt unnatural and constrained within the context of the overall virtual space. This result is not inconsistent with techniques of cinema, in which low angle framing is often reserved to serve a specific narrative function [Bordwell 1990].

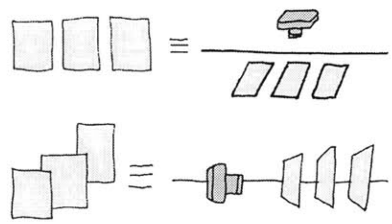

Figure 13. Camera orientation is often inferred by the spatial arrangement of objects.

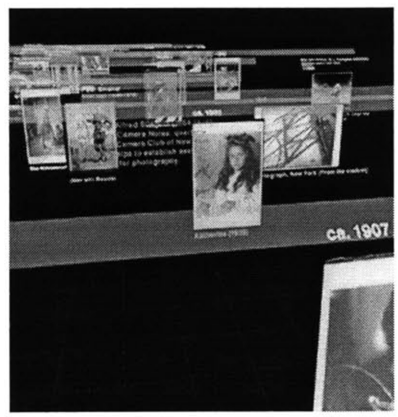

Figure 14. Obscuring of photographs by bars indicating the lives of photographers.

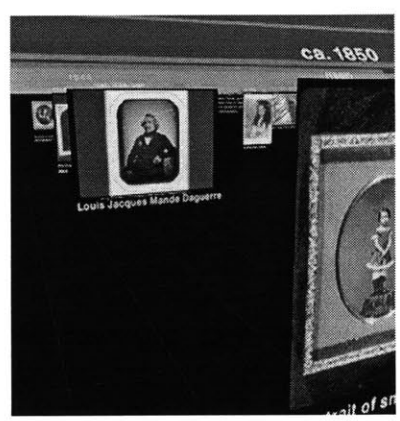

Figure 15. Low-angle view of the photographers unobscured by bars.

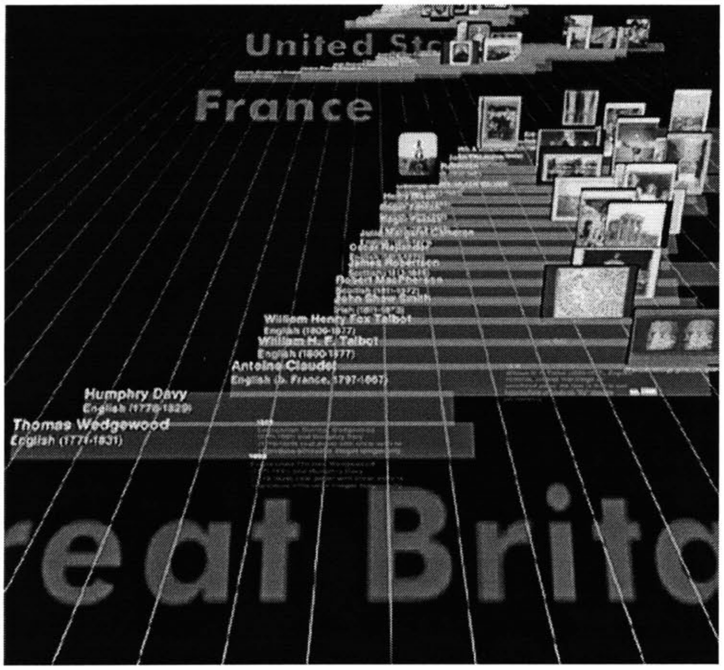

Figure 16. Fourth prototype: time is shown along the X-axis while the Y-axis shows photographers arranged primarily by the country in which they lived and secondarily by their date of birth

Fourth Prototype

A fourth prototype is very similar to the third prototype. It is improved by organizing the Y-axis by the nationality of photographers, with secondary organization within each nationality group by the photographers' date of birth (Figure 16). The problem of the bars obscuring the photographs has been remedied by standing the images on the bars rather than hanging them from the bars.

The design process of constructing a well-proportioned final visual structure included a rapid series of adjustments to the spacing and relative scale of the graphical elements. The basic unit of spatial measurement was chosen to be equivalent to one year. All graphical elements are scaled and placed relative to the images. Rather than place the images in a predetermined landscape, the landscape is built around the images.

In order to determine the absolute scale of the images, the images were experimentally placed in the three-dimensional space according to the unit of measurement chosen, and qualitatively evaluated.

In order to arrive at the final proportions and spacing of the photographers bars, many design iterations varied the width of the bars (1) relative to the size of the photographs and text placed along the bars, and (2) relative to the other photographers bars as they fit into the overall virtual space. Another important consideration in deciding the width of the photographers bars was the aspect ratio of the camera. Early iterations experimented with a 1:1 ratio; the final aspect ratio of 5:4 was chosen so that the resulting image would fit the proportion of the monitor and fill the entire screen.

Scale of typographic elements was determined in part by the semantic relationships among pieces of information. For example, detailed information,such as the nationality of a photographer, is smaller than the information it is detailing, i.e. the name of the photographer. Other factors include qualitative judgments as to visual balance and proportion in relation to other graphical elements.

The visual structure provides a framework for organizing a large and detailed virtual space. Point-of-view helps to communicate different aspects of the visual structure. Point-of-view (POV) is a term used in cinema to describe a shot which is optically and contextually subjective, i.e. a shot in which cues such as angle and framing impel the viewer to understand the shot as seen through the eyes of a particular character. I will use the term POV to refer to a view onto the scene in which cues such as angle and framing impel the viewer to understand the information according to a particular visual organizational structure or level of detail. For example, Figure 25 shows the scene from the point-of-view of nationality. Other POVs include that of a particular decade or that of a particular photographer. POV is a powerful way to convey information in a dynamic virtual space. Using POV, one stable visual structure can afford the user many different ways of understanding the visual information.

Upon first consideration, the assignment of the time dimension to the X-axis in this prototype seems less than optimal, since it does not take advantage of the spatial time metaphor for "going ahead" and "looking back." However, using the technique of POV, one can move forward in time similarly to the way one moved forward in time in the first prototype. This can be achieved by repositioning the camera so that it is pointed in the X direction, then moving the camera forward along the X-axis.

3.1.2 Visualizing Dense Information

This section addresses the problem of visualizing clusters of data without obscuring information. The visual structure of the fourth prototype (described in the previous section) includes accurate placement of photographs according to their date and their photographer. This accurate placement has the problematic side-affect of overlap between images which were created during the same year or within a few years of each other by the same photographer. This overlap, illustrated in Figure 17, is undesirable because (1) it obscures parts of images and (2) it makes clusters of images in one area appear as one image. However, the visual effect of showing clusters of photographs is desirable in that it communicates prolific periods in the life of the photographer.

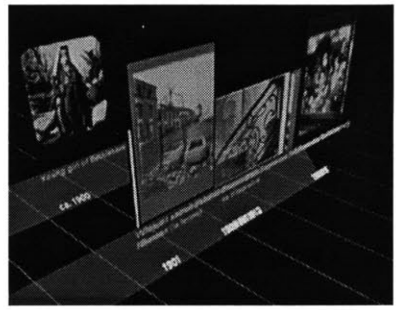

Figure 17. The problem of overlap among photographs by the same photographer.

One solution to the overlap problem is to place the images accurately along the time axis, but to space them out in another dimension. This other dimension could be the Z-axis, or even the Y-axis in conjunction with the use of parallax. This solution was rejected because it distracted from the overall visual structure already established. An alternative solution which was observed to be successful was placing the images at an angle, slightly spaced, so that the edges of the images are clear at all times, and the visible clusters of images more accurately convey areas of density (Figure 18).

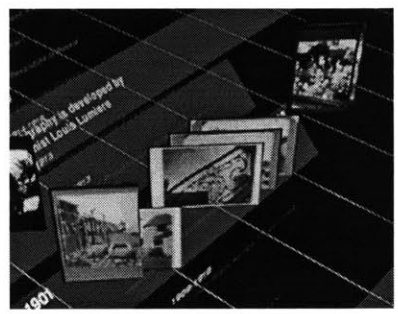

Figure 18. Solution to the overlap problem in which photographs are displayed at an angle.

However, placing images at an angle still hides large parts of some photographs. A solution to this problem is to allow the user to interact with photographs in which she is interested. Clicking on the edge of a particular photograph to express interest causes the photograph to slide out and display itself with obstruction to the viewer (Figure 19).

This approach brings out another interaction question—how do the images return to their original placement? Does the photograph simply pause for a few seconds and then slide back into its slot automatically, or does the user need to activate its return by clicking on the image? Each approach has advantages. In the first approach, the viewer is spared the trouble of active response. However, the viewer is troubled to go through the process again if the time during which the photograph is displayed is insufficient. In the sec-

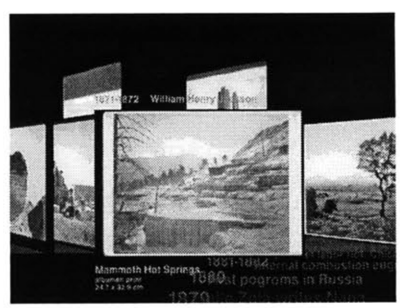

Figure 19. A photograph displays itself for the user, complete with detailed information about its photographer, date, process, and size.

ond approach, the user has the chore of putting each photograph away when she is finished viewing it; however, her pace can be leisurely. Furthermore, the selected photograph can function as a bookmark for later reference, or a history of her viewing of the database.

3.1.3 Providing Stability

This section examines the problems related to providing stability in a dynamic environment.

The visual structure of the fourth prototype is well defined. However, in the course of interacting with this space by manipulating the virtual camera, several people who evaluated the work noted that the graphical objects which comprised the visual structure often gave the impression that they were "floating" in the virtual space (Figure 20) The floating graphical objects were not perceived as anchored or grounded in one particular location.

This lack of visual stability of the graphical objects is a serious problem in two ways. First, it undermines the carefully constructed visual structure. Second, it allows motion to take on meanings which were not intended by the designer, leading to miscommunication and breakdown of visual structure.

If graphical objects are perceived as being free to move about a virtual space, then the motion of the camera over a stable structure of graphical objects can be interpreted as a stable camera viewing moving objects. I will call this *the*

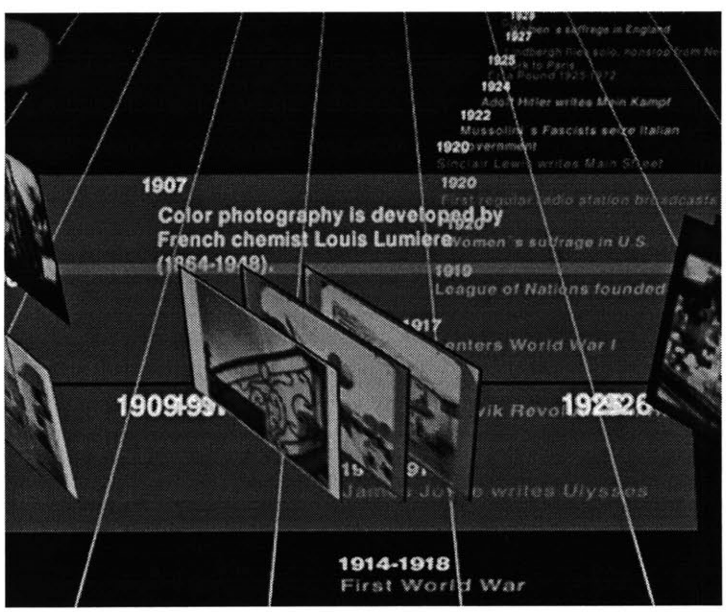

Figure 20. The problem of lack of stability: spatial differences are difficult to understand and graphical objects appear to "float."

problem of relative motion (i.e. motion in which the visual result is identical whether it is the camera or the object which is moving).

How can the problem of relative motion be avoided? In an abstract virtual space, we cannot use cues from the natural environment to determine whether the camera is moving or the graphical objects are moving. Instead, we rely on our expectations. For example, the previous section described the animated movement of a photograph in response to a mouse click. When the user clicks on the photograph, there is an expectation that the click will cause the photograph to respond to the interaction. When the photograph slides out, there is not even the slightest doubt that the photograph is moving while its surroundings are stable, instead of the inverse. By giving a graphical object the visual quality of stability, the designer can lead the user to expect that the object will remain stable and to interpret relative motion as movement of the camera.

How can the designer give the visual quality of stability? Experimental design solutions included the following approaches:

1. Defining a stable floor indicated by two-dimensional grids/lines and/or shadows of graphical elements (Figure 21)
2. Orienting important structural elements along an invisible three-dimensional grid
3. Using color to help indicate depth
4. Framing (dependent upon how many graphical objects are in the frame)
5. Providing a heads-up display with crosshairs and/or text
6. Highlighting the area which is the focus of attention

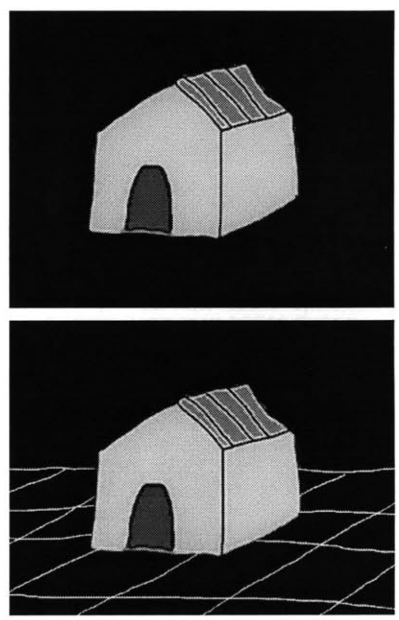

Figure 21. Grounding an object. In the top illustration, the doghouse floats in space. in the bottom illustration, the doghouse is grounded on a floor by adding a two-dimensional grid.

The first solution, defining a floor, was observed to be effective when moving the camera at right angles to the lines indicating demarcations along the time dimension. In this particular implementation, in which the Y-axis does not represent any quantifiable dimension, a grid would be inappropriate. However, as a purely visual solution, it was briefly implemented and found to be successful. I was not able to create a prototype which used shadows to indicate a floor. Unfortunately, the rendering of shadows is not supported by the Performer library at this point in time and is extremely costly in terms of processing.

The second solution of utilizing an invisible three-dimensional grid was implemented by simply rotating the bars indicating the lives of the photographers 45 degrees, so that they lay flat, parallel to the floor. Compare Figure 20 and Figure 22. In Figure 22, the bars feel much more stable and thus are far more successful as structuring elements.

The third solution, using color to help indicate depth, is also demonstrated by comparing Figure 20 and Figure 22. Both solutions show typographic elements representing historically contextual events which are placed between the images and the floor. In the Figure 20, the placement of these typographic elements is unclear. Because of their brightness, they seem to hover in the foreground. Figure 22 shows that this effect can be avoided by choosing colors of lower saturation and value.

The fourth solution involves controlling camera movement to reduce the effects of relative motion. It was discovered that the feeling that the graphical objects were moving was at least partially dependent upon how many graphical objects (especially structuring elements such as the photographers' bars) are in the frame at the time. When the camera is

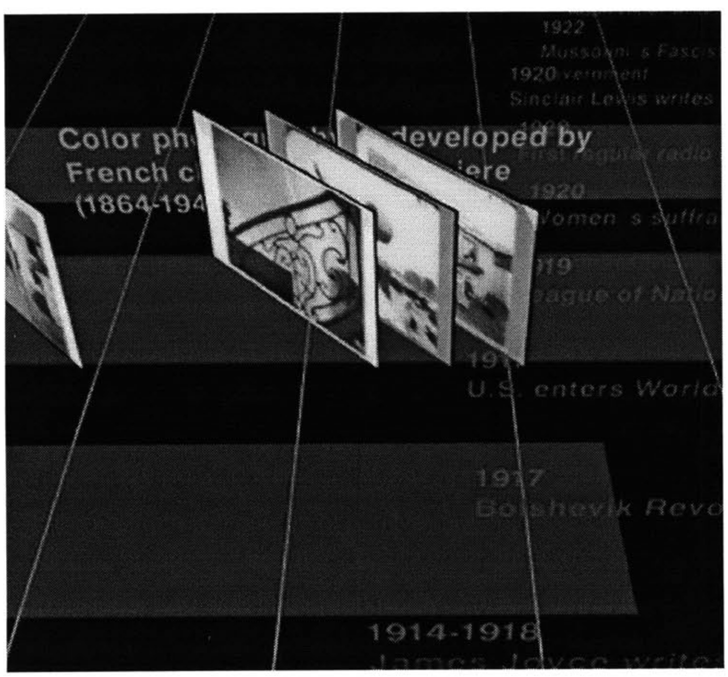

Figure 22. Increased visual stability. Color is used to better indicate depth, and bars are laid parallel to the ground.

placed at a closeup range at which very few graphical objects appear in the frame, the structural elements are less effective at conveying structure and the graphical objects seem less stable. When the camera is placed at a medium range at which many graphical objects appear in the frame, the structural elements seem more effective. This could be due to an expectation that the movement of a few graphical elements is far more likely to happen than the movement of a great number of visually structured graphical elements.

The fifth solution, providing a heads-up display with crosshairs and/or text, is borrowed from the realm of video games and flight simulators. The use of crosshairs (Figure 23) was not found to be particularly helpful in avoiding relative motion problems, and was furthermore found to be distracting. The use of text to indicate the year in the center of the screen was, on the other hand, found to be extremely helpful in orientation. As the user moves, the text indicating the year is continuously updated to reflect her current location in time.

The sixth solution, highlighting the center of attention, was implemented in the form of perpendicular bars which ran parallel to the floor and which crossed in the center of attention. This solution, shown in Figure 24, did little to ease the problem of relative motion.

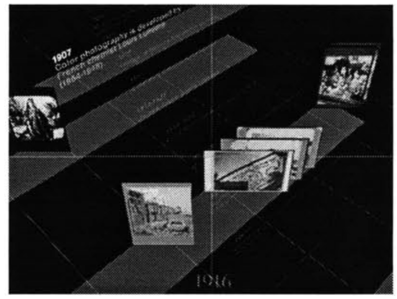

Figure 23. Experimenting with crosshairs.

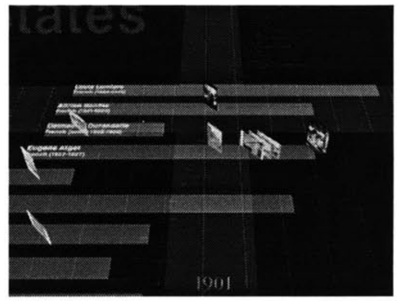

Figure 24. Experimenting with bars to indicate center of attention.

3.2 Overview and Detail

The problem of providing overview while allowing easy access to more detailed levels of information is approached by visually filtering the information. In this solution, all the graphical objects which represent different levels of detail exist simultaneously. By visually filtering the information, the designer can avoid information overload and support smooth transition between overview and detail. This section describes a subtle range of experimental solutions designed to strike a delicate balance between visually indicating the possibility of accessing infinite amounts of information, and providing clear overviews of the visual structure and the different dimensions intrinsic to the information.

3.2.1 Providing Overview

In approaching the problem of providing a visual overview, I began by applying metaphor. Taking the word "overview" literally, I designed the visual structure in such a way that positioning the camera at a high angle, looking down over the information, would result in an uncluttered (filtered) view of the visual structure showing the photographers and countries represented in the database (Figure 25). This approach takes advantage of the fact that, when viewed from the top, flat vertical elements become invisible. By placing the bars indicating the lives of the photographers at a 45 degree angle to the X-Y plane (Figure 26), I intended to allow the user to view the bars both from the top, to get an overview uncluttered by photographs and more detailed typographic elements, and from the front, to see the photographs and more detailed typographic elements. As the circular movement of the camera from a top view to a front view is animated on the screen, there is a smooth transition between two levels of detail as the photographs gradually become more visible.

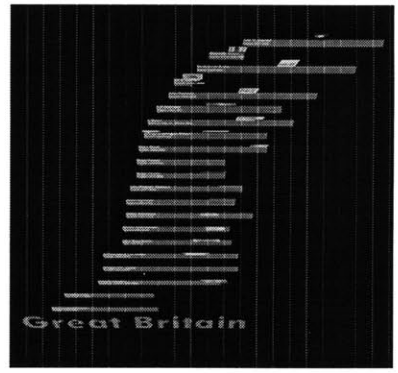

Figure 25. Overview of the photographers active in Great Britain who are represented in the database.

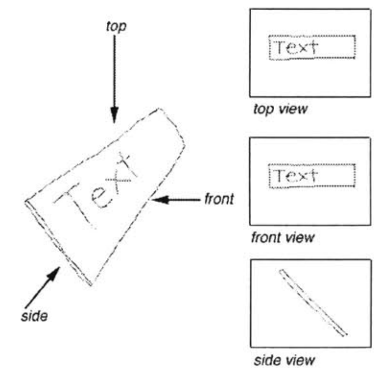

Figure 26. Bars placed at a 45 degree angle can be viewed both from the top and from the front.

This type of overview was observed to be successful in communicating information about the photographers and countries represented by the database. Especially successful is the smooth transitioning from overview to more detail through animated camera movement.

Other types of overview which it is desirable to communicate include overviews of photographers by their genre, of photographs by their process or subject matter, or of a particular year or decade. How can the chosen visual structure support these other types of overviews? One solution is to alter the visual structure in response to user query. Another solution is to maintain one visual structure and use visual filtering techniques to highlight the pertinent information. Because of the strong emphasis this research places on providing and maintaining visual structure as an aid to orientation and comprehension of the information in context, I chose to focus on exploring the second solution.

3.2.2 Filtering

This section explores the use of different visual techniques for filtering information in order to avoid visual overload and to provide different types of overviews. Three experimental solutions were identified:

1. use of transparency and resolution in drawing a graphical object as a function of its distance from the camera
2. use of point-of-view/framing to reduce the amount of information at a given time
3. use of color and transparency to highlight information
4. use of temporal presentation to form sequences of the information

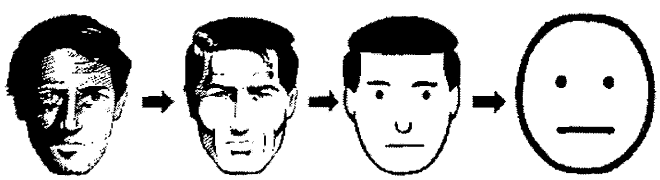

Figure 27. Four levels of abstraction, moving from less abstract to more abstract (left to right). Taken from McCloud's Understanding Comics.

The first solution uses a technique which I will refer to as level of abstraction. Using this technique, a graphical object is drawn at a level of abstraction that is a function of its distance from the camera. Figure 27 illustrates the implementation of four levels of abstraction. As an object moves farther away from the camera (or as the camera moves farther away from the object), the object is drawn at a higher level of abstraction (Figure 28). This increases rendering speed and parallels human vision. To avoid distracting visual artifacts when transitioning from one level of abstraction to the next, there is a transition zone in which the two levels of abstraction are smoothly blended. One way in which this research uses level of abstraction is in the presentation of the photographic images. When the camera is close to a photograph, it is presented at its highest resolution. As the camera moves away from the photograph, the photograph smoothly transitions to a lower resolution version of the same image. Not only does this reduce the rendering time, but it reduces the amount of information presented to the user and helps to focus the user's attention.

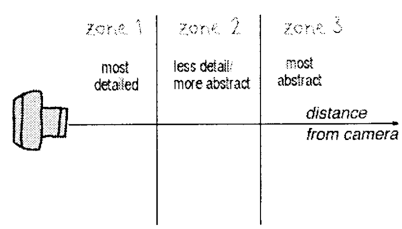

Figure 28. The level of abstraction of a graphical element is determined by its distance from the camera.

Early experiments applying level of abstraction to typographic elements used three zones. The first level presented the typographic element at its highest resolution, as an anti-aliased image. The second level presented filled polygons in the shape of the individual letters. The third level presented rectangles to indicate whole lines of text. This approach proved slightly distracting, as the visual transitioning was a bit clunky and slow in terms of rendering time. A second round of experiments tried using two levels of abstraction in which the first level was the highest resolution anti-aliased font and the second level was invisible, i.e. the typographic element gradually became completely transparent. After carefully adjusting the distances at which different typographic elements transitioned between the two levels of abstraction, I found this approach to be successful. Not only did it greatly reduce rendering time, but it also provided a means by which I could filter out individual typographic elements when providing an overview, and smoothly transition into more levels of detail as the camera moved closer an area of more detailed information.

The second experimental solution to the problem of visually filtering the information is the use of point-of-view to reduce the amount of information at a given time. One example of this solution, literally manipulating the camera to look down over the visual structure, was described in the preceding section and shown in Figure 25. Another example, used to give overview of one particular photographer, is to move the camera so that it frames information about that photographer and minimizes other extraneous information.

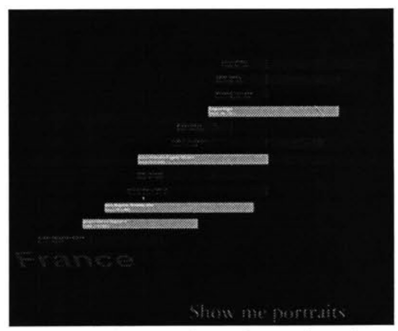

Figure 29. An overview of portraiture in France as represented by this database.

The third solution, using color and transparency to highlight information, was applied to providing an overview in response to user query. When the user queries the system for a particular topic, graphical elements which are not related to that topic animate a change in transparency from 100% opacity to 20% opacity. Thus, the graphical elements that correspond to the information sought by the user stand out from the others by virtue of their brightness. Then, by moving the camera into the default "overview" position, the user can get an overview of the particular area of interest while at the same time still operating within a consistent overall visual structure (Figures 29 and 30).

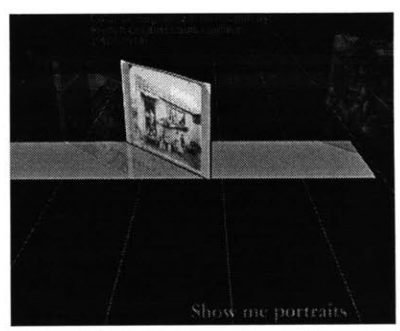

Figure 30. An overview of portraiture highlights Atget's one portrait that can be found in this database.

The fourth experimental solution, using temporal presentation, was also applied in response to user query.

Temporal presentation is a sequential presentation of the graphical elements which are indicated as pertinent to the user's query. This temporal presentation can be achieved by (1) a guided tour in which the camera continuously weaves its way through the virtual space so that it traces a path among the pertinent graphical elements, pausing briefly at each; (2) a sequence of shots of the pertinent graphical elements; or (3) a combination of (1) and (2), in which a guided tour is conceived as a sequence of animated camera movements.

I experimented with several variations of each of these three methods. For example, one variation involved presenting the highlighted graphical elements in chronological order, while a different variation presented them in order of their positioning in the virtual space from front to back, left to right. Another variation involved moving the camera to different types of overview positions (which can be referred to in cinematic terms as medium shots, medium long shots, or long shots) before diving in and showing each highlighted graphical element. This overview could be shown by animating the camera movement to the medium long shot, or by cutting to form a sequence. Extending this, I experimented with moving the camera to an overview of each country before diving in and presenting all of the highlighted graphical elements associated with that country. This can also be shown by animating the camera movement or by cutting.

Each solution had its own strengths and weaknesses. Variations in which the camera weaves through the photographs were strong in providing local context, but sometimes it was difficult to place particular photographs in the larger context of the database. While weaving through the photographs to present them in chronological order is extremely helpful in communicating change over time (for example, how portraiture changed over time), it is weak in communicating nationality and it sometimes involves an excessive amount of camera movement depending upon the particular query.

Cutting between photographs presents the materials quickly, but at the expense of the communication of context. This approach might be suitable for a user who is familiar with the contents of the database. Combining sequences of close-up shots of individual photographs with cuts to overview positions was observed to be better at communicating context. Given familiarity with the virtual space and its structure, cuts to overview (or establishing) shots were not observed to be awkward or disorienting.

Combinations in which the camera alternated among weaving through the photographs and moving to overview positions were observed to be better at providing a more global context, but quickly became slow and tedious.

A solution which could balance these tradeoffs is to allow the user to choose the style of presentation which is most appropriate to her goals.

3.2.3 Showing Detail

This section explores the problem of how to present more detailed information in context. In this section I will often refer to the level of detail or level of specificity of the information. The level of detail of information refers to the intrinsic organization of the information based upon its relative specificity. For example, a chronology may contain information on the levels of detail of millennium, centuries, decades, years, and days. Information about the Ice Age is on the level of detail of millennium, whereas information about the Industrial Revolution is on the level of detail of centuries, and last year's Earth Day is on the level of detail of days. This is not to say that one event may not appear in several levels of detail; however, when presenting a huge database, control of the the semantic detail of the information being presented at any one time is vital.

Many commercial multimedia products, such as history CD-ROMs, present information from a single database. A common means of accessing more detailed information on a particular topic uses point-click interaction to trigger a hyperlink jump to a new, static graphical presentation of more detailed information. This type of hyper-jumping to different levels of detail is problematic in that its visual discontinuity makes it difficult to understand the context of the detailed information (i.e. what is the broader topic to which this information is related? how detailed is this information?) Furthermore, point-click interaction is usually multi-purpose in nature, i.e. the action of point-clicking does not always result in the presentation of detailed information. The experiments in this section search for a new way to navigate through information which is organized hierarchically by its level of specificity.

The possibility of linking to remote, specialized databases by different authors to access more detailed information currently exists and will become more common in the future. This possibility raises visualization questions: How can information from remote databases be visualized within the context of the main database? How can this extra information be visualized in a way which allows the user to easily navigate through the different levels of detail? Again, avoiding disorienting hyperjumps in favor of smooth visual transitions is highly desirable.

My approach to visualizing the movement between levels of detail is to combine scale and infinite zoom to allow smooth visual transition (Figure 31). This approach was inspired by the Information Landscape research of the Visible Language Workshop [Abrams 1994]. One of the most compelling aspects of the Information Landscape is its use of scale and infinite zoom in presenting hierarchically nested information about the Media Lab. The Media Lab is organized into three sections, each of which contains several research groups. Each research group contains faculty, students, and staff. When looking at graphical information at one level of this hierarchy, one can see the next level of the hierarchy as an area of tiny graphical elements. As the camera moves into that area, the elements grow increasingly larger until they become legible and fill the frame. This process is repeated for all the levels of information available.

Figure 31. Two views of detailed information about the photographer Eugene Atget, including a critical essay, a timeline of his life, a portrait by Berenice Abbott, and photograph of his studio.

Why is this approach is so compelling? Detail is a small or subordinate part of a whole. In physical objects, detail is literally smaller in relation to the whole. When visualizing detail in an abstract manner, we apply what we know to be true physically. Metaphor reflects this—the act of examining something in more detail is often verbalized as "Let's take a closer look" or "Let's dive in."

One way in which the Information Landscape is less successful is that it fails to convey which level of detail the user is currently experiencing. Both viewers and navigators often become disoriented. This disorientation is easily cured by pulling back the camera. I propose that there are two basic reason for this disorientation: (1) lack of visual structure (addressed in section 3.1), and (2) visual homogeneity across different levels of detail.

By visual homogeneity I mean that there are no visually distinguishing features by which a user can easily identify or approximate her current level of detail. To solve this problem, I suggest using visual modes in the form of a gradated color scheme to indicate levels of detail. For example, when presenting the lowest level of detail, the overview, the background is black. As the user moves into areas of more detail, the background color shifts to increasingly lighter shades of gray, giving at least a rough approximation of the scale of the graphical elements being presented.

3.3 Historical Context

Section 1.1 discussed the importance of presenting events within their historical context. This section explores how to provide and maintain a dynamic historical context.

3.3.1 Visualizing Historical Context

The events that form the historical context of this visualization are *not* part of the photographic history database. Instead, they constitute a second database of major world events which was created for the purpose of providing historical context for this visualization. As such, they constitute a subjective world history which represents the views of the art historians from whom they were gathered, and my own views as editor. In visualizing contextual events, I began from the premise that it is important to visually distinguish between the different sources of information and their inherent subjectivities.

In deciding how to incorporate these contextual events in the already visually structured space, I took a metaphorical approach and located these "background" events literally in the background. In addition to simply placing the events in the background, careful attention was applied to make the events look as if they were in the background.*

* *this is described in Section 3.1.3*

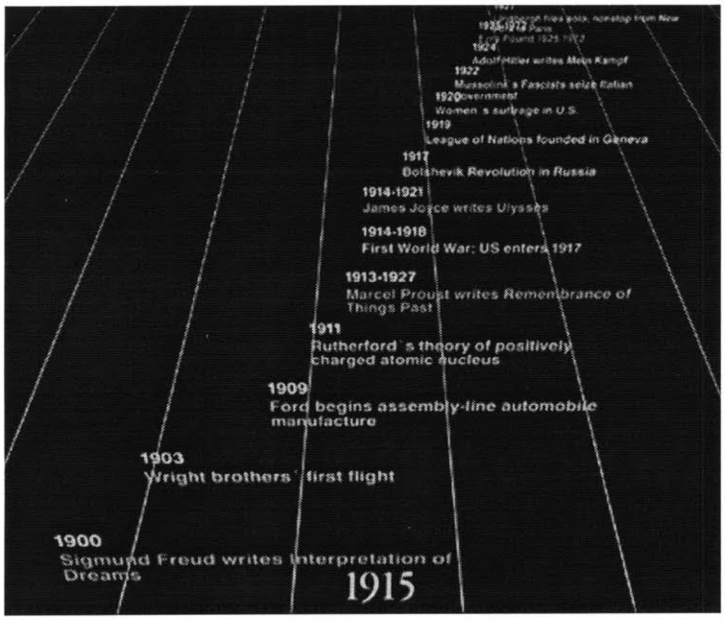

Figure 32. Historical context of the year 1915. Events and their associated dates are expressed typographically and color-coded to indicate category. Graphical elements are arranged along the X-axis by date and spaced along the Y-axis to avoid overlap.

I present contextual events as a chronologically arranged series of typographic elements composed of the year and the description of the event (Figure 32). Each event is located along the X-axis (the time axis) according to its date of occurrence and at regular intervals along the Y-axis to avoid legibility problems due to overlap. Rather than separate events spatially into their different categories of science and technology, literature and culture, and politics, I differentiate among types of events solely through color-coding.

3.3.2 Maintaining Historical Context

Because the database presented is so large and the virtual camera is so dynamic, it is impossible to place the contextual events in such a way that they are always visible (and legible) in the background. Rather than conceive the contextual events as anchored in one location, I imagined them as being able to dynamically relocate themselves in response to the user's interaction. I will refer to this visualization technique as *sliding context*. In this visualization, the typographic elements representing the contextual events slide along the Y-axis as the user moves the camera, thus continuously maintaining the appropriate historical context.

For example, if the user's attention is centered upon the year 1910, the contextual events presented will reflect the year 1910 and the years immediately before and after 1910.

As the user's attention shifts to 1915 (i.e. as the user pans right), the contextual events will transpose themselves so that they reflect the years around 1915. Figure 33 shows three frames which illustrate the shift in historical context that occurs as the user moves from 1910 to 1920. By following the transposition of the First World War across these three frames, one can see how the context slides into place in response to the user's interaction.

What form does this movement take? I experimented with (1) animating the movement of the contextual events, and (2) updating the position of the contextual events so that they simply appear in their new position.

The first solution, animating the movement, seemed to be the obvious choice in that it provided a smooth transition and better communicated exactly how the events were transposing. Experimentation showed that although animation did indeed achieve these goals, in the long term it was deemed unsuccessful. The motion of the background distracted user's attention away from the foreground information and imbued the background information with undue importance.

The second option, updating the position of the contextual events so that they simply appear in their new position, seemed to confuse some users at first. However, after a few seconds of interaction, most users understood that the contextual events were updating in response to their camera movements, and preferred this option over animation.

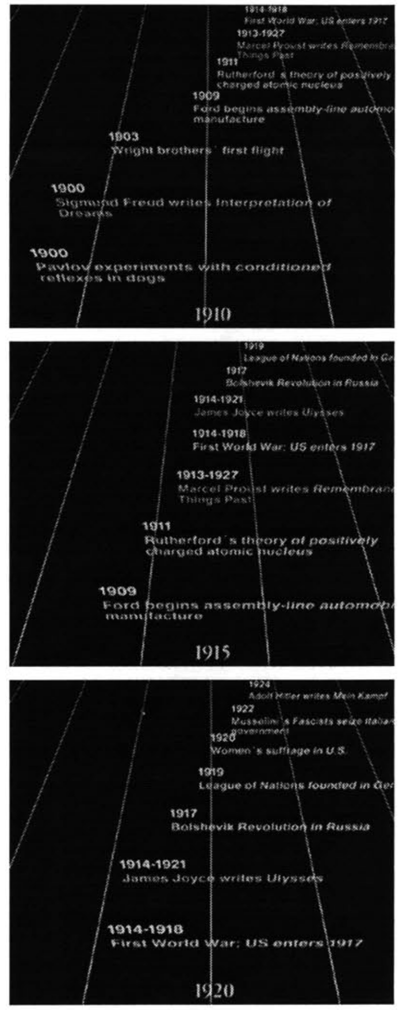

Figure 33. Three frames showing the shift in historical context as the user moves from 1910 to 1920.

Each historical database has its own intrinsic information structure and must be visualized with its own particular communicative goals in mind. This research experimented with the visualization of one database of text and images from the history of photography in order to identify and to discuss design issues, and to provide general guidelines for designing the interactive visualization of historical information.

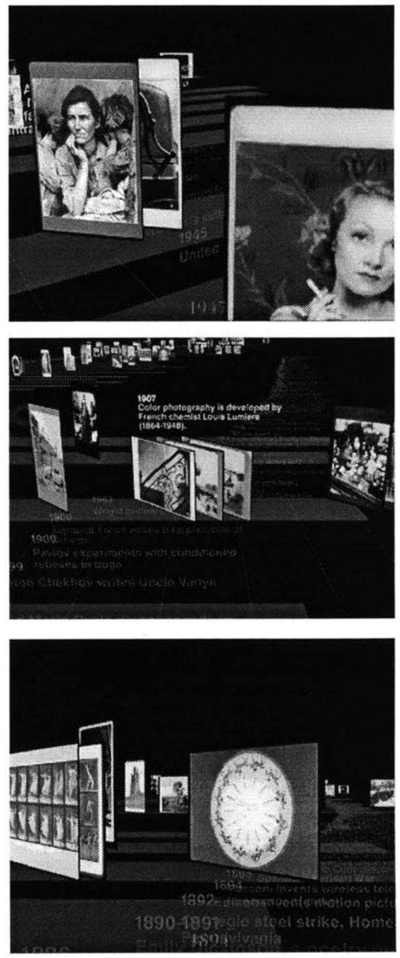

Figure 34. Screen images from the photographic history visualization.

4.1 Design Issues and Solutions

This research identified the following design issues which were found to be particularly relevant to the three-dimensional visualization of historical information:

1. providing visual structure in a three-dimensional interactive environment
2. stability, or maintaining visual structure in a dynamic environment
3. articulating point-of-view through dynamic visualization
4. using visual filtering techniques to provide overviews
5. using scale to communicate levels of detail
6. providing historical context

This research has shown that the design of a visual structure is vital to the successful communication of information in an abstract virtual space. The design of the visual structure must communicate the intrinsic structure of the information and also provide a framework which allows the interactive communication of overview, detail, and historical context. Furthermore, the designer of the visual structure must also provide a feeling of stability given the dynamic nature of the environment. Useful approaches to designing a stable visual structure include defining a floor to help the user orient herself in the three-dimensional space; aligning important structural elements to an invisible three-dimensional grid; and the use of color to help indicate depth.

This research explored techniques for communicating different aspects of the information while maintaining a single stable visual structure. Useful techniques for bringing out different aspects of the information in response to user query include the temporary use of selective transparency to fade away less relevant graphical elements, and the use of framing and sequencing in the temporal presentation of selected graphical elements.

Techniques which were found to be useful in allowing the user to move among different levels of detail of the information include level of abstraction and infinite zoom in conjunction with the use of relative scale (i.e. a graphical element's size is mapped to its the level of detail, so that the higher level of detail it is, the smaller it appears).

In the communication of historical context, this research proposes the "sliding context," a group of important contextual events which is positioned in the background and continuously repositions itself to reflect the years currently being examined by the user.

This research has shown that spatial metaphor can be useful in designing an abstract three-dimensional information space. The camera metaphor was found to be crucial in the design of this large, dynamic space. Use of the camera metaphor engendered the use of cinematic language, both verbal and visual, which were found to be central in describing, designing, and evaluating the visualization.

4.2 Future Work

This section addresses future work which would continue and expand the investigations begun by this research. The first subsection, 4.2.1, touches upon further visualization work and work in other areas of design research. The second subsection, 4.2.2, poses interesting questions when considering expanding the scope of this research to include issues of subjectivity. The third subsection, 4.2.3, proposes applications and other areas of research which could utilize the findings of this thesis.

4.2.1 More Design Experiments

The visualization was designed and run on a Silicon Graphics Onyx workstation with Reality Engine 2 graphics hardware. Although this equipment is the current state-of-the-art in terms of three-dimensional graphics rendering, there are still performance limitations in terms of rendering speed. The design of this visualization was constrained by the number of graphical elements (texture-mapped polygons) that can be drawn simultaneously. In future explorations on faster hardware, I would like to experiment with presenting even more graphical information in order to experiment with more complex data. Furthermore, given faster hardware I would like to implement visual techniques such as the simulation of focus and blur, which can help to direct the user's attention, and the use of shadows to help anchor graphical elements.

Another hardware-related constraint was the lack of audio. In future research, the addition of sound is an extremely rich area for exploration. Audio cues could be used to aid in orientation in the three-dimensional space, and audio narration could add editorial or contextual information. For example, the voice of the curator of this photographic database could describe particular images of interest as the user was being taken on a guided tour of the database.

I would also like to spend more time designing ways in which temporal connections could be made across space. For example, events such as the invention of color photography are depicted only once in this database and is spatially separated from many of the photographs. A better solution would be to somehow give a visual indication of this singular event across space.

There are several other areas which can provide interesting topics for future research. Although this research focuses on visualization, another very important (and problematic) area is interaction. Allowing the user to easily navigate in an abstract three-dimensional environment is not a trivial task. Future types of interaction which are of interest to the author include multi-modal approaches combining speech with gesture.

Another interesting question is how to allow the user to take notes or save photographs and/or text in the context of this three-dimensional virtual space. The ability to select and place images side-by-side on a virtual "light table" could be very helpful in comparing and contrasting individual works which are separated by time and space in the visualization.

Other related areas include applying automatic layout principles to generalize the visual structure across different databases with different information structures; visualizing links to remote databases; and examining issues related to the weaving of story-telling and narrative into the visual structure.

4.2.2 Expressing Subjectivity

As described in Section 1.2 *Approach*, the process of narrowing down the scope of this research included making decisions as to the nature of the information being presented. Some of the most interesting questions which could be addressed in future work arise from reconsidering the problem without making these simplifying assumptions.

The research presented in this thesis makes the simplifying assumption that the historical information presented is static. However, historical information is not static. Future areas for investigation include: How can we visualize dynamic historical information? How can we visualize a historical database as it changes over time?

In narrowing the scope of this research, it was also decided that the visual presentation of the database information should be approached as objectively as possible. An important question for future research is, How can a visualization reflect subjectivity and the different points of view of different schools of thought in terms of influence and classification of artists and events into movements?

4.2.3 Applications

Applications which could utilize the results of this design research include interfaces for visualizing historical databases or image collections, virtual museums (described in Chapter 2), or interactive documentaries in which documentary-style presentations complete with audio narration could take place within a virtual space which could be also be explored interactively.

6 References

Abrams, Janet. "Muriel Cooper's Visible Wisdom," *I.D.* Sept./Oct. 1994.

Ahlberg, Christopher, and Shneiderman, Ben. "Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays." *Proceedings of CHI'94*. 1994.

Bederson, Benjamin B. and James D. Hollan. "PA3D: A Multiscale Hierarchical Sketchpad." *UIST'93.* 1993.

Bordwell, David and Kristin Thompson. *Film Art: An Introduction*. Third Edition. New York: McGraw-Hill, 1990.

Card, Scott K., George G. Robertson, and Jock D. Mackinlay. "The Information Visualizer, an Information Workspace." *Proceedings of CHI'91*, 1991.

Cassell, Justine. MIT Media Laboratory. Conversation and correspondence, 1994.

Colby, Grace, and Laura Scholl. "Transparency and Blur as Selective Cue for Complex Information." *Proceedings of SPIE'92*. 1992.

Cooper, Muriel. "Computers and Design." *Design Quarterly*, Volume 142, pp. 22-31. 1989.

Drucker, Steven M. *Intelligent camera control for graphical environments*. Ph.D. Thesis at the Massachusetts Institute of Technology, 1994.

Eames, the Office of Charles and Ray. *Powers of ten; a film dealing with the relative size of things in the universe and the effect of adding another zero*. [Motion picture]. Santa Monica, Calif.: Pyramid Films, 1978.

Eskind, Andrew. George Eastman House, Rochester, New York. Personal communication. 1994-1995.

Eugène Atget. The Aperture history of photography series, volume 14. Millerton, N.Y.: Aperture, 1980.

Friedman, William. *About Time: Inventing the Fourth Dimension*. Cambridge: MIT Press, 1990.

Gardner, Helen, Horst de la Croix, Richard G. Tansey, and Diane Kirkpatrick. *Gardner's Art Through the Ages*. 9th ed. San Diego: Harcourt Brace Jovanovich, 1991.

Hartt, Frederick. *Art: A History of Painting, Sculpture, and Architecture*. 3rd edition. New York: H.N. Abrams, 1989.

Heibert, Kenneth. *Graphic Design Processes*. New York: Van Nostrand Reinhold, 1992.

Ishizaki, Suguru. *Typographic Performance*, Ph.D. thesis at the Massachusetts Institute of Technology, to be completed 1996.

Janson, H. W. *History of Art*. 4th edition revised and expanded by Anthony F. Janson. Englewood Cliffs, N.J.: H.N. Abrams, 1991.

Kinney, Leila. MIT Department of Architecture, Program in History, Theory, and Criticism. Conversation, 1995.

Lakoff, George and Mark Johnson. *Metaphors We Live By*. Chicago: The University of Chicago Press, 1980.

Levin, Iris, and Dan Zakay, Editors. *Time and Human Cognition: A Life-Span Perspective*. Amsterdam: Elsevier Science Publishers B.V., 1989.

Lokuge, Ishantha and Suguru Ishizaki. "GeoSpace: An Interactive Visualization System for Exploring Complex Information Spaces." *CHI'95*, Denver, CO, May 7-11, 1995.

Macar, Francoise, Viviane Pouthas, and William J. Friedman, Editors. *Time, Action, and Cognition: Towards Bridging the Gap*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1992.

McCloud, Scott. *Understanding Comics*. Northampton, MA: Tundra Publishing Ltc., 1993.

Mitchell, W.J.T., Editor. *On Narrative*. Chicago: The University of Chicago Press, 1980.

Mitchell, William J. *City of Bits: Space, Place, and the Infobahn*. Cambridge, MA: MIT Press, 1995.

Newhall, Beaumont. *The history of photography: from 1839 to the present*. New York: Museum of Modern Art, 1982.

Powell, Corey S. "The Rights Stuff: Buying and selling art in a digital world," *Scientific American*. Volume 272, Number1, January 1995.

Rennison, Earl. "Galaxy of News: An Approach to Visualizing and Understanding Expansive News Landscapes." *UIST*, Marina del Rey, CA. November 2-4, 1994.

Rosenblum, Naomi. *A World History of Photography*. New York: Abbeville Press, 1984.

Sabiston, Bob. *Infinite Zoom Animation*. Research of the Visible Language Workshop Videotape. 1993.

Small, David, Suguru Ishizaki and Muriel Cooper. "Typographic Space." *CHI'94 Companion,* 1994.

Strausfeld, Lisa. "Financial Viewpoints: Using Point-of-view to Enable Understanding of Information." *CHI'95 Proceedings*, Denver, CO, May 7-11, 1995.

Strausfeld, Lisa, and Earl Rennison. "The Millennium Project." *COSIT'95.* Vienna, Austria. 1995.

Szarkowski, John, and Maria Morris Hambourg. *The Work of Atget*. New York : Museum of Modern Art, 1981.

Tarnas, Richard. *The Passion of the Western Mind: Understanding the Ideas That Have Shaped Our World View.* New York: Harmony Books, 1991.

Tufte, Edward R. *Envisioning Information*. Cheshire, CT: Graphics Press, 1990.

Tufte, Edward R. *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press, 1983.

Wong, Yin Yin. *Temporal Typography*, M.S. Thesis at the Massachusetts Institute of Technology, 1995.