

**In compliance with the
Canadian Privacy Legislation
some supporting forms
may have been removed from
this dissertation.**

**While these forms may be included
in the document page count,
their removal does not represent
any loss of content from the dissertation.**

University of Alberta

Affordances of Prospect
For Academic Users of Interpretively-Tagged Text Collections

by

Stanley Ruecker



A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the
requirements for the degree of Doctor of Philosophy

in

Humanities Computing

Department of English
Department of Art and Design

Edmonton, Alberta

Fall 2003



National Library
of Canada

Bibliothèque nationale
du Canada

Acquisitions and
Bibliographic Services

Acquisitons et
services bibliographiques

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file *Votre référence*

ISBN: 0-612-88040-0

Our file *Notre référence*

ISBN: 0-612-88040-0

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

Canada

University of Alberta

Library Release Form

Name of Author: Stanley Ruecker

Title of Thesis: Affordances of Prospect For Academic Users of Interpretively-Tagged Text Collections

Degree: Doctor of Philosophy

Year this Degree Granted: 2003

Permission is hereby granted to the University of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves all other publication and other rights in association with the copyright in the thesis, and except as herein before provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatever without the author's prior written permission.

22 August 2003

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled Affordances of Prospect For Academic Users of Interpretively-Tagged Text Collections submitted by Stanley Ruecker in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Humanities Computing.

20 August 2003

ABSTRACT

Drawing on previous research on habitat theory (Appleton 1975) and the ecological perception of affordances (J.J. Gibson 1979), this dissertation strengthens the theoretical basis for further research into the development and use of rich-prospect interfaces, where some meaningful representation of every item in a collection is an intrinsic part of the interface. It also: a) analyses some of the details of applying rich-prospect principles to computer interfaces, and in particular to interpretively-tagged text collections; b) examines some methods for evaluating the new affordances made possible by rich-prospect interfaces; and c) suggests some strategies designers might use in carrying out the design of rich-prospect interfaces, including the need to work in a participatory manner in order to develop an appropriate set of item representations and tools for manipulating the display.

KEYWORDS

Prospect, Affordance, Computer-Human Interfaces, Browsing Interfaces, Orlando Project

ACKNOWLEDGMENTS

The author wishes to acknowledge the support of the Social Sciences and Humanities Research Council of Canada, the Department of English, and the Department of Art and Design at the University of Alberta, as well as the Orlando Project.

CONTENTS

Introduction: Goals	1
Rationale	3
Terminology	5
Structure of the Document	8
Chapter 1: The Analog Affordances of Prospect.....	12
Literature Review: Affordances.....	12
Affordances: Ontology of Affordances and Effectivities	14
Affordances: Perception	16
Affordances: Intention	20
Affordances: Learning.....	23
Affordances: Nesting.....	25
Affordances: Sequencing	26
Affordances: Using	27
Affordances: Static, Kinematic, and Dynamic.....	28
Affordances: Modality	29
Affordances: Features	29
Affordances: Reflexivity	30
Affordances: Relationship to Gestalt	31
Affordances: Pleasure	32
Literature Review: prospect.....	33
Universalism.....	34
Prospect and Crime.....	35
Prospect and Surveillance.....	36
Prospect and Mental Fatigue.....	37
Prospect, Biophilia, and Healing	37
Prospect and Aesthetic Appreciation.....	39
Operationalizing Affordances	40
Previous Theoretical Definitions: Dynamic Systems.....	42
Previous Operational Definitions: π Numbers.....	43
The Affordances in Question Are Not Physical.....	44
The Affordances Are Physical, But Bodily Dimensions Are Not Relevant.....	44
The Goal Is to Analyse All the Relevant Information about Affordances	45
Relational Factors of Affordances	45
The Strength of Affordances as a Vector Space.....	50
Interaction of the Vectors.....	50
Vector Values	51
Vector Anchors.....	53
Tacit Capacity	53
Situated Potential	54
Awareness.....	54
Motivation.....	55
Ability.....	55
Preference	55
Contextual Support	56
Agential Support.....	56
Affordance Vector Worksheet.....	57
Chapter 2: The Digital Affordances of Prospect.....	60
Composite Affordances.....	60
The Landscape Metaphor.....	62
Maps	63
Virtual Datascapes	65
Topic Maps.....	66
Wayfinding	68
Panoramas	69

Depth of Field	71
Rich-Prospect Interfaces.....	72
Meaningful Representations of Items	72
Users	73
Form	75
Content.....	78
Relationship.....	79
Hybrids	81
Amount of Information on Display.....	82
Insights About the Collection.....	83
Insights about Content.....	84
Insights about Structure.....	85
Insights about Context.....	86
Insights about Features.....	87
Insights about Limitations.....	87
Insights about Connections.....	88
Insights about Trends.....	88
Insights about Anomalies.....	88
Insights about Navigation.....	89
Reminders.....	89
Reassurance.....	90
Reduced Helplessness.....	90
Prospect-Related Interface Tools.....	90
Prospect-Related Interface Tools: Zooming.....	91
Prospect-Related Interface Tools: Panning.....	94
Prospect-Related Interface Tools: Sorting.....	94
Prospect-Related Interface Tools: Selecting.....	96
Prospect-Related Interface Tools: Grouping.....	97
Prospect-Related Interface Tools: Subsetting.....	98
Prospect-Related Interface Tools: Renaming.....	99
Prospect-Related Interface Tools: Annotating.....	99
Prospect-Related Interface Tools: Opening.....	100
Prospect-Related Interface Tools: Structuring.....	101
Incorporation of Prior Affordances	101
Typeahead Searching.....	103
Characteristics of Candidate Collections.....	104
The Possible Uses of the Collection	104
The Number of Items in the Collection	105
The Characteristics of the Individual Items	106
The Degree of Homogeneity Among Items.....	106
The Possibility of Providing Some Meaningful Representation of Each Item.....	108
The Extent of the Markup of the Collection	109
Design Issues for Rich-Prospect Interfaces.....	109
Hierarchies and Taxonomies	109
Screen Size	111
Persistence of Display.....	114
Priming.....	115
Ventral vs. Dorsal Stream Perception.....	115
Mental Models.....	115
Sequential vs. Spatial Prospect	116
Inter-Affordance Effects.....	117
Interaction Histories.....	117
Co-ordinating Multiple Views	117
Performance	118
Characteristics of Candidate Tasks	119
Chapter 3: Prospect on Interpretively-Tagged Text Collections.....	120
Chapter Outline	121

Textual Markup	122
Tagset Definition as Interpretation	123
Document Tagging as Interpretation	125
Rich Prospect on the Tagset	127
Rich Prospect on the Tagset: Content	130
Rich Prospect on the Tagset: Structure	130
Rich Prospect on the Tagset: Context	131
Rich Prospect on the Tagset: Features	132
Rich Prospect on the Tagset: Limitations	133
Rich Prospect on the Tagset: Connections	134
Rich Prospect on the Tagset: Trends	134
Rich Prospect on the Tagset: Anomalies	135
Rich Prospect on the Tagset: Reminders	135
Rich Prospect on the Tagset: Reassurance	136
Rich Prospect on the Tagset: Reduced Helplessness	137
Rich Prospect on the Tagging	137
Rich Prospect on the Tagging: Tagged Content	138
Rich Prospect on the Tagging: Tagging Choices	139
Rich Prospect on the Tagging: Tag Frequency	139
Rich Prospect on the Tagging: Tag Density	140
Rich Prospect on the Tagging: Tag Distribution	140
Rich Prospect on the Tagging: Tagging Consistency	142
Help Systems and Tutorials	142
Visual Culture	143
Prospect as a Visual Position	145
Adoption of New Technology	146
Constraints	147
Natural Mappings	150
Information Visualization	151
Meaningful Representation of the Tagset	151
Vocabulary	153
Picklists	154
Structure	155
Attributes and Attribute Values	156
Meaningful Representation of the Tagging	159
Access to the Tags While Browsing	161
Access to the Tags While Reading	163
Chapter 4: Prospect-Based Interfaces and the Orlando Project	167
Chapter Outline	168
The Complexities of Orlando	169
Orlando and Prospect	170
Orlando Tagsets	171
What might the user gain by having prospect on the Orlando tagsets?	171
How might prospect on the Orlando tagset be provided?	172
Should the presentation of the three tagsets keep them distinct?	173
How could tagset prospect interact with collection prospect?	174
Orlando Names	175
Extending the Rich-Prospect Name Display	176
Author Names and the Name Authority List	180
Displaying Author Names	181
Picklists	181
Microtext Picklists	181
Walls of Text	182
Panoramas	183
Cross-references	185
Are there alternatives to naming documents by author?	185
Use of Images to Represent Documents	186
Heterogeneous Displays	187

	Homogeneous Representation.....	188
	Homogeneous Items	188
	Visual Structure.....	190
	Indexical Cues	190
Orlando Dates.....		191
	Dates and Chronologies.....	192
	<Date: Certainty>	193
	Displaying <Date: Certainty>.....	194
	<ChronStruct>.....	195
	Displaying <ChronStruct>.....	195
	Timelines	195
	Scattergrams.....	197
	Sequential Prospect.....	199
	Rich Prospect.....	200
	<ChronStruct: Relevance>.....	201
	Displaying <ChronStruct: Relevance>	203
	<ChronStruct: ChronColumn>.....	206
	Displaying <ChronStruct: ChronColumn>	207
	Attaching Visual Cues to Individual Items	207
	Grouping Items.....	208
Prospect-Related Tools for Orlando.....		210
	Sorting.....	211
	Grouping.....	212
	Interaction Histories.....	214
Academic Users and Usability.....		216
Conclusion		217
Chapter 5: Summary and Conclusions.....		218
	Strengthening the Theoretical Grounds for Rich-Prospect Interfaces.....	218
	Applying Rich Prospect to Computer Interfaces	219
	Applying Rich Prospect to Interpretively-Tagged Text Collections	219
	Evaluating New Affordances.....	220
	The Design of Rich-Prospect Interfaces.....	221
	Conclusions.....	223
Chapter 6: Further Research		224
	Further Research: Digital Collections.....	224
	Empirical Determination of Pi Numbers for Rich-Prospect Displays	225
	Evaluation of Affordance Strength for the Affordances Involving Prospect	225
	Alternative Meaningful Representations in Rich-Prospect Interfaces	226
	Effects of Non-Persistent vs. Persistent Display on Perception of Prospect	226
	Structural Priming in Rich-Prospect Interfaces	226
	Effects of Sequential vs. Spatial Prospect.....	227
	Further Research: Textual Markup.....	227
	Effect on Pi Numbers of Multiple Simultaneous Prospect Views	227
	Evaluation of Affordance Strength for Prospect on the Tagset and the Tagging.....	228
	Further Research: The Orlando Project.....	228
	Combining Document Types in Heterogeneous Displays.....	228
	Comparing Results of Automatically-Generated Event Labels	229
	Document Access and its Correlation with Interaction Histories	229
	User Research on Prospect-Based Interfaces.....	229
Bibliography		231
Appendix A: Technical Considerations		243

FIGURES

PAGE

Chapter 1: The Analog Affordances of Prospect

Figure 1.01	The dorsal and ventral visual streams, from the eye to the thalamus at the centre, to the primary visual cortex at the back. The dorsal stream then proceeds to the posterior parietal lobes, while the ventral stream moves to the inferior temporal lobes.	17
Figure 1.02	In the Müller-Lyer illusion, the first two horizontal lines appear to be different lengths although they are the same length. Ellis et al. used the third version, with arrows pointing the same way, which changes the perception of the centre of the line rather than the perception of its length.	20
Figure 1.03	In the Ponzo illusion, the bar appears to be wedge shaped although it is actually rectangular. The perceived centre of mass of the bar is therefore affected.	20

Chapter 2: The Digital Affordances of Prospect

Figure 2.01	A sequence of maps can serve as a visual narrative. Monmonier's map of railroad lines on the Delmarva Peninsula, from 1869 to 1991, illustrates the changing commitment to rail which had its peak expression just after the turn of the century (Monmonier 1993, p. 216)	64
Figure 2.02	A topic map shows the available topics in a collection and groups the information according to the relationships among ideas. This map shows the six categories of anti-infective agents known as cephalosporins.	66
Figure 2.03	If the topic map is reconfigured to rely on the Gestalt principle of proximity, the result can be a more compact display that leaves room for additional information. In this case the number of manufacturers of each anti-infective has been added, and the total is shown for the entire class of drug, which is indicated by a larger font in a different shade of gray, rather than by central position as in the topic map.	67
Figure 2.04	The interface to the Nemo project, which accesses documents related to Electricité de France. The repetition of icons gives a low information return on investment (Hascoët and Soinard 1998).	76
Figure 2.05	Unlike the interface to the Nemo project, the data mountain uses a unique graphic for each document, although the high degree of overlap renders many of the images inamenable to viewing (Czerwinski et al. 1999).	77
Figure 2.06	The Photomesa interface provides the user with a wall of thumbnail versions of photographs which are perhaps surprisingly accessible to browsing, given the complexity of their initial visual impact (Bederson 2001).	79

Figure 2.07	Starfield displays are like entity-relationship (E-R) diagrams in that the relationship between individual items is considered primary. This starfield display shows a building's cooling system as a central point representing each fan, with surrounding points indicating the fan temperature as either too hot, too cool, or just right (the original is colour-coded red, blue or green) (Johnson Controls 2002).	80
Figure 2.08	The Spotfire Decisionsite for Functional Genomics contains a variety of tools, types of displays, and simultaneous multiple views to allow geneticists to work with complex genetic information (Spotfire 2002).	81
Figure 2.09	The tool palettes in Fractal Design's Poser collapse into tabs at the bottom and right edges of the screen (left). When expanded (right), these tabs dramatically increase the functionality of the software. However, for users who are not familiar with the visual language of the program, these tabs can easily be overlooked. One possible solution is to provide prospect on the tool bars by having them collapse in an animation during the start of the program.	92
Figure 2.10	A fisheye menu system allows the user to obtain prospect on the entire list of options but selectively magnify them at the point where a choice of items is being made. This screenshot shows the same menu at three different insertion points (Bederson 2000).	92
Figure 2.11	The film score in Adobe Premiere can be displayed in various time increments, which allows the user to focus in on parts of the film or see the entire score at once. Similar functions are available in many programs that require a time-related display, such as programs that deal with digital music.	93
Figure 2.12	Baudisch's toggle map for selection of television stations in Germany has many admirable prospect features: the switches are indicated by buttons on the text rather than by additional graphical elements; the items are grouped by geographical location; and the context of placement within Germany is indicated by superimposing the toggle map on a map of the country.	97
Figure 2.13	Electronic paper as conceived by Nick Sheridan of Xerox PARC in the early 1970s and manufactured in prototype rolls by 3Com in the late 1990s.	112
Figure 2.14	TextArc provides a visually striking means of investigating word frequency, distribution and co-occurrence in a long text. This display shows <i>Alice in Wonderland</i> .	113
Figure 2.15	The ISO standards for signage suggest a minimum font size related to average viewing distance, as well as placement within certain ranges depending on the contents of the sign (Frascara 1984).	114

Chapter 3: Prospect on Interpretively-Tagged Text Collections

Figure 3.01	Kartoo selects web sites with large numbers of documents on a particular topic, then displays those sites as central nodes in a network of search results. Sites that fall beneath the threshold are not part of the display, which means that Kartoo usually privileges institutional sites over individual ones. Clicking on a node brings up further information about the documents it contains.	141
Figure 3.02	By providing constraints on the user's input in accordance with the required syntax for retrieval, a search interface quickly becomes visually and conceptually complicated. This set of interfaces shows the progressive addition of features relating to searches on contents, tags, and attributes in an interpretively-tagged system.	148
Figure 3.03	Text layering can be used to keep the display of the tagset and its associated information in a compact form. Here the tags are shown as the largest text items, with the attributes of each tag superimposed in a smaller font, and the various pre-defined attribute values are smaller still and shown next to the attributes.	157
Figure 3.04	A variation on text layering that is not as visually complex but is less compact is to cluster the tags, attributes, and attribute values by proximity.	158
Figure 3.05	An example of an illustration showing a complex ecology of information, this full-prospect display shows 1000 document titles, a tagset of 75 tags, and two documents opened to show comparative placement of the two tags currently of interest to the user. In this case, the tags are reproduced near the open document thumbnails, and small geometric shapes are used to show rough placement in the document. Ideally, each of the items in the display would be an object that could be manipulated by the user, and the various panes could be repositioned to allow access the material underneath.	162
Figure 3.06	Adobe Acrobat shows a thumbnail display of the current document. This sidebar can serve both as a form of prospect on the contents and a means of navigation. A slightly larger form might be modifiable to show tag type and placement.	164

Chapter 4: Prospect-Based Interfaces and the Orlando Project

Figure 4.01	The default Google web search engine is an example of a retrieval interface with no prospect.	167
Figure 4.02	This wall of text shows a list of author names from Orlando, in alphabetical order by author last name. A fisheye lens effect is used to allow the entire display to fit at one time on the screen.	182

Figure 4.03	This panorama contains roughly 12,000 names from the Orlando Project, of which the smallest strip shows approximately half at any one time. The names are arranged in alphabetical order in columns. The zooming feature is continuous, and corresponds to the user's movement of the mouse up or down on the screen. The strip is shown here at three different levels of magnification.	184
Figure 4.04	Visual representations of the values for the attribute Certainty, which is used with the various Date tags in Orlando. The tags might be similarly grouped.	194
Figure 4.05	The Nowslider provides a visual prediction of the future based on a current state of knowledge. As the user changes the state of current knowledge in the system by sliding the thumb along the bottom timeline, the temporal model on display also changes (Drucker and Nowviskie 2003).	197
Figure 4.06	This scattergram shows one point for each of 3600 events. The user is able to select subsets of the points by moving the vertical bars at the endpoints of the selection. In this case, 500 events are shown as currently selected.	198
Figure 4.07	A radio button interface to the <ChronStruct: Relevance> attribute values would allow the user who is familiar with the terminology to select an appropriate choice. For users unfamiliar with the terminology, some additional experience or explanation might be necessary.	203
Figure 4.08	Since the four possible values for <ChronStruct: Relevance> are additive, one appropriate solution is to allow the user to select them by choosing among nested buttons where the outer choices automatically include the inner choices. The degree of grey on each button is supposed to reinforce the idea of additive selection.	204
Figure 4.09	Grouping tags by Relevance values provides the user with an immediate impression of the significance of the different values.	205
Figure 4.10	Here the four <ChronStruct: ChronColumn> values are used to organize a table of chronological results. The horizontal bars indicate that some of the columns have been collapsed to the left. The heading of each column indicates the value. Dragging the vertical strip associated with each heading will expand that column and collapse the others, although some representation is always visible for all four possible values.	209

INTRODUCTION: GOALS

This dissertation examines the potential value of a particular kind of overview of document collections. The discussion is largely in terms of the new opportunities for action that can be made available to the user by building on the basis of such an overview, although the analysis also includes some advantages that are more directly related to perception than to action.

In more technical terminology, the ultimate purpose of this dissertation is to strengthen the theoretical justification for further research into the development and application of user-centered, domain-specific, rich-prospect browsing interfaces for interpretively-tagged text collections.

A user-centered interface is one where the design has been informed by the involvement of a particular community of users, whose understanding, both conscious and tacit, goals, common practices, and needs form the criteria by which the interface is developed and evaluated.

A domain-specific interface attempts to accommodate the underlying structure of the material it expresses. In a rich-prospect interface, a meaningful representation of each item in the collection is intrinsic to the interface.

A browsing interface is intended to support someone seeking to understand, interpret, or systematize the material in a domain. Browsing interfaces stand in contradistinction to retrieval interfaces, where the goal is to support someone looking to find a particular item whose retrieval characteristics are well-specified (by author, title, keyword, publication, publication date, and so on).

Interpretively-tagged text collections are those where some text markup system has been used to provide an invisible layer of information about the text, where this information provides more than structural formatting.

Within this ultimate purpose of strengthening the case for rich-prospect interfaces, there are three related sub-goals. The first goal is to outline the potential contribution of prospect in browsing interfaces to interpretively-tagged text collections, and possibly to other forms of electronic collection as well. Interpretively-tagged text collections are singled out because they are a special case that involves an additional layer or layers of complexity beyond the explicit content of the collection. The material that has been tagged (when it extends past the level of structural tagging) is amenable to searching, not only by string, but also by tag. As such, these collections represent an important opportunity to study the possibility of providing the user with multiple levels of meaningful representation of the individual items, as an intrinsic part of the interface that makes access to the collection possible. If the case can be justified for research into prospect-based interfaces for these

kinds of collections, where the provision of some form of prospect seems strongly suggested, then perhaps the case can also be extended to justify similar research on behalf of the users of collections that are not tagged so deeply, or perhaps not tagged at all.

The second goal of the dissertation is to examine the means by which the new or extended affordances of prospect in such interfaces might be amenable to study within the constraints of a particular use of the interface by a given user with a particular agenda in a specific context. By piling these criteria one on another, it becomes clear that a lab-based comparison of interface features might form a component of such studies, but will not be sufficient on its own terms to meet the larger brief. Some contextual framework needs to be introduced that will accommodate as much of the relevant information as possible, but will also include a level of granularity that is appropriate for each given case. That is, the study of rich-prospect interfaces should include observations at the following levels, which are arranged from the coarsest level of granularity to the finest level:

- each of the new affordances *per se*, within the context of potential tasks, but divorced as much as possible from the visual and technical implementation
- the affordances as implemented, which will necessarily include the context of a given collection and its contents, and possible tasks specifically related to that collection
- the affordances as implemented, but within the framework of a set of particular tasks as understood and carried out by an identified user
- the affordances as used in a context of actual work in an environment established by the user's office, computer support services, colleagues, and so on.

The third goal is to deploy whatever insights are developed so that they can contribute where appropriate to the design of rich-prospect browsing interfaces. It is widely recognized in the design community that there is a chronic applicability gap (Mitchell 1993, p. 36) between the information available from user-based studies and the information necessary for a designer to develop a valid solution to a particular brief. If means can be construed for the analysis of affordances in general, and these methods can be shown to be useful in the case of affordances of prospect in particular, then perhaps the applicability gap can be reduced in certain cases.

RATIONALE

The case for further research into these areas requires strengthening, because it represents an opportunity to expand the range of tools and perceptual advantages available to people accessing electronic materials. There is no question that excellent research has already been done in this area (e.g. Pirolli et al, 1996, Shneiderman et al. 2002, Wexelblat and Maes 1999). However, more work remains to be done. Rich-prospect browsing interfaces have the potential to facilitate the selection and organization of collection items in ways that contribute both to an understanding of the structure and contents of a given collection, and to the exploration of that collection in terms of the generalized areas of interest (as opposed to specific retrieval targets) that are of significance for a given researcher.

One of the primary ways in which prospect-based interfaces are unique is that they provide affordances that are not found in other kinds of interfaces. An affordance is an opportunity for action. The idea was formulated by J. J. Gibson (1979) as a way of attempting to find an alternative position in psychology to the schools of behaviourism and mentalism. J. J. Gibson felt that, because biological organisms are involved, the normal methods of cognition are directly related to activity in the environment of the perceiver. To distinguish in a somewhat arbitrary way between mental awareness and subsequent action was therefore to miss the point: perception is fundamentally coupled with action. That it is coupled with action in a given environment by a particular creature results in some additional complexity in the theory, which has been one of the grounds for investigation and discussion by researchers in ecological psychology. By placing the significance of the perceptual event in the relationship between the organism and its environment, J. J. Gibson emphasized that it is less helpful to analyze perception as a set of algorithmic steps, than it is to examine it as part of the dynamic process of interaction.

In order to help make the case that prospect-based interfaces provide affordances not found in other kinds of interfaces, one approach is to look at J. J. Gibson's concept of affordances as it is currently understood both in the computing science community and among the ecological psychologists who have built on J. J. Gibson's work. By examining how affordances have come to be understood, studied, and measured in other areas, it is possible to suggest methods for studying and measuring the affordances of prospect-based interfaces. This literature also provides a detailed

analysis of the concept of affordances and their various features (e.g. Bingham 2000, Chemero 2000 and 2001, E.J. Gibson 2000, Hecht 2000, Heft 1989, Höge 1990, Norman 1990, Warren 1984).

However, there is a fundamental difficulty with investigating new or increased affordances in computer interfaces, because, on the one hand, to compare two interfaces that have the same affordance provided by different methods is to compare the methods and not to address the question of the affordance being new. On the other hand, to attempt to compare an interface that provides a particular affordance with one that does not is to risk committing a category error, as though the comparison was between, for example, a colour and a shape. The two categories are not commensurate; they are not amenable to comparison (Ryle 1949, pp. 16-24).

One possible solution to this Catch-22 is to adopt a method of comparison that is based on the component factors of affordance strength, which, when taken together, can provide some indication of how various people value different opportunities for action within the context of a particular research activity. Further discussion of a proposed component model of this kind is found at the end of Chapter 1.

In connection with the study of new affordances, it may also be useful to distinguish among different categories of people using the system, based on their degree of active interest in the domain knowledge that forms the basis of the collection for which the interface has been designed.

To take a hypothetical example, if a designer is interested in helping people learn to shop for cars, one method is to design the interface, then ask participants to pretend they are interested in buying a car. Another method would be to recruit participants to examine the interface who are currently interested in buying a car and who have active experience of the current process, the search collections that are available, and their interfaces. Such participants also have the specific domain knowledge they have been collecting. They therefore have a variety of characteristics that distinguish them from the first group. A third option would be to bring in study participants who are professional buyers of cars – for instance, those who buy cars for car rental dealerships. This third group would bring another level of domain knowledge to bear, deriving from their prolonged experience and expertise in the field. By emphasizing the role of previous domain experience in the selection of the study participants, it may be possible to obtain help and advice that would not otherwise be available to the designer.

TERMINOLOGY

One question that might arise is whether or not it is useful to adopt the term “affordance” in this discussion, in preference to more common alternatives, such as “functionality.” The concept of affordance is important here because it helps to establish a semantic space related to the fluid mediation of understanding that occurs between people and their environment, as opposed to the unmarked term “functional,” which can tend to narrow the discourse in the direction of a simpler, goal-driven activity. Affordances naturally expand in a more multivalent way. As Bingham (2000) points out, it is normal for people to perceive a wide range of possible uses for an object with a given set of properties in the analog world.

A knife could provide an opportunity for cutting, hammering, driving a screw, chiseling, scraping, forking, reflecting light, branding, throwing a projectile, drawing a straight edge, measuring a length, picking one’s teeth, cleaning one’s nails, scratching a message, and so on, ad infinitum. (Bingham 2000, p. 34).

The “function” of a knife, on the other hand, is to cut. This distinction between single function and multivalent affordance will become important during the discussion, in Chapter 2, of the design of digital interfaces and their components, where the question has been raised by some members of the interface design community as to whether it is helpful to speak of affordances at all in the digital world.

Affordances are also a kind of interface. An affordance is by definition an opportunity that exists for action in the environment of a particular perceiver. As such, affordances cannot properly be considered as attributes of either the environment or the perceiver. It is the combination of the creature and its environment that is being specified.

The affordances of the environment are what it offers the animal, what it provides or furnishes, either for good or ill. . . . I mean by it something that refers to both the environment and the animal in a way that no existing term does. It implies the complementarity of the animal and the environment. (J. J. Gibson 1979, p. 127)

A function, however, seems somehow to exist independently of the people who will be using it. For these reasons, discussing functionality is not the same as discussing affordances, and attempts to establish measurements of the two concepts are going to diverge as, on the one hand, it is natural to

consider primarily the details of performance (how well does this knife cut?) while, on the other hand, the direction is toward considering the larger domain of the connections among the object, its properties, the person using it, the mode of use, and the environmental circumstances (under what conditions is this person using the knife to reflect light? Is it a satisfactory reflector under these conditions for this person at that time? Could it be better or worse? Would making it better or worse have an effect on any of the other affordances of the knife?).

A second distinction that should be addressed is the one between the marked term “prospect” and the more widely used word “overview.” Like “affordance,” “prospect” has a connotation of human involvement with an environment, albeit a relationship in which the person is attempting to achieve an increased degree of success through gaining a larger perspective on the landscape. “Overview,” on the other hand, tends to suggest a subject-object relation in which the observer is independent of the situation. The distinction derives in part from the difference between the hermeneutic or phenomenological position of Winograd and Flores and the more established positivist or rationalist discourse of mainstream scientific inquiry. Describing their orientation, Winograd and Flores say: “It emphasizes those areas of human experience where individual interpretation and intuitive understanding (as opposed to logical deduction and conscious reflection) play a central role” (Winograd and Flores 1986, p. 9).

The term “prospect” also has an intellectual pedigree that draws in related concepts from the appreciation of landscape painting. In this respect, it is part of a tradition that has parallels to human-computer interfaces. Like an interface displayed on a screen, a landscape painting is usually a representational image on a two-dimensional plane, which has to be understood and interpreted by the perceiver. A landscape painting usually sits, like a computer screen, inside a conventional frame that helps to differentiate it from the surrounding optic array. Both the landscape painting and the interface can have greater or lesser degrees of “realism” in the sense of displaying objects that have visual counterparts in the analog environment.

The term “prospect” is also valuable in that it has been adopted by some of the ecological psychologists interested in questions relating to human interaction with the larger environment. So not only does “prospect” connect with theoretical discussions of landscape painting, but it also has formed part of the discourse concerning actual landscapes. This connection is interesting in the sense that

interfaces and their components have in the past sometimes been profitably related to analog counterparts, especially within the design paradigm of the graphical user interface originally designed for the Xerox Star system, which introduced the desktop metaphor with its file folders, pointing arrows as a cursor shape, buttons, check boxes, and so on (Bewley et al. 1983).

The distinction between the analog and digital worlds is not as straightforward as it might at first seem, and the discussion that follows is explicitly framed in terms of the interaction of the two realms. In addition to the technical distinction that operates at the nanoscopic level and differentiates analog from digital in terms of the format of data – on the one hand continuous and on the other binary – there are the grossly perceivable differences wherein the “computer” world is simply not the “real” world. However, many of the sensations, perceptions, and actions in the two domains are parallel to each other, and the interactions between the analog and the digital worlds are many and complex.

At the current point in its evolution, one way in which the digital world is distinct from the analog is that it provides such a limited number of possibilities for interactions with the perceiver. The limitations are in some respects a reflection of the state of computing technology and commerce. For example, good holographic displays or inexpensive wallpaper monitors are simply not currently available to most users. In another sense, however, the restrictions are arbitrary ones primarily mediated by the restricted capacities of both the hardware and software interfaces. Here the issue is less one of pure technological carrying capacity than of unnecessarily restricted design. For example, the technology has existed to make dialog boxes translucent for as long as there have been dialog boxes, yet the default dialog box is still opaque, and therefore still occludes the information behind it on the screen, to which it would often be useful to refer for judgment. Opaque dialog boxes are not a limitation of the technology – they are a limitation of the design (Harrison et al. 1995).

How such limitations come to be entrenched in the culture of computing is a study in its own right, but one interesting suggestion has been put forward by the human factors and usability expert Pat Jordan. According to Jordan, the pragmatics of product development are such that privilege is often accorded to ideas based on whether or not they cross a certain threshold of difficulty in implementation. Improvements or modifications to existing designs that might make a profound difference to the user are often overlooked by developers because they would be too simple to

implement. The argument goes that nobody can justify a computing project that has the goal of making dialogs semi-transparent, because the work involved from a programming perspective is probably less than half a person-day, including testing. So dialog boxes continue opaque, until such time as the modification can be piggy-backed onto a larger, related development project by someone who feels the change is sufficiently important to serve as an advocate for it (Jordan 1999).

A parallel situation exists with respect to the design of browsing interfaces. The computing community has made great strides in the underlying technology of retrieval. Interface researchers have also been active, but the concurrence of initiatives has not been such that acknowledged standards are widely available for the comparison and evaluation of interface effectiveness. The proposal of a metric that might be applied in this way is one of the goals of this dissertation.

STRUCTURE OF THE DOCUMENT

The structure of the document that follows is based on a narrowing focus of attention, starting with the analog world of landscape and ending with the details of the Orlando project. Orlando is an integrated history of women's writing in the British Isles, which consists of a set of electronic documents that have been written and tagged with a set of five SGML tagsets specifically designed for the project. The tags in Orlando allow people doing the markup to specify a wide range of details, including interpretive information that is not present in the text itself.¹

The study of new affordances in interfaces to interpretively-tagged text collections such as Orlando leads inevitably to questions of the study of affordances in interfaces to other kinds of text collections. Consideration of the affordances of digital interfaces leads to questions about the nature of affordances in general, and also to questions about the terms in which affordances have been studied in the analog world.

This document therefore begins with a study of the literature on affordances as they are currently understood within the domain of ecological psychology, where a lively discussion has been underway for a couple of decades into the details of what affordances are, how they are perceived, and

¹ I would like to acknowledge the support of the co-investigators on the Orlando Project, without whom this dissertation would not have been possible. The volume authors are: Susan Brown, Patricia Clements (Director), and Isobel Grundy. The other co-investigators are Rebecca Cameron, Renee Elio, and Jo-Ann Wallace. Many other people associated with the Orlando Project have also provided guidance and support. These people include Sharon Balazs, Terry Butler, Jane Haslett, Susan Hockey, and Jeanne Wood.

how they are related to similar phenomena such as events, environmental features (as opposed to object properties), and cognitive phenomena (such as the Gestalt tendencies). Within that literature, it is possible to identify several threads dealing with the affordances of prospect on a landscape. A seminal book in this area is Appleton (1975), which discusses the features of landscape painting through the lens of habitat theory. Appleton's contribution is significant to the extent that it provides a conceptual framework that strengthens the case for the existence of a human desire to find cognitive reassurance in some form of prospect. In addition to a survey of the relevant literature, this chapter proposes a contribution in the form of operationalizing the concept of the strength of affordances. It does so by creating a vector space that defines the factors of the relationship between the perceiver and the environment that are relevant to the pragmatic evaluation of a particular affordance for a given individual in a specific context.

The second chapter dramatically narrows the focus of discussion, from the entire analog world of perception to the specifics of digital interfaces designed for browsing. The suggestion is made that the significance of prospect may reside in large part in its role as a necessary component of composite affordances. The chapter then addresses some of the implications of implementing the prospect metaphor in the digital environment as a literal interpretation of the analog landscape. The remainder of the chapter focusses more specifically on rich-prospect interfaces and their characteristics.

Chapter 3 narrows the discussion still further, by carrying the insights developed in Chapters 1 and 2 into the domain of interfaces to interpretively-tagged text collections. The argument is that if browsing strategies are to be supported for such collections, it will be necessary to find methods of providing prospect not only on the collection itself, but also on the tagging system. These tags are by default invisible to the reader, and are too obtrusive to be usefully displayed in text designed for continuous reading, but can serve as the basis for enhanced methods of browsing and retrieval.

A further complication is introduced in cases where the tagging system includes pre-defined tag attributes which are invisible to the reader. For example, the Orlando project has a tag defined for use in marking up documents about the writing careers of British women writers. The tag allows the person doing markup to specify mode of publication; the tag has an attribute "publicationMode" which can have one of the predefined values: self-publication, privatelyPrinted, limitedEdition, pirated, subscription, and dedication (Clements et al. 2003). The values of the tag attributes can either be pre-defined as a form of

controlled vocabulary, can be specified in terms of form but not content (e.g., dates within this attribute must be numeric and use the universal date form YYYY/MM/DD), or can be left open to accept any value in any form.

Chapter 4 draws on the previous chapters by applying the discussion to the Orlando project at the University of Alberta and Guelph University. Orlando is implemented as a collection of SGML documents structured around five custom tagging systems (in SGML, these are designated as Document Type Definitions or DTDs). Orlando is an example of an interpretively-tagged text collection in that the tags have been defined to include information on more than document formatting, providing in addition a wealth of detail in the form of cross-references within the collection, references to information external to the collection, and standardized markup that interprets sub-sections of the documents, in many instances at a level of granularity as small as individual words.

The DTDs in Orlando include definitions for more than 250 tags, with more than 600 attributes associated with the tags. For example, the tag identifying genre – <TGENRENAME> – has an attribute “GENREREG” which allows the tagger to specify the genre.² The values of the attribute “GENREREG” are not pre-defined, but a search on the textbase in its current form indicates that more than 60 different genres have been identified so far. In other cases, the possible attribute values have been specified, resulting in a form of controlled search vocabulary.

The Orlando project provides an opportunity for the design of alternative interfaces for an actual collection which can be used to study the effects of increased prospect for the community of Orlando users, many of whom will be academics working in the area of women’s writing in the British Isles. Orlando is a significant part of this dissertation because it provides the opportunity to construct specific examples of prototypical interfaces designed to address some of the questions raised earlier about the nature and use of rich-prospect interfaces. Orlando also represents an actual area of domain knowledge as constructed and delivered to a reasonably well-defined community of users in an interpretively-tagged electronic text format. Since it is one of the contentions discussed in this

² Tags and tag attributes in SGML and XML are case sensitive. That is, tags called <Tgenrename> and <TGENRENAME> would be understood by the system as different tags. However, for the purposes of making this document easier to read, mixed case has been used in many instances for tags and attributes that are actually defined in full upper case. Exceptions occur primarily where the tags or attributes are included within a direct quotation from the Orlando Project.

document that research into prospect-based interfaces should be carried out in the context of a particular domain and community of users, the Orlando project provides a concrete testing ground.

Chapter five presents a summary and conclusions. Chapter six deals with topics for further research. In some cases, these projects will require custom interface solutions, while in others the research is predicated on a need to develop greater understanding of the ways in which a particular user community works with the existing tools to meet their research goals. In other cases, it will be useful to develop research tools intended to address specific questions about the provision of prospect.

CHAPTER 1: THE ANALOG AFFORDANCES OF PROSPECT

In order to begin the discussion, it will be helpful to look at how affordances have been studied in a real-world or analog context, then how the methods of studying affordances might be extended to an approach that can also be applied to the study of computer-human interfaces, and finally, how the concept of prospect can be combined with the idea of affordances to spell out some of the affordances of prospect in the analog world.

LITERATURE REVIEW: AFFORDANCES

An important fact about the affordances of the environment is that they are in a sense objective, real, and physical, unlike values and meanings, which are often supposed to be subjective, phenomenal, and mental. But, actually, an affordance is neither an objective property nor a subjective property; or it is both if you like. An affordance cuts across the dichotomy of subjective-objective and helps us to understand its inadequacy. It is equally a fact of the environment and a fact of behavior. It is both physical and psychical, yet neither. An affordance points both ways, to the environment and to the observer. (J. J. Gibson 1979, p. 129).

The concept of affordances has undergone some significant developments since it was first developed by J. J. Gibson in the first half of the 20th century. In originally choosing to use the term “affordance,” J. J. Gibson was relating his ideas to a concept suggested in 1926 by the German phenomenologist Kurt Lewin (“Aufforderungscharakter”). J. J. Gibson was also influenced by the ideas of Kurt Koffka, a Gestalt psychologist who had been J. J. Gibson’s colleague at Smith College during the 1930s, and who used the term “demand-character” to describe the relationship between the perceiver and the environment (J. J. Gibson 1979, pp. 138-9). J. J. Gibson objected to Lewin and Koffka on the grounds that they described affordances primarily as phenomenological or psychological in nature, while he felt it was important to stress that affordances were relational.

J. J. Gibson expanded on the idea in his now-classic book *The Ecological Approach to Visual Perception* (1979; 1986). Theories based on perception of an affordance are distinct from other theories of perception in that the affordance represents an acknowledgment of an interface between the perceiver and the environment which consists of the possibilities for action in that environment on the part of the perceiver. In fact, it might be said that the concept of affordances confounds the distinction between perceiver and environment. Affordances involve both an environmental property and some capacity of the perceiver to use that property for an action. J. J. Gibson also argued that

affordances could be directly perceived, rather than being constructed by the perceiver using smaller individual pieces of visual or other perceptual information.

The concept of affordances is one of the most controversial aspects of J. J. Gibson's work. By emphasizing the process of direct perception, he was choosing to ignore the possibility that significant levels of mental activity were required by the perceiver. As Ullman (1980) points out, J. J. Gibson can be understood as adopting a two-level model of perception, where the highest level directly represents information about the opportunities for action in the environment, and the lowest level consists of the physiological mechanisms that provide the information. These physiological mechanisms do not rely exclusively on mental activity, but are rather a result of the actions of the organism as a whole. There is some evidence to suggest, for instance, that the dynamic movements of the eye during foveal saccades are essential to the perception of contrast. The retina also contains specialized receptors which respond only to particular kinds of light. The eye would therefore appear to be not so much a static receptor – a kind of camera connected to the brain – as it is an active part of the processing system for visual information.

In that it proposes a human perceptual system composed of a paired mechanical and higher order mental process, J. J. Gibson's two-level model is similar to previous psychological models of the Graz and Würzburg schools (Koffka 1935, pp. 559-60). Faced with the question of how sensations become construed as shapes, the Graz school introduced the concept of a higher mental function they called "production," which served as a label for the end result of a process that was not elaborated further. Similarly, the Würzburg school, in looking at how memory develops from associations, suggested a higher mental function called the "determining tendency."

These schools were explicitly criticized by Koffka as being vitalistic – which is to say they implied that a principle something like the soul was required to explain human mental capacities. J. J. Gibson's model of direct perception of affordances might be similarly accused of implying vitalism, although the counter-argument can also be made that purposeful perception and subsequent action might be developed as species characteristics through natural selection.

It has also been suggested that the two-level model fails to account for all the facts. Research involving perceptual misperceptions and ambiguities or illusions suggests that there should be an intermediate level of study, dealing with algorithmic processes and possible internal representations that can form the basis for mental transformations of perceived objects (Ullman 1980, pp. 379-381).

Although followers of J. J. Gibson have tended to ignore this algorithmic arena of research for programmatic reasons, there seems to be no reason to reject research from other groups that might inform this area (for example, see “perception,” below).

The theory of affordances may therefore have flaws or inherent limitations, but it has nonetheless played a significant role in a wide range of fruitful research and debate. Researchers have looked at a variety of the issues relating to affordances, including the following areas:

- ontology
- perception
- intention
- learning
- nesting
- sequencing
- using
- static, kinematic, and dynamic
- modality
- features
- reflexivity
- relationship to Gestalt
- pleasure

Affordances: Ontology of Affordances and Effectivities

One of the fundamental questions that needs to be addressed in any discussion of research involving affordances is whether or not they exist, or perhaps more precisely, in what way they can be considered as existing. According to J. J. Gibson, an affordance is a perceptual primitive; although it is possible to subdivide it into details of perception related to the optics of surfaces, to undertake that subdivision is misleading because the perceiver does not construct an awareness out of smaller visual components, but rather experiences it as a complete whole.

It is clear that the human visual system perceives surfaces. This aspect had been widely studied by perceptual psychologists. J. J. Gibson proposed extending the significance of surfaces by equating them with direct awareness of what actions their perception suggests, or in his terms, what they afford the perceiver (J. J. Gibson 1979, p. 127).

Prior to J. J. Gibson's ecological theory, the field of psychology could largely be understood as divided into two camps, which had their roots in Descartes and the duality of mind and body. In their psychological guise, these themes were expressed as behaviourism and mentalism, depending on the research emphasis placed, on the one hand, on physical responses and activities, and, on the other hand, on mental constructs and processes. Yet J. J. Gibson emphasized that affordances are not based on a subject-object duality, but are to be understood as forming a middle ground between the organism and its environment. It seems clear that he was attempting to establish an alternative ground for research that ascribed to neither of the two existing camps.

Despite this orientation, one modification of J. J. Gibson's ideas that has been suggested relates to an expansion of the mechanisms involved in the role of the perceiver. The implication is that the original formulations were not completely spelled out in all their details (Turvey and Shaw 1979), and that, in fact, the meaning of the word "affordances" needed to be shifted slightly, so that an affordance is not the interface between perceiver and environment, but rather exists as a property of an object or of the environment, independent of the perceiver. New factors are therefore introduced to account for the role of the perceiver. These new factors are effectivities (or sometimes abilities) and intentions, which represent, respectively, the capacity of the organism to perceive and make use of the affordance available, and the motivation or goal of the organism that may bring it to the point of taking advantage of a perceived affordance.

The term *effectivity* is offered to complement the term *affordance*, and it is defined subject to revision as follows: The effectivity of any living thing is a specific combination of the functions of its tissues and organs taken with reference to an environment (Turvey and Shaw 1979, pp. 9-10).

Factoring an affordance into the aspects that pertain to the object and the aspects that pertain to the perceiver seems like a promising approach to take in attempting to operationalize the concept of affordance for the purposes of research. One problem, however, with Turvey and Shaw's approach is that it would still be useful to retain some term for the designation of the relational aspect. It might be useful to adopt a second set of terms to deal with the environmental properties, which could be substituted for Turvey and Shaw's "affordance," leaving that term as the over-arching specification of the larger interrelation. It seems likely that at least two terms would be required. The first term would deal with the actual value of the object as it offers a particular potential function. A hammer, for

instance, offers a very good potential for pounding. A screwdriver, on the other hand, has only a limited use in this area, primarily through inversion from its normal position in the hand and repurposing of the handle as a form of hammer. The potential of the hammer for pounding would therefore be said to exceed the potential of the screwdriver for pounding.

The second term would deal with the situated potential of the object and its property. It is not very useful to say that a hammer offers better opportunities for pounding than a screwdriver does, if all that is available at the moment is the screwdriver.

There is a sense, however, in which the act of factoring the affordance, on the one hand, and the effectivity, on the other, simply re-introduces the subject-object duality that J. J. Gibson was seeking to reject in the first place (Sanders 1997, p. 104). J. J. Gibson's point was that a rationalism that depends on the existence of subject and object misconstrues the nature of visual perception by not accounting for the central role of the perceiver as an active participant in the environment. In place of this duality, he therefore placed a form of visual perception which provides the perceiver with information related to successfully continued existence in the environment, rather than with alternative conceptions of visual perception, such as the one that suggests that the function of visual perception is to provide faithful images that internally reproduce the external world.

Given that the visual spectrum comprises such a tiny segment of the electromagnetic spectrum, and that the human mechanisms for perception of even that tiny segment have their intrinsic limitations, it would be difficult to make the case that human visual perception provides anything but a small sample of the available environmental information. Whether it is better to understand this sample as being primarily representative of some external reality, or simply as one of the perceptual components that form the basis for human action, is the question that J. J. Gibson addressed with the concept of affordances.

Affordances: Perception

In terms of human visual perception, the current mainstream neurophysiological stance is that the system is based on the development and exploitation of information, as opposed to being based on the capture and storage of data. That is to say, the process and mechanism of vision consists not of image transmission, where an accurate image of the outside world has somehow been sent through the machinery to be recorded in the brain as a faithful image of the outside world, but rather of the extraction of optical information from the environment in a form that is useful to the organism

(Livingstone 2002, p. 24). This optical information is manipulated at each step in a complex path from the moment that light impinges on the photoreceptors on the retina, on through to the thalamus in the centre of the brain, and from there to the primary visual cortex at the back.

After the primary visual cortex, there is some fairly convincing further evidence for the existence of two structurally distinct but related mechanisms in the higher processing areas of the brain (Milner and Goodale 1995). The first stream, which is associated with vision for conceptualization, follows a ventral path forward from the primary visual cortex to the inferior temporal lobes. The second stream, which since Milner and Goodale has become associated with vision for action, follows a dorsal path from the primary visual cortex to the posterior parietal lobes (see Figure 1.01). The evidence for the existence of these two streams and their associated functions comes from studies involving two different kinds of participants: those who have suffered damage to their brains and those who have not.

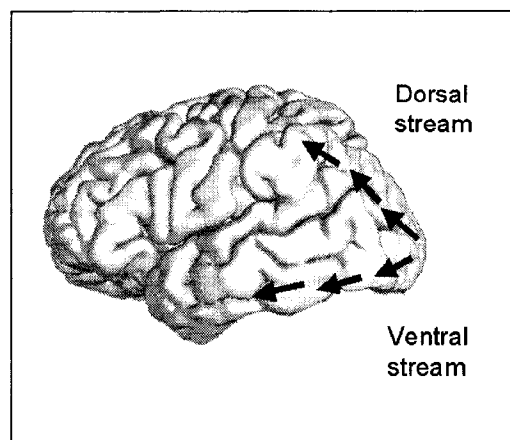


Figure 1.01 The dorsal and ventral visual streams. Information travels from the eye to the thalamus at the centre, to the primary visual cortex at the back. The dorsal stream then proceeds to the posterior parietal lobes, while the ventral stream moves to the inferior temporal lobes.

In the former class is the woman D.F., who suffered damage to the ventral stream and experienced visual agnosia after a case of carbon monoxide poisoning. She could use everyday objects but had problems identifying them and their characteristics. The opposite condition is called optical ataxia, where patients with damage to the dorsal stream can verbally describe common objects, but have trouble using them (Michaels 2000, p. 243).

In the latter class (experiments on people without brain damage) are results that show: disparities between description and action for visual illusions; effects on perception but not action from visual masks which appear after the original stimulus; effects on action but not perception of visual precues; and effects on action but not perception of target repositioning during saccadic eye movement. Michaels (2000) adds to the list a set of experiments involving either judgment or action, based on the same visual stimulus, in which head movement varied according to the type of task. Michaels asked participants to perform two different kinds of tasks related to a circle of light in a dark room. In one case, the task required participants to judge whether the ball represented by the light would land in front of the participant or behind. In the other case, the task was to judge if they should step forward or backward in order to catch the ball. In the former task, involving judgment but not action, the participants kept their heads level. In the latter task, involving action, participants tracked the ball with their heads.

Van der Kamp et al. (2001, p. 168) similarly report that the interceptive timing of hand closure in one-handed catching experiments does not appear to rely on the optical variables that distinguish time to contact (i.e., on perceptual information), but rather on a combination of the relevant rates of change (i.e., on a more complex form of information appropriate for action).

If this distinction between visual streams for different purposes exists, there are several implications for the study of affordances. Firstly, since different information is available from the two different visual streams, it will be necessary to set up experiments that collect information based both on action, or at least on reports of imagined action, as well as on user reports involving conceptualization.

Secondly, the kinds of experiments called for may vary according to the stream being studied. Michaels suggests, for example, that one of the differences between the two streams is that dorsal stream (action-related) visual information may be tacit, while ventral stream (conceptualization-related) visual information may be explicit. She also suggests that there is possibly some distinction between the two streams in terms of their time scale. D.F. experienced increasing difficulty when she had to delay her actions. In Michaels's phrase, "the dorsal stream seems to be very much a use-it-or-lose-it system" (Michaels 2000, p. 252).

Here is Michaels's complete list of ways in which the two streams may vary:

- the information is likely to be different
- the phenomenological experiences may be different
- the principles of learning may be different
- the mechanisms of information detection might be different
- they may operate on different time scales
- they may differ as to the importance of spatial viewpoint
- vision for action may be tacit while vision for perception is explicit

(Michaels 2000, pp. 252-3)

The distinction between the dorsal and ventral streams may not, however, be as clear-cut as Milner and Goodale suggest. Kotchoubey (2000), for instance, points out that neurophysiological evidence has traditionally been found to support whatever the current psychological theories required it to support, since basically everything in the brain can be shown to be connected to everything else one way or another. He also makes the suggestion that, since the studies of D.F. relied on verbal reporting of her conceptualization activity, it is not possible to distinguish in her case between perception and speech, which confounds the clear distinction between vision for conceptualization and vision for action by turning the reporting of conceptualization into a second kind of action.

Some of the supporting experimental evidence has, however, been revisited by subsequent research projects, which in general have confirmed that there seems to be some reproducible difference between vision for reporting and vision for action, but that the details still need to be investigated. Ellis et al. (1999), for instance, carried out experiments using, respectively, a modified form of the Müller-Lyer illusion (Figure 1.02) and the Ponzo illusion (Figure 1.03). The Müller-Lyer illusion uses arrowheads at the end of a line to create an illusion of extended or reduced length. In the Ellis experiments, both arrowheads pointed in the same direction, which can cause perceivers to misjudge the centre of mass. The Ponzo illusion similarly causes mistaken impressions of centre of mass by laying a rectangular shape on a background of converging lines so that the rectangle appears to be wedge-shaped. In both cases participants significantly misjudged the centre of mass, both in the situation where the judgment was indicated by verbally directing someone else to place a mark and in the situation where the judgment was indicated by picking up the bar. However, the latter judgment – the one indicated by the action of the participant – was found to be significantly more veridical than the former.

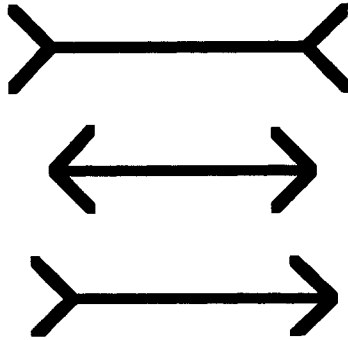


Figure 1.02 In the Müller-Lyer illusion, the first two horizontal lines appear to be different lengths although they are the same length. Ellis et al. used the third version, with arrows pointing the same way, which changes the perception of the centre of the line rather than the perception of its length.

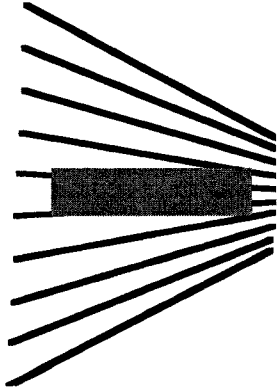


Figure 1.03 In the Ponzo illusion, the bar appears to be wedge shaped although it is actually rectangular. The perceived centre of mass of the bar is therefore affected.

One interesting avenue of future research might involve attempts to identify and study situations in which the ventral stream perception is more accurate than the dorsal stream perception. If the ventral stream is, under certain conditions, superior, then there should be cases where action is significantly influenced by an illusion that is less effective on perception. That is, participants should not necessarily verbally identify an illusion that nonetheless interferes with their actions.

Affordances: Intention

Another potential factor in the perception of affordances relates to the current intentions of the perceiver. When people look at objects, it seems clear that, if they are hoping to accomplish some predetermined tasks with them, they are more likely than otherwise to perceive whether or not the task

can be undertaken using affordances of the object. For example, it is one thing to look for a coffee cup when the goal is to have a cup of coffee. It is somewhat different to see a coffee cup when the goal is to get ready for bed. In the former case, the intention to get coffee makes the situated potentials of the coffee cup consciously significant to the perceiver. The cup affords holding hot liquid, grasping, and drinking. In the latter case, the irrelevance of these affordances to the task at hand means that they are given only cursory, if any, attention.

J. J. Gibson made clear that one of the distinctions between his theory of affordances and previous ideas by Gestaltists such as Koffka was that affordances were to be understood as invariants that did not rely on user intention:

The affordance of something does not change as the need of the observer changes. The observer may or may not perceive or attend to the affordance, according to his needs, but the affordance, being invariant, is always there to be perceived. An affordance is not bestowed upon an object by a need of an observer and his act of perceiving it. The object offers what it does because it is what it is (J. J. Gibson 1979, pp. 138-9).

J. J. Gibson's stance on this issue of user need suggests, among other things, that attraction should not be equated with perception. It also might be understood to suggest that there is an objective quality to the affordance – that it is a quality of the environment, rather than a fact about the interface between the perceiver and the environment. The case can be made, however, that J. J. Gibson's purpose was not to re-open the question of subject-object duality, but rather to prevent the extreme of mentalism in which the emphasis in the relation shifts entirely to the side of the perceiver. The term he uses for the object's role in the interaction is the somewhat active verb "offer," which might be understood to imply that there is a perceiver receiving the offer. Since English syntax is predicated on an inherent dualism of subject-object distinctions, confusions of this kind are inevitable when discussions of relations are the focus.

The question still remains whether or not, in perceiving the cup at all, people also immediately and commensurately perceive all of its affordances (Hecht 2000, p. 59). One test case in this situation is the infant. As Sanders (1997) points out, a baby in proximity to an electron microscope will perceive a wide range of affordances. There will be knobs for turning and shiny surfaces that reflect, there will be some removable parts that may or may not afford swallowing, and, depending on the strength of the infant and leverage conditions of the microscope, there is always the

possibility that the device will afford tipping over. The infant will not, however, be consciously aware of the primary affordance of the electron microscope, which is to visually magnify down to a molecular scale, nor of the related affordances, such as the possibility of winning a Nobel prize (Sanders 1997, pp. 107-8).

The infant's limitations, on the other hand, are not necessarily a deciding factor in the question of whether perception of affordances is holistic or not, because those limitations mean that the device does not afford those actions for that child at that time. It does seem clear that perception of affordances cannot be said to be holistic in the sense that a perceiver is immediately aware of all possible affordances for all possible perceivers, although certainly it is the case that some affordances can be perceived on behalf of other organisms on some occasions. A dog owner, for example, can perceive the affordance of a dog dish for holding dog food for the consumption of the dog, even though the owner has no personal intention of eating the dog food out of the dish.

There is also some evidence to suggest that intention does influence perception. For example, Hommel (1993) reports two experiments designed to investigate the Simon effect (where stimulus-response times are influenced by spatial information that is irrelevant to the task). In a typical experiment on the Simon effect, participants might be asked to respond to a binary stimulus, consisting of a high or low auditory tone, by pushing an appropriate left or right key on a panel in front of them. If each key is associated with a light that comes on when the key is pushed, then there are a total of three spatial objects in the experiment: the source of the tone; the keys; and the lights. By varying the placement of these objects, it is possible to show that response times are faster when the objects are physically associated, even when physical placement is irrelevant to the task.

In light of these and other related results, Hommel was interested in finding out whether the mental model of the perceiver concerning the task could influence the Simon effect. He found that the effect could be inverted by explaining to different groups of participants that their task was either to press a key in response to the stimulus, or else to turn on a light as the response. Depending on the nature of the instruction, the location of either the key or the light became the relevant factor, even though the actual action was identical to an external observer.

The implication of these studies for research design for user interface performance seems straightforward. In order to understand the details of user interactions, it is necessary to establish that the mental models held at the time of the experiments concerning the task are also well-understood

and documented as part of the study. It also seems probable that pre-existing mental models based on relevant domain expertise will be a significant factor, whether that expertise relates to content, procedures, or previous experience with interfaces used for research in the given field. It is therefore likely that results will vary based on whether the participants in the study are currently active in the domain for which the interface has been developed.

Affordances: Learning

J. J. Gibson acknowledged explicitly that people have to learn to recognize and use affordances, beginning, as E. Gibson points out, with an exploratory toolbox that is limited to a few basic functions, such as sucking and looking (E. Gibson 2000, p. 55). How people proceed from there has been the subject of educational theorists for centuries. Within that larger terrain, however, there have been some research projects looking specifically at the learning of affordances. In their study of expert, novice, and inexpert wall climbers, for example, Boschker et al. (2002) identified a number of the factors that differentiate those groups. Expert climbers were able to recall more information that was specifically relevant to the task by clustering it according to the climbing affordances of the wall, whereas inexpert climbers focused on less-significant features and spoke in structural terms rather than in terms of climbing opportunities. Climbing walls use two kinds of holds: footholds, which are too small and smooth to afford grasping, and hand holds, which can also afford standing. In reconstructing a climbing wall with an easy lower section, critical middle section, and difficult top section, experts focused on learning first the position and orientation of the hand holds (which are more crucial to success). They also concentrated first on the critical middle section and difficult upper section, which were the sections that presented the greatest climbing challenges. Inexperts, on the other hand, did not differentiate among the sections of the wall, and treated all holds as equally important. Finally, expert climbers tended to perform climbing gestures or movements during their explanations of the climbing choreography, while inexperts did not use their bodies during their explanations (Boschker et al. 2002, p. 34).

Body theorists insist that learning as a field of activity is not confined to cognitive processes, but that the body itself is something that is learned within the context of a particular culture and environment. In a seminal article in that field, Mauss (1935) compiles an impressive list of body techniques that vary by culture, including walking, running, dancing, marching, swimming, jumping, climbing, descending, holding, throwing, washing, spitting, eating, drinking, massaging, and

reproducing. His own education in France in the late Victorian period included learning to swallow water and spit it out again while swimming: “In my day swimmers thought of themselves as a kind of steam-boat” (Mauss 1935, p. 71). It seems likely that a swimming technique of this kind, as opposed to a technique where the water is not swallowed, would tend to influence the detection of affordances for swimming toward bodies of water that were clear enough to be safely ingested. He also tells the story of British troops in the first World War who were working in alternating shifts with French troops in digging trenches. The army was obliged to provide different spades to each group, because the English could not dig with French spades and vice versa. The learned techniques of the body can therefore have profound effects on both the perception and use of affordances. A corollary of this observation is that it is possible to introduce new affordances to people, provided they are educated to recognize and use them. Without appropriate education, it is not reasonable to expect people to be able to climb the wall, use a new kind of spade, or otherwise behave in an expert manner with respect to a given affordance.

On his list of the learned techniques of the body, Mauss also lists education in vision, which is not one of the topics he elaborates. Some degree of visual perception is inherent from birth, but the differentiation of the visual field by the infant is part of the natural development of the child. The relationship between development and education, however, is a subject of debate among educators. As Vygotsky (1978, p. 80) points out, various theorists interested in education have adopted each of the possible positions, including the idea that development necessarily precedes learning – Piaget and Binet – that the two are actual synonymous – William James – and that they interdigitate, with one feeding the other, then the reverse – Koffka. Whichever the actual situation, the learning or development of visual perception should be classed among the learned techniques of the body which are implicated in the perception of affordances.

Mauss also includes education in composure, or the deliberate suspension of activity. The relationship between inactivity and affordances has received some attention by ecological psychologists, who appear uncertain what status should be given to inaction as a form of action. Since the behaviour of someone who is choosing not to act on an available affordance may be indistinguishable from the behaviour of someone who is unaware of the affordance, or of someone who is aware but lacks the ability to use it, the problem is a complex one to analyze. Proponents of certain forms of inaction, however, would stress that the force of intention and volition are significant,

and that the resulting effects on the environment are also important. An example might be an ecological awareness resulting in unwillingness to purchase or consume products with a negative environmental impact. Another example would be in the historical case of Mahatma Gandhi, whose principle of satyagraha or nonviolent resistance resulted in dramatic cultural changes in India and Britain in the twentieth century. Finally, inaction in a context of strong environmental support for action is a clear indication of volition, as in the cases of people acting collectively to oppose corporate interests by holding a workers' strike.

Affordances: Nesting

Although J. J. Gibson suggests that affordances should be treated as perceptual primitives, it is possible to distinguish among different kinds of affordance, based on the manner in which they interact with other affordances. One such interaction is the nesting of affordances, where several different affordances are intrinsically related to each other by being grouped together spatially. Within this larger category of nested affordances, there are sub-categories, including: invisible nesting; metonymic nesting; and nesting across different planes of experience.

The first kind of nesting involves a combination of visible and invisible affordances, where some of the affordances in a nested group are initially invisible but become apparent upon investigation. An example of this kind of nested affordances is a doorknob. It is not always possible to determine by visual examination whether a particular doorknob is locked or unlocked, or whether it should be turned clockwise, counterclockwise, or either. But for people with the appropriate physiology and experience, the doorknob does afford grasping – that much information is available to visual examination – and the hand that grasps it can be used to determine whether it also affords turning and, if so, in which directions.

A second kind of nested affordance is one where the presence of the entire collection can be signaled by the visual presence of an object or object property that stands in a metonymic relation to the whole. An example of this kind of object is a printed book. Whereas a doorknob might be locked or unlocked, and those two conditions represent a state that is one of the affordances of the doorknob, there is usually no corresponding mystery about a book. If a person is literate in the language and has the appropriate visual acuity and lighting conditions, then a book affords grasping, opening, and reading. In most cases, the cues for language are available in a printed form on the cover or spine, so it is not necessary for the perceiver to open the book in order to decide whether or not it is printed in a

language he or she reads. The language used for printing the spine or cover of the book is therefore an example where an object property, as opposed to the entire object, serves in this iconic or metonymic fashion.

A related but distinct kind of nesting occurs in cases where several affordances occur simultaneously at different cognitive or experiential planes. For example, a cat may afford petting by its owner; the petting affords pleasure for the cat; the petting affords pleasure for the owner; the petting and the cat's pleasure afford a sense of companionship for the cat owner (and arguably for the cat, too). The pet-ability of the cat is a mechanical affordance. The pleasure of the two creatures involved is an affective affordance. The companionship is a social affordance. It is possible to have any of these affordances without the others. The cat may still afford companionship even if it is not currently in the mood for being petted. The cat may also afford petting but fail to experience pleasure, and so on. The cat is also unlike the book in that its willingness to afford petting in the first place is volitional – the book cannot actively resist reading.

Given that affordances can be nested in these various ways, it is not necessary to perceive all the details of an affordance in order to be able to identify and use it. In the case of invisible affordances, such as the locked or unlocked doorknob, it is only necessary to perceive that the doorknob affords grasping and either to know or guess that it may afford turning. In the case of the book, it is not necessary to know ahead of time the various mental states that reading the book will afford – it is only necessary to realize that it affords reading. With respect to petting the cat, it is not necessary to anticipate that the petting may result in a sense of companionship – it is only necessary for either the owner or the cat to initiate the negotiation and see where it leads.

Affordances: Sequencing

The complex nesting of affordances involved in petting the cat introduces a related concept that deals not so much with the nesting of one affordance inside another as with either the changing nature of a given affordance or else the sequential relationship of different affordances across time. To continue the example of the doorknob, the turning of the doorknob may introduce another affordance, namely the affordance that the door has for opening. The movement of the door will reach a point where the doorway it has previously blocked is now cleared, and the doorway will begin to afford entrance. This sequential unfolding of affordances allows the perceiver to interact in a continuous manner with the environment (Bingham 2000, p. 31).

There are similar examples in the natural world directly related to the affordances of prospect. The perceiver of a landscape from a perspective of prominence does not necessarily see the details of the landscape, but the details are not essential to the value of the prospect. It is sufficient to be able to identify areas of potential shelter, danger, food, water, and so on. Upon entering the environment, the prior experience of prospect will contribute to wayfinding, helping to guide the perceiver into the desired situations, as for example in approaching a stream in order to get a drink of water. Although the general path might have been observable from a position of prospect, the details of approach to the stream will not necessarily have formed part of the information available, and may have to be worked out once the perceiver has sufficiently advanced toward the water. As in the case of nested affordances, sequential affordances therefore do not require complete perception, but are amenable to exploration once any component affordance has been recognized.

Affordances: Using

Once a perceiver begins to make use of an affordance, the situation can quickly become complex. For one thing, it is in the nature of affordances that, for the most part, they allow for multiple behaviours. Bingham points out, for example, that a floor which affords support for locomotion for a human adult does not necessarily predetermine the form of locomotion that a given perceiver will adopt. A person may crawl, skip, walk, or dance, and may do any of these actions efficiently or inefficiently, gracefully or gracelessly, and at different possible speeds (Bingham 2000, p. 31).

The distinction between affordance and behaviour is therefore significant in several ways. First, it is the case that the latter is ontologically dependent on the former: every behaviour is predicated on the existence of an affordance that makes it possible, even if the affordance should be a property of the organism. Second, while affordances as a complex whole are not subject to training, behaviours are; a person can perceive that a wall has an affordance for climbing, but initially be unable to climb the wall. After training, the affordance remains the same, but the behaviour changes.

Behaviour is also distinct from ability or effectivity, which is the potential for action on the part of the perceiver. For people who have sought to factor affordances as an approach to operationalizing them for the purposes of research, effectivities are one of the perceiver-side factors. Effectivities differ from behaviour in that behaviour is action, whereas effectivity is potential.

Another factor in behaviour is, therefore, that even though it may be based on a general effectivity and intention that are characteristic of the agent, the actual behaviour is not necessarily

predictable or consistent. A professional ice skater can slip in the middle of an international competition; a person who has learned to punch a falling ball can have inconsistent results in actually accomplishing the movement and striking the ball. The ball nonetheless affords punching and the person has the necessary effectivity and intention; only the behaviour is unsuccessful.

Affordances: Static, Kinematic, and Dynamic

Human beings can perceive affordances that derive from information far more complicated than simple properties of objects or the environment. One of the complexities of perception that the human visual system affords has to do with detecting affordances of objects that are either in uniform motion or under accelerated motion.

Static affordances are those which do not involve objects in motion, although the case has been put that no perception is truly static from the perspective of the perceiver, because the nature of the human eye dictates that vision involves frequent foveal saccades. There is also a tendency for people to move during information-seeking behaviours. Given these caveats, however, there is still a valid taxonomic distinction based on the role played by the motion of the object. For example, to perceive that a ball sitting motionless on the floor is of a size that affords one-handed grasping is to perceive a static affordance of the ball.

The next level of complexity is in the case of objects in uniform motion (that is, not subject to accelerations indicated by changes of motion or velocity). To perceive that a ball rolled along the floor by another person affords trapping between the knees (ignoring the somewhat more complex effects of gravity and friction, which are actually accelerations rather than kinematic effects) is to perceive a kinematic affordance of the ball. Another way of describing this second kind of affordance is to say that it is the first derivative of the position of the ball as it changes over time.

Finally, human beings can perceive affordances that are derived from objects under acceleration. To perceive that a flyball falling from the sky affords catching in a baseball glove is to perceive a dynamic affordance of the ball, since the position is not constant, as in the static affordance of grasping, nor is the position changing at a uniform rate, as in the kinematic affordance of the rolling ball. Instead, the ball is subject to acceleration due to gravity. Another way of describing this kind of affordance is to say that it is the second derivative of the position of the ball (as it changes over the square of time, or accelerates). People are able to perceive all of these different kinds of affordances,

and more. With respect to the design or implementation of new affordances in computer-human interfaces, it should therefore not be necessary to restrict the discussion to static features.

Affordances: Modality

Although much of the research on perception of affordances deals with visual perception, J. J. Gibson and others have not altogether neglected the role of the other perceptual systems. Perception of affordances can therefore be understood to occur across various sensory modes.

For example, a person in the autumn who is deciding whether or not to wear a winter coat might begin by looking out the window to see what the weather is like, then extend the process of information exploration by putting a hand against the glass, and complete the process by going to the front door, stepping outside to feel the air, and concurrently listening to the wind. In this case, the variety of sensory modes used (vision, haptics in several forms, sound) is helpful in determining whether the weather affords prolonged exposure of the body without additional protection.

In fact, one of the tenets of ecological psychology is that intermodal information is often fundamental to action. Not only does someone, while carrying out an action, perceive visually, but the action itself will often involve senses such as touch, smell, and sound, as well as bodily awareness (proprioception) and the awareness of the physical surround (exteroception), all of which contribute to the recognition and use of the larger affordance.

Affordances: Features

One of the intriguing characteristics of affordances is that, while they have been primarily defined and discussed by researchers subsequent to J. J. Gibson in terms of the relationship between perceivers and object properties, an entire class of affordances exists independent of discrete objects. These affordances exist as properties of the environment, or perhaps, to use J. J. Gibson's taxonomy, as properties associated with the medium (air), other substances (water and various solids), or places. As Chemero (2001, p. 114) points out, many affordances of this class are signaled in speech or writing by feature-placing sentences such as "It's hot in here," or "It looks like rain," where the intention of the communication is to identify a feature of the environment that has implications for human activity but is not directly associated with any particular object (Strawson 1959, pp. 202ff, 214ff).

In the domain of landscape perception, or more precisely in the field of prospect on landscapes, the perception of features is primary to the experience. There are any number of potential

features that are observable within the composite of prospect affordances, from the condition of the weather to the current state of development of this year's crop. Some of the more common affordances that features provide to the perceiver relate to wayfinding. Wayfinding has implications for everything from map design to traffic safety to web navigation, and has therefore been widely studied in a variety of contexts. In one project intended to outline the potential implications of analog methods of wayfinding for digital environments, Vinson (1999) identified five types of features that are typically used by people in the process of navigating in the real world. These navigational features, which may be useful in the design of computer-human interfaces with some form of prospect, are: paths, edges, districts, nodes, and landmarks.

Affordances: Reflexivity

Affordances are not restricted to aspects of interactions between a perceiver and natural object or properties of the natural environment. Human artifacts also provide affordances for human beings (and other animals), and in this way a mutuality relation is established between the artifacts and the people by virtue of the affordances.

Human artifacts are to be found at a variety of ontological levels. By applying the concept of human factors to the ontology of artifacts, it is possible to formulate a taxonomy that includes the physical, cognitive, interpersonal, and cultural. Each of these levels of artifact has the potential to serve within a reflexive cycle that enables people to define themselves, or provides a context for definition which is continuously available to processes of modification (Pickering 2000, p. 74). These categories are not necessarily mutually exclusive, since affordances can be nested across experiential levels (as in the case of the petted cat).

The simplest form of artifact to understand in this context is the physical. Physical artifacts can range from those that are microscopic (perfumes, for example), to many at the scale appropriate for grasping (e.g., hand tools), through to those that create an entirely constructed environment. By building cities, for example, people have significantly modified their surroundings, and living within that new urban landscape has consequences for how people understand themselves and their behaviours. The built environment also includes a whole range of new opportunities for action.

Cognitive artifacts are those which have no physical form, but represent the consequences of human activity in the mental sphere. Cognitive artifacts are to a large extent the consequences of learning, and include language, philosophy, intellectual skills, and so on. It seems uncontentious to

claim that language has a reflexive effect on the person who uses the language. Related classes of artifacts are those which are imaginative, metaphoric, or symbolic. An example of a symbolic affordance provided by an artifact would be the affordance of a sense of domestic security provided for a child by that child's favourite blanket.

Interpersonal artifacts that are neither physical nor cognitive include emotional or affective states that develop through involvement with other perceivers. An example of an interpersonal artifact is the spontaneous arising of compassion felt in observing a suffering animal. Compassion is an interpersonal artifact in the sense that it requires an object in order for it to arise in the perceiver; once compassion has been experienced, it can have further consequences for the actions of the person, such as the affordance of compassion to increase the inclination for the person to act in an altruistic manner.¹

Cultural artifacts are those which are created and maintained at the larger level of society. Examples of cultural artifacts include institutions or collective forms of activity and their mechanisms, such as businesses or governmental bodies, legislation, news, marketing and other broadcast phenomena, and so on. An example of a cultural artifact is the internet community, where the individuals together form a collective that can take on an active role in providing new affordances to the members. An example of such an affordance is the idea of spam-blocking by vote, where e-mail users in a given group agree to pool their opinions about the messages they receive to determine which will be filtered at the server level on behalf of the collective (Spamnet 2002).

Affordances: Relationship to Gestalt

The concept of affordances had its genesis in J. J. Gibson's interest in the ideas of the Gestalt psychologists, and the intellectual descendants of the Gestalt school continue to take an interest in J. J. Gibson's idea of affordances. One provocative suggestion is that a new Gestalt tendency should be identified to account for the human ability to perceive complex nested or sequenced affordances. The argument is that, in the same way that the human systems of visual perception have a tendency to fill in the missing pieces under a variety of conditions (for example, in visual association of items in proximity, closure of incomplete outlines, association of objects that are in alignment, and implied

¹ The generation and subsequent cultivation of compassion is the distinct basis of the mental exercises employed by Tibetan Mahayana Buddhists; in this context, the affordance would have to be said to be cultural as well as interpersonal.

relations among similar objects), so the human perceptual systems that allow for perception of affordances have a tendency to create a Gestalt or holistic impression of those affordances that are related either by proximity or sequence over time (Van Leeuwen and Stins 1994).

For Van Leeuwen and Stins, the mechanism of this holistic perception is related to the compounding of affordances across multiple orders of complexity. Tools provide a good example. A pair of pliers affords grasping, which in Van Leeuwen's system is a simple, first-order affordance. The primary purpose of a pair of pliers, however, involves considerably more than its ability to afford grasping. There are several other first-order affordances, such as opening, closing, applying pressure, and gripping. In a particular situation, there may be other affordances that are also first-order, such as reachability from the perceiver's current position, or visibility among the other tools in the toolbox or workshop. The second-order affordance of the pair of pliers, however, is that it affords the tight holding and squeezing of objects that fit within the jaws. For a person with the requisite knowledge of the tool, to perceive a pair of pliers is therefore to perceive an entire array of both first-order and second-order affordances.

A relationship between the affordances of prospect and refuge and the Gestalt tendencies has also been outlined by Nelson et al. (2001, p. 323). Having established that participants correlated completeness of the canopy of a tree with both its fecundity and visual attractiveness, Nelson et al. suggest that the Gestalt figural tendencies, and in particular the principle of closure, may be vestigial mechanisms related to perception of survival-related affordances.

Affordances: Pleasure

In an unpublished manuscript dating from approximately the same period as the first edition of *The Ecological Approach to Visual Perception*, J. J. Gibson briefly discusses the question of how affordances are related to pleasure (J. J. Gibson 1979u). He distinguishes three kinds of pleasure related to the viewing of surfaces, depending on whether the surface has an affordance, stands for other things, or invites inspection for its own sake. This taxonomy relates in part to a taxonomy he proposes of modifications to artifacts, whether to modify their affordances, display additional information, or enhance appearance.

The idea that perception of affordances might relate to pleasure or satisfaction seems like a natural outcome of the role of action in human and other life. One of the areas in which the relationship has been further developed is in the work of Appleton, who initiated what he called

“habitat theory” as a means of discussing the relationship between the perceiver and certain characteristics of an environment, including an overview of what it may afford.

LITERATURE REVIEW: PROSPECT

The concept of prospect was first introduced by Appleton (1975), who was interested in aesthetic appreciation of landscape painting. He began with the question: “what is it that we like about landscape, and why do we like it?” (Appleton 1975, p. 1). His approach was to identify, within habitat theory, two features of landscape that are directly related to survival for people and animals in a natural environment: prospect and refuge: “Where he has an unimpeded opportunity to see we can call it a prospect. Where he has an opportunity to hide, a refuge” (Appleton 1975, p. 73).

Using these twin concepts as a lens, Appleton examined comments published by art critics in the western world who were looking at European paintings of landscape, and was able to identify and elaborate on the themes of prospect and refuge using their work. His contention is that these features of the landscape, which once had survival value, remain as atavistic tendencies toward certain preferences. These tendencies contribute significantly to the appreciation of those artistic representations that include reference to the appropriate landscape features in some form, either as direct representations or as symbolic elements. In this formulation, various configurations are possible, based on how the symbols of prospect and refuge are deployed in a picture. In some cases, the image will be prospect-dominant, in others refuge-dominant, and in still others there will be a balance. Appleton also introduces a third landscape feature – hazard – which he uses to account for symbols that indicate the sublime. He emphasizes that the impact of these symbols is not necessarily related to the rational strength of their connection to what they symbolize:

In just the same way the symbolic representation of danger may be only vaguely and quite irrationally related to a real danger; a ‘refuge’ may afford no real guarantee of security, and a ‘prospect’ which visually satisfies the observer that his immediate environment is free from danger, may be permeated with radiation hazards or alive with poisonous snakes. Yet the symbolic impact of these environmental phenomena can induce in us a sense either of ease and satisfaction or of unease and disturbance, and it is on these emotional responses rather than on the real potency of the danger, the refuge or the prospect that our aesthetic reactions will depend. (Appleton 1975, p. 81)

Spires, for example, have a strong visual structure indicating elevation over the surrounding territory. It is not usual for people to ascend spires, some of which are actually inaccessible to human beings for reasons of physical construction; others are inaccessible through policy; most often, there is no purpose in climbing to the top of a spire. The actual use of a spire for obtaining prospect, however, is not important in recognizing and acknowledging the spire as a strong symbol of prospect (Appleton 1975, p. 90).

From the perspective of ecological psychology, Appleton implicitly identifies a number of affordances of real and symbolic prospect. These include affordances for creating emotional states in the viewer such as ease and satisfaction, or conversely, unease and disturbance. Some of the survival-related affordances of prospect that Appleton mentions include advantages in hunting, seeking shelter, identifying positions of concealment, and exploring (Appleton 1975, pp. 70-1, 175). Appleton also mentions explicitly interpersonal affordances, such as surveillance activities related to the establishment and maintenance of territory (Appleton 1975, p. 41). In cases where the prospect includes elements of the sublime, affordances may also be available for the experience of emotional states such as astonishment, admiration, reverence, or respect (Appleton 1975, pp. 28-9). As mentioned earlier, wayfinding is also an affordance of prospect.

Universalism

One of the fundamental objections to Appleton's formulation is that it is predicated on a universalism in human response which is currently unfashionable in academic circles, particularly among post-colonialists. The idea that there exist certain basic truths which apply to all human beings was, in the eighteenth century, a positive force that was wielded politically by members of the anti-slavery movement. Subsequent generations, however, found that universalism was more often than not adopted as an excuse, not for increased humanitarianism, but rather for various forms of cultural imperialism. The underlying argument was that if all human beings are basically the same, then their manifest differences must be the result of ignorance, misunderstanding, or outright wickedness, and should be corrected.

Appleton relies on biology as the basis for his universalism. By basing his theory on the survival value of prospect and refuge, Appleton suggests that natural selection has played a significant role in allowing the continued survival of those members of the species who were able to identify and capitalize on situations where these two factors were crucial. People who were unable to appreciate

prospect and refuge were theoretically killed before they were old enough to breed, and their inadequate genes were removed from the genome.

It does not seem difficult, however, to posit circumstances in which prospect would not be available, and therefore could not be a significant survival factor, or in which its detection and employment would not be essential for the survival of the individual. Among geese, for instance, it is common for an experienced leader to provide guidance for the rest of the flock in finding water, food, and shelter. It seems reasonable that prehistoric groups of human beings might similarly have relied on previous experience, either individual or collective, rather than on the serendipitous availability of a prominence that afforded prospect to each person in the group. Survival value of group membership would therefore be the predominant factor. Although one of the tools used by one of the leaders might include knowing where to obtain prospect on the area, other successful members might not even be aware of it.

However, the suggestion that prospect and refuge are universally relevant due to human biology is not without merit. Leaving aside natural selection for a moment, it is true that human beings are biological organisms, bipedal, with two highly-specialized eyes on the same side of the head and a tremendous amount of brain capacity dedicated to the processes of visual perception. This physical conformation suggests that certain kinds of environments are going to be privileged by this creature, where plenty of visual information is available in the front and the unobserved back of the head is protected. Appleton's prospect and refuge meet this description nicely.

Subsequent studies of actual landscapes and their perception have looked at potential affordances of prospect that extend beyond the ones originally identified by Appleton. Since the landscapes under investigation are often those involving trees, the researchers are usually interested in some aspect of the biophilia hypothesis, which suggests that people benefit in a variety of ways from exposure to other living things and natural environments. As a consequence, prospect has seldom been isolated as a single significant factor, although it is often implicated in the findings, and deserves to be given closer attention in future research.

Prospect and Crime

One of the extended affordances of prospect, within the context of high-canopy foliage in an urban setting, involves its relationship to the correlation between the presence of trees and the occurrence of crime. In a study of a subsidized housing project in Chicago, Kuo and Sullivan (2001) found that

apartment blocks surrounded by high-canopy trees received significantly fewer police crime reports than neighbouring blocks without trees, even though residents were not involved in maintaining the trees and were randomly assigned to the buildings. This result contradicts both conventional wisdom and previous studies regarding the relationship between trees and crime, where urban vegetation has been understood as affording concealment for potential criminals. One of the factors identified by Kuo and Sullivan was the absence, in this situation, of a significant understory, since the trees were mature specimens of deciduous species with a high canopy and, aside from lawns, the area underneath had been kept clear of growth.

Kuo and Sullivan suggest that the combination of prospect and deciduous foliage provided two affordances that contributed to the lower rates of reported crime: an affordance for the perception on the part of potential criminals of an increased likelihood of resident surveillance, and a reduction of mental fatigue, three components of which have been positively linked to aggression. Both of these affordances are related to prospect. In the case of surveillance, the unobstructed view is a necessary component (although the actual situation is more complicated – see “prospect and surveillance,” below). In the case of reduction of mental fatigue, prospect also appears to play a role, since some of the sub-factors in positive reporting of self-affect in studies of landscape preference include that the field of view be high-depth, spatially open, and natural (Ulrich 1993, p. 83).

Prospect and Surveillance

Although prospect and surveillance might be read as synonyms, in the context of Kuo and Sullivan’s work the former serves the latter as a component of a more complex affordance. To have a situation of complete prospect would be to have no trees or other obstructions to vision whatsoever– an environment which was correlated with a high crime rate in the neighbourhood under study. To have reduced prospect is to have trees with a heavy understory, which was not the situation near any of the dwellings. To have high prospect in the presence of high-canopy trees is the particular environment that correlated to lower crime rates.

Kuo and Sullivan conjecture that one of the deterrents to crime in the areas with high prospect and trees might have been that potential criminals experienced an increased expectation of resident surveillance. The expectation could have been based on the actual increased presence of more residents outdoors enjoying the high canopy foliage, or else on the possibility that residents would be more inclined to look out through their windows at the view, or, finally, that the presence of the trees

indicated a higher degree of resident care of the grounds, which would include natural surveillance by the people during their caretaking activities.

Prospect and Mental Fatigue

The other possibility suggested by Kuo and Sullivan is that the presence of high-canopy trees helped to reduce the rate of reported crimes because potential criminals were less prone to mental fatigue, which contributes to aggressive activity. The three precursors to violence related to mental fatigue are: inattention, irritability, and reduction of impulse control. Mitigation of these factors has been linked by previous researchers to perception of vegetation (Kaplan and Kaplan 1989, pp. 177-200).

Since the crime rates used in their study are based on police records, a third possibility not discussed by Kuo and Sullivan is that the presence of high canopy trees in some way influenced, not the commission of crimes, but rather the reporting of crimes to the police. Since the reports include both those filed by citizens and those filed by an officer, some confounding effect may be possible that accounts for the lower rate of reporting. Perhaps, for instance, the areas with fewer trees are interpreted by the patrolling officers as rougher parts of the neighbourhood and therefore receive correspondingly higher levels of attention, resulting in more reports being filed by officers for those areas, while in fact the actual crime rate across all areas is constant.

The possible relationship among prospect, vegetation, and reduction of mental fatigue is nonetheless interesting, and deserves further attention. It seems reasonable to assume that, in circumstances where vegetation is dense enough to reduce prospect and afford concealment of potential perpetrators of crime, one of the consequences might be an increase in mental strain for the potential victims (and hence the classic wisdom and multiple studies linking dense vegetation with increased fear of crime). Prospect therefore may play a role in the reduction of mental fatigue in particular situations, although it might be said to be a necessary but not sufficient condition. It also seems likely that prospect onto a group of criminals armed and waiting for victims would not provide reduction in mental fatigue. Much depends on what the prospect reveals.

Prospect, Biophilia, and Healing

There are a variety of reasons for believing that people can benefit, under certain conditions, from some form of contact with nature. There is historical evidence of human interest in maintaining contact with nature in urban environments dating back thousands of years, from the hanging gardens

of Babylon (built c. 575 B.C.E.) to the ancient gardens of China. There are several modern studies which correlate the perception of vegetation with human well-being (Whitehouse et al. 2001, pp. 301-2). Folk wisdom supports this correlation, insofar as it is traditional to take cut flowers or living plants to people recovering in hospitals. In the early 1980s, the term “biophilia” was coined to express the possibility that there is a genetic predisposition to respond positively to other living things, including certain forms of landscape (Wilson 1984). In the case of North Americans, Europeans and Asians, studies confirm a marked preference for natural versus constructed environments, with savanna-like landscapes taking precedence over other environments (Ulrich 1993, pp. 90-4).

One of the research agendas deriving from the biophilia hypothesis relates to the human ability to learn some adaptive responses to natural features quickly, learn them vicariously, and forget them less readily, than similarly adaptive responses to artificial objects. Since the negative (or biophobic) responses are easier to control in the laboratory, these have received more research attention than the positive (or biophilic) responses. The theory is that people are biologically prepared to learn some responses which have been significant for survival in the past. Instances of biologically prepared learning include startle reactions to hazardous natural creatures like snakes and spiders, in distinction to hazardous human artifacts like handguns and electric wiring.

Ulrich (1993, p. 88) suggests that biophilia might include three positive responses to natural landscapes: liking/approach responses; restoration or stress recovery responses; and enhanced high-order cognitive functioning in non-urgent tasks. A related possibility is that prospect can be a component in the affordance of a view that contributes to physical healing. In a study of medical records for post-operative patients who had undergone gall bladder surgery, Ulrich (1984) examined various factors indicating level of recovery. Two groups of patients were compared: those who spent a week recovering in rooms with a window that provided a view of a brick wall, and those who spent their recovery week in rooms with a view onto a grove of deciduous trees. The latter group recovered significantly faster, resorted to fewer doses of high-strength analgesics, and received fewer negative comments on their charts from nurses.

From the perspective of analyzing the role of prospect, this study is not adequate, because it did not compare views of equivalent prospect with different content, but rather compared a view having prospect on vegetation with a view that had no prospect and a built structure. One interesting finding of a related study, however, does point out that a proximate natural environment (in this case a

garden in a children's hospital) can only contribute to well-being if it is known about and utilized by the patients and their families (Whitehouse et al. 2001). In this respect, prospect (perhaps through window views) would have been one means of encouraging awareness of the garden.

At issue here is the degree to which the term "affordance" can reasonably be extended to include the effects of an environment that are not entirely under the conscious control of the perceiver. If an affordance is an opportunity for action for a particular perceiver, perhaps it is unwarranted to suggest that healing, or reduction of mental fatigue, should actually be classed as affordances. However, given a different object related to healing, for example an acetaminophen tablet, it would not be as controversial to say that the drug affords relief from headaches. There is an element of implied volition in that the person has to swallow the pill, but the same could be said of the environment that affords reduction of mental fatigue, in that the person might choose to enter that environment. Cases where people either deliberately leave an urban setting in search of more natural surroundings, or conversely, where people seek natural elements even under the most difficult urban conditions, have both been put forward as likely targets for research on biophilia (Kahn 1999, pp. 113-4).

J. J. Gibson also uses breathing as an example of the affordances of the medium air (J. J. Gibson 1979, p. 130), which suggests that not every affordance has to be under conscious control. Surely if breathing is an activity of living organisms, then healing is also an activity – and as such there must be affordances for it.

Prospect and Aesthetic Appreciation

Appleton's interest in prospect, refuge, and hazard was related to the mechanisms of human aesthetic appreciation, so there is a sense in which the theory is predicated on the relationship among certain kinds of prospect involving particular configurations (and sometimes representational objects), and aesthetic appreciation. Appleton's primary research involved the appreciation of landscape paintings, although his theories were formulated in the context of appreciation of actual landscape. Subsequent researchers have further expanded on the connection between aesthetic appreciation in the two realms, emphasizing that both are complex phenomena, but that despite the obvious differences (e.g., that art is intentionally constructed by people and nature is not), there are many parallels (Matthews 2001).

The question remains, however, of the relationship between landscape painting and affordances. One of the vexing problems in the ecological approach to visual perception is the clear

distinction possible to most viewers, most of the time, between an actual environment and a painting or other pictorial representation. The Renaissance preoccupation with the *trompe l'oeil* is the exception that strengthens the rule; if it were normal for people to be confused as to whether they were seeing a real scene or a painting of a scene, the circumstances under which they could be confused would not have been a source of such fascination.

One answer to the question is suggested by Höge (1990, pp. 111-3), who points out that a primary distinction between the real scene and the pictorial one is that the latter does not actually have the affordances that its imagery suggests. Like Ullman (1980) and other of J. J. Gibson's critics, Höge is interested in re-emphasizing the algorithmic level of perceptual analysis – in this case, in the domain of aesthetic appreciation. To that end, Höge presents findings that demonstrate the influence of pre-induced emotional states on the interpretation of paintings, where participants described the same subjects in opposite terms depending on whether their induced mood was elated or depressed. Höge's observation that pictures are affordance-free might, however, be equally taken as evidence for the validity of J. J. Gibson's assertion that the perception of affordances is a form of direct perception that is, more often than not, *unmistaken in identifying opportunities for action*. The perceptual differences between real scenes and pictorial ones might therefore relate to the primacy of the dorsal perceptual system, in the former case, and the ventral stream, in the latter. If so, Höge's study suggests that it might be possible to demonstrate a stronger link between emotional response and the ventral stream than between emotion and dorsal perception. Such a finding would lend neurophysiological support to studies of expert response to crisis situations, which suggest that one of the factors that distinguishes experienced personnel from those without experience is the ability to postpone emotional reaction until the crisis is over.

OPERATIONALIZING AFFORDANCES

It [the discussion of the relationship between events and affordances (Stoffregen 2000)] certainly has pointed out a serious challenge that the ecological community should take on:

Operationalize the concept of affordance, continue J. J. Gibson's work" (Hecht 2000, p. 62).

Although Hecht's rallying cry is directed at ecological psychologists, it is equally significant for designers. What designers require is a means of discussing how well a particular affordance is met in a given context, so that improvements can be made appropriately to the features that provide that affordance. It would also be useful to have methods for discussing various affordances using a

common yardstick, so that attention and resources could be deployed strategically in developing the most important affordances. Finally, an affordance yardstick could be used in discussions related to whether or not a particular affordance is already being adequately met by existing methods; if the case can be made that it is, then further research could be directed to areas that have more possibility for improvement. An operational definition of affordances, or rather, of affordance strength, could therefore be potentially useful to designers working on a single feature of a product; designers working on multiple features of a product; and managers faced with deciding where resources should be allocated across multiple product lines.

One system for describing the spectrum of possible improvements to existing product designs has currency in the design community as forming a brief summary of the history of product development in the Western world. The questions that describe each position on the spectrum are: “is this product functional?” “is this product usable?” and “is this product a pleasure to use?”

Functional designs are those which can be used to perform a task. They have the necessary affordance, although that is all that can be said for them. A rock the size of a baseball can be used to pound a nail, although it is easy to think of better tools than the rock. Because the rock has not been specifically designed as a tool for pounding, it may not have the right combination of a section suitable for grasping and another flat surface suitable for striking with. It may tend to disintegrate under repeated use. It also requires considerable effort to use for pounding. A high-tech example of a functional design is the standard VCR. It can be used for videotaping from the television, but the interface is notoriously complex and difficult to understand and use.

The next stage is where function is augmented by usability. At the usable stage, the relevant question is no longer “can this be made to provide the affordance required?” or “does this work?” but rather “how well does this work for this person in this context at this time?” A hammer is a more usable alternative to the rock, because the moment arm provided by the handle allows the user to do the job with the expenditure of considerably less effort. The handle of the hammer may also be easier to grasp and hold than the rock was. Since the striking surface in the hammer is removed to a distance from the grasp of the user, there is also a reduction in the possibility of injury. In the case of the VCR, the more usable version is the one with on-screen programming, where the interface is the television rather than the VCR panel. An on-screen interface allows feedback at a larger size, which may be a significant factor in improving usability for some people. It can take advantage of the remote control

as a physical interface device, which allows the user to interact through a system that may be more familiar because it is more frequently used than the VCR panel. It also allows for virtual rather than physical design, which can have practical implications in terms of reducing manufacturing restrictions on the development and implementation of improved design.

The final stage in the chronology of tool development is where usability has been sufficiently established that attention can be given over to pleasure. In the case of the hammer, the comfort of the grip now becomes a central design issue, since every available hammer already has a proper handle and appropriate head. The hammer as an all-purpose pounding device may be superseded by specialty hammers which accommodate various user needs in terms of arm strength or grasp pressure. A pleasurable generation of VCRs might include features such as automatic TV program detection and capture, so the user could in effect ask the system to record the next episode of the program of interest, regardless of which channel broadcasts it or when.

Another purpose for an operational definition of affordance strength might therefore be in contributing to the understanding of where a given tool sits on the spectrum of functional, usable, and pleasurable, why it is positioned in a particular place on that spectrum, and perhaps most importantly, to what extent the affordance as a whole matters either to a given perceiver or to a larger segment of society.

In order to operationalize the concept of affordance strength, it is useful to determine the extent to which it is first necessary to formalize the concept. Previous work in this area has ranged from the attempt to establish theoretical mathematical models of affordances as dynamic systems (Van Leeuwen and Stins 1994) to pragmatic definitions of affordances in terms of measurements of the relevant physical dimensions of the perceiver and some aspect of the environment (Warren 1984).

Previous Theoretical Definitions: Dynamic Systems

Van Leeuwen and Stins (1994) outline a strategy for describing affordances as parameters in an environment that is modelled as a complex dynamic system. One goal of this formalization is to establish a role for perceiver intention, which, for Van Leeuwen and Stins, has a critical role to play as a factor of higher-order affordances. Another goal of the dynamic system model is to incorporate reflexivity, where the environment is influenced by previous actions of other people, and that influence can result in new or modified affordances. The significance of various time scales is also included in the dynamic system, since different parameters can be treated either at a microscopic level

(with minimal effects on the larger activities of the organism or its environment) or at a control level, where changes to the larger dynamic system related to the affordance are played out with lasting effects to the organism or even to its evolution (depending on which scale is used).

The primary disadvantage of the model of Van Leeuwen and Stins is that it is largely unsuited to function as a tool for the designer. The model is primarily intended for understanding the role of affordances in the actions and evolution of an organism in a particular environment. As such, it is not specifically related to the question of how an individual affordance can be analysed in its own right, although it is interesting to think that it may provide the basis for a mathematical model that could be implemented as a computer simulation of the more complex activities of an organism.

Previous Operational Definitions: π Numbers

Since the purpose of affordances is to discuss the relationship between the perceiver and the possibilities for action in the environment, some of the previous studies of affordances have introduced the idea of measuring affordances by using a dimensionless ratio, where the numerator stands for some property of the environment and the denominator signifies a measurement of the corresponding effectivity of the perceiver using the same units of measurement (so that the units cancel each other out). These dimensionless numbers are called π numbers.

The classic study of this kind is Warren (1984), who examined perception of climbability of stairs. The ratio used in this case was the riser height over the leg length of the perceiver. Warren determined that a critical threshold boundary could be identified to distinguish stairs that were climbable from those that did not afford climbing. This threshold occurred at $\pi = 0.88$, which is to say, at the point where the height of the riser was 0.88 of the length of the perceiver's leg. He also wanted to find out whether participants were capable of visual estimates of optimal ranges for minimizing energy expenditure in stair climbing (they were), and whether the resulting π number would be a constant independent of observer height (it was – about 1/4) (Warren 1984, pp. 698-9).

π numbers have also been found for a number of different animals, including limpets (at what size do they stop fleeing from predatory whelks and start attacking?), frogs (how does their body size relate to their willingness to try jumping through an aperture?), and praying mantises (what is the relationship between reach and prey radius for high-frequency attack responses?). The concept of measuring thresholds of affordance using ratios of relevant physical qualities is useful, but it has limited applicability in the following situations:

- the affordances in question are not physical
- the affordances are physical, but bodily dimensions are not relevant
- the goal is to analyse all the relevant information about affordances, rather than to determine a threshold level between affordance and non-affordance, or to study the perception of levels of optimum activity.

The Affordances in Question Are Not Physical

One shortcoming of Warren's approach is that the dimensionless number π is not appropriate in cases where the affordances in question are not physical. By definition, this excludes cognitive, interpersonal, and cultural affordances. The affordances of language alone suggest that these areas comprise a significant realm of human opportunities for action, which any operational description of affordance should accommodate.

The Affordances Are Physical, But Bodily Dimensions Are Not Relevant

Ratios of object to body scale are also not useful for situations where the affordance is physical but the dimensions of the body of the perceiver are not a primary factor. Air, for example, usually affords breathing, but the question of the dimensions of the person doing the breathing do not seem particularly relevant. It may be possible, however, to adapt the idea of π numbers to accommodate cases where the measurement still involves some relationship between human capacity and a feature of the environment. The suggestion has been made, for instance, that time-to-impact studies might be usefully recontextualized in terms of some metric such as escape margin or catch margin, either of which would situate the time to impact within a particular framework relevant to a human participant (Hecht 2000, p. 60), although not necessarily to the size of a particular part of the body.

It is also possible to extend the idea of measurement into more complex physical forms. Air requires a certain amount of oxygen in order to support respiration, and it needs to be at a certain pressure and cannot contain various lethal constituents above given thresholds of human tolerance, and so on. The human respiratory system also has subcomponents that all need to be present and working above a given threshold of capacity needed to sustain life, and the body itself needs to support the respiratory system. Given this complexity of the subject and object, it might be possible to construct some composite numbers that adequately describe the dynamics of the affordance of air for

breathing; this mathematical model, however, would still address only the physical dimensions of the situation.

The Goal Is to Analyse All the Relevant Information about Affordances

Finally, and perhaps most importantly, any model restricted to physical measurement seems potentially to ignore a range of significant information, including, in the case of stair climbability, such factors as the perceiver's age, state of health and strength, energy levels, intentions, goals, narrowness of the stairs, the presence or absence of a railing or banister to lean on for support, potential hazards such as frost or liquid spills, staircase landings or other possible resting areas, and so on.

Relational Factors of Affordances

It may be possible, however, to design a more general form of π number that meets these objections. In keeping with the spirit of J. J. Gibson's rejection of dualism, the first principle in such a definition should be that the factors involved represent in some way, not the individual characteristics of the two participants, but aspects of the relationship. In this respect, the classic π numbers used by Warren and others have the disadvantage of being measurements of the first kind, while the relationship between the numbers is indicated by their use in a ratio. It is possible, that is, to measure leg length without reference to riser height, and vice versa. Π numbers, as defined by Warren, are therefore primarily measurements of the subject and object, and only in a secondary sense is the relationship suggested by forming a combination of the two in a ratio.

Ideally, an operational definition of the strength of affordances should allow only for measurements of various aspects of the relationship, rather than measurements of the perceiver and some quality of the object or feature of the environment being perceived. These measurements of relational factors should be specific enough to capture the various kinds of information that are relevant to both the perception and use of the affordance by a given perceiver at a given time, and yet general enough that they will apply to all the different kinds of affordances, whether static, kinematic, dynamic, physical, cognitive, interpersonal, cultural, or any other.

In order to factor the strength of affordances appropriately, it is therefore necessary first to establish which are the necessary components that are specifically related to the relational nature of the perceiver and the perceived environmental feature. It is also necessary to establish to a satisfactory

degree that the list of components is sufficient to serve as a pragmatically useful measurement of a situated potential for action for a given perceiver, at a given time.

A factor is relational if it does not make sense to discuss it outside the context of a particular affordance. For example, the primary affordance of a pen is that it can be used to write on a piece of paper. Within the context of using a pen for writing, it is reasonable to talk about whether or not it has ink in it, and if so how much ink, and whether or not the pen allows the ink to flow out in a smooth stream onto the paper. It is also reasonable to ask whether the pen affords grasping by a particular person who wishes to write with it. However, if the person is looking for a pen in order to use the side of it as a straight edge in order to draw a straight line, then the amount of ink in the pen and how it flows are irrelevant. If the affordance of the pen is that it can serve as a straight edge, then graspability is still a relevant factor, along with the length of the straight section and the smoothness of the pen shaft; the ink levels and flow characteristics, however, become irrelevant.

Given the need to specify the significant relational factors that characterize the strength of an affordance, it is possible to distinguish eight factors that together represent the relational aspects of the object, the perceiver, and the dynamics of the context. These factors together can be used to create a vector space that defines the relational aspects of affordance strength in an operational way.

For example, if a given adult wishes to keep dry while walking two blocks in the rain, the unfactored affordance of the object is the twin capacity to be carried while walking and, simultaneously, keep someone dry. The object in question might be anything that is large enough to cover at least the top surface of the head, light enough to be held up there, and impervious to rain. A range of objects are possible, from specialized devices like umbrellas, to makeshift ones such as newspapers or briefcases, to the objects of last resort, such as the back of the coat pulled up over the back of the head or a covering made of the two hands. Although it is possible to measure objective features of the various candidate objects, such as their size, weight, slope, imperviousness to water, tendency to sustain water damage, monetary value, and so on, each of these features is only important in this situation because the person wants to stay dry while walking two blocks in the rain. For practical purposes, it may therefore be sufficient to aggregate these features into one larger relational factor that represents how well the object can perform the task at hand.

The first necessary factor is therefore the tacit capacity of the object to provide the affordance in situations of the kind being studied. In this case, the tacit capacity of the umbrella in

situations where a person needs to walk two blocks in the rain while staying dry would be very high, while the tacit capacity of, for example, a wrench, would be zero. The wrench has an excellent tacit capacity for other types of actions. In fact, because it is a specialized tool (like the umbrella), it has a primary affordance. But for the work at hand it is useless.

It is possible but not necessarily helpful to subdivide the tacit capacity into sub-features such as the weight of the umbrella or the slope of the dome or the nature of its fabric, since the perception of the tacit capacity is in a sense given. Every adult knows that umbrellas have this affordance; that it is, in fact, their primary affordance. The direct perception of the affordance is also a central point of J. J. Gibson's approach. It would therefore only be helpful to address these sub-features in cases where the tacit capacity is open to contention. An example of this kind of situation might be at the occasion of the original purchase of the umbrella, where factors such as expense vs. utility may need to be considered.

The second necessary relational factor is the situated potential of the object, not generally in circumstances of the kind under investigation, but in one particular situation at one particular time. It is all very well for the person about to walk in the rain to realize that an umbrella has an excellent tacit capacity for keeping a person dry, when at the point of setting out there is no umbrella available, or the umbrella that is available is torn.

These two factors – tacit capacity and situated potential – are relational attributes where the attention of the researcher is directed toward the object or environment and its relevant affordances for action. There are other factors that treat the relational aspects of the agent, where the researcher's attention is directed at what have been called the perceiver's effectivities.

The first of these factors is awareness. For the person about to walk in the rain, a perfectly good umbrella might be sitting to hand, but if the person is distracted or confused or in a rush, the umbrella might not be perceived, and for all of its high tacit capacity and situated potential, the umbrella still stays dry while the person gets wet.

The second factor is motivation. If the person in question wants to walk in the rain and would prefer not to get wet but does not really mind it all that much, that person's tendency to seek and adopt an available affordance is significantly reduced in comparison with the person who hates getting wet, has just had a cold, and is wearing clothes that will be damaged by the rain. The former person may casually take up an available umbrella if one were available, since the tacit capacity and situated

potential are high enough that the action has an appropriately low resource load. If only a newspaper is available, the lower tacit capacity might be such that the person would prefer to simply get rained on. For the latter person, it is likely that the high motivation and absence of an umbrella would lead to extremes of behaviour such as deciding not to walk but take a taxi instead, or perhaps going back into the building to see if an umbrella could be found somewhere.

Like many of the other factors, motivation is a composite of a wide range of sub-factors, including the whole complex terrain of personality traits and their expression under various circumstances; previous experience or behavioural conditioning; and perception of risk and the tendency to either accept or avoid it when perceived. In spite of the complexity of the terrain, however, it is not unreasonable to ask someone with respect to a given scenario: “how motivated would you say you would be to carry out such and such an action, on a scale of zero to five?”

The third relational factor that is associated with the perceiver is ability. For a person with a physical disability that makes grasping difficult or lifting the arm problematic, the option of carrying anything above the head may simply not be available. In this case, all the other factors may be present, including an umbrella with high tacit capacity and an excellent situated potential, a strong awareness of the umbrella on the part of the perceiver and a correspondingly strong motivation to use it. But inability to grasp the handle renders the affordance zero for this particular person at this particular time. Ability is related to a variety of issues discussed earlier, including the sociocultural aspect of the perceiver being able to recognize and use new affordances through training. Like the person who has difficulty grasping, the infant who has not yet learned the use of hands is not able to either recognize or use a grasping affordance. Another factor in ability is the current condition of the perceiver: a person suffering from extremes of fatigue, hunger, or thirst, for example, is less able than the same person when not so afflicted. Talent, natural proclivity, and intelligence of various kinds are also involved, especially if intelligence is construed in the broader sense suggested by recent educational researchers, who have identified as many as nine distinct kinds of intelligence.

The last factor related to the perceiver represents the role played by individual preference. All other factors being equal or even roughly equal, it is often the case that individual adoption of affordances depends at least to some extent on established preferences. In the case of the person who wants to stay dry in the rain, if there are two umbrellas available and one is a favorite, that will probably be the one that gets employed. Preference can be based on any one of a dozen sub-factors,

ranging from aesthetic considerations to interpersonal influence to previous personal experience. Preference is distinct, however, from ability, and although preference is related to motivation, the two are not equivalent. A person might be highly motivated, for example, to perform an action that should probably not be correctly characterized as a preference, as when soldiers fling themselves on live grenades in order to save the lives of their comrades.

The final factors in the proposed vector space are needed in order to adequately account for features of the situation that are relevant but are not directly related to the relationship between the perceiver and the object. They stand instead for the relationship between the affordance and its context. The first of these factors is contextual support, where factors in the environment that are not part of the affordance have an influence one way or the other on the perceiver's interaction with the affordance. There are a wide range of possible contextual supports, including aspects of the situation that are physical, cognitive, and environmental, and the precise nature of the contextual supports in a given situation should be outlined during the process of analysing the affordance as a whole.

In the example of someone who wishes to stay dry in the rain, the contextual factors would include environmental facts such as how hard it is raining, whether it is warm or cold outside, how hard the wind is blowing and in what fashion, and so on. If it were raining hard and was cold enough that the rain was almost turning into sleet, and the wind was blowing hard in a fairly horizontal direction, then this context renders the umbrella's affordances virtually useless. On the other hand, if the sun is shining through the rain and it seems likely to clear within a couple of minutes, the perceiver's motivation to find an umbrella or some other object with appropriate affordances may be dramatically reduced in favor of the strategy of waiting for the rain to stop.

The definition of the context factor as one providing support is important in order to keep the list of vectors homogeneous. An alternative definition might use the idea that contextual factors should be characterized in terms of their interference with a particular affordance. However, since the other factors are all framed as positive elements in the affordance, it makes sense to approach the context in the same way.

The other feature that has not been accounted for yet in an explicit form is the role of other agents in the scenario. Contextual support includes all those factors (excluding the affordance itself) that are present in the environment at the time of the perceiver becoming involved with the affordance. Agential support, on the other hand, includes those features relating to the roles of the

other people, animals, insects, and so on who are also potentially part of the situation. Agents are distinct from other factors of the environment in that they have agency, which is to say volition, goals, and actions of their own, which may have some bearing either directly or indirectly on the particular affordance.

For instance, for the person who wishes to stay dry in the rain, it may turn out that there are other people present who also wish to walk outside. One of them might be elderly or frail and lacking an umbrella, in which case our perceiver could be motivated to behave altruistically and turn over the superior affordance of the umbrella to the other perceiver, choosing instead an inferior solution such as a folded newspaper.

As in the case of contextual support, agential support is defined for the purpose of an operational definition of affordance as a positive factor, in keeping with the definitions of the other factors.

The Strength of Affordances as a Vector Space

One operational definition of affordance strength is to take these relevant factors about the relationship and its supports and use them to create a vector space. Vector spaces are a way of positioning information on a multi-dimensional co-ordinate system by providing an ordered set of numerical values, each of which stands for a significant dimension of the larger universe under consideration. In this case, the proposed universe is the one containing, not all possible affordances, but rather all possible strengths of affordances. The proposed grid exists in eight dimensions, corresponding to the relational factors that have been discussed. In equation form, the vector for affordance strength would therefore be as follows:

$$\text{affordance strength} = (\text{tacit capacity, situated potential, awareness, motivation, ability, preference, contextual support, agential support}).$$

Interaction of the Vectors

Although, from one perspective, the factors involved in the affordance vector are relatively independent of one another, in another sense it is possible to identify mechanisms whereby the various factors interact. For example, there is a potential inverse correspondence between tacit capacity and ability, in cases where improvement to the capacity of the object to provide a given affordance

increases the complexity or novelty of the object in such a way that the ability of the user is adversely affected.

A classic example is provided by the cockpits of jet aircraft, some of which are sufficiently complex that they have reached the thresholds of human capacity to monitor all the relevant instruments. Although the tacit capacity of the cockpit instrumentation to afford information to the pilots is very high, the ability of a non-pilot to receive the information and carry out the appropriate actions is reduced in proportion to that capacity. Only through extensive training and experience have qualified pilots been able to develop a level of ability that corresponds to the tacit capacity of the instrumentation.

Another pair of factors that can influence each other are motivation and agential support. A given person can either be encouraged to act or discouraged from acting by other people. Agential support can also have a paradoxical inverse effect, as in the case when a person acts out of a sense of rebellion against social expectations. Although a particular behaviour might have strong agential support in the form of interpersonal or cultural approval, for the person motivated by a spirit of rebellion, the very strength of the agential support can serve to reinforce the negative preference.

Vector Values

There are several options available for assigning values to the different factors in the affordance strength vector, but perhaps the simplest method is to choose a common Likert scale that can be used for all the factors. Each of the items might be rated, for example, on a scale of 0 to 5, where 0 means the affordance factor is such that the entire affordance is rendered null, and 5 means that the affordance factor is as strong as it needs to be for all practical purposes.

Likert scales are a form of ordinal (or ordered) scale, which means they are useful in discerning difference. However, because it is difficult to establish that various respondents agree as to the precise meaning of the anchor values, Likert scales are not usually treated as interval scales. That is, the distance between a zero and a one is not necessarily the same as the distance between a two and three.²

² If the intervals between the individual items were equal to each other, then it would be meaningful to calculate a mean score for all respondents. An example of a scale of this kind would be one developed using the Thurstone method of equal-appearing intervals. The Thurstone method assembles statements from participants on a common topic, then collects ratings that place the statements at equal intervals on a scale.

The primary advantage of a Likert scale is that it is easy to apply. The primary disadvantage is that it reduces what may be fairly complex qualitative information into a simple number, rather than preserving the complexity. It also requires that the evaluation be carried out in terms of a choice between one whole number and another, rather than as a point on a full continuum.

In short, the simplicity of Likert scales is simultaneously their strength and their weakness. One means of reducing the weakness is to capture additional information that is more qualitative by allowing the people using the scale to provide comments, either as written addenda to each question or else in the form of an interview.

In terms of the design of the Likert scale, there are several decisions that need to be made concerning the relative values of each item. If the same numeric scale is used for the different factors, then they each count as equivalent elements in the whole ranking. An alternative strategy would be to weight some of the factors so they contribute either more or less than the others. One means of adjusting the weights only slightly would be to adopt different Likert scales for various factors. Another stronger weighting strategy would be to use the same scale but add a multiplier. There seem to be no obvious *a priori* reasons for choosing to weight one factor over another, although subsequent research may indicate that such a system would be more accurate.

It would also be possible, of course, to establish more complex research criteria related to each of the factors, so that values might be assigned through decomposition of each factor into sub-factors that were subjected to rigorous study, then aggregated to create a total. The introduction of this additional level of complexity should be reserved, however, until such time as the simpler method proves insufficient.

Using the vector space based on a six-point Likert scale, it is possible to have an individual person evaluate a particular affordance in a given situation. The assessment will be more convincing, however, if it is performed by a larger number of people who have equivalent characteristics in the relevant aspects of their profiles. It may also be useful to have ratings both from the actual participants and from observers, who might provide a form of reality check.

Vector Anchors

In defining Likert scales for the various factors, one approach is to label each of the numbers on the scale with its own anchor text. This strategy provides the user with a maximum amount of specific information concerning the intended meaning of each value. In some cases, however, it is preferable to label only the extreme ends of the scale. Providing evaluators with only the extreme anchors can result in some minor variation in interpretation of the intermediate values, but has the advantage of making the task less demanding. If they are not required to read the text on each value, study participants are able to react more naturally to the implicit ranking suggested by the numbers. In an ideal situation, the task would be even further simplified by having the same anchors apply throughout the vector space. However, because the factors in the affordance strength vector differ from each other quite dramatically, the following discussion provides the finest level of granularity, with anchors spelled out explicitly for each point on the scales for the various factors.

Tacit Capacity

In many cases, tacit capacity may be one of the most difficult of the affordances to evaluate. Where the object in question is a dedicated tool with a single primary function, the situation is relatively straightforward, but even in cases of this kind the individual variation among different evaluators may prove to be significant. Part of the reason for predicting disparity among perceivers is that there is a mainstream cultural bias toward emphasizing product feature variation in western capitalism. Minor differences among dedicated devices form part of the niche approach to marketing that drives the economy of the western world, and as such they tend to receive a high degree of attention. It may prove difficult, in fact, to separate evaluation of tacit capacity and individual preference. Given these reservations, a Likert scale for tacit capacity might use the following anchors:

- 0 – useless
- 1 – very poor
- 2 – poor
- 3 – acceptable
- 4 – good
- 5 – great

Situated Potential

Inter-evaluator perceptions of situated potential, on the other hand, seem likely to vary less significantly, since there is no comparable cultural mechanism in place to emphasize different values for what is ready to hand. The evaluation of the situated potential of an affordance primarily consists of its proximity to the perceiver, although there are possible confounding circumstances in special cases, as when a tool can be seen but not grasped because a fence is in the way, or where it is visible but out of immediate reach on a high shelf. A Likert scale for situated potential might use the following anchors:

- 0 – not available
- 1 – available with extreme effort
- 2 – available with considerable effort
- 3 – available with some effort
- 4 – easily available
- 5 – effortlessly available

Awareness

Awareness is also a fairly complicated factor. The argument could be made that it is inappropriate to suggest a scale for awareness at all, since it is by nature a condition with two possible states – either a perceiver is aware or unaware, and that is the end of it. However, any model of awareness needs to account for phenomena such as priming, the tip of the tongue effect, and the various shades of suspicion leading to full conviction. For example, a person might have a nagging feeling that there is a screwdriver in the house, based on the priming of an unconscious or subconscious memory of having seen a screwdriver somewhere recently. Alternatively, the person might know there is a screwdriver in the house but not know exactly where it is. Finally, the person might suspect that there is a screwdriver in the top left drawer of the kitchen counter without having a full conviction that there is a screwdriver there. Given this range of possibilities, a Likert scale for awareness might incorporate the following anchors:

- 0 – completely unaware
- 1 – unconsciously aware
- 2 – consciously suspicious
- 3 – dawning awareness

4 – developing certainty

5 – fully aware

Motivation

A Likert scale for motivation would need to allow for conditions ranging from a degree of motivation that is effectively non-existent through to a strong, immediate desire to accomplish the action in question. The suggested anchors are:

0 – will not act

1 – will act under coercion

2 – grudgingly willing

3 – willing

4 – highly motivated

5 – absolutely determined

Ability

Ability is a complex factor that may involve: prior learning; experience; physical or mental qualities such as dexterity, strength, or determination; age; health; and even predilection and talent. A Likert scale for ability might use the following anchors:

0 – incapable

1 – beginner

2 – novice

3 – intermediate

4 – advanced

5 – expert

Preference

Preference might at first seem relatively straightforward, but the elicitation of preferences is actually a study in itself. One of the problems is that people are not necessarily conscious of their own preferences and will have a tendency to rate themselves in ways that are significantly different from their observable behaviors. Self-image plays a role. It is also necessary to avoid the observer expectancy effect, where participants in a study attempt to second-guess the researcher by providing

an answer that will be correct or pleasing. A simple Likert scale for preference might use the following anchors:

- 0 – avoid at all costs
- 1 – avoid if possible
- 2 – grudgingly acceptable
- 3 – acceptable
- 4 – preferable
- 5 – strongly preferable

Contextual Support

Contextual support is difficult to define in detail and is subject to so many potential factors that it would be impossible to list them all. Within a given circumstance, however, it should be possible for an evaluator to identify the contextual features that seem to be relevant for a particular affordance and assign them a composite value. The following Likert scale might provide a framework:

- 0 – complete interference
- 1 – partial interference
- 2 – minor interference
- 3 – neutral
- 4 – partial encouragement
- 5 – full encouragement

Agential Support

As with contextual support, agential support is composed of any number of possible sub-components relating to the activities of other people or sentient creatures. A Likert scale similar to the one used for contextual support may prove useful, consisting of the following anchors:

- 0 – aggressive interference
- 1 – partial interference
- 2 – minor interference
- 3 – neutral
- 4 – partial encouragement
- 5 – active encouragement

Affordance Vector Worksheet

Based on the anchors described above, a generic affordance evaluation worksheet could be constructed to serve as the basis for studies of a particular design, object property, or environmental feature. The questions indicated below might be an appropriate starting point for a worksheet for an existing interface feature. The sections for comments allow some qualitative information to also be collected.

One method of using this strategy would be to have participants paired, with one using the feature and the other watching. The participants would fill out separate worksheets, in order to provide the judgments of both the user and someone observing the user.

For the following items, please rate the feature on a scale of 0 to 5.

1. tacit capacity

How well would you say the feature works?

0 – useless 1 – very poor 2 – poor 3 – acceptable 4 – good 5 – great

tacit capacity comments: _____

2. situated potential

How easy would you say the feature is to access?

0 – not available 1 – available with extreme effort 2 – available with considerable effort 3 – available with some effort 4 – easily available 5 – effortlessly available

situated potential comments: _____

3. awareness

How aware were you of the feature and its use?

0 – completely	1 – subliminally	2 – consciously	3 – dawning	4 – developing	5 – fully
unaware	aware	suspicious	awareness	certainty	aware

awareness comments: _____

4. motivation

How strongly would you be motivated to use such a feature?

0 – will not	1 – will act	2 – grudgingly	3 – willing	4 – highly	5 – absolutely
act	under	willing		motivated	determined
	coercion				

motivation comments: _____

5. ability

How would you rate yourself as a user of this kind of feature?

0 – incapable	1 – beginner	2 – novice	3 – intermediate	4 – advanced	5 – expert
---------------	--------------	------------	------------------	--------------	------------

ability comments: _____

6. preference

If you had a choice of this among other features that provided the same function, how would you rate your personal preference of this feature?

0 – avoid at	1 – avoid if	2 – grudgingly	3 – acceptable	4 – preferable	5 – strongly
all costs	possible	acceptable			preferable

preference comments: _____

7. contextual support

How well is this feature supported by the context in which you would normally be using it?

0 – complete 1 – partial 2 – minor 3 – neutral 4 – partial 5 – full
interference *interference* *interference* *encouragement* *encouragement*

contextual support comments: _____

8. agential support

How much support would you have from other people if you wanted to use this feature?

0 – aggressive 1 – partial 2 – minor 3 – neutral 4 – partial 5 – active
interference *interference* *interference* *encouragement* *encouragement*

agential support comments: _____

CHAPTER 2: THE DIGITAL AFFORDANCES OF PROSPECT

... interface design is a fragmented endeavor in which design solutions lack any substantive coherence within or across work domains. Affordance theory has the potential to change that by bringing to the endeavor a unifying theoretical structure (Lintern 2000, p. 68).

There are a variety of new opportunities for action that can be made available to users of digital collections through rich-prospect interfaces. These new affordances are based on the direct visible presence of information about the contents, structure, and other significant features of a collection, such as how it was understood by its developers, how it has been organized, and, in some cases, how it has been encoded with additional interpretive material that is not contained in the actual text. Although this visible information can make a significant difference in terms of user perception, by itself it is not sufficient to provide many new affordances. In order to be of greatest value, a rich-prospect interface needs to also provide a set of appropriate tools that can take advantage of the visible representation of the items in the collection and the other collection features.

This chapter is subdivided into two sections. The first section begins with a brief discussion of the composite affordances that are related to prospect, and continues by examining some of the implications of interpreting prospect as implying a relatively literal digital implementation of the landscape metaphor. The second half of the chapter deals more precisely with the issues relating to rich-prospect interfaces, including the meaningful representation of items, the kinds of insights about the collection that are potentially made available, interface tools related to prospect, the incorporation of prior affordances, characteristics of candidate collections, and the discussion of a variety of rich-prospect interface design issues.

COMPOSITE AFFORDANCES

The opportunities for action related to prospect in the analog world are not exclusively provided by prospect itself. Instead, prospect is usually involved as a component of a larger, composite affordance. For instance, Appleton (1975) identifies a number of survival-related activities that can be facilitated by prospect, including the identification of potential sources of:

- water
- food
- shelter
- concealment

Wayfinding is another affordance related to prospect in the analog world, where opportunities for subsequent navigation of the environment may be provided by an unobstructed view on the terrain. Some of the other affordances involving prospect relate to territoriality, where the perceiver is watching over a scene in order to subsequently assert or re-assert control. The opportunities for action that a territorial perceiver has available once a potential encroacher has been located are dependent to a large extent on what other resources are available, and those opportunities have to be gauged in comparison with the resources available to the interloper.

Finally, there are analog affordances of prospect onto natural landscapes, especially of the savanna type, that relate to the biophilia hypothesis. Biophilia is a concept formulated by biologist E. O. Wilson (1984), which suggests that human beings may be biologically prepared to learn about particular aspects of their environment that were at some point crucial for survival. These affordances relate to potential health benefits of prospect onto natural environments of an attractive kind, various affective responses (either of attraction or repulsion) to various kinds of environments (which responses may be related to biologically-prepared learning), and consequences of prospect in terms of the reduction of some of the sub-factors of mental fatigue, such as a lessened tendency to engage in acts of aggression.

Each of these affordances is a composite that includes prospect, the characteristics of the landscape, and the intentions of the perceiver. The absence of any of the elements can render the affordance null. For example, it makes no sense to speak of actions of a territorial kind if the perceiver does not consider the landscape in a proprietorial way. Similarly, if there are no sources of food in a given location, then prospect on that location does not afford the identification of potential food sources (although it can afford perception of their absence). Finally, a perceiver who is looking for resources may be able to identify food, water, shelter, and so on using other means than looking out from a position of prominence, but without the element of prospect, the information will tend to be less complete because it lacks framing within the larger context.

In this respect, prospect is not a sufficient condition for any of the affordances listed. For some of the affordances, prospect is not even a necessary condition. However, when the provisions of the landscape and the intentions of the perceiver do coincide, then the availability of prospect as the third factor in the equation does provide a significant dimension that is not otherwise present. Similarly in the case of the affordances related to biophilia – some degree of prospect is beneficial. To

have a view onto a deciduous grove may ease the mind; conversely, to have a view onto the trunk of a tree growing directly outside the window may be a source of frustration.

A similar principle holds true for the affordances that are related to prospect in the digital environment. Prospect should perhaps most usefully be considered as a component of a number of larger composite affordances, which will vary significantly based on the nature of the perceiver's intentions and the characteristics of the landscape – or in this case, the intentions of the user of the interface and the characteristics of the collection that the interface is intended to provide prospect into. The detectable presence of affordances can also, of course, serve as a form of encouragement to undertake actions that were not necessarily part of carrying out the original intention.

THE LANDSCAPE METAPHOR

The currently prevailing metaphor of the graphical user interface is the one developed for the Xerox Star – the desktop metaphor. It includes icons that represent files and folders, which function in a manner analogous to real-world files and folders. It also includes elements that are distinct to the computer interface and are not analogous to real-world items: for example, menu bars, pop-up dialog boxes and help balloons, and scrollbars. The office desktop metaphor allows the user to learn to work in the digital environment by analogy with the real-world environment. It also helps to establish the discourse for the computer as a tool that is appropriate to be in the office, as opposed to alternative discourses which placed the computer in the laboratory, where it performed esoteric calculations for physicists and mathematicians, or in the military, where it served as a tool for calculating ballistics, as well as for enciphering and deciphering intelligence communications.

If the landscape metaphor is implemented on the computer within the constraints of a fairly literal interpretation, where various visual elements on the screen correspond to visual elements in a landscape, then it seems likely that it will meet with resistance from the business community. The landscape cannot serve as an adequate substitute for the business office, because the discourse of the landscape is not an office discourse. It would be no more appropriate to put a form of landscape on the office desk than it would be to try to incorporate some form of physical landscape into the business world, and move the desks, as it were, out under the open sky. Where the landscape has seen widespread development in the digital environment is in computer games, which do not suffer from the necessity of sustaining a level of perceived seriousness, but are free instead to function as fully-fledged digital environments.

The discourse of the computer as an office-based tool has also been relaxed somewhat in the last twenty-five years through the efforts of various manufacturers and retailers to create a position for the computer as a home-based device. One of the most successful strategies in the last ten years has been to emphasize the role of the computer as a recreational and informational tool for the home, often designed specifically as a means for accessing the internet. This positioning strategy provides the computing industry with leverage into a variety of established consumer areas, including entertainment, news, shopping, communication, and so on. The colourful Macintosh shellders are an example of the marketing of the computer as a home-based, rather than office-based device. If the landscape metaphor is to find widespread adoption, it is most likely to be in relation to the computer as a device for the home.

However, even assuming that the desktop metaphor continues to predominate, there are still a number of ways in which less-literal interpretations of the landscape metaphor, and in particular the corresponding affordances that involve prospect, have been and may continue to be adapted for use without the need to radically challenge the prevailing discourse.

Maps

A map is a prospect-based artifact. It provides in a portable form a part of what the view from a prominence provides in a fixed form. In addition to being portable, it has the advantages that it can be tailored to emphasize particular information about the landscape, through various techniques that vary optical weight or tonality, as in the use of contrast, colour, textual labels, and scale. Its primary disadvantages are similar to the disadvantages of prospect in general – namely, that the level of granularity of detail is not always sufficient to meet the intentions of the perceiver.

There is, in maps as in the analog world, a natural trade-off between prospect and detail, which relates in the case of the map to the size and complexity of the printed artifact, and in the case of the analog world to the visibility of various details of the landscape from a given range and position. The kinds of information available from the prospect-based artifact and the situation of analog prospect may also differ, depending on the intentions of the people who made the map. As a general rule, maps are intended for wayfinding, which means that location cues are more significant than affordance cues. A perceiver in a position of prominence over a landscape may look at features such as potential sources of shelter, danger, food, or water, with the idea in mind that some of these affordances may prove helpful. On the other hand, a person in viewing a map may also be looking for

sources of shelter, danger, food, or water, but will be seeing them primarily from the perspective of location, which is the main information the map can provide. Cues as to quality of the potential sources are not necessarily present in the prospect-based artifact.

In summary, a map presents selected information in a standardized format, with conventional visual elements. An alternative prospect-based artifact is the aerial or satellite photograph, which often requires expert interpretation to distinguish the different component elements and what they signify. Aerial photographs contain a richness of visual detail that can paradoxically serve to obscure the prospect.

Maps have been designed and used for a wide range of purposes, from their default functions in wayfinding to specialist uses in summarizing political, economic, military, climatic and other data, sometimes in time series that indicate movements of people or resources over extended periods. The juxtaposition of flow maps that express the same kind of information at discrete intervals can be used to create compelling visual narratives (Figure 2.01), especially when they are combined with caption texts that highlight for the reader the significance of the images (Monmonier 1993, pp. 189, 242).

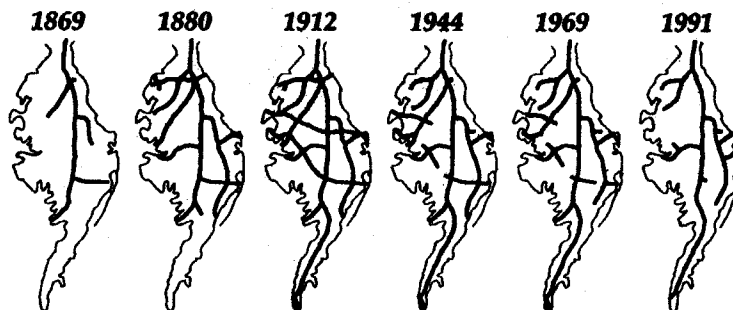


Figure 2.01 A sequence of maps can serve as a visual narrative. Monmonier's map of railroad lines on the Delmarva Peninsula, from 1869 to 1991, illustrates the changing commitment to rail which had its peak expression just after the turn of the century (Monmonier 1993, p. 216). © 1993 by The University of Chicago. All rights reserved.

On a standard road or city map, the information that has been selected is specifically intended to assist in wayfinding. Different qualities of road surface are emphasized by thickness of line; rivers and other natural features are often colour-coded and marked by contours and labels; landmarks are sometimes specified by name; the directions are spelled out; scale is indicated; and so on.

What is missing from a map are all the elements that are considered unimportant for wayfinding, or which are too detailed or otherwise problematic to be practical. Individual dwellings,

for example, are usually not visible, and if they are visible they are not labelled. The purpose of the map is to direct the user to the correct block, after which the standard lot-numbering systems used for urban planning (at least in most North American cities) are expected to serve.

Virtual Datascapes

In the realm of interfaces onto collections of documents, the concept of a prospect-based artifact analogous to the map has been investigated in a number of ways. The most direct implementation is through the design of three-dimensional virtual landscapes, where digital elements are substituted for conventional landscape features. A document might be represented, for instance, by a structure that visually resembles a vertical block, which in its proportions resembles a featureless building. This data building can be scaled according to the size of the document being represented. It can also be juxtaposed with other documents to create a cityscape of data buildings.

There are a number of difficulties that need to be addressed before this strategy can be widely adopted. First, it directly attempts to replace the desktop metaphor with a landscape metaphor, which brings with it the connotations of the computer game and the inappropriate juxtaposition of the landscape with the office. Second, it is at one and the same time a form of visual overkill and visual impoverishment. A virtual landscape is a complex form of visual information, and that complexity is unnecessary, given the kinds of data that can be conveyed using the metaphor. Document size is a significant fact about a document, but from the perspective of the importance of the actual content of the document, it is not a very important fact. To scale data buildings according to the length of the documents they represent is to visually emphasize a comparatively trivial aspect of the collection. Building size, however, does not seem to provide an immediate analogy to any other document feature.

The virtual landscape is a form of visual impoverishment because although it resembles an aerial photograph (or perhaps an aerial video), the selective nature of the information makes it more like a map. Analog buildings have structural characteristics that are significant and interpretable. Digital buildings might have selected characteristics that provide additional information about the documents they represent. They might, for example, be color-coded. But the natural coupling between the appearance of buildings and their functions as potential human environments does not carry forward into the virtual world of datascapes.

Topic Maps

An alternative prospect-based artifact is the topic map, which is a form of entity-relationship (E-R) diagram. A topic map displays in a visual form the topics a document or collection contains and the relationships among the topics (Figure 2.02). The topic map has an advantage over the datascape in that the information it can represent is more significant than document size or document type – a topic map is an index to the significant content. In cases where the topics are linked to the documents they represent, a topic map also has the capacity of taking the user instantly to a desired destination.

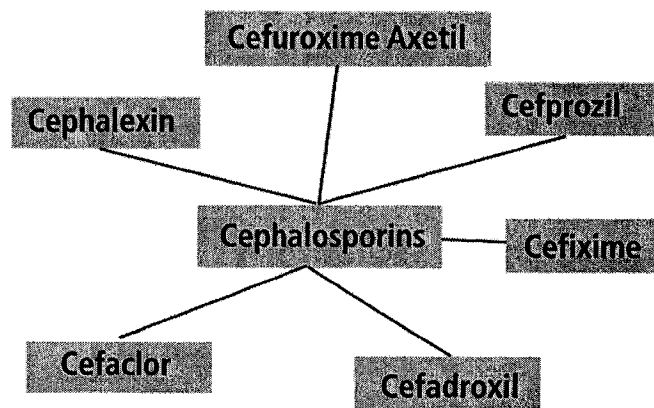


Figure 2.02 A topic map shows the available topics in a collection and groups the information according to the relationships among ideas. This map shows the six categories of anti-infective agents known as cephalosporins.

In spite of their many admirable qualities, however, topic maps also have several limitations. First of all, from a visual perspective the form contains features that are often unaesthetic and sometimes redundant. The information is usually composed of lines, boxes, and text. The boxes indicate entities and the lines indicate relationships among the entities. However, this system is visually redundant along both these dimensions. The text items already indicate entities, and the Gestalt tendency to associate items in physical proximity with each other means that if related text items are juxtaposed, people will understand that they are related. As in the case of the datascape, the amount of visual emphasis given to relatively trivial or unnecessary information renders the standard ER diagram less useful than it otherwise might be. If the redundant visual clutter is removed, then additional meaningful information can be added in the form of numbers that indicate quantities of

references or numbers of documents in the collection, and the text can be resized to cue the perceiver to relative numbers at a glance (Figure 2.03).

Cephalosporins – total 60
Cefaclor (23) Cefuroxime Axetil (8)
Cefadroxil (3) Cefprozil (4)
Cephalexin (20) Cefixime (2)

Figure 2.03 If the topic map is reconfigured to rely on the Gestalt principle of proximity, the result can be a more compact display that leaves room for additional information. In this case the number of manufacturers of each anti-infective has been added, and the total is shown for the entire class of drug, which is indicated by a larger font in a different shade of gray, rather than by central position as in the topic map.

Leaving aside the question of its visual form, the topic map is also limited in terms of the kind of information it displays. A zoomable digital map intended for wayfinding in an analog city contains a variety of different kinds of relevant information. It has some lines that represent roads, others that represent regions, and still others for natural features, such as rivers or hills. It has a grid system of its own, usually overlaid on the grid system used by the city. It has an alphabetized index of street names. It has indicators of direction and scale. In comparison, the topic map onto a data collection shows only topics and their relationships. Forms augmented with numbers may also indicate how many documents of each kind are available in the collection. Versions with clickable entries may link to the actual documents, which is an excellent feature, but only if the indexed topics are meaningful representations of the documents for a given user, and only if the numbers of documents are not overwhelming. If, for example, a single topic can be found in a thousand documents, a clickable topic map entry is not going to provide a very strong affordance.

In its strongest implementation, the topic map combines the advantages of a good index and keyword catalog, coupled with the dynamic opportunities for linking directly to the specified material that are made possible because the map leads to a digital collection. Topic maps are sufficiently useful that they have been adopted as an ISO standard, with an associated SGML document type definition (DTD) (Ontopia 2002). For the purposes of wayfinding within a document collection, however, there are more opportunities available.

Wayfinding

Vinson (1999) identifies five features that are often used by people for wayfinding in an analog environment, and suggests that analogous features might be useful in designing navigation strategies for the digital world. These features are paths, edges, districts, nodes, and landmarks. Most of these items are typically represented on maps. Paths are often shown as roads or public transit lines, and districts are often marked with outlines and text labels. Nodes are places where paths converge, which are visible on most maps. Some landmarks may be given, in the form of outlines and labels indicating prominent buildings or statues. Like nodes, edges – the boundary conditions between different kinds of landscape features – are often indicated but are not given particular emphasis.

In the world of digital collections, the concept of paths was introduced roughly fifty years before the invention of HTML. Bush (1945) suggested that a method might be developed to help over-taxed post-war scientists stay current in their literature, whereby specialists in what he called “trail blazing” would be able to create and store associated materials, in much the way an anthologizer compiles physical items:

When the user is building a trail, he names it, inserts the name in his code book, and taps it out on his keyboard. Before him are the two items to be joined, projected onto adjacent viewing positions.... Thereafter, at any time, when one of these items is in view, the other can be instantly recalled merely by tapping a button below the corresponding code space. Moreover, when numerous items have been thus joined together to form a trail, they can be reviewed in turn, rapidly or slowly, by deflecting a lever like that used for turning the pages of a book. It is exactly as though the physical items had been gathered together from widely separated sources and bound together to form a new book. It is more than this, for any item can be joined into numerous trails (Bush 1945, p. 107).

There are any number of hypertext or hypermedia authoring systems that resemble Bush’s memex. Researchers at Brown University have developed a series of such tools, including:

- the Hypertext Editing System (HES) (1968)
- the File Retrieval and Editing System (FRESS) (1969)
- Intermedia (1985)
- Storyspace (1992) (Keep et al. 2000)

These tools varied from one another in terms of technical implementation and range of features, but shared the capability of allowing the user to construct hyperlinked sets of information. Storyspace, for example, contained a menu item called Roadmap, which displayed a local map of paths.

Bush's ideas have also been implemented by Shipman et al. (2000) in a system designed for use in high school classrooms. Users are able to create an organizing metastructure that combines existing web pages and annotations. This metastructure can be stored for subsequent use by other people, although there are issues related to copyright and the volatility of web materials which still remain to be addressed. If the system stores the pages, it not only infringes on copyright, but the pages may also become outdated. If the system does not store the pages, they may change, disappear, or move to a new server or URL, rendering the pathways that include them obsolete.

Bush's hypothetical memex and the Walden's Paths system of Shipman et al. allow users to create conceptual pathways through an electronic collection. They are not, however, particularly visual implementations of the idea of paths. Bush did not elaborate on interface ideas for the memex, and the published screen shots of Walden's Paths suggest that the web browser is the visual model adopted by the designers, with many boxes of text overlaid on each other.

An alternative display that does attempt to provide a visual implementation is the one described by Roussinov et al. (1999). Their map is intended to visually represent documents as clusters of colour-coded and labelled icons on a grid. The user can open any of the items shown, and also has the ability to modify the map in various ways (for example, by rating the items shown for relevance, or removing irrelevant items altogether).

Digital versions of Vinson's other landscape features (edges, districts, nodes, and landmarks) are somewhat more difficult to identify within the desktop model, although they are fairly common in computer games and datascares.

Panoramas

Interface panoramas are horizontally scrolling display fields that have typically been implemented based on 360 degrees, or an entire circle, of view. Many early implementations were designed for the purpose of displaying interior or exterior spaces by stitching together photographs; the technology has subsequently been extended to display other kinds of data as well. Panoramas are now used as interface tools showing a wide range of data, from network traffic to galleries of student art. By making individual items within the panorama clickable, the designer has the opportunity to make the view

into an access tool, whether to more panoramas, larger format images, or any other kind of data files associated with the links.

Panoramas can be any height, and the contents typically scroll left or right at a speed determined by the current offset of the cursor off the centre line. In some versions the panorama may also scroll vertically based on cursor offset up or down from the top or bottom of the pane containing the panorama; other versions have the panorama expand from a thin to thick display pane based on the vertical offset; still others have no vertical effects whatsoever. Earlier versions involving photos that had been amalgamated, as for example with QuickTime Virtual Reality (QTVR) often had an inadvertent fisheye lens effect which distorted the view and emphasized its artificiality. Panoramas created with alternative technologies, such as Director or Flash, seldom show this kind of distortion.

In any case, as a strategy for providing prospect, the panorama has several advantages over the standard vertically-scrolling window. First, it is controlled with cursor position rather than with a specialized device such as a scrollbar with arrows and a thumb. The effect of cursor position is arguably easier to identify and learn to use than a scrollbar, although it may be frustrating in cases where the response is sufficiently slow to create a time lag for the user, which might tend to perceptually decouple the stimulus from its feedback. The second advantage of the panorama is that it forms a complete circle, which means the user does not have to reverse the scrolling effect in order to arrive back at the beginning of the display. This continuous visual loop allows the user to pan around the entire view in either direction, at a speed that is under direct control, which quickly provides an overview of the entire display. Finally, panoramas are horizontal, which means looking at a panorama can be understood as analogous to obtaining prospect on a horizon. Although the interface is limited by the interactions of the mouse, the metaphor at work is that the user is standing at the center of a prominence and can look around at everything.

In combination, these three factors – ease of use; continuous looping under user control; and resemblance to an analog panorama – make interface panoramas a strong candidate technology for providing prospect. On the negative side of the scale are the features that sometimes make the panorama difficult to use, because it has several functions active at the same time. For instance, since cursor displacement off centre determines both direction and speed, if someone wants to click an object that is visible on either of the sides of the panorama, the object will appear to run away from the cursor. People can learn to get around this problem by always moving items to be selected to the centre

of the panorama, where positioning the cursor will simultaneously stop the motion and allow for clicking. However, the more static quality of normal window pane movements, where the user is used to seeing motionless contents unless one of the dedicated sliding tools is being employed, have created an enculturated expectation that objects will not skitter away when approached. Another design solution is therefore to add horizontal scroll bars or some other dedicated tool to control the scrolling of the panorama, which essentially converts it into a very wide scrolling window with the additional feature of wrapping back to the beginning rather than stopping at each end.

Depth of Field

In the analog world, prospect necessarily involves some depth of field. To have a view from a window onto a grove of deciduous trees is an experience involving prospect. To have a view from a window onto the trunk of a tree growing directly beside the building is to have an experience of thwarted prospect. One of the differences is depth of field. Landscape painting similarly attempts to suggest depth of field through a variety of techniques involving factors such as focal length, scale, perspective, foreground and background cues, occlusion of distant objects by closer ones, atmospheric effects such as colour change or blurring, and so on.

In the digital environment, some depth cues are common while others are seldom seen. Occlusion, for instance, is a default behaviour of windows, where the currently active window is intended to sit in front of any others that are open. Icons, however, are a bit more complex. Under the protocols for “drag and drop,” the user is often able to trigger an application icon by placing a data icon on top of it. Folder icons, on the other hand, have the default behavior of ingesting other icons of any kind that are placed on them. The ingested icons disappear from view until the folder icon is opened. Under some conditions, it is also possible for one icon to simply occlude another, as can sometimes happen when folder contents that have been displayed as a list are subsequently displayed as a set of icons. From the user’s perspective, these icon behaviors are sufficiently unrelated to depth cues that it is more realistic to discuss them as characteristics of icons rather than as metaphoric treatments of a virtual third dimension.

Some experimental designs, however, have attempted to make use of three dimensionality as a means of providing prospect without sacrificing too much screen space. An example is the kind of interface where items appear to advance and recede, either as individual objects or as parts of a larger rotating whole. In general, however, the desktop and the application window are usually treated as flat

surfaces. By implementing interfaces in three apparent dimensions, the developer has the opportunity to take advantage of a much larger display environment. However, there is the risk of creating too literal an interpretation of the landscape metaphor. A limited solution might therefore rely primarily on size changes that do not include perspectival narrowing, coupled with occlusion – provided that the occlusion is not implemented to such a degree as to eliminate prospect.

RICH-PROSPECT INTERFACES

A rich-prospect interface is one in which a meaningful representation of every item in the collection is an intrinsic part of the visual display that allows the user to access the collection. Ideally, this form of display serves as the basis for a set of tools that can be used for sorting, subsetting, grouping, and otherwise manipulating the information shown, in ways that are useful for a particular user of the collection. A variety of interface technologies have been developed to provide various forms of prospect, but the value of rich prospect has yet to be widely recognized.

In general, any prospect-based interface should address three fundamental questions for the user. These questions relate to the affordances of the interface and the tools that are provided with it.

They are:

- what am I looking at?
- why would I want to look at it?
- what can I do with it?

The answers to these questions are related first of all to the provision of meaningful representations of the items in the collection.

Meaningful Representations of Items

In order for an item to be represented in a meaningful way, as opposed to simply being represented, it is necessary that the designer be familiar with the people who will be using the system, and understand both how they will immediately perceive what they see, and how in the process of working with the interface they will construct an understanding from the materials they have available. It is equally necessary for the designer to understand the nature of the material itself, since the construction of meaningful representations must occur with respect to the contents of the collection.

Users

It is now widely recognized that to design anything is to be involved in an act of communication, and that to communicate effectively requires some common terrain that is recognized by both interlocutors. Language itself is such a terrain, but is only part of the larger environment that also includes the presuppositions of the various parties, their personal experience, the public history of which they are a part, and so on. In the field of industrial design, the Environmental Design Research Association (EDRA) was founded in 1969 to promote better understanding of product users to help inform the design process. In visual communication design, recognition of the central role of the user has been slowly growing, and various methods for involving the user have either been developed from first principles or imported from the social sciences (Frascara 1997, pp. 33-59). However, in spite of this affirmative stance, in practice the actual interactions between the designer and the end user are often limited for very good reasons involving the needs of both groups.

Designers may need to know about the intended users of a system, but there are often no such people readily available. Designers may want to know about the intended users, but the brief often assumes that someone else will be responsible for letting them know what they need to know. It may even be the case that management of the project requires limiting the contact of the designers from the end users in order to prevent one of the most serious problems a project can face – namely, scope creep, wherein the bounds of the design are modified or expanded as the project proceeds, resulting in a project that can never be completed, or at least never completed within the constraints of available time and budget. Finally, designers may have users available for study, but may simply not have the time or the expertise to find out what is necessary.

From the perspective of the users, if the system being designed is a new system in any substantial way, there may not be an existing body of users to draw upon. If there are users, they may feel that they do not have the expertise necessary to contribute to the design of an interface – that the work is in the domain of the expertise of others, namely the interface designers. Finally, many people who may in fact have the expertise to help also have other commitments of their time and resources that preclude them from serving as guinea pigs for interface designers.

The result is what Mitchell (1993, p. 36) and others have referred to as “the applicability gap,” where the information that is available is either not appropriate or not used by the designer when the work of creating the design actually begins.

Given the sometimes overwhelming problems of finding and understanding actual users, many studies make use of study participants who happen to be available, such as students or administrative staff. This approach has the value of at least involving actual people interacting with the designer's ideas. Another strategy, even less connected to actual user-centred design, but very useful as a way of managing client expectations, is the creation of user profiles, where fictitious people are substituted for actual users (Fleming 1998, pp. 8-9). Discussions of user needs can then be held in the context of the characteristics and needs of the invented person, which serves to reduce the chances for deadlock which sometimes arise between the designer and the client, because there is a third party (albeit a fictitious one) to be referenced in any decision. Since this third party is an invention of the designer, it can be given whatever characteristics seem appropriate to the task at hand.

Some studies, however, are based on projects where the actual users have been involved in an iterative design that responds to their feedback with revisions to the system. An example of such a project is the Alexandria Digital Library (ADL), which consists of a geographic database containing a variety of information about various points on the surface of the earth. Researchers with the ADL worked extensively with three target user groups: earth scientists, information specialists, and educators (Hill et al. 2000, p. 250). The partial list of requirements that derived from these users has eight categories, which are extensive enough that they might be used as a general summary of system features:

- search functions
- session management
- result display
- user workspace
- holdings visualization
- user help functions
- usability features
- data distribution.

The ADL researchers emphasized that the design of a system for use by a particular community is essentially different from the design of a system that will showcase its own capabilities. These differences include both content and interface (Hill et al. 2000, p. 257). It is interesting to note, however, that in spite of the nature of the content and the extensive user

participation, the ADL is not an example of project that provides the user with prospect on the contents of the collection.

Form

There are several different strategies available to use in providing prospect. One method is to use the form of the visual material as an indication, not of the content, but of the digital nature of the material that is being offered to the user. For example, in the desktop metaphor, there are standard icons that represent documents and folders, and variations of the icons are used to indicate whether a given document is an application or a data file that belongs to a particular application. Some research projects have attempted to leverage this existing visual vocabulary for use in browsing interfaces.

There are two problems with this approach. First, the icons tend to be quite large, since they were originally intended to draw the user's attention to files on the desktop, rather than having been designed to work together as a complex display. Second, and more importantly, the icons do not provide a significant level of return on investment in visual terms. To see a thousand icons, each representing a data document, is to perceive a complex pattern composed of identical elements. This kind of display can provide, in fact, an instance of the sublime of repetition, where sheer numbers of identical or near-identical items can trigger an emotional response in the viewer. Unfortunately, however, all it conveys in information terms is that there are many identical data files available. In order for the files to be differentiated from each other, it is necessary to add textual labels. The icons become redundant once one is given per section, and the purpose of the display is therefore largely fulfilled by the content rather than by the form, even though the form is allocated a large portion of screen space (Figure 2.04).




Figure 2.04 The interface to the Nemo project, which accesses documents related to Electricité de France. The repetition of icons gives a low information return on investment (Hascoët and Soinard 1998).

Even in systems where the icons differ from each other in a significant degree, there are still details of display that can render the interface more or less useful. In the Data Mountain visual interface, for example, web sites are shown as thumbnail images using a snapshot of the actual home page of the URL. Each image is therefore unique, but the interface allows a nontrivial amount of visual occlusion between images, rendering all but the front image difficult to interpret (Figure 2.05). The Data Mountain visualization is arguably based on content rather form. However, since the thumbnail often reduces the content to the extent that little or no text is actually legible, the point is open to debate.




Figure 2.05 Unlike the interface to the Nemo project, the data mountain uses a unique graphic for each document, although the high degree of overlap renders many of the images inamenable to viewing (Czerwinski et al. 1999).

Another form-based display is the one used by the Alexandria Digital Library, where geocentric information is made available to the users through the use of a visual footprint, which consists of an outline superimposed on the surface of a map. The contours of the outline in this case are significant, since the superimposition indicates the region of interest to the user. These visual footprints allow the user to query any region of the globe, independent of the name of the region, which simplifies the query in cases where the user may be uncertain of the spelling or where the designers of the system have not included all of the valid spellings (which often vary for geographic regions by dialect, language, and source). For example, the city known in the English-speaking world as Copenhagen is referred to by the people who live there as København. The difference in spelling – particularly in the initial consonant – means that an alphabetical listing of place names using the word “Copenhagen” might represent a barrier to København residents interested in finding out about their area.

Visual footprints also allow queries on regions which do not have names. For example, if a user of the system were to draw an irregular polygon around several cities, it is possible that the collective area indicated would not be known by a unique name. Similarly, a section chosen from the middle of a lake, ocean, or desert is unlikely to have its own name, yet the system of geo-information may contain relevant data concerning its climate, wildlife, topography or other features.

Based on the principle implemented in the ADL, good candidate collections for queries and displays based on visual forms are those which have some pre-existing visual vocabulary that can be used as the basis for the system. Other systems may also be able to effectively adapt visual forms

through the use of metaphor, symbolism, and other sign systems which associate meaning with visual material (as opposed, for example, to text), but in cases where the collection itself is naturally associated with visual materials, the use of visual access methods is in alignment with the underlying content domain.

In the case of ADL, for example, the map of the globe or the region under study is the pre-existing visual element that serves as the basis for the footprints, which are themselves a type of visual query on the system. The user of the ADL collection does not retrieve information by typing text: instead, the material is accessed by placing a bounded region onto a globe. A hypothetical example of a related kind of pre-existing visual information might be in an interface to a collection of retail goods for the home or office, where the visual basis for the query could be a floorplan that allowed shoppers to quickly narrow their interest from a position of prospect on the entire building to the details of a particular room or area.

Content

Although form-based systems do exist, the most common method of providing meaning to the user is through displays based on content. The primary kind of content-based display uses text, which has the advantage of potentially conveying a maximum amount of meaning, but the disadvantage that it is only accessible to people who are literate in the language and share a common orthography. It is also not uncommon for text itself to be waste space in terms of the information it conveys, either because it is repeated unnecessarily or because it does not sufficiently differentiate the items it represents. Unnecessarily repeated text often appears as labels which are intended to structure the display and make the user aware of the kinds of data available. In cases where elements of the same kind are repeated, these labels quickly become redundant. Insufficiently differentiated text occurs in cases where a representation which is supposed to distinguish one item from another instead serves to indicate similarity. An example might be in a keyword listing where the same keyword has been applied to every item shown, and is used as a part of the display of each item, rather than appearing as a key to the whole page. European archival records of the soldiers killed in the First World War, for instance, will sometimes list the names next to a conventional military designation such as “killed in action,” which can continue for page after page of entries until the reader is numbed by the sheer repetition.

For displays where only a few items are shown, inefficient text is not necessarily a serious problem, although it can become a source of irritation over time rather than through repetition at one time. For displays intended to provide prospect, however, it seems clear that redundancies should be avoided wherever possible in order to maximize the effective use of the limited screen real estate.

It is also important to note that meaningful representation by content does not necessarily imply a textual representation. Collections of images, for instance, or video clips, might be better represented by thumbnail images than by textual labels (Figure 2.06). The problem with visual representations is that they may need to be comparatively large in order to be distinguishable, which has implications for the design in terms of screen real estate. There are also limits to what a perceiver can tolerate in terms of visual complexity or simple overload, although the details of these limits and how they can be addressed through various strategies (such as selection, grouping, subsetting, and so on) require further research.




Figure 2.06 The Photomesa interface provides the user with a wall of thumbnail versions of photographs which are perhaps surprisingly accessible to browsing, given the complexity of their initial visual impact (Bederson 2001).

Relationship

For a rich-prospect display to convey a maximum amount of meaning in a form that is readily understandable, one strategy is to emphasize relationships among the collection items rather than either the content or the form. Starfield displays are one form of relationship display where form and content have been reduced, often to a single pixel, in order to provide as simple a visual presentation as possible. The information conveyed in a starfield interface is therefore primarily in the form of relational positioning, with some selected document characteristics used to group individual collection items into larger aggregates (Figure 2.07).

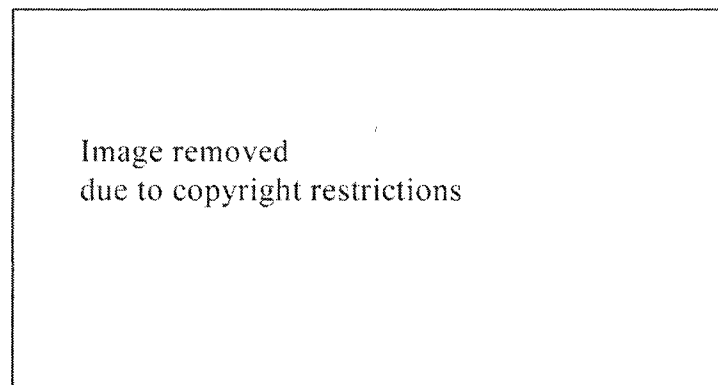


Figure 2.07 Starfield displays are like entity-relationship (E-R) diagrams in that the relationship between individual items is considered primary. This starfield display shows a building's cooling system as a central point representing each fan, with surrounding points indicating the fan temperature as either too hot, too cool, or just right (the original is colour-coded red, blue or green) (Johnson Controls 2002).

The primary disadvantage of starfield interfaces is that in order to keep the size within reasonable bounds, the individual items are not meaningful in themselves. There are essentially two solutions to this problem – make the individual items meaningful, or else associate them in some accessible way with other kinds of representation that are meaningful, whether in the current window or in an associated one. The second solution is the one that has been pursued most extensively by the commercial manufacturers of starfield software, where various kinds of display are simultaneously presented to the user (Figure 2.08).



Figure 2.08 The Spotfire Decisionsite for Functional Genomics contains a variety of tools, types of displays, and simultaneous multiple views to allow geneticists to work with complex genetic information (Spotfire 2002).

In a somewhat different context, the synthesis of various kinds of information is also at the heart of a form of classification system known as facet analytical theory. Originally formulated by Ranganathan in the 1930s, facet analysis is a relatively complex approach to knowledge representation governed by three planes, 46 canons, 13 postulates, and 22 principles, which have been subsequently modified and adapted by other researchers in the library sciences (Spiteri 1998). The three planes of facet analysis represent, respectively, the need to divide a subject area into its component parts, choose appropriate terminology, and create a notation that preserves the notion of the components. The components must be mutually exclusive, so that individual items in the collection can be uniquely represented by combining the terms (Broughton 1998).

Hybrids

People are capable of perceiving and using a wide variety of complex nested and sequential analog affordances, so there is no *a priori* reason for rejecting the possibility of creating complex digital affordances. The most powerful tools are also often the tools that are most flexible – that is, they

provide the user with the greatest number of affordances, including the possibility of using the tool in ways that the designer did not anticipate.

Hybrid forms of the meaningful display of items have the potential to open up additional affordances by combining the relevant features of each of the specialized forms. If content, form, and relationship can all be deployed strategically together in order to convey meaning, along with a range of tools to allow various opportunities to act on that meaning, it may be possible to expand the benefits to the user of the interface beyond the affordances available through a representation that relies on only one of these methods.

Amount of Information on Display

No matter what kind of display is appropriate for a given user in a particular context, the question remains as to how much information is necessary or potentially useful. There is some need to manage the limited amount of available screen space, and independent of the screen space available, the cognitive demands of a rich-prospect form of display on the user are potentially quite high. A tradeoff therefore exists between the choice to display as much information as possible, in the hopes that it will prove useful to someone, and the structuring of the information in such a way as to increase prospect.

Many of the current web search engines and document retrieval systems provide the user with a list of search results that scrolls vertically and consists of individual items that are approximately three lines each in length. An arbitrary display limit is usually set, with items over the limit either not available at all (e.g. ACM) or else available in subsequent screens (e.g. Google). These three choices – to display the results in (1) some detail using (2) a vertical list of (3) limited length – all serve to reduce the amount of prospect available to the user. An alternative strategy with more prospect might create a structured display of very short representations, perhaps clustered around relevance rating, and numbering at least in the thousands, which would give the user as much as two orders of magnitude more information to work with. The presentation of so much information of course requires a variety of strategies to make it manageable rather than overwhelming, but visually structuring large amounts of information in useful ways is a challenge the visual communication design community has taken on for several generations, and many strategies are available.

Insights About the Collection

“The most successful designs are not those that try to fully model the domain in which they operate, but those that are ‘in alignment’ with the fundamental structure of that domain, and that allow for modification and evolution to generate new structural coupling.” (Winograd and Flores 1986, p. 53)

Rich-prospect browsing interfaces are potentially important because they may allow designers to create the basis for a kind of information access that is congenial to many people. They may also provide opportunities for actions that are not possible using interfaces that do not provide some form of prospect. Some of these new affordances may only emerge during the course of research; it is, however, possible to postulate what some of them might be through considering the kinds of actions that could be made available in conjunction with a rich-prospect interface.

The possible actions are in turn related to two factors: user insight into the collection; and the interface tools that have been provided to allow the user to do something with that insight. The insights available to the user are primarily related to indicating the bounds of discourse that have inevitably been established by the collection – that is, the terms under which the items have been collected, labelled, categorized, and otherwise organized. These areas of direct insight can be grouped into the following categories:

- contents
- structure
- context
- features
- limitations
- connections
- trends
- anomalies
- navigation
- reminders
- reassurance
- reduced helplessness

The ways in which these factors relate to the opportunities for action provided by the system will differ significantly based on the factor involved, but one of the primary felicity conditions for each of the affordances is the availability of prospect on the collection.

Insights about Content

By providing a meaningful representation of every item in the collection, a rich-prospect interface allows the user to directly perceive what is available. The user is not dependent on previous experience with the collection, or on having read explanatory material about it, although both of these might of course be useful. Simply by glancing at the items, the user is able to ascertain with some degree of certainty what the collection is about, how large it is, and whether or not it can contribute to the purpose at hand. In cases where the rich-prospect is provided in combination with a search function, direct insight into the contents may also help to establish an appropriate search vocabulary (Pirolli et al. 1996).

It should be emphasized, however, that direct insight is possible only in cases where the system uses terms that the user would consider relevant. As a hypothetical example, in a collection of prescription drugs, it might be useful to organize the display according to the type of drug if the user is a doctor or other medical professional. An organizing scheme in this case might use categories such as “cephalosporins,” or “aromatic glycerol ethers.” If the same collection were being designed for access by people suffering from some medical condition, however, it might be useful to provide an organizing scheme that used categories based on the disease or other medical problem. In this case, the collection might have categories such as “sinus infection” or “back pain.” If the patient attempted to make use of a rich-prospect interface designed for the doctor, it would not necessarily be possible to distinguish which drugs might be suitable for which kinds of medical conditions.

The development of meaningful representations of content items might also draw on facet analytical theory, in the sense that the system might construct representations by combining multiple organizing principles into a single composite term. Ranganathan’s Idea Plane Canon of Relevance emphasizes that the facets used as components for such a term should align with the intention of the collection: its purpose, subject and scope (Spiteri 1998). In the hypothetical case of a collection designed for patients looking for information about medical conditions, for example, a faceted description might include a composite representation that included the part of the body afflicted, the

medical condition, and the cost. These three terms represent three facets that are mutually exclusive, may be of interest to the user of such a collection, and uniquely identify the collection items.

It is also possible for a rich-prospect interface to actually misrepresent the collection. For instance, if a collection of commentaries on philosophers were organized by the names of the primary authors and their works, some users might interpret the interface as implying that the collection contains the primary materials, when in fact it consists entirely of secondary critical material. Guarding against potential misinterpretations based on alternative presuppositions is arguably one of the most difficult tasks of the designer, because the nature of the problem stems from disparities that are not necessarily explicit either for the designers or the users. Extensive user involvement in the design process, and in testing, can help to forestall these kinds of situations, and iterative approaches to development can help reduce the impact of any which do occur.

Insights about Structure

For interfaces to digital collections, there are two distinct structures involved. First is the structure of the collection itself, in terms of the kinds of documents that it contains and the ways in which they are conceptually organized by the collectors or designers of the system. A group of people responsible for designing a digital conference proceedings, for instance, might decide that the digital papers should be collected in groups according to the session of the conference in which they were originally presented. Alternatively, if the papers related to topics that were of national interest, they might decide to organize the documents according to the national affiliations of the authors, with papers from the U.K. in one section and papers from Malaysia or the U.S. in another. Or the papers might be organized by length, with the full papers in one section and the poster sessions in another. Any number of different organizational schemes are possible.

The second structure relates to the interface. Independent of how the underlying documents have been organized, the interface designer has another opportunity to provide an organizing principle, which might reflect the understanding of the people who created the collection, but which might also reflect alternative understandings, such as those of the users.

If the interface provides an appropriate structure of the first kind, the user can be provided with potential insights into the nature of the collection. For example, a collection of text documents might consist of items that deal with the same subject matter, but that have been written with different audiences in mind. If the interface is designed with the subject matter rather than the audience

as a central organizing principle, then the user would have an immediate cue to the fact that certain documents that might otherwise appear to be unrelated in fact have something very central in common with each other – namely, their reference to a common subject. If the interface were to be designed the other way – that is, organized by audience – then the common reference to subject would be occluded, but the different audiences could be made immediately evident as the central theme of each document cluster.

Whether or not the interface reflects the underlying organization of the material, depending on the user, some organizing principles are going to be more useful than others. The subdivision of the Amazon.com site into different kinds of products is an example of an interface that provides structural information about the elements in the collection. The user can choose to search the entire product line, or can limit the search to books, videos, CDs, and so on. In a rich-prospect form of interface, these categories could serve as an organizing principle for the display.

Product format, however, is not necessarily an organizing principle that provides the most useful kind of information for the user. An alternative strategy might involve clustering the information by topic areas, with visual cues within the topic cluster used to specify format. A topic cluster for use in Amazon.com might be a subject area such as gardening, with all the available materials, whether books, videos, garden tools, or seeds, shown in proximity to each other. A user looking for a particular plant might therefore find the seeds for it shown in relation to books about how to grow the plant, tools to use in working with the plant, and paintings that feature it as a subject matter. If users are given the facility to create and store structural groupings that can serve as interaction histories for other users, it may be the case that the collection of materials on that particular plant were not created as part of the original system design, but are part of a legacy of structural suggestions made by previous users of the system.

Insights about Context

There are many collections in the world, whether analog or digital, and some are more clearly defined than others. Depending on the nature of the collection and its status in the culture, it may not always be straightforward for users to determine what kind of collection they are currently investigating. Some collections are immediately recognizable for what they are, because they have become enculturated as collections. A phone book, for example, is an artifact that has a strongly enculturated identity. People familiar with phone systems immediately recognize a phone book because of its size

and poor paper quality, which are consequences of having to annually replace high print runs of large quantities of data for mass distribution.

On the web there are some cues to the likely reliability of information, such as the taxonomy of URLs. In general, U.S. sites that contain the designation .edu and sites elsewhere marked .ac are associated with academic institutions, while .com sites are commercial or personal enterprises. Sites related to public bodies sometimes, but not always, use a country abbreviation such as .ca, which is useful in any case for placing the site geographically.

URLs are also not the only source of insight into how reliable a source may be. There are various branding strategies, such as institutional identities, and in some cases there may also be explanatory text that provides the potential user with some idea of the scope and coverage of the site.

However, if a rich-prospect browsing interface is used, the individual items in the collection are immediately present to the user within the context of the larger collection, which provides the users with clues as to what kind of collection they have found.

Insights about Features

A collection can have any number of attributes in addition to the content items and the visual language. These attributes can in turn form part of the rich-prospect interface, allowing for their direct perception by the user. An example is the presence or absence of an interpretive tagging system, along with its potential complexities in terms of the definitions of the tags, and also in the use of attributes on the tags and the values of those attributes, all of which are features that can be used for retrieval purposes by the computer, but can also serve as components of a rich-prospect interface. Rich-prospect interfaces for tagged collections will be discussed at greater length in the next chapter.

Insights about Limitations

Just as prospect can allow the user to identify the strengths of a collection, either in terms of the significant clusters of documents contained or the individual items being sought, so can a prospect-based interface allow the user to identify areas where the collection is not going to be useful or may be useful only with extra effort. For example, if a collection contains a number of documents intended to market electronic products, there may be a combination of promotional items and technical specifications. If a user looking to troubleshoot the product finds a prospect display of the marketing materials, it will be immediately apparent that troubleshooting advice is not part of the collection.

Similarly, a rich-prospect interface can indicate not just that certain kinds of documents are missing, but also that the intentions of the designers of the collection are either going to make the current browsing task easier or more difficult to perform, depending on how the presuppositions on which the design was based either correspond or fail to correspond to the presuppositions of the user. A typical example might be in the use of keywords as part of the meaningful representation of the items in the collection, where the keywords chosen by the people responsible for creating the index will not necessarily correspond to the definitions used for the same concepts by the people seeking to access the collection. Structuring the display as clusters of document titles around each keyword may be one way to suggest to the user the way in which particular keywords have been applied in that system.

Insights about Connections

If an interface places different meaningful representations together in the display, the Gestalt tendency of proximity will encourage users to consider potential connections among the items. The organization of the items will naturally tend to strengthen or weaken this tendency. If the display is arranged chronologically, for example, the user may be able to identify items that are part of a thematic interest of a particular era, or people who were contemporaries, or form some sense of historical narrative such as can be achieved through examining a visual timeline or other sequence.

Insights about Trends

Independent of the structure of the interface or the structure of the underlying collection, there may also be trends in the collection that are potentially significant to the potential user but would not be obvious to someone just looking at the individual documents. For example, a collection arranged chronologically may prove to have strong holdings in one particular period but very few holdings in another. A chronological rich-prospect display would make that difference immediately apparent to the user, since the number of items showing in the historical period with a lot of holdings would form a comparatively larger group on the screen.

Insights about Anomalies

With a rich-prospect interface, the user may have the opportunity to identify individual items or groups of items which seem to be out of place in the collection or are in some other way anomalous.

As Shneiderman et al. (1992) point out, one of the common activities of information foraging in western culture involves looking for bargains. For a user interested in identifying an item that can be

purchased at a discount, a collection interface that uses price as a structuring principle would quickly allow identification of possible bargains. An even more useful organization of the interface, however, would be one that emphasized the comparison of items that were similar across every dimension except price. For example, a list of houses organized by street or neighborhood could potentially show anomalous pricing more clearly than it could be shown by an interface organized by price range.

Insights about Navigation

If the design of the rich-prospect interface is such that it contains information about the structure of the underlying collection, then the interface also has the potential to serve as a navigational aid. As Winograd and Flores (1986) point out, an even more optimal situation is one in which the user has the opportunity to either modify existing strategies for communicating with the collection, or else has some means of establishing new ones.

Some interesting possibilities have been investigated by previous researchers. Wexelblat and Maes (1999), for instance, developed a suite of “footprint” tools to provide interaction histories, both for the current user and for subsequent users who might want to take advantage of previous work: “One of the primary benefits of interaction history is to give newcomers the benefits of work done in the past.” (Wexelblat and Maes 1999, p. 217). The record of past work in a footprint can include paths through the collection, although in order to increase their usefulness for others, it is helpful to find ways of conveying not just where they went, but also who did it, why they did what they did, and how the history was created (that is, automatically by the system, or subject to selection or editing by the user).

Reminders

If some meaningful representation of items is available to the user, there is the possibility that the person will look at the representation and be reminded of collection items that are of potential interest, either because the user knew about them at some point and has forgotten, or else because they are something new that would not have occurred to the user if the system had not offered them up for observation. This affordance is a digital analog to the opportunity available to the library patron who scans the stacks looking for items that might be related to the title already found, or are otherwise of interest. In the case of the library, the affordance is made available through the organization of the shelves by subject.

Reassurance

If the prospect display is attached to the results of a search process, it has the possibility of providing the user with a means of understanding the search results within their context. For instance, a dictionary search using some of the versions of the online Oxford English Dictionary results in the display not only of the word found (or not found), but also of the dozen or more words that occur alphabetically before and after the target word. This strategy can help to reduce the consequences of some minor spelling difficulties by providing the user with a picklist of alternative words that begin with a character string similar to the characters that begin the word being sought. It can also suggest related words that might vary slightly from the target, provided that the spelling begins with a similar string, as it frequently does in English.

Reduced Helplessness

With no prospect on a collection, a user who has no idea where to begin can be left feeling helpless. If the collection has a rich-prospect interface, a user may not be able to figure out any of the tools available, but at the very least there is some meaningful representation of the collection to be examined. There is the cognitive reassurance that there are actually items in the collection, and if the representation is meaningful to the user, there is the additional reassurance that the collection is either a good choice for further investigation or may not contain the kinds of items being sought. Although reduction in helplessness is not an affordance *per se*, it might be understood either as a felicity condition or as a component of other opportunities for action, such as the opportunity to continue using the collection.

Prospect-Related Interface Tools

The second factor related to the affordances of prospect lies in the additional tools provided to deal with the prospect in various ways. Some of these tools are also useful in contexts where the interface is not based on rich prospect, although the functioning from the user's perspective may vary significantly because of the differences between the two kinds of interface. The tools that are potentially useful in a rich-prospect context include:

- zooming
- panning

- sorting
- selecting
- grouping
- subsetting
- renaming
- annotating
- opening
- structuring

Prospect-Related Interface Tools: Zooming

Screen real estate is an issue in the design of a rich-prospect interface, and magnification methods are one clear means of allowing the user to move from an overview to a detailed view at various levels of granularity under user control. Strategies that involve zooming include fully collapsing the view through various stages ending in an icon or other representation (as in the collapsed window bar at the bottom of the screen in Windows environments); selective zooming through fisheyes; and the use of three dimensional representations, where some objects recede in the virtual distance, while others advance.

Zooming through collapsing the view has the advantage that it requires a minimum amount of room on the screen in order to provide the user with a visual cue that something is present. It can be confusing, however, for users who are unfamiliar with the system and do not realize that the visual cues correspond to larger items (Figure 2.09). In its most extreme case, visual collapse of elements can result in them being hidden from the user altogether – typically these methods place the retrieval system under a menu or associate it with a keystroke or key sequence, which can be useful for a sophisticated user but disorienting for a novice. Adobe Photoshop, for example, allows the user to temporarily hide all the tool palettes by pushing a tab. This feature allows a clear view of the image on the current working area, but requires that the user know the key combination that will bring the palettes back.

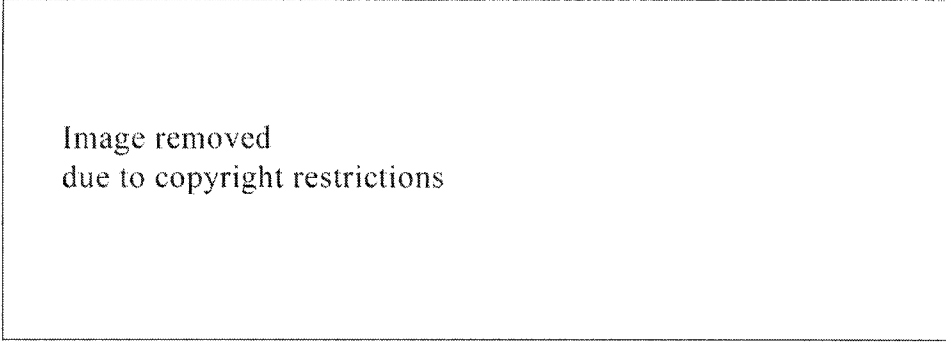


Image removed
due to copyright restrictions

Figure 2.09 The tool palettes in Fractal Design's Poser collapse into tabs at the bottom and right edges of the screen (left). When expanded (right), these tabs dramatically increase the functionality of the software. However, for users who are not familiar with the visual language of the program, these tabs can easily be overlooked. One possible solution is to provide prospect on the tool bars by having them collapse in an animation during the start of the program.

Selective zooming through fisheyes has the advantage that the items on display are constantly present to the user, which helps to prevent disorientation and provides a form of prospect (Figure 2.10). Fisheyes have the disadvantage, however, that they only allow expansion of a part of the display at a given time.




Image removed
due to copyright restrictions

Figure 2.10 A fisheye menu system allows the user to obtain prospect on the entire list of options but selectively magnify them at the point where a choice of items is being made. This screenshot shows the same menu at three different insertion points (Bederson 2000).

One way to avoid the possible disorientation caused by selectively collapsing the view is through a magnification strategy that changes the entire display. An example of this kind of zooming is in Adobe Premiere, where the user can expand the time scale of the movie score by moving a slider that is associated with the larger view (Figure 2.11). This strategy has the advantage of allowing the system to animate the change rather than requiring a dramatic shift from one scale to another. It is still possible for the user to experience disorientation, since the expansion mechanism has to clearly maintain the current insertion point. If the view expands across multiple scales, however, from the largest overview to the closest detail, the insertion point cannot be clearly indicated, since what is a point at the least magnification becomes an area when expanded. One strategy (not used by Premiere) would therefore be to allow the line indicating the insertion point to visually widen as the view expands. The ability to reset the insertion point size would then need to be added, to allow the user to see the minimal insertion line, regardless of the current scale.

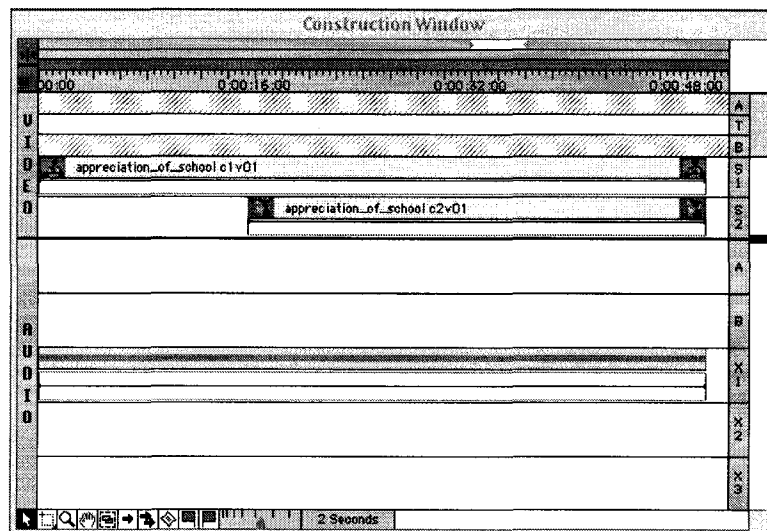


Figure 2.11

The film score in Adobe Premiere can be displayed in various time increments, which allows the user to focus in on parts of the film or see the entire score at once. Similar functions are available in many programs that require a time-related display, such as programs that deal with digital music. Adobe product screen shot reprinted with permission from Adobe Systems Incorporated.

Prospect-Related Interface Tools: Panning

Panning functions are often associated with zooming, since the user often requires some means of moving over or through the display. Panning can take various forms, including:

- implicit panning through positioning of the mouse (as in panoramas)
- specialized tools such as the repositioning hand that allows users to move the larger workspace within the viewing frame, and
- objects like the standard window scrollbars, with their directional arrows and thumb.

All of these solutions are in some respects expressions of the limitations of the keyboard and mouse. There are a wide variety of other options that become available with alternative hardware, from video game controllers that allow complex navigation in three-dimensional environments, to steering wheels, joysticks, digital gloves, and positional trackers or sensors. Like the landscape metaphor itself, the hardware and software devices that simplify interaction in three virtual dimensions have so far found very limited implementation in the office environment, perhaps because they are so strongly associated with the discourse of digital games.

Prospect-Related Interface Tools: Sorting

If it is possible to make the representation of the individual items meaningful to the user, it is also possible to make the arrangement of the items meaningful. A common example might be the case where the display has been sorted so that the items are in alphabetical or chronological order. In those cases, the user is able to directly perceive the organizing scheme, and can therefore use that knowledge to help in locating items where the approximate spelling is known but the exact spelling used by the system is uncertain, or where the spelling might be influenced for retrieval purposes by features which are often considered trivial by human beings, but represent difficulties for search engines – such as capitalization and lemmatization. Lemmas are words that have been inflected in some way to indicate conjugation (for verbs) or declension (for nouns), which results in words that are not of a form that is identical to the word being sought. For example, the user is looking for “chase” and the document has “chasing.”

In addition to sorting the display alphabetically or chronologically, it is also possible to establish other sorting criteria which are potentially useful for particular users of a given collection. For example, in collections which consist of technical papers, it is sometimes useful to be able to see which papers have been most popular with previous readers. Citation-based text archives provide this

function by indicating how often a given article has been cited by other articles in the collection (CiteSeer 2002). The number of citations is understood in these cases as an index to the significance of the article being cited. A rich-prospect interface to a citations-based collection might therefore sort the articles by number of citations.

A more specialized form of sorting might be provided in the case where an interaction history is available. If the system stores information on document access by user, it would be possible to give an individual heavy user of a collection a display sorted by frequency of previous use. In order to provide this kind of information, it is necessary to maintain user profiles by document over time, which has several implications for record-keeping in the system: should the user log in, or is it sufficient to have the system recognize the computer? If the latter, then what about cases of shared or public-access computers? If the former, then the system is introducing an extra step between the user and the information, which may be in some cases a significant deterrent to using the system at all.

An additional layer of complexity is added if the system is going to share an interaction history gained from one user of the system with other users. Amazon.com, for example, provides prospective book purchasers with information on related titles that have been of interest to other people who bought the current book. Since the company is a retailer, there are no serious implications to this sharing of information among users. However, in cases where the collection is an academic archive of primary materials, identifying connections among items is one of the professional activities performed by academics in the course of their research. To offer previous connections made by one academic to other users of the system might therefore introduce issues of academic privacy.

Another complex form of context is provided, not by a chronological sort *per se*, but rather by the choice of representation of the chronological sort. If the items in a chronology are used to create a timeline, it is possible to emphasize or de-emphasize individual items, create visual connections, and even generate what are essentially narratives or themes, through the visual presentation of the items. Each of the visual effects available to the designer therefore needs to be carefully considered in the framework of the agendas to be served by a given collection.

Prospect-Related Interface Tools: Selecting

In order to work with any subset of the items shown in a rich-prospect interface, it is necessary for the user to be able to select individual items or subsets. A wide variety of selection mechanisms are part of the standard desktop environment, including:

- menus: usually visible across the top of the window
- picklists: called up by the menu items or by clicking on a selection box that expands
- check boxes: to allow multiple values in one field
- radio buttons: to provide the choices of mutually-exclusive options in one field
- rollovers: it is possible to have elements react to the presence of the cursor
- clicking: it is normal to have elements respond to clicking with the cursor
- double-clicking: secondary behaviours can sometimes be triggered with two clicks in quick succession – selecting entire words or phrases in MS-Word is an example of this treatment being applied to use in selection
- right and left clicking: these can have different effects if two buttons are available and the software accommodates their use
- shift-clicking: allows the user to select multiple discontinuous items
- dragging a selection area: allows the user to select contiguous groups of items

In addition to these standard options, there have been research efforts to develop other strategies for object selection that work by modifying, extending, or supplementing the existing approaches. Baudisch (nd), for instance, describes the potential application of painting metaphors for item selection in cases where hundreds of items are involved (Figure 2.12), allowing for rapid and discontinuous selection of items. In its original formulation, the painting metaphor was intended for use with toggle maps, which are groups of check boxes, each of which has only one of two possible states – on or off. Baudisch also expands the idea for use in cases where multiple selection states are available to the user, by having the painting tool apply shades of gray that are incrementally darkened as the cursor is passed repeatedly over the area.




Figure 2.12 Baudisch's toggle map for selection of television stations in Germany has many admirable prospect features: the switches are indicated by buttons on the text rather than by additional graphical elements; the items are grouped by geographical location; and the context of placement within Germany is indicated by superimposing the toggle map on a map of the country.

Prospect-Related Interface Tools: Grouping

In a related but distinct area of visual presentation are those tools available to the user for grouping items together. In some cases it may be possible to determine in an *a priori* manner some of the ways in which items might be usefully grouped on behalf of the user. In other cases it may be equally useful to provide the user some means of creating new groups of items. The two situations are also not mutually exclusive.

As with sorting, *a priori* grouping might be performed according to standard schemes such as the alphabetical or chronological. Provided that the items grouped are not collapsed into invisibility, the interface will retain its nature as a rich-prospect interface. Particularly if the user has the ability to re-group what has been previously grouped or to control the collapse and expansion of the groups that have been defined, the affordances of prospect are not necessarily compromised and can in fact be supplemented.

One implementation of this idea is in the Scatter/Gather interface developed by Pirolli et al. (1996), where users are able to manipulate the items in a starfield display in order to create related

groups. In this case, the groups might be organized according to any scheme that seems appropriate. For example, a user may be involved in sorting through web sites, and have one group for sites that have been visited and found interesting, another group for sites that have been visited and proven uninteresting, and a third group for sites yet to be visited. Another user may choose to create groups that represent sources of the sites, with commercial sites in one group, academic sites in another, and personal web pages in a third.

Special attention should be given to the visual format of groups, because it can have consequences both in terms of user perception and also in terms of the allocation of the limited screen real estate. Grouping can be indicated by physical connection (as in lines connecting related items); additional graphical elements (boxes or other shapes underlying related items); proximity; color-coding; and similarity of appearance in terms of form, texture, size, or any other visual attribute. Alignment on a grid can also be used to suggest grouping.

Prospect-Related Interface Tools: Subsetting

Grouping is a form of subsetting, but grouping implies that the items grouped stay visible on the screen. Subsetting, on the other hand, has the connotation of reduction, although it is of course possible to create subsetting functions that allow the items that fall outside the subset to still remain visible. In a rich-prospect interface, this strategy would have the advantage of not disrupting the prospect while at the same time allowing the user to focus attention on some part of the collection.

An example might be an interface that allows the user to view the material in alphabetical order, with some indication of the alphabet visible. In a dictionary or phone book, for example, there is not only the larger structure of the alphabet, but also the guide words in the header, which give a more precise indication of the range covered on each page. If a rich-prospect interface were to employ columns of alphabetized text, a similar use of column headers might be useful in terms of identifying subsets of the collection that are of particular interest.

In the interface design community, one tool that is sometimes applied to the problem of subsetting is the interval slider, where a small horizontal or vertical bar represents the entire collection, and thumbs on the bar are positioned in order to select a subset of the total. On window sliders there is typically only a single thumb, since the goal is to specify the location of an insertion point. On subsetting sliders there are often two thumbs, which can be used to indicate the start and end points of the selection.

An alternative form of interval slider was developed by Eick (1994), who applied a painting metaphor in place of the thumbs. In Eick's model, sections of the bars are selected using a paint tool that can be applied to discontinuous portions, to create arbitrary selection groups for display. In a rich-prospect interface, a paint-based interval slider could be applied as a very flexible form of auxiliary selection device.

Prospect-Related Interface Tools: Renaming

One way of providing alternative browsing opportunities for the user without having to create multiple interfaces is to allow the user to select the meaningful representation of items from a list of options. This strategy is likely to be most effective in cases where there is a one-to-one correspondence between the various alternatives. For example, a collection of novels that provides alternative access by the names of the authors and the titles of the books has the disadvantage that one author might have written multiple titles in the collection. To allow the user to convert from a display of author names to a display of document titles is potentially disorienting, since the latter display will contain many more unique items than the former display. One solution to this problem would be to have both the author names and the titles listed together to form a meaningful representation that is a composite. Another possibility is to list the number of documents available for each author as a number placed next to the author's name, in which case the conversion to a display showing titles could derive from the numbers indicated. A third strategy would be to animate the conversion so that the transitions from author to title and back again are clearly shown.

Prospect-Related Interface Tools: Annotating

If the user is able to see the entire collection represented at once, and is able to sort and subset the material into various groups, it is also likely that some form of annotation would be helpful. At its basic level, this annotation function should allow the user to label the groups; at a slightly more sophisticated level, it should provide the ability to insert text, sound, images, or whatever the user desires at any point in the display, in order to help make sense of the whole. Within the context of a given collection and group of users, it may also be useful to consider having the annotations indicated by some form of visual cue, and to provide the users with the ability to switch them from visible to invisible.

There is also the possibility of the annotations of one user being persistent across sessions with the collection, which means that the system has to store not only the annotations, but also a user profile. Current strategies include having the user log in to a database or having the system automatically recognize the user's computer, which is not particularly helpful for users who do not always use the same computer. A hybrid approach is therefore to store a client certificate for the computer that provides the username, but still requires a password from anyone who wants to log in under that name.

Finally, there is the option of allowing the annotations of one user to be accessed by other users. Like interaction histories dealing with structure, interaction histories based on annotation have the potential to create new ways of understanding the material, independent of the discourse established by the original designers. Interaction histories also have their limitations – a primary one being that they are only as insightful as the people who create them. As with any kind of public system, such as bulletin boards, listservs, or chat groups, it may therefore be helpful in some situations to have a human moderator involved, so that the system does not deteriorate rather than develop through use.

Prospect-Related Interface Tools: Opening

Since the rich-prospect interface uses a simple representation of each item in the collection, there is the possibility that the user may wish to open the representation, either selectively for a subset of the collection or else for the entire display. The degree of expansion might be made available through a series of increments. Kaugars (1998) discusses a multi-scale text visualization that has four increments: closed; thumbnail; semi-open; and fully open. A rich-prospect interface using this strategy might logically be positioned between the closed and thumbnail versions of display, in which case it could be designed to provide various levels of meaningful representation.

For example, a display might use a single word to represent each item in the collection, but each of these words or some set of them could be expanded under user control to replace the single words with a list of keywords or a phrase. Further expansion might replace the phrases with sentences or short abstracts; then the short abstracts could be replaced with full abstracts, and so on until the full documents are open. If the selection mechanism is provided in a way that is relatively intuitive and simple to use, a fluid change from one form to another could be made available, so that the incremental steps are not disorienting to the user.

Prospect-Related Interface Tools: Structuring

The structure of the display of the meaningful representations is a significant part of a rich-prospect interface. In addition to methods for visually associating some items with others, there is also the possibility of arranging the items within some larger structure that has been designed specifically to make the user's work of examining the display easier.

One of the common structuring strategies for text items is to arrange them in columns rather than as a block of text. If the items are sorted alphabetically, they can also be marked with guides that indicate the first and last words in each column, or the range of the characters in the alphabet that the column represents. These elements perform a function similar to the page headers in a dictionary or the phone book, allowing the user to look through the display more quickly by scanning the headers than would be possible by looking only at the alphabetical list of items.

A related structuring strategy that applies to graphical objects as well as text is to arrange the display using a grid system. Grids have the advantage of allowing the designer to visually associate items through alignment, even in cases where the items might not be in immediate proximity on the screen. A typical example might be in the header or footer of a text document, where the author's name or article title might be flush left while the page number is flush right, but because these items occur on the same line, and are clearly outside the body text, the reader automatically associates them as both being part of the header or footer. Grid systems were widely employed by print designers for much of the past century, and have always been a basic feature of text layout programs. However, they are not yet strongly associated with the design of computer interfaces, perhaps in part because they have not been implemented as a standard component of web design applications. Unlike in the case of layout programs, interface design applications have not been derived from a tradition that includes the historical relationship between the technology of printing and the use of grid systems.

Incorporation of Prior Affordances

A rich-prospect browsing interface may not be the interface of choice for every user on every occasion. Interfaces are by definition the mediating software between an application or a data collection and the person using the application or the collection. Different tasks therefore call for different kinds of interfaces.

This fact also holds true in the analog world, although the logistical difficulties and costs involved in making multiple physical interfaces available for most tools have often been prohibitive.

The controls on a car, for instance, do not vary according to the intentions of the driver or the situation on the road. Navigating in city traffic, driving hundreds of miles of straight highway on a clear summer day, and rolling down a twisted mountain trail in a blizzard all use the same interface to the car, and the driver is required to adapt.

The constraints, however, are not so severe in the digital environment. There are no comparable physical reasons why there could not be different kinds of interfaces to electronic text collections depending on the different kinds of users or user needs. With such a strong cultural default in place for the analog world, there may be other reasons why alternative interfaces would not be acceptable to users. For one thing, there is the problem of having to identify and select among interfaces, unless the system does it automatically. If the system does not do it automatically, then the user has an extra step at the beginning of every task – namely, to identify the various options available and choose the appropriate one. It seems likely that the default interface would therefore be the one most frequently chosen. If the system does choose automatically, it may sometimes choose wrong, potentially leaving the user feeling frustrated or helpless.

A third solution is therefore to make the functions available in search interfaces also available in browsing interfaces. To provide existing affordances by reapplying existing technologies in a new context does not seem like an unreasonable approach, and certainly to allow users to search a rich-prospect display by typing words into a keyhole search field does not compromise the new affordances of prospect. In fact, because the meaningful representation of every item in the collection is available for feedback, there are some increased opportunities made available. The same reasoning holds true for a variety of strategies used in search interfaces, including the use of indexes, keywords, and relevance ratings, just to name a few.

An example of a commercial interface that could be repurposed in this way is the one used by Amazon.com. The current Amazon interface includes limited prospect in the form of a tab system that allows users to focus the search within different product areas. The search function provides a list of results, each of which contains a variety of information, including standard fields such as bibliographical material and details of pricing and delivery, as well as related information of a less-standard kind, such as sample pages, reader reviews, author statements, and a list of similar titles that might be of interest because they were part of purchase orders by other customers that also included the current book.

What Amazon does not currently contain is a system that allows the user to browse through a display of all the titles available. Since the number of possible products numbers in the millions, the entire collection is likely too large to be a good candidate for a rich-prospect interface. However, within a particular genre or subject area, it may be possible that there are subsets of the entire collection that could be represented in some rich-prospect form. Whereas a single one of the current search results often extend beyond the length of a screen, in a rich-prospect interface, the individual items would be represented in a form short enough that a thousand or more of them might fit on the screen without the user having to scroll down to see the entire set. The designers of the system could then provide a variety of tools or enhancements for manipulating the display, which would depend in part on how the collection items have been identified and indexed within the underlying database.

Typeahead Searching

One of the possible enhancements to an existing search function is through the rich-prospect interface equivalent of the typeahead, where the current search string as it is being created moves the insertion point on the display to match the text. The user of a search with a typeahead function therefore has live feedback on the success of the search even as it proceeds. If the colour or some other visual feature of the found string is also changed by the interface, the user also has a visual cue to identify the current position of the cursor. Typeahead functions have found commercial application in several document search systems, as well as some internet browsers (Mozilla 2002).

Typeahead searching, however, can only work on rich-prospect displays in cases where the item being identified is in the same category as the item being displayed as the meaningful representation of the collection. For example, if the collection is expressed as author names, and the user is searching by author, the system is providing appropriate feedback. However, if the rich-prospect interface is displaying authors and the user wishes to search by titles or keywords, the display is not helping the process. In this case, there are several alternatives. First, the system might respond by locating the appropriate title and highlighting the name of the author. The feedback would not match the input string, which is a serious problem. However, the user who has confidence in the system might nonetheless be able to understand that the items being highlighted or subsetted are those meeting the search criteria, even if the display is not the same. An alternative strategy is to have the system change the form of display to match the kind of search the user is performing. A combination of these strategies might be the most flexible solution, with the user able to specify the form of

display independent of the form of the search, but with the system providing an optional prompt for cases where the search and display do not match. A third option is to have the browsing interface deactivated when the search string fails to correspond to the display, under the assumption that the user is not interested in watching the contents of the browsing interface, but simply intends to perform a straight search. Further research is necessary.

Characteristics of Candidate Collections

Although the details remain to be discovered in general, and will likely vary significantly from one case to another, some collections are going to be better candidates than others for rich-prospect browsing interfaces. The relevant characteristics that need to be studied are:

- the possible uses of the collection
- the number of items in the collection
- the characteristics of the individual items
- the degree of homogeneity among items
- the possibility of providing some homogeneous meaningful representation of each item
- the extent of the markup of the collection

The Possible Uses of the Collection

Some collections may have been created for a specific purpose that precludes the necessity of any user ever wanting prospect over them. For example, a set of technical specifications for a manufacturer might be labelled with part numbers that are found in an index somewhere and used to retrieve the specific documents currently being required by the technical staff. Within the constraints of that environment and those users, the need for a rich-prospect interface showing a representation of all the technical materials seems minimal, especially if it were to consist of the relatively meaningless document numbers.

However, even in such an extreme case it is possible to suggest possible scenarios involving users and tasks that might find prospect useful on such a collection. For someone in management, for instance, it might be helpful to have an overview of the technical documentation, especially if the representation of the items in that case included additional information such as cost or maintenance cycles or sales totals. For someone in charge of the technical documentation system, a rich-prospect interface might help to provide reassurance that all the parts are where they should be—that no

documents have been mislaid or overwritten, especially if the display were to contain additional information on items such as a date and time stamp for most recent update, current file size in some meaningful units, or current status.

Similarly, the potential usefulness of prospect on something like a dictionary is relatively limited, since the primary function of the dictionary is to facilitate retrieval of information about a single word at a time. However, even in the case of a dictionary, some degree of prospect can be beneficial in certain scenarios. For instance, when the user is uncertain about the spelling of a word, a list of the words surrounding the word being sought can provide some cognitive reassurance, either that the correct word has been located, or that variants may be available that differ in relatively minor ways, such as in their inflectional morphology.

The Number of Items in the Collection

There are undoubtedly limits to what a human perceiver can integrate from a rich-prospect interface in a useful way, but those limits will likely vary according to a number of factors such as learning, experience, visual acuity, and motivation. Monitor size is also an issue, of course. A 21-inch monitor full of text, without vertical scrolling, can hold in excess of 2000 words of 12 point single-spaced Palatino, which is a reasonable size and font for screen display for most users. A 10-inch laptop monitor, on the other hand, can display roughly an order of magnitude fewer words – slightly in excess of 200. If one of the criteria of the design is that the prospect should not involve vertical scrolling, a good candidate collection for the laptop might therefore be one that contains only 200 items or fewer. There does not seem to be, however, any *a priori* reason to disallow vertical scrolling from a prospect-based interface. There does not seem to be any *a priori* reason to disallow horizontal scrolling either, for that matter. When viewing analog prospects, people do not find it unusual to have to turn their heads or even turn their bodies around in order to scan the horizon. What would be required in the interface, however, is some visual cue that there is more information available outside the current display.

The naive limits on text display mentioned above are not necessarily realistic either. The designer of a rich-prospect interface is able to employ any number of techniques to structure the information in ways that make it more accessible—some of these techniques may allow increased prospect on larger collections without compromising the advantages that accrue to the strategy of showing a meaningful representation of each item. It seems likely, however, that an upper limit on

the number of items that can be reasonably displayed using current desktop monitors might be ten thousand items at most.

The Characteristics of the Individual Items

The principle of return on investment for both the designer and the user can be applied in considering the kinds of collections that are good candidates. One way of applying this principle is to examine the individual items in the collection in terms of how useful they might potentially be. If, for example, the collection is fairly small and consists of very short items, such as single sentences or paragraphs, or small structured records, it may be possible to create a display that shows the entire contents of the collection rather than some meaningful representation. On the other hand, if the collection has short items in their hundreds of thousands or millions, a search system may be the optimum solution, and browsing solutions may not be possible. Finally, if the collection consists of fairly large items that are individually rich sources of information, then the overhead involved in designing a rich-prospect interface may be more easy to justify.

Some kinds of data may also lend themselves more readily than others to the creation of meaningful representations, although in general all kinds of information are routinely catalogued, indexed, and displayed in one form or another in library collections or on the web. An extreme case might be a collection of artifacts obtained in an archeological site, which might contain everything from pot shards to bones and inscriptions. If an archivist is tasked with recording diverse collections of artifacts, ranging from physical objects of unknown purpose to texts in undeciphered languages, it is necessary to create some form of useful labels, if nothing else than as indexes to a set of images or objects. These labels can also be used in a rich-prospect interface, although they will only be as meaningful there as they are elsewhere.

The Degree of Homogeneity Among Items

Within any given digital collection there can be a wide range of items that are not necessarily of the same class or in the same form. There might be, for example, sound files, video clips, text documents of various kinds, and digital images in any number of formats. Even collections of text documents can contain diverse kinds of items. General Electric Energy Services, for example, has in its research and development area the mandate of creating electrical substation automation hardware and software. Each component of the system has a set of associated text documents, including in-house testing reports,

technical documentation intended for client use, and marketing materials. In order to provide a meaningful representation of every item in this collection, it may be necessary to indicate in some way not only the content, but also the system components and the intended audience.

One means of providing some homogeneity is through a meta-tagging system that provides a similar structure for the information about each document, which is stored along with the documents. The Dublin Core, for example, consists of a set of fifteen meta-tags that can be used as part of a document header to provide the information needed to characterize a digital document for cataloguing purposes. These tags are:

- title
- creator
- subject
- description
- publisher
- contributor
- date
- type
- format
- identifier
- source
- language
- relation
- coverage
- rights (Dublin Core 2002)

The contents of any or all of these tags could be used as the basis for a rich-prospect display, depending on the information needs and intentions of the user. One disadvantage of the Dublin Core, however, is that the information tagged is quite general in nature, which limits the options available to the interface designer.

There are more complex encoding standards, such as the Metadata Encoding and Transmission Standard (METS), which is an XML schema developed by the Library of Congress (METS 2003). A METS document may include tags in the following five areas:

- Descriptive Metadata
- Administrative Metadata
- File Groups
- Structural Map
- Behavior (METS 2003)

Material from any one of these sections may be useful in developing ways of representing heterogeneous documents.

Other meta-tagging systems for document definition include the MARC encoding standard, which, like METS, was defined for use by library scientists, and COCOA, which was used by the Oxford Concordance Program and was later extended for use in TACT (Hockey 2000, p. 27). Tagging grammars such as Standard Generalized Markup Language (SGML) and eXtensible Markup Languages (XML) also allow developers to define tagging systems which can contain meta-tags for document definition.

The Possibility of Providing Some Meaningful Representation of Each Item

If a meta-system has been used and contains information that is meaningful to the user, the rich-prospect interface can be based on these kinds of tags. Whether meta-data is available or not, it is necessary to consider what the user brings to a given task in terms of prior knowledge about the field and expectations of what is appropriate or useful. For someone unfamiliar with law, for instance, it might seem reasonable to access a collection of case documents by the judge involved. Each case requires a judge; the judge's name is included in every document; and the decisions that set different kinds of precedent might reasonably be expected to cluster around particular judges. However, in the legal field, precedent cases are not conventionally accessed by judge, but rather by the names of the plaintiff and defendant. A collection of cases that used an interface based on the names of judges would therefore likely be useless to lawyers.

Different searches require different kinds of information: this fact is widely acknowledged in the design of search interfaces to library collections, where it is not uncommon to have different interfaces to allow access by author, title, publisher information, or keywords – that is, based on the various meta-tags that have been used to define the document records. Similarly, it may be useful to have different kinds of rich-prospect display for different kinds of browsing activity.

Another solution is to provide the user with a variety of information about each item. This strategy has been widely implemented by web browsers, which respond to the search string with a long list of possible links. Each link typically includes two or three lines of text, which means that screen space is sacrificed in the hopes that some of the information will be relevant enough to help the user decide which sites to access.

The Extent of the Markup of the Collection

Some digital collections consist of documents that have been tagged using a markup system such as those defined with SGML or XML. These kinds of collections are a special case, because they contain not only the information available to a reader of the text, but also information that is in some respects hidden from the reader by being contained in the tags, the attributes on the tags, and the values of the attributes. Depending on the complexity of the markup system that has been developed and applied, a collection of documents might have relatively simple information relating to formatting, or quite sophisticated information in the form of hermeneutic interpretations of the material contained in the tags, or some level of encoding in between.

Although any level of encoding is potentially useful in developing a rich-prospect display, the more sophisticated levels of interpretive encoding are particularly interesting opportunities to make the hidden intelligence in the tags available for perusal and use by the people accessing the collection. The next chapter will look in detail at the implications of rich-prospect interfaces in collections that have been textually encoded.

Design Issues for Rich-Prospect Interfaces

The primary problem with any rich-prospect interface is that to show so much information at one time is to invite disaster in the form of overwhelming the user. Designers working on interfaces based on rich prospect will therefore have to pay special attention to strategies for eliminating the sense of being overwhelmed by the display.

Hierarchies and Taxonomies

One method that does not provide rich prospect but can provide partial prospect and has been widely implemented is to categorize information according to some meta-schema, which allows users who know the system to traverse the collection efficiently. Well-known schemas include library

cataloguing systems such as the Dewey Decimal system and the Library of Congress subject headings; the biological taxonomy of Linnaeus; and chemistry's periodic table of the elements.

The problem with using a hierarchy, indexing system, or other taxonomy is that the information is effectively hidden behind the meta-schema. For people who are not familiar with the taxonomy or who do not ascribe to the presuppositions under which it has been constructed, the system can become a barrier rather than a tool. An example is the problem faced by academics in the late twentieth century working in Queer studies, who were interested in the history of sexuality, and in particular the issues of construction of gender and the development of the concepts of gay and lesbian, and their expression in literature and culture. The standard library cataloguing systems do not include the keywords "gay," "lesbian," or "queer" and it is therefore necessary for scholars in this field to attempt to identify appropriate texts by formulating alternative queries using the keywords that are available.

In the retrieval community, the twin concepts of precision and recall have been defined to express the degree to which the documents in a particular collection are amenable to being correctly located. Precision and recall are both ratios: precision is the number of correct documents retrieved over the total number of documents retrieved; recall is the number of correct documents retrieved over the number of correct documents available in the collection.

Unfortunately, in some cases the taxonomy defined for the documents is not appropriate for their content, or the people undertaking the indexing are not able to provide keywords that will allow other people to retrieve the documents, or the user is simply not able to make use of the taxonomy in searching, because although it might actually represent the documents in the collection and have been implemented properly, it does not coincide with the user's information requirements.

One strategy for addressing this problem is the application of facet analytical theory, where the content domain is divided into logical categories that are mutually exclusive. These categories can then be synthetically joined to form composite representations of each item in a collection related to that domain (Maple 1995).

Several methods of addressing these problems have also been developed using automated indexing systems; latent semantic indexing is one such strategy. Another solution is to attempt to profile documents by statistical methods. N-grams, for example, which are based on counts of fixed-

length sub-strings in a document, have sometimes been used as an entirely automated system of indexing documents without any need to address the semantics (Liu et al. 2000).

All of these methods, however, primarily use the computer as a retrieval tool, which filters the data for the user, rather than attempting to create an interface to the data that the user can employ in directly browsing the electronic collection.

It is possible, of course, for an existing hierarchy or classification system to be implemented as a visual component of an interface. A common strategy, for instance, is to provide 26 links that each represent a letter of the English alphabet. For a collection where items are represented by author names, under each letter will be found the documents that were written by an author whose name begins with the letter. This strategy has the advantage of subdividing the collection so that the user is not required to view lengthy lists. It has the disadvantage of providing no immediate prospect on the entire contents of the collection, so at a glance it is impossible to determine how many documents are available under each letter, or in fact whether there are any documents available at all. Although this interface may seem like an extreme case, it shares these limitations with many other forms of hierarchical display that attempt to conquer by dividing.

Screen Size

Although simply increasing the size of the display seems like an obvious solution, there are limits to what it can accomplish. Certainly the default 14-inch or 17-inch desktop monitor is not an optimum device for viewing large amounts of data. A 21-inch display, or two of them placed side-by-side, begins to meet the brief. A display the size of a large window or small wall should also be manageable for most perceivers. Various manufacturers have experimented with large screens and how they might be designed to provide adequate resolution without excessive weight or cost. Wallpaper displays have been developed in prototype by both Xerox and E-Ink, and consist of rolls of electronic paper that could be used to create displays of any conceivable size. The prototypes have existed for several years – Xerox had a partnership with 3Com to produce sample sheets in the late 1990s (Figure 2.13).

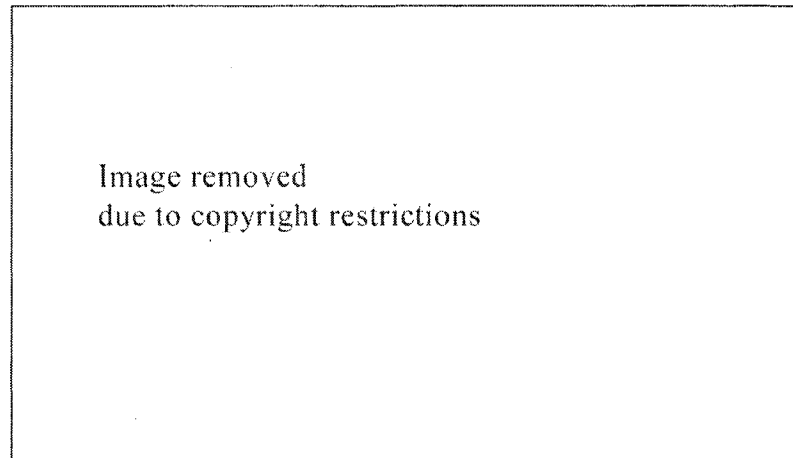


Figure 2.13 Electronic paper as conceived by Nick Sheridan of Xerox PARC in the early 1970s and manufactured in prototype rolls by 3Com in the late 1990s.

Other experimental large forms of electronic text display have been designed and created, often with the intention of providing increased forms of prospect or creating other new affordances. There is a prototype wall of electronic text on display in Xerox PARC, which combines an overview on the main wall with detailed information shown on sliding panels that change their content based on their present location on the larger wall (Xerox PARC 2001).

An example of an interface that presupposes a large screen is the one used by TextArc – a word frequency and collocation program that prints a long text (e.g. a novel) in a spiral around the outside of the display, then positions each word that appears more than once in the text inside the spiral, with links to its actual occurrences appearing when the word is selected (Figure 2.14). Multiple words can be selected at once, allowing users to quickly identify patterns based on co-occurrence of terms (TextArc 2002).


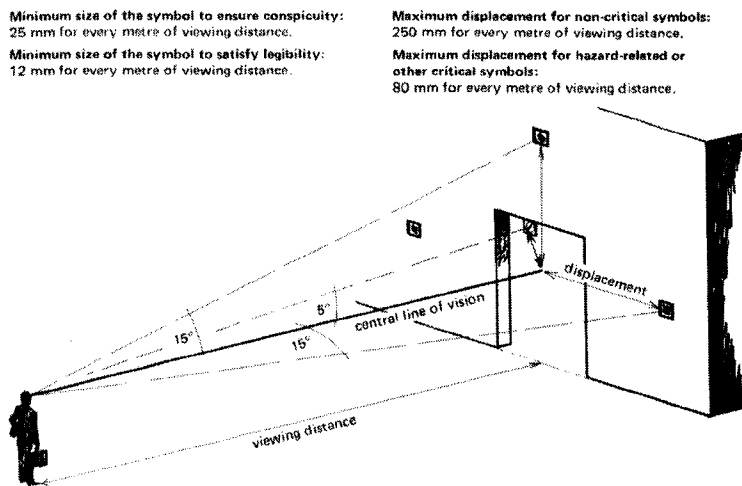


Figure 2.14 TextArc provides a visually striking means of investigating word frequency, distribution and co-occurrence in a long text. This display shows *Alice in Wonderland*.

However, simply providing all the data on a display the size of a domestic interior wall is not necessarily going to give the user a sense of prospect on a collection. The visual representation of the data needs to be designed in such a way that the minor features are seen as minor and the major features stand out. Human foveal saccades tend to cluster on areas of high contrast, such as edges between dark and light. Attention is drawn to these kinds of areas. Size matters. So does colour. There is a wide range of techniques for the visual construction of information, from the use of a grid system for layout to the tendency for the eye to take directional cues from the shapes of objects. Optimum line lengths have also been studied, at least for printed text, where what is at issue is the point at which readers are still able to accurately monitor line starts to prevent reading errors caused by skipping lines or re-reading the current or previous ones. In rich-prospect interfaces, these visual communication design techniques need to be applied so that the perceiver is able to make sense of the prospect quickly.

A related issue has to do with the limits on human visual acuity. The ratio of text height to viewing distance is another form of pi number, similar to the ones calculated by Warren (1984) for stair climbability. The ISO standards for public signage suggest that there should be 12 mm of image height and 4.5 mm of text height for every metre of viewing distance (Figure 2.15). These standards are based on the Snellen chart used by optometrists to study vision. In order to survey the contents of a wall-sized display, it is necessary for the perceiver to stand at some remove. As the size of the display increases, the perceiver needs to stand further back in order to be able to survey all of it at once. Another pi number could

therefore be calculated, comparing the size of the wall or other display with the amount of graphical or textual information it could contain at a size that is readable for a viewer able to survey all of it at once.



The terms and definitions taken from ISO 7239:1984, Fig. 6, Development and principles for application of public information symbols, are reproduced with the permission of the International Organization for Standardization, ISO. This standard can be obtained from any ISO member and from the Web site of the ISO Central Secretariat at the following address: www.iso.org. Copyright remains with ISO.

Figure 2.15 The ISO standards for signage suggest a minimum font size related to average viewing distance, as well as placement within certain ranges depending on the contents of the sign (Frascara 1984).

Persistence of Display

If the user is actively engaged with the rich-prospect interface, using various tools to reorganize or structure the meaningful representations of collection items, or subsetting or grouping them in some way, there is a question as to how the display should respond in terms of items that are not currently selected.

There are basically three possibilities. The first possibility is that the material that is not within the current selection disappears from the screen, leaving the user with an intermediate result screen that only shows partial prospect. The second possibility is that the unselected items remain visible, but the selected items are differentiated in some way, such as by colour-coding, highlighting, or removal to a section of the screen distinct from the rest of the display. The third possibility is that the unselected items as individual items disappear, but the user is given a visual cue of their continued presence, such as an icon at the bottom of the screen that can be expanded to recall the rich prospect. Further research will be required to determine which of these strategies is best under which conditions, or whether they are equally useful. One means of evaluating them would be to create affordance strength

vectors for each of the different interfaces, where the optimum form of display might be related to the subsequent task, such as adding items to the existing subset, changing the current selection in some way, or continuing to narrow the search by incremental grouping.

Priming

Human beings are able to locate and identify items more quickly if they have been primed to identify them by previous exposure, even if the people do not have a conscious memory of having seen or heard the precue (Baars 1997, pp. 118-9, 170). The strategy of attempting to prime users with some form of fleeting image could prove useful to the interface design community, especially if the contents of the visual priming were related to the structure of the collection.

For example, in a rich-prospect interface that was organized in columns like a phone book, it might be possible to load the data in two steps, with the first increment showing only the column or section headings that provide the larger framework, and the second step filling in the data. Further research is necessary.

Ventral vs. Dorsal Stream Perception

Milner and Goodale (1995) suggest that there are two streams that are used for visually processing information in the human brain, and that one stream relates primarily to concept formation, while the other relates primarily to opportunities for action. If there are two distinct but interacting mechanisms, then it may be possible to design an interface in such a way as to facilitate either action or reflection, depending on the nature of the task. In addition to the possible implications for design, there are also implications for the study of interfaces and their affordances. For example, if dorsal perception (for action) is primarily tacit, while ventral perception is explicit (Michaels 2000, pp. 252-3), then it may happen that affordance strength vectors based on user reporting will be less accurate than affordance strength vectors based on evaluations by a third-party observer. Further research is necessary.

Mental Models

The mental model of the user in undertaking a task can have measurable effects on performance. In a study of wheel rotation responses, Guiard (1983) asked participants to control a cursor using a joystick, where the response direction was counter-intuitive: moving the joystick to the left moved the cursor to the right, and vice versa. One group was instructed in the mechanics of the response – namely, that the task was to control a cursor using a joystick – while the other group was told that

the joystick was actually affixed to the underside of a steering wheel. The task for the two groups was identical, but the group with the steering wheel metaphor performed significantly better than the group who had not been provided with the metaphor.

Although metaphors are often considered as comparatively esoteric artifacts belonging primarily to the realm of literary expression, Lakoff (1980) makes the strong case that metaphoric thinking is in fact a widespread strategy and might correctly be understood as a fundamental part of human cognition. Drawing on examples from English diction and idiom, Lakoff demonstrates that metaphors structure a wide range of language, and by implication, thought. Metaphors are therefore a potentially powerful tool for the interface designer attempting to create intuitive electronic artifacts, although as Stubblefield (1998) points out, there is a necessary degree of caution required to ensure that the developers and users share a common understanding of the implications both of the metaphor itself and of the consistency of its implementation in a particular system.

The classic use of a metaphor to create a mental model for interface tasks is the computer desktop. However, the strategy of providing the user with a mental model for a task is amenable to extension into a wide variety of possible activities, including the use of rich-prospect browsing interfaces, where provision of a mental model appropriate to the interface, collection, or task might help to reduce the sense of visual overload.

Sequential vs. Spatial Prospect

Some previous researchers have suggested that a form of prospect is possible through a combination of an index and a sequential display. Ahlberg and Shneiderman (1994), for example, presented the Alphaslides, which was a form of horizontal scrollbar with an internal index consisting of letters of the alphabet. The letters were spaced according to the number of documents in the collection, giving the user some limited sense of prospect. The primary strategy, however, was to have the titles of the items in the collection appear in rapid sequence in a display placed just above the slider. The items appear and disappear as the user moves the mouse, so it is possible to flash quickly through an alphabetical sequence. Novice users could locate a film title out of a collection of ten thousand titles in an average of 24 seconds, which according to Ahlberg and Shneiderman compares favourably to menu selection systems containing an order of magnitude fewer entries.

Sequential display occurs in many systems that attempt to provide prospect in spite of limited screen space. Vertically-scrolling windows and panoramas, for instance, both employ a form of

sequential display, as do interfaces that use selective magnification as a tool. Ahlberg and Schneiderman (1994) also mention the possibility of using the Times Square strategy of having text scroll past the user rather than having it appear in rapid sequence.

The question is whether the prospect provided by these means is adequate to create the various new affordances that are potentially available from more spatial forms. Further research is necessary.

Inter-Affordance Effects

In designing to provide new affordances, there is always the possibility that existing affordances will be affected in some significant way, and previous affordances will be reduced or lost as the new affordances are made available. In software development projects in general, unintended consequences of incremental changes can be guarded against in various ways, including modular design and the practice of re-testing against a standard testbed that grows as the application expands.

Interaction Histories

In addition to providing prospect on a collection, it is also possible to provide a form of prospect related to interaction history, or the activities of previous users related either to individual documents in a collection or to the rich-prospect display as a whole. Interaction histories have been discussed earlier in this chapter in terms of both structure and annotation.

In addition, Hill et al. (1992) presented an interesting set of widgets designed to give prospect on document editing and reading. Based on interval sliders, the edit wear and read wear scroll bars show internal lines that corresponded to areas of the document that have been edited or read. These marks allow users to see at a glance which areas of a document have received the most attention from previous editors and readers. In cases where discrete lines of marks are created to correspond to different periods, they also show which areas of the document have received attention most recently. The concept of graphically presenting edit or read wear could also be extended to entire collections of documents, as suggested in the case of web browsers by Wexelblat and Maes (1999).

Co-ordinating Multiple Views

Another natural extension of the idea of prospect in an interface is to provide multiple views that show the collection at different levels of granularity. These views are usually displayed simultaneously, but may also be shown consecutively. Baldonado et al. (2000) suggest eight rules to govern the development of interfaces that incorporate multiple views of either kind:

- **diversity:** multiple views are appropriate under the following conditions: when the collection has diverse attributes; allows for diverse models or levels of abstraction; or contains different genres. Multiple views can also help when the user profiles are diverse
- **complementarity:** multiple views are useful for collections where different views can reveal patterns or disparities.
- **decomposition:** use multiple views to help the user divide and conquer complex data.
- **parsimony:** since multiple views add complexity for both the designer and the user, they should be used sparingly.
- **space/time resource optimization:** there should be a return on investment for both the designers and the users
- **self-evidence:** perceptual cues such as highlighting or coupling should be used to keep relationships between the views as clear as possible, although coupling needs to be judged against difficulty and speed, and should not be unidirectional.
- **consistency:** the interface for each view should use the same features in the same ways.
- **attention management:** the interface should use perceptual cues to help direct the user's attention appropriately

In the case of rich-prospect interfaces, the allocation of time and space would need to be considered as a fairly central issue, since the rich-prospect display alone would likely require significant screen space, and any windows displayed at the same time would create an additional demand where the demand is already heavy.

Performance

System performance for rich-prospect interfaces is also going to be an ongoing issue. If the user needs to wait for the system to download a screen full of data before the process of looking for documents can even begin, frustration is going to result in many cases. As the internet matures, these performance issues may become less significant, provided that the nature of the collections does not also mature into forms that involve larger representations. It is currently within the constraints of the technology to download a screen full of text relatively quickly; to download a screen full of video thumbnails, however, would still pose a problem for most users.

Characteristics of Candidate Tasks

Although one of the advantages of a rich-prospect interface is that it should allow ready access to collection items even to people initially unfamiliar with the collection, the various features of such an interface and the tools that might go with it will have a learning curve. User motivation to work with the rich-prospect interface will vary, however, depending on a number of factors related to the user and the task.

Just as some collections will be better candidates than others for the development of a rich-prospect interface, so some task characteristics, within the constraints of a particular user at a particular time, will be better suited to the use of such an interface. For example, users who have an understanding of the collection and its significant features that is congruent to the presuppositions of the designers will tend to find a rich-prospect interface more useful than users who do not share the same presuppositions.

Another user characteristic that might be useful is previous positive experiences in using rich-prospect interfaces, or conversely, previous negative experiences in using interfaces without some form of prospect. Such a statement, however, could be made about any form of technology. Unique to the rich-prospect approach is the need for the designer to assist users in considering screens full of information as an opportunity rather than a source of frustration or intimidation.

CHAPTER 3: PROSPECT ON INTERPRETIVELY-TAGGED TEXT COLLECTIONS

It is clearer now than ever that inserting markup in a text is an act of interpretation. This raises questions of what the interpretation is and therefore who is doing the markup (Hockey 2000, p. 48).

The most common kind of digital text collections contain only the actual texts of the various documents, stored either as entire books or else subdivided into chapters or other sub-divisions. In some collections, however, these digital texts have been augmented with a layer of textual markup that is normally invisible to the reader, but which can serve as information the system can use either to format the document for display or to supplement the search function. An even smaller subset of digital text collections has an interpretive level of markup, which provides information beyond what is required for formatting or retrieval. The value of interpretively-tagged text collections is that they contain a level of markup that represents the contribution of intelligent judgment by people who have read (or, in the case of the Orlando Project, written) the texts. This markup encapsulates human reading and analysis of the text for subsequent readers.

Textual markup requires a tagset, which can be defined using one of the markup grammars such as Standard Generalized Markup Language (SGML) or eXtensible Markup Language (XML). SGML and XML share many features, including the capacity to define not only the tags, but also attributes that are associated with the tags. Tag attributes are the technical means of providing an interpretive level of tagging, since the values contained in the attributes allow the tagger to attach information not present in the text. A common example of a tag attribute is the Standard attribute on the TEI <Name> tag, which can be used to specify a single spelling for someone's name, regardless of how the person is identified in the text that the tag marks.

Another common feature of the tagsets definable with SGML and XML is that the tags form a nested hierarchy which can be resolved to a standard tree structure. A tag tree is useful because it can be used to facilitate rapid traversal of tagged documents, which means the system can check the syntax of tagged documents for oversights or errors. However, a side effect of this restriction is that SGML and XML tagsets can be used to mark instances where tags interleave, but only through repurposing a function originally intended to allow simultaneous use of multiple tagsets, or else through including a null tag that acts as a marker, and also reduces the effectiveness of syntax checking.

For example, if someone wanted to tag each complete sentence in a manuscript, but also wanted to tag a point where the author had crossed out the end of one sentence and the beginning of the next, a tagset defined in SGML or XML would have to either use a null tag or rely on a second

interwoven tagset (Renear 1996). Researchers working with texts written by the philosopher Wittgenstein have therefore created an alternative markup grammar called the Multi-Element Code System (MECS), which allows tagging overlapping elements (Wittgenstein Archives 1998). MECS has the disadvantage of not allowing automatic syntax checking through traversal of a tree hierarchy. It may therefore also have restrictions in terms of generating rich-prospect forms of display based on the tagging of individual documents, although the details of such limits may only become apparent in the design of a rich-prospect interface for an actual MECS system.

CHAPTER OUTLINE

In general, it is possible to divide the options for providing rich prospect on an interpretively-tagged collection into three domains: contents, tagset, and tagging. Since the provision of prospect on the contents of a collection can be extended to any digital collection – whether or not it has been tagged – the provision of rich prospect with respect to content has already been discussed in the previous chapter. The topics addressed in this chapter are rich prospect on the tagset, and rich prospect on the way the tagset has been implemented.

This chapter is divided into two main sections: Textual Markup, and Information Visualization. The first section addresses several issues, including: tagging as an act of interpretation; the possible new opportunities for action provided by a rich-prospect interface to the tagset; and the possible value of having some form of prospect on the actual tagging of the documents. It also examines several related issues, including the role played by visual culture; the relationship between rich-prospect interfaces and complexity; and how rich prospect relates to the concepts of constraint and natural mapping. The contents of this first section might be summarized as a response to the question: why is prospect on the markup, as opposed to prospect on the contents, potentially useful? The answer relates in part to the kinds of information that the user might obtain by having prospect on the tagset, and how these kinds of information might be applied in understanding and accessing a collection. It also relates to the question of how the user is able to gain confidence with using a collection through having various assurances of what the collection contains, as well as assurances about how it has been understood by the people who created it, and how that understanding might translate into various approaches to accessing the materials.

The second section – Information Visualization – takes as its starting point the assumption that rich prospect on the markup of a collection is going to be useful, and examines some of the

concepts involved in attempting to provide it, including a discussion of strategies that have been used for documents without markup but which may prove useful by extension, as well as the strengths and weaknesses of the various approaches, both for a tagset as an entity in its own right, and also for the implementation of the tagset in a given collection. Whereas the first section of this chapter addresses the question “why?”, the second section is primarily directed at the question “how?”

TEXTUAL MARKUP

There is a sense in which textual markup can be said to exist whenever anyone creates an electronic text document that contains formatting. The markup in this case is everything about the visual appearance of the document that is not strictly text, including the font, the style of the font (e.g. normal, bold or italic), the page header or footer, the margins, the indentations, and so on (Burnard 1995). This broad definition of textual markup is useful in that it calls attention to the common nature of text formatting across documents, and therefore also to the logic of attempting to standardize the formatting commands. However, the narrower and more conventional understanding of textual markup is that it consists of a standard system that can be used to specify information about a document by inserting tags around sections of the text inside the document. These tags are generally intended to standardize the content for purposes of formatting and automated searching. It is consequently the normal practice to keep the tags invisible to the reader.

Standardized textual markup systems, such as SGML and XML, had their genesis as a means of addressing cross-platform formatting issues. The problem consisted of a proliferation of means of textual production that had not been standardized. There were potential compatibility issues among any set of components, including the computer monitor and other hardware, text processing software, the printer (which included its own configuration of hardware, software, and firmware), digital fonts (both for display on screen and for printing), and a variety of related utilities. Professional printing facilities introduced another layer of complexity, since they used their own configurations of equipment, software, and fonts, which were often unavailable for the desktop environment. For people using mainframe equipment, text processing had its own set of proprietary tools and requirements. Finally, there were many non-trivial problems related to typefaces, where the provision or absence of diacritics, mathematical symbols, and characters from non-Roman alphabets had also not been standardized.

It was felt that if descriptive markup could be inserted into text, many of these cross-platform compatibility problems could be reduced or eliminated. Coombs et al. (1987) suggest a taxonomy of

six categories of textual markup: punctuational, presentational, procedural, descriptive, referential, and meta-markup. The preferred category for adoption in a standard markup system is descriptive markup, where the user does not directly specify formatting (e.g. “this text should be in a bold typeface”), but instead specifies the structure of the document or the intention of the author (e.g. “this text is a major heading”). The translation from the descriptive markup to the formatting capabilities of a given platform could then be carried out as an intermediate step by any software that was capable of interpreting the markup. HTML is an example of a markup language that was originally used in this way. As Price (1998) summarizes it, descriptive tags are nouns that express the nature of the text, while procedural tags are verbs that express how the text should be processed.

Tagset Definition as Interpretation

In discussing markup languages, it is useful to distinguish among three different but related entities. First are the markup grammars, such as SGML, XML, or MECS, which are standard systems for defining sets of tags. Next is a particular set of tags for use in a given collection. This set of tags, which may also be called the tagset or Document Type Definition (DTD), can consist of any combination of tags, attributes on the tags, and pre-defined values for the attributes. HTML is an example of a tagset. Finally are the instances of the tags as they have been applied in marking up a particular document.¹

As Hockey (2000, p. 48) points out, marking up a document always involves an act of interpretation. A significant phase of that interpretive act takes place when the tagset is defined. The tagset establishes the bounds of discourse within which the people doing the tagging are going to work. If the tagset contains a tag for a particular element of interest, the tagger has a tool at hand for identifying that element whenever it occurs. As a hypothetical example, if the tagset contains an element called <firstname>, then whenever the tagger finds the first name of a person, the system allows it to be marked. If the tagset does not contain a tag for a particular element, then the tagger either has to fudge the system in such a way that the element can still be tagged, or arrange for a modification to the tagset, or else leave the element unmarked.

¹ In fact, Price (1998, p 174) identifies a total of 11 different kinds of markup: the SGML declaration; the document type declaration; entity declarations; notations; element declarations; attributes; comments; marked sections; short references; links; and system-dependent processing instructions. Of these various categories, only links as a special kind of tags are discussed near the end of the chapter.

Fudging the system is problematic in that one of the primary purposes of tagging is to facilitate retrieval through the imposition of standard forms in free text. Any use of the tagset in non-standard ways may therefore defeat the intention of tagging the collection in the first place.

Modifying the tagset is problematic in that the change may imply the need to revisit previously-tagged documents in order to identify instances of the new tag. If the tagset has not been implemented in a way that is as consistent as possible both within and across documents, the tagging of the collection is once again compromised.

The simplest solution for the tagger who encounters material that is not defined in the tagset is therefore to leave that material unmarked, even though it might actually be of use or interest to people accessing the collection. An even greater loss of potential value can occur when the tagset includes predefined attribute values that limit the tagger from attaching new material that may be significant. In this case the limitation is not in the tagset, but rather stems from a restriction on the attribute values, which in turn derives from the natural desire on the part of the designers of the tagset to maintain as much consistency as possible to facilitate retrieval.

The choices made in the course of defining the tagset are therefore an indication of how the people responsible for tagging a collection of documents understood the collection. The tagset is one indication of what is considered significant enough to be marked for standard retrieval. It is also possible to interpret the various tags as belonging to levels of interpretation. Ruecker (2002) suggests a taxonomy of six kinds of markup:

- raw data: unmarked text, esp. in cases where a string is marked in one instance but the same string is unmarked in others
- descriptive markup: primarily intended for use in document formatting, as defined by Coombs et al.
- meta tags: classify the document, as in the Dublin Core
- internal glosses: standardize content for easier retrieval, often through the use of tag attributes and their values
- external glosses: add new content not found in the text, again primarily through tag attributes and values
- hermeneutic: interpret the content

A given tagset may contain any configuration of tags at any of the levels: it is not necessary for a tagset to contain tags at each level, nor is it necessary for lower levels of tagging to be present in order for higher levels to be used. There has been, however, a tendency for developers to concentrate their efforts at the first four levels. The Orlando Project is one exception to this tendency, where the three tagsets defined for the project contain tags at all six levels, including external glosses and hermeneutic tags.

Another issue related to the definition of tagsets is the extent to which information is available as content, tags on the content, tag attributes, or attribute values. For instance, in the Orlando Project the people tagging the documents were also responsible for writing the documents. They therefore had the opportunity not just to standardize the content by applying tags, but also by using text forms that would be amenable to searching as text strings. However, for stylistic reasons it was often more appropriate to use a Standard attribute on a tag, which allowed the writer to vary the content while still providing the system with the information necessary for retrieval.

In other cases, the designers of the tagset may have to make choices as to which tag attributes have fixed value lists, and which attributes can take any value. An example of this kind of choice in the Orlando Project is the Standard attribute on the <GenreName> tag, which was originally defined to take any value. However, while the collection was being finalized for publication, this attribute was redefined to take a fixed value list, which was composed based on the list of genres that had been developed by the taggers during almost ten years of document development.

Document Tagging as Interpretation

Within the bounds of discourse set by the tagset, there is another layer of interpretation that is carried out by the people who apply the tagset within a given document or collection. As the taggers implement the principles embedded in the tagset, various decisions are necessary which may constitute a non-trivial degree of interpretation. To continue the earlier hypothetical example, in applying the <firstname> tag, if the document contained the name “J. K. Rowling,” would the tagger leave the name unmarked, since no first name is specified, or would it be better to mark the letter “J,” or would the right choice be to mark the letter “J” and include the author’s actual first name in an attribute associated with the tag?

In some projects these kinds of questions may be addressed in documentation that is developed either in conjunction with the tagset, or else in iterative form as the questions arise (or

both), in order to establish best practices for the project. Since one of the primary functions of markup is to provide a degree of standardization, these kinds of guidelines for taggers represent another potential source of information about the interpretations that might be expected in a given collection. The extent to which this information can or should be made available to the users of the collection is another area of potential research. In cases where it may be reasonable to attempt converting the best practices guidelines for taggers into some form of useful overview for the users of the collection, it seems reasonable to suggest that this kind of information would therefore be a candidate for inclusion in a rich-prospect display of the tagset.

Although clearly the external glossing or hermeneutic tags are interpretive in nature, there is a sense in which even the imposition of descriptive markup is a form of interpretation. Examples from the analog world include the division of Shakespeare's plays into acts and scenes by later editors (the First Folio of 1623 lists only lines), and the division of Old Testament texts into chapters and verses, when the Hebrew originals are in continuous unpunctuated text, often with consonants only, since the vowels took up precious space and could be filled in as required during the reading of the texts by the original writers and their community. In the history of book design, there was also a transition from physical scrolls to codices, where the scrolls were divided into discrete pages. These various ways of indicating subdivisions of text, often accompanied by the addition of page numbers, are useful in several ways to the reader.

First of all, descriptive subdivision and pagination allow rapid non-sequential access to parts of the text, either by someone reading the material for the first time under the direction of someone familiar with it, or else by someone revisiting the text. Second, they provide a standard means of referring to parts of texts, either in speaking or writing about them. Subdivision - or descriptive markup - is, however, an intrusion into the original text of additional material that can have significant consequences to the perception of the reader. For example, consider the chapter that ends on a dramatic note or cliffhanger. The reader often feels urged to read in chapter-sized sections, and on encountering a cliffhanger ending, finds extra motivation to begin a new chapter. If the chapters were sub-divided differently, this effect would not take place.

Although even the imposition of descriptive markup may constitute a form of interpretation, if the tagset contains tags at higher levels, the degree of interpretation by the tagger or by some other expert who guides the tagging is correspondingly greater. For example, in a hypothetical collection of

documents related to legal decisions, it would be possible to have an external gloss tag that identified precedent cases that were relevant for the current decision. The choice by the tagger to mark a given case with the related legal precedents represents a significant addition of information about the decision. In order for the tagger to be competent to apply such a tag, it would be necessary for someone with adequate legal knowledge to make the connection between the current case and the precedent cases. The necessity for this domain expert to be involved in order for the tagging to be accurate is one of the indications that additional intelligence is being encoded in the markup.

Rich Prospect on the Tagset

The degree to which textual markup can be used to facilitate retrieval is to some extent dependent on the level of the tags. However, some research has indicated that even a descriptive level of tagging can be useful in providing the user with improved forms of retrieval. Myaeng et al. (1998) developed a retrieval system based on structural SGML markup which not only improved overall retrieval performance at the document level, but also added new affordances for the retrieval of subsections of the documents. Depending on the tagset, some elements proved more useful than others in facilitating these retrieval functions. This is only natural, since tags indicating structural elements such as chapter or section headings often provide keywords that indicate the content of the subsequent material, whereas tags indicating structural features such as paragraph breaks are comparatively devoid of indexical content.

For interpretively-tagged collections, the value of the tagging in facilitating improved retrieval mechanisms seems indisputable (although future studies will be required to determine how significant tagging proves to be within the constraints of a given community of users of a particular collection and its interface). An interpretively-tagged collection by definition contains a number of tags that can be guaranteed to contain information that in some way represents and perhaps even provides additional insights into the content. Rather than needing to rely on structural cues as to what is important, the interpretive tagging makes the connection explicit between the content and the tags, in terms that an automated retrieval system can access.

Given this improved functionality, the question can be raised whether there is any advantage to the user in being able to obtain prospect on the collection's tagset and tagging. It is at this point that it is important to emphasize the interpretive nature of both the tagset definition and the act of tagging. If the sole purpose of document markup is to facilitate retrieval, either of the entire document or of some

portion of the document, then the retrieval system needs to access the tags, but the user does not necessarily need to even be aware they exist. For example, a user looking for information about J. K. Rowling might enter the search term “Joanne Rowling,” and the <firstname> tag would allow the retrieval system to make the connection and return the reference, even if the text actually says “J. K.” rather than “Joanne.” The invisible tagging system within the document would in this case provide a standardized form for the retrieval system to access, and the user can therefore obtain the desired search result without consciously invoking the markup system.

However, even in a retrieval context it may be possible to make the case that access to the tagset would help the user to perform more successful or more accurate searches, since information about the tagset can potentially provide the user with insight into the retrieval mechanisms and how they are being supported by the markup. If the query mechanism is defined appropriately, it may be possible for the user to construct queries through selecting appropriate tags, filling in possible values for the attributes or contents or both, then submitting the resulting query to the system.

Various strategies are possible for presenting the tag information for use within a query constructor. One approach is to provide the user with a query wizard, or stepwise procedure that walks the user through the process. Query wizards are particularly useful for systems designed to facilitate what is sometimes called “day one performance,” where users are either accessing the system for the first time or else use it infrequently enough that they might as well be using it for the first time (Karat 1997). Automated banking machines are an example of a technology that has been designed for optimal day one performance, and many of the interfaces to banking machines use the equivalent of a wizard, where the user is asked to provide one piece of information at a time, often by selecting from pre-defined lists. For sophisticated users, however, there is often a problem with systems designed to facilitate day one performance—namely, that the system can be irritatingly slow and systematic in its sequential approach to the task. Telephone navigation systems are one example of a day-one performance design that is notorious in this regard.

As an alternative to query construction by stepwise procedures or wizards, there is also the possibility of providing the user with access to the various pieces of the system, coupled with examples that illustrate the principle of how final queries should appear. The user is then able to create queries directly from whole cloth, rather than relying on an automated system to assemble them from individual elements. In either case, the display of the tags, attributes, and attribute values that are available in the

system can potentially contribute to the process of query construction, provided that the display is meaningful to the user and that the use of it is evident in the manifest affordances of the interface tools.

It is also possible to extend the purpose of the document markup to include more than just its role in facilitating retrieval. Both the tagset and its application in a particular document can be used to provide opportunities for the user to understand how a collection has been interpreted, both by the people who defined the tagset, and by the people who did the tagging. This level of information can help the user to choose strategies for accessing the collection, either through retrieval by query, retrieval by directly opening documents, or retrieval mediated in some other fashion by the browsing interface (for instance, by providing direct access to the tag contents at various passage levels, as opposed to always opening the entire document).

Direct insight into the tagging system, like direct insight into the contents of a collection, has the possibility of providing new affordances in a variety of areas, including the following:

- contents: what tags, attributes, and attribute values have been defined?
- structure: how are the tags nested within a hierarchy?
- context: for a given section of content, what choices of tags were available?
- features: does the tagset contain any characteristics that are unique, surprising, or particularly helpful to the user?
- limitations: to what extent can the tagset facilitate searches in the area of interest to the user?
- connections: does the structure of the tagset and the definitions of the individual tags suggest new ways of viewing any of the material?
- trends: what kind of themes are discernible from the pattern of the tags?
- anomalies: are there any highly unique tags, attributes, or attribute values that fall outside the larger pattern?
- reminders: do some of the tags suggest interpretations of the content that would otherwise have been forgotten?
- reassurance: insofar as the tagset provides insight into the discourse of the organizers of the collection, it may serve to reassure users that the system either matches or else fails to match their expectations.
- reduced helplessness: for people unfamiliar with the system or the content domain, the tagset may provide some framework for understanding the material.

Rich Prospect on the Tagset: Content

The contents of the tagset include the tags, the attributes, and any pre-defined values on the attributes. Each of these contents may be part of a separate representation, but it makes sense that they be combined into a single form, since the attributes and their values are dependent on the tags they are associated with.

If the user has insight into the tagset, it is possible first of all to begin to interpret the markup in terms of its potential value as a retrieval aid. Assuming that the interpretation of the collection indicated by the tagset coincides with the retrieval needs of the user, then there is also the opportunity to use the information about the fixed list of attribute values as a kind of fixed search vocabulary, with some degree of certainty that the retrieval system will recognize the terms.

Neither of these functions are available in retrieval systems that do not make the tagset accessible to the user. A user of a system with limited prospect or no prospect at all may still gain some insight into the value of a tagging system for retrieval purposes by noticing that certain kinds of searches yield better results than other kinds, but the details of how that mechanism is assisting the search will remain obscure. A search vocabulary based on the tagset may also be made available through methods other than a rich-prospect interface (as for instance by using a thesaurus with a lookup table and suggestion mechanism, or by giving the user a picklist of terms), but then each of the other advantages of the rich-prospect approach may not necessarily also be provided.

Rich Prospect on the Tagset: Structure

If the user of the collection can examine the structure of the tagset, that information may help guide the process of searching for relevant information. Knowledge of the structure is particularly relevant in cases where similar kinds of tags have been defined at different points in the hierarchy. As a hypothetical example, a collection dealing with automotive parts might have the following nested tags in the tagset:

```
<engine parts>
  <4 cylinder>
    <carburetors>
  <6 cylinder>
    <carburetors>
  <8 cylinder>
    <carburetors>
```

The contents of these tags might consist of part numbers, part names, or descriptions of specific automotive parts, or perhaps of manufacturer or retailer information. Alternatively, the contents might consist of troubleshooting routines relevant to each kind of tag. The routines generally relevant to 4-cylinder engines would be marked with the <4 cylinder> tag, while those procedures specific to 4-cylinder carburetors would be identified with the subtag of <4 cylinder> called <carburetors>.

For a user interested in carburetors, searching for anything tagged with a <carburetors> tag would therefore return parts or procedures relating to all three kinds of engines. If, however, the user were only interested in carburetors for engines with 6 cylinders, knowing the structure of the tagset would allow a narrower search that specified a nested tagging of both engine size and carburetor.

In terms of insight into the interpretation applied to the collection, seeing this tagset might suggest two ideas to the reader. First, it would appear that the people responsible for the collection distinguish engine parts as a class distinct from other kinds of automotive parts. Second, since <engine parts> is a high-level tag in the hierarchy, what is primary about an engine – what distinguishes one kind of engine from another – is the number of cylinders.

Organizing the material in this way is not, however, the only way to organize a collection dealing with automotive parts. Another collection on the same topic might ignore distinctions of this kind altogether, and focus instead on manufacturers and product lines. A third collection might subdivide the information according to year of manufacture, a fourth by the geographic warehouse location where the parts are physically stored, and so on. Each of these organizational schemas is potentially significant to the user looking at the collection, and each schema could be intrinsic to the tagset. Showing the user the tagset in some form is therefore one means of providing useful information about the way the collection has been organized.

Rich Prospect on the Tagset: Context

A related kind of information is the framing of a particular tag in the context of other similar tags. Depending on the tagset, it may happen that different tags have been created in order to make distinctions among conceptually similar items. If the user is given access to some form of the tagset, the existence of these distinctions may become evident.

In the example above, the automotive parts collection has three different carburetor tags, or at least one carburetor tag that may occur inside three other tags (SGML and XML have syntactic means

of indicating details such as whether a tag is mandatory or optional, or in what places it can nest). The existence of three tags instead of one suggests that different kinds of carburetors are potentially going to be identified in the collection.

The automotive parts tagset may also contain tags for parts related to fuel injection systems, which are an alternative to carburetors. If the user of the collection sees the <carburetors> tag near the <fuel injectors> tag, the proximity may cue an awareness that the collection holds material on both kinds of fuel system parts. For those users who may not have made the logical connection between fuel injection systems and carburetors, this structural information may also cue an awareness of the similarity between the two kinds of devices.

Rich Prospect on the Tagset: Features

The features of a tagset might be considered as a subclass of the tags, attributes, or attribute values, in that a features represents characteristics that are in some way surprising or unique. The features may potentially be useful either for retrieval purposes or for developing a new insight about the collection.

One form of uniqueness has to do with the level of tagging that has been defined. In a tagset that consists primarily of tags at the descriptive level, it may be possible that some subset of the tags have been defined at the higher level of external glosses or hermeneutic tags. These tags might be considered a special feature of the collection. In the case of a tagged copy of the primary text of a play, for example, the tagset may include tags to indicate acts, scenes, and lines. It may also contain tags that can be used to indicate alternative versions of the text. This material, which in a printed edition might appear in footnotes or some other form of critical apparatus, may represent a significant addition to the collection, the presence of which would be signaled by the existence of the tags intended to mark this kind of information. By providing the users with insight into the tagset, there is the possibility that such features will be drawn to their attention.

In the case of tags with significant attributes, the feature might consist of the range of values the attributes indicate as possibilities. To take a hypothetical example, a tag in a collection dealing with gardening tools might contain a tag to indicate that something is a kind of <handsaw>. A typical collection might provide attributes on the <handsaw> tag such as <handsaw:manual>, <handsaw:circular>, and <handsaw:jigsaw>. This taxonomy of handsaws, however, is not the only one possible. An alternative taxonomy might instead have attributes related to the nation of origin of the saw. The tag in that case might include attributes such as <handsaw:American>.

<handsaw:Danish>, and <handsaw:Japanese>. This list of attributes suggests that the designers of this collection saw handsaws in a way that is not typical, and these documents may therefore prove worthy of further investigation by someone looking for information about a handsaw that is out of the ordinary.

Rich Prospect on the Tagset: Limitations

Direct insight into the tagset can also provide the user with the opportunity of identifying what the limitations of the markup on a collection are going to be. The most obvious limitation is in the case where a tag simply does not exist to mark a component of interest to the reader.

For example, in the TACT markup used on Ovid's *Metamorphosis*, the developers created tags to indicate names, eros, and violence (McCarty 1991). These three tags allow for a number of insights into the book. For example, the details of how a character is named in the *Metamorphosis* are one indication of that character's status. Naming by relationships (especially parents and offspring), variations in name and title, and naming by role all help to distinguish minor from major characters. McCarty points to Medusa as an interesting case in point. By these criteria, she is a minor character as a woman, but becomes a major character as a monster.

For readers interested in these three attributes, the TACT tagset of the *Metamorphosis* suggests that the markup will potentially be helpful – although the details of the implementation of the tagset are also going to be a significant factor. For readers interested in other features of the book, such as for instance the role of music, the TACT markup may not prove to be as helpful.

Another possible limitation in the tagset is in the case where a tag does exist, but it has been constrained in some way that makes it less useful than it might otherwise have been. This situation could occur in the definition of a fixed set of attributes on a tag, where the attributes do not include the choice of interest to a particular reader. A hypothetical example might be the case of a tag that marks people's names and includes an attribute for professional qualifications, using fixed values that designate university degrees of various kinds. If the user is searching for someone with an occupational designation based on membership in a professional organization, this list is not going to be helpful, because the attribute's value list does not contain the right class of choices.

Rich Prospect on the Tagset: Connections

It is fairly straightforward to envision the possible identification of new connections among meaningful representations of content items. In these cases, what the perceiver is doing is identifying similarity, perhaps through some form of content or through proximity or both, as when two people are associated through being born in the same year. This initial connection can potentially lead to further inquiry, as in examining the course their lives took – what schools did they attend and what did they study; what careers did they pursue and were the highlights of the one career in any way connected to the other, and so on. The identification of connections is a relatively simple kind of synthetic reasoning, which can be used to develop further insights.

In the case of a rich-prospect display of a tagset, the possibility of the user making useful connections among tags will depend in part on the characteristics of the display. If the tagset can be organized in various forms that have some intrinsic usefulness, then the chances will increase that the user will be able to use the display to identify possible connections.

One simple means of organizing tags is alphabetically by meaningful representation. If the tags are represented, for example, by synonyms or phrases in English, then organizing the display alphabetically may help the user to identify tags that have similar meanings.

An alternative strategy might be to organize the tags explicitly by higher-level semantic categories, creating a structure of nodes with portions of the tagset clustered nearby. It may similarly be possible to create interaction histories in cases where a user has grouped the tags into sets that are meaningful for some particular task. To be of optimum use for subsequent users, this kind of interaction history may require some explanatory text from the user who serves as designer.

Rich Prospect on the Tagset: Trends

The identification of connections among tags may allow the user to begin to draw patterns of significance in the tagset, which may or may not equate to patterns of significance in the actual tagging of the collection or in the contents of the collection independent of how the tags have been applied. In a similar vein, the identification of trends in the tagset may help the user to postulate the existence of related trends, either in the application of the tags or in the actual content or both.

Trends might be identified through the distribution of topic areas in the tagset as well as through the elaboration of the individual tags. A trend may exist wherever the tagset gives evidence of having received special attention from the designers. For example, in the TACT tagset described for

Ovid's *Metamorphosis*, the variations on the <name> tag, which can include a wide range of possible terms that would not normally be considered names in the strict sense, suggests that names in the *Metamorphosis* were a preoccupation of the designers of the tagset. As it turns out, the indication of the significance of names suggested by the elaborations of the <names> tag was actually borne out in the tagging practices of the project, with the result that names form a significant subset of the tagged text (McCarty 1992).

A rich-prospect interface showing some form of the *Metamorphosis* tagset may therefore cue a reader to the significance that names have in the book. Whether the reader were already familiar with the text or were coming to it for the first time, this insight into how the designers of the tags thought about names might suggest a useful line of further investigation, which would be facilitated by the tagging of the document.

Rich Prospect on the Tagset: Anomalies

Anomalous tags are those which fall outside the general trend of the rest of the tagset. They might be tags that are comparatively independent of the rest of the hierarchy, or they might be tags that occur at multiple locations within the tagset. What is important about them is that they stand out in some way from the rest, and therefore may indicate something anomalous about the collection that the designers considered important enough to capture with its own tag.

Alternatively, of course, the anomaly might turn out to be a tag that was never widely used when the time came to actually perform the tagging. Only subsequent analysis of the tagged collection, either through searching on the anomalous tag or through examining a rich-prospect display of tagged elements will determine the extent to which an anomalous tag was implemented.

One way for the interface designer to help protect the user against disappointments in these cases is to ensure that the representation of the tagset includes only those tags that were actually used in the collection, or were used to a significant degree.

Rich Prospect on the Tagset: Reminders

Just as a meaningful representation of the contents of the collection can cue a perceiver to the presence of items that might otherwise have been forgotten, so the provision of rich-prospect forms of display of the tagset can suggest ways of examining the contents that might not otherwise have come to mind.

For example, in a collection of military history, a tagset might be defined that contains tags for indicating different kinds of military resources. Someone using such a collection might be interested to notice that the tagset for the First World War contained a tag for marking instances of <horses>, which still had a significant role to play in that war, despite what were then fairly recent developments in more automated forms of military transport. Such a reminder could lead to subsequent investigation of the various ways in which horses were used in the different military actions, either as draft animals or in cavalry, or in some other capacity indicated by the taggers.

Rich Prospect on the Tagset: Reassurance

Having a meaningful representation of content items can help instill confidence in the user that the collection is the right one to be looking at. A meaningful representation of the tagset may also provide a sense of reassurance to the user, if for example the collection of tags and their associated attributes are the kinds of information the user is hoping to be able to find in the collection.

For example, in the case of a hypothetical collection of geographical information, if there were tags to indicate items such as city names, populations, major industries, and so on, then a user who is interested in looking at the collection for information about the activities of the human occupants may be reassured by the definition of the tags. On the other hand, for a user interested primarily in indigenous wildlife, this tagset might suggest that the collection is not going to be particularly helpful.

There are several limitations to this sense of reassurance. First, it should be pointed out that since the tagset is not the same as the implementation of the tags in a given collection, the reassurance is of necessity conditional on how the tagset has actually been applied. Similarly, if the definition of the tags themselves is open to interpretation by the person looking at the tagset, it is possible that the presuppositions of that person will not correspond to those of the people who designed the tagset, with the result that there may occur some misunderstandings as to how a particular tag was intended.

As a hypothetical example, a collection of secondary material on a literary topic might contain a tag called <title>, intended to identify the titles of published books. If the person looking at the tagset interprets the <title> tag as potentially applying to article titles, then the reassurance that the tagset has a provision for marking that kind of information is going to prove unfounded.

Rich Prospect on the Tagset: Reduced Helplessness

In looking at unfamiliar material in a new context, it is possible for the user to become disoriented or confused, especially in cases where the affordances of the interface are very limited. A rich-prospect interface that gives some representation of the tagset can potentially help to forestall this kind of difficulty, by providing the user with some sense of the way the designers of the tagset understood the material.

The understanding of the designers is communicated indirectly through the tags that have been included and the attributes and values that have been defined for them. Since other considerations (such as the resources available for the tagging effort) may have played a role in the definition of the tags, this understanding may only be partial. However, insofar as it does represent the bounds of what may actually have been carried out in the tagging, the tagset remains a potentially significant artifact in its own right.

Rich Prospect on the Tagging

There are three distinct kinds of prospect possible in a collection that has been interpretively tagged:

- prospect on the contents of the collection
- prospect on the tagset
- prospect on the tagging

The third category consists in some respects of a union of the previous two, but since the implementation of the tagset in a given collection is subject to the interpretations of the taggers, it is not tenable to suppose that knowing about the contents and about the tagset is the same thing as knowing about the way the tags have been applied.

For example, suppose there were a collection dealing with medical information, with a tag for <drugname>. It is possible that the purpose of the tag was such that not every drug named in the collection is tagged. In fact, <drugname> could very well have been defined for some highly specific use only – for instance, as the tag to mark the proper chemical name of a drug the first time it is defined in conjunction with its chemical formula. The <drugname> tag would therefore never be applied where the drug in question is identified by the trade name or manufacturer's name.

Providing rich prospect on the tagging is a problem that is more complex than the provision of rich prospect on either the tagset or the contents of a collection, since for every document there are going to be multiple tags. The prospect therefore needs to display a meaningful representation of

smaller pieces of a document, rather than a meaningful display where each element represents an entire document. Depending on the number of tags in each document, the problem of the number of elements in the rich-prospect form of display can expand by as much as three orders of magnitude.

The implication is therefore that the prospect display of the actual tagging of a collection may contain too many elements, even for a relatively small collection. In such cases, however, there is the possibility of subsetting the display, either by showing prospect on the tags in one document or only a few documents at a time, or else by showing prospect on a particular tag as it has been used across documents.

Whichever form of display is chosen, prospect on the tagging of a collection has the possibility to provide the user with a sense of how the tagset was applied. Some of the significant features specifically related to the implementation of the tagset in a collection are:

- tagged content
- tagging choices
- tag frequency
- tag density
- tag distribution
- tagging consistency

Rich Prospect on the Tagging: Tagged Content

One of the most important features of a tagged collection is the individual quality of the items that have been marked. The user's expenditure of resources in time and energy in examining a collection is rewarded if the designers have facilitated the discovery of significant information by implementing a useful set of tags and attributes. The usefulness of the tags and their implementation needs to be understood within the context of a particular user with a given task, but it is definitely possible to have markup in a collection that is going to prove to be of little value to anyone, because it does not represent a significant investment of tagger intelligence in the collection.

An example of an insignificant tag is the descriptive-level HTML tag `<p>`, which marks the ends of paragraphs. Paragraph breaks are important because they subdivide continuous text to facilitate reading. In some forms of writing they also often serve to indicate a change of topic. But since every paragraph in a document needs to be marked with a terminal `<p>`, there is little value in the tag for the purposes of either interpreting or accessing the collection.

A hypothetical example of a tag at the difference pole might be one intended to identify the thesis statement in each document in a collection of research articles. The <thesis> tag would be applied only once for each document in the collection, and would identify a very significant piece of text. Other examples for the same collection might be tags intended to mark sentences that contain a <conclusion> or <recommendation>.

Rich Prospect on the Tagging: Tagging Choices

A tagger using an interpretive tagset must constantly be making non-trivial decisions in terms of what is important enough to tag and what should be left. If the choice is made to tag everything that might possibly be included in the tag, the usefulness of the collection may be compromised, since the significant occurrences may be lost in a sea of insignificant ones.

The problem is similar to the one that has been perennially faced by book indexers, who are responsible for identifying the material in a book that might reasonably serve as access points for people looking for information on a particular topic. If indexing were a simple matter of marking every occurrence of a particular keyword, then it could easily be automated. However, on the contrary, indexing involves constant value judgments of the content, in order to decide which of the occurrences of a particular keyword are in a context significant enough to warrant drawing the reader's attention to the material. If an index expands to the point where a single entry has dozens of instances, it begins to lose its usefulness.

The tagging in a document collection can be similarly complex, involving the judgments of the taggers as to which words deserve to be tagged and which should be left unmarked. The ultimate basis of the decision is most likely to be the value of drawing the reader's attention to the material, just as it is for the indexers. However, a tagged document can contain more instances of a particular tag than a printed index can reasonably sustain, since the ability of the system to find and present the material is so much more efficient in the digital case than in the print format.

Rich Prospect on the Tagging: Tag Frequency

Assuming that the tagging has been done in such a way that the tagger has discriminated between significant information and similar information that is not as significant, another important feature of the tagging of a collection is therefore the frequency with which a tag has been deployed, either in a single document or across multiple documents. Tag frequency has the potential to serve as an index of

what the tagger found most important about a given document or set of documents, given the larger range of what topics are actually covered.

One means of measuring tag frequency might be in relation to other tags. For example, out of every 100 tags in a given document, knowing how many were a particular tag might be an indication of the attention given in that document to the kind of information the tag was designed to identify, in direct numeric comparison with the other kinds of information the document contains. However, it should be pointed out that numeric comparison is not necessarily the most important indicator from a pragmatic perspective, since it is possible that the topic in question is covered very well in a short scope, which may have required only a single tag to mark it, or that key information is all that is required by the user, in which case multiple instances of examples and elaboration would not be helpful for that user. The user in these cases is interested in finding out whether the key information is present, rather than whether the topic area has been covered in depth. However, in some circumstances, the frequency of reference may serve as an indication that a particular document or cluster of documents deserves further scrutiny.

Rich Prospect on the Tagging: Tag Density

Just as tag frequency in a document may prove to be an index to the kinds of content that predominate, so the density of the tags in a particular document or section of a document might indicate an area of special interest. For example, in a chapter of 5,000 words, if the first 3,000 words only contain a dozen tags but the last 2,000 words contain hundreds of tags, either there is a problem with the tagging, or else the latter portion is where the most interesting material can be found.

On the other hand, a low tag density may also be an interesting indicator, if it serves to draw the user's attention to an anomalous document that does not fit into the framework that has been established by the definition of the tagset. For some users – particularly those who are not necessarily in full agreement with the presuppositions of the people who defined the tagset – this kind of anomalous document may prove an interesting point to begin looking more closely at the collection.

Rich Prospect on the Tagging: Tag Distribution

Within a given document or across multiple documents in a collection, some form of prospect on the tagging might help to indicate how the concepts are distributed. Certain documents may have a heavy concentration of one kind of tag, while other documents have another kind predominating. These

variations in tag distribution may help the reader to understand which documents in the collection are worth investigating further.

A similar process can sometimes occur with web search engines, when multiple search results point to different locations within a single site. Searching on the keyword “visualization,” for instance, inevitably produces multiple pages from the University of Maryland, where Schneiderman, Ahlberg, and their colleagues have been working on visualization for more than two decades. This density of documents on a given topic at a single site is one indication of the quantity of work being done and reported on that site. Providing a visual form to facilitate access to this kind of information is one of the goals of the Kartoo search engine, which provides a type of entity-relationship diagram that relates search keywords to prominent sites (Figure 3.01).

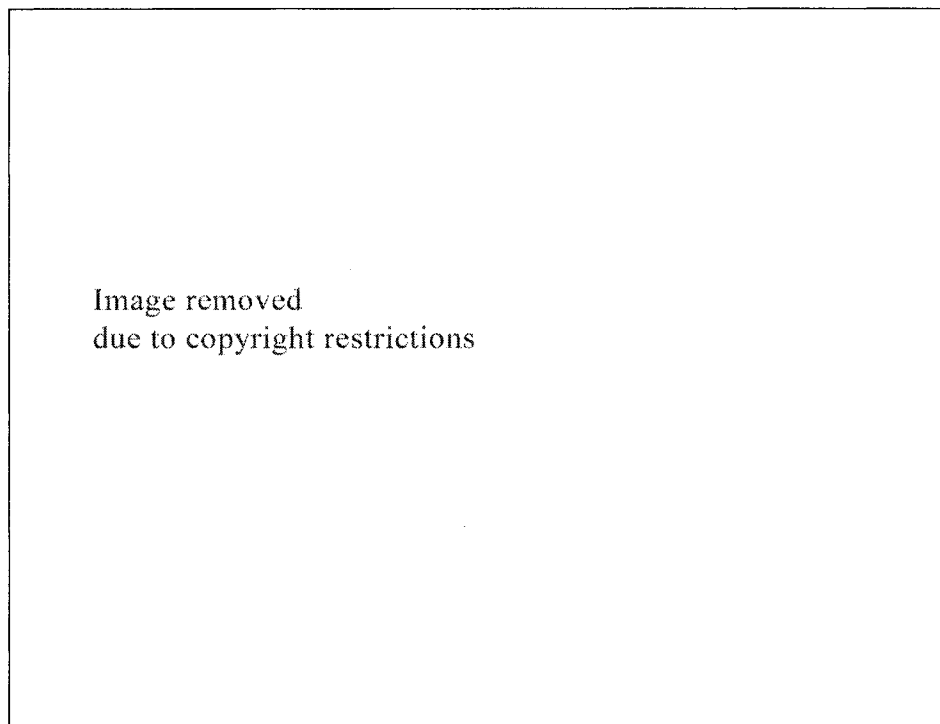


Figure 3.01 Kartoo selects web sites with large numbers of documents on a particular topic, then displays those sites as central nodes in a network of search results. Sites that fall beneath the threshold are not part of the display, which means that Kartoo usually privileges institutional sites over individual ones. Clicking on a node brings up further information about the documents it contains.

Rich Prospect on the Tagging: Tagging Consistency

The choice of whether or not a given piece of text should be tagged is often difficult, although experience and domain expertise can help to make the decisions more consistent and accurate. With that in mind, if there is some accessible means for the reader to compare actual document content with the tagging of the document, it may be possible to evaluate the degree to which the tagging in the collection can be trusted to be consistent. Consistency is significant because it represents the extent to which the tagging of the document will actually be able to draw the reader's attention to the most relevant material.

A simple form of this evaluation will occur naturally whenever a reader encounters two instances of the same keyword or phrase, where one is tagged and the other is not. If it is readily evident that the untagged version is not significant enough to be tagged, then the reader's sense that the quality of the tagging is adequate will likely be reinforced.

An example of this kind of situation might be where the smallest item a system returns in response to a query on a tag is an entire paragraph. If a text string occurs twice in that paragraph, it should only be necessary to tag either the first or the most important occurrence, since otherwise the system will draw the reader's attention to the same paragraph twice.

Help Systems and Tutorials

Many of the individual affordances of rich-prospect interfaces that show some representation of the content, tagset, or tagging can also be provided using other strategies than prospect. One way to provide the user with a variety of insights into the collection and how its designers understood it is to write these insights explicitly into an appropriate set of texts and attach them as a help system.

Another method is to create a tutorial that the user can work through in order to gain some understanding of the system.

Help systems and tutorials are useful tools, but they have the disadvantage that they are not part of the situated activity of using a collection of documents. They are an additional task that has to be undertaken. The default activity for most people is to begin using the system, then turn to tutorials and help systems only when the situated activity fails (Suchman 1987, p. 54). For many people, it is preferable to use a system at less than optimum levels rather than exert the additional effort of gaining a better understanding of how the system might be used.

This reluctance is not unreasonable, since the less-obvious features of a system might be logically interpreted as a form of communication from the designers that the features are less obvious because they are less important. There are further levels of potential difficulty in that the design of the help system or tutorial may be less than optimal, resulting in increased user frustration at a time when frustration is already serving as a motivator to seek help. Finally, the possible return on investment is unknown, since it may turn out that the system includes a wide range of useful features that were not immediately obvious, but can be learned from the tutorial or help function – or it may not.

In spite of these limitations, help systems and tutorials are an important component, if for no other reason than they provide some safety net: in the case of help systems in situations where the user is unable to continue without support, and in the case of tutorials when the user does not know how to begin. However, insofar as the system is capable of providing an interface and corresponding set of tools that give the user the desired opportunities for both gaining information about the collection and for acting with respect to that information, the need for the safety net is reduced.

Visual Culture

In addition to complexity, another issue that should be addressed in any discussion of rich-prospect interfaces is the role of visual culture and positioning. Within the visual communication design community, the concept of visual culture has been growing over the past half century, in reaction against the modernist principle that design could be judged with reference to a universal ideal. For the modernist designers, there was a scale of design quality that was applicable independent of the circumstances of deployment and audience. Modernism in general tended to emphasize simplicity and geometry, in opposition to ornamentation and organic forms. Coupled with a growing industrial capacity to mass-produce identical items using materials such as plastic and aluminum, modernist design standards became widely accepted and implemented throughout the western world.

In response to this universal standard, proponents of the concept of visual positioning emphasize that in order to be most effective, design needs to acknowledge that there is a complex universe of visual cultures, and that to position a design solution for one visual culture may result in it being inappropriate for alternative visual cultures. Multiple design solutions are therefore necessary, in dependence on the number of different visual cultures being addressed.

The marketing community has an analogous concern in creating communications collateral for distribution in support of a particular product. Recognizing that the market as a whole actually

consists of many smaller markets, one strategy among marketers has been to consider the economy as a kind of complex ecology, with a wide variety of market niches available. Each niche can be understood, not as an unsatisfied functional requirement, but rather as a position within the ecology. These positions consist of the mental awareness of a particular idea, slogan, concept, or identity in the mind of the consumer. In this ecological paradigm, marketing consists of efforts to capture and maintain a position in the consumer's mind, and therefore within some portion of consumer culture.

Visual positioning resembles market positioning in that both efforts require the designers or communicators to understand as much as possible how the interlocutors in the communicational exchange are likely to interpret the material they encounter. Both kinds of positioning therefore necessitate studying the actual user – in the one case in terms of visual environment, awareness, and preferences, and in the other in terms of commodity awareness and interpretation. The two kinds of positioning differ from each primarily in terms of scope and intention. For the marketer, the goal is to generate product awareness and consumer demand for a particular marketable item or service. For the visual communication designer, the goal is not so much to capture and hold a niche, but rather to increase the likelihood of successful communication by adopting an appropriate visual vocabulary, since visual vocabulary, like lexical vocabulary, is a kind of subtext to the actual communication, which in some cases may communicate more effectively than the contents of the message do. The subtext created by choice of vocabulary often communicates information about the pragmatics of the exchange, suggesting possible identities and agendas rather than the semantics, or what is actually being said.

For people working outside the modernist paradigm, an effort to adopt the correct visual position for a given visual culture may seem like a necessity. However, developing multiple solutions instead of a single solution requires a significant commitment to the principle of visual culture, because the design and development efforts must be increased, if not as a direct multiple of the number of solutions, then at least in a relative proportion. Complexities of access are also introduced, insofar as the system must somehow differentiate members of one user community from those of another, not only in terms of functional requirements, but also in terms of correct positioning of visual culture. However, to maintain that one size fits all, in interfaces as in anything else, is to risk alienating some segment of the potential user community.

Prospect as a Visual Position

In the case of rich-prospect displays, the wealth of information is going to be an element of any visual language. To provide the user with a search screen on the one hand and a browsing display that may be showing thousands of elements on the other is to position the two interfaces at some remove, regardless of other similarities that might exist in terms of font choice, colours, visual elements, placement or orientation of elements, and so on.

In addition, there are some further elements which have a high likelihood of occurrence in rich-prospect interfaces, simply because the large number of items on display will necessitate strategies to make the whole manageable, while at the same time providing the user with a sense of prospect. Strategies of this kind might include various methods of making use of the third virtual dimension, perhaps through scaling or superimposition or visual occlusion of portions of some elements by others. Methods of collapsing and expanding items are also likely to be useful in rich-prospect displays, as are means of grouping items under user control in order to organize the contents of the display.

However, given that there are going to be certain strategies that are more likely to find implementation in rich-prospect browsing interfaces than in other kinds of interfaces, there is still the issue of whether or not it is appropriate or even possible to provide various visual positions for a single interpretively-tagged collection with a given set of display options, or whether the complicated nature of the information on display is inevitably going to dwarf the significance of the other visual elements to the point that the contribution they make to the overall visual language is minimal. A particular interface might show, for instance, several thousand elements that are the meaningful representations of documents in the collection, and that display may be coupled with another display of hundreds of items that represent the tagset. From the point of view of the user, that mass of information may constitute the primary visual element of the site, although there may still be sufficient room for variation in the way the items are displayed, the way the tools are indicated, and the way the other elements of the site are created and deployed, that the designer will be able to override the visual effect of all that information to visually position rich-prospect browsing interfaces for various audiences.

Adoption of New Technology

In addition to the difficulties presented by the need for visual positioning, there are also variations across different user communities in terms of the experience of different mechanisms of user interaction. Among the people interested in accessing an interpretively-tagged text collection, it is likely that there will be some who have existing skills involving various interface options. Even among people who have such skills, however, there are going to be some who are more adaptable or more willing to investigate possible new affordances, and others who will tend to stick more exclusively with what they already know.

It is not yet possible to determine with any precision the likelihood of a particular technology finding an appropriate niche and becoming adopted by a group of users: there are many cultural, interpersonal, and individual factors involved. However, some researchers have singled out various relevant aspects of the problem. For instance, with respect to the diffusion of technology, Rogers (1983) proposes five factors that contribute to the rate of adoption of innovations:

- relative advantage over previous technologies
- compatibility with existing values, experiences, and needs
- simplicity of the new technology to understand and use
- trialability in which the user can engage in harmless limited experiments
- observability of the new technology as it is being used by others (Rogers 1983, p. 207)

In terms of these five factors, rich-prospect forms of interface may fare quite well. Their relative advantages include the list of new affordances provided by both the meaningful representations of collection items, as well as those provided by rich prospect on the tagset and on the tagging. Their compatibility with existing values should be fairly strong, provided that the arguments about the genetic predisposition for human beings to obtain prospect are valid, and that these arguments in turn are applicable in the electronic realm. Rich prospect is relatively simple to understand, although it may also turn out to be comparatively intimidating. It is trialable in the sense that manipulations of the display will not tend to result in information disappearing, but rather in information being subject to reorganization. Finally, it is a form of interface that is highly observable, provided that the potential user is in a position to see other people working with it.

Constraints

As Norman (1990, p. 82ff) points out, one means of helping the user to carry out a task is to make it difficult or impossible to carry out the wrong task by mistake. Constraints in the analog world are conditions on the affordances of an object that restrict the user to certain kinds of actions, as when a door can be pulled open but not pushed. The appropriate way for a designer to signal constraints is by providing the user with an unambiguous interface. In the case of the door that cannot be pushed, the handle should be shaped in such a way as to suggest the action of pulling. Conversely, for a door that only affords pushing, the interface should be a flat plate or a breakout bar that either does not allow or does not immediately suggest pulling.

In the digital world, constraints are often the rule rather than the exception: the problem is that because the digital environment is deliberately constructed by designers, the affordances rather than the constraints are what need to be developed. The constraints come all too naturally, as anyone who has used a system with a command-line interface and its unforgiving syntactic requirements can testify.

However, even in a graphical user interface, the invisible constraints can be a source of considerable frustration, rather than a guide to what behaviours are appropriate. For example, a standard search box on a retrieval interface allows the user to enter a search string. In some cases there may even be syntactic restrictions, based on the features of the retrieval software, which are invisible to the user. A typical example might be the need to include the word “and” between terms that are supposed to both be present in any target document, and a simple space for words where either one or the other or both might be present. This constraint on syntax, however, is not necessarily simple for the designer to indicate. In some cases it may be possible to provide each of the syntactic elements separately, so that the user has a perceptual cue that they exist. But knowing they exist is not the same as knowing how to use them: complex logic in searches requires some training and experience before it makes sense, and in any case the interface can quickly become intimidating (Figure 3.02).

Simple search

author

title

Search with tags

author

title

keywords

tag

attribute

attribute value

Search with tags in English

author

title

keywords

standard

tag

attribute

attribute value

advanced

Advanced search with Booleans

author

and or not title

and or not keywords

standard

and or not tag

and or not attribute

and or not attribute value

advanced

Figure 3.02 By providing constraints on the user’s input in accordance with the required syntax for retrieval, a search interface quickly becomes visually and conceptually complicated. This set of interfaces shows the progressive addition of features relating to searches on contents, tags, and attributes in an interpretively-tagged system.

Another strategy is to simply not indicate that the constraint on retrieval syntax exists, on the assumption that the user will figure out that the default is a logical “or” and that if what the system should really be providing is a logical “and,” there may be some syntax available to provide that. The user could then either experiment with the retrieval system or refer to a help function.

In the case of a rich-prospect form of interface involving the contents, tags, attributes, attribute values, and other kinds of information displayed as various forms of meaningful representations, one actual constraint on the system is that it will contain the information being represented. Another constraint is that it will not contain information that is not represented. If the design of the browsing system is such that the markup portions of the display can be used as picklists for creating queries, then the constraint of a kind of tag being present or not present in the display is an unambiguous representation of a constraint actually present in the system – a constraint that may otherwise be invisible to the user.

Similarly, if the display indicates the structural form of the tagset, then a constraint is indicated to the user as to where tags can occur within the nested structure of definitions. To take a hypothetical example, in a collection of materials related to movies, if the tagset has been defined according to country of origin, with genre tags potentially occurring within countries, then the user might find that there are tags for <American> and within <American> there are tags for <gangster> films, but that the system does not include any tag to indicate <Canadian> <gangster> films. Using a search mechanism based on this markup system, it is therefore not possible to search for a Canadian gangster film: the tagset does not contain any such designation. Gangster films may have been made in Canada – the constraint is not on the film industry, but rather on the markup system of this particular hypothetical collection. In fact, the constraint is not even necessarily on the contents of the collection: it is possible that there are Canadian gangster films present. But the lack of an appropriate tag in the markup system means that they will have been tagged as something else, and the user may have to search further for an appropriate tag that might have been used instead.

Conversely, of course, the presence of the tag in the tagset does not necessarily indicate that there are instances of the tag having been used anywhere in the collection. It may turn out that there are no American gangster films present either, but that the designers of the tagset, knowing that such films exist, wanted to make provision for their future acquisition. Another possibility is that the tagset was designed to be as complete as possible and every genre is available for every nation,

regardless of the actual genres produced by the film industry in those nations and regardless of the actual contents of the collection. What therefore might be the best practice in this case is for the meaningful representation of the tagset to either eliminate unused tags altogether from the display, or else to indicate them in some way as inactive (perhaps by graying them out, which is the standard visual syntax for inactive menu items).

Natural Mappings

Another common design principle that can be applied to the visualization of complex information is natural mapping, where the form of the display is intended to be congruent to the form of the information being displayed. An example of natural mapping in the analog world is when the controls for the burners on a stove are placed on a horizontal surface in a pattern that matches the position of the burners on the stovetop. An alternative natural mapping would place the burner control in proximity to the burner. An unnatural mapping would put the burner controls on a vertical surface in a straight line, although the burners themselves are in a rectangular pattern on a horizontal surface. Instructions as to which control matches which burner are essential in cases of unnatural mapping, and redundant when the mapping is natural (Norman 1990, pp. 75-8).

In the case of a digital interface, the concept of natural mapping is not as clear, since the information being displayed does not have a physical form or configuration to begin with. However, there are some instances in which the principle does hold. For example, Norman (1993, pp. 69-72) points out that some information is naturally associated with a spectrum or a scale. Survey results using Likert scales would be one such case. Another case might be in election results, where the numbers of votes cast in each province or electoral district might be of interest. To display this kind of scaled information, Norman suggests that substitutive displays are unnatural, while additive displays are natural. Using different colours or different kinds of patterns, for instance, to represent different values on a scale would be to use an unnatural substitutive display. To show the same data using different shades of gray, with the lightest shade corresponding to the lowest value and the darkest shade corresponding to the highest, would be to use a natural mapping, since there is a visual correspondence between the additive kind of data in the scale and the additive kind of shading.

There are several possible ways in which the concept of natural mappings might apply in the case of rich-prospect displays of files, tagsets, and tagging. For instance, a display might begin by showing some meaningful representations of files. This display could then be combined with a display

that indicates which tags in the tagset can be found in which documents. By placing the two kinds of information in proximity, the relationship between the files and the tags used in them would form a kind of natural mapping.

Alternatively, a display that begins with the tagset might be combined with another display that lists meaningful representations of the documents where each tag might be found. The visual results of the two strategies would be quite different, but in both cases if the display clearly relates the individual tags or other elements in the tagset with the documents where they can be found, then the display would express a natural mapping between the tagset elements and the documents. On the other hand, if the system does not allow the user to have an overview that includes both kinds of information simultaneously, or in which the information is not visually connected, then there is no natural mapping of this kind present in that particular display.

A similar instance of natural mapping might occur in the display of the attributes of the tags, if these attributes (and in some cases their values) are shown in relation to the tags to which they refer. An extra level of mapping could be added by also providing the contents of the tags, or a meaningful representation of the contents, again shown in relation to the tags and attributes.

INFORMATION VISUALIZATION

...we don't design systems merely to replace human work, but to enhance human capabilities to do productive work (Karat 1997)

Given that rich-prospect displays of the tagset and the tagging in a given collection may provide the user with helpful information, the problem still remains of how the designer might go about dealing with the issues raised by attempting to display so much information in a way that is manageable. By keeping the principle in mind that the goal of the interface is not just to make the information available, but also to do so in order to provide the user with new affordances, the designer may be able to generate solutions that enhance the work environment rather than unnecessarily complicating it.

Meaningful Representation of the Tagset

Although it may be beneficial to provide the user with a method of obtaining prospect on the tagset, the means of displaying what may be a large and complicated amount of information in a useful form is a non-trivial problem. In terms of the approximate scale of the tagset, it is not likely that even the

most complex ones would include thousands of tags – more typically the tags of even a fairly sophisticated project would number only in the hundreds. Given that many of the tags may contain several attributes, the size of the total set of both tags and attributes could be increased into a thousand or two. If the pre-defined attribute lists are also included, there may be another marked increase in size, but not so much as an order of magnitude, since it is unlikely that the average attribute would have a value list ten items long. So even the most complicated tagset, including attributes and attribute value lists, may still only comprise two or three thousand items, and many will number fewer than a thousand.

There is, however, another complexity to be considered in addition to the numbers of elements. A tagset is also hierarchical, with some tags occurring inside other tags. Tags may also be either mandatory or optional. Attributes similarly may be mandatory or optional. Finally, it is also possible to define tags recursively, so that a tag is allowed to contain another instance of the tag that contains it. Recursive definition allows the indefinite nesting of tags: the actual depth of tagging in collections where recursion is possible can only be determined by examining the tags that have been implemented. All of these details about the tagset are potentially significant in that they will have formed constraints on the way the tagging of the collection has been carried out. Any rich-prospect display of a tagset may therefore need to indicate not only the tags, attributes, and attribute values, but also the relationships among them.

For users who are familiar with SGML or XML, it may be possible to simply provide the opportunity to look directly at the definitions of the tags. Technically sophisticated users of this kind would then be able to examine the definitions and come to understand what had been done. However, it is not a reasonable assumption that most users of interpretively-tagged text collections are going to be proficient in reading markup grammars. For the users who are not, it is necessary to find means of making the tagset accessible not as an SGML or XML document *per se*, but rather through some form of meaningful representation.

With respect to the meaningful representations of content items for rich-prospect interfaces, knowledge of the user is important, because meaning does not exist in a vacuum. A similar necessity holds in the case of meaningful representations of the tagset, where it will be necessary to acquire some knowledge of the user in order to make the presentation meaningful.

There are some conventions related to the syntax of tagging languages which may be sufficiently widespread to be useful. Anyone who has worked in HTML will recognize angular braces as the delimiters of a tag. One solution is therefore to leverage off this existing knowledge and present each item of the tagset within angular braces. This strategy seems the least likely to result in confusion on the part of the user as to whether the display is showing a set of tags or some other kind of information.

Given that the tagset is going to contain at most a few thousand items, and that each of the items is relatively small – that is, not an entire file in itself, but at most a phrase in length – then it may be possible that the user can obtain prospect on the tagset by examining some form of display that contains representations of each of the items. However, the details of how this display should be constructed will need to be carefully considered if they are to provide the user with some useful understanding (that is, some sense of prospect), rather than simply with a view of everything in the tagset. As Gershon et al. (1998) put it, “many interesting classes of information have no natural and obvious physical representation. A key research problem is to discover new visual metaphors...”

Vocabulary

There is also a potential difficulty with the terms that have been used or created to define the tags, which may very well be idiosyncratic to the developers of the tagset, and therefore not necessarily accessible to anyone else without translation. The simplest solution for the developer is to make this the user’s problem, and to provide some form of help system or glossary for people who are interested in understanding the tags. The problem with this strategy is that it may leave the user either confused or frustrated, and in cases where the tags seem to be interpretable by the user but have actually been defined in such a way that an obvious interpretation is incorrect, the display may actually end up misleading the user altogether.

Another solution is therefore to provide the user not with a direct representation of the tagset itself, but rather with insight into a representation of the tagset created by substituting meaningful synonyms for the actual tag names. The difficulty to be faced in this strategy is the need of providing adequate synonyms for tag names that may in fact represent a combination of fairly complicated definitions and the related decision-making processes for best practices in the tagging project. For some tags it may therefore be necessary to provide phrases rather than single words, which increases the size and complexity of the display.

The difficulty of achieving an appropriate vocabulary is well documented in several domains, including in particular the fields of library science and information retrieval. The vocabulary problem is twofold: on the one hand, different people use different words to express identical or at least similar ideas, and on the other hand, different people use the same words to express dissimilar ideas. The vocabulary problem is not amenable to solution by mandate, since what is involved are various people and their habits of thought, and those are not easy to change. It is a simpler approach to try to adapt the systems so that in spite of the user's preference of terms, the software is able to interpret correctly the intention behind the terms and produce the desired result. Strategies include everything from latent semantic indexing, where the software attempts to cluster terms by their meanings, to thesauri of various kinds, such as those based on WordNet, to markup technologies, where use of terminology in the document is regularized, not in the texts themselves, but in an additional layer of information provided through the manual addition of tags. It is ironic that the display of the tagset, however, can reintroduce the vocabulary problem into the equation – albeit this time at a meta-level.

One solution that has been shown to be useful in the realm of information retrieval is the use of relevance feedback, where the person using the retrieval interface is given the opportunity to specify which items and in some cases which search terms from a first iteration resulted in appropriate documents. Given the ability of the system to store these relevance choices and reuse them in subsequent queries, it is possible to improve performance on retrieval. If the system has mechanisms in place to store the relevance feedback as a kind of interaction history, it may also be possible to use the terms for suggestion on similar subsequent queries by other users.

Whether or not relevance feedback could be applied in the area of rich-prospect display of tagsets depends in part on the ways in which the tags are displayed in the first place, as well as on how they are used for the purposes of formulating queries. A further consideration is going to be the extent to which the user is able to indicate success or failure of the query. A related question is whether, given a display of a tagset that has been modified according to feedback from one user during a particular query, that modified display will prove useful both for query construction by subsequent users, and also as a display for purposes other than the formulation of queries.

Picklists

The obvious candidate technologies for displaying any list of items are the pull-down menu and the picklist (where the latter is distinguished from the former by the addition of a vertical scrollbar). The

advantage of these solutions is that they are standard parts of existing GUI development systems (and hence easy to implement, with an enculturated user base); and they are easy to understand, consisting of a list of mutually-exclusive terms that are amenable to selection. However, they are not necessarily an optimum solution for several reasons. First of all, in their default configuration they do not allow the user to select multiple discontinuous items. Secondly, and for the purposes of prospect most importantly, they do not allow the user to see more than a few dozen items at a time, and even those items are not typically structured in any complex fashion.

One means of alleviating the constraint on the number of items that can be shown simultaneously, and hence providing some form of prospect on a picklist, is through the use of a fisheye lens, which magnifies the current items and those in proximity, but leaves other items at a very small scale. Coupled with some additional features such as an indexing system, fisheye menus can provide prospect on lists of a length similar to those likely to be encountered in displaying a tagset. In cases where it may also be useful, however, to provide the user with some sense of the structure of the tagset, a list will not be adequate.

Structure

SGML tools such as DynaText display the tagset as a tree diagram that can extend at 10 or 12 points for metres of virtual screen space, making prospect on a regular monitor simply impossible to obtain. The purpose of this kind of display is to show the details of the hierarchy, rather than to provide prospect on the entire tagset. If the user is interested in tracing the connecting lines between a higher-order tag and the tags it may contain, the view can be collapsed and expanded accordingly. But to see a meaningful display of the entire tagset is impossible for any DTD containing more than a few dozen tags.

Any representation of the tagset designed to provide prospect should make optimum use of available screen space, while still presenting the user with sufficient details into its structure. The tension between the need to avoid wasting screen space while at the same time presenting hierarchical details may require some inventiveness on the part of the designer.

One possible strategy is to make use of the virtual third dimension, so that parts or all of the display can recede into the distance or be brought forward for inspection of details. If the material is presented in such a way that different parts of it become resolvable at different virtual distances, then

the user may have the opportunity to understand some of the structural qualities of the tagset before having to deal with the intricacies of the identities of individual tags and their associated information.

Another possibility is to have the tagset displayed using the close proximity of elements, where higher-level tags occur in physical relation to the tags they contain, and those tags are further represented in visual proximity to the tags they enclose, with possibly some variations in font size or colour used to help distinguish one level of tag from another. The important feature of this kind of display is that there should be some means of blocking together those tags that have been defined at a similar level, rather than presenting them as lists, without suggesting that the blocking is actually a kind of nesting.

Another possibility is to in a sense restructure the tagset for display, by omitting altogether those levels of tag that are not of particular semantic value. For example, if the tagset contains higher-level tags that indicate structural elements such as paragraphs, and these paragraph tags can contain several possible optional tags that indicate the topic or content of the paragraphs in some way, then those optional tags with some semantic significance should be considered for display, but the paragraph tags might be omitted. Sacrificing the strict display of the hierarchy in order to emphasize the semantically significant portions would naturally have no effect on the retrieval mechanism, which may still, for example, return the paragraph-level tags in response to a query. But for the user who is attempting to gain a sense of prospect on the tagset, displaying only those tags which are potentially meaningful may reduce the cognitive load of first needing to sort the syntactic tags from the semantic ones, before going through the list of semantic ones for those of particular interest. Displaying only the semantically significant tags may also serve to reduce the complexity of the hierarchy, by collapsing it into individual segments that can be treated as more or less autonomous units.

Attributes and Attribute Values

An additional complexity arises with respect to those tags that have associated attributes and attribute values. In these cases there is the possibility of attempting to find a means of displaying all of the information in some form, so that the user is able to obtain a prospect on the full definition of the tagset.

One possible strategy for showing both the tags and their attributes is to duplicate the tags as many times as there are attributes, so that the display of the tagset becomes in some respects a display of the attribute list, with the appropriate tags serving as prefixes to the attributes. An even finer level

of granularity is to extend this strategy one layer deeper, and provide duplicates of both the tags and the attributes as prefixes to the attribute values.

In visual terms, the redundancy of the prefixes in the solutions above is an area amenable to improvement, since visual redundancy can quickly become visual noise. Redundant information does have the potential advantage of being a source of reassurance: people know what they are looking at, because the entire section of attribute values has the same attribute and tag name attached to each entry. However, it is necessary to balance that potential reassurance against other methods that may make more efficient use of the limited screen space.

One possible visual solution that removes redundancy is to use text layering, setting attribute values near larger text that names the attributes, which is in turn placed on text that is larger still and names the tags themselves. The result can be visually complex, but is relatively compact (Figure 3.03).

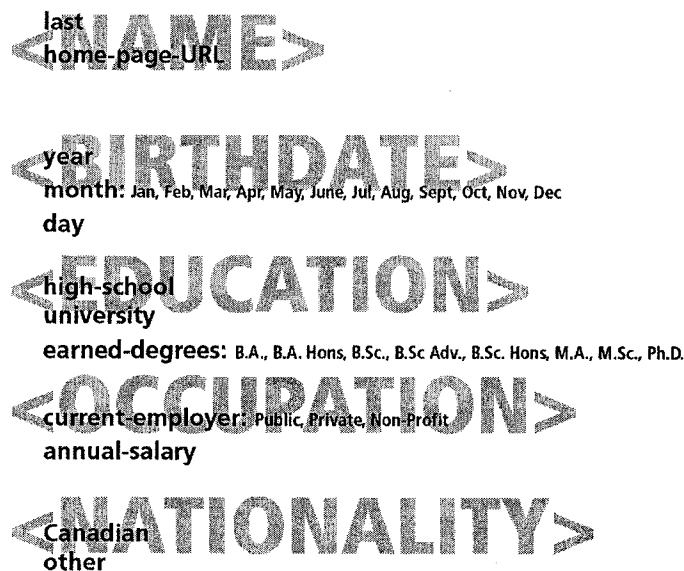


Figure 3.03 Text layering can be used to keep the display of the tagset and its associated information in a compact form. Here the tags are shown as the largest text items, with the attributes of each tag superimposed in a smaller font, and the various pre-defined attribute values are smaller still and shown next to the attributes.

A slightly less complicated version involves clustering the attribute values, attributes, and tags in groups that are associated by proximity but distinguished from each other by some other visual quality, such as font size or font colour (Figure 3.04).

NAME**last****home-page-URL****BIRTHDATE****year****month**

Jan, Feb, Mar, Apr, May, June, Jul, Aug, Sept, Oct, Nov, Dec

day**EDUCATION****high-school****university****earned-degrees**

B.A., B.A. Hons, B.Sc., B.Sc Adv., B.Sc. Hons, M.A., M.Sc., Ph.D.

OCCUPATION**current-employer**

Public, Private, Non-Profit

annual-salary**NATIONALITY****Canadian****other****dual**

Figure 3.04 A variation on text layering that is not as visually complex but is less compact is to cluster the tags, attributes, and attribute values by proximity.

Another option is to display the attributes or the values not as a form of prospect display, but rather as a feature of the tags that can be expressed by clicking on the tag name or by rolling over the area. The advantage of this approach is that it preserves screen space; the disadvantage is that it becomes difficult to get a sense of how attributes and fixed lists of attribute values have been deployed across the tagset.

In terms of the significance of this kind of information, there are several possibilities based on how a tagset may have been defined. It may be the case that there are attributes attached to many of the tags; it may be the case that very few of the tags have attributes. In either case the attributes might be relatively trivial, consisting, for example, of unique identifier numbers for the tags; or they might

be relatively important, listing significant taxonomies that have been used to structure the data according to some interpretive framework. It seems reasonable to suggest that in collections where the attributes or their value lists are potentially quite significant to the user, the case for attempting to provide some form of prospect on the attributes is also strongest.

Meaningful Representation of the Tagging

The structure of verbal language, however, offers a limited capacity to convey information. In the long run, this has limited our capacity to understand serious problems of a physical or social nature, due to the lack of ability of the verbal language to promote the perception of contexts, complexity, and simultaneity, in other words, due to its lack of ability to promote thinking in terms of ecologies of information. (Frascara 2001).

The difficulties presented by the decision to provide a meaningful display of the tagging of a collection are also not trivial, because each document might contain hundreds or even thousands of tags, and the collection may contain many thousands of documents. The potential size of the display therefore far surpasses rich prospect on either the tagset or the collection, which shows either representations of the tagset and its related information, or else representations of entire files.

For the purposes of providing prospect on a collection, it may be sufficient to limit the use of rich-prospect displays to those showing the content and the tagset, and allow access to the actual tagging only through retrieval tools, rather than through another rich-prospect interface. However, insofar as it may be possible to provide a rich-prospect interface to the tagging itself, the potential exists to provide the user with several advantages similar to those associated with rich prospect on the collection and the tagset, but at a finer level of granularity. This increased level of detail may be significant for some users, even at the cost of a larger or more-complex visual presentation.

Another possible reason for providing prospect on the tagging is that there are aspects of the information inherent in the tagging that cannot be extrapolated from knowledge of the contents and the tagset – namely, to what extent and with what degree of consistency have the tagging practices indicated for the collection actually been carried out?

Tagging includes all of the information together: the content, the tags, the attributes, and the values of the attributes. To show all of this material at one time would be to show not only the entire contents of the collection, but also to include the markup that would otherwise be invisible to the reader. Simply showing all the information at once would not, however, meet the brief in terms of

creating an impression of rich prospect for the user, since the basis of rich-prospect interfaces is that they involve some shorthand meaningful representation that stands in place of the actual material. The user needs to be able to stand back, as it were, and obtain an overview, without having to sort through all the details, while at the same time having some confidence that the details are readily available. In the best case, the user would also have the ability to shift perspectives in order to see the tagging in different ways. For example, in a collection with different levels of tagging, it might be useful to be able to emphasize one level at a time, even though all the levels might be present. The internal glossing tags, for instance, might be made to stand out from the rest of the tagging at one point, with the same display shifting under user control to emphasize hermeneutic tags at some other time.

Given that the prospect needs to provide some kind of shorthand representation of the various items, and that there are at least three kinds of items to be shown, the first requirement for any system attempting to provide prospect on all this information is that the different kinds of data need to be visually distinguished from each other. Items can be distinguished by colour, size, font family, font style, location, grouping, additional marks (such as the angular brackets for tags), and so on. The best combination of these cues in a given interface will provide unambiguous information as to the kind of each of the elements on display. The choice of cues should also be related to the visual language of the interface as a whole, so that the user is aware that the rich-prospect display is related to the rest of the site. The tools for manipulating the data should also be constructed in such a manner as to be consistent as possible across the project.

Rich prospect on the tagging of the collection also involves the problem that the amount of text contained by any given tag can vary widely. Some tags, such as the Dublin Core meta-tags, might surround the entire document, and serve as a means of identifying the whole. An example of this kind of tag is the HTML `<body>` tag, which is used to indicate the main content of an HTML page. Other tags may include any level of granularity from major chapters, paragraphs, sentences, and phrases, down to the smallest tags that may contain single words, or even may contain no content at all, in the case of tags that have been placed in the text to mark a place and serve as a reference point, or perhaps to add an attribute value.

In order for the user to obtain an overview of the tagging of a document, it may therefore be necessary to include some means of indicating how much text each tag contains. Given all of the

possible kinds of information in a document that are related to the tagging, the list of components of a rich-prospect form of display on the tagging becomes quite long. It includes:

- the tag
- the document containing the tag
- higher-level tags containing the tag
- any attributes on the tag
- any attribute values associated with the attributes
- the contents of the tag
- the size of the contents of the tag
- the tag's location in the document
- the location of the tag in the tagset hierarchy (to precisely identify the kind of tag)
- the frequency of occurrence of this tag in the document and collection
- the overall frequency of occurrence of all tags (for comparison purposes)

Access to the Tags While Browsing

With so much potential information to display, one strategy is to subset the information. If the system is able to identify the tag the user is currently interested in, the entire display can be structured around providing prospect on that tag. In this case, the tag would serve as a frame for the entire display, while the actual contents of the display would show the meaningful representations of documents, associated (most likely by proximity) to the attributes on the tag. Additional information might include numbers indicating frequency of the tag in each document (for the attribute serving as the cluster center), and ratios showing tag count against total numbers of tags in the document.

As a hypothetical example, if the user were interested in the tag used to mark intertexts in a literary collection, the system might show a display of only those documents that contain the intertext tag. If the attribute of the tag included the title of the intertext, the documents might be grouped by that attribute, so that documents would cluster around the title of the common document they all reference.

If a drill-down function is also provided, the user would be able to modify the form of the display, so that instead of showing the meaningful representations of entire documents, the prospect would switch to showing some representation of the contents of the tags.

Another possibility is that the information is not subsetting, but that the display takes on the nature of a complex ecology of information, where the various elements are all present, but different optical weight is given to various features depending on the nature of the task. For example, the user might be interested in examining the comparative occurrence and placement of two kinds of tags within a particular group of documents. This information could be shown as a set of superimposed texts, with the two tags shown as insertion points in the relevant documents, which are expanded for the purpose, while the other documents and tags remain as meaningful representations that have been reduced in size, shifted to the background, or screened down in intensity (Figure 3.05).



Figure 3.05

An example of an illustration showing a complex ecology of information, this rich-prospect display shows 1000 document titles, a tagset of 75 tags, and two documents opened to show comparative placement of the two tags currently of interest to the user. In this case, the tags are reproduced near the open document thumbnails, and small geometric shapes are used to show rough placement in the document. Ideally, each of the items in the display would be an object that could be manipulated by the user, and the various panes could be repositioned to allow access the material underneath.

The provision of a function that allows the user to shift perspective relates to an aspect of prospect described by Appleton in terms of the difference between primary and secondary vantage points. A primary vantage point is the one currently in use by the perceiver; secondary vantage points, on the other hand, represent further options for observation that increase the value of the prospect:

All direct prospects are views actually achieved by the observer from his position of observation, which we can call his primary vantage-point. A very important role, however, is discharged by other potential vantage-points. It must be remembered that the satisfaction of seeing is only a part of the satisfaction of achieving an advantageous position within one's habitat, and clearly the belief that one's field of vision can be further extended if one moves to another observation-point will accentuate the sensation of environmental advantage (Appleton 1975, p. 89).

An example of a secondary vantage point in nature is the horizon line, which symbolically suggests that there is more to be seen if only the perspective could be changed to allow increased depth of view. In the digital environment of the interpretively-tagged text collection, depending on the design and on the users, the existence of additional tools for manipulating the interface by shifting the perspective on various elements may be analogous to the horizon line in that they similarly represent changes in observation point on the prospect display.

Access to the Tags While Reading

In addition to a rich-prospect display of the collection, it may also be useful to consider methods for providing the user with some form of prospect within the contents of an individual document. Although the ability to scan over the contents is obviously an affordance of the viewing or editing software that opens the file for use, the screen size restricts the view to a few hundred lines at most, with the result that the user has to either scroll through the material or page through it, and an overview of the material is not necessarily available. In the case of an interpretively-tagged collection, there is also the question of how best to show the tags to the user in a way that will be useful in providing a sense of what additional information has been incorporated through the tagging.

One strategy for providing some prospect within a tagged document might be to provide a parallel-column approach, where the information about the tagging is shown as a sidebar to the contents of the document. Adobe Acrobat has a display that shows thumbnails of the pages in a vertical navigation bar down the side of the page (Figure 3.06). It may be possible to adapt that kind of display to indicate tag placement and kind, in cases where the density of the tags is not prohibitive. One of the visual problems to be overcome is that tags surround text. It is therefore necessary to show not just the insertion point where a tag begins, but also the tag's end point. Another problem is that tags may nest within each other, which makes it necessary to show not only tags around text, but also tags around other tags. A document with even a

moderate degree of tagging is not amenable to reading with all the tags expanded, so simple text display of the tagged text is not an adequate solution, and more sophisticated visualization methods must be found.

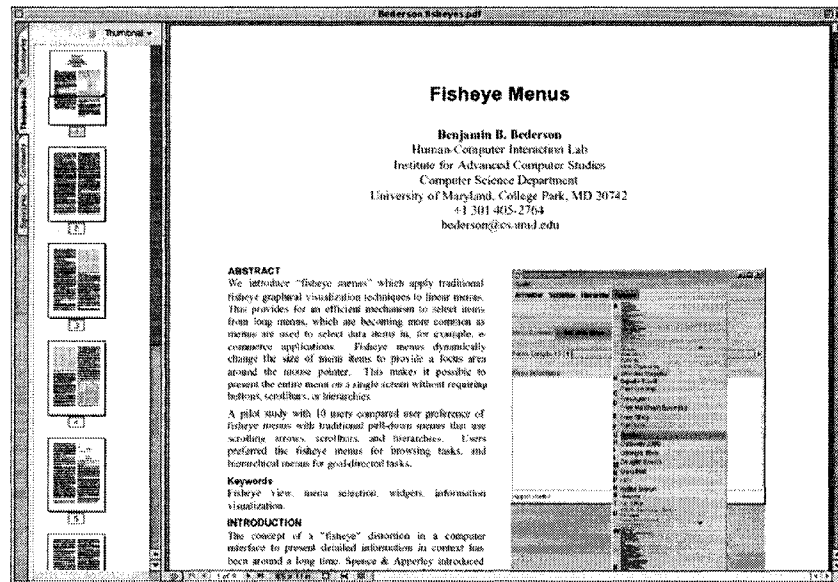


Figure 3.06 Adobe Acrobat shows a thumbnail display of the current document. This sidebar can serve both as a form of prospect on the contents and a means of navigation. A slightly larger form might be modifiable to show tag type and placement. Adobe product screen shot reprinted with permission from Adobe Systems Incorporated.

Another possibility is to have the tag tree for the document shown in a sidebar, with rollovers on the actual text causing highlights to appear in the hierarchy. By running the cursor across the written text, the user would be able to see where various tags occur in the document, without the distraction caused by having the actual tags appear in the text being read. The problem with this strategy is that most readers do not use the mouse to follow the text while reading; it introduces a new and somewhat demanding task for the reader to perform. The highlighted text in the hierarchy display would also pose a potential distraction that might interfere with the reader's immersion in the reading, rather than facilitate an understanding of the document.

There is also the question of why the reader might be interested in knowing about the tags once the document is being displayed. If the primary function of the tagging is to help the reader find the right text, then once the text is found, the tagging becomes irrelevant. However, this is where the interpretive extent of the tagging becomes a significant issue. If additional information has actually

been inserted in the document by the taggers, then the reader may find it useful to have the ability to access that information while reading.

If the purpose of examining the tags is to identify cases where additional interesting information has been encoded, one possibility is to provide the user only with those tags that are at the highest levels of the tag taxonomy: the external glosses and hermeneutic tags. If, on the other hand, the purpose of examining the tags is not restricted to looking for additional information, but relates perhaps to an interest in how the information has been standardized for the purposes of searching by the computer, then the user may also be interested in looking at the internal glossing tags. The kind of information that would be gained in this way would be of potential assistance in formulating queries: for example, it might lead to an expansion or consolidation of the user's search vocabulary. Yet another possibility is that the user may want to see the remaining portions of the tagging system, such as the meta-tags or even the descriptive markup, in order to understand in the first case how the documents have been defined as documents and in the second how the various visible kinds of formatting have been associated by the system with the structural definitions of the different parts.

If the system has been designed in such a way that these various levels of tags can be accessed as taxonomic groups by the user, then the ability to switch on the display of a particular level of tag can be implemented as a tool for the user. On the other hand, if these kinds of internal groupings are not available, the system may be able to parse the definitions in order to create such groups as a form of higher-order index. This parsing stage could also serve the dual function of filtering instances of repeated tags within the current unit of information. As a hypothetical example, the reader might be interested in looking at individual paragraphs, but the taggers marked a particular piece of text not just the first time it occurred in a paragraph, but whenever it occurred in individual sentences. Unless the system accommodates this disjunction between tagging and retrieval, there may therefore be cases where the same paragraph is retrieved multiple times, because the tag of interest is repeated more than once.

Finally, the possibility also exists of allowing the user to examine the tagset at a rich-prospect level and create groups of tags for subsequent display when actually reading one of the documents, rather than relying on automatically-generated groupings from the system. This method has the advantage of giving the user additional control over what kind of tag information is displayed,

although it carries the cost to the user of having to create the groups of tags that will be of interest. Such a feature would therefore likely be of most value to users who are quite experienced with the collection and tagset.

However, once this kind of information has been developed by one user, there is always the possibility of having the system store it in some form and make it accessible to subsequent users as a form of interaction history. The existence of such stored interactions is another kind of information that the system can contain which may be of interest to the user. The catalog of available interaction histories is therefore another candidate kind of information for rich-prospect display.

CHAPTER 4: PROSPECT-BASED INTERFACES AND THE ORLANDO PROJECT

In the early stages of a new technology, people tend to think that its purpose is merely to replace and improve on something they already know. The promise of the new is thought to be quantitative: the new thing will do the old job faster, more efficiently, and more cheaply.... Tools, however, are perceptual agents. A new tool is not just a bigger lever and a more secure fulcrum, rather a new way of conceptualizing the world.... (McCarty 1991).

Rich-prospect interfaces provide their users with opportunities for action that are not available from search engines that do not provide prospect. For users looking for a well-defined target document, search interfaces with no prospect may be a good solution. However, for users looking for an understanding of a collection and how the various components comprising it interact, rich-prospect interfaces have the potential to be a better solution.

There is, however, a significant difference between the two kinds of interfaces in terms of their design. Whereas a search interface is primarily a front end to an algorithm, and can therefore be as simple as a single search box (Figure 4.01), a rich-prospect interface is a front end to the entire contents of a collection. In the case of tagged collections, the interface may also contain representations of the tagset, including the attributes and their values, and the tagging, which can be used to subdivide each document into myriad pieces. The design process for a rich-prospect interface therefore involves identifying strategies to help the user make sense of a large amount of structured information, as well as providing tools for working with the information in ways that will prove useful.

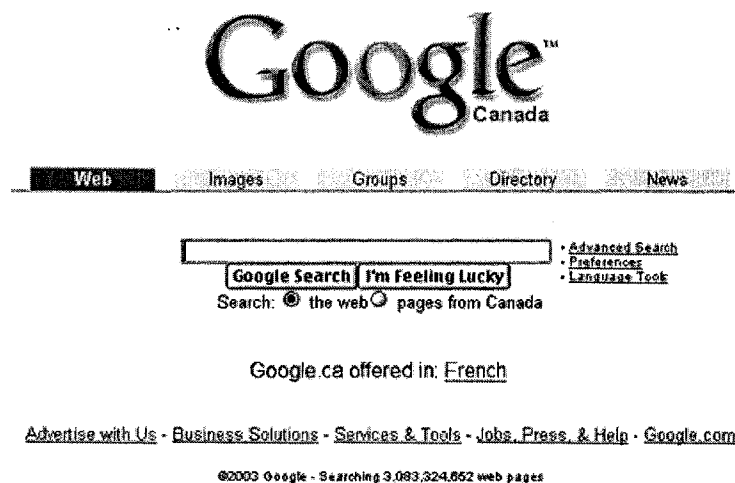


Figure 4.01 The default Google web search engine is an example of a retrieval interface with no prospect.

There is a sense in which a rich-prospect interface mediates between the designers of the collection and the collection's readers. The readers bring their goals in accessing the collection, domain knowledge, prior experience, visual preferences, and presuppositions of various kinds. The designers, on the other hand, have an understanding of the material, how it has been organized and tagged, and how it can serve to forward whatever scholarly agendas may be intrinsic to the process of creation. The developers of tagged text collections may be fully aware that they have been engaged in an act of interpretation, as is the case in the Orlando Project:

In theorizing and pushing the limits of descriptive markup, we recognize that the world we wish to label cannot be viewed objectively; rather, we are presenting that world as contextualized in a specific time and place, a world seen through the critical lens of the projects' researchers (Fisher, 1998).

From the interface designer's perspective, part of the job is therefore to find methods of making the critical presuppositions of the collection's developers explicitly available to the users of the collection, in terms that the users are able to interpret. There may also be practical limitations relating to what can be delivered in a reliable and timely manner using the current technology. For a brief discussion of these constraints on the first release of the Orlando Project, see Appendix A: Technical Considerations.

CHAPTER OUTLINE

The following chapter examines some of the concepts relating to rich-prospect interfaces in relation to the contents and tagging of the Orlando collection. The first part of the chapter gives a brief overview of the Orlando Project, and summarizes the characteristics of the collection that make it a good candidate for rich-prospect designs. The second half of the chapter looks in particular at two of the key components of Orlando – author names, and <ChronStructs>, or pieces of text associated with dates – which have been defined in the tagsets and implemented in the tagging in such a way as to necessitate careful analysis of the details and how they influence options for providing prospect on the collection.

It is in the details of the tags, attributes, and attribute values for an actual collection, such as Orlando, that the principles of rich-prospect interface design come into contact with the kinds of constraints and conditions that need to be addressed as an intrinsic part of the design process. Complicated as they may be, these details serve to test, validate, and refine the concepts in a way that is otherwise impossible.

THE COMPLEXITIES OF ORLANDO

The Orlando Project is an integrated history of women's writing in the British Isles. Developed by collaborative research teams working at the University of Alberta and the University of Guelph, Orlando currently contains documents on several hundred individual authors, as well as more than ten thousand historical events. The documents are divided into five categories, which correspond to the five tagsets developed by the project: biography, writing, events, topics, and bibliography.

Each of the authors represented in the collection has two documents: a biography, which contains information about critical events and activities throughout the author's life; and a writing document, which is essentially a mini-biography specifically dealing with the author's writing and publication activities, including summary discussions of major works.

In addition to the paired documents associated with each author, there are at present a dozen short topic documents, and also more than ten thousand small text items, called "events," that contain information intended primarily to provide historical context. Events are typically quite short, consisting of a paragraph or two and a date, while writing and biography documents are considerably longer, although usually not exceeding 4,000 words each. The events documents are understood as an essential part of Orlando's agenda of placing women's writing within a framework of relevant detail:

For women's writing such informative contexts have been lacking, and one result is that this writing is often dehistoricized and seen in essentialist gendered terms which impoverish response (Grundy et al. 2000).

Major themes in the design of the Orlando collection include: people, texts, chronologies, organizations, places, culture, and politics, as well as literary, social, and family connections. The tagsets support each of these thematic areas (and more), because the tags have been designed and implemented with the intention of facilitating support for access to documents or sections of documents within these larger umbrella themes.

The tagging on the Orlando project occurred concurrently with the writing of the documents, which is another way in which this project differs from other markup projects in humanities computing. The usual process is for a project to digitize or collect existing documents, whether primary or secondary, then insert markup. By writing and tagging simultaneously, the Orlando authors had the ability to write the material in such a way that it was more amenable to tagging.

Orlando and Prospect

From the perspective of investigating the value of prospect as a design strategy, the Orlando collection (and in particular its biocritical material) has several features that make it an excellent test case. These features include:

- the size of the collection
- the size of the documents
- the homogeneity of the documents
- the interpretive level of tagging
- the characteristics of the users

Orlando contains biocritical documents numbering in the hundreds or low thousands, rather than in the tens of thousands or hundreds of thousands. The size of the collection is, therefore, within the right range. The number of documents is important because, if there are too few documents, the potential of rich-prospect strategies is not likely to be fully realizable; there is likely not going to be much use, for example, in having a technology that allows the user to group items in a collection of a few dozen documents. On the other hand, if there are too many documents, the chances of the user finding some meaningful way to interact with them as an entirety are reduced. There seems little chance that a rich-prospect display will ever prove useful for the entire index of web pages in Google, for example, since the current index contains more than 3 billion items.

Documents are of an appropriate size if they are longer than a few paragraphs, but still under twenty pages. Length of document is important because, if the documents are too short, the return on investment for the user is reduced, while, if the document is too long, the chances are greater that it will not be possible to represent it with some single meaningful representation.

A degree of homogeneity in document content means that the documents can be readily represented by a common form – in this case, the names of the authors who are receiving biocritical treatment by the project. However, the Orlando Project does not perfectly meet this criterion, since representation by author's name is actually appropriate only for the writing and biography documents, and not for the events documents. So there is, in fact, a need to have two separate strategies for representing documents. The presence of the events documents is an essential part of the collection, however, because it fulfills one of the agendas of the collection developers: namely, to situate the history of women's writing and lives within an explicit historical context.

On the other hand, the interpretive level of tagging in Orlando establishes it as an exemplary project – one of the most ambitious SGML-tagged collections of text undertaken to date in the humanities. It is not unique in its use of an elaborate tagset with a range of tag types, but it is a good example of the kind of extensive and elaborate encoding that creates many opportunities for the interface designer. The Orlando Project provides a test bed that can be used to examine the detailed issues of conveying complicated information (the collection contents) that interacts in various ways with information of a different kind, that is at least equally complicated (the tags, attributes, and attribute values).

Finally, the audience of the collection is also fairly well-defined, so that research into the design and application of rich-prospect forms of interface can occur within the context of use of the collection by academics who are interested in the history of women's writing.

ORLANDO TAGSETS

If the user is to have access to the Orlando tagsets, several design issues need to be addressed. The first two issues relate to the general questions “why?” and “how?” while the last two issues deal with questions of how the various components might interact with each other. The issues are:

- What might the user gain by having prospect on the Orlando tagsets?
- How might prospect on the Orlando tagset be provided?
- Should the presentation of the tagsets keep them distinct?
- How could tagset prospect interact with collection prospect?

What might the user gain by having prospect on the Orlando tagsets?

The original purpose for tagging a document with an SGML-defined tagset is twofold: so that the markup can serve as a support tool for a retrieval algorithm; and so that the markup can be used as a means of facilitating formatting. To make the tagsets available to the users in any form is, therefore, to re-purpose the tagsets to a new use or set of uses, which include making visible some of the organizing principles of the collection, as well as the ways in which the designers of the tagsets understood the material being created. One potential benefit for the user in gaining such an understanding is that it may serve as a form of education in the content domain.

Depending on the tools that are provided, the user may also benefit from the opportunity to use the tagset in formulating queries. In cases where the nature of the query is congruent with the tagset, knowing which tags are available can potentially help the user to refine the query to make the

best use of the tagged material. If the system also provides the user with feedback as to the actual query being formulated (most likely in Structured Query Language), it may also be possible for sophisticated users to learn to formulate or modify queries using the appropriate syntax. Additional features such as a query formulation wizard, which walks the user in steps through the process, can assist in the process of growing familiar with query formulation through examining the tagset.

Knowledge of the tags, attributes, and attribute values that were used can also suggest, not just more accurate forms of previous queries, but also new queries that might otherwise not come to mind. In this context, the tags, attributes, and attribute values become visible cues to the kinds of information available in the collection.

How might prospect on the Orlando tagset be provided?

Most tagsets are hierarchical, and the Orlando tagsets are no exception. The position of a tag within the hierarchy may in some cases be significant enough that the user would benefit from knowing that position. Some of the possible strategies for visually structuring a tagset display were discussed in the previous chapter. Options for creating a display that indicates hierarchical position include the use of layering, clustering, and tree diagrams. However, because tree diagrams can quickly extend past the boundaries of a normal screen, they should be visually optimized where possible in order to save space and allow for greater prospect on the whole.

In applying these concepts to Orlando in particular, a further difficulty may arise from the fact that the Orlando tag definitions are, in some cases, recursive. That is, there are tags of type A which can contain subtags of type B, which in turn can contain subtags of type A. Complete display of the tagsets would therefore require some indication of the places where recursion is possible.

Display of the tagging, on the other hand, would need to indicate the places where recursion has been applied. Further research is necessary to determine how often this has occurred in the Orlando documents, and also to understand whether the user would benefit from a display that indicates the precise location of each tag within the recursive hierarchy, or whether the display of sections of the hierarchy as a kind of local nesting of tags would be sufficiently informative.

In cases where hierarchy is less significant to the user, some of the options discussed below for providing prospect on the documents in the collection could also be applied to the problem of

providing prospect on the tagset (see “Displaying Author Names”). In order of increasing prospect, these options include:

- picklists
- microtext picklists
- walls of text
- panoramas

As with prospect on the contents, prospect on the tagset can be structured in ways that are more or less complex. It may also be useful in some cases to provide people with tools that can be used to manipulate the display of the tagset, either through sorting or subsetting or grouping the tags according to criteria relevant to the task at hand.

Should the presentation of the tagsets keep them distinct?

The design process for tagset displays in a rich-prospect interface should give attention to the unique aspects of the particular tagset. One question that arises in the case of Orlando is whether or not it is useful to differentiate between the tagsets used by the project. The display might, for instance, make the point that the different document types have each been tagged using a different tagset. On the other hand, the display might merge the tagsets into one larger meaningful representation, since there is considerable overlap, especially between the tagsets for the writing and biography documents (the tagsets for events and topics are considerably smaller, and the tagging correspondingly less complex). The various solutions will each result in a different user understanding of the collection, which implies that the interface designer should address this issue with the developers of the collection.

The first solution has the advantage of keeping the use of the tagsets in constructing searching criteria clearly in line with the design of the collection. For example, if the user is looking specifically at Biography documents, then if the display does not distinguish between the tagsets, the situation may arise where searches are being performed on Biography documents using tags that do not occur in them. Maintaining an alignment between the form of the interface and the form of the collection is useful in helping the reader come to an understanding of the collection. However, in this case, the replication of some of the tags across the tagsets results in a degree of redundancy which might prove irritating or confusing to some people.

A more flexible solution might, therefore, be to have the display of the tagsets change automatically to accommodate the various kinds of searches. For someone interested in looking at the

Biography documents, a toggle on the display used to constrain the search could also trigger a change in the display of the tagsets so that only the Biography tags are visible. If the person were interested in all kinds of documents, then the complete amalgamated tagset could be shown. The advantage of this strategy is that the options available conform to the current environment. The disadvantage is that the appearance and disappearance of interface options can be disorienting, and also tends to restrict the reader from easily coming to an understanding of the larger system and the tools it contains.

There is a third form of display that circumvents the problem of confronting the user with an interface where a choice in one area generates unexpected changes in other areas. Such a display would be one that constrains the search options and the representation of the tagset so that rather than invisibly linking them, the system would show them in parallel. In computing terms, the difference is between a modal solution, where the current activity limits the range of possible actions, and a modeless one, where the user is not constrained by the environment of the current activity. The classic modal situation is the dialog box that appears and suspends access to anything else on the screen, requiring a response from the user before any other activity can proceed. In general, the computing community has recognized that modal situations should be avoided wherever possible, although modality in menu choices remains quite a common design feature.

How could tagset prospect interact with collection prospect?

If both the collection's documents and the collection's tagsets were made available to the user through some form of rich-prospect display, the resulting material would be too complicated to fit, at a legible font size, onto a standard monitor screen.

However, if the material were to be presented as microtext, then it would be possible to provide simultaneous display of both the tagsets and the collection. If the two forms of display were allowed to interact, then the user would be able to associate documents in the collection with the tags they contain. If the attributes and attribute values were also made available, then the degree of complexity would increase, but so would the potential functionality. Finally, if the user can also provide search criteria, or make use of other features to sort, subset, and group the material, and have those features interact with all the other forms of display, the result is a rich-prospect interface on the entire collection, augmented by some tools that make use of the new opportunities for action provided by having some meaningful representation available. This kind of interface has the potential to allow

people insight into the collection, and also to provide them with a variety of new affordances that assist in working with general areas of research interest.

The details of how best to display each kind of material, how to visually represent the interactions between the different kinds, and how to provide tools for generating the interactions will all be significant decisions. It is going to be necessary to investigate the user community in order to establish the extent to which these prospect-related strategies can work, and also to determine which visual formats are most conducive to people learning and using the system.

One possible strategy is to provide the various displays as separate windows or dialog boxes, which can be opened or closed by the user in much the same way that tool palettes are open or closed in programs related to digital imagemaking. Another solution may be to provide the user with a set of wizards that break the process into sequential steps. The various strategies are not mutually exclusive, but can coexist in the same interface.

ORLANDO NAMES

The users' interests will have to be brought into contact with our purposes and intentions, the story we want to tell, the emphases we wish to make, the misconceptions (some of a monumental nature) we wish to redress (Butler 1998).

The primary organizing scheme of the collection is biocritical, with supporting documentation that is historical. The user's focus may involve either of these perspectives, depending on whether the emphasis is on the individual author or the contextual events. It is also possible that the user may shift perspective in the course of a single use of the collection, first examining, for example, the documents relating to a particular author or authors, then following some historical thread from one of the documents out into the larger events collection, spending some time looking into the history, then perhaps returning to the author. The number of possible paths through the collection materials is open to the interests of the individual researcher. The central organizing principle of the collection is nonetheless the biocritical material, since it is in the biography and writing documents that the Orlando project is making its principal contribution to literary scholarship. The events are not primarily intended to be new contributions to history, although some of the ones focusing on women's writing or the other activities of women may very well serve that function. In general, however, the events as they currently stand are designed to be contextualizing material for the biographies and writing documents.

Given the central place of these document pairs in the organization of the project, it is necessary and reasonable to allow the user access to the collection through either a search or a browsing function that gives author names as at least one of the meaningful representations of documents. The biography and writing documents were written with the intention that in many cases they would be treated as reading texts in their entirety: the user would find an author of interest, call up the full texts of the biography and writing documents, and read through them.

Extending the Rich-Prospect Name Display

Finding a particular author by name may turn out, however, to be difficult. First of all, not every author used a single name during their entire writing career, and some authors used many different names. Even the most common cases of authors with pen names can leave the reader in some doubt as to where to look for material: should someone looking for information on George Eliot, for instance, look under Eliot or under Mary Ann Evans?

For purposes of facilitating retrieval, the Orlando developers have included a Standard attribute on the <Name> tag, so that taggers could specify a consistent name throughout the collection. The taggers identify standard names by consulting a Name Authority List for the project, which was developed by the project textbase manager based on the following set of authorities, in order:

- i. The Orlando document archive catalog
- ii. *The Feminist companion to literature in English: women writers from the Middle Ages to the present*: for women writers
- iii. *The Oxford companion to English literature* (5th ed): for male writers
- iv. *Dictionary of National Biography*: for British non-writers (except for those with peerage title)
- v. British Library Catalogue online
- vi. *Everyman's Encyclopedia /Encyclopedia Britannica*
- vii. Library of Congress authority files
- viii. G.E. Cokayne, Complete Peerage
- ix. The volume authors (Clements et al. 2003a)

Standard names do not have to be attached just to pen names, however. They have also been used for identifying people where the actual text only gives an oblique reference. As a hypothetical

example, the text might read “Nancy Mitford’s sister was also a writer” and the <Name> tag on “sister” provides the standard name Jessica Mitford, to distinguish her from the three other Mitford sisters who were not writers.

In addition to the <Name> tag, there is also a <personName> tag, which is used at the beginning of each biocritical document to clearly identify the author under discussion. The definition of the Orlando <personName> tag provides for a number of possible subtags, which are listed in the project glossary as follows:

“PersonName is a Div1 content element. It has the following sub-elements to capture specific names:

- surname
- birthname
- professionalTitle
- indexed
- married
- nickname
- pseudonym
- religious
- royal
- selfConstructed
- styled
- titled” (Clements et al. 2003)

These details of tagging practice become significant for the design of an interface that makes the names visible to the user. If the designer wants the rich-prospect display to include not just a single authoritative name for each person, but rather all of the names used for that person throughout the collection, the list becomes fairly complicated. It might include the contents of the <personName> and <Name> tags, their attributes, and their subtags. The result might, therefore, include a list of non-unique identifiers such as “brother,” “sister,” “mother” and so on.

Avoiding the problem of identifying people by non-unique names is the purpose of using a name authority list in the first place. However, this solution does not accommodate two cases: the situation where a reader is interested in a common name shared by multiple people; and the case where the reader has a particular search target in mind, but has only a vague sense of who the person might

be. As a hypothetical example, the user might be thinking “I would like to find a particular woman writer from the Renaissance. All I can remember is that she was the mother of another woman writer, and they both wrote plays.” In order to allow for this kind of search, it might be useful to allow the user to see the full display of the contents of all name tags, perhaps including a few words on either side of the tag to provide context, as is done in concordances. In any case, the display would include a number of non-unique identifiers.

The choice of whether or not it is appropriate to use a particular set of non-unique identifiers rests on how they have been tagged. There are three possible scenarios: they may prove to be too common in the collection to be useful for differentiating items; they may turn out to be too inconsistently applied to be of any real use; or, they may turn out to be both consistently applied and uncommon enough to be of value.

One class of those identifiers which may prove too common to be useful are those relating to families. For example, many of the women discussed in the collection are someone’s mother or someone’s sister. If the taggers have consistently attached a <Name> tag and standard attribute to uses of the words “mother” and “sister,” there may simply be too many of them for the designation to be helpful in differentiating items (although the identifiers may, in this case, prove useful for grouping items).

In terms of those identifiers which are too inconsistently tagged to be useful, it may turn out to be the case that the Orlando taggers have marked with a <Name> tag some, but not many, of the instances of familial roles. If that were the case, then for the few that have been marked, there are two possible states: the identification might be significant, or it might be trivial. The choice of whether or not to draw on such identifiers as components in the interface name display would therefore need to be determined by looking at the actual implementation of the tagging across all the documents, in order to see if some logical system has been applied in the choice of when a familial role should receive a <Name> tag. Further research is required.

In the final class, the non-unique identifiers may prove both consistently applied and uncommon enough to be of value. For example, here is a fairly typical <Name> tag from the biography document of Henrietta Battier: “<NAME STANDARD=“Russell, William,, Lord”>Lord Russell</NAME>” (Clements et al. 2003). As it happens, Lord William Russell lived in the late 18th century. But he is identified in the text that the reader sees as “Lord Russell,” and William Russell is only one of several Lords Russell that have held the title over the generations. If the name

display includes both the <Standard> attribute value and the contents of the tag, there will be both the unique standard name “Lord William Russell” and the non-unique name “Lord Russell.” If another of the Lords Russell is mentioned in the collection somewhere, then the link specified by “Lord Russell” would need to point to the references to both people. However, being able to identify all the Lords Russell in the collection at one time may be useful to some researchers.

In summary, a display that results from using some massaged form of the text inside the <Name> and <personName> tags would show more entries than there are documents. That is, there would be a many-to-one relationship between document names and documents. This form of display might help facilitate retrieval by people who are unfamiliar with the variations of naming that might apply to someone they are interested in finding. It might also be useful for people who are looking for groups of names that fall into some recognizable class that would otherwise be difficult to identify.

Providing people with some means of switching the display between one-to-one and many-to-one representations would provide both affordances. For example, a user might be interested in finding all the women in the Orlando biocritical materials who held the title “Lady.” For some of these writers, the title may be part of the standard designation. For others, it may appear in the text of a <Name> tag but would not necessarily form part of the standard designation. If the display could be expanded from the form where it shows a one-to-one relationship between document titles and contents to a form where it shows a many-to-one relationship, it might be possible to provide the user with some means to find and group the entries of both kinds. On the other hand, once the reader has identified particular people of interest, switching the display back to one-to-one would reduce the complexity at the point when it is no longer required.

In its optimum form, the system would provide the user with a means of changing the display between its three or four potential forms, with the default display being the one that shows one meaningful representation per document (or, in the case of Orlando, one meaningful representation per document pair). The choices would be:

- show standard author name only (one per document)
- show all possible forms of author names (likely more than one per document)
- show all possible author names and oblique references (perhaps several per document)
- show all possible author names and oblique references in context (concordance style)

A dialog of this kind could be designed to apply either to the entire display or to some pre-selected subset.

Author Names and the Name Authority List

Another factor complicating the use of names in Orlando is that the authors who have received biocritical treatment are not the only people mentioned. There are currently approximately ten names listed in the name authority file for every author with a pair of biography and writing documents. These names include various historical figures, male writers, and the colleagues, relatives, friends, and associates of the women authors.

In some cases the names signify people who have received electronic treatment in some other collection. In other cases there are names of people who are mentioned in more than one document.

Showing the full Name Authority List as an interface object might be useful to readers interested in pursuing links outside the Orlando collection, or to readers looking for relationships between authors who have received biocritical treatment, or to readers interested in relationships between authors and other people mentioned.

However, it is necessary to provide some clear idea to the user that most of the names on the list do not have biocritical documents associated with them. The two options available to the designer are to either mark individual items to indicate their status in the collection, or to group the items according to status. Issues related to marking individual items are discussed below (see “Displaying <ChronStruct: ChronColumn>”).

There are several variants within the idea of grouping. One strategy might be to group names on the list according to the names of the authors in whose documents they occur. An alternative that would be equally interesting would be to provide a cluster of names of biocritical authors around each of the other names, to indicate for a given person where in the collection his or her name is mentioned. Such a list could then be sorted by frequency, and the reader could begin to identify people who are not British women writers who have nonetheless had some significance in the writing or biographical histories of more than one woman writer.

For the designer, it is perhaps not entirely necessary to make a definitive choice between the various options for individually marking items or grouping them, since one of the purposes of a rich-prospect interface is to provide the user with tools that allow restructuring the display in various ways that are helpful to the task at hand. However, it may nonetheless be useful to provide a default solution, in which case one of the criteria for the decision should be to support the most frequent tasks carried out by users of the collection.

Displaying Author Names

There are several possible ways of making the list of authors available to the users of the Orlando collection. Although the differences between some of the methods might appear in theory to be trivial, in application the effect on the perceptions of the reader can be significantly different given even a relatively minor change.

For example, having a display where the font size is not adjustable and the default is slightly too small for a given person's visual acuity can be an excruciating experience. Given the same display with the addition of a facility to modify the font size to something appropriate can remove several obstacles to the use of the site, including:

- the actual mismatch between font size and acuity
- the sense of the user being helpless
- the sense that the helplessness could have been alleviated relatively easily, and hence that the situation is unfair
- the anger that results from being in a situation that is unfair when one feels helpless

Rather than risk introducing the user to experiences of this kind, it would be better to examine each of the alternative strategies, in order to identify details that may contribute to or detract from their use by a particular community of people accessing the collection.

Picklists

One method of displaying names would be to include them in a picklist or menu. These options have the advantage of being familiar to most users of graphical user interfaces. However, given that the eventual list of author names alone could number in the thousands, the length of the menu becomes difficult for the user to manage, and after the first few dozen names it does not provide a very effective sense of prospect.

Microtext Picklists

Microtext, on the other hand, can provide a method of displaying more text at one time. The user can see the list in its entirety, but access it by moving the fisheye lens or other magnification device across the list. Microtext menus have been studied in other contexts and found to be useful, although the inclusion of various ancillary features (such as a temporary "locking" function associated with the list) have not always proven to be distinctly beneficial for some users (Bederson 2000).

Walls of Text

A picklist, whether at a legible font size or as a microtext list, is restricted by vertical placement of items into a column, with the result that it does not take maximum advantage of the potentially available screen real estate. A display that shows the names as a block of text across the entire screen can contain more items, or the same number of items at a larger font size. Sorting the names can help the user to traverse the display more easily (Figure 4.02).

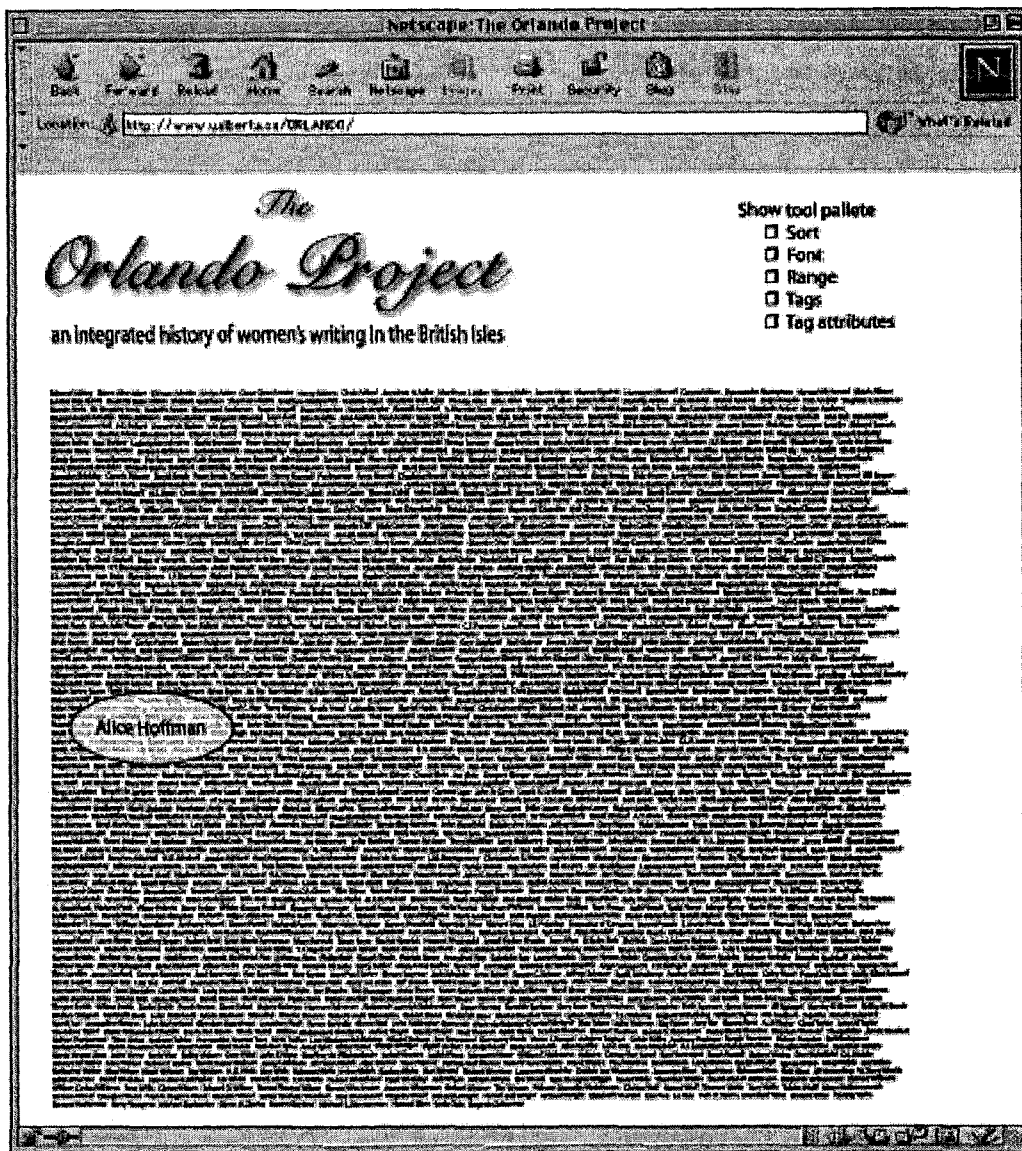


Figure 4.02 This wall of text shows a list of author names from Orlando, in alphabetical order by author last name. A fisheye lens effect is used to allow the entire display to fit at one time on the screen.

As with the picklist, the wall of text can be displayed at a font size under user control. It can alternatively be equipped with a magnifying or fisheye lens to allow the user to examine parts of the display in a selective manner, while keeping the remaining contents static on the screen. It may also be valuable in some situations to organize the names into columns, especially if column headers are provided as a form of visual indexing that can help the user to traverse the display more easily.

Panoramas

While the wall of text is an interesting solution to the problem of displaying hundreds of names, it is not necessarily the best solution for representing thousands of names, since the limited screen size is a factor that needs to be accommodated.

The next logical step is therefore to extend the wall of text into a full panorama, or horizontal strip. Virtual panoramas also have the merit of being analogous to physical panoramas, and therefore associated with prospect on a landscape.

If the ends of the panorama strip are virtually connected, then the panorama can be rotated either to the left or right, and the user is constrained from losing the panorama by scrolling one end of it off the screen. If the panorama is sorted in a way that is immediately obvious (for example, alphabetically – or by date with the dates visible), then the user is less likely to become disoriented. If it also contains some kind of distinguishing feature indicating to the user when the circle has been completed, there is less chance of the user accidentally circling back over the same material.

If the text on the panorama is arranged in columns, then indexical headings can be used to aid navigation. In addition, if the names are arranged in alphabetical order, and if there are enough columns per letter of the alphabet, then large versions of the letters can be added as a visual cue to allow the user to quickly and easily move through the display.

Finally, if the panorama is implemented in an appropriate technology, it may also be amenable to zooming, which will allow the user to see more of the strip at one time (Figure 4.03).

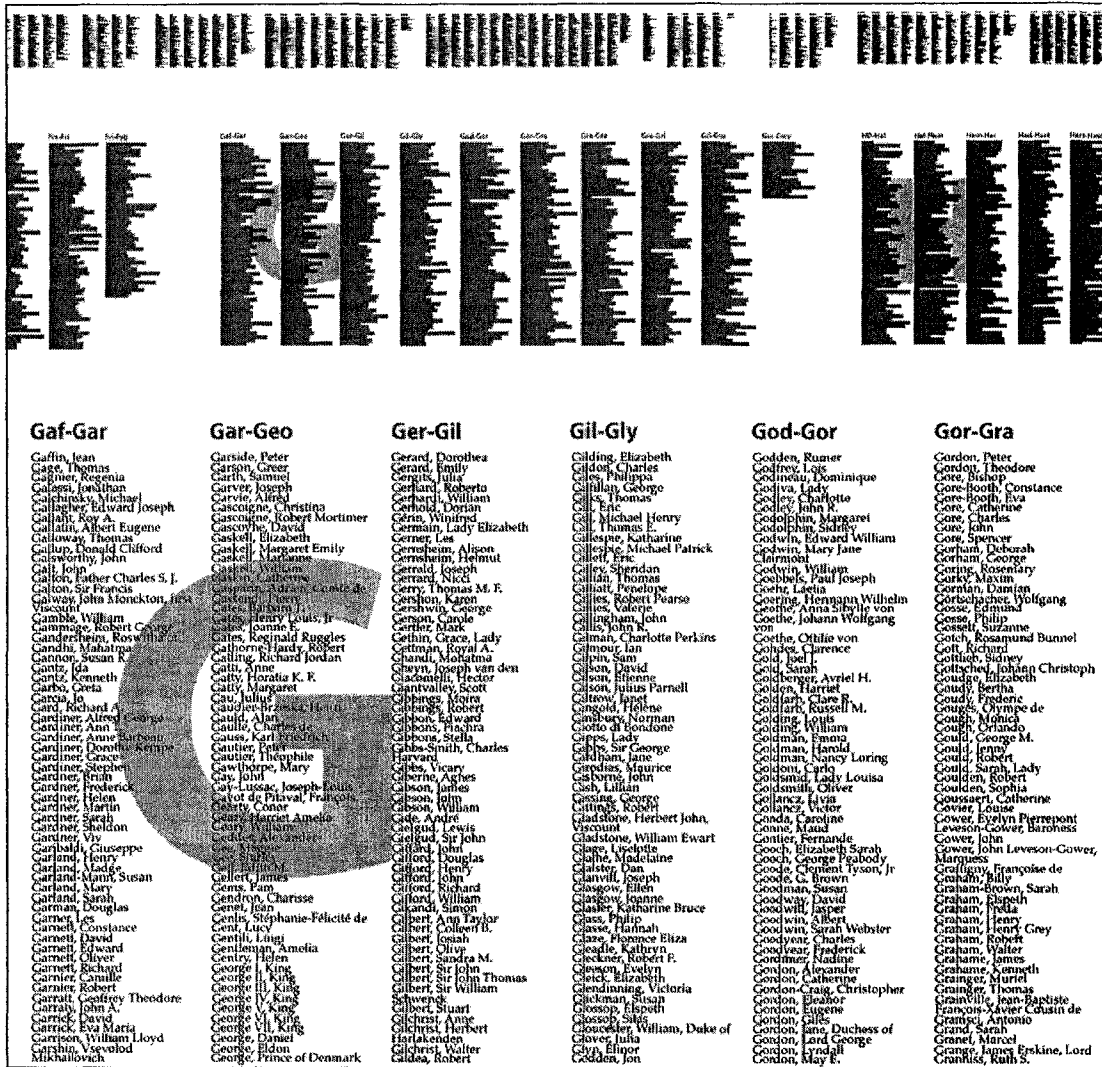


Figure 4.03

This panorama contains roughly 12,000 names from the Orlando Project, of which the smallest strip shows approximately half at any one time. The names are arranged in alphabetical order in columns. The zooming feature is continuous, and corresponds to the user's movement of the mouse up or down on the screen. The strip is shown here at three different levels of magnification.

Cross-references

The Orlando biocritical and events documents contain frequent references to people, places, and organizations. In many cases, the people mentioned in one document are the main subject of another document in the Orlando collection. A basic means of providing access to the writers mentioned in a given document is to have the system identify each instance of an author's name as a live link that can take the user to the relevant biocritical materials in the collection.

It is also possible to provide a hyperlink prospect list that would consist of all the author names collected from throughout the document and displayed in one location (Cameron 2003). This kind of hyperlink list provides the reader not only with a simple method for moving to other points in the collection, but also with an overview of the literary figures who are in some way related to the current author.

The problem with this kind of list as a source of information is that in its simplest form it does not provide the user with any idea of the relationship between the current author and the other authors in the list. A document might, for example, name another author in a casual reference, even though that author has little or no historical connection to the author who is the topic of the current biocritical text. On the other hand, the same document might name another author who had a significant historical connection to the writer in question.

In order to be more useful as a prospect list, the author hyperlinks might therefore also include a brief explanation of relationship, perhaps through display of items from a predefined value list. The Orlando tagging allows for identification of people in various ways, either as family, colleagues, friends, and so on, which could serve as the basis for the labels on the author hyperlinks.

Are there alternatives to naming documents by author?

The logical choice for biocritical document display in Orlando is by the names of the authors. However, since the documents have been heavily tagged, it is also possible for the system to deliver portions of many documents rather than complete biography and writing pairs. It may therefore be possible to find other meaningful representations of what are essentially composite documents relating to a particular tag. The result would be a form of display where author names are still used, but they have been placed within a context that is provided by the tag. This strategy would be most appropriate in cases where there are many documents in the collection that contain the tag being sought.

For example, a user might be interested in reading through the collection materials from the perspective of the history of the education of the women authors represented. One way to proceed would be to open each of the biography documents and read the appropriate sections. Since Orlando has been interpretively tagged, there is a better option, which is to search for the <Education> tag and review the results. The vast majority of the biography documents will contain an <Education> tag, so a rich-prospect display that showed a meaningful representation of the <Education> tags by listing the author's names would not differ dramatically from a rich-prospect display of all the authors in the collection.

What might be useful, however, would be to modify the display so that the authors' names are grouped according to some relevant criteria. One display might show, for instance, the list of authors organized by the schools they attended, so the reader might be able to make connections between individual educational facilities and the people who studied there. A further organizing principle might be to sort the groups within each school into chronological order, so that it would be apparent at a glance how many of the authors might have come into contact with each other while at school together, leading perhaps to further investigation into their subsequent contacts.

For researchers interested in educational politics, for example, another strategy might be to group the authors not by the name of the school, but rather by the name of the city or town where the school was located. Since larger centres will naturally tend to have more schools, patterns may begin to emerge in this case relating to rural vs. urban education, which could in turn be used as the basis for further investigation into the characteristics of writers and their work within the larger patterns of British settlement.

Use of Images to Represent Documents

For some kinds of collections, it may be useful to construct the interface using some graphical rather than textual representations of documents. This strategy will be most promising in collections where some unique image is available to represent the contents of each document and the images can be combined with each of them at a relatively tiny scale to allow as many as possible to fit in the browser window at once. A hypothetical example might be a display of book covers at a book retail site, where the screen would be tiled with images of the upper boards or dust jackets of all the books which meet some given criteria.

Representing documents by image may not, however, be a viable approach in Orlando. Since the Orlando collection currently consists exclusively of text documents, it may be difficult to generate a meaningful graphical representation for each document that is distinct enough to allow it to serve as an index to document content. Providing a display that emphasizes images would also be fundamentally misleading about the nature of the materials.

One possible exception to this generalization would be in the case where the face of the author is available and is widely recognizable. Since many of the writers represented in the Orlando collection are in the process of being recovered from some degree of historical obscurity, there may be many cases where an image is not available, and would not be widely recognized if it were.

However, if it were the case that images were readily available for every writer, and that imagery formed some substantial component of the collection, then it might be possible to combine image and text for a more effective prospect display. In effect it would be somewhat similar to the photo display interfaces developed by Shneiderman et al. (2002), where the user is presented with a panel full of thumbnail images, each of which represents an author in the collection. This kind of display could be manipulated by a set of tools similar to those provided for use in the text-based rich-prospect display, allowing the user to group, organize, sort, and subset the representations – working with the display in addition to working with the contents of the documents.

Heterogeneous Displays

A rich-prospect display that has been sorted in some way, and perhaps also augmented with other kinds of organizational features such as grouping, visual structure, or indexical information will assist the user in moving through the representations of objects in order to find the ones of interest. To put it briefly, organizational features can facilitate visual searching. The mere existence of some structure to the information, does not, however, preclude a user from being able to browse through the display in order to see what it contains in general, rather than whether or not it contains a specific, pre-identified target.

If browsing is actually the primary purpose of an interface, then it may not be necessary to provide all of the elements of rich prospect. The following features may turn out to be optional:

- homogeneous representation
- homogeneous items
- visual structure
- indexical cues

Homogeneous Representation

As discussed in previous chapters, the basic form of rich-prospect display is one that shows a common representation for all documents in the collection. For example, if documents are to be represented by author names, then there should be an author name attached to each document, and nothing but an author name should appear in the display.

There may be situations, however, where different kinds of documents could best be represented in different ways. For example, in the Orlando Project, the documents that contain biocritical materials might be represented by author names, while the documents that contain contextualizing historical material might be represented by a keyword or composite representation. Methods for generating appropriate composites might be developed using facet analytical techniques. Further research is required.

The advantage of this kind of heterogeneous display is that it would allow for cases where a collection contains different kinds of documents, without forcing some labelling principle to be used consistently despite the characteristics of the collection.

The disadvantage of a heterogeneous display is that it will necessarily complicate the prospect, making it more difficult to sort documents or group them according to some standard criterion that relies on a homogeneous representation. To sort the entries in a display of author names alphabetically, for example, makes finding a given author relatively simple. If the same display contains monograph titles as well as authors, the result may be confusing, especially in cases where proper nouns have been used as titles. Additional research in this area would be useful.

Homogeneous Items

Since digital information takes many forms, there are many possible kinds of files that might be part of a digital collection. A heterogeneous set of documents might include the following kinds of information, or others:

- digital images
- sound files
- digital video
- text files
- text files with markup
- spreadsheets
- databases

Each of these kinds of files can also exist in various forms. For example, there are dozens of digital image formats, some of which are proprietary to a particular piece of software from a single vendor, while others have been standardized. Within the various formats there are also possible variations, so that two images that are both Adobe Photoshop native files might nonetheless differ by being layered or flattened, include clipping paths or not, be represented as bitmaps or vectors, and either contain embedded font definitions or else rely on the system to provide them. They may also differ according to the version of the software that created them.

In addition to the variations between kinds of data and within the formats available for a given kind, individual files might also contain information from several formats, so that a text file, for instance, may contain images, sound or video clips, and tabular information originally derived from a spreadsheet or database. In some cases it may be possible to determine a primary information type, while in others it may be necessary to simply indicate that several types are present.

Each of these variations, permutations, and combinations may provide opportunities for constructing some form of meaningful representation of the document. They may also represent obstacles to accurately representing the document, since the designer is faced with the task of not only indicating the content, but perhaps also of indicating the form of the content.

For users with some degree of sophistication, indications of document type are available in the desktop environment, primarily through document icons and file extensions. For users who are less sophisticated, there are also protocols to automatically associate documents with relevant applications. It may therefore be useful in some heterogeneous collections to include file extensions, for example, on the words or phrases that are used to represent the documents.

In its current state, Orlando contains text documents that have been encoded in SGML. It does not contain files of other types. However, future developments in Orlando or in other projects with interpretive tagging may extend the collection into file types that are not encoded text, in which case the role of expressing the meta-data about document type within prospect may need to be addressed.

Visual Structure

A prospect display can be structured in any one of a number of ways. One basic form consists of spaces or bullets between elements in order to separate them visually. Tables or columns can provide structure, as can grid systems, which logically extend the concept of the table.

Allowing people the opportunity to simultaneously sort and group the meaningful representations of items opens another whole area of possibility for creating useful visual structures. For example, a list of names that has been sorted alphabetically can then be grouped by letters of the alphabet. A list of authors sorted by birth date or date of first publication can then be grouped by century. The principle is that the sorting can be carried out at a level of granularity that is finer than the level of granularity applied to the group. A related method is to group related items, then sort them within their groups. For example, authors might be grouped by century of first publication, but organized alphabetically by last name within each group. Given the complexity of the tagging and the numbers of attribute value lists in Orlando, the opportunities for sorting and grouping either tagged text sections or else entire documents are endless.

However, visual grouping and sorting are not the only methods that might be useful. An alternative to grouping is through the use of network diagrams that can associate items by visually representing some logical relationship between them. Entity-relationship diagrams would provide an example of structure by network diagram, as would topic maps.

However, in spite of the range of possibilities for providing structure, it is also possible to create displays with very little if any organizing principles, where the representations of items appear as if at random on the screen.

Unstructured displays in the physical world include items such as sales bins in retail stores, where items are tossed into a common container and consumers browse through the pile. The browsing can be casual, or in some cases quite focused, as when the consumer has a sense of not knowing what might be found, but has a conviction that the right choice will be obvious when it is found. An unstructured digital display may provide analogous opportunities for people looking for serendipitous objects of interest.

Indexical Cues

In a display of items that has been sorted, it is useful for the designer to provide some indexing information that serves to group related items by a relevant criterion, so that the user can quickly

narrow in on the information of interest. In a phone book, for example, each page header contains the first and last names on the page. Since the names are sorted in alphabetical order throughout the book, the user can quickly identify the page of interest when looking for a particular person and number. A similar principle applies to rich-prospect interfaces, where the user can be provided with indexing information to help make sense of the larger interface.

In some cases, however, it may not be possible or useful to provide indexical cues. An example might be in the heterogeneous display of materials where no sorting criteria are obvious. If a display is not amenable to sorting, then it does not lend itself to addition of indexical cues to help the user traverse the material.

ORLANDO DATES

In formulating the principles under which they would subsequently operate, the Orlando Project developers made a decision to provide the users of the collection with the ability to view as much of the material as was appropriate in formats that have been arranged chronologically: “As perhaps the most vital tool for relating historical events and processes to each other, and to the over-arching narrative, we have chosen chronology” (Grundy et al. 2000). This decision has had far-reaching consequences, both for the tagset and for the tagging on the project, because in order to make chronologies available to the reader, it is necessary to attach dates wherever possible. The tags used for this are <Date>, <DateRange>, and <DateStruct>. Dates are, however, only important insofar as they are associated with a block of text. The tag that creates this association in the Orlando Project is <ChronStruct>.

<ChronStruct> is in some senses a fundamental building block of the Orlando tagsets. It occurs in the events tagset as <ChronEvent>, but the purpose of the tags is similar in that both the <ChronStruct> and the <ChronEvent> hold together a date, some tagged text, and the bibliographical references associated with the text. To simplify the following discussion, the term <ChronStruct> will therefore be used to signify both tags.

As far as the user reading the collection contents is concerned, the <ChronStruct> itself is an empty container: it does not directly contain any text. Instead, it contains subtags that contain text. It also contains attributes that are useful in displaying material that has been extracted from the collection to be displayed in chronologies.

The following <ChronStruct> occurs in the biography of Mary Somerville, a Scottish mathematician and scientist who lived from 1780 to 1872:

```
<CHRONSTRUCT RELEVANCE="SELECTIVE"
CHRONCOLUMN="BRITISHWOMENWRITERS" RESP="CJH"> <DATESTRUCT
VALUE="1825-06-"> <SEASON> Summer</SEASON> <YEAR> 1825</YEAR>
</DATESTRUCT> <CHRONPROSE> MS undertook her first scientific investigation: she
designed and conducted a number of experiments to determine the effect of light on
magnetism.</CHRONPROSE> <BIBCIIT PLACEHOLDER="Patterson, Mary Fairfax, 213"
DBREF="7510"> 213</BIBCIIT> </CHRONSTRUCT> (Clements et al. 2003).
```

From the reader's perspective, the experiment occurred in the summer of 1825. From the perspective of the Value attribute on the <DateStruct>, the experiment occurred in June of 1825, which would allow the system to sort this <ChronStruct> to appear in a chronology at the beginning of the summer.

The following instructions to taggers emphasize the nature of <ChronStruct>s as extractable units:

Because chronStructs may be removed from the documents in which they were created and be placed alongside unrelated information, always make sure that you put enough information in a chronStruct such that it will make sense when read out of context. Make sure that any important names, dates, places, or orgNames are tagged inside a chronStruct. Also, do not use pronouns in a chronStruct unless their referent is also present (Clements et al. 2003a).

Dates and Chronologies

In terms of the design of the tagset, the Orlando designers were aware at an early stage that it would not be a simple matter to provide accurate dates for every piece of significant information. The project materials cover centuries of women's writing and historical events. Some of this material could be associated with a single day that is part of the historical record, while in other cases the events might have taken place on a single day, but the recorded account does not provide an accurate indication of which day that might be. In other cases, the events span a range of time, the endpoints of which might be very precise (marked, for example, by the signing of a treaty or the publication of an article in a daily newspaper) or only approximate, or there may be a range that only contains a start date (ex.

“By 1900 women accounted for twelve percent of the library staff in Britain whereas in America ninety-five percent of library staff were women.” (Grundy et al. 2000)).

Each of the possible date configurations has implications for the way the system is going to construct and arrange chronologies. For example, if a <ChronStruct> specifies the month “May,” the algorithm could sort that piece of text anywhere in the month – to the beginning, middle, or end. In the case of a month, the position is not particularly critical in terms of accuracy, but if the date specifies only a year, there is some considerable difference between the beginning and end of a year, and even more difference in the case of a decade or a century.

If the <ChronStruct> explicitly includes a date range, there is a similar problem of deciding how to position the material. The default solution is to use the earliest date in the range, but in cases where a number of other <ChronStructs> are also visible, the reader can lose track of the number of texts that should be understood as occurring during the same period. Chronological searches on Orlando currently sort in the following order:

- year-only dates
- year/month dates
- year/month/day dates

A given year may have few or many <ChronStructs>. 1621, for example, currently has 15 items in its full chronology, while 1921 has 70 year-only dates and 62 others that are either year/month or year/month/day – an order of magnitude difference (Grundy et al. 2000).

Complications involving date accuracy and format are, however, not the only complications in the Orlando tagset and its use of dates.

<Date: Certainty>

In addition to the accuracy with which a date or date range can be specified, there are also indications of how the user is to interpret the degree to which the date is reliable. Since much of the Orlando material is derived from historical sources, there is a range of certainty involved both in the original materials and in the reliability of the reporting. For example, someone remembering an event from twenty years past will usually be less accurate to the day than someone recording an event that happened only yesterday. Some sources are also more consistently reliable than others.

The tags relating to dates in the Orlando Project are therefore equipped with the attribute “Certainty,” which provides the tagger with the facility to indicate the reliability of each date. In the

case of <DateRange>, there are separate Certainty attributes for both the “From” and “To” parts of the equation. The predefined values of the Certainty attribute are: By, Cert (certain), C (circa), Roughly Dated, and After.

Displaying <Date: Certainty>

These values represent an interesting challenge in terms of interface design, since although they are all relevant attribute values for Certainty, they are not syntactically nor semantically in the same class with each other. The user who wishes to understand and use the Certainty attribute values is therefore required to make a different mental adjustment for each of them. Three of the values – Cert, C, and Roughly Dated – might be visualized as concentric circles around a point in time. Appropriate synonyms for these values might be, respectively: confident, approximate, and rough estimate. The other two values – By and After – might be visualized respectively as a line with an endpoint and a starting point with a line (Figure 4.04).

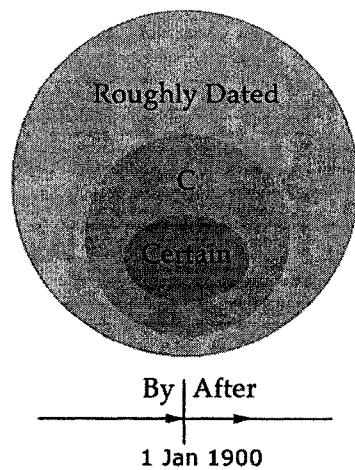


Figure 4.04 Visual representations of the values for the attribute Certainty, which is used with the various Date tags in Orlando. The tags might be similarly grouped.

The Certainty attribute is important not only because it provides significant information about each date, but also because it provides an example of an attribute whose values could be used to structure the display through either grouping or subsetting the items.

<ChronStruct>

The Orlando Project includes biographical information and details about the writing and publishing careers of hundreds of writers, as well as historical information to provide context. That is, the information in the Orlando Project is historical. There is therefore a significant investment in the project in the provision of dates for various items, whether those consist of entries within a biography or writing document or of entries in the events database. It is possible to construct a chronology of women writers who are represented in the collection. It is also possible to construct a wide range of alternative chronologies based on events and sections extracted from the biocritical documents that have been dated. As has been previously mentioned, the primary tag used in extraction and display of chronologies is <ChronStruct>. What has yet to be discussed are the methods available to the designer for visually representing chronologies.

Displaying <ChronStruct>

Various strategies exist for displaying chronological material. Some of these have been developed for print and repurposed for electronic media, while others are primarily electronic both in origin and use. These strategies include:

- timelines
- scattergrams
- sequential prospect
- rich prospect

Timelines

The standard technology for displaying chronological material is the historical timeline, where a directed horizontal line is used to indicate sequence in time, and individual events are indicated either through parallel lines that represent duration, or through perpendicular lines that represent punctive events. Either may be labelled with a brief descriptive text. Explanatory material, usually quite brief for reasons of conserving space, is also sometimes available, as are images that can serve to provide additional information and may also help to orient the viewer.

Timelines have a long history as a print technology, and their re-purposing for digital displays can draw on the existing visual vocabulary. Additional factors come into play, however, since digital timelines can be generated by the user or automatically by the system, rather than exclusively by the designer. Issues of selection and preference and visual weight that would

normally have been under the control of the designer therefore become available as options for the reader. In cases where the timelines can be stored as a form of interaction history, the reader also has the opportunity to communicate with subsequent readers. The use of interactive timelines to convey insights into chronological materials is one of the most exciting areas of possible future research on interaction histories in collections like Orlando.

Timelines are exciting in part because they are a form of visual narrative that is relatively accessible to everyone. Their primary constraints, especially when designed for use on a monitor, are their size and complexity. It is difficult to fit much information on a horizontal strip that will sit within the margins of a browser.

One solution is therefore to provide the user with a magnification strategy, so the timeline can be scaled, either through the addition or removal of secondary events, or else through physically changing the size of the display through some process of magnification or its reverse.

Another recent development in the use of electronic timelines is their application to the display of temporal modelling, where alternate outcomes can be shown as modifications to the timeline (Drucker and Nowvskie 2003). In the Catastrophic Nowslider Demo, the user chooses points on a timeslider that represent the current state of information available to the heroine of a narrative. As the information point shifts, so does the temporal model (Figure 4.05).

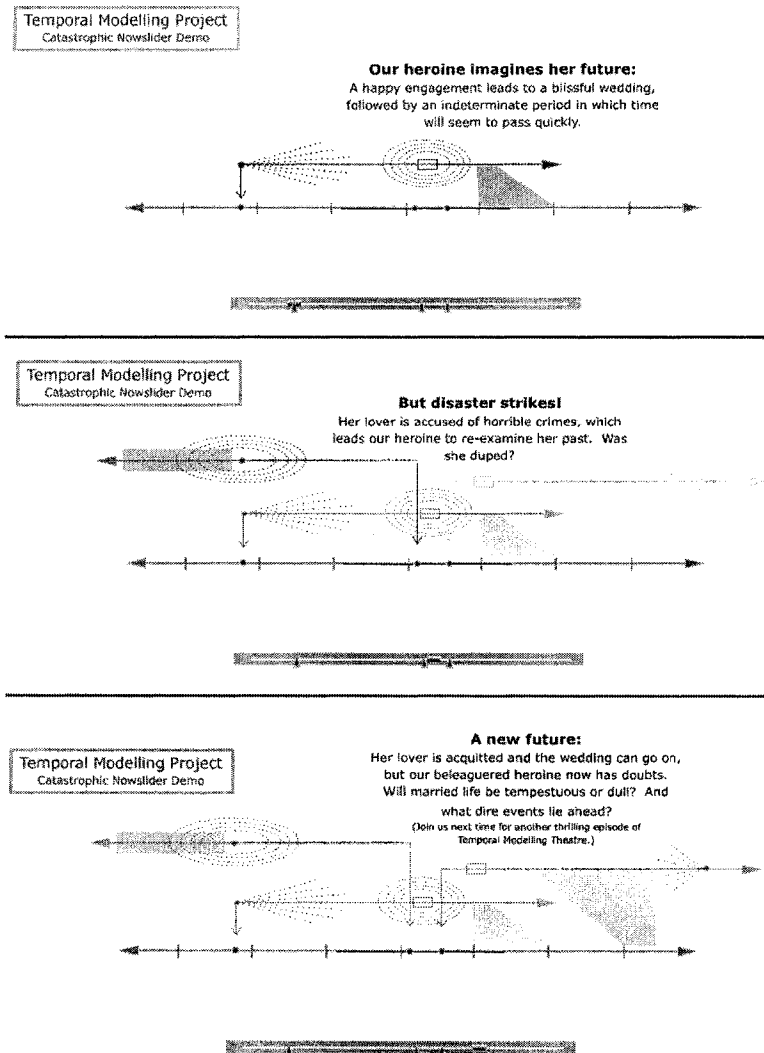


Figure 4.05 The Nowslider provides a visual prediction of the future based on a current state of knowledge. As the user changes the state of current knowledge in the system by sliding the thumb along the bottom timeline, the temporal model on display also changes (Drucker and Nowviskie 2003).

Scattergrams

Another means of providing prospect on a chronology is to create a plot of points, each of which represents a single entry or event in the chronology. The distribution of points along the horizontal axis indicates how many events occur in the collection at each time.

Like timelines, scattergrams can be used at various scales, with the display collapsing individual points into aggregate points as the timescale increases. Alternatively, the number of points

can remain fixed but the vertical size of the scattergram can increase as the horizontal scale decreases, in order to accommodate increased stacking of the event points.

As a means of accessing a collection, a scattergram can be used to select subsets of events inside a range set by the user. The texts represented by the points can then be collected into a subsequent display for further refinement, perhaps through first changing their representation into some form that is more meaningful than a point. Alternatively, the selected points can be used directly as a collection of items to expand for reading (Figure 4.06).

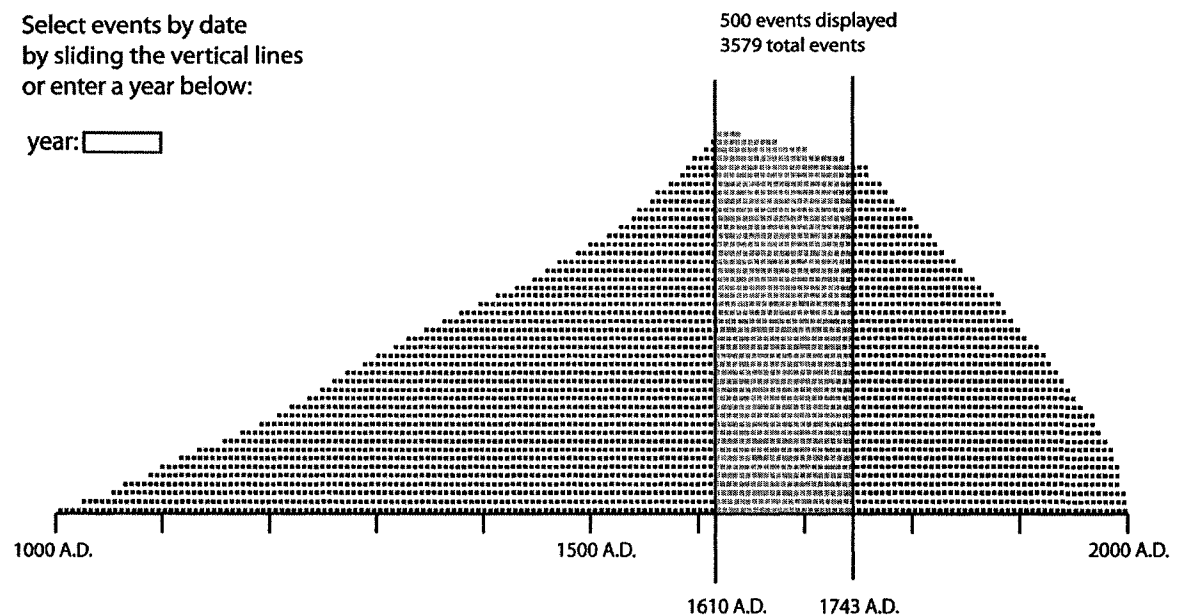


Figure 4.06 This scattergram shows one point for each of 3600 events. The user is able to select subsets of the points by moving the vertical bars at the endpoints of the selection. In this case, 500 events are shown as currently selected.

Scattergrams have the advantage over rich-prospect displays showing meaningful representations of items in that they are relatively compact, and as in rich-prospect displays, if the relationship is one-to-one between points and collection entries, the display can give the user some sense of the structure of the collection in terms of the amount of material available for each period in the chronology. However, scattergrams have the disadvantage that the points themselves are not intrinsically meaningful.

Some meaning can be applied to the individual points, primarily through colour coding, since the single pixels are not amenable to differentiation by shape. However, if the scattergram is

implemented in such a way that the user can magnify it, then there is the possibility of having the individual points expand into larger representations that could be meaningful, either through shape or labelling. As the scattergram has been magnified, it will become less compact, but will transform into what is essentially another rich-prospect display.

Sequential Prospect

Having all of the items on display at one time inevitably requires the use of strategies to accommodate the limited screen real estate. A method that sidesteps this necessity is to have some form of sequential prospect, where the user is able to scan through a representation of the collection items by viewing them one at a time in quick succession.

An example of a sequential prospect tool is the range slider developed by Ahlberg and Shneiderman (1994), which allows the user to move a horizontal thumb in order to view a lengthy list of entries. The technical obstacle to be overcome in the use of sequential prospect sliders is that for fairly large collections, the position of the thumb on the bar itself is not an appropriate means of setting the location in the collection, since the length of the bar would need to extend well beyond the sides of the screen. The suggested alternative approximates the position of the thumb but with a much finer level of granularity, by using the position of the mouse to determine which item to display. Since the mouse movement can be coupled fairly loosely to the thumb movement, even a fairly small slider can be used to traverse collections numbering in the tens of thousands of items.

In the case of the Orlando collection, such a device might be used to provide prospect on several different kinds of information. For instance, a pair of sliders might be used to display on one hand the list of tags in the tagset, and on the other hand a matched list of all texts found in the tag. An alternative pair of sliders could be used to show respectively all available tag attributes and their attribute values. A slider could also be used to display all of the names in the collection, all of the dates in the collection, or all of the documents in the collection.

Sequential prospect has the strong advantage of not requiring excessive amounts of screen space while still providing the user with some means of looking directly at collection contents, tagsets, and so on. If the sliders are also amenable to different kinds of sorting, then the user would have the opportunity to determine the order in which the items are going to appear. For example, the same slider might be used to show the names of the authors receiving biocritical treatment, first in alphabetical order, then in chronological order by date of birth. Additional indexical cues might be

added in support of each kind of sorting, so that for example if the slider is horizontal and the name appears above it, then the letter of the alphabet or relevant date might simultaneously appear beneath the slider.

Rich Prospect

In order for an interface to have a rich-prospect form of display for chronological data, it is necessary to show some meaningful representation of every chronological item, either within the entire collection or within the current date range of interest. Chronological items are generally quite brief, consisting largely of single sentences or short paragraphs. One solution is therefore to provide chronological material as a complete listing of <ChronStruct> contents.

The disadvantage of using the entire entries is that even single sentences can quickly fill the available screen space, especially when it is necessary to provide additional line spacing between items to indicate that they are not part of the same entry. In order to take maximum advantage of the available screen area, it is therefore preferable to find some means of representing chronological events in an abbreviated form.

A basic strategy would be to represent the items in a chronology as dates. However, the meaning inherent in a date is only a small part of the event. In cases where the events occur simultaneously or in quick succession, the dates may either need to be refined to an unreasonable degree in order to distinguish the events, or else a single date may have to be used to access multiple events.

In order to provide more information about the events and to avoid the one-to-many relationship between interface items and chronological events, it would therefore be more useful to create a display representation that included both the date and a brief keyword, phrase, or title to label the event. If the keywords are not unique, then they would be useful in grouping or subsetting a larger display into sections related to various topics of potential interest to the reader. For example, the keyword “suffrage” might be associated with events in the Orlando Project relating to the securing of votes for women.

However, once the items marked “suffrage” are grouped or extracted as a subset, it is no longer useful to mark them with that non-unique keyword, because every item in the display would use the same word. For purposes of distinguishing between items in the same group, it would be useful to provide a second, unique keyword or phrase which could be used to replace or supplement the

non-unique keyword and date. Since the point of creating the representation of items is to save screen space, substituting the unique keyword for the non-unique one may be the best option, with the non-unique keyword perhaps being moved to a position that indicates that it applies to the entire group or subset of representations.

The disadvantage of keywords is that they are labour-intensive to apply and maintain, since each event must be keyworded at both the unique and non-unique levels. The list of keywords also needs to be established in such a way that changes are kept to a minimum, since the addition of new keywords would require that someone review previously keyworded events in order to see if the new keyword also applies.

Attaching a keyword is also an act of interpretation that is analogous to the interpretation involved in attaching textual markup. One solution is therefore to apply as many keywords as possible. However, if their purpose is to simplify the display, then a long list of keywords is not going to be any more useful than a descriptive phrase might be, since both involve several words to describe a single event.

Since many of the events in the Orlando Project relate to historical activities of people or organizations, one possible strategy would be to use the existing tagging to generate descriptive text for representing the items. The representation would then consist of a date, the contents or standard attribute contents of one of the other core tags such as name, place, or orgname (which would be in most cases non-unique), and a tag selected from a list of potentially relevant ones. For example, one event might be described as date, name, and a tag relating to life stages: 1879, Annie Kenney, birth. Another event might be displayed as date, place, and a tag relating to historical activity: 18 June 1815, Waterloo, battle.

Using the existing tagging to generate representations has the advantage that it can be automated and does not rely on the maintenance of keyword lists and their application. However, there may be cases where the algorithms for selection are not going to result in meaningful text that genuinely represents the major contents of an events. Further research in this area would be useful.

<ChronStruct: Relevance>

In order to provide an idea of how important an individual event was in the grand scheme of history, the <ChronStruct> tag includes an attribute for relevance. The relevance attribute has four possible values: Selective, Period, Decade, and Comprehensive. They are in increasing order of magnitude of results. The system is designed in such a way that searching for <ChronStruct: Relevance: Period>

will actually return not just the Period items, but also the <ChronStruct> paragraphs that were marked with <ChronStruct: Relevance: Selective>. Similar treatment is given to each of the subsequent attribute values, so that searching, for example, for <ChronStruct: Relevance: Comprehensive> will return all <ChronStruct>s.

In addition to the relative scale, the semantics of the attribute values are also significant. The first value – Selective – is used to mark only those items which the project personnel consider essential to a basic chronology. The paragraphs describing landmark events in an author's life, such as birth, death, and major writing or publishing activities (such as first and last publication, or publication of the most-famous works), are all marked with <ChronStruct: Relevance: Selective>. If the user searches the collection for a particular author and constrains the search for only the selective ChronStructs, the result will be a brief sketch of the highlights of the author's life and writing career, along with major contemporaneous world events.

The next value – Period – is used for material that might be appropriate for a standard undergraduate university course, as for example a course in Renaissance literature. Period also indicates material that falls within identifiable historical eras that are not necessarily equivalent to the ones usually applied to literary studies. For example, if a user were interested in writing activities during the War of the Roses, the Period attribute would be appropriate.

The third possible attribute value for <ChronStruct: Relevance> is Decade, which is used to locate details surrounding a particular historical event or relatively short span of time. For example, while a user interested in women's suffrage would likely want to search for Relevance: Period, a user interested in the first incarceration of suffragists such as Millicent Garrett Fawcett might prefer to search using Relevance: Decade.

The final possible value for <ChronStruct: Relevance> is Comprehensive, which is used to mark material that is significant in the biography or writing career of an author, but which is not necessarily of historical importance. Examples might include dates of starting or leaving a particular job, or dates marking the birth or death of parents, spouses, or children.

The <ChronStruct: Relevance> attribute and its four possible values are significant because they will constrain the results that the user can expect to obtain from a given date search. However, the details of how they have been defined and implemented represent a potential obstacle to the user, which is exacerbated by the fact there is no standard terminology available to indicate what the attribute values signify.

Displaying <ChronStruct: Relevance>

The attributes on <ChronStruct: Relevance> will determine the size of the set returned to the user by a chronology search; it is therefore necessary that the user be able to specify which of the four options are appropriate for a given search. One default solution is to provide the user with a set of radio buttons, which are a standard GUI method of allowing mutually-exclusive choices (Figure 4.07).




Figure 4.07 A radio button interface to the <ChronStruct: Relevance> attribute values would allow the user who is familiar with the terminology to select an appropriate choice. For users unfamiliar with the terminology, some additional experience or explanation might be necessary.

However, a radio button choice on a search screen constrains the user to one selection at a time, which indicates that the values are mutually exclusive. Since this is not the case, even though a set of radio buttons could be re-purposed to provide the user with the correct result, the meaning of the selection tool is fundamentally misleading.

Another standard selection tool is the set of check boxes. Check boxes allow the user to have multiple simultaneous selections. In an interface that does not specify how the multiple choices are to be combined by the search engine, the selection is ambiguous. On the one hand, choosing more than one item might mean that they all need to be present in the result (a logical AND). On the other hand, choosing more than one item might mean that any one of them should be present, but that it is not necessarily for them all to be present (a logical OR). A third logical possibility is even more difficult to communicate – this is the logical XOR, or the case where one or the other but not both items should be present in the search results.

Check boxes also do not indicate to the user that the values themselves are additive: instead, they are assumed to be distinct from each other. The difference between a radio button and a check box is simply that the radio buttons only allow one choice at a time, while the check boxes allow multiple choices. The visual syntax of the two devices therefore indicates that the one has a constraint that the other does not have.

In the case of a selection mechanism for the interface to a search engine, another possibility would be to use a slider that moves between the anchors “selective” and “comprehensive.” However, like radio buttons and check boxes, sliders have an enculturated semantics – in this case, one that suggests a continuum. Since the values available for <ChronStruct: Relevance> consist of four discrete possibilities, a slider sends the wrong message to the user.

It is possible, however, to develop prospect-related solutions for providing the user with the necessary functionality, without requiring that the user understand the <ChronStruct: Relevance> attribute values.

The key point to be made with respect to <ChronStruct: Relevance> is that the values are additive. In order for the interface to indicate to the user the proper relation between the values, it is therefore useful to consider alternatives that are also additive.

The <ChronStruct: Relevance> attribute values may therefore be a case where the existing interface options are not appropriate. What is required is that the user understand that selecting each of the available values in turn would generate an expanding set of results, with the fewest results occurring at “selective” and the most results at “comprehensive” (given that other search criteria remain constant).

One appropriate solution is therefore to show the choices as a set of nested buttons, with “selective” in the centre and “comprehensive” as the label on the largest button. Constraining the buttons so that the inner ones are automatically selected when the user chooses an outer one makes the choice clear, even when the choices are available only as part of a search interface (Figure 4.08).¹




Figure 4.08

Since the four possible values for <ChronStruct: Relevance> are additive, one appropriate solution is to allow the user to select them by choosing among nested buttons where the outer choices automatically include the inner choices. The degree of grey on each button is supposed to reinforce the idea of additive selection.

¹ It is interesting to note that this solution was developed independently by different members of the Orlando delivery team.

With an additive selection mechanism, the user will know that choosing <ChronStruct: Relevance: Period> will provide more results than choosing <ChronStruct: Relevance: Selective>, but there is no indication of how many results there may be in either case.

If, however, the interface involved is one where a rich-prospect version of the tagging in the collection is available, the display of an additive grouping of tags makes the understanding of the Relevance values intuitively available to the user (Figure 4.09). The names of the relevance attributes are not essential to the display.

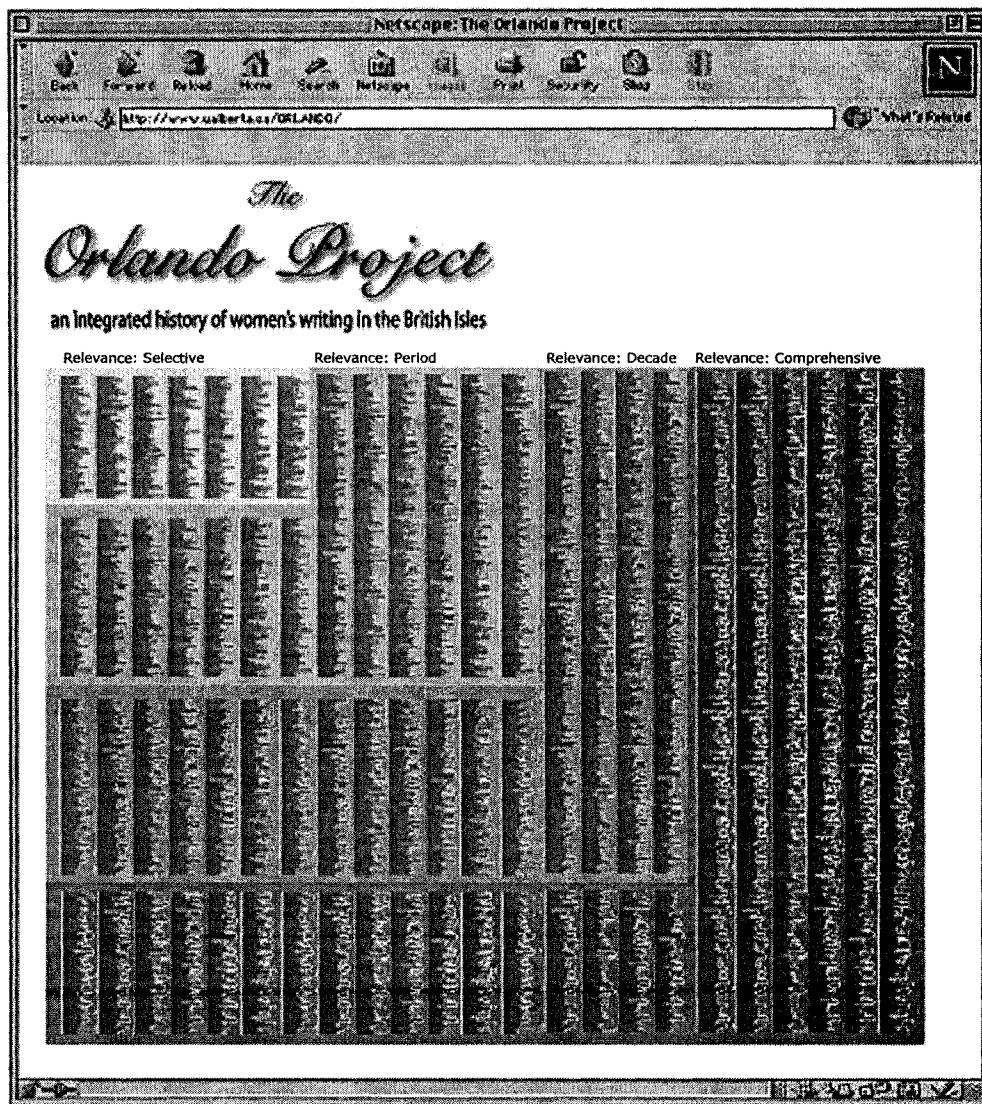


Figure 4.09 Grouping tags by Relevance values provides the user with an immediate impression of the significance of the different values.

<ChronStruct: ChronColumn>

Like Relevance, ChronColumn is an essential attribute for users of the Orlando Project interested in retrieving, viewing, and otherwise working with material arranged in chronological order. Also like Relevance, the ChronColumn attribute has four possible values, which in this case are: British Women's Writing; Writing Climate; Social Climate; and National International. However, unlike the additive attribute values for Relevance, the ChronColumn values are used to mark information that is mutually exclusive. That is, for example, a <ChronStruct: ChronColumn: British Women's Writing> is not a subset of a <ChronStruct: ChronColumn: Social Climate>: the attribute values are used to distinguish between different kinds of material.

- The majority of the ChronStructs in the collection are about British Women Writers.
- Writing Climate marks equivalent material for male writers and women writers who are not British. It also marks anything else related to the literary industry.
- Social Climate, on the other hand, is the attribute value used to signify information on topics of historical interest which are outside the bounds of the literary. Events dealing with science, law, fashion, and so on would all be marked with Social Climate.
- Finally, National International is the attribute value for events related to areas such as military or political history.

The following text, a ChronStruct from the writing document of Christabel Pankhurst, is an example of a passage that has been marked with the ChronColumn attribute "British Women's Writing."

15 October 1908 CP gave a speech at the St James's Hall titled The Militant Methods of the N.W.S.P.U., which was published verbatim by The Woman's Press the same year. (Clements et al. 2003)

Here is the identical passage, with all of its tags visible:

```
<CHRONSTRUCT RESP="KDC" CHRONCOLUMN="BRITISHWOMENWRITERS"
RELEVANCE="SELECTIVE"> <DATE> 15 October 1908</DATE> <CHRONPROSE>
CP gave a speech at the <PLACE> <PLACENAME> St James's Hall</PLACENAME>
<SETTLEMENT REG="London"> </SETTLEMENT> </PLACE> titled <TITLE
TITLETYPE="MONOGRAPHIC"> The Militant Methods of the <ORGNOME
STANDARD="Women's Social and Political Union"> N.W.S.P.U.</ORGNOME>
</TITLE> , which was published verbatim by The <ORGNOME> Woman's
```

Press</ORGNAM> the same year.</CHRONPROSE> <BIBCIT
 PLACEHOLDER="Pankhurst, Militant Methods 34" DBREF="7998"> 34</BIBCIT>
 <BIBCIT PLACEHOLDER="OCLC" DBREF="1709"> </BIBCIT> </CHRONSTRUCT>
 (Clements et al. 2003)

A user interested in working with an events chronology might wish to combine results from any of the four possible ChronColumn values, or select a single value as the focus of attention. The Project focus is emphasized by the default chronological sort, which puts British Women Writer events at the top of any list of events that share the same date.

Displaying <ChronStruct: ChronColumn>

In order to facilitate searching for any combination of ChronColumn values, the standard interface tool that is most appropriate is a set of checkboxes, so that the values can be mixed and matched. The default setting for the checkboxes could either be to return just the British Women Writer ChronStructs, or else all the ChronStructs, depending on the preference of the collection designers. Since checkboxes are familiar to GUI users in general, as long as the interface shows the set of checkboxes, the user is able to modify the selection before running the search.

From the perspective of prospect, indicating which of the four ChronColumn attribute values has been applied to each of the items showing in a display is somewhat problematic. There are two basic classes of solution. The first involves associating the individual meaningful representations of items with some characteristic that indicates the ChronColumn attribute value. The second involves organizing the display in such a way that items with the same ChronColumn values are visually grouped together.

Attaching Visual Cues to Individual Items

In the first category, possible solutions include the use of icons or text. Either of these elements could be further differentiated through secondary visual attributes such as the application of colour. In the case of fonts, morphological changes could also be applied, consisting either of different fonts or different styles of the same font (e.g. bold, italic, oblique).

If the strategy is to use an iconic representation of each of the values, and attach the icon to each element in the display, the result may, in some cases, be a considerable amount of repetition of the same icons, since there may be hundreds or even thousands of items on display, each of which shares the same ChronColumn value. Another potential problem with icons is that their meaning is

not always simple to establish, especially in the case of relatively complicated terms such as the ChronColumn values. In order to create meaningful iconic representations of the four values, it would be necessary to study a reasonable sample of users, and perhaps develop a system that displayed different icons for users from different cultures, since it is fairly well established that cross-cultural implementation of icons often results in confusion about the meanings.

If the choice is to use text labels, similar problems may arise in terms of repetition and possible misinterpretation of meaning. Text may also require more screen real estate than icons.

Colour-coding presents several difficulties. First of all, the use of four different colours could pose problems for some readers, who may find some of the colours less congenial than others. There is also the problem that colour is not in itself intrinsically meaningful, so that the user who has difficulty in associating meaning with an icon or a text label may find colour-coding even more difficult to interpret or remember. Since many people only have access to printing in black and white, there is the added logistical problem that printouts might not preserve the colour distinction. Finally, there is the problem that some percentage of the population is going to have difficulty with any system that relies heavily on distinctions based on colour, because they are not able to perceive the colours distinctly, or in some cases at all. The principle of inclusive design suggests that these people should wherever possible not require specialized equipment or strategies, but should be accommodated in the original design.

Font and other morphological variations share many of the problems associated with colour: fonts are not intrinsically meaningful, which adds an arbitrary memory demand on the user. A display using four fonts or font styles simultaneously may also be unattractive or difficult to read.

Grouping Items

Grouping the items according to ChronColumn, on the other hand, can be more or less effective depending on the details of how the groups are arranged. For example, if each ChronColumn value were to be assigned to a column in a 4-column table, the result would be a fairly clear indication of which texts belonged to which value. However, the resulting four columns would each consist of relatively small horizontal portions of a standard monitor, making reading difficult and potentially irritating, especially in cases where screen space was lost to one or two entries in one column when another column extended to dozens.

A possible solution to the screen allocation problem in a tabular display would be to make the columns readily collapsible and expandable, so the user would be able to choose whether to view them in parallel or to view each of them in turn. If some indications were also available to suggest which column belonged to which value, and perhaps also to suggest how many items were in each column, the display using columns might be relatively simple and intuitive to use (Figure 4.10).

British Women's Writing	Writing Climate	Social Climate	National International
			3 September 1791 France became a constitutional monarchy, in which (as in England) suffrage was restricted to property-owners.
			11 August 1792 The French Legislative Assembly voted to establish a National Convention elected by universal manhood suffrage .
			October 1793 A British Convention on manhood suffrage was held in Edinburgh, bringing together delegates from the London Corresponding Society and Scottish delegates of the Society for Constitutional Information. Edinburgh authorities closed the meeting and arrested its leaders. Arguments that the convention was treasonous (as the government sought to prove) were heard at the trials, in Scotland, of William Skirving, secretary of the convention, and Maurice Margarot and Joseph Gerrald, delegates from London. All three were found guilty and sentenced to three years' transportation.
			1797 Charles James Fox hinted, as the merest possibility, the idea that educated women might appropriately vote. He was explicitly in favour of universal manhood suffrage .
			7 June 1837 The London Working Men's Association issued a Six Point petition in conjunction with six Radical MPs. The six points were manhood suffrage , annual parliaments, the ballot, payment of MPs, equal electoral districts, and the abolition of the property qualifications for parliament.
			8 May 1838 The "People's Charter" was published in London by the London Working Men's Association. The six points of the Charter were: universal manhood suffrage ; the abolition of property qualifications; voting by secret ballot; annual parliaments; salaries for Members of Parliament; and the equalization of electoral districts. The Charter became a rallying point for the first major working-class political movement, known as Chartism, which attracted more followers than any other in the century until the rise of socialism. An early draft of the Charter had included universal suffrage , but it was feared that giving the women the vote would prove too contentious and the provision was removed. Women were nevertheless quite active in the Chartist movement, although popular participation dwindled after the mid-forties.

Figure 4.10

Here the four <ChronStruct: ChronColumn> values are used to organize a table of chronological results. The horizontal bars indicate that some of the columns have been collapsed to the left. The heading of each column indicates the value. Dragging the vertical strip associated with each heading will expand that column and collapse the others, although some representation is always visible for all four possible values.

PROSPECT-RELATED TOOLS FOR ORLANDO

Some of the advantages to the user in having a prospect-based interface derive from the visible presence of the meaningful representations of the collection items. These advantages include the immediate observability of what the collection contains, with the implications that can be derived from those observations as to how the designers of the collection understood those contents, as well as the cognitive reassurance that the items being sought either are or are not present in the collection.

A prospect-based interface also lends itself to manipulation in a number of ways, each of which provides the user with some additional functionality. In order to allow the user a number of opportunities for action, the designer must include the appropriate tools for the user to work with the display. Two of the most important manipulations are those designed to sort the data and to group it.

The most appropriate use for the various kinds of information found in the markup is in some ways dependent on the characteristics of the information. There will be some information that is more suitable for sorting the display of the document representations, and some information that is more appropriate for grouping items in the display.

In order to be useful for sorting purposes, a tag or tag attribute should have several characteristics:

- the tag should occur only once in every document
- the information marked by the tag or included in the attribute should be different in each document
- the information should be meaningful to the user (unique tag identification numbers, for instance, are not a good candidate attribute value for sorting)

In order to be useful for grouping purposes, a tag or tag attribute should meet the following criteria:

- the tag should occur only once in every document
- the range of values should be restricted to a relatively small number of possibilities

In cases where the tags, rather than the documents, form the basis for the rich-prospect display, then the first criterion (that tags should occur only once in every document) of course becomes irrelevant.

Sorting

There are a variety of possible sorting criteria for a rich-prospect display of the contents of the Orlando collection. The most obvious way to sort the rich-prospect display of biocritical documents is alphabetically by author's last name. Another fairly straightforward idea would be to organize the display chronologically by date of author's birth. These two sorting schemes are useful in several scenarios, including respectively for people looking through the collection for a particular author (especially in cases where the spelling is uncertain), and for seeing what historical period a particular author falls into, along with the other authors who were her contemporaries.

However, because of the complexity of the tagging in Orlando, there are many other possibilities, both for sorting and subsorting. Some of these sorting schemes are quite directly related to characteristics of the authors or their publishing careers, while others utilize other information encoded in the collection.

Examples of this latter kind that have already been discussed include sorting according to the contents of the <Education> tag, and according to the attribute values of <ChronStruct: Relevance> and <ChronStruct: ChronColumn>.

Sorting the prospect display in a meaningful way is important in that it can allow the user to quickly narrow a visual search down to a few items of particular interest. To have the display sorted according to an appropriate criterion is essential in those cases where the user is hoping to obtain some cognitive reassurance regarding the identification of a single item or a group of related items. For example, someone looking for information on the author Mary Davys might be uncertain of the spelling of her last name, and not know her first name at all. With a retrieval interface to the Orlando Project, running a search on Davys but spelling it "Davies" would result in the retrieval of the biocritical documents on Emily Davies. The user may be uncertain whether the correct result has been obtained or not, and may waste time looking through the materials the system has returned, before becoming aware that the intended search target and the actual search result are not the same person.

If the display contains meaningful representations of every pair of biocritical documents by author name, and is furthermore sorted alphabetically by the last name of the author, the same user would be able to look closely at the section of the display that contains both Davies and Davys. If the user sees the two names in proximity, chances are increased that he or she will mentally register that the collection contains two authors with similar names, and the chance of accessing the wrong one by mistake will be reduced.

Grouping

Some information in the tagging will be primarily useful for sorting the documents in the collection, while other information will work best as the basis for creating groups of documents or tags. The difference is in the nature of the information, rather than in the choices available to the user. For example, any tag that has an attribute that is meaningful in its own right and also has a fixed list of attribute values is a good candidate for the creation of groups, because the fixed value list provides an organizing scheme.

On the other hand, an attribute that has values that are not meaningful to the user (such as an attribute intended to attach a unique identification number to each tag) is not going to be useful for the purposes of grouping. An identification attribute would also not be helpful because the list of values does not fall naturally into groups: these attributes would have to be grouped (if at all) by some larger organizing principle.

In order to make use of identification attributes for grouping, one solution would therefore be to include a set of prefixes on the numbers, which might have meaning according to some predefined code. Another strategy might be to group the numbers by one of their digits.

Some attributes that do not have fixed value lists may also still turn out to be useful for grouping tags or documents, depending on the nature of the actual data that occurs in the tagging as it has been implemented in the project. It may happen, to take a hypothetical example, that the Standard attribute on the <Orgname> tag will have been implemented in such a way that the organizations designated fall into groups. These groups might consist of tags or documents that all mention the same organizations.

On the other hand, the groups might be formed from some higher-level organizing principle that can be parsed from the contents of the <Orgname: Standard> attribute. This strategy in its simplest implementation would involve looking for words such as “school” or “Inc.” which could be used as the basis for creating groups or subsets of the tags or documents. A more-sophisticated approach might combine the selected words with a thesaurus of synonyms, so that the parser could identify instances that are not syntactically identical but nonetheless suggest the same meaning.

For example, the following <Orgname> tag occurs in the writing document of Anna Maria Bennett, who was a novelist in the latter half of the 18th century:

<ORGNAME STANDARD="Minerva Press" REG="Minerva Press">Minerva</ORGNAME> issued two works by another Bennett whose name may (like various inauthentic Radcliffes) be a publisher's fiction... (Clements et al. 2003).

For someone interested in grouping together all the publishers discussed in the Orlando collection, the text marked by the <Orgname> tag would not be useful, since it says simply "Minerva." However, the <Orgname: Standard> attribute contains the words "Minerva Press." Someone looking for publishers might therefore retrieve a list parsed from the attribute contents which has identified the word "Press" as a possible synonym for "Publisher."

A further refinement would be to combine this kind of parsing of attribute values with a similar algorithm to parse the tag contents, so that even in cases where the value that has been entered into the <Orgname: Standard> attribute is not sufficient to identify an instance of a particular kind of element, the text in the <Orgname> tag might provide the appropriate information.

The different kinds of grouping strategies are going to be more or less precise. Grouping based on a fixed list of attribute values is going to be as accurate as the implementation of the tags and attributes. Grouping based on selected words in the attributes requires the addition of a parser, which can introduce errors of unintentional omission in cases where the morphology of the selected word differs from the standards acceptable to the parser. Problems might arise in this case through archaic or foreign spelling, as well as through inflections in English. A stemming algorithm, which allows the parser to identify items that differ by standard inflections, can help reduce the problem, but the risk is nonetheless present that some items might be assigned to the wrong groups or else not included in groups to which they should belong.

To provide the user with as much control over the process as possible, it may be useful to place the mechanism of the identification process under user control, in which case the system would allow people to choose from the following options:

- exact matches
- matches with stemming
- matches with stemming and thesaurus items

Any of these options could be combined with the choice as to whether the system should examine attribute values, tag contents, or both.

There may also be cases to resolve where two or more words are found in the same attribute. A probable solution in this situation is to allow many-to-one relationships between group membership and actual tags or documents in the collection.

Interaction Histories

Another kind of tool that is not necessarily tied directly to prospect is the provision of some form of interaction histories, where a subsequent user is allowed to benefit from the work of a previous user. Cases where the interaction history is related to prospect would include histories that retain manipulations of a rich-prospect display by a user to create new configurations of sorting or grouping the items, as well as custom ways of labelling the material.

In any interaction histories, the issue arises as to the means by which a given interaction is stored, described, and perhaps also vetted for content or quality. Ideally, each history would be subject to review by a competent editor, who could ensure that the user's activity has been carried out in a complete and accurate manner. In order to be able to understand what the user was attempting to accomplish, it may also be useful to provide some means for annotating an interaction, as well as providing it with some meaningful title. Each of these items could then be provided by a user interested in creating a history item for subsequent people to access, with the system reviewer providing a safety net.

If the logistics of having a person involved in the review process prove unmanageable, it is also possible for interaction histories to be created automatically by the system to record any significant interaction by a user. However, in order for subsequent users to have a list of interactions that are meaningful, such records either need to be of actions that are self-explanatory, or else of actions that have been labelled by someone in order to make their meaning clear. An example of this kind of interaction is in the Amazon.com lists of books that were also purchased by people who purchased the current book being displayed. Since the books are grouped by purchaser, the system can automatically create connections between each book in the group and the entire set of books purchased.

In cases where there are sufficient numbers of similar histories for some statistical operations to be applied, another layer of sophistication might be added to calculate and report levels of significance or other metrics. In the case of the Amazon book recommendations, for instance, the list

of books might be ranked according to frequency, or if enough duplication takes place across the system as a whole, a threshold might be set before a particular book is included in the group.

With respect to the Orlando collection, an analogous interaction history might consist of grouping authors according to the documents opened by previous users. The group of related links might be labelled "People who read about this author also read about the following authors." This heading would be followed by a list of author names with links to their documents. What remains to be determined is whether or not people accessing the collection would find previous interaction records of this kind interesting and useful. Further research is required.

One problem, however, with automatically-generated histories in general is that the system has no means for capturing the intention of the user. In cases where the intention is fairly idiosyncratic, the information captured might turn out to be useless to subsequent users. The system also cannot easily differentiate optimum interactions from garbage. For example, one user might access a set of authors based on some criteria that are significant and represent a widely-accepted understanding of the collection, such as dividing authors by genre and literary period. Another user might randomly select half a dozen author names simply to get a sense of the kind of material in the collection. If the "other authors accessed" history is automatically generated in both cases, subsequent users have no way of knowing that the second person was not engaged in an activity that would be useful to anyone else. Worse still, if the two interactions are merged into a single group of "other authors accessed," the random list may corrupt the significant list with extraneous entries.

One solution to this problem is therefore to have the system record the interaction for subsequent review by the person responsible, who could choose to label, annotate, and store it, or else ask the system to delete it. This solution also has the advantage of allowing users to maintain some level of privacy in their use of the collection.

A reduced form of interaction history might also be provided through a sub-system that identifies each user uniquely, either through a user identification and password protocol, or through a machine-based client certificate or cookie, or some combination of the two. In this case, the interaction history might consist of records or past activity provided to benefit individual users on subsequent visits, rather than to benefit all subsequent users. Although this solution has the disadvantage of not providing additional functionality to all subsequent users, it has the advantage that users do not have to worry that their actions are being recorded contrary to their own wishes.

Private interaction histories can also be a source of information for other users, in the case where the original user is willing to provide others with the necessary username and password. Examples of this kind of activity might include school teachers who locate a group of related biocritical and events materials and save them as a personal record, then pass the access information to students in order that they can also view the collected information.

From the perspective of a design standard that attempts to develop prospect for the users of a collection, the existence of interaction histories of any form is another opportunity to express in a meaningful representation at the level of the interface some form of the available information. Prospect on interaction histories might be provided through any choice of the methods described for providing prospect on the contents or the tagset or the tagging, depending on the complexity and details of the particular interaction histories, as well as on the relevant characteristics of the user community and of the design language adopted for the rest of the interface.

ACADEMIC USERS AND USABILITY

The primary users of the Orlando Project are likely going to be academics at some level, since the complexity of the material, the depth of the treatment, and the style of writing are all appropriate for a university audience. There may also be teachers and students in high schools or middle schools who find Orlando materials interesting and useful, although perhaps also challenging.

Usability testing procedures, and in particular the user-centred usability principles that outline the process for creating computer interfaces, have been developed and have undergone subsequent refinement for several decades. Gould and Lewis (1985) outline three key principles that characterize the design process to create usable interfaces:

1. Early focus on users and tasks, including cognitive, behavioral, anthropometric, and attitudinal characteristics, as well as the nature of the work
2. Empirical measurement using simulations to do real work
3. Iterative design that is not just fine-tuning, but is intrinsic to the project

However, because of the changing capacity of the technology and the expanding sophistication of the user community, as well as changes within the interface design community itself, the kinds of interfaces resulting from the design process have changed substantially. As a consequence, an interface that might be considered exemplary from a usability perspective at one point in time may prove to have significant shortcomings at a later date. The changing landscape of interface

standards and user experience does not, however, imply that the process for designing usable interfaces is inadequate. On the contrary, it is possible to interpret the slow migration of expectations as to what is acceptable as a sign of the success of the process – as an indication of the movement from functional to usable to pleasurable.

In the case of the design of the interfaces for Orlando, the needs, expectations, previous experiences, and other characteristics of the academic and other possible user communities all need to be discovered and accommodated in the design process.

CONCLUSION

The Orlando Project represents an important opportunity to study the design of rich-prospect interfaces, because it is a collection containing an appropriate number of documents (in the low thousands), with a homogeneous content (British women writers). It also provides the basis for examining the use of rich-prospect strategies as a means of repurposing the tagsets and tagging. Textual markup systems are primarily intended to facilitate formatting and retrieval. However, through a combination of their expression in rich-prospect interfaces and the provision of related tools for manipulating the display, the tagsets and tagging can be brought into the service of helping the reader understand the structure and contents of a collection, as well as providing a method for engaging in tasks related to examining a generalized area of research interest. The understanding made available through a rich-prospect interface can also have implications in terms of improved opportunities for the reader to carry out tasks related to formatting and retrieval.

However, rich-prospect interfaces may not be appropriate for all kinds of collections or for all kinds of data. With respect to the Orlando Project, for example, rich-prospect strategies may not be viable for simultaneously displaying the tagging across all the documents in the collection, since the numbers of tags are several orders of magnitude greater than the numbers of documents. On the other hand, in terms of the collection contents and tagsets, there are many potential advantages to the reader in having access to a rich-prospect interface and related tools, both in the availability of new affordances, and in the provision of new perceptual opportunities.

CHAPTER 5: SUMMARY AND CONCLUSIONS

The primary goal of this dissertation was to strengthen the theoretical basis for further research into the development and use of rich-prospect interfaces (that is, interfaces where some meaningful representation of every item in a collection is an intrinsic part of the interface). There were also three secondary goals, namely to examine some of the details of applying rich-prospect principles to computer interfaces, and in particular to interpretively-tagged text collections; to consider some methods for evaluating the new affordances made possible by rich-prospect interfaces; and to suggest some strategies designers might use in carrying out the design of rich-prospect interfaces.

STRENGTHENING THE THEORETICAL GROUNDS FOR RICH-PROSPECT INTERFACES

The process followed was to draw on the intersection between evaluation of landscape painting and habitat theory, as formulated by Appleton (1975), and to examine the implications of Appleton's ideas for computer interfaces from the perspective of J. Gibson's ecological approach to visual perception (1979).

J. Gibson suggests that people are able to directly perceive opportunities for action in the environment. Appleton's idea is that people have a predilection for being able to obtain prospect on a landscape. If both theories are correct, then people who are able to obtain prospect should also be able to directly perceive at least some of the opportunities for action that prospect makes available, although it is also understood that perception and adoption of affordances in general requires some degree of prior learning.

There are also differences to be considered between various kinds of opportunity for action. Some actions are sequential, as for example when a person leaves home to go out and buy a newspaper. Some actions are nested, as when someone grasps and turns a doorknob in order to open the door. Most actions need to be learned, as does the ability to perceive that they are possible, and there are significant differences in learning based on culture, interpersonal factors, and individual characteristics of the learner such as capacity, previous experience, and so on. The literature on affordances includes discussion and debate of over a dozen such topic areas, ranging from the ontological status of affordances to the nuances of intention in use, all within the context of either the natural environment or the built one, where the creation of new affordances is part of the reflexive cycle of affordance and perception that is intrinsic to human culture and development.

Situated within this larger framework, the creation and learning of new opportunities for action in the digital environment does not mark a dramatic change in human behaviour. If some of

these affordances relate to existing perceptual predilections in people, then those affordances should have the advantage of being built on strengths that have been long established.

APPLYING RICH PROSPECT TO COMPUTER INTERFACES

Functions that are already available through interfaces with no prospect can also be available in interfaces with rich prospect. These functions include various forms of searching, either through simple text string comparisons or else through more sophisticated information retrieval algorithms that involve stemming, indexing, semantic clustering, and so on.

In addition, rich-prospect interfaces make possible several new opportunities for action. The new affordances discussed in this dissertation include those for manipulating the rich-prospect display of meaningful representations of content items through zooming, panning, sorting, selecting, grouping, subsetting, renaming, annotating, opening, and structuring the items. Various technologies have been designed over the years by researchers and developers interested in facilitating each of these functions, although not always with respect to interfaces that could be strictly called rich prospect. The review of this literature on visualization technologies yields strategies ranging from fisheye menus, which can be used to scan over areas of microtext (Bederson 2000), to the PhotoFinder toolkit, which provides the user with a wall of tiny photos and related utilities, as the interface to a digital photo archive (Shneiderman et al. 2002).

There are also several perceptual features that do not represent opportunities for action *per se*, but which are nonetheless of potential significance to users. The perceptual features that have been discussed in this dissertation in terms of rich-prospect interfaces are those that permit direct insight into contents, structure, context, features, limitations, connections, trends, anomalies, navigation, reminders, reassurance, and a reduced sense of helplessness.

APPLYING RICH PROSPECT TO INTERPRETIVELY-TAGGED TEXT COLLECTIONS

While there are a variety of new potential affordances and perceptual opportunities provided by rich-prospect interfaces to digital collections, the degree of complexity increases when the principle of providing prospect is applied to interpretively-tagged text collections. A collection with an interpretive level of tagging is one where information is included in the tags that is otherwise not available in the text that the tags are marking.

In order to provide rich prospect on a tagged collection, it is necessary to consider not only the display of the contents of the collection, but also the display of the tags, tag attributes, and the

values contained in the attributes. Since one potential use of tagged text is to allow the user to extract relevant sections of documents, it may also be necessary to consider some means of providing prospect on segments of documents, rather than treating each document as a single entity.

Each of these components of the interpretively-tagged text collection may lend itself to more than one strategy for providing prospect. It may be useful in some instances, for example, to provide a rich-prospect form of display of the tagset itself, independent of the way in which it has been applied in the documents. Display of the tagset may provide perceptual features that give insight into the nature of the collection and how it has been understood by the people who developed it. It may also provide opportunities for action, by allowing the user to manipulate the display in various ways, or by using components from the display in the construction of queries on the collection.

Rich prospect on the tagging of a collection, on the other hand, may turn out in many cases to be above the level of manageable size, simply by requiring the display of too many items. Strategies may therefore need to be adopted to provide other forms of prospect, involving subsets of the collection, extracted portions of multiple documents, or the display of the tagging as it has been applied in individual documents that have been opened for reading.

EVALUATING NEW AFFORDANCES

The primary difficulty in the evaluation of new affordances is to avoid committing a category error: to keep from comparing apples and oranges. By definition, new affordances are opportunities for action that were not previously available. In order to compare an interface that offers new affordances to an interface without them, it is therefore necessary to first determine whether the new affordances are of interest or potential benefit to the users of a particular collection. It is then necessary to determine the degree to which the new affordances are valuable as they have been implemented in the interface.

The following affordance strength vector space contains factors that have been singled out as being potentially relevant in the discussion of the relative merits of various affordances. Each of the factors deals with the relation between the person and the object in a particular environment. By associating numeric ratings with the different factors, it is possible to arrive at an affordance strength vector number that can be used as an indication of affordance strength. If users also provide comments related to each of the factors in the vector space, it may be possible to discover details that could not be captured by a simple numerical rating. The vector space is as follows:

Affordance strength = (tacit capacity, situated potential, awareness, ability, motivation, preference, contextual support, agential support)

Tacit capacity is the degree to which the object can provide the affordance in general. A wrench, for example, has no tacit capacity to serve as an umbrella.

Situated potential is the degree to which the object can provide the affordance under the given circumstances. An umbrella in general has a high tacit capacity to stop rain, but if no umbrella is to hand then the situated potential is zero.

Awareness represents the degree to which a person is conscious of an affordance. One person may have an umbrella in hand, while another merely wonders if there is one in the house somewhere.

Ability represents the degree to which a person is able to make use of a tacit affordance. A child may know that an umbrella would help keep off the rain, but be uncertain how to open one.

Motivation is a complex factor that includes a wide range of subfactors, which together establish the degree to which a person is interested in making use of a potential affordance.

Preference is another complex factor that is distinct from ability and motivation, and yet can play a pivotal role in the choice of whether or not to adopt an affordance.

Contextual support summarizes all the environmental factors that are not properly attributable to the direct relationship between the perceiver and the affordance, and yet nonetheless are significant. These might range from lighting conditions to the direction the wind is blowing.

Agential support is the degree to which the presence or behaviour of other people or agents in the environment may influence the actions of the perceiver.

THE DESIGN OF RICH-PROSPECT INTERFACES

The design process for rich-prospect interfaces will involve some mandatory activities, including the need to establish some appropriate means of representing every item in the collection and to determine in what ways the contents or tagging of the collection lend themselves to the provision of various methods for manipulating the display through sorting, grouping, subsetting and so on. It may also be useful to establish which of the potential new affordances are of particular significance for a given set of users of a collection, whether through applying the factors in the affordance strength vector or through some other means.

In addition, it should be noted that not all collections are going to be amenable to rich-prospect display, since some may contain too many items, or the items may not be homogeneous

enough for there to be a single means of representing them. In these latter cases it may be possible to identify more complex forms of representation that combine tags, attributes, or contents.

Examination of the tagset and tagging in the Orlando Project indicates that even the most straightforward approaches to providing some meaningful representation of every item in the collection can quickly result in a number of complexities. Orlando is an integrated history of women's writing, and therefore makes extensive use of the <Name> tag and the tag for dated text:

<ChronStruct>. However even a tag as seemingly straightforward as <Name> can prove complicated in practice. In Orlando, for example, not every name in the collection is the name of an author, and not every reference to an author who is represented in the collection appears as the text of a given <Name> tag – in some cases, the tagged text contains only an indirect or oblique reference to the person, and the contents of the <Name: standard> attribute are essential. There are also anonymous authors, authors with names that are the same as other authors, and authors with pseudonyms. For people mentioned but not represented in the collection, there are those whose names appear in relation to only one author who is represented, those who appear in relation to multiple authors, and those who were important historical figures in their own right, involved in activities that may appear as part of the Events database.

These details complicate the use of names as a form of meaningful representation, and need to be taken into consideration by the designer hoping to provide a rich-prospect display for Orlando based on the <Name> tag and its attributes. A number of approaches to providing a rich-prospect display based on the <Name> tag are possible, including (in increasing order of sophistication) picklists, microtexts, walls of text, and panoramas:

The <ChronStruct> tag is more complex still, since it involves several mandatory subtags, including <Date> or <DateRange>. All three of these tags have attributes whose values are significant with respect to using <ChronStruct> information in a rich-prospect display of chronological materials. These include attributes to indicate the scope of historical relevance, the degree of certainty with which the date has been supplied, and the domain of the material with respect to the primary intention of the collection – namely whether or not the material is about a British woman writer. Dates may also be more or less complete, leading to the necessity for decisions as to how partial dates should be sequenced in chronologies. Methods of providing prospect on these chronological materials include the use of timelines, scattergrams, sequential prospect, and rich prospect. One criterion which

the designer might apply in selecting among these methods is the extent to which they provide relevant new affordances and other perceptual advantages to the reader.

CONCLUSIONS

Rich-prospect forms of interface, where some meaningful representation of every document or other relevant dimension of the collection is an intrinsic part of the interface, have the potential to provide the user with a number of new perceptual experiences and new opportunities for action involving the displayed items. Having prospect on a collection may relate to a human predilection for having prospect on a landscape, in which case there may exist an association between seeing an overview and understanding some of the advantages it could provide. If the designer of the interface can facilitate this understanding and these advantages through the provision of appropriate tools for the user to apply in working with the display, then the benefits of having prospect on a collection may outweigh the complexities of having many items showing in the interface, at least for some kinds of collections.

The degree to which the new opportunities for perception and action weigh against the potentially intimidating numbers of items in the display is going to be subject to a number of factors related to the nature of the material in the collection and the characteristics of the user. However, common analog artifacts such as maps, phone books, dictionaries, and encyclopedias lend support to the belief that, given the right conditions, people are able to manage large amounts of information. Interpretively-tagged text collections such as the Orlando Project's integrated history of women's writing in the British Isles are strong candidates for this kind of interface research, since the sophistication of the tagged material and the potentially complex requirements of the academic users may necessitate provision in some form of the kinds of affordances that can be made possible through rich-prospect interfaces. Additional research based on user studies of Orlando material, or of other interpretively-tagged text collections, assuming that each of these could be made available through various kinds of interfaces, would be a useful next phase of research.

CHAPTER 6: FURTHER RESEARCH

The areas of possible further research identified in this dissertation have been subdivided according to the chapter to which they refer. The introduction to each section below also suggests where these projects may relate to some of the objectives of the dissertation. Although many of the projects relate to more than one of the objectives, in many cases only the most significant objective is listed. However, for those research topics where more than one objective is of primary importance, more than one objective may be shown. In summary, the objectives of this dissertation were:

- Strengthening the Theoretical Grounds for Rich-Prospect Interfaces
- Applying Rich Prospect to Computer Interfaces
- Applying Rich Prospect to Interpretively-Tagged Text Collections
- Evaluating New Affordances
- Strategies for the Design of Rich-Prospect Interfaces

FURTHER RESEARCH: DIGITAL COLLECTIONS

A number of areas have been suggested throughout the discussion of digital collections where further research is required. These areas include:

Research Topic	Objective
Empirical Determination of Pi Numbers for Rich-Prospect Displays	<ul style="list-style-type: none"> • Applying Rich Prospect to Computer Interfaces • Applying Rich Prospect to Interpretively-Tagged Text Collections
Evaluation of Affordance Strength for the Affordances Involving Prospect	<ul style="list-style-type: none"> • Evaluating New Affordances
Alternative Meaningful Representations in Rich-Prospect Interfaces	<ul style="list-style-type: none"> • Strategies for the Design of Rich-Prospect Interfaces • Evaluating New Affordances
Effects of Non-Persistent vs. Persistent Display on Perception of Prospect	<ul style="list-style-type: none"> • Strategies for the Design of Rich-Prospect Interfaces • Evaluating New Affordances

Structural Priming in Rich-Prospect Interfaces	<ul style="list-style-type: none"> • Strategies for the Design of Rich-Prospect Interfaces
Effects of Sequential vs. Spatial Prospect	<ul style="list-style-type: none"> • Strategies for the Design of Rich-Prospect Interfaces

Empirical Determination of Pi Numbers for Rich-Prospect Displays

The amount of information that can be comprehended in some meaningful way by a user of a rich-prospect interface is related to a number of factors such as visual acuity, previous experience, confidence, motivation, and so on. However, within the parameters of a given user, collection, interface, and task, it should be possible to determine a pi number for information display – that is, a point at which a given strategy has put too much information in front of the user at once, and the user experiences a sense of information overload. Various strategies to reduce this sense of overload could then be developed and tested. These strategies might include methods of sorting, selecting, grouping, subsetting and so on as applied to the rich-prospect display. Each of the resulting variations would then need to be evaluated independently. The result should be a list of strategies that will allow designers to manipulate large displays of information in ways that make them easier for the user to accept.

Evaluation of Affordance Strength for the Affordances Involving Prospect

Each of the new affordances identified above can be evaluated using the vector space model suggested in Chapter 1. To be meaningful, this kind of evaluation needs to take place within the constraints of a given user community using an interface to a particular collection. The affordance strength vector has eight factors, as follows:

Affordance strength = (tacit capacity, situated potential, awareness, ability, motivation, preference, contextual support, agential support)

The affordances that involve prospect are related to insights in the following areas:

- contents
- structure
- context
- features

- limitations
- connections
- trends
- anomalies
- navigation
- reminders
- reassurance
- reduced helplessness

In order to evaluate whether self-reporting varies from the reporting of observers, evaluators from both groups should be involved.

Alternative Meaningful Representations in Rich-Prospect Interfaces

If the user has the opportunity to choose the representation used in a rich-prospect interface, within the constraints of the use of a particular collection by a given set of users, it may be possible to identify patterns of preference. Records would need to be kept of user selection of interface elements. A related study might examine the relationship between choice of tools for manipulating the display and the kind of representation chosen.

Effects of Non-Persistent vs. Persistent Display on Perception of Prospect

Affordance strength vectors could be created for the different means of treating the meaningful representation of items once the user has begun to manipulate the display. The system might respond, for instance, by visually modifying items that are not currently selected in any of the following ways:

- changing some visual feature (such as colour or intensity)
- grouping and moving them to the side
- collapsing them into an icon at the bottom of the screen.

Each of these strategies could be examined with respect to the current user task, such as searching, sorting, grouping, and so on.

Structural Priming in Rich-Prospect Interfaces

In order to emphasize structural features in a rich-prospect display, it would be possible to load the interface in two stages – the first stage emphasizing the structure, and the second stage filling in the

contents. The question is whether this strategy provides the user with any demonstrable benefits over simply loading the entire interface at once.

Effects of Sequential vs. Spatial Prospect

It is possible to provide prospect on a collection in a spatial form – that is, with all the meaningful representations of the collection elements displayed at one time. An alternative form of display is sequential, with the items either appearing in a fixed location one after another or else scrolling past the user on a marquee. This study would look at the three strategies for providing prospect in terms of their perception by the users and their possible effects on selection tasks. Various tools to facilitate use of each kind of display would also need to be considered. It may be possible to provide some insight by having users and observers create affordance strength vectors for each kind of display.

FURTHER RESEARCH: TEXTUAL MARKUP

Rich-prospect interfaces that include not only the content of the collection but also the tagset represent an area of research that has not yet been well explored in the literature. A similar statement can be made about research into rich-prospect interfaces that include the tagging as it has been applied in the documents. Possible research areas include the following:

Research Topic	Objective
Effect on Pi Numbers of Multiple Simultaneous Prospect Views	<ul style="list-style-type: none"> Evaluating New Affordances
Evaluation of Affordance Strength for Prospect on the Tagset and the Tagging	<ul style="list-style-type: none"> Evaluating New Affordances

Effect on Pi Numbers of Multiple Simultaneous Prospect Views

Once research results are available for the amount of information that is manageable within the constraints of a given user, collection, and display (see Chapter 2: the Digital Affordances of Prospect), the next step is to examine the way in which additional information about the tagset and the application of the tagset affect the perception of what is visually too much or too complex.

Evaluation of Affordance Strength for Prospect on the Tagset and the Tagging

Just as it is possible to examine the affordance strength for new affordances related to displays of the content of a collection, it is also possible to have people using the collection and observers of those people independently evaluate the strength of new affordances related to prospect on the tagset and the tagging. Each of the possible advantages of having prospect on the tagset or the tagging should be evaluated, either independently or as composites, depending on which form is most appropriate.

Finally, it may also be worthwhile to examine partial prospect on the tagging, in the form of displays that show only selected tags rather than the entire tagset.

FURTHER RESEARCH: THE ORLANDO PROJECT

The following areas of interest have been identified as possible topics for further research on the use of prospect in the interfaces to the Orlando collection and other collections that have been tagged at an interpretive level:

Research Topic	Objective
Combining Document Types in Heterogeneous Displays	<ul style="list-style-type: none"> • Applying Rich Prospect to Interpretively-Tagged Text Collections
Comparing Results of Automatically-Generated Event Labels	<ul style="list-style-type: none"> • Evaluating New Affordances
Document Access and its Correlation with Interaction Histories	<ul style="list-style-type: none"> • Evaluating New Affordances
User Research on Prospect-Based Interfaces	<ul style="list-style-type: none"> • Applying Rich Prospect to Interpretively-Tagged Text Collections • Evaluating New Affordances

Combining Document Types in Heterogeneous Displays

The Orlando collection contains three distinct types of documents: writing histories, biographies, and events. The former two document types can be represented by author name, since a single author is the subject of both documents. Events, however, are more difficult to represent, since they can refer to a

wide range of topics rather than strictly to individual writers. The tagging on the project also affords extraction and display of segments of biocritical documents which may not be related exclusively to authors. Is it useful to construct prospect displays of these materials, where the representations of the documents might not be homogeneous? How might such displays be structured in order to make them most useful to a particular user engaged in a specified task?

Comparing Results of Automatically-Generated Event Labels

A prospect list for events requires some means of representing the events in a consistent, brief manner. One possibility is to provide three components: a date, a keyword or short phrase based on text from one of the other core tags, and a second keyword or phrase based on the name of a tag present in the event. Different algorithms would result in different labels for the events, and some labels are going to be more accurate representations than others will be. It may therefore be useful to identify several different strategies for extracting event labels from existing tagging and compare the results manually to determine which if any is the most accurate. The results of this study could then be applied in the definitions of future tagsets for collections that have items resembling the Orlando events.

Document Access and its Correlation with Interaction Histories

One way to study the effects of automatic creation of interaction histories would be to run two versions of the project interface in parallel. One interface would contain the interaction history groupings and the other interface would not. The system could then record document access from each kind of interface, providing some indication of whether subsequent users were more likely to access authors when their names appeared on the related list. If the system also contained user profiles through a password or subscription system, it may also be possible to identify common characteristics of the people who tended to make use of this feature. Future designs for those people could then be carried out in such a way as to include similar functions.

User Research on Prospect-Based Interfaces

The following topics are all areas for further investigation of the relationship between the users and the Orlando collection, as mediated by the details of particular interfaces.

- query formulation comparison: how many queries involve tags, attributes, and attribute values? How many involve explicit identification of nested tags?

- learning to use new affordances of prospect-based browsing tools. How quickly did people learn? How eager were they to learn? Is there a difference between people in the domain and not in the domain?
- empirical studies of pi. At what point are there too many items showing? How does this relate to design, screen size, visual acuity, previous experience?
- comparisons of prospect methods
 - ◊ words vs. icons
 - ◊ depth cues vs. flat displays
 - ◊ zooming panoramas vs. panoramas at a fixed size
 - ◊ all items showing vs. a subset
 - ◊ all items showing vs. a hierarchy
 - ◊ all items showing vs. a recursive hierarchy
 - ◊ all the hierarchy showing vs. a subset of the hierarchy
 - ◊ horizontal vs. vertical scrolling
 - ◊ clustering vs. ER-style diagrams
 - ◊ effects of varying the meaningful representation of items dynamically
 - ◊ effects of immediate prospect (entry screen) vs. delayed prospect (some subsequent screen)

BIBLIOGRAPHY

- Abiteboul, S. (1999). On Views and XML. *ACM SIGMOD Record*. 28.4:30-38. NY: ACM Press.
<http://doi.acm.org/10.1145/344816.344853>
- Ahlberg, C. and Shneiderman, B. (1994). The alphaslider: a compact and rapid selector. *Conference proceedings on Human factors in computing systems: "celebrating interdependence."* pp. 365-371. <http://doi.acm.org/10.1145/191666.191790>
- Appleton, J. (1975). *The Experience of Landscape*. London: John Wiley and Sons.
- Appleton, J. (1988). Prospects and refuges revisited. In *Environmental aesthetics: Theory, Research, and Applications*. Cambridge: Cambridge University Press. pp. 27-44.
- Baldonado, M.Q.W., Woodruff, A, and Kuchinsky, A. (2000). Guidelines for Using Multiple Views in Information Visualization. *Proceedings of the Working Conference on Advanced Visual Interfaces*. pp. 110-9. <http://doi.acm.org/10.1145/345513.345271>
- Baars, Bernard J. (1997). *In the Theatre of Consciousness: The Workspace of the Mind*. NY & Oxford: Oxford University Press.
- Baudisch, P. (nd). Don't Click – Paint! Applying the Painting Metaphor to Query Interfaces and Personalization. <http://ipsi.fhg.de/~baudisch/publications/1999-Baudisch-HCI99review-DontClickPaint.pdf>
- Bederson, B. B. (2000). Fisheye menus. *Proceedings of the 13th annual ACM symposium on user interface software and technology*. pp. 217-225. <http://doi.acm.org/10.1145/354401.354782>
- Bederson, B. B. (2001). PhotoMesa: a zoomable image browser using quantum treemaps and bubblemaps. *Proceedings of the 14th annual ACM symposium on User interface software and technology*. pp. 71-80. <http://doi.acm.org/10.1145/502348.502359>
- Bewley, W. L., Roberts, T. L., Schroit, D., and Verplank, W. L. (1983). Human factors testing in the design of Xerox's 8010 "Star" office workstation. *Proceedings of 1983 Conference on Human Factors in Computing Systems*. pp. 72-77.
- Bingham, G.P. (2000). Events (like objects) are things, can have affordance properties, and can be perceived. *Ecological Psychology*. 12.1:29–36.
- Boschker, M. S.J., Bakker, F.C. and Michaels, C.F. (2002). Memory for the functional characteristics of climbing walls: Perceiving affordances. *Journal of Motor Behavior*. 34.1:25–36.
- Bradley, J. and Rockwell, G. (1994). What Scientific Visualization Teaches Us about Text Analysis. Presented at the ALLC/ACH conference. Paris.
<http://www.humanities.mcmaster.ca/~grockwel/ictpaper/paris.htm>

- Broughton, V. (1998). "The revision process in UDC: An examination of the systematic auxiliary of 'Point-of-View' using facet analytical methods". Paper given at the IFLA General Conference Amsterdam. <http://www.ifla.org/IV/ifla64/103-158e.htm>
- Brown, S., Clements, P., Grundy, I., Butler, T., Hockey, S., Fisher, S., Carter, K., Harvey, K. and Wood, J. (1997). Tag Team: Computing, Collaborators, and the History of Women's Writing in the British Isles. Plenary at the annual meeting of the Association of Canadian College and University Teachers of English. <http://www.ualberta.ca/ORLANDO/o-ACUT97.htm>
- Brown, S., Clements, P., Grundy, I., Hockey, S., Butler, T., Fisher, S., Binhammer, K., and Wood, J. (1997). The Orlando Project: Building Digital Resources for an Integrated History of Women's Writing in the British Isles. Digital Resources in the Humanities Conference. <http://www.ualberta.ca/ORLANDO/o-DRH97.htm>
- Brunson, L., Kuo, F., and Sullivan, W. C. (2001). Resident appropriation of defensible space in public housing: implications for safety and community. *Environment and Behavior*. 33.5:626-652.
- Burnard, L. (1995). What is SGML and How Does It Help? *Computers and the Humanities*. 29:41-50. <http://www.tei-c.org.uk/Vault/ED/EDW25/>
- Burnard, L. (1996). SGML on the Web: Too Little Too Soon, or Too Much Too Late? Presented at "A Decade of Power", the third annual conference of the Belgium-Luxembourgian SGML Users Group. Brussels 30-31 October 1996. <http://users.ox.ac.uk/~lou/Belux/>
- Bush, V. (1945). As we may think. *The Atlantic Monthly*. 176.1:101-8. <http://www.theatlantic.com/unbound/flashbks/computer/bushf.htm>
- Butler, T. (1998). Eluding the Chains of Technology: Computer Systems for Orlando. Fisher, S. (1997). SGML and the Orlando Project: Theorizing Descriptive Markup. In Hockey, S., Butler, T., Brown, S., and Fisher, S. The Orlando Project: Humanities Computing in Conversation with Literary History. ACH Session. <http://www.ualberta.ca/ORLANDO/o-ACH97.htm>

- Butler, T., Fisher, S., Hockey, S., Coulombe, G., Clements, P., Brown, S., Grundy, I., Carter, K., Harvey, K., and Wood, J. (1999). Abstract: Can a Team Tag Consistently? Experiences on the Orlando Project. ACH-ALLC 1999 Joint International Conference, Charlottesville, VA.
<http://www.ualberta.ca/ORLANDO/o-ACH99abs.html>
- Cameron, R. (2003). Unpublished Orlando Project internal correspondence (e-mail).
- Chemero, A. (2000). What events are. *Ecological Psychology*. 12.1:37–42.
- Chemero, A. (2001). What we perceive when we perceive affordances: Commentary on Michaels (2000) “Information, perception, and action” *Ecological Psychology*. 13.2:116–6.
- CiteSeer (2002). NEC Research Institute. <http://citeseer.nj.nec.com/>
- Clements, P. (Director), Brown, S., Grundy, I., Elio, R., Cameron, R., and Wallace, J. (2003). The Orlando Project: An Integrated History of Women’s Writing in the British Isles. Electronic Textbase. Publication pending.
- Clements, P. (Director), Brown, S., Grundy, I., Elio, R., Cameron, R., and Wallace, J. (2003a). The Orlando Project: An Integrated History of Women’s Writing in the British Isles. Unpublished internal project documentation.
- Coombs, J. H., Renear, A. H., and DeRose, S. J. (1987). Markup Systems and the Future of Scholarly Text Processing. *Communications of the ACM*. 30.11:933-47.
<http://doi.acm.org/10.1145/32206.32209>
- Czerwinski, M., Dumais, S., Robertson, G., Dziadosz, S., Tiernan, S., and van Dantzich, M. (1999). Visualizing implicit queries for information management and retrieval. *Proceeding of the CHI 99 conference on Human factors in computing systems: the CHI is the limit*. pp. 560-7.
<http://doi.acm.org/10.1145/302979.303158>
- Dixon, P. and O’Reilly, T. (1999). Appearance, Form, and the Retrieval of Prior Knowledge. *Proceedings of Design and the Social Sciences*. Edmonton.
- Drucker, J. and Nowviskie, B. (2003). The Temporal Modelling Project.
<http://jefferson.village.virginia.edu/time/project/index.html> and
<http://jefferson.village.virginia.edu/time/storyboard/>
- Dublin Core. (2002). Dublin Core Metadata Initiative.
<http://dublincore.org/documents/1999/07/02/dces/>

- Eick, S. G. (1994). Data Visualization Sliders. *Proceedings of the 7th annual ACM symposium on User interface software and technology*. pp. 119-120.
<http://doi.acm.org/10.1145/192426.192472>
- Ellis, R. R., Flanagan, J.R. and Lederman, S. J. (1999). The influence of visual illusions on grasp position. *Experimental Brain Research*. 125. 109–114.
- Feldman, R., and Hirsh, H. (1997). Finding Associations in Collections of Text. In Michalski, R.S., Bratko, I. and Kubat, M. *Machine Learning and Data Mining: Methods and Applications*. NY: J. Wiley. pp. 223-240.
- Fisher, S. (1997). SGML and the Orlando Project: Theorizing Descriptive Markup. In Hockey, S., Butler, T., Brown, S., and Fisher, S. *The Orlando Project: Humanities Computing in Conversation with Literary History*. ACH Session. <http://www.ualberta.ca/ORLANDO/o-ACH97.htm>
- Fleming, J. (1998). *Web Navigation: Designing the User Experience*. Sebastopol, CA: O'Reilly and Associates, Inc.
- Fowler, R. H., Fowler, W. A. L., Wilson, B. A. (1991). Integrating Query, Thesaurus, and Documents through a Common Visual Representation. *Proceedings of the 14th annual international ACM SIGIR conference on Research and development in information retrieval*. NY: ACM Press. pp. 142-151. <http://doi.acm.org/10.1145/122860.122874>
- Frascara, J. (1984). *ISO Technical Report 7239, Development and Application Principles for Public Information Symbols*. Geneva: ISO.
- Frascara, J. (1997). *User-Centred Graphic Design: Mass Communications and Social Change*. London: Taylor & Francis.
- Frascara, J. (2001). Diagramming as a Way of Thinking Ecologically. *Visible Language*. 35.2:164-177.
- Fraser, B., Roberts, J., Pianosi, G., Alencar, P., Cowan, D., German, D., and Nova, L. (1999). Dynamic Views of SGML Tagged Documents. *Proceedings of the 17th annual international conference on Computer documentation*. New Orleans, Louisiana, United States.
<http://doi.acm.org/10.1145/318372.318564>
- Freeman, E. and Gelernter, D. (1996). Lifestreams: A Storage Model for Personal Data. *SIGMOD Record*. 25.1:80-86. <http://doi.acm.org/10.1145/381854.381893>

- Gershon, N., Eick, S. G., and Card, S. (1998). Information Visualization. *ACM Interactions* 5.2:9-15. <http://doi.acm.org/10.1145/274430.274432>
- Gibson, E. J. (2000). Where is the information for affordances? *Ecological Psychology*. 12.1:53–56.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton-Mifflin.
- Gibson, J. J. (1979u). A Study in the Psychology of Decorative Art. In *The Purple Perils*.
Unpublished manuscript. <http://www.ksu.edu/psych/farris/gibson/files/decorative.html>
- Goodale, A. D. and Milner, M. A. (1995). *The Visual Brain in Action*. Oxford: Oxford University Press.
- Gould, J. D. and Lewis, C. (1985). Designing for Usability: Key Principles and What Designers Think. *Communications of the ACM*. 28.3: 300-311. <http://doi.acm.org/10.1145/3166.3170>
- Grundy, I., with Clements, P., Balasz, S., Brown, S., Rebecca Cameron, R. Carter, K., Elio, R., and Gomboc, D. (2001). Delivering Childbirth: Orlando Project Encoding. Presented at the British Eighteenth- and Nineteenth- Century Women Writers conference, Lawrence, Kansas. <http://www.ualberta.ca/ORLANDO/Childbirth.htm>
- Grundy, I., Clements, P., Brown, S. Butler, T., Cameron, R., Coulombe, G., Fisher, S., and Wood, J. (2000). Dates and ChronStructs: Dynamic Chronology in the Orlando Project. *Literary and Linguistic Computing*. 15:3 (2000), 265-289.
<http://www.ualberta.ca/ORLANDO/Chronolpaper.htm>
- Hagerhall, C. M. (2001). Consensus in landscape preference judgements. *Journal of Environmental Psychology*. 21:83–92.
- Harman, D. (1993). Overview of the first TREC conference. *Proceedings of the 16th annual international ACM SIGIR conference on Research and development in information retrieval*. <http://doi.acm.org/10.1145/160688.160692>. pp. 36-47.
- Harrison, B. L., Ishii, H., Vicente, K. J., and Buxton, W. A. S. (1995). Transparent Layered User Interfaces: An Evaluation of a Display Design to Enhance Focused and Divided Attention. *Proceedings of the 1995 conference on Computer-Human Interaction*.
http://www.acm.org/sigchi/chi95/Electronic/documnts/papers/blh_bdy.htm
- Hascoët, M. and Soinard, X. (1998). Using maps as a user interface to a digital library. *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*. pp. 339-340. <http://doi.acm.org/10.1145/290941.291028>
- Hecht, H. (2000). Are events and affordances commensurate terms? *Ecological Psychology*. 12.1:57–63.

- Heft, H. (1989). Affordances and the body: an intentional analysis of Gibson's ecological approach to visual perception. *Journal for the Theory of Social Behaviour*. 19.1:1-30.
- Hill, W. C., Hollan, J. D., Wroblewski, D., McCandless, T. (1992). Edit Wear and Read Wear. *Conference proceedings on Human factors in computing systems*. pp. 3-9.
<http://doi.acm.org/10.1145/142750.142751>
- Hill, L. L., Carver, L., Larsgaard, M., Dolin, R., Smith, T. R., Frew, J., and Rae, M. A. (2000). Alexandria digital library: user evaluation studies and system design. *Journal of the American Society for Information Science*. 51.3: 246-259. <http://www3.interscience.wiley.com/cgi-bin/issuetoc?ID=69501275>
- Hockey, S. (1999). Is There a Computer in this Class? Presented at "Is Humanities Computing an Academic Discipline?" An Interdisciplinary Seminar. Institute for Advanced Technology in the Humanities. Virginia. <http://www.iath.virginia.edu/hcs/hockey.html>
- Hockey, S. (2000). *Electronic Texts in the Humanities*. Oxford: Oxford University Press.
- Höge, H. (1990). Ecological perception and aesthetics: pictures are affordance-free. In Klaus Landwehr (Ed.). *Ecological Perception Research, Visual Communication, and Aesthetics*. NY: Springer-Verlag.
- Hommel, B. (1993). Inverting the Simon effect by intention: determinants of direction and extent of effects of irrelevant spatial information. *Psychological Research*. 55:270-279.
- Hull, D. (1993). "Using statistical testing in the evaluation of retrieval experiments." *Proceedings of the 16th annual international ACM SIGIR conference on Research and development in information retrieval*. <http://www.acm.org/pubs/citations/proceedings/ir/160688/p329-hull/>
- Johnson Controls. (2002). Data Visualization: Starfield Display.
<http://www.johnsoncontrols.com/metasys/newproducts/starfielddisplay.htm>
- Jordan, P. (2000). *Designing Pleasurable Products: An Introduction to the New Human Factors*. London: Taylor and Francis.
- Jordan, P. (1999). Personal Interview. Edmonton, Alberta.
- Kahn, P. H., Jr. (1999). *The Human Relationship with Nature: Development and Culture*. Cambridge, Mass.: MIT Press.
- Kaplan, R. and Kaplan, S. (1989). *The Experience of Nature: A Psychological Perspective*. Cambridge: Cambridge University Press.

- Karat, J. (1997). Evolving the Scope of User-Centered Design. *Communications of the ACM*. 40.7:33-38. <http://doi.acm.org/10.1145/256175.256181>
- Kaugars, K. (1998). Integrated multi scale text retrieval visualization. *Proceedings of the conference on CHI 98 summary: human factors in computing systems*. pp. 307-8. <http://doi.acm.org/10.1145/286498.286782>
- Keep, C., McLaughlin, T., and Parmar, R. (2000). The Electronic Labyrinth. University of Virginia. <http://www.iath.virginia.edu/elab/elab.html>
- Koffka, K. (1935). *Principles of Gestalt Psychology*. NY: Harcourt, Brace, and World, Inc.
- Kotchoubey, B. (2000). About hens and eggs – perception and action, ecology and neuroscience: a reply to Michaels (2000). *Ecological Psychology*. 13.2:123–133.
- Kulikowich, J.M. and Young, M.F. (2001). Locating an ecological psychology methodology for situated action. *Journal of the Learning Sciences*. 12:165–202.
- Kuo, F. E. and Sullivan, W. C. (2001). Environment and crime in the inner city: does vegetation reduce crime? *Environment and Behavior*. 33.3:343–367.
- Lakoff, G. (1980). *Metaphors we live by*. Chicago: The University of Chicago Press.
- Lintern, G. (2000). An affordance-based perspective on human-machine interface design. *Ecological Psychology*. 12.1:65–9.
- Liu, Y. H., Dantzig, P., Sachs, M., Corey, J. T., Hinnebusch, M. T., Damashek, M., and Cohen, J. (2000). Visualizing document classification: A search aid for the digital library. *Journal of the American Society for Information Science*. 51.3: 216-227. <http://www3.interscience.wiley.com/cgi-bin/issuetoc?ID=69501275>
- Livingstone, M. (2002). *Vision and Art: the Biology of Seeing*. NY: Harry N. Abrams, Inc.
- Maple, A. (1995). Faceted Access: A Review of the Literature. Presented at the Music Library Association Annual Meeting. <http://www-sul.stanford.edu/depts/music/mlatest/BCC/BCC-Historical/95WGFAM2.html>
- Matthews, P. M. (2001). Aesthetic appreciation of art and nature. *British Journal of Aesthetics*. 41.4:395–410.
- Mauss, M. (1935, 1973). Techniques of the body. In Roy Bailey et al. (Eds.). *Economy and Society*. London: Routledge and Kegan Paul.

- McCarty, W. (1991). Finding Implicit Patterns in Ovid's *Metamorphoses* with TACT. *CCH Working Papers*. v1. <http://www.kcl.ac.uk/humanities/cch/chwp/mccarty/index.html>
- Metadata Encoding and Transmission Standard. (2003). Library of Congress.
<http://www.loc.gov/standards/mets/>
- Michaels, C. F. (2000). Information, perception, and action: what should ecological psychologists learn from Milner and Goodale (1995)? *Ecological Psychology*. 12.3:241–258.
- Milner, A. D. and Goodale, M. A. (1995). *The Visual Brain in Action*. Oxford: Oxford University Press.
- Mitchell, T. C. (1993). *Redefining Designing: From form to experience*. New York: Van Nostrand Reinhold.
- Monmonier, M. (1993). *Mapping It Out: Expository Cartography for the Humanities and Social Sciences*. Chicago: University of Chicago Press.
- Moser, G., Ratiu, E., and Fleury-Bahi, G. (2002). Appropriation and interpersonal relationships: from dwelling to city through the neighborhood. *Environment and Behavior*. 34.1:122–136.
- Mozilla. (2002). Type Ahead Find.
<http://www.mozilla.org/projects/ui/accessibility/typeaheadfind.html>
- Myaeng, S. H., Jang, D. H., Kim, M. S. and Zhoo, A. C. (1998). A Flexible Model for Retrieval of SGML Documents. <http://doi.acm.org/10.1145/290941.290980>
- Nelson, T., Johnson, T., Strong, M. and Rudakewich, G. (2001). Perception of tree canopy. *Journal of Environmental Psychology*. 21.3:315–324.
- Norgard, B. Kim, Y., Buckland, M. Chen, A., and Larson, R. (1999). Advanced Search Technologies for Unfamiliar Metadata (demonstration abstract). Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval.
<http://doi.acm.org/10.1145/312624.312752>
- Norman, D. A. (1990). *The Design of Everyday Things*. NY: Doubleday.
- Norman, D. A. (1993). *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Reading, Massachusetts: Addison-Wesley Publishing Company.
- Ontopia. (2002). Ontopia: Solutions for Managing Knowledge and Information. <http://www.ontopia.net/>

- Pavani, F., Boscagli, I., Benvenuti, F., Rabuffetti, M., and Farne, A. (1999). Are perception and action affected differently by the Titchener circles illusion? *Experimental Brain Research*. 127. 95–101.
- Pickering, J. (2000). On the proper treatment of affordance: formality or mutuality? *Ecological Psychology*. 12.1:71–77.
- Pirolli, P., Schank, P., Hearst, M., and Diehl, C. (1996). Scatter/ Gather Browsing Communicates the Topic Structure of a Very Large Text Collection. *Proceedings of the SIGCHI conference on Human factors in computing systems: common ground*. pp. 213-220.
<http://doi.acm.org/10.1145/238386.238489>
- Post, D. L., Geiselman, E. E. and Goodyear, C. D. (1999). Benefits of Color Coding Weapons Symbology for an Airborne Helmet-Mounted Display. *Human Factors*. 41.4:515-23.
- Price, R. (1998). Beyond SGML. *Proceedings of the third ACM conference on Digital libraries*. 172-181. <http://doi.acm.org/10.1145/276675.276694>
- Prindle, S. S., Carello, C., and Turvey, M. T. (1980). How wrong is Gibson? *Behavioral and Brain Sciences*. 3:394-5.
- Renear, A., Mylonas, E., and Durand, D. (1996). Refining our Notion of What Text Really Is: The Problem of Overlapping Hierarchies. In *Research in Humanities Computing*. N. Ide and S. Hockey (Eds.). Oxford: Oxford University Press.
<http://www.stg.brown.edu/resources/stg/monographs/ohco.html>
- Rogers, E. M. (1983). *Diffusion of Innovations*. 3rd ed. NY: Free Press.
- Roussinov, D., Tolle, K., Ramsey, M., McQuaid, M., and Chen, H. (1999). Visualizing Internet search results with adaptive self-organizing maps. *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*. p. 336. <http://doi.acm.org/10.1145/312624.312773>
- Ruecker, S. (2002). A conceptual taxonomy of SGML tags: based on the tagging in the Orlando Project. *Special issue of the Journal of Applied Systems Studies (JASS) on WEB Information Systems Applications*.
<http://computer.org/proceedings/wise/0577/vol%202/05772002abs.htm>
- Ryle, G. (1949). *The concept of mind*. NY: Harper and Row.
- Sanders, J.T. (1997). An ontology of affordances. *Ecological Psychology*. 9.1:97.

- Sanderson, M. and Croft, B. (1999). Deriving concept hierarchies from text. *Proceedings of the 22nd annual international ACM SIGIR conference on research and development in information retrieval*. pp. 206-213. <http://doi.acm.org/10.1145/312624.312679>
- Shipman III, F. M., Furuta, R., Brenner, D., Chung, C. C., and Hsieh, H. W. (2000). Guided paths through Web-based collections: Design, experiences, and adaptations. *Journal of the American Society for Information Science*. 51.3: 260-272.
<http://www3.interscience.wiley.com/cgi-bin/issuetoc?ID=69501275>
- Shneiderman, B., Kang, H., Kules, B., Plaisant, C., Rose, A., and Rucheir, R. (2002). PhotoFinder Project. <http://www.cs.umd.edu/hcil/photolib/>
- Shneiderman, B. Williamson, C. and Ahlberg, C. (1992). Dynamic Queries: Database Searching by Direct Manipulation. *Proceedings of the SIGCHI conference on human factors in computing systems*. pp. 669-670. <http://doi.acm.org/10.1145/142750.143082>
- Singhal, A, and Kaszkiel, M. (2001). A case study in web search using TREC algorithms. *Proceedings of the tenth international conference on World Wide Web*. pp. 708-716.
<http://doi.acm.org/10.1145/371920.372186>
- Smeets, J. B. J. and Brenner, E. (2001). Perception and action are inseparable. *Ecological Psychology*. 13.2:163-166.
- Smith, B. C. (1996). *On the Origin of Objects*. Cambridge: MIT Press.
- Spamnet. (2002). <http://cloudmark.com/products/spamnet/>.
- Spiteri, L. (1998). A Simplified Model for Facet Analysis. *Canadian Journal of Information and Library Science*. 23:1-30. Reprinted at
http://aifia.org/pg/a_simplified_model_for_facet_analysis.php
- Spotfire. (2002). Spotfire for Academics: Spotfire Decisionsite for Functional Genomics.
<http://www.spotfire.com/academics/products-fg.htm>
- Stoffregen, T. A. (2000). Affordances and events. *Ecological Psychology*. 12.1:1-28.
- Strawson, P. F. (1959). *Individuals*. London: Methuen.
- Stubblefield, W. A. (1998). "Patterns of Change in Design Metaphor: A Case Study." *Conference proceedings on Human factors in computing systems*. Los Angeles, CA USA. 73-80.
<http://www.acm.org/pubs/citations/proceedings/chi/274644/p73-stubblefield/>

- Suchman, L. A. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge: Cambridge University Press.
- TextArc. (2002). TextArc: An alternate way to view a text. <http://www.textarc.org/>
- Turvey, M. T. and Shaw, R. S. (1979). The primacy of perceiving: An ecological reformulation of perception for understanding memory. In L.G. Nilsson (Ed.), *Perspectives on memory research*. Hillsdale, NJ: Erlbaum.
- Ullman, S. (1980). Against direct perception. *Behavioral and Brain Sciences*. 3:373–381.
- Ulrich, R. S. (1984). View from a window may influence recovery from surgery. *Science*. 224:420-1.
- Ulrich, R. S. (1993). Biophilia, biophobia, and natural landscapes. In E. O. Wilson and S. R. Kellert, (Eds). *The Biophilia Hypothesis*. Washington, DC: Island Press. pp. 73–137.
- van der Kamp, J., Savelsbergh, G. J. P. and Rosengren, K. S. (2001). The separation of action and perception and the issue of affordances. *Ecological Psychology*. 2.427:167–172.
- Vanleeuwen, C. and Stins, J. (1994). Perceivable information or – the happy marriage between ecological psychology and gestalt. *Philosophical Psychology*. 2.772:267–285.
- Vicente, K. J. (2002). Ecological Interface Design: Progress and Challenges. *Human Factors*. 44.1:62-78. http://web2.infotrac-custom.com/pdfserve/get_item/1/Sf006d5w4_1/SB389_01.pdf
- Vinson, N. G. (1999). Design guidelines for landmarks to support navigation in virtual environments. *Proceedings of the CHI 99 conference on human factors in computing systems: the CHI is the limit*. <http://www.acm.org/pubs/citations/proceedings/chi/302979/p278-vinson/>
- Vygotsky, L. S. (1978). *Mind in society: the development of higher psychological processes*. Cambridge, MA: Harvard University Press.
- Walker, G. and Atkinson, S. D. (1991). Information Access in the Humanities: Perils and Pitfalls. *Library Hi Tech*. 33.9:1, pp. 23-34.
- Warren, W. H., Jr. (1984). Perceiving affordances: visual guidance of stair climbing. *Journal of Experimental Psychology*. 10.5:683–703.
- Wexelblat, A. and Maes, P. (1999). “Footprints: History-Rich Tools for Information Foraging.” *Proceedings of the CHI 99 conference on Human factors in computing systems: the CHI is the limit*. Pittsburgh, PA USA. pp. 270–277. <http://doi.acm.org/10.1145/302979.303060>

- Whitehouse, S., Varni, J. W., Seid, M., Cooper-Marcus, C., Ensberg, M. J., Jacobs, J. R., and Mehlenbeck, R. S. (2001). Evaluating a children's hospital garden environment: utilization and consumer satisfaction. *Environmental Psychology*. 21: 301–314.
- Williams, K. J. H., and Cary, J. (2002). Landscape preferences, ecological quality, and biodiversity protection. *Environment and Behavior*. 34.2:257–274.
- Wilson, E. O. (1984). *Biophilia*. Cambridge, Mass.: Harvard University Press.
- Winograd, T. and Flores, F. (1986). *Understanding computers and cognition: a new foundation for design*. Norwood, NJ: Ablex Pub. Corp.
- Withagen, R. and Michaels, C. F. (1999). An ecological approach to cognitive (im)penetrability. *Behavioral and Brain Sciences*. 22.3:399–400.
- Wittgenstein Archives at the University of Bergen. (1998). Bergen, Norway: Advanced Computing in the Humanities at the University of Bergen. <http://www.hit.uib.no/hit/ach/wab.html>
- Xerox PARC. (2001). Experiments in the future of reading: reading wall. <http://www2.parc.com/red/projects/xfr/readingwall.html>

APPENDIX A: TECHNICAL CONSIDERATIONS

Most of the solutions discussed in this dissertation could be implemented in some form using contemporary technology. There are a few areas, however, where either the hardware, software, or domain contents would need to be extended to allow one of the strategies for providing rich prospect to be applied. For example, rich prospect for interpretively-tagged text collections requires that the collection be tagged with a level of markup that extends beyond what is required for formatting. Although the number of such collections is growing, the vast majority of electronic archives do not contain textual markup beyond what is available in HTML. Similarly, techniques that involve displaying material on large screens may eventually prove most useful on screens that are larger than those currently available.

A more serious consideration, however, is that the value of rich prospect interfaces lies in their ability to make collections accessible to academic users. Technology for providing designs involving rich prospect therefore needs to be readily and consistently available through web browsers. However, variations in browser capacity across different systems and generations of browsers are still a significant barrier. Panoramas, for example, can currently be implemented in at least three different technologies: flash, shockwave, and javascript. All three of these formats can be read by some browsers, but it is not the case that all browsers support any one of the formats. Even for those browsers which can display one or more of the technologies, it is often necessary for the user to add the capacity by downloading and installing a browser plug-in.

It is therefore not currently possible to deliver a web panorama that can consistently be read by all web browsers, or even by the majority of browsers. Since academics are not necessarily using the most recent equipment, it may also be the case that this subgroup of the user community for electronic collections has less technological capacity than that available to for example, the subgroup formed by design students.

Another limitation is related to network bandwidth, which can constrain deployment of solutions that would be viable on an individual computer, but which are too slow across some of the lower-capacity network access methods.

Given these circumstances, the first release of the Orlando Project has incorporated only those forms of prospect that can be delivered with comparative confidence across a range of browsers and network connection speeds. These include features that are less than optimal in terms of the levels of prospect they provide: features such as picklists, scrolling lists, and hierarchical displays, although there is also one unstructured wall of text (on the home page).

Hopefully, as the installed baseline of technology rises over time, future releases of Orlando will be able to incorporate more sophisticated methods of presenting and manipulating rich-prospect interface features.

Affordances of Prospect
 Stan Ruecker
 Copyright Clearance for Images

The thesis contains the following images. Copyright clearance activity status is indicated on the right.

Figure	Activity	Permission status
1.01	Brain illustrating Milner and Goodale, from Darren Burke. Request sent Sept 22, 2003. Darren wrote Sept 29 to say he took it off Goodale's site. Email request to Goodale Oct 1, Yes with image change Oct 7.	yes with change of image
1.02	Author	yes
1.03	Author	yes
2.01	U of Chicago Press fax. Letter written for Monmonier map 26 Aug 2003.	yes aug 29: grant number 57014 with appropriate text
2.02	Author	yes
2.03	Author	yes
2.04	ACM sent email Sept 15	no reply--blank inserted Oct 15
2.05	ACM sent email Sept 15	no reply--blank inserted Oct 15
2.06	Bederson's Photomesa with the cat. Request with jpeg sent Sept 22. He wrote to say yes Sept 23.	yes
2.07	Starfield display from Johnson Controls. Request sent to on Sept 15, along with letter and jpeg.	no reply--blank inserted Oct 15
2.08	Spotfire Functional Genomics interface. Request sent to n on Sept. 19, along with letter and jpeg. Response Sept 22 from saying the contact person at Spotfire is John Babers. Request to him Sept 22.	no reply--blank inserted Oct 15
2.09	Curious Labs now owns Poser. We need to email their legal department. Email sent Sept 19th.	No reply-blank inserted Oct 15
2.10	ACM sent email Sept 15	No reply—blank inserted Oct 15
2.11	Adobe allows clearance for Premiere screen shots given appropriate notice	yes with appropriate text
2.12	Baudisch's toggle map. Request with jpeg sent to on Sept 22. He wrote later the same day to say very kindly yes.	yes
2.13	Xerox PARC. Request for info with appended letter and images sent Sept 19 to Reply from Mimi.Gardner within minutes, saying she had forwarded it appropriately. Tracy Kugelman wrote Sept 23 to say I should write to Gyricon. I sent an email and jpeg to n on Sept 25.	No reply-blank inserted Oct 15
2.14	TextArc. Request and jpeg sent to Sept 22. Reply Sept 27: yes with change.	yes with change of image

2.15	ISO Jorge's image. Request sent with no jpeg on Sept 22 to the site: http://www.iso.ch/iso/en/xsite/contact/01enquirysevice/014copyright/copyrightForm.html . I got an email from Gisela Helberling on Sept 23 asking for the page and figure number. I sent her those in an email along with the jpeg on Sept 25.	yes with appropriate text
3.01	Request sent to _____ on Sept 19, along with jpeg of the image.	No reply-blank inserted Oct 15
3.02	Author	yes
3.03	Author	yes
3.04	Author	yes
3.05	Author	yes
3.06	Adobe allows clearance for Acrobat screen shots given appropriate notice	yes with appropriate text
4.01	Google fax form downloaded and filled out, image printed, and letter drafted 24 Aug 2003. Faxed to Google by Milena, 26 Aug 2003. Yes by email sept 5	yes
4.02	cleared Orlando 13 March 2003	yes
4.03	cleared Orlando 13 March 2003	yes
4.04	cleared Orlando 13 March 2003	yes
4.05	copyright permission granted in an email from Johanna Drucker 26 Aug 2003	yes
4.06	cleared Orlando 13 March 2003	yes
4.07	cleared Orlando 13 March 2003	yes
4.08	cleared Orlando 13 March 2003	yes
4.09	cleared Orlando 13 March 2003	yes
4.10	cleared Orlando 13 March 2003	yes

no reply yet 8
yes 26 of 34



Mr. Stan Ruecker

Your ref.
Our ref. COPYRIGHT

Date 2003-09-26

Dissertation for the University of Alberta – Request for Copyright

Dear Mr. Ruecker,

With reference to your e-mail of 26 September 2003, we are happy to grant you permission for the reproduction of **ISO 7239:1984, Fig. 6, Development and principles for application of public information symbols** as per the attached request which forms an integral part of this copyright agreement.

This permission is granted on the condition that due credit be given to ISO along the following lines:

“The terms and definitions taken from ISO 7239:1984, Fig. 6, Development and principles for application of public information symbols, are reproduced with the permission of the International Organization for Standardization, ISO. This standard can be obtained from any ISO member and from the Web site of the ISO Central Secretariat at the following address: www.iso.org. Copyright remains with ISO.”

Yours sincerely,

Jacques-Olivier Chabot
Director
General Services and Marketing

The University of Chicago Press

Permission Grant

Date: August 29, 2003
Grant number: 57014

Stan Ruecker

Re: your request dated 8/26/03

Dear Stan Ruecker,

Thank you for your request for permission to use material from the publication(s) of the University of Chicago Press. Permission is granted for use as stated below. Unless specifically granted below, this permission does not allow the use of our material in any other edition or by any additional means of reproduction including (by way of example) motion pictures, sound tapes, electronic formats, and phonograph records; nor does this permission cover book clubs, translations, abridgment, or selections which may be made of the publication. No subsequent use may be made without additional approval. This permission is subject to the following terms:

1. On each copy of the selection full credit must be given to the book or journal, to the author (as well as to the series, editor, or translator, if any), and to the University of Chicago as publisher. In addition, the acknowledgment must include the identical copyright notice as it appears in our publication.
2. This permission does not apply to any part of the selection which is independently copyrighted or which bears a separate source notation. The responsibility for determining the source of the material rests with the prospective publisher of the quoted material.
3. This permission grant for the material listed on the invoice below is provided GRATIS.
4. This permission covers publication of one edition of the work up to **250** copies.
5. Permission granted is non-exclusive and, unless otherwise stated, is valid **throughout the world in the English language only**.
6. This permission is void if:
 - a. More than 10% of your forthcoming work, exclusive of index and bibliography, is composed of University of Chicago Press material.
 - b. publication does not take place within 24 months of the date on this form.
7. This permission includes the use of Braille, large type, or other editions of your work by nonprofit organizations solely for the use by the visually handicapped, provided no fees are charged. The limitations listed in Clause 4, above, do not apply to such use.

Reference

Material Requested

1 fig. (p.216) from M. Monmonier, MAPPING IT OUT: EXPOSITORY
CARTOGRAPHY FOR THE HUMANITIES AND SOCIAL SCIENCES

To appear in a PhD dissertation by Stan Ruecker, to be submitted to the University of Alberta, Dept. of Art and Design.

Approved by: _

_____, Perry Cartwright, Rights and Permissions

WebMail - RE: copyright permission request



Delete

File

New

Reply

Reply All

Forward

Previous

Next

Index

Help

Logout

Date Sent: October 7, 2003 1:32 PM**From:** Mel Goodale >**Add to Address Book****To:** Stan Ruecker**Subject:** RE: copyright permission request**Status:** Urgent New

Hi

Since that image is an amalgam of my own work and that of another illustrator, I would prefer that you use another image. Here is one that does the job.

The brain shown here is a reconstruction of an actual brain rather than the inaccurate drawing shown on the website.

Cheers
Mel Goodale

At 06:13 PM 10/1/2003 -0600, you wrote:

>Dear Dr. Goodale,

>

>I am writing to ask for permission to reproduce an image from your work
>in my interdisciplinary PhD dissertation for the University of Alberta. The
>dissertation is entitled Affordances of Prospect for Academic Users of
>Interpretatively-tagged Text Collections.

>

>I require written copyright permission specifically for the four bound
>copies of the dissertation at the University of Alberta, and to accompany
>the copy that is sent to the National Library of Canada for microfilming.
>If I

>am not able to obtain permission, there is provision to include a textual
>description in place of the image.

>

>The image I am hoping to use is an illustration of the brain, showing the
>ventral and dorsal streams. It is from the top of your web page
><http://www.ssc.uwo.ca/psychology/faculty/goodale/research/>. If you are
>willing to let me use this image, please reply by email to that effect. If I

>am

>able to find a publisher for my dissertation, I will contact you again to ask
>for permission for that further use.

>

WebMail - RE: copyright permission request

Delete File New Reply Reply All Forward Previous Next Index Help Logout

Date Sent: September 26, 2003 4:06 PM

From: W Bradford Paley **Add to Address Book**

To: Stan Ruecker

Subject: RE: copyright permission request

Status: Urgent New

Hi Stan,

You have my permission to use any image on the TextArc.org site, so long as all of the pixels are represented as I put them on the site. You do not have permission to use the attached JPEG.

The JPEG you sent was reduced and almost meaningless since none of the words were visible. I would prefer not to have my work represented in this way. I checked my site, and believe that you got the image from the TextArc Overview PDF--I am sorry for the lousy compression in that file--I have corrected it. Thanks for bringing it to my attention.

Since you are doing bound, print copies, I believe that any of the attached images will work. Please let me know if they won't.

The legend or credit for the image(s) should be:
"TextArc, W. Bradford Paley, <http://textarc.org>"

Good luck with your dissertation.

Best,
Brad

At 07:52 PM 9/22/2003 -0600, you wrote:
>Dear Brad,
>
>I am writing to ask for permission to reproduce an image from your work
>in my interdisciplinary PhD dissertation for the University of Alberta. The
>dissertation is entitled Affordances of Prospect for Academic Users of
>Interpretatively-tagged Text Collections.
>
>I require written copyright permission specifically for the four bound
>copies of the dissertation at the University of Alberta, and to accompany
>the copy that is sent to the National Library of Canada for microfilming.
>If I

WebMail - RE: copyright permission request

Delete	File	New	Reply	Reply All	Forward	Previous	Next	Index	Help	Logout
--------	------	-----	-------	-----------	---------	----------	------	-------	------	--------

Date Sent: September 22, 2003 10:50 PM

From: "Patrick Baudisch" <

Add to Address Book

To: Stan Ruecker <



Cc:



Subject: RE: copyright permission request

Status:

Urgent New












You are so totally welcome to use that picture. Make sure to play with the applets at <http://www.patrickbaudisch.com/projects/paintableinterfaces>

Good luck with your thesis,

Pat

> Dear Patrick,
 >
 > I am writing to ask for permission to reproduce an image from your work
 > in my interdisciplinary PhD dissertation for the University of Alberta.
 > The
 > dissertation is entitled Affordances of Prospect for Academic Users of
 > Interpretatively-tagged Text Collections.
 >
 > I require written copyright permission specifically for the four bound
 > copies of the dissertation at the University of Alberta, and to accompany
 > the copy that is sent to the National Library of Canada for microfilming.
 > If I
 > am not able to obtain permission, there is provision to include a textual
 > description in place of the image.
 >
 > The image I am hoping to use is attached as a jpeg. It is a screen shot
 > showing the toggle map from your paper "Don't click-paint!" at
 > <http://ipsi.fhg.de/~baudisch/publications/1999-Baudisch-HCI99review-Don'tClickPaint.pdf>. If you are willing to let me use this image, please
 > reply by
 > email to that effect. If I am able to find a publisher for my
 > dissertation, I
 > will
 > contact you again to ask for permission for that further use.
 >
 > Thank you for your attention to this matter.
 >
 > Sincerely yours,
 > Stan Ruecker
 >
 > *****
 > Stan Ruecker, PhD
 > Department of Art and Design
 > University of Alberta

WebMail - RE: copyright permission request

Date Sent: September 23, 2003 11:57 AM

From: "Bederson, Ben" < > [Add to Address Book](#)

To: Stan Ruecker

Subject: RE: copyright permission request

Status: Urgent New

That is fine. You may have permission as requested below - thanks for letting me know.

- Ben

+-----+

| Asst. Prof. Ben Bederson Director, Human-Computer Interaction Lab |
 | Computer Science Department |
 | www.cs.umd.edu/~bederson ing |

+-----+

> -----Original Message-----
 > From: Stan Ruecker
 > Sent: Monday, September 22, 2003 9:06 PM
 > To:
 > Subject: copyright permission request
 >
 > Dear Dr. Bederson,
 >
 > I am writing to ask for permission to reproduce an image from
 > your work in my interdisciplinary PhD dissertation for the
 > University of Alberta. The dissertation is entitled
 > Affordances of Prospect for Academic Users of
 > Interpretatively-tagged Text Collections.
 >
 > I require written copyright permission specifically for the
 > four bound copies of the dissertation at the University of
 > Alberta, and to accompany the copy that is sent to the
 > National Library of Canada for microfilming. If I am not able
 > to obtain permission, there is provision to include a textual
 > description in place of the image.
 >
 > The image I am hoping to use is attached as a jpeg. It is a
 > screen shot showing your Photomesa interface. If you are
 > willing to let me use this image, please reply by email to
 > that effect. If I am able to find a publisher for my
 > dissertation, I will contact you again to ask for permission
 > for that further use.
 >
 > Thank you for your attention to this matter.
 >

Date: Tue, 26 Aug 2003 17:41:51 -0400
Subject: Re: copyright permission request

format=flowed
v551)

Status: RCV


Stan,












Permission granted. No problem.

Johanna Drucker


On Sunday, August 24, 2003, at 02:19 PM, Stan Ruecker wrote:

> Dear Johanna,
>
> I am writing to ask for permission to reproduce an image from your work
> in my interdisciplinary PhD dissertation for the University of
> Alberta. The
> dissertation is entitled Affordances of Prospect for Academic Users of
> Interpretatively-tagged Text Collections, and was successfully defended
> on 20 August 2003.
>
> I require written copyright permission specifically for the four bound
> copies of the dissertation at the University of Alberta, and to
> accompany
> the copy that is sent to the National Library of Canada for
> microfilming. If I
> am not able to obtain permission, there is provision to include a
> textual
> description in place of the image.
>
> The image I am hoping to use is attached as a jpeg. It is a composite
> screen shot illustrating the Nowslider. If you are willing to let me
> use this
> image, please reply by email to that effect. If I am able to find a
> publisher
> for my dissertation, I will contact you again to ask for permission
> for that
> further use.
>
> Thank you for your attention to this matter.
>
> Sincerely yours,
> Stan Ruecker

WebMail - Your Google Permission Request 

 Delete  File  New  Reply  Reply All  Forward  Previous  Next  Index  Help  Logout

Date Sent: September 5, 2003 4:15 PM

From: "Amy Kieraldo-Sanchez"  [Add to Address Book](#)

To: Amy Kieraldo-Sanchez <

Subject: Your Google Permission Request

Status: Urgent New

Hello:

Thank you for writing to request permission to use Google in your upcoming project. We have approved your request and require no further information.

Thanks again for writing - and keep using Google!

Cheers,

Amy Kieraldo-Sanchez
Google



Secure mode active.

Powered by Infinite Mobile Delivery v2.6 -- © Copyright 1995-2002 by Captaris

Adobe policy

Screen shots

Screen shots are individual print screens contained within Adobe software applications. For an example of a screen shot, [click here](#). You may use screen shots in advertising, on documentation (including educational brochures), in tutorial books, in videotapes, or on Web sites, provided you adhere to the following guidelines:

1. Your use may not be obscene or pornographic, and may not be disparaging, defamatory, or libelous to Adobe, any of its products, or any other person or entity.
2. Your use may not directly or indirectly imply Adobe sponsorship, affiliation, or endorsement of your product or service.
3. You may not use the screen shot in a comparative advertisement.
4. You may not alter the screen shot in any way except to resize the screen shot. Your use must contain the entire screen shot. You may not use portions of the screen shots and no portions of a screen shot may be included in your product user interface.
5. You may not use screen shots from Adobe beta products or other products that have not been commercially released by Adobe.
6. You may not use screen shots that contain third-party content.
7. You must include the following copyright attribution statement: "Adobe product screen shot(s) reprinted with permission from Adobe Systems Incorporated."
8. If your use includes references to an Adobe product, the full name of the product must be referenced at the first and most prominent mention (Adobe® GoLive®). When referencing any Adobe trademarks, please mark with a ™ or ® as indicated in the [Adobe trademark database for general distribution](#). If you are distributing materials exclusively in Japan, please [contact Adobe](#) for a copy of the appropriate Adobe Trademark information for Japan.
9. If your use includes references to any Adobe trademarks, you should include the following trademark attribution statement: "[List of marks used, with "Adobe" first, if used, followed by other Adobe marks used, in alphabetical order] are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries."
10. Screen shots must be reproduced in their entirety. Permission may be given in certain limited circumstances, however, to use specific portions of a screen shot. For permission to use a screen shot in this manner, [click here](#).

If your use meets the above criteria, no further written permission is required. For use by the press in articles or other news reports, please go to the [Adobe Press Room](#). All other uses require written permission from Adobe, including use of screen shots in or on a television program, movie, or commercial. If further written permission is required, please complete the [online form](#) and submit your request by [e-mail](#). Please allow two weeks for your request to be processed.

Please note that Adobe cannot give permission for use of any third-party content (artwork) that may be contained within an Adobe screen shot. To determine whether an image contained within an Adobe screen shot belongs to a third party, please refer to the guidelines on [stock art](#), [clip art](#), and [photographs](#).

Date: Thu, 13 Mar 2003 09:27:49 -0700
To: Stan Ruecker <sruecker@ualberta.ca>
From: Patricia Clements <Patricia.Clements@ualberta.ca>
Subject: Re: request for permissions
In-Reply-To: <3E73891E@webmail.ualberta.ca>
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"; format=flowed
Status: RO

Stan, this note gives you permission to use the 9 illustrations listed below in your dissertation. Should you wish to publish these at a later stage, you would need to ask permission specifically for that purpose. All best, Pat

At 05:24 PM 11/03/03 -0700, you wrote:

>Hi Pat,

>

>Attached is a PDF file that has the 9 illustrations I am hoping to use in the
>last chapter of my dissertation. I've also included the captions. These
>images refer to material from Orlando, so as we had discussed, I am
>sending them to you to look over, before I send the chapter to anyone
>else. Please let me know what you think. Your word is law.

>

>yrs,

>Stan

(There were actually only 8 images included. They were):

Figure 4.02	This wall of text shows a list of author names from Orlando, in alphabetical order by author last name. A fisheye lens effect is used to allow the entire display to fit at one time on the screen.	182
Figure 4.03	This panorama contains roughly 12,000 names from the Orlando Project, of which the smallest strip shows approximately half at any one time. The names are arranged in alphabetical order in columns. The zooming feature is continuous, and corresponds to the user's movement of the mouse up or down on the screen. The strip is shown here at three different levels of magnification.	184
Figure 4.04	Visual representations of the values for the attribute Certainty, which is used with the various Date tags in Orlando. The tags might be similarly grouped.	194
Figure 4.06	This scattergram shows one point for each of 3600 events. The user is able to select subsets of the points by moving the vertical bars at the endpoints of the selection. In this case, 500 events are shown as currently selected.	198
Figure 4.07	A radio button interface to the <ChronStruct: Relevance> attribute values would allow the user who is familiar with the terminology to select an appropriate choice. For users unfamiliar with the terminology, some additional experience or explanation might be necessary.	203

Figure 4.08	Since the four possible values for <ChronStruct: Relevance> are additive, one appropriate solution is to allow the user to select them by choosing among nested buttons where the outer choices automatically include the inner choices. The degree of grey on each button is supposed to reinforce the idea of additive selection.	204
Figure 4.09	Grouping tags by Relevance values provides the user with an immediate impression of the significance of the different values.	205
Figure 4.10	Here the four <ChronStruct: ChronColumn> values are used to organize a table of chronological results. The horizontal bars indicate that some of the columns have been collapsed to the left. The heading of each column indicates the value. Dragging the vertical strip associated with each heading will expand that column and collapse the others, although some representation is always visible for all four possible values.	209