

Exploring Information Visualization Use Patterns in Casual Contexts

by

David William Sprague

B.Sc., Queen's University, 1998

B.Sc., Queen's University, 2001

M.Sc., University of British Columbia, 2006

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Computer Science

© David William Sprague, 2011

University of Victoria

All rights reserved. This dissertation may not be reproduced in whole or in part, by photocopying or other means, without the permission of the author.

Exploring Information Visualization Use Patterns in Casual Contexts

by

David William Sprague

B.Sc., Queen's University, 1998

B.Sc., Queen's University, 2001

M.Sc., University of British Columbia, 2006

Supervisory Committee

Dr. Melanie Tory, Supervisor
(Department of Computer Science)

Dr. Margaret-Anne Storey, Departmental Member
(Department of Computer Science)

Dr. Amy A. Gooch, Departmental Member
(Department of Computer Science)

Dr. Jim Tanaka, Outside Member
(Department of Psychology)

Supervisory Committee

Dr. Melanie Tory, Supervisor
(Department of Computer Science)

Dr. Margaret-Anne Storey, Departmental Member
(Department of Computer Science)

Dr. Amy A. Gooch, Departmental Member
(Department of Computer Science)

Dr. Jim Tanaka, Outside Member
(Department of Psychology)

ABSTRACT

This dissertation describes a series of studies conducted to explore why people use information visualizations during their non-work time (*casual InfoVis*) and which factors are critical for visualization adoption and long duration use. I also model typical casual InfoVis usage patterns and provide a framework for future hypothesis testing. Each study explored a different facet of casual InfoVis research and each built on lessons from the previous studies. The first study explored the development and evaluation of a casual InfoVis system, PartyVote, and how visualizations can be used to aid informal group social interactions. Results from the evaluation indicate that the system successfully helped give people a more equal share in choosing music during social gatherings and people could strategically choose music, but social pressures did not constrain behaviors or reduce cheating as much as expected. The complexity of factors affecting PartyVote use led to a pseudo-experiment evaluating the appeal of motion based data encoding. Study results indicated that participants formed distinct opinion-based groups and motion data encoding was only considered appealing to less than half of the participants. Utility was a critical factor for half the participants, but a sizable group still preferred motion use, despite knowing that it reduced system utility. My final study examined how people encountered and used visual representations of data (artifacts) during their non-work time. The artifact

study led me to develop the Promoter / Inhibitor Motivation Model (PIMM) of casual visualization interaction. PIMM subsequently helps explain results encountered during the first two studies. The model provides a framework for future casual InfoVis investigations and identifies potential shortfalls and areas of concern when conducting casual InfoVis research. PIMM should also help guide future casual InfoVis system designs.

Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	v
List of Tables	x
List of Figures	xi
Acknowledgements	xiii
Dedication	xiv
1 Introduction	1
1.1 Background	1
1.2 What is Casual InfoVis?	3
1.2.1 Casual InfoVis Example	3
1.2.2 Important Tasks During Leisure Time	4
1.2.3 Similar Terms	5
1.3 Research Goals	8
1.4 Research Scope	10
1.5 Overview of Dissertation	11
2 Previous Work	14
2.1 Visualization Uses & Motivators	14
2.1.1 Insights, Learning & Understanding	15
2.1.2 Tools & Externalized Cognition	17
2.1.3 Communication	17
2.1.4 Other Uses	18

2.2	Empirical Evidence	19
2.2.1	Non-System Factors	20
2.2.2	User Expertise	20
2.2.3	Data Type	20
2.3	Style vs. Utility	21
2.3.1	Style Priority	22
2.3.2	Utility Emphasis	23
3	Methods for Investigating Casual Information Visualizations	24
3.1	Research Questions	25
3.1.1	Motivation	25
3.1.2	Long-Term Adoption	25
3.1.3	Immediate Usability	26
3.1.4	Accommodating Diverse Goals	26
3.1.5	Usage Changes over Time	26
3.2	Research Methods	27
3.2.1	Controlled Analysis	27
3.2.2	Qualitative Analysis: Realistic + Small Group	31
3.2.3	Analysis of Real World Systems: Realistic + Large Groups	34
3.2.4	Additional Approaches	37
3.3	Conclusion	37
4	Study 1: Music Selection using the PartyVote Democratic Jukebox	38
4.1	Introduction	38
4.2	Problem Area	40
4.3	Previous Work	43
4.3.1	Co-present Collaborative Systems	43
4.3.2	Group Music Selection Systems	43
4.4	PartyVote System Overview	45
4.4.1	System Interface	45
4.4.2	Algorithms	47
4.4.3	Implementation	49
4.5	Case Studies	50
4.5.1	Case Study 1: Coffee House	50
4.5.2	Case Study 2: House Party	51

4.6	Discussion	56
4.6.1	Design Implications	58
4.6.2	Lessons Learned	60
4.7	Revisiting the Research Goals	61
4.8	Conclusions and Future Work	63
5	Study 2: Motion in Casual InfoVis and the Interrelation Between Personality, Performance, and Preference	65
5.1	Introduction	65
5.1.1	Motivation for Study Approach	66
5.2	Previous Literature	67
5.2.1	Benefits and Drawbacks of Motion	67
5.2.2	Uses of Motion	68
5.3	Research Overview	69
5.3.1	Instructions, Training, and Consent	70
5.4	Study 2A: Glyph Comparison Task	71
5.4.1	Method	71
5.4.2	Hypothesis	73
5.4.3	Results	74
5.4.4	Discussion	75
5.5	Study 2B: Visual Search Task	76
5.5.1	Method	77
5.5.2	Hypothesis	78
5.5.3	Results and Discussion	79
5.6	Study 2C: Playlist Generation + Interviews	79
5.6.1	Guiding Questions	81
5.6.2	Findings	81
5.6.3	Discussion	84
5.7	Study Limitations	88
5.8	Conclusions & Future Work	89
6	Study 3: Visual Representations of Data During Non-Work Time - Modeling User Goals and Regulated Motivations	91
6.1	Introduction	91
6.2	Artifacts vs. Visualizations	92

6.3	Casual Information Visualization Studies	93
6.3.1	Factors Influencing Visualization Use	93
6.3.2	Research Approach	96
6.3.3	Participants	96
6.3.4	Cultural Probes	97
6.3.5	Artifact Definition and Instructions to Participants	98
6.3.6	Data Analysis	99
6.3.7	Data Cleaning	100
6.3.8	Model Generation	100
6.4	Findings	101
6.4.1	Artifacts Reported	101
6.4.2	The Promoter / Inhibitor Motivation Model (PIMM)	103
6.4.3	Regulating Factors	109
6.4.4	Stages of Artifact Use	112
6.5	Discussion	117
6.5.1	Design Considerations	118
6.5.2	Evaluation Implications	121
6.5.3	Caveats	121
6.5.4	Study Limitations	122
6.6	Related Computer Use and Motivation Theories	122
6.6.1	Motivation Theories	123
6.6.2	Goal Setting Theory	124
6.7	Future Work	125
6.7.1	PIMM in Other Contexts	125
6.8	Conclusion	126
7	Revisiting the Earlier Studies	131
7.1	Revisiting PartyVote	131
7.1.1	PartyVote According to PIMM	131
7.1.2	Potential Improvements	133
7.1.3	Redesigning the PartyVote Study	135
7.2	Revisiting the Motion Encoding Studies	136
7.2.1	Redesigning the Motion Studies	138
7.3	Conclusion	138

8	Conclusions	139
8.1	Extending the Research	139
8.2	Contributions	140
8.3	Final thoughts	142
	Bibliography	144
A	Additional PartyVote Information	159
B	Motion Study Additional Information	168
C	Artifact Study Additional Information	182

List of Tables

Table 5.1	The effects of question type and distractions on mean trial time and accuracy.	74
Table 5.2	Subject groups and their distinguishing traits	90
Table 6.1	Definitions provided to clarify points of discussion in Chapter 6.	128
Table 6.2	Regulating factors identified via interviews and journal entries. A matrix identifies which goals and stages of artifact use regulating factors were associated with.	129
Table 6.3	Artifact examples from of the PIMM stages of use and an associated description of activities.	130

List of Figures

Figure 1.1 Stamen Design’s digg swarm shows recent articles and their connections with moving icons.	2
Figure 1.2 The Baby Name Voyager, showing all names starting with ”Ja” between 1880 and 2007 (<i>www.babynamewizard.com</i>).	4
Figure 1.3 An Euler diagram showing the theorized set of all currently conceivable visualizations and iconic examples of different categories.	5
Figure 1.4 A flow chart of my dissertation.	12
Figure 4.1 A screen shot of the PartyVote system demonstrating the space-themed visualization window and the selection window.	41
Figure 4.2 Alterations to the potentially playable song region based on a vote for ‘2+2=5’ by Radiohead, indicated by a red arrow.	47
Figure 4.3 The distribution of user interaction events throughout the party.	53
Figure 5.1 The ordering of motion encoding sub-studies for each research session.	69
Figure 5.2 Song glyph representing the song ‘Night Clubbing’ by Iggy Pop. Beats per minute and beat strength encodings differ between conditions: static (A), motion (B), redundant (C), and extraneous (D).	71
Figure 5.3 Glyph comparison task results showing the interaction between question type and conditions on mean trial time (milliseconds) and correctness (%/100).	74
Figure 5.4 Glyph comparison task results showing the interaction between the presence of a distraction factor and the question type on mean trial time (milliseconds) and correctness (%/100).	75
Figure 5.5 The visualization used during the visual search tasks and playlist creation tasks.	78

Figure 6.1 An assortment of submitted target artifacts and their primary goals.	101
Figure 6.2 The proposed relationship between source factors, regulating factors, goals and artifact use.	104
Figure 6.3 Hierarchy of participant goals.	108
Figure 6.4 Stages of artifact use. Artifacts continue to be used until inhibitors exceed promoters.	113
Figure 6.5 Examples of recognition stage artifacts.	114
Figure 6.6 Examples of Short-duration Single Use Stage (SS) artifacts. . .	115
Figure 6.7 Examples of Long-duration Single Use Stage (LS) artifacts. . .	115
Figure 6.8 Examples of Short-duration Repeat Use Stage (SR) artifacts. . .	116
Figure 6.9 Examples of Long-duration Repeat Use Stage (LR) artifacts. . .	117
Figure 6.10 PIMM when the primary goals surrounding visualization /system use are uniform across users.	127

ACKNOWLEDGEMENTS

My doctoral work was filled with numerous difficulties. The support and understanding of the following people allowed me to ‘muddle through’ and finish. I would like to thank:

Bill and Gerry Sprague: My parents, who have given my unconditional support throughout my studies.

My wife, Mélanie Sprague: Your love and understanding have meant the world to me. Thank you for the occasional kick in the backside to get things done. I probably would not have finished without your help.

My friends and family: There are too many people to thank individually but your advice and help (and occasional participation in my studies) was invaluable.

Dr. Melanie Tory: My supervisor, who has given me more valuable advice than I can count. You’ve always been understanding, accommodating, and available to help. I wouldn’t be the researcher I am today without you.

My supervisory committee: You each provided a unique and important perspective for my research, yet each of you were incredibly supportive of the various hiccups and changes to my research focus. Both myself and my new family appreciate your understanding.

The GRAND research network and NSERC for funding me during my doctoral studies and making this research possible.

David Sprague

To my loving wife Mélanie

I may permanently be a student on Benjamin's birth certificate, but this should be the last dissertation I ever have to write. Thank you for your love, support and understanding. It's now time for me to get a "real" job.

Chapter 1

Introduction

This dissertation explores factors that motivate people to use computer generated visual representations of abstract data (casual information visualizations) during leisure time or when randomly encountered. I assert that the lack of financial or job related incentives means that casual information visualizations may require a different design emphasis and motivational factors may differ from work related visualizations. In particular, I will identify how usage goals ultimately frame the interpretation of different system features and system use duration patterns. My three doctoral research studies address casual information visualization motivations from three complementary perspectives: I developed a casual visualization and tested how it was used, I empirically tested the importance of a single commonly used aesthetic factor (motion) on task performance and user preferences, and I devised a grounded theory derived model based on how visual representations of data are currently being used.

1.1 Background

Research in Information Visualization (*InfoVis*) frequently focuses on professional data analysis in fields such as biology, physics, and marketing. Facilitating insights, improving task efficiency, and enhancing data comprehension are all important for professional productivity and system effectiveness metrics are easily defined and framed according to these goals. Subsequently, visualization utility metrics can be considered tightly coupled with user motivations in work contexts. A potential question that then arises is: what about more ‘casual’ uses for InfoVis (which we will label *casual InfoVis*)? Large datasets available to people, such as medical histories,

email archives, and music collections, could be visualized to benefit people in non-work related contexts. In fact, many such systems already exist, although they have not (yet) been widely adopted by the public. I believe that in non-work contexts the motivation to use a visualization is not always apparent and requires research. Visualization design guidelines can be derived from this work, hopefully improving future user adoption and making casual visualizations more accessible.

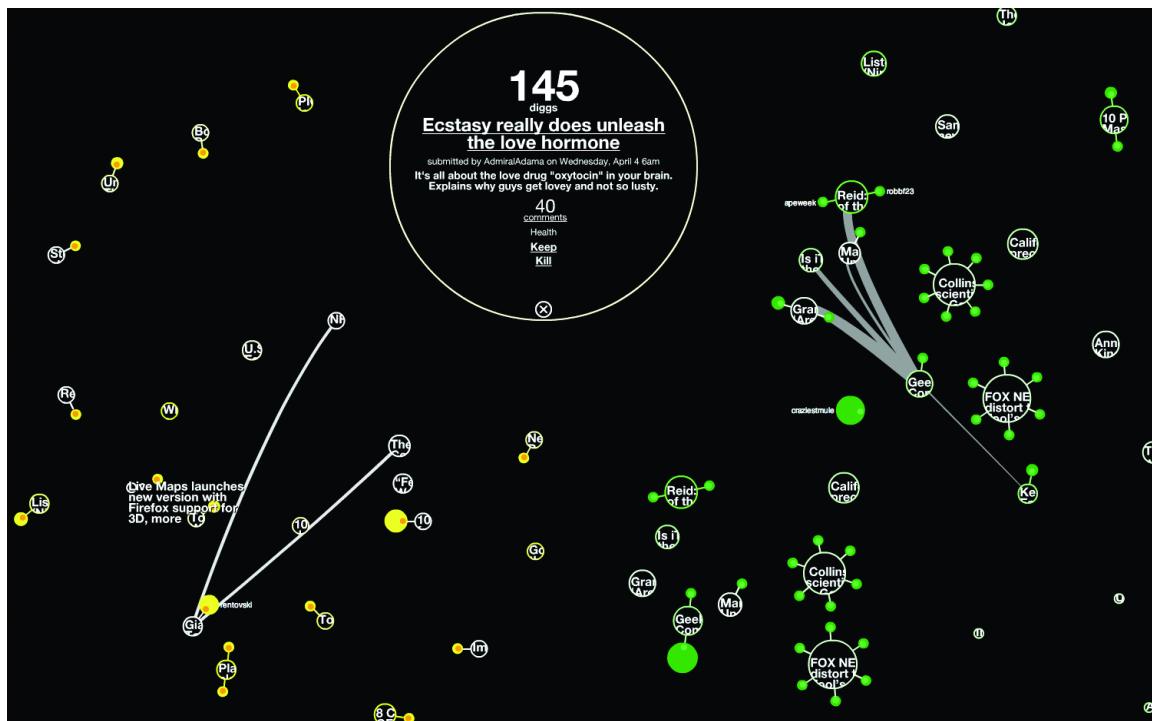


Figure 1.1: Stamen Design’s digg swarm (labs.stamen.com/swarm) shows recent articles and their connections with moving icons. We consider this visualization an example of increased visual appeal potentially at the expense of practical utility.

Why would people with no financial or job-related incentives take the time to learn to explore visual representations of data? If job related productivity is not a critical factor, what factors are critical to initial use, prolonged use, or repeat use? For example, why might some users use the digg lab visualizations [34] when the traditional text list is still available (see Figure 1.1)? Who examined the Guardian’s visualization of twitter message term frequency used during each FIFA World Cup 2010 match and why [118]? The World Cup visualization displayed the most common terms tweeted during each soccer match and therefore potentially indicates what important events occurred and when. Why might a user want to see this information

rather than watching match highlights or reading a sports article? It is currently unclear why people may use visualizations away from work, especially if there is a time and effort cost to using a new system and benefits are not obvious.

Out of curiosity, I asked two women with no computer science training to look at the digg visualizations by Stamen design (see Figure 1.1). To my surprise, comments changed from “sooo cute” and “I like [the] animation” during the first minute to “annoying” and “cute but not practical” less than five minutes later. Motion rapidly went from being appealing and motivating to being distracting and discouraging. Perhaps simply getting eyes on the screen is insufficient. But what makes a visualization successful in informal contexts, and if we don’t know, how do we find out?

1.2 What is Casual InfoVis?

Card *et al.*’s [23] commonly given definition of ‘traditional’ InfoVis is that such systems are interactive visual representations of abstract data used in order to amplify cognition. Alternately, Munzner proposes that Information Visualization is the study of visual representations of data where the spatial mapping of data to a computer screen is chosen rather than associated with the data [109]. We define *casual Information Visualization* as information visualizations used to interpret or understand data in the absence of financial or work related incentives, and used during leisure time, doing chores, or when randomly encountered. Leisure time is subsequently categorized as time periods not allocated for necessary daily tasks and responsibilities. Leisure activities can include pursuing a hobby, playing a sport, and watching television. For my dissertation, the term casual does not refer to casually or haphazardly designed visualizations but instead focuses on the casual context of use. My casual InfoVis definition could be expanded to include any abstract data visualization used outside of work related contexts (essentially non-work InfoVis), but I felt this was insufficiently focused and based on the negation of a situation (non-work). For clarity, I define motivation as the desire or drive to act in certain way or perform a certain task based on an external stimuli or mental process.

1.2.1 Casual InfoVis Example

I consider the Baby Name Voyager by Wattenberg to be an iconic and illustrative example of a casual visualization (see Figure 1.2)[139]. Wattenberg’s web-based vi-

sualization of American baby names chosen since 1880 is simplistic and intuitive, yet colourful and enticing. Unlike traditional InfoVis systems, productivity is not the primary concern surrounding system use and job related system use seems unlikely. Prolonged system use appeared motivated by personal importance, curiosity, playful engagement, and social influences [139]. However, it is unclear how these factors interrelate with traditional usability metrics.




Figure 1.2: The Baby Name Voyager, showing all names starting with "Ja" between 1880 and 2007 (www.babynamewizard.com). Girl names are pink and boy names are blue. Baby Name Voyager is an iconic example of effective casual information visualization.

1.2.2 Important Tasks During Leisure Time

The distinction between casual and traditional visualizations can sometimes be ambiguous. For example, some casual activities can ultimately aid work tasks, and some non-work information (e.g. personal medical visualizations) may be taken even more seriously than work-related visualizations. Nonetheless, I chose to focus on non-work

activities and leisure activities specifically because I believed that I would encounter a wider variety of motivating factors in that context and for the most part, user motivations cannot be assumed.

1.2.3 Similar Terms

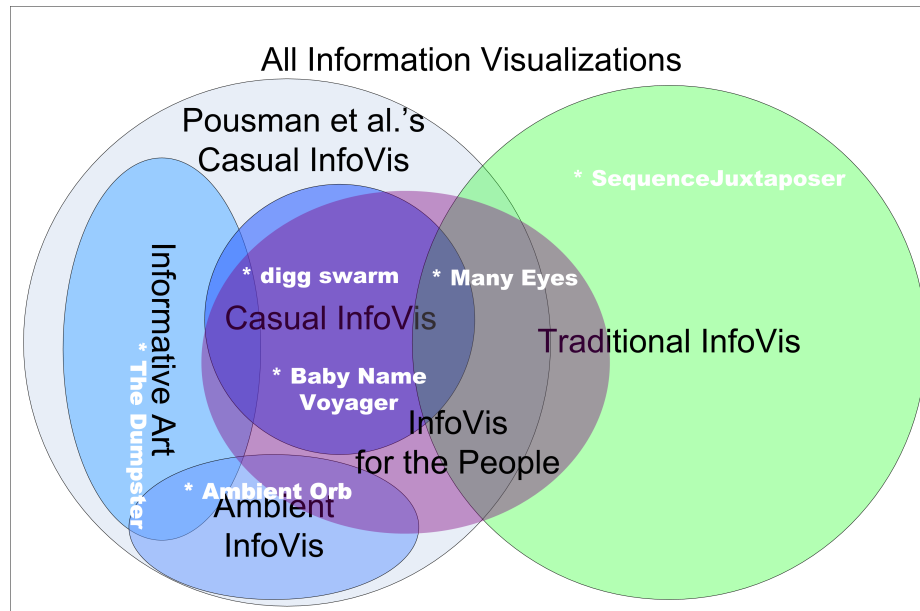


Figure 1.3: A Venn diagram showing the theorized set of all currently conceivable visualizations. The diagram includes the set of traditional work based visualizations and Pousman *et al.*'s [106] Casual InfoVis set. My dissertation will concentrate less on Informative Art and Ambient Visualizations and more on systems used during free time or randomly encountered that help people perform an expected task. Iconic visualization examples are written in white text ([34, 79, 120, 106, 135, 139]). InfoVis for the People [31] was never precisely defined and therefore an exact outline for the region has not been given. InfoVis for the Masses is defined based on the possibility of mass adoption and would be difficult to illustrate with these sets.

Previous research has used several different terms and definitions to describe a similar set of visualizations. I classified casual InfoVis based on its use during leisure time and hence an associated lack of financial or job based incentives. My categorization is somewhat unique and reflects my particular research focus on user motivation. Other terms such as 'InfoVis for the Masses', 'aesthetic / pragmatic visualizations', 'vernacular visualizations', and 'Information Visualization for the People' will be discussed to help frame the current research approach. These similar terms overlap somewhat

in terms of their pragmatic/artistic dimensions [69], the user’s need for formal training, the size of the user population, the importance of personally meaningful data, and the inclusion of non-analytical insights (see Figure 1.3).

Pousman *et al.*’s [106] definition of casual InfoVis is based on visualization edge cases such as informative art, hedonic visualizations, and social visualizations. They noted four primary differences between casual and traditional InfoVis used to demarcate the research area: casual InfoVis includes a wider user population including non-experts, systems are frequently used repeatedly but for a short duration each time, data is often personally relevant rather than work related, and visualizations provide insights that are non-analytical. Although I could have used Pousman *et al.*’s [106] definition, my chosen classification provided a clearer (although some might argue arbitrary) differentiator, particularly relevant for my research focus.

Pousman *et al.*’s definition also includes visualizations as artistic and reflective outlets, which I will de-emphasize, effectively ignoring informative art and the benefits of the ‘sublime’ to expand a user’s horizons [69]. Data driven aesthetics such as the Munterbund book visualization [108], and informative art such as ‘We Feel Fine’ [52], although considered casual InfoVis by Pousman’s definition, will not be discussed except in relation to the task of data analysis. Artistic uses for visualizations would expand the scope of this research too greatly. Unless explicitly stated, all uses of the term casual InfoVis will refer to my provided definition. Making such a distinction between casual InfoVis, informative art, and data driven aesthetics may seem vague and a matter of philosophy but I feel such a clear demarcation is necessary for my analysis. Evaluating the merits of an artistic work also appear to require an entirely different research approach and philosophy. Kosara [69] similarly differentiates InfoVis along aesthetic lines and proposes pragmatic and artistic visualization sub-categories. Although I focus on pragmatic visualizations, both of Kosara’s categories are well explained by the PIMM model (discussed in Chapter 6). I originally referred to the set of visualizations I was interested in exploring as Pragmatic Casual InfoVis, but this seemed to overly limit the research areas explored during the artifact study (Chapter 6), hence I expanded my focus.

Lau *et al.* [76] sorted various visualizations and information art works to derive a new category, ‘Information Aesthetics’, defined as the equal focus on interaction, data, and aesthetics. Thus, Information Aesthetics systems are evaluated both in terms of functionality and artistic merit (artistic influence and meaningfulness). Lau *et al.* also mapped visualizations and informative art along ‘mapping focus’ and

‘data focus’ dimensions. The mapping dimension (direct vs. interpretive) essentially identifies the amount of data ambiguity. The data dimension focuses on data insights and tool-based analysis (intrinsic data) compared to appreciated and interpreted data (extrinsic data) that may invoke personal reflection. As previously stated, I am focusing on pragmatic casual InfoVis applications and thus predominately exploring direct intrinsic data representations. Furthermore, Lau *et al.*’s chosen systems are not necessarily a representative sampling of all visualizations, nor do the authors address mutual exclusivity issues between pragmatism and artistic expression. Finally, the apparent correlation between data focus and mapping focus dimensions suggests that these categories are likely not orthogonal, and could predominately be represented using a single dimension. For these reasons I chose not to use Lau *et al.*’s taxonomy.

Viégas and Wattenberg [133] propose the term *vernacular visualization* to describe visualizations developed outside of the academic research community and adopted by the general public. The primary vernacular visualization example is the tag cloud; a variety of text or type styles and sizes arranged in a single view. Viégas and Wattenberg did not provide additional examples of vernacular visualizations (if other successful ones exist) but the concept of non-academically created visualizations designed to be understood by non-experts seems closely related to the current research.

Danziger [31] proposes ‘*Information Visualization for the People*’ (which I will abbreviate IVftP), a classification similar to casual InfoVis and defined as user-centered InfoVis designed in a way that the average user can both do relevant analysis and enjoy the experience of interacting with information. Productivity for IVftP is considered less a priority since visualizations are meant to be enjoyed and presumably not used for work tasks. IVftP also includes informative art and ambient visualizations like Pousman *et al.*’s definition, but the exact categorization criteria are too vague to compare with our casual InfoVis definition.

Finally, ‘*InfoVis for the Masses*’ appears to be a set of systems closely related to casual Vis, with extremely large potential user populations and when formal system training may not be assumed. The large population also means that user backgrounds, interests, and demographics are more diverse than those of most InfoVis systems. The Many Eyes web site may be considered the iconic example for this category [135]. InfoVis for the Masses does not appear to have a succinct definition, but I assume that casual InfoVis is not a perfect subset: business intelligence systems are used by large groups of business people without data analysis experience while a visualization for a particular hobbyist group may not have a large user population.

Despite the different definitions, the above terms often refer to a section of the design space where a large population of non-expert users use visualizations for activities outside of traditional work based or efficiency critical tasks. Casual InfoVis continues this tradition. I do not assume that casual users are non-trained in visualization use, but I expect that this will be true in the general case (most people are not visual data analysis experts). Thus, novice vs. professional differentiators, in terms of training, are not being explored in my dissertation. However, highly ‘compact’ data representations may not be comprehensible to first time and untrained users and thus may be an important casual InfoVis factor [116].

1.3 Research Goals

As a community, we have made numerous statements about designing for the nonprofessional ‘masses.’ Some researchers believe we should teach InfoVis best practices; tools such as Many Eyes [135] epitomize this approach. Researchers such as Few [39] and Kosara [72] claim that traditional productivity and efficiency metrics continue to be key components for casual InfoVis. Thus the same metrics and design practices for traditional visualizations should still be used for casual visualizations and utility should not be sacrificed for appeal or aesthetics. The InfoVis Manifesto [80] asserts that all effective visualizations should make aesthetics secondary to functionality and knowledge discovery. However, as Kosara points out, many or most ‘InfoVis for the masses’ data sets are being visualized but not analyzed [70]. He suspects that this might be due to the ill defined, open-ended nature of analysis, which can be unappealing (and independent of productivity or aesthetics). If data analysis is not appealing, why might people want to explore data in their leisure time?

Others assume that playfulness, aesthetics, and style are at least as important as functionality, that it is critical to get ‘eyes on the screen’, or that information’s beauty serves its own purpose (for example, see Flowing Data’s blog¹). Researchers, such as Danziger [31] and Cawthon *et al.* [24] suggest that new success metrics, such as enjoyment and aesthetic appeal, are required given that user priorities and goals are different in casual contexts. New York Times visualizations are a popular compromise between these extremes [17].

At the start of my doctoral work, I adhered to the seemingly common assumption

¹www.flowingdata.com

that fun and functionality were effectively mutually exclusive [72, 71]. As such, identifying whether participants sought utility or entertainment was a primary research interest for me. Fun /functionality trade-offs, however, appear to be time dependent, overly simplistic and one-dimensional based on my experiences. Discovering the extra complexity surrounding casual InfoVis use helped to ultimately guide my research's progression.

Currently there is limited research describing casual visualization contexts, and subsequently how they differ from traditional use contexts. Critical factors surrounding casual InfoVis use cannot be identified until researchers understand casual visualization use and motivations better. This forms a central goal to my dissertation.

I claim that casual InfoVis factors (later defined as *regulating factors*) are highly contingent on context and user specific characteristics [141]. My dissertation attempts to empirically identify regulating factors, model their interrelation corresponding to casual InfoVis use, and propose design considerations based on empirical findings.

My dissertation is designed to address the following research goals:

1. **Explore why people use visualizations when not motivated by financial incentives or work related reasons:** I wanted to explore what factors influence why people use visualizations in their free time or when randomly encountered. Even when a system provides a clear purpose, user motivations are not necessarily implied. Do people find data exploration appealing or do they only examine data when there is a clear benefit to themselves? If a visualization is being used for fun, what system features will promote or inhibit system use? This goal is predominately addressed in Chapter 6.
2. **Investigate motivations, goals, and use contexts to design a model of casual InfoVis interaction:** Developing a more comprehensive means to evaluate and test casual InfoVis systems requires a model of user interactions and motivations. My Promoter / Inhibitor Motivation Model provides a means to predict system adoption and use patterns based on known or predicted user goals and system features. This model is described in Chapter 6.
3. **Examine how perceptions of utility, aesthetics, and entertainment interrelate with regards to system use:** It is not clear if opinions about fun, functionality and beauty interrelate, although previous discussions have sug-

gested that productivity and artistic or expressionist visualizations are necessarily mutually exclusive. More direct empirical evidence of pragmatic /artistic trade-offs can help clarify what users want in casual InfoVis systems.

4. **Identify methods for evaluating and researching casual InfoVis:** Although not a primary contribution of my dissertation, the three research approaches I used provide informal feedback as to which approaches work best for casual InfoVis research and what important casual InfoVis specific pitfalls are associated with each research method.

One cannot go in blindly testing developed systems in a controlled manner as key metrics such as satisfaction, desire to use the system, and perceived usefulness are all highly contingent on user characteristics and context. Deployed system users may also predominately be early system adopters or otherwise not representative of the general population. Hence, I approached casual InfoVis testing using a variety of methods in an attempt to tease out key factors from a variety of research directions. Furthermore, these experiences taught me important lessons associated with casual InfoVis system testing, including designing appropriate empirical testing techniques, qualitative exploration approaches, and system deployment considerations.

1.4 Research Scope

The current research focuses primarily on what I believe is the key difference between casual and traditional InfoVis: user motivation. As such, the studies were constrained to exploring how people use casual visualizations and their thought processes surrounding use. I particularly concentrated on user motivation, goals, and what they hoped to accomplish using the system.

Clarifying constraints to the current research is critical given the general subject matter and the number of visualizations involved. My dissertation is not an exhaustive examination of all or even most casual InfoVis systems. Furthermore, in most cases I will not be comparing systems and the quality of their designs. I will also not focus on discussing the effects of casual InfoVis adoption on our society. Instead, I will concentrate on how casual InfoVis affects individuals and why the visualizations were used.

My dissertation will also not delve deeply in topics such as ambient visualizations, informative art, and community visualizations. Perhaps more importantly, the

artistic merits of any visualization will not be discussed in detail as this requires entirely different evaluation methods and seems outside the scope of most visualizations. Instead, artistic incentives will be generally ignored and examining the ‘sublime’ is assumed to be a relatively rare use of casual InfoVis. Although artistic appeal may act as a system incentive, the focus on more pragmatic visualizations means my results rarely focus on artistic factors.

1.5 Overview of Dissertation

The remainder of my dissertation has been separated into seven chapters (see Figure 1.4). First, I will present an overview of the relevant casual InfoVis literature and how these studies frame my research. Chapter 3 provides an overview of different effective research techniques that can be used to evaluate casual InfoVis systems and the contexts of their use. These suggestions derive from both previous research and lessons learned from my own three studies. Chapters 4 to 6 then present the three primary doctoral research endeavors. My research approach was not decided in advance; rather lessons from each stage led me to explore different facets of casual InfoVis design. I then generated new research questions based on previous results. Chapter 4 will discuss an early attempt at designing and evaluating a casual InfoVis system and the lessons I learned. The system evaluation provided insights into both how casual InfoVis systems can be effective, and revealed some interrelations between user opinions and visualization use (such as the difficulties people have with abstract data concepts). Interaction styles, aesthetics and user motivations were too difficult to isolate with a novel visualization used during an uncontrolled social gathering. Hence common assumptions made about casual InfoVis design (such as the importance of aesthetics) were not validated. Chapter 5 describes my controlled study examining the effects of motion data encoding on user preference and performance. The motion study was an attempt to isolate an aesthetic factor and its effects in a more controlled single user environment. I learned that the conditions appropriate for casual InfoVis evaluations (such as high level tasks rather than step by step instructions) enable too much inter-personal variance. However, I still noted a distinct user opinion clustering associated with gender and field of study. Identifying how user characteristics, context, and system features influenced user motivations, opinions, and usage behaviors seemed a critical next step, given these results. Chapter 6 explores the more holistic artifact study conducted to leverage the previous studies and provide a

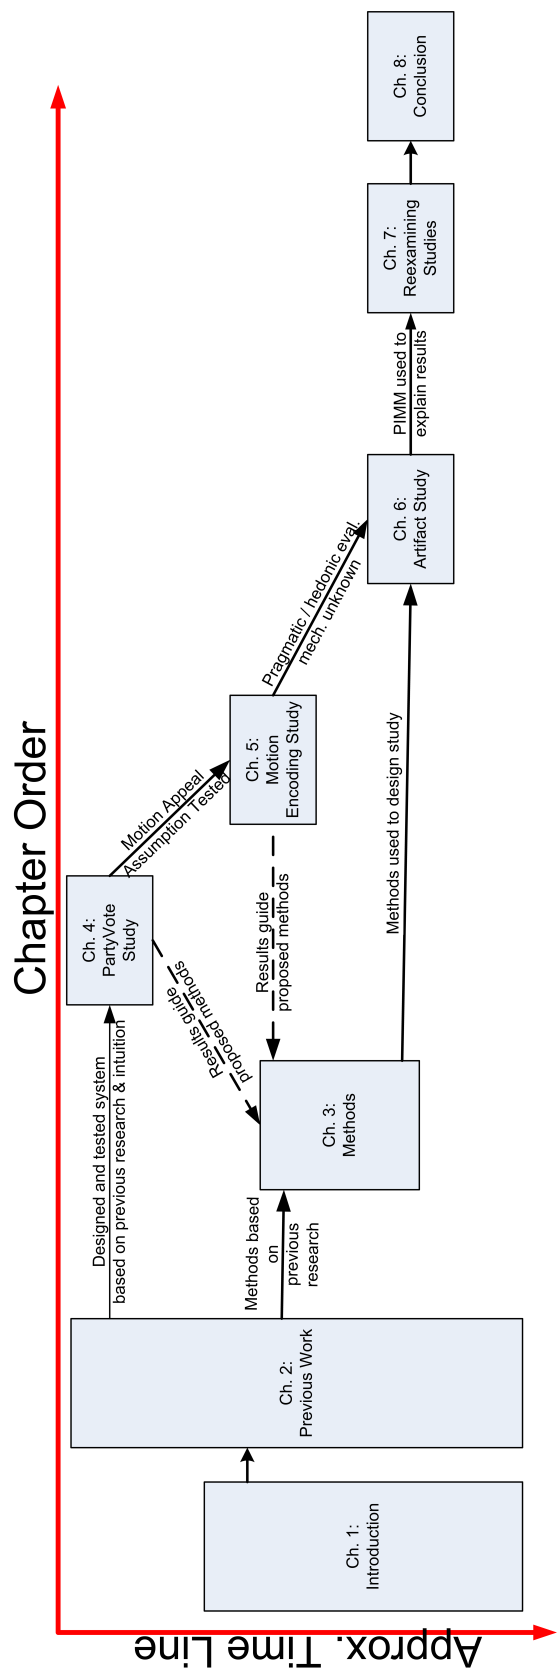


Figure 1.4: A flow chart of my dissertation. After a brief examination of previous literature (Chapter 2) and a discussion of casual InfoVis research methods (Chapter 3), I will describe the three primary studies for this dissertation. Each study builds upon lessons from previous studies (indicated on connecting arrows). I will then re-examine these studies based on what was learned (Chapter 7) and finish with the conclusion.

model of casual InfoVis use. The artifact study explores the visual representations of data people encounter in their daily lives and how people interpret and manage this information. I subsequently use these findings to propose the *Promoter / Inhibitor Motivation Model (PIMM)* of casual InfoVis use. PIMM suggests that intermediary factors (such as user goals) exist between objectively measurable factors (such as a user's education level) and visualization use. Hence, despite controlling what user groups study or the context surrounding system use, a direct relationship between system factors and activity patterns is unlikely. The PIMM model provides a novel framework with which to explore casual InfoVis use patterns. The three research approaches also provide a diverse and complementary set of experiences, better enabling me to identify and articulate research methods that may be successful in the future. In Chapter 7, I re-examine the findings from the first two studies using the PIMM framework. Finally, I conclude and propose future work in Chapter 8.

Chapter 2

Previous Work

To address my research questions, it is first critical to understand which casual visualizations currently exist and what questions previous research has addressed. Although the current discussion of previous work is not exhaustive, this discussion is designed to give a sufficient overview to provide a current ‘state of the field’ and what has been attempted. Philosophical contentions in the research community about ways to design casual InfoVis systems also helped motivate my doctoral work. Previous work predominately relating to one of my studies will be discussed in the corresponding chapter.

2.1 Visualization Uses & Motivators

Casual InfoVis systems have existed for a relatively long time. Early systems such as the Dynamic HomeFinder real-estate visualizer [3] and the FilmFinder collaborative music selection visualization [2], while compelling, were not formally evaluated or provided to the public. However, large visualized data sets useful to people outside of work time were relatively rare. The advent of the Internet and the popularization of digital libraries has led to a proliferation of data available to the average person, providing more opportunities for casual visualizations [102]. Visualizations can provide numerous benefits to everyday users, not just professionals. The reasons for casual InfoVis use identified below are based on previous results and our results from the artifact study (see Chapter 6).

Traditional visualization benefits can be categorized into three dimensions: insight generation, externalized cognition, and communication. Heer *et al.* [57] similarly

propose that collaborative visualizations can serve two primary purposes: communication and analysis / exploration (corresponding with insight generation). It is not clear whether casual visualization will share these dimensions, although previous discussions and initial studies suggest that entertainment and aesthetic appeal could pose as unique casual InfoVis appeal factors [68, 76, 106, 131, 134]. Hence, a new categorization of uses seems warranted for casual InfoVis.

2.1.1 Insights, Learning & Understanding

Insight Generation

Perhaps the most commonly assumed use for casual InfoVis systems is the more traditional use for visualizations: generating new insights and discoveries about a data set. *Insight generation* generally refers to the ability to acquire new knowledge or understandings using a visualization such as a genetics visualization revealing the biological underpinnings to a medical problem. For example, the visualization of tweets during FIFA world cup matches presumably provides the possibility of learning about how each match progressed, what twitter users were interested in at the time, and commonalities they had between each other [118]. The SLife system could also be considered a means to aiding insight generation by examining one's digital traces. The Expense Visualizer, is designed to promote Canadian governmental transparency by visualizing government spending and providing this data to the public for analysis [41]. Similarly, a Washington Post visualization titled 'Top Secret America' displayed the proliferation of intelligence agencies created after September 11, 2001 and the difficulties understanding their bureaucratic structure [101]. Many Bills allows people to discover facts about past and present U.S. congressional bills [7]. Finally, the site 'Patients Like Me' enables people to both find other people suffering from the same disease or medical problem and to visualize symptoms and treatments [86]. Patients can compare and filter treatments and symptoms to potentially identify the best course of action for treating their illness. The Patients Like Me visualization is a clear example of how casual InfoVis systems with personally meaningful information can be as important to the user as traditional visualizations, despite the lack of financial incentives.

Curiosity

Curiosity, or the drive to discover new information seems to be a common motivating factor determining casual InfoVis use. I presume, based on Wattenberg’s [139] results, that many initial users of the Baby Names Voyager examined the system out of curiosity. Similarly, I noted during an informal exploratory examination of people using casual visualizations, that curiosity was the primary motivation guiding interest in the digg labs visualizations [34]. Pousman’s Imprint system relies extensively on curiosity to encourage people to explore the printing habits of themselves and their colleagues [103]. Peer recommendations and curiosity about the system also likely promoted the use of novel but not particularly utilitarian visualizations such as the ‘Web Trends Map’ by Information Architects [64] and the periodic table of visualizations [78]. The Web Trends Map displays the most popular web sites as stops on the Tokyo subway map with train lines representing different site topics. The periodic table of visualizations uses the shape and format of the periodic table of elements to interactively describe a visualization taxonomy.

Digital Traces & Reflection

User activities often leave *digital traces* or records about what that person has done. These traces can both enable people to reflect on their own behavior and to understand the actions of other users. *Reflection* can be considered the thorough analysis and contemplation of information to understand one’s self, an action, event, or thing. For example, Slife allows users to analyze the proportion of time spent using different computer programs [121] while ‘ThePrinter’ visualizes communal printer activities enabling people to reflect on their printing habits, resource consumption, and the activities of their colleagues when taking a break from work. GMail inboxes can also be visualized, enabling people to determine trends in their lives and who they are in contact with [144]. Some systems, such as Bloom.io’s ‘Fizz’ visualization [18], are designed to be “playful, explorable, visually compelling views on personally relevant information from services like Facebook, Twitter, Instagram, and iTunes.” Recent work by Bartram *et al.* [12] demonstrates digital trace analysis for self-reflection on a much larger scale, providing visual feedback about energy use to the inhabitants of a ‘smart home.’

2.1.2 Tools & Externalized Cognition

Visualizations can also provide an opportunity to aid the speed at which people acquire information or help users perform a specific task. I consider such efficiency improving visualizations *tools*, although the term *cognitive aids* could also be used. For example, an ultrasound image or a visualization of a glaucoma patient’s optic nerve can help expedite a patient’s understanding about their current state of health. Hard drive content visualizations can act as tools for determining disk usage and hidden uses of disk space [114]. The TakeAway’s ‘Vote Prediction Tracker’ visualized the number of electoral college wins expected for each 2008 presidential candidate based on the news outlet [8]. Although some insights can be gained by the Vote Prediction Tracker, this visualization provides a rapid means to disseminate predictions from numerous news sources and determine the overall consensus about expected election results. Visualizations can also enable people to quickly explore personal data patterns and habits, such as a heat map of Belgium, enabling citizens to view their home’s energy use and compare that with others [1].

Externalized cognition, as defined by Card *et al.* [23], is the use of an artifact or external object to reduce mental effort and memory demands when performing a task. For example, a pen and paper can aid someone doing arithmetic. Thus visualizations that help externalize cognition can perhaps be viewed as a tools subset when the primary task is not data analysis. For example, the web accessible ‘Baby Name Voyager’ provides access to name trends without the need for users to remember particularities; they can simply look up the names again as necessary. Naming trends are also easier to comprehend than examining the raw data. Google’s ‘Ngram viewer’¹ allows the comparison of multiple terms used in literature over time. Although these graphs can provide insights, displaying multiple term trends on the same graph also simplifies cross-referencing. Commonly encountered static data representations such as pie charts, bar graphs, and time tables (eg. bus schedules) can also be considered effective means for externalizing cognition.

2.1.3 Communication

Casual visualizations are also often used for persuasion or as a means to communicate an idea. For example, ‘Many Eyes’ has been specifically designed to promote community involvements and discussions over data sets [135]. The Many Eyes web

¹<http://ngrams.googlelabs.com/>

site permits anyone to upload and visualize their own data using online tools, however, that information must also become freely available (and editable) to all other users. Furthermore, Viegas *et al.* [134] identified that Many Eyes visualizations were frequently used for persuasion. Hans Rosling’s [112] GapMinder world population visualization clearly illustrates this application. His talk uses data to illustrate a particular point he is trying to make rather than providing viewers with the open exploration possibilities associated with insight generation tasks. Other visualizations are designed to communicate an idea or tell a story such as General Electric’s visualization of aging world populations [37]. Although this visualization enables people to discover demographic differences between countries, a major emphasis of the system seems to be to illustrate how the world’s elderly population is growing at the fastest rate ever. The PartyVote system described in Chapter 4 is another example of a communication based visualization as the information is designed to inform users about the music choices of other users. Heer *et al.* [57] provide an extensive survey of communication based visualizations.

2.1.4 Other Uses

Personal Expression

We will note in Chapter 6 that personal expression and personal image construction can be driving factors behind visualization creation. We consider *personal expression* to be any activity used to externalize someone’s feelings or personality characteristics. Personal expression can be used as a personal artistic outlet or a means of communication with others. Keystrokes was specifically designed to generate unique visualizations based on the text and typing style of users, providing an image unique to the user and the underlying tone of the writing [91]. The authors subsequently suggest these images could then be included as a form of expressive email signature. Many Eyes provides another outlet for visualization personal expression. All users can generate and customize their visualizations thus making new content a means of self-expression [134].

Social Interaction

Systems like the Baby Name Voyager also relied on blog communities, not only to promote initial use, but also to promote repeat use. Blog users frequently discussed

their findings and offered challenges to other users, thereby providing a motivation for repeat system use and non-egocentric data analysis. Similarly, systems such as Imprint [104] and PartyVote [123] provide opportunities for users to examine the actions and behaviors of other users, which can facilitate the use of social pressures to constrain behaviors. Visualized peer data may also aid or frame communications between people using the system. Communication benefits are discussed in detail by Heer *et al.* [57] and in Chapter 4.

Ludic Engagment

Applications with a strong *ludic* (Latin for playful) emphasis provide no direct work related utility but may be used by people for entertainment [88]. Such systems may include the playful emphasis of the digg labs visualizations [34], ‘GraphJam’ humorous charts [63], and ‘the Periodic Table of Visualizations’ [78]. Viegas *et al.* [134] similarly reported Many Eyes being used for similar entertainment purposes such as a participant identifying the contents of his freezer for data analysis. Similarly, the academic project ‘Sashay’ [98] is designed to facilitate ‘wonderment’ by visualizing how cellular telephone tower use changes over time for people nearby.

Aesthetics and the Sublime

Finally, despite generally limiting our research to casual InfoVis systems, we would be remiss to not identify the importance of aesthetics, artistic appeal, and the appeal of what Kosara [69] refers to as ‘the sublime.’ For such systems, visual aesthetics, ambiguity, and a reduced insight identification rate can be critical to attracting users and may be mutually exclusive with utility and practicality. For more complete list of such systems see work by Pousman [105], Kosara [69], and Ghahary [45].

2.2 Empirical Evidence

Despite the above assumed motivations for using casual InfoVis, direct empirical evidence examining the importance of each of these factors is lacking. For example, I found little to no evidence of ludic visualizations motivating casual users (such systems exist but motivation was not tested). Some sites such as Many Eyes and the Baby Name Voyager have been released to the public, and users have articulated social and communication based motivation [134]. However, even in such circumstances, the

strength of different motivating factors cannot be compared overall, as each system is ideally suited for particular tasks. Furthermore, there may be differences between system adopters and the general public in terms of what motivates them. Hence the reason the current research was conducted.

2.2.1 Non-System Factors

User tasks and goals can be affected by a variety of influences besides system factors including: the user, the context of use, and the data being visualized. I define *context* as the set of circumstances or influences excluding the user, visualization, or data represented. Hence, context includes the physical environment, time of day, and the user’s schedule. Contextual factors have not been directly discussed in the literature, perhaps due to the intractable number of elements involved. My three studies do examine context, but continue to not exhaustively describe or categorize all use contexts.

2.2.2 User Expertise

The effort required to learn how to use a visualization means a system’s success can often be contingent on the user rather than the system itself. Heer *et al.* [57] propose a simple yet effective user categorization: novice, savvy and expert users. Novice users are considered people with little to no data analysis experience, savvy users have extensive, often self-taught data analysis experience but this is not the primary focus of their jobs or training, and expert users have extensive training and are often visualization designers and professional data analysts. Novice users and savvy users are the primary casual InfoVis groups, as such people are not formally trained in the creation and visual analysis of data. Systems designed for novice users seemed particularly correlated to the casual InfoVis definition.

2.2.3 Data Type

Heer *et al.* [57] also propose three overlapping data categories corresponding to their meaning to the user: personal data, community data, and scientific data. Personal data refers to visualizations involving data exclusively concerning the individual user such as email visualizations and jogging programs that display a person’s progress over time [62]. Community data involves information directly relevant to a group

of users such as the Baby Name Voyager. Finally, scientific data is not personally relevant to the user but the user analyzes it anyways. The authors presumed that such systems would be used for a job (particularly for scientific research), but their definition does not exclude other reasons such as what Luther *et al.* [82] term “data junkies.” We presume most casual InfoVis research involves personal and social data sets and this was informally confirmed during our three studies (described later). Heer *et al.* [57] also claim that personal meaning is a clear motivator for both social visualizations and visualizations in general; a claim we will examine in more detail in Chapter 6.

2.3 Style vs. Utility

For several years, I have labeled casual InfoVis systems that sacrificed usability and comprehension for style, fun, and aesthetic considerations as *InfoFizz*. This term was used to compare such systems to cola drinks, which are appealing but lack nutritional benefits. In fact, based on previous discussions by others about casual InfoVis, a utility to style dimension was a major component of my earlier explorations. I define a *utility benefit* as an objective task performance improvement (such as task efficiency) associated with using a visualization. *Style* can be considered a design decision intended to make a visualization’s appearance or interactive behavior relatively unique, potentially improving aesthetic judgments.

With opposing schools of thought regarding casual InfoVis, it appears that empirical and observational evidence is necessary to clarify what factors are important for motivating casual InfoVis use. The studies described in Chapters 4 to 6 are an initial step. However, after three user studies conducted for this dissertation, it appears that a style/utility dimension is not entirely an accurate or beneficial framework to describe casual visualizations. Style and utility do not appear to be mutually exclusive, nor do they appear to be perfectly orthogonal. Thus, the traditional style against utility comparisons may not be accurate. However, it is worth carefully considering aesthetics and utility separately. In the end, my studies suggest that characteristics influencing system use may be dependent on what the system is being used for rather than any form of absolute rules. Despite these findings, the style versus utility discussion was prominent in many casual InfoVis discussions; some of which are described below.

2.3.1 Style Priority

We note that InfoFizz systems, although lacking productivity benefits, may provide an underestimated service; they attract users and are pleasing to use. Like popular art or popular music, InfoFizz systems may be dismissed by trained professionals but may gently introduce the average person to the field. Thus current InfoVis best practices may not be ideally suited for casual contexts.

Furthermore, utility and productivity may not be a user's primary concern in casual contexts. After all, professional drivers are primarily concerned with how a car performs but the average driver may be swayed by something as simple as the paint colour. Norman's [93] three-process theory of human judgment, suggests that only behavioral processes involve logic, functionality, and practical utility. Visceral and reflective processes, instead, are subconscious and based on emotion and previous experiences respectively. A user's judgment about a visualization, is therefore not based solely on practical judgments of productivity.

Danziger's recent thesis advocates that casual InfoVis systems need to consider both system appeal and what the visualization is being used for. Many commercial visualizations have taken the 'not just functionality' concept further, actually sacrificing productivity for visualization style and 'fun factor.' For example, the music visualization created by Burst Labs shows a limited subset of unmarked circles constantly floating in a circular enclosure [22]. Each circle represents a song and upon mouse over events, song information is displayed and the song may be played. Although more music data could be shown, this may intimidate initial users. Instead, the animation may have a calming effect or be visually appealing. Fizz provides another prime example of style over utility [18]. Circles in the interface drift and bounce off each other to emphasize the playful nature of the system, even if such movements make the data more difficult to interpret. Despite the stylistic additions, however, Fizz is designed to aid reflective insights about one's friends. Other popular style priority examples include the Web Trend Map 2008 [64], the Periodic Table of Visualizations [78], and the digg swarm [34]. The Periodic Table of Visualizations provides an overview of previously deployed visualizations, constrained to a tabular format resembling the periodic table to provide an overview of visualization research [78]. The digg visualizations all visualize recent digg messages and their popularity [34]. The digg visualizations rely heavily on animations to attract attention and indicate changes in popularity and content. For more 'InfoFizz' examples, see visu-

alization discussions by ‘Smashing Magazine’ [122] and ‘ReadWriteWeb’ [99].

2.3.2 Utility Emphasis

Stephen Few’s capstone speech at InfoVis 2007 [39] thoroughly analyzed the emerging InfoVis for the Masses area and the role researchers can play in guiding its development. Few focused on the application of traditional InfoVis systems to real-world settings, with no regard on what motivates non-professional visualization use. Few also implied that visualization researchers should lead by example, in order to ensure visualizations are developed the ‘right way’ with an emphasis on productivity and display efficiency.

Similarly, Tufte [126] has long decried extraneous information in charts and visualizations, deeming the extra information as ‘chart junk.’ For interactive visualizations, this term only accounts for some of the potentially extraneous or misleading information. For example, animations solely for the sake of aesthetics can also be distracting. InfoVis for the masses systems such as Tableau public² and Many Eyes use traditional InfoVis best practices.

One critical concern when examining why people use casual InfoVis systems is the importance of context on the user’s motivation to use the system. In the next chapter I will explore a variety of different but complementary methods (organized by contextual factors) for evaluating casual InfoVis systems.

²www.tableausoftware.com

Chapter 3

Methods for Investigating Casual Information Visualizations

How do traditional and casual InfoVis methods differ, and is such a distinction actually necessary? We believe that these two forms of InfoVis need different approaches because they involve different user goals and motivational factors. Casual users do not need to use a system for their job, they are unlikely to get formal training or read the system manual, and from my experience, they can be quite resistant to using a new system. Thus, getting ‘eyes on the screen’ can be problematic. Furthermore, most casual InfoVis systems are relatively new and have no proven track record of repeat use so a meta-analysis of existing successful and unsuccessful systems is not currently possible or appropriate. This stands in contrast with the numerous successfully adopted traditional visualizations.

Methods used for serious games (video games for training or learning), virtual reality, and human-computer interaction communities have all similarly examined the dichotomy between entertainment and work, and this research can help guide casual InfoVis study directions. However, casual InfoVis concentrates in a unique way on facilitating human data processing, meaning the communities are not identical. Hence a discussion about casual InfoVis methods appears warranted. This chapter directly addresses my fourth primary research goal (see Chapter 1): Identifying methods for evaluating and researching casual InfoVis.

3.1 Research Questions

Casual InfoVis provides a unique research area, not because it is used in less ‘serious’ tasks but because, as a tool, its usage patterns might be more multifaceted. This seemingly subtle difference leads to at least five general research questions fundamental to casual InfoVis (not all of which are addressed in my dissertation):

1. What motivates system use?
2. What factors influence long-term adoption?
3. How important is immediate usability, and how can it be provided?
4. How do we ensure a system can adapt and accommodate diverse user needs and goals?
5. How does system use change over time?

To explore these new questions, however, new or less common research methods are needed. Each of these questions are discussed in greater detail below.

3.1.1 Motivation

If job responsibilities and financial incentives do not drive system use, why do people use visualizations such as the Baby Name Voyager [139]? Intrinsic factors (factors inherent in the data and user) promoting system use include personal interest, curiosity, personal challenges, control, boredom, aesthetic appeal, and novelty. Extrinsic factors (factors external to the data and user) include social pressures, peer recommendations, competition, and recognition seeking. Controlling or precisely quantifying these factors can be difficult, so qualitative evaluations are critical. The contextual nature of extrinsic and intrinsic factors also means that laboratory studies might be too artificial. The key point here is that we cannot assume what motivates tool use (despite what the visualization may be designed for), so motivational factors and realistic contexts should be a major part of user evaluations. Evaluating contextually based motivations is a central facet of my dissertation.

3.1.2 Long-Term Adoption

The challenge of long-term adoption also arises with traditional InfoVis systems. However, long-term adoption is perhaps more crucial to evaluate for casual InfoVis,

given that demonstrated tool effectiveness is not necessarily the predominant factor influencing continued use.

3.1.3 Immediate Usability

Professionals such as geneticists frequently expect training time when a system is introduced. However, previous research on casual shared screen visualizations has found a critical need for immediate usability [59]. Immediate usability reduces user frustration and possible embarrassment, which act as barriers to use. Initial frustrations will likely affect all casual systems. Exactly how initial system confusion and intimidation affect user exploration and how to provide immediate usability are open research questions.

3.1.4 Accommodating Diverse Goals

The lack of formal training, the diverse subject pool, and less utilitarian nature of tasks performed will likely require more varied tasks and approaches to system use than in traditional InfoVis. I noted this diversity of uses during my evaluation of the PartyVote music visualization system (Chapter 4) and my motion-encoding experiment (Chapter 5). For example, with PartyVote some people wanted to influence the music in their favor, some wanted to learn about other party attendees, and some just wanted to satisfy their curiosity. Designing for and evaluating this diversity of goals can be difficult. My artifact study's findings (Chapter 6) suggest that constraining possible user goals via the interface's design could also improve user satisfaction with a visualization.

3.1.5 Usage Changes over Time

Casual-usage patterns can be different, manuals can be ignored, and novelty can be important, thus casual InfoVis usage over time might differ from traditional InfoVis. Do trade-offs exist between initial appeal and long-term use? How can a system adapt to users' changing needs and experience levels? Note that although this is an interesting research question, system usage over time is not the focus of my dissertation. I briefly examine this in Chapter 6.

3.2 Research Methods

On the basis of the above research questions, I will examine three research approaches - controlled analysis, contextual/exploratory analysis, and the analysis of real-world systems - and how they can be adopted for casual InfoVis. Special attention will be paid to how these techniques test user motivation. A fourth ‘Additional Approaches’ section will describe promising but less orthodox techniques not used for my dissertation. For each research method, the typical number of participants and degree of experimenter control will be identified on the headings. For a more detailed discussion on research methods, see thorough discussions on the subject by Creswell [27] and Leedy & Ormrod [77].

For many online discussions about the requirements and needs of visualizations away from work (and thus user motivations), many of the proposed suggestions are based on either traditional InfoVis results or intuition [31, 40, 68, 99, 130]. However, there is starting to be a body of empirical work to verify and test many of these claims. A key issue, however, in conducting casual InfoVis research is the balance of contextual factors and scientific rigor. I will discuss many of these study types, their benefits and limitations, and previous research conducted using them. The three studies conducted for my dissertation will also serve as illustrations, although these studies will be discussed in further detail in subsequent chapters.

3.2.1 Controlled Analysis

Traditionally, researchers have quantitatively evaluated InfoVis techniques in highly controlled environments, using metrics such as time, error rate, and insight generation [100]. However, these approaches may not always evaluate whether a system meets a casual user’s goals, motivations, and needs. It can be difficult to conduct controlled empirical studies to evaluate casual InfoVis systems since experimental pressures to use such systems may influence participant behaviors when they would not normally. Further, researchers can often offer help and aid with regards to learning how a new system operates. For example, von Ooi [95] asked participants to use two Facebook applications to evaluate the CASSM evaluation method. However, by providing instructions on how to use each system and enforcing a set time using it, system learning barriers to entry and immediate user motivation were effectively removed.

Casual InfoVis usage is complex and constrained to context. Therefore, to evaluate PartyVote in a realistic environment, we examined participants in an uncontrolled

natural environment. We used system logs to examine usage patterns but relied on interviews to help clarify user opinions, motivations, and approaches. A more controlled experiment would not likely have revealed observations such as strategic voting, adherence to social norms, and infrequent bursts of system use over the four-hour observation period. Furthermore, it was clear that the coffee house study was not sufficiently correct contextually to evaluate PartyVote.

Laboratory Studies: Controlled + Small Group

Despite the above contextual issues, controlled studies, in comparison to non-controlled studies, permit greater generalizability, stronger conclusions, and make statistical probability tests of quantitative data justifiable. Controlled studies for casual InfoVis seem most appropriate when examining system factors outside of the context of use. For example Bartram *et al.* [11] examined the use of motion based data encoding for icons (termed ‘moticons’) outside of realistic tasks where such icons would be used. Instead, the ability to process motion based icon information was tested in a controlled way. Appeal associated with motion based data encoding (Chapter 5) used a semi-controlled method. My motion-encoding experiment collected log data during the playlist generation task, measured glyph perception and performance abilities using the timed glyph comparison task, and collected user opinions during half-hour interviews at the end. My hybrid / mixed method approach let me connect task actions with user opinions and approaches. However, work by Reilly and Inkpen [107] demonstrated that contextual factors can greatly influence a visualization’s use patterns. They demonstrated that for a transformation between different maps, a ‘juxtaposing’ transition was most effective in a controlled ‘white room’ setting while a morphing transition was most effective in a noisier public atrium. Thus, not only can context and environment unexpectedly affect laboratory results, but examining how context of use changes performance metrics is an interesting and potentially fruitful research direction.

Red-Herring Experiments: Maintaining true double-blind experimental conditions can be difficult, particularly when the tested systems are clearly novel to subjects. The lack of double-blindness is particularly problematic when conscientious subjects use a system purely to help out the experimenters. A potential solution is a ‘red-herring’ study, which obfuscates the primary research question by focusing on another, seemingly more important task. For example, during my motion-encoding

experiment, most subjects claimed to be unaware that my primary research interest was motion encoding's effect. Instead, most subjects thought I was testing the music system itself.

Online Evaluations: Controlled + Large Groups

Frequently, controlled studies examine small groups, however, online testing permits larger subject pools by sacrificing experimental controls. *Online experiments* might provide an ideal environment for some casual InfoVis evaluations [24]. First, the diversity and number of subjects recruited can potentially be larger, although study control is generally sacrificed to acquire this. Second, online testing's loosely controlled nature might suit casual InfoVis experiments in terms of context. Third, on-line test-beds enable activity logging and can dynamically adjust to participant actions. Finally, the self-motivation or procrastination characteristics likely associated with online research participation probably match traits found in a casual InfoVis system's early adopters / initial users. As such, online evaluations act like a half-way between controlled empirical testing and 'in the wild' system evaluations. For example, Dörk *et al.* [35] evaluated their VisGets blog mapping system by enabling participants to log onto a secure system, solve a series of problems using the system and then fill out an opinion questionnaire. Unfortunately, results from this online evaluation have not been published yet, but the approach was still relevant. Cawthon and Vande Moere [24] evaluated the appreciation of aesthetics in visualizations by asking 285 online participants to choose between pairs of static visualization images and to perform information retrieval tasks (such as "find the largest value"). The authors measured speed of task completion, accuracy rate, task abandonment and latency of erroneous response. The authors concluded that bold colours, a lack of 3D, and 'clean' displays improve system appeal, which can lead to improved system use. Furthermore, they noted a correlation between latency in task abandonment and incorrect response times with reported aesthetic measures. Unfortunately, the static nature of the images meant that changes of perspective and interaction style were not considered, which may have substantially affected results. Furthermore, participants describing a visual preference for one visualization over another does not necessarily entail that they would be more likely to use one system over the other, although, based on their results, aesthetics may affect system abandonment.

Online research, however, offers a major drawback: in-depth interviews and quali-

tative data are more difficult to collect with distributed subjects. Questionnaire data only partially addresses this issue.

Questionnaires and Surveys: Controlled + Large Groups

Maximizing the number of people involved in a study can help minimize individual quirks and biases and examine a far larger portion of a target population than otherwise possible [82]. A mix of closed and open questions on questionnaires enables a simultaneous mix of theory support and exploratory analysis. Responses to close-ended questions can provide support for theories while open-ended questions allow people to expand their responses and provide an explanation for the close-ended responses [77]. Questionnaires enable geographically distributed groups to be examined such as open source communities. For example, Roberts *et al.* [110] examined how well the *Self-Determination Model* explained behavioral patterns in open source projects. Questionnaires were provided to active community members along with an analysis of how community involvement and code submissions were correlated with extrinsic and intrinsic rewards. Contrary to expectations, Roberts *et al.* found no evidence indicating that extrinsic motivations decreased intrinsic motivations. Instead, pay was associated with status motivations and the motivation to participate in the project. Questionnaires, however, may limit user responses to pre-formatted questions and restrict details about why people took the actions they did. Less formatted surveys (with paragraph style questions), however, can lead to a lack of convergence amongst user opinions if the subject pool is not sufficiently large.

System evaluations, such as KeyStrokes' questionnaire, also reduce the costs associated with giving feedback often associated with system evaluations [91]. The KeyStrokes questionnaire was distributed to people using the system during two conference demonstrations, and, as such, time limitations would have presumably limited participant responses. KeyStrokes users frequently had differing goals when using the system, sometimes seeking utilitarian benefits for the visualization (being able to reverse engineer a text based on the visualization), some played games using the system, while others enjoyed the personal touch a generated image provided each email. Utility/style evaluations will be explored more in Chapter 5, while the effects of goals on key system features will be discussed in Chapter 6.

Concerns

On the basis of our results, I suggest extensive piloting to find the balance between the ecological validity of tasks and experimental control. Exercising greater experimental control might remove factors integral to real-world use. For example, the Baby Name Voyager’s appeal is often due to curiosity, the desire to procrastinate, and social interactions, but a lab study asking people to use the system would ignore these factors. To see these complexities, researchers must determine the task’s ecological validity, which can be difficult in a controlled setting and is susceptible to inadvertent experimenter influence. Thus, red-herring tasks and online experiments seem well suited for casual InfoVis. Deployed real-world systems (discussed later in this chapter) can also address these issues.

I also discovered an important lesson about casual InfoVis during my doctoral work: user actions can be extremely diverse, even for guided tasks such as playlist creation (used during the motion encoding study). Such diversity is not inherently bad; it led to unexpected findings such as the clustering of subjects based on their opinions and task performances. However, with any sufficiently complex task, diverse user behaviors can negatively affect quantitative performance measures. In contrast, qualitative methods embrace the complexity and diversity of user behaviors and help us better understand them.

To help manage this diversity, I suggest sometimes using hybrid or mixed method qualitative / quantitative approaches, such as the methods used during my motion-encoding experiment. In that study, subjects did not behave as expected; for instance, when I asked them to choose exercise music, they occasionally chose cool-down music. My mixed method approach let me explore user decisions and actions case-by-case. Alternately, if experimental control is required, little to no attempt should be made to address ecological validity. Instead, a constructive positivist approach may work, testing key component aspects of the system to build an understanding of the whole. Such an approach worked for Bartram *et al.* [11].

3.2.2 Qualitative Analysis: Realistic + Small Group

Qualitative evaluation of InfoVis is an emerging trend, although some qualitative approaches such as interviews and direct observation (such as job shadowing) have been used by the research community for some time [27]. Researchers have recently proposed *multidimensional in-depth long-term case studies* (MILCs)[117] and *grounded*

evaluations [65] as effective InfoVis methods. These approaches could generate rich, in-depth understanding of complex user activities, motivations, and opinions. Also, you can achieve greater realism because with qualitative methods you don't have to experimentally control tasks. A grounded theory inspired approach - which I consider a qualitative approach - is used in our Artifact study (see Chapter 6). I will briefly discuss the three qualitative approaches used for my dissertation. Note that I have chosen to describe two data collection techniques (rather than the associated methods) and one method with a unique collection and analysis approach. I have found that ethnographic studies [110] and case studies [125] are perhaps the most frequently used research methods in InfoVis, and for this dissertation, only case studies were used. Case studies attempt to investigate a particular person, technology, or event in depth while ethnographic studies examine a group of people with a shared culture (such as open source developers [110]). The current research predominately explored how a particular set of technologies (visualizations and infographics) were being used. For other methods, please see the cited references above.

Interviews (Case Studies / Ethnographies)

Interviews rely on the investigation of (typically) a small group of participants to explore their previous experiences and thought processes. Thus interviews tend to work well for case study and ethnographic methods [77]. User experiences are also presumed to have been in a realistic environment. Unstructured interviews, unlike questionnaires, enable researchers to adapt questions to explore a common research topic, even when user experiences are truly disparate. However, interviews typically rely on participant introspection, which is affected by recall biases, rationalizations of non-rational tasks, and faulty memory. Interviews can also severely bias results based on how questions are formulated. I used interviews for all three dissertation studies.

Direct Observation (Case Studies /Ethnographies)

Direct observations typically avoid the recall issues associated with interviews but the researcher's presence can influence behaviors. For example, when job shadowing a participant to see how he or she uses email at work [119], the person's actual actions and approaches can be noted. Discussions during the task can illustrate that person's thought processes, although this too is subject to rationalizations and self-reflection

biases. Direct observations tend to work well for both case studies and ethnography studies. Unfortunately, irregular and unscheduled events can be difficult to observe without spending a prohibitive amount of time with each participant. Hence, direct observation was used for the PartyVote study but was not applicable for the Artifact study.

Grounded Theory

Grounded theory studies aim to generate a theory by systematically analyzing ‘grounded’ (real-world) qualitative data. Such studies typically obtain data through interviews, although any data source (including journal entries) can be used. Different data sources can complement one another and improve the results’ reliability (*triangulation*). Data analysis typically involves manually coding (or *tagging*) video data and interview transcripts. When the current data seems insufficient, additional data is collected and analyzed. Data collection typically ceases when new tags are not identified (*saturation*). Patterns and trends between tags are then iteratively derived as a process called *theoretical coding*. The term *axial coding* is also frequently used for the same process, although the approach details differ somewhat. Patterns and trends then form the basis of deriving new models and theories. For details on conducting grounded-theory-based studies, consult “Grounded Evaluation of Information Visualizations” [65] or a qualitative-methods text. A grounded theory approach was used for our Artifact study.

Caveats and Concerns

When implementing these approaches, care must be taken during both data collection and analysis. For example, experimenter actions and instructions can strongly influence user behaviors. Similarly, faulty recall and experimenter prompting can influence subject self-reporting. For the Artifact study, I attempted to avoid recall biases by having participants fill out a daily survey. However, it was evident that instructions when providing artifact kits to new participants did affect how they approached the study. Finally, the actions and opinions of one or a few users do not necessarily represent the broader user population.

3.2.3 Analysis of Real World Systems: Realistic + Large Groups

Deployed systems provide a unique opportunity to examine visualizations as they are used in the real world. Developing systems for real world use with substantial activity logging and (optionally) follow-up questionnaires and interviews can effectively illustrate system usage patterns. System evaluations do not replace more in-depth qualitative measures because opinion and motivational inferences should not be based on log data. Furthermore, this approach does not provide extensive control (relying instead on frequencies and correlations). Instead, deployed system analysis acts as a compromise between exploratory and confirmatory work. Although the data collected is typically quantifiable, log data is inherently noisy, and isolating patterns in it might be difficult. However, if you want to observe system use ‘in the wild’ without subject self-reporting biases, this approach is crucial.

Testing real-world applications requires a system to be completely developed. Therefore, if the system to test does not exist, the development and deployment time for real-world analysis can be prohibitive. However, if a system is available for study, the user pool can be considerably larger and more diverse because incentives to participate are fundamental to the application. Furthermore, users are the intended audience by definition, although they may be early adopters. Perhaps most important, the drive or influential factor behind the system’s use is the system itself. Guilt, financial incentives, and even altruism toward the experimenters are unlikely to influence subject behaviors. Also, system features can be manipulated over time, user actions can be logged, public comments about the system (e.g. blogs) can be analyzed, the researcher may have better access to the user community, and system changes that correspond with a change in use result in a stronger causal implication.

Examples of Analyzed Deployed Systems

Martin Wattenberg has provided a unique but effective example of the deployed system analysis approach [139]. The Name Voyager was released to the general public and became a minor Internet fad, sparking discussions on many community posting boards and blogs. Wattenberg collected comments about the system (from web pages, blogs, and discussion boards) and analyzed them for use patterns. Collected comments helped illuminate how system adopters used the system and illustrated the social yet personal nature of the Name Voyager. Amongst his results, Watten-

berg noted that users fell into ‘roles’ similar to Bartle’s [9] user groupings: achievers, socializers, explorers, and killers. Achievers tend to focus on goal accomplishment, socializers concentrate on social interactions and developing inter-personal relationships, explorers seek to discover new information in the data, and killers use the system for more combative social interactions (such as publicly mocking a particular name or trend). User comments also revealed that peer recommendations and challenges helped establish a community of repeat users but frequently users examined their own names or the names of close relatives initially. Egocentric searching reveals the importance of personally relevant information, potentially used as a way for people to situate themselves within the data.

Similarly, the site Many Eyes was analyzed using both system data and by analyzing blog and web page references to the system [134]. Results indicated that visualizations were used in five distinct ways: analysis, social interaction, data mirroring, expression (called “sending a message”), and innovative uses (which could be considered a form of expression). The five uses were mirrored in how blog authors used the system. We will identify similar categories during the Artifact study (Chapter 6). The authors also concluded that data analysis (in the form of observations, questions, and hypotheses) occurs on Many Eyes, but at a lower frequency than their previous sense.us work [58]. A related Many Eyes study examined the thought processes of ten bloggers and ten non-bloggers with similar use patterns using phone interviews [30]. System use patterns and regularity were determined via system logs. Blogger interviews revealed that participants used the system for personal reasons (such as their research or hobbies), to share data with others, to start a conversation, and to contribute to their community or group. These reasons tended to map well with the (above) five Many Eyes uses and our own research observations.

Some prototypes are released to the public in a limited manner. A mail visualization called ‘Themail’ was evaluated using participants recruited through university and research laboratory mailing lists [132]. The visualization and analysis tool were distributed via email enabling users to examine their own data and avoiding privacy concerns. Subsequently, volunteers were interviewed to identify what they used the system for. In the case of Themail, the system was used as a ‘digital mirror’ enabling users to reflect on their email activities, both in terms of ‘needles’ (searches for exact messages) and ‘haystacks’ (identifying trends among the emails). Use patterns also indicate a strong utility emphasis surrounding system use.

Some casual InfoVis systems have been evaluated in person, within the context of

their use (thus they can also be considered direct observation studies). For example, the ‘EMDialog’ visualization informed attendees of the Glenbow Museum in Calgary about Emily Carr’s life and work [59]. The designers assumed the visualization would be used only once for a short duration, and so it was designed to be intuitive, visually appealing, curiosity evoking, and engaging. EMDialog was evaluated by letting museum attendees use the system as usual while researchers took field notes a distance away, two to four hours a day over a fifteen day period. Volunteer attendees also filled out questionnaires available during the exhibit’s run. The authors concluded that attendees were motivated to approach their system for three primary reasons: the display technology, the visualization’s appearance, and seeing other people use the system (the so-called *honey pot effect* [21]). Hinrichs *et al.*’s findings suggest that visual appeal and social factors are strong motivators for public displays. Field notes also indicated that people who did not find ‘meaningful’ information (photos and statements) left the visualization, suggesting that initial successes with the system were critical for longer term use. Demotivating aspects of the visualization included technological intimidation, the lack of instructions, and the ‘performance’ pressures associated with using a public display.

Concerns

Developing fully deployed systems costs considerable time and effort. Robust system development is difficult, and no guarantees exist that people will use the system or that the logged data will reveal meaningful trends. Furthermore, any deployed system should address a sufficiently unique area so as to provide a natural incentive for use.

During my doctoral research, I considered implementing and deploying an on-line political visualization that would offer unique insight into Canadian politics, thereby motivating its use. Political InfoVis is not novel; two examples are Christopher Healey’s 2008 US election visualizations¹ and Takeaway’s prediction tracker². I believed that new access to the data would attract attention and that the system could provide useful information that was previously unavailable. All system activities would have been time stamped and associated with an IP address to help isolate individual users or households. Log data would have let me determine what features initially attract users, how usage patterns change over time, and the rate of return use. I could have also identified common usage mistakes and system elements that

¹www.csc.ncsu.edu/faculty/healey/US_election

²<http://vote2008.thetakeaway.org/2008/09/20/track-the-electoral-college-vote-predictions>

were ignored. However, my collaborators and I decided that the time commitment required to implement such a system (by all group members) and the substantial risk that the system would not be widely adopted outweighed potential benefits of this approach. My former research direction is only mentioned to illustrate the potential pitfalls with a system development research approach.

3.2.4 Additional Approaches

A recent paper by Heer [55] addresses a variety of novel (and currently rarely used) ways to test InfoVis for the Masses. First, both Heer and Kosara [69] propose the idea of visualization criticism to evaluate new systems ‘in the wild.’ In much the way art is critiqued, visualizations could be evaluated and critiqued as an acceptable publication, thereby providing incentives to the academic community without requiring developers to design based on a set style. Casual InfoVis systems could be critiqued by the general public in addition to academics with a meta-review of comments providing conclusions.

The second approach advocated by Heer is so-called ‘crowdsourcing’ or using mechanical turk systems to outsource an experiment to crowds of people. In this manner, participants would be paid a modest monetary amount to perform simple perceptual experiments or provide preference data. Crowdsourcing appears similar to online testing, however, subject pools are expected to be much greater with far less work involved per participant. These approaches have not been thoroughly validated or tested but they may offer promising ways to evaluate casual InfoVis systems.

3.3 Conclusion

Although casual InfoVis shares many similarities with traditional InfoVis, performance metrics alone seem insufficient in this new context. The methods identified in this chapter were a first step in answering my research questions and potentially helpful for other casual InfoVis researchers. In the next chapters, I will discuss the three primary studies constituting my doctoral work, which should help illuminate how these casual InfoVis research methods can be utilized.

Chapter 4

Study 1: Music Selection using the PartyVote Democratic Jukebox

4.1 Introduction

My initial project for this dissertation involved exploring how a visualization can be used to mediate group social interactions and to improve perceived fairness for group music selection. As such, we developed (but did not deploy) a democratic music jukebox system called *PartyVote*; designed to give each attendee at an informal social gathering the opportunity to help select the music. Music was selected specifically to provide the greatest overlap between participant choices. The visualization itself helped users find music, understand how others voted and how the system operated, and mediate individual behaviors by leveraging social norms. Two user studies were conducted to test the system in a realistic use environment.

I began my investigation of casual InfoVis design somewhat blindly, by developing a novel music visualization based on traditional InfoVis best practices and intuition (not always both at the same time). For example, a space travel metaphor was used to improve the appeal and ‘fun factor’ of the system, while a traditional multi-dimensional scaling to two dimensions was used to position music on a planar surface. The system was then tested in two contexts which quickly revealed to us benefits and short-comings in both our system design and our evaluation method. Please note that I will use first-person plural pronouns when describing the work performed for all three studies (Chapters 4 to 6) since I had research collaborators in each case.

Three primary research goals were addressed by the PartyVote study in the fol-

lowing ways (PartyVote specific sub-questions are listed as sub-points):

1. **Explore why people use visualizations during leisure time (Goal 1):** Why might someone want to use PartyVote rather than letting others choose the music played? We explored why people chose to use the system and which factors (eg. social pressures) influenced participation using questionnaires, log activities, and interviews. Influencing factors are identified throughout the discussion section.
 - **How important is initial usability and system simplicity?** An initial concern surrounding the PartyVote design was user intimidation and confusion due to the novel way music was chosen. Hence, difficulties initially using the system were noted, as were approaches people took to understand the visualization.
2. **Perceptions of utility, aesthetics and entertainment relating to system use(Goal 3):** PartyVote provides a variety of different ways someone can approach system use. Entertainment or playful features like the space metaphor were provided, and the system itself could be used to play strategic ‘voting games’. PartyVote can also simplify music selection at a party, improving utility. We predominately examined ways the system was used via interviews.
3. **Casual InfoVis Methods (Goal 4):** The PartyVote study demonstrated an initial approach on how casual InfoVis systems can be contextually evaluated. Study results also helped illustrate short-comings and potential improvements with our approach.
 - **How can one conduct a realistic evaluation of casual InfoVis systems when system novelty and the presence of the researcher influence use patterns?** We attempted to examine PartyVote in a realistic setting by testing it during regular social gatherings with people who knew me (a participant observer approach). Log data provided quantifiable metrics while interviews enabled us to explore unexpected results.

The above research goals also led me to a number of important but unexpected nuances associated with casual InfoVis research, which will be discussed in the “Lessons Learned” section. I will also examine results particular to PartyVote, given that the evaluations were initially conceived for usability testing. The work described in the

remainder of this chapter has been published, in part, at the conference on Advanced Visual Interfaces 2008 [123].

4.2 Problem Area

Conflicts in informal social environments such as house parties often arise from differences in individual preferences. Unlike music sharing, choosing music for a group involves making compromises between each user’s individual tastes. Music is frequently chosen by a party’s host to avoid conflict. This leaves music selection in the hands of an individual. Although some hosts allow anyone to help determine the music played, the time required to repeatedly select songs or albums is often unappealing. Small groups of self-designated disk jockeys (DJs) can take over music selection responsibilities, also limiting decisions to a select few users. Playing music that everyone knows can frustrate more ‘musically adventurous’ party-goers. Conversely, playing less popular music can annoy people who want to hear familiar songs.

We designed PartyVote as a system that provides established groups with a simple democratic mechanism for selecting music at social events (see Figure 4.1). We define a *party* as a social gathering with the primary purpose of socialization and where pre-recorded music is played. PartyVote is ideally designed for parties involving established social groups of friends. Established social groups tend to behave by implicit or explicit rules or *norms* agreed upon by group members [29]. For example, obstructing others from choosing music is probably contrary to group music selection norms. Adherence to norms via peer pressure is more probable in small to moderate sized groups, as undesirable behaviors are more evident [29, 83]. A visualization system for a casual setting ideally would be intuitive, informative and would ideally support unwritten social rules to ensure fair system use while enabling individuals to express their preferences. Our system accomplished this by improving *user visibility* - information presented to everyone about an individual’s actions. Although PartyVote could be used in larger groups or with strangers, we did not believe group norms would constrain voting behaviors and conflicts may have arisen. Our design goals for PartyVote included:

1. **Support individuals & appease the group:** The system was designed to act as a ‘discount disk jockey’, taking requests and playing music that would

1: Selection radio buttons

2: Song selection windows

3: Vote button

4: Search box and search button

5: Music playback control buttons

6: Currently playing song information

7: Stars (non-playable songs)

8: Planets (playable songs)

9: Space ship (position of playing song)

10: Pop-up song information

11. Emile Campagne song "Un Canadien Errant"

147 songs available

Artist Album Genre

Thalia Zerkon
The Afghan Whigs
The Age Of Electric
The Allman Brothers Band
The Arcane Warnings
The Beatsteaks
The Beatles
The Blountown Rats
The Bourbon Talmadge Choir
The Chemical Brothers
The Constantines
The Corn Sisters
The Cranberries
The Crystal Method
The Dears
The Dears
The Girl Next Door
The Grapes Of Wrath
The Hidden Cameras
The Jan Spencer Blues Explosion
The Juliana Hatfield Three
The Lemonheads
The Magnetic Fields
The National
The National
The Organ
The Orphan
The Philosopher Kings
The Prodigy
The Ramblers
The Roots
The Roots
The Roots

Because
Carry That Weight
Come Together
Golden Shambers
Her Majesty
I Wanna Be Like The Sun
I Want You (She's So Heavy)
Maxwell's Silver Hammer
Mean Mr. Mustard
Octopus's Garden
Ob-La-Di, Ob-La-Da
Poker Face
She Came In Through The Bathroom Window
Something
Sun King
The End
The End
Give Me Your Money
Woolly Bully
A Day In The Life
Blackbird
Cry Baby Cry
Don't Pass Me By
Gass Onon
Good Night
Happiness Is A Warm Gun
Helter Skelter
Honey Pie
I Wanna Take A Trip With You
No So, I Tired
Julia
Junk
Mean Mr. Mustard
Mother Nature's Sun

Display song weighting by planet size

Adj. Normative Song Sim. Metadata

Vote 3

Reset Votes

Search Type search item here

GO!

5

6 Song Rank! The Suburbs | Artist: Ben Folds | Album: Rocker! The Suburbs | Genre: Alternative & Punk | Total Vote Weight: 1.4 | User Vote Weight: Any | Row 59

Peated: 0.03

The End (2:19)
The Beatles
Abbey Road
Genre: Pop
weight: 0.13
BIS: 0.03, Scott: 1.0, Pete: 0.02

Figure 4.1: A screen shot of the PartyVote system demonstrating the space-themed visualization window (points 6-10) and the selection window (points 1-5). The visualization matches the system state at the end of case study 2 with the contested song identified (11). System components are identified on the right.

appease the most people.

2. **Leverage existing social dynamics:** We assumed that system use would be constrained by peer pressure in established groups. The system allowed users to see each other’s votes and their influence on the music selection. Awareness information subsequently allowed social pressures to be applied.
3. **Minimize & simplify necessary interactions:** The time commitment for choosing music was minimized. For example, users did not need to choose more than one song. The system was also not designed to be the center of attention. The visualization was designed to be transparent, fun, and intuitive.
4. **Use common hardware & personal music:** Small casual social gatherings are more likely to have a PC, speakers, and a local digital music library than to have a tabletop display or other specialized equipment.

PartyVote allowed each participant at a party to choose a song, album, artist, or genre from a local digital music collection. Each voter was guaranteed that at least one song from their choice would be played within two hours. Each song in the collection was given a weight dictating the probability that it would be played. User votes increase the weight of similar songs and defined the boundaries of the potentially playable song region in the collection’s music information space (the black region in Figure 4.1). Songs outside of this area were not played. Each voter determined at least one song that was played while the remaining songs were chosen to appease the greatest number of people.

The current research offers three main contributions. First, mediating social conflicts and decision making using a minimal commitment voting mechanism is a novel approach to group music selection with substantial benefits. Second, PartyVote uses visualized social awareness cues, enabling peer pressure to enforce system fairness. No previous group music jukeboxes used this approach. Third, the two case studies evaluating PartyVote revealed interesting and unexpected ways that users interacted with our system and each other. Study results introduce new research topics and guided subsequent research during my dissertation.

In the following section we discuss related research. The “PartyVote System Overview” section provides an in-depth description of the system. The Case Studies section describes two user studies we conducted to initially evaluate PartyVote. We

then discuss design implications, revisit the doctoral research goals, and conclude with future work.

4.3 Previous Work

PartyVote builds on previous research involving synchronous co-located (co-present) computer assisted collaborations, digital music system design and music recommender systems. Although music sharing applications are similar to PartyVote in some ways, the pervasive co-present nature of a party means that musical compromises for all participants are a primary consideration. Furthermore, we consider music sharing between two users with independent sound sources (headphones or personal speakers) to be practicing remote music sharing as users are not hearing the same sound signal [13]. Asynchronous music sharing systems such as Kazaa [47] and remote synchronous music sharing like iTunes music sharing [136], both permit individuals to choose the music they want to hear. Consequently, we focus our literature review primarily on co-present music sharing.

4.3.1 Co-present Collaborative Systems

Previous Computer Supported Cooperative Work (CSCW) research for small casual groups, such as the notification collage by Greenberg & Rounding [51], semi-public display research by Huang and Mynatt [61], and work by Morris *et al.* [87], all identify a need for awareness information. These systems rely heavily on group norms to ensure the systems are not abused. Volda *et al.*'s [136] work on iTunes music sharing suggests that the awareness cues about one's music library constitute a form of social image construction. Participants in her study edited and modified the music they shared with the group based on how they believed others would judge them. Computer supported social mediation systems such as Meme Tags [19] and Ticket2Talk [85] were designed for large parties. Our research similarly looks at social mediation, but we focus on smaller groups in which most people already know each other.

4.3.2 Group Music Selection Systems

There are three common approaches to computer mediated music selection for synchronous co-located groups: playlist generation systems, recommender systems, and

collaborative music jukeboxes. Playlist generation systems create song sequences based on either a user's tastes or musical flow. Systems like last.fm [75] and Pandora.com [97] use user feedback, similar to votes, to refine and guide this playlist generation. These systems, however, provide no feedback as to how music is selected and are designed for single users.

Recommender systems have perhaps explored computer mediated group music selection the most thoroughly. Such systems use artificial intelligence techniques and knowledge about users to formulate recommendations. Intelligent systems such as musicFX [84], Adaptive Radio [25] and Flytrap [28] rely on user profiles. These profiles can be used by a collaborative filtering algorithm or other heuristic to find a best compromise for the group [20]. User profiles are generated by collecting a person's opinion (or vote) about each item in a set. User profiling incurs a large time cost to acquire a reasonable number of votes. Also, user profiles rarely consider the user's current mood, which may be critical to finding appropriate music. Furthermore, McCarthy and Anagnost [84] reported that some users discovered how their algorithm worked and constrained their interests to force other users to listen to their top preferences. Jameson [66] suggests using transparent decision mechanisms and providing user awareness information to avoid such manipulations. PartyVote addresses both of these suggestions.

Collaborative jukebox systems enable users to jointly select music. Existing systems require frequent user interactions. Jukola [94] is a democratic MP3 jukebox designed for use in public places, such as a coffee shop. Jukola relies on users to nominate songs for playing and participants vote from four songs as to what they wish to hear next. The researchers found that a staff member used the system unfairly by repeatedly skipping songs. We believe this was due to the loosely knit structure of the group, a lack of participant visibility, and the absence of a peer pressure mechanism to mediate conflicts. By contrast, PartyVote's visualization allowed users to see each other's influence on the music selection and thereby influence their peer's choices. MUSICtable [124] is a collaborative music library visualization system, designed for private social gatherings. Songs in the system's database are visually represented using a static two-dimensional (2D) geographic map metaphor, similar to PartyVote. Unlike PartyVote, MUSICtable required regular user interactions, and it was designed to be the center of attention at a party. Finally, recent work by Landers [74] follows up on PartyVote's design, providing a group music selection method via voting but requires participants to connect with the system via wireless devices such as smart

phones. Landers' 'Party Jukebox' approach attempts to simplify and clarify the user interface compared to PartyVote, enabling individuals to generate music playlists that can be shared with the group. Party Jukebox's primary usability issue, however, was the inability for people without wireless devices to contribute; an issue PartyVote was actively designed to avoid by using a single interface.

4.4 PartyVote System Overview

PartyVote was designed for use with a regular keyboard, mouse, and monitor. Our system discussion focuses on our three research priorities: the visualization, the voting mechanism, and how we expected the design to affect user decisions and group dynamics. System interaction techniques are only discussed as needed.

4.4.1 System Interface

The PartyVote interface consisted of two main sections: a visualization window and a text based selection window (see Figure 4.1). Clicking an entry in the selection window resulted in corresponding items in the visualization window being selected and vice versa (i.e. brushing and linking).

Selection Window

The selection window permitted 'iTunes-like' text-based music browsing by artist, album, song title, or genre. Selecting an artist, album or genre in the left selection window resulted in all songs in that sub-list being displayed in the right selection window. Selected items could be voted for using the 'vote' button, which required the user to input a user name. Votes for sets (albums, artists and genres) resulted in one randomly chosen song being guaranteed and the weight of that vote being distributed throughout the set. Text queries returned all songs with matching text in the song name, artist name, or album fields. Conventional pause/play, stop and skip buttons provided music control.

Visualization Window

PartyVote used people's votes to weigh and cull the set of potentially playable songs during a party. Weighting was done by determining a similarity between each pair of

songs, and positioning the songs in 2D space so that similar songs were close together. A space travel analogy provided a playful, intuitive and informal atmosphere for the system, while still conveying similarity-based point distances. Initially, all songs were represented as *stars*: songs that have no probability of being played. Each star was drawn as two intersecting light gray line segments so it faded toward the gray background. Each star had a fixed 2D position, calculated a priori to improve system performance.

Votes defined the playable region of music, which was shown as a convex polygon or the convex hull including all guaranteed songs. Each song in this region, or *planet*, was provided a weight. Votes increased the weight of nearby songs. Planets were either guaranteed to be played (*guaranteed songs*) or potentially playable (*potential songs*). Guaranteed song planets were distinguished from other planets by an orbiting moon or Saturn-like ring. A planet's size represented its weight, and its colour matched the vote that influenced its weighting the most. If a vote v_1 represented by the colour red contributes 0.3 to a song's total weight of 0.4, that planet was coloured red. Thus, users could quickly determine the main influence a vote had and who's vote caused a song to be played. The non-zero weights contributed by the five most influential users were displayed in text pop-up information (see Figure 4.1). Played songs had a darker planet colour. Twenty five distinct voting colours were used before repetition occurred.

The currently playing song was identified by a space ship, which flew to the planet when the song began playing. Song and voting information about the current song was displayed at the bottom of the application. Pop-up song information appeared when the mouse rolled over a planet or star, displaying song metadata, weight, and voting information. Left clicking a planet or star selected it. For the studies, visualization zooming was toggled by a right mouse click, but was later controlled by rotating the middle mouse wheel.

PartyVote's visualization aimed to provide system transparency. Figure 4.2 demonstrates how the potentially playable song region changes following a vote. By visualizing how songs were chosen, users could make informed choices. Indecisive voters could choose to vote for popular music to appease the group. A user who disliked a music group could avoid voting for music in that band's region. Some people voted for a song simply because they liked it, while others chose a song to increase the probability of a region of songs being played. Users could change their vote by simply voting again using the same user name. The system could check for valid user

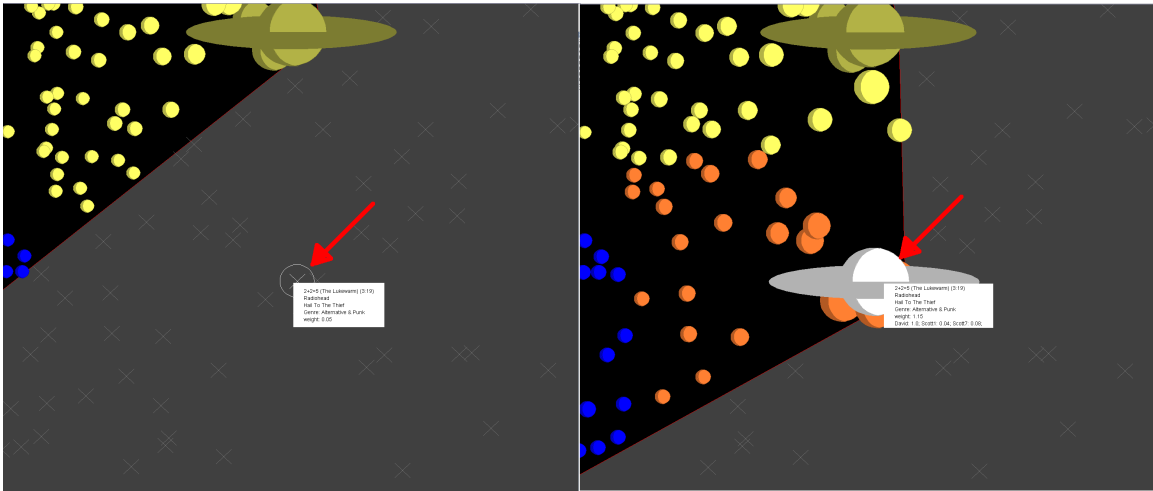


Figure 4.2: Alterations to the potentially playable song region based on a vote for “2+2=5” by Radiohead, indicated by a red arrow. Songs whose weights are primarily determined by the new vote are orange. The selected song is white.

names in a text file. By providing each party attendee a valid name at the start, multiple votes per user are preventable. This validity check was not enabled in our case studies.

A planet’s weight and the three primary contributing vote weights were provided as song text information. Voting weights potentially helped users find people with similar musical interests, identified people breaking group norms, and permitted strategic voting.

4.4.2 Algorithms

Multidimensional Scaling

If a similarity metric (or distance) between two songs is calculated, n songs requires $n - 1$ dimensional space for song icons to be positioned correctly in all cases. Thus, compromises to the layout must be made in 2D space. Techniques such as principal component analysis, self-organizing maps, and multidimensional scaling (MDS) are frequently used to map multidimensional spaces to 2D [89, 96]. We chose to use MDS because it is efficient, I expected it to provide a more globally correct solution, and it permits the use of any distance metric. Problems with the use of MDS (such as the loss of important patterns when reducing dimension) will be discussed in the ‘Lessons Learned’ section. Sound similarity and music metadata were used to calculate song

distances. Song layout implementation details are discussed later.

Voting Algorithm

Each user vote affected music selection in two ways: at least one song per vote was guaranteed to be played and music similar to the selection was more likely to be played. A vote indicated the selection of a set of one or more songs. The guaranteed song was randomly chosen from this set. Each vote had a weight of 1.0 and this weight was distributed evenly across all songs in the set.

The Playable Music Area

A convex hull of the guaranteed song positions was calculated using a 2D Graham's Scan and defined the potentially playable area of music [33]. A convex hull defined the minimal convex polygon bounding all user choices. Points within the hull defined a compromise between the votes. Numerous music regions could be used, however, we felt the convex hull resulted in more musical compromises between users and led to music discoveries. Similarly, a higher dimensional Graham's Scan could have been used but calculations would be extremely slow given a typical library size and the number of song similarity dimensions.

For each song voted for in the playable area, similar songs were weighed according to the formula

$$W_j = \sum_{i=1}^N W_i \div D_{ij}$$

D_{ij} is the Euclidean distance between songs i and j , W_i is the weight given to song i and N is the set of songs in the library. Thus, songs frequently voted for and songs near popular music are more likely to be played. We assumed this approach would satisfy most listeners while boundary voters got to listen to at least one song of their choice.

Music Playlists

Playlists were generated every time a new vote was cast. Guaranteed songs were played within two hours of voting, ensuring users heard the music they liked while delaying the time before they wish to vote again. In retrospect and based on user feedback, we believe guaranteed songs should be played within a half hour to promote positive reinforcement of the system for new users. Potential songs filled in the remainder of the playlist based on song weight.

4.4.3 Implementation

PartyVote was designed for typical hardware and personal digital music libraries. The PartyVote system was written using Java Swing. MP3 files were played using JLayer. In the user studies, the program was run on a 1.66 GHz Intel Centrino Duo Core laptop with Windows XP and 1GB of RAM. The mean refresh rate was 34 frames per second.

We tested the system using a music library of 3364 songs: the primary author's personal music collection. The library has 289 albums, 308 artists, and contains 19 genres. Because we designed PartyVote to be used on a standard PC with an individual's personal music collection, we believe this library is representative of real use.

Song Layout

Song metadata was collected using iTunes and saved in a text file. Metadata was weighted with artist, album, genre, song title and user rating weighed heavily. Other metadata included composer, song duration, year, track number, play count, file format, time last played, date added to the library, and bit rate. Each dimension was normalized between 0.0 (a perfect match) and 1.0 (no match). String pair differences were either Boolean or $1.0 - (\# \text{ of common characters} / \text{ text length})$. Numerical fields used absolute differences. Errors in the metadata such as the mis-spelling of a band's name, resulted in a less correct layout. We did not correct the metadata since we expect most music libraries to have similar problems.

Sound similarity was calculated using bextract and the Marsyas sound analysis library [129]. The first 120 seconds of each song were analyzed using Mel Frequency Cepstral Coefficients and Short Time Fourier Transforms to provide 68 attributes. Dimensions were normalized and Euclidean distances between pairs of song vectors were calculated.

Normalized metadata and sound similarity distances were combined in a variety of ways and visually tested before the first study to ensure a good layout. Layouts were calculated using metadata only, sound similarity only, the metadata and sound similarity mean, evenly weighted metadata, and album/artist/song metadata only to name a few. The squared normalized average was ultimately chosen to remove variability differences between sound similarity and metadata distances and to emphasize tightly clustered song groups. In the resulting layout, music genres were tightly clus-

tered, multiple copies of the same song with different metadata were positioned close to each other, and music from eclectic artists like Radiohead were more distributed than music from consistent sounding bands like Interpol.

Songs were laid out in 2D using MDSteer [140] without steering. MDSteer is designed to handle large data sets like our 3364^2 distance matrix. Overlapping points were slightly shifted programmatically to eliminate occlusion. Layout files were pre-computed and read at PartyVote’s startup.

4.5 Case Studies

Two mixed method case studies were conducted to evaluate PartyVote in a realistic environment. Study 1 looked at a large, loosely knit group who used the system for approximately one hour. Although loosely knit groups are not our target users, study 1 was designed to evaluate usability and gain initial feedback from a large number of people. Study 2 looked at a small established social group who used the system for five hours, and was designed to assess real system use. The case studies also provided initial feedback on music voting effectiveness, informing future design decisions. Both studies were designed to evaluate PartyVote’s use in a realistic environment using a participant observer inspired approach [26]. Video cameras or a stranger at a party can alter group behaviors, confounding results [107]. For both studies, ethics forms were signed when participants arrived (see Appendix A).

The PartyVote studies were conceived to evaluate realistic system use as the system itself was considered our primary research contribution. Thus social interaction confounds were unavoidable and our analysis focused on system usability rather than our higher level research questions presented in Chapter 1.

4.5.1 Case Study 1: Coffee House

Case study 1 took place during a monthly social gathering called a ‘coffee house’ for students, staff, and faculty members at the University of Victoria’s computer science department. Thirty five people attended and each person signed a consent form upon arrival provided they wanted to interact with the system. PartyVote played background music for the event, and attendees were encouraged, but not required, to use the system. The laptop computer described in the implementation section was used with an additional 19 inch LCD monitor and computer speaker set. Program

events were logged and a two-page questionnaire was given to people as they left. Questionnaires consisted of Likert scale questions ranging from one to five with a sixth “not applicable” option, and long answer questions (see Appendix A). This style of case study is similar to previous evaluations of Meme tags [19] and Vister [56].

Unfortunately, most people at this event treated PartyVote as a system demonstration. People tended to investigate the system rather than use it to mediate music choice, potentially because previous coffee house events had involved research demonstrations. Some participants blindly voted and one person intentionally increased the convex hull to see what would happen.

The coffee house study did provide some interesting results, however. Twenty four of the attendees used the PartyVote system and seventeen users voted one or more times. Eight people voted two or more times. People reported they enjoyed 71% of the party’s music on average ($N = 24, S.D. = 31.87$). We found no significant correlation between using the system and reported music enjoyment. However, Pearson correlations revealed that music recognition was correlated to the number of votes cast by the user ($r(29) = .401, p < .031$) and how active a role users played at selecting music at traditional parties ($r(29) = .473, p < .010$). This result suggests that users who recognized the library’s music vote more or that voting improved recognition. These same people were more likely to select music at other parties according to our survey. A negative correlation between feelings of empowerment and difficulties finding recognizable music was found ($r(16) = -.522, p < .038$). The perceived effect each vote had on the music played correlated with music enjoyment ($r(21) = .574, p < .007$).

4.5.2 Case Study 2: House Party

A second study was conducted to determine how a small established group would use the PartyVote system at a real party. The study occurred on a Friday night between 8pm and 1am. The party was held in a two-bedroom apartment and took place during one of the semi-regular house parties held by the observed group. No attempt was made to curb activities such as food and alcohol consumption, although participants remained sober. The system setup was nearly identical to study 1, except an external monitor and external speakers were not used.

Participants

Six people plus the primary author attended the party. Four participants knew the author in advance but only the host knew the author well. All participants signed a consent form when they arrived at the party, however, they were permitted to decline participation and still attend the party. We permitted this option since people presumably arrived for the party rather than the study. Participant names have been replaced with pseudonyms.

Amy and Scott are computer engineers, Pete and Bill are electrical engineers, Judy is a carpenter, and Mary is a staff member at a local university. Amy, Scott, Peter and Bill individually discussed their experiences in semi-structured interviews one week after the party. Our analysis centers primarily on these four participants. Recorded and transcribed interviews lasted between 18 and 59 minutes with a mean duration of 36 minutes.

Amy was 28 years old and regularly hosted house parties. She regularly played pre-made mixed CDs in hopes that the diverse music would appease everyone. Amy has modest musical training in piano and likes alternative music.

Scott was 32 years old and regularly attended Amy's parties. Scott rarely helped choose music played at parties due to the time involved and claimed to not concentrate on the music. Scott stated that he is generally indifferent to new music until he has heard it several times. Scott has no formal musical education and likes alternative and rock music.

Pete was 42 years old and occasionally attended Amy's parties. He has no musical training. Musical diversity was most stressed by Pete and he claimed to like all music genres.

Bill was 49 years old and regularly attended Amy's parties. Bill did not usually choose music played at parties because he believed his musical tastes do not overlap with most groups. Bill is a classical guitarist but has played in a rock band and a big band for over ten years. Bill had a professed dislike of popular music.

Findings

Amy, Pete and Bill were present during the first hour of the party. Scott arrived at approximately 9pm while Mary and Judy arrived at 11pm and midnight respectively. Bursts of system interactions occurred approximately when new users arrived (see Figure 4.3). Figure 4.3 also shows that no users interacted with the system for 45,

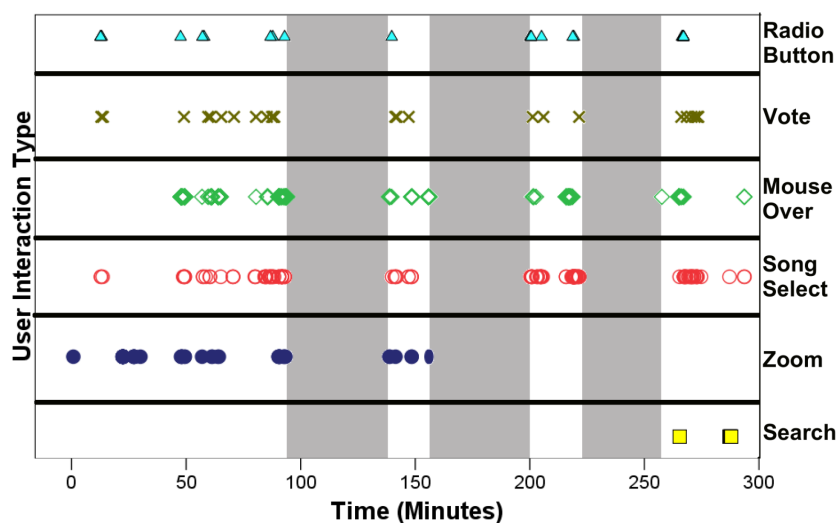


Figure 4.3: The distribution of user interaction events throughout the party. Note the three extended periods in gray when no user activities were recorded.

44, and 36 minute periods. Fifteen periods of inactivity greater than five minutes were identified. The figure illustrates that the system was essentially ignored for a great deal of time at the party and was often not the central focus. System use therefore matched our system design expectations. Voting occurred in bursts with people voting multiple times, suggesting that participants were interested in how each other voted, and may have attempted to influence the music played in their favor. Amy, when asked how frequently she used the system, stated:

“It gets addictive. Especially when you hear the crappy songs from other people...I was ignoring it when people were ignoring it or when the music was good.”

Opinions about the music library and the music played varied widely. Pete, who wanted musical diversity, liked a third of the music played and disliked none of it. Amy enjoyed 80-90% of the music played and disliked the 10% she attributed to Bill’s selections. Bill’s dislike of popular music may explain his modest enjoyment of 20-30% and dislike of 20-30% of the music. Finally, Scott was indifferent about 70% of the music corresponding with his indifference to unfamiliar music. All users felt their vote affected the music.

Social Pressure and Music Choice

Several participants reacted negatively to an Emile Campagne Quebec folk song nominated by Bill. Amy and Scott vocalized their dislike of the song immediately and asked who voted for it. Bill immediately confessed but expressed no regret. This event was a key point of discussion. Amy stated during her interview:

“I think for the most [part] I enjoyed [the music] except for [Bill]’s votes. When songs were selected from his part I didn’t like it because he selected like jazz and French songs.”

Similarly, Scott stated, “Some people seem to have some strange choices. Like those French songs.”

Conversely, Bill did not enjoy the music chosen by Amy and Scott but stated he liked Pete’s selection since “he chose the Pink Floyd thing.” Pete stated that he did not notice how others voted at all. Amy and Scott particularly enjoyed each other’s music selections, which is unsurprising given their overlapping tastes. All four interviewees stated that they did not feel any social pressure to change their vote. When asked why he voted for Emile Campagne, Bill stated “I’m comfortable with the people. In a different group, I wouldn’t have.” Bill’s experience demonstrates that peer pressure did not affect behaviors as much as we expected.

Multiple Votes

Participants were informed that multiple votes with the same user name resulted in vote replacement. Multiple names per user soon became the norm and this policy did not seem to create conflict, despite altering fairness. Instead, people informed each other that a new user name was required for each vote. Pete was the first to vote twice and did so without talking to the group. In the end, Pete had five distinct user names, Bill had three and Scott had two. Amy used two names but voted again after her song was played. In retrospect, it seems clear that some users may want to interact with the system more than we expected. It is important to provide a mechanism to allow varied levels of system use while remaining fair. Users should be able to more precisely articulate their areas of interest if they desire, but a single vote interaction should continue to be an option.

Voting Strategies

Only Amy claimed to have strategically voted to alter the influence of other users.

“At one point I knew that [Bill]’s votes were in the corner and I tried to find songs that were the opposite corner of it. Something that I liked though. I didn’t just blindly vote for something that was the opposite corner.”

Given Amy’s strategy to actively negate Bill’s votes, it’s not surprising that Bill, when asked what he disliked about the system, said:

“The ability for people to play demographic games or statistical games. Where you look at how somebody’s voting and you vote counter to that. You’re not voting your taste. You’re voting against somebody else’s taste. Or the group’s taste.”

Bill said that voting games could be amusing sometimes, but he was concerned about ‘music bullying’.

Bill himself voted by defining the musical area he was interested in. He found three or four songs he cared about and assumed this would nominate the desired musical area. Bill’s mental model of the PartyVote system was probably slightly incorrect. Marking an area of interest succeeded in making desirable songs playable, but songs between distantly spaced votes are unlikely to be played.

Scott and Pete simply voted for representative songs in a style of music they wanted to hear, without trying to affect other votes. However, Pete attempted to change the potential song space (enlarge the convex hull) to increase the musical diversity

In the end, the visualization was used in a variety of ways. It was used extensively by Amy and not at all by Scott. Pete used the visualization simply to see the music diversity and Bill used it to identify the area he was interested in.

Interface Limitations

Several interface limitations were identified. Planet colour was unused except by Amy, who was more familiar with the system than the others. A legend mapping colour to votes may clarify this feature. Text based song weight information was also unused by these first time users. Planet size and positions made sense to all users except Mary.

Scott, however, had difficulties imagining how numerous dimensions of information could be represented in 2D. Both Scott and Amy felt there were 'inaccuracies' in the layout since distances between some songs were not what they expected. Scott was also annoyed by the accuracy of the similarity-based layout, since his vote for Nirvana led to a neighboring Sarah McLachlan song being played. Similarly, Amy did not believe the layout accurately represented song similarity since Nine Inch Nails (electronica) and Nirvana (rock) were in distinctly different regions but she considered them both to be 'alternative'. Thus, developing a mental model for how the system operates was problematic at times. We suggest that the visualization's axes should convey meaning to aid understanding. I will discuss this further in the 'Lessons Learned' section.

4.6 Discussion

Here we discuss how well PartyVote meets the system design goals that we presented in the 'Problem Area' section.

Support Individuals & Appease the Group

PartyVote provides an equal say to each user, supporting both individual preferences and the group's overall music choice. Because we did not make an experimental comparison, we cannot directly determine whether PartyVote increases music enjoyment. That said, participants indicated they enjoyed 76% of music during the first study and disliked less than 12% of the music on average during the second study. Given the unconventional music in the music library, these values seem quite high. We believe that group tastes were generally satisfied. Individual preferences were also supported. Users who voted believed that they affected the music selected. Moreover, all users at the house party voted even though three of the interviewees rarely selected music at parties.

PartyVote is designed to provide one vote to one user, ensuring fairness. Multiple votes per user potentially disrupt the balance between individuals and the group and allows self-designated DJ behaviors to occur. Pete's five votes gave him more influence on the music than Mary's one vote. Further system refinements are required to handle this issue. I suggest enabling multiple votes to define music set with each song in the set gaining equal weight.

Leverage Existing Social Dynamics

Designing for social conventions is more nuanced than we expected. Bill decided to vote in opposition to the group's musical taste because he was comfortable with the others. Pete voted multiple times without concern about what others would say. Amy and Scott felt comfortable vocalizing their dislike of Bill's selection. Users may feel comfortable experimenting with the system because social sanctioning is uncommon in their group. Hence, contrary to our expectations, group dynamics alone may not constrain user behaviors. It seems that the circumstances surrounding system use help dictate what is considered appropriate behavior. The importance of context on system use is discussed much more thoroughly later but this result helped guide our subsequent studies.

Minimize & Simplify Necessary Interactions

Observations from our studies suggest that our interface was generally intuitive but stood to be improved. Most people intuitively understood basic system use, including how similarity based on planet position worked, and how voting affected nearby music. Planet size and positions mapped intuitively to weights and similarity for most people. Planet colour, song layout, and the voting procedure could be improved. In particular, participants were confused as to how music similarity could be calculated, and how it was presented in two dimensional space. Furthermore, they expected the two dimensional axes to have a meaning. MDS therefore does not seem to be an ideal method for positioning songs. MDS also results in potentially important relationships being removed as the dimensionality decreases. Worse yet, the relationships removed cannot be actively chosen but instead are random. Choosing two meaningful dimensions with which to distribute the music (or enable people to choose their dimensions) may be preferable. Further testing is required.

If the perceived individual costs to using PartyVote are less than traditional music selection methods and perceived benefits are similar, we can consider our system a success. When asked why he would use the PartyVote system but not choose music at a conventional party, Scott responded:

“Well at the same time, the time commitment is only one time, it's not the whole evening changing music the whole time. It's totally different.”

Scott's comment is encouraging. We also note that some participants wanted

to use the system more than we expected, voting multiple times. Permitting more than minimal interactions is necessary to address user diversity. However, prolonged periods with no system use did demonstrate that minimal interactions were required.

Use Common Hardware & Personal Music

PartyVote was tested using a laptop and a large personal music library. System performance should scale linearly with the library size if layouts are pre-computed. Thus our approach should remain viable for conventional hardware as music collections grow.

4.6.1 Design Implications

PartyVote generally allowed individuals to express their music preferences while also appeasing the group. The visualization was informative for people that used it, enabling strategic voting. However, we observed several unexpected user behaviors that should be considered in future designs.

More Than Minimal Interactions

A large annoyance to party attendees was the method of voting. Users in study 2 were not initially informed that each new vote with the same user name replaced the earlier vote. Five of six users attempted to vote multiple times and Bill stated that he wished to define an area of acceptable music. The repeat voting norm suggests that votes using a single name should be concatenated and each user should be allowed multiple votes. More importantly, some people actively seek to refine their choice or play with the system more than others. Hence, the system should support minimal interactions for some users and allow more in-depth use for others. To enable people to play with the system or satisfy their curiosity without immediately influencing the music, we suggest adding ‘undo’ functionality and the ability to preview what affects a vote would have in subsequent system designs.

Future versions of PartyVote will allow small groups to have more guaranteed songs per user. Participants suggested that users should log in to vote and see their selections. Finally, when someone leaves the party, the host or that user should be able to remove their vote. House party attendees also wanted vote weights to decay over time.

Peer Pressure

PartyVote is designed to use group norms and peer pressure to ensure fairness. Bill suggested that traditional social pressures can be a blessing and a curse. Music bullying can potentially occur when one or more users determine the music played either by manipulating the system or by socially dominating other users. Bill believed our visualization may increase this possibility since users can easily identify how others voted and they can devise intelligent ways to nullify another person's vote. Bill suggested an optional visualization and anonymous voting would prevent music bullying in groups where this is a problem. Alternately, voting could take place before the event, making the visualization simply informative.

Social Interactions

When Amy discussed the differences between CD based music selection and choosing music using PartyVote, she suggested that PartyVote may isolate users from the rest of the party:

“If you look at CDs, it's easier to interact with the rest of the party than if you're at the computer by yourself. Because if you're with the CDs you can say [things] like 'what is that' and they can like sing it for you or hum it...which I have done.”

Amy's comment raises an interesting question: how can we integrate a computer system into the social fabric of a party without demanding that attention is paid to it? Large screen displays could make voting less private, but this may exacerbate social pressures and conflicts. A large screen may also distract users. However, presenting information to all users would permit voting collaborations and prevent social isolation when using the system. Users may have entertainment rather than utility based goals and therefore the importance of different system features may change. Displaying the 'winning' or most popular songs may help reduce social isolation and be entertaining. Some participants in both studies discussed PartyVote in competitive or game-like terms.

Alternately, mobile phones and tablet PCs could be shared amongst party attendees. Although this opens up the possibility of a device becoming lost, informal passing and sharing of of a portable device would enable people to vote without

moving to a particular location. Voting would not prevent relaxing on the couch or socializing in the kitchen.

Group norms were difficult to predict. Optional system constraints may allow a host to dictate the style of interaction similar to putting away a CD collection before a party to prevent friends from changing the music. House party participants wanted the ability to veto other votes and vote multiple times. Voting may be considered a game or personal statement of self. Providing options like anonymous voting and the ability to veto would allow customizations for the expected group behaviors.

4.6.2 Lessons Learned

Ultimately, what participants wanted to accomplish with PartyVote appeared highly reliant on what their system goals were. People playing games with the voting mechanism may want to view the ‘winning’ vote while people just wanting to help determine the music with minimal work may simply want to be able to find their favorite song quickly. Some people voted simply to determine how the system worked. Furthermore, not everyone wanted to thoroughly examine the music list or visualization. The key is to either identify what people wish to use the system for and design for those uses, or to constrain the system to one particular form of activity (such as a game based voting activity), thereby constraining user behaviors.

PartyVote’s interface was perceived as complicated and complex to some people particularly in terms of the 2D plane having no conceptual meaning. System complexity seems to have led to avoidance by some and playing with the system for others. Ideas such as multidimensional scaling seem to be overly complex for casual use and people wished to assign some logical meaning to interface elements (like the music map axes) they did not understand. Conversely, some system aspects such as planet colour and voting weights were effectively ignored by many participants. Immediate usability appears to require the immediate establishment of a clear and comprehensible mental model. Novelty also seemed to be a major promoter of system ‘playing’.

The ‘playful’ elements of the system were perhaps the most surprising findings for PartyVote. Although we did not expect the space travel analogy to be a major point of appeal, we expected it to be aesthetically pleasing or entertaining. Instead, people did not seem to care about it at all and the extra complexity may have hindered system understanding. Although PartyVote was used in informal environments, the

system itself was being used as a tool to choose music or out of curiosity. The system itself was not the entertainment for most people. In retrospect, this seems obvious, however, many current casual InfoVis systems appear to make this same mistake [22, 34].

Designers cannot expect people to spend much time to figure out how the system operates. Our participants looked at the system and if they did not understand its functionality they either ignored features, took trial and error actions, or asked others how it worked. The trial and error actions taken by some means undo functionality and administrative rights might be important aspects for group casual InfoVis systems. Undo functionality would reduce or eliminate the costs associated with exploring a system.

The PartyVote studies both clearly demonstrated that social factors can influence system use and that these factors are extremely complex and difficult to predict. PartyVote was used in spurts during the house party, suggesting that the actions of one person motivated others to use the system, either to counter another person’s vote (implied) or because they were reminded to use the system (explicitly stated). However, social norms did not constrain user behaviors as expected, with each participant perceiving socially acceptable actions differently. Repeated evaluations of PartyVote might make these social norms more explicit but our assumption that established social norms would translate to system behaviors does not appear necessarily justified. Furthermore, given how recent casual InfoVis research is, it seems prudent that future studies should concentrate on single user applications when possible.

4.7 Revisiting the Research Goals

The PartyVote studies also enabled us to address three of my four primary research goals. The optional nature of PartyVote’s visualization means there is substantial overlap between these goals and traditional human-computer interaction research, since data comprehension is not necessarily the primary concern surrounding system use. However, this neither negates the importance of our findings nor are we aware of similar patterns reported in the HCI literature:

1. **Explore why people use visualizations during leisure time:** As mentioned previously, we were surprised by what people wanted to use PartyVote for (goals) and how nuanced group norms were for influencing motivation. For

example, Amy performed strategic voting, Scott effectively ignored the visualization, and Bill was unconcerned about his friends disliking his chosen music due to his comfort with them. During the house party, PartyVote was used for choosing music and exploring the music collection, but was never directly a source of entertainment. Utility and learning based goals therefore seemed to motivate users the most. Social factors also seemed to motivate system use, as evidenced by the temporal voting pattern observed during the house party. This is similar to results reported by Hinrichs *et al.* [59]. The two studies also clearly illustrated how contexts of use could affect use patterns.

- **How important is initial usability and system simplicity?** This question was not a primary consideration of the study, although our results strongly suggest that initial usability can greatly influence visualization use. Two of the house party participants were initially reluctant to use PartyVote because they did not think they would get it. We also noted several examples of playful behavior and experimentation with the system during the two studies. We expect that these actions would be reduced with subsequent PartyVote use. However, these findings are quite informal and more direct initial usability and simplicity testing is needed.

2. **Perceptions of utility, aesthetics and entertainment relating to system use:** PartyVote provides a variety of different ways someone can approach system use. Entertainment or playful features like the space metaphor were provided, and the system itself could be used to play strategic ‘voting games’. PartyVote can also simplify music selection at a party, improving utility. Interestingly, many users chose to ignore peer information and simply voted. Visualization aesthetics and playfulness did not seem to influence participant opinions and motivations, but utility and system complexity did.
3. **Casual InfoVis Methods:** We learned important context lessons about how to conduct a casual InfoVis system study. We hypothesize that the coffee house study was not effective for evaluating the system because the social group was not sufficiently close, the study duration was too small, and the traditional associations between the coffee house and demos meant the system itself was the focus of activity (not the music or socializing). The subsequent study addressed these issues but novelty effects potentially still affected behaviors. Thus repeat

realistic system evaluations using the same participants may be ideal. The PartyVote studies also clarified the importance of several casual InfoVis metrics for me. System logs, although useful, were considerably noisier than expected. Quantifiable metrics such as number of clicks, and number of songs selected were not applicable because each user approached system use differently (including minimal use) and these patterns changed over the study’s duration. Interviews, however, were extremely important as system use patterns differed from expectations yet questions could be tailored to the individual. Quantifiable metrics may be useful in future studies, but systems catering to a wide variety of user goals should perhaps be explored with more qualitative methodologies.

- **How can one conduct a realistic evaluation of casual InfoVis systems when system novelty and the presence of the researcher influence use patterns?** The initial one hour evaluation of PartyVote revealed that many participants played with the system to understand it or out of curiosity rather than to use it to select music, indicating a clear system novelty effect. The second evaluation subsequently examined PartyVote being used by a small group of friends in a realistic setting where the researcher’s influence was minimized and the long study duration reduced novelty effects. Study results, however, suggest that the novel nature of the system continued to influence participants during the house party. I propose that repeat evaluations with the same participants would ideally minimize research and novelty influences to motivation.

4.8 Conclusions and Future Work

PartyVote provides a lightweight mechanism for established social groups to choose music at a party, relying on group norms and participant visibility to ensure fair system use. Guaranteed songs and song weightings allow the majority to determine the general style of music played while safeguarding individual choice. Unlike previously reported recommender systems and music jukeboxes, no user profiles are needed and constant user vigilance is not required.

PartyVote successfully enabled participants to influence the music during our two case studies. Even users who rarely choose music at conventional parties contributed to music selection using PartyVote. However, our evaluation of PartyVote quickly

revealed complexities to Casual InfoVis system design unique to the area, particularly in terms of what people wanted to use the system for. Group norms also developed in unexpected ways. Our results suggest that multiple votes and varying amounts of system use should be supported. Users also did not constrain their votes due to peer pressure as much as we expected and concerns over music bullying arose. More testing is needed to clarify our findings but our initial results suggest that a transparent music voting system may be effective for choosing music at parties.

Chapter 5

Study 2: Motion in Casual InfoVis and the Interrelation Between Personality, Performance, and Preference

5.1 Introduction

The PartyVote study described previously left numerous unanswered questions about the ideal design principles for casual InfoVis. For example, it was unclear whether the space ship analogy or planet motion aided the system's 'fun factor' or were just distracting. Based on existing casual InfoVis system designs, utility may not be the sole concern associated with system use. For example, the Burst Labs music player [22] and the digg visualizations [34] presumably use motion to increase system appeal at the possible expense of performance.

For the current set of studies, I sought to explore the performance costs and potential benefits of designing with a more hedonic or enjoyment based emphasis. More precisely, my supervisor and I wondered if users would choose a system that has known performance costs to receive some other benefit. We chose to study motion because it has been commonly used in casual InfoVis, presumably to increase system appeal, yet it has been shown to be distracting and not particularly effective for data encoding. Bartram *et al.* [11] and Robertson *et al.* [111] both noted that motion was appealing to users but was less effective at conveying data than static encodings.

The current set of studies was also complementary to our previous PartyVote study, as motion was previously assumed to increase system appeal.

Study 2A was designed to verify that the motion encoding in our glyphs was indeed readable, albeit with a notable performance decrease. Study 2B explored how potentially distracting motion could be for an visual search task unrelated to the motion encoding. Study 2C then qualitatively examined how user preferences and opinions are affected in situations where motion is visually appealing yet causes a utility decrease. Thus we were not studying potential performance benefits of motion data encoding; rather we looked at fun/functionality trade-offs using motion as the context.

5.1.1 Motivation for Study Approach

Results from the PartyVote study revealed that some of the intuition based recommendations for casual InfoVis design may not be justified. In particular, PartyVote’s space metaphor and use of motion were designed to make the system visually appealing, yet many participants found these features either benign or annoying. I wanted to empirically test the potential costs and benefits of using a common casual InfoVis design choice, thereby clarifying when such a feature is appropriate. A series of such studies could then elucidate user motivations and provide clear design recommendations for future casual InfoVis systems.

The subsequent results from the motion study, however, motivated me to shift my research focus to a holistic exploration of factors influencing visualization use. The motion encoding study, while retrospectively too tightly focused for an early investigation of my research goals, provides a illustrative example of how a specific regulating factor (motion aesthetic appeal) is influenced by user goals. Regulating factors will be discussed in more detail in Chapter 6.

Most of the previous motion research has been performed in highly controlled environments, independent of the final proposed application domain. This approach is effective for determining pre-attentive cognitive abilities and the task independent effects of motion. However, we wished to examine complex user-system interactions with cost / benefit trade-offs. We felt that a more application specific approach was necessary but we still wanted a degree of control to permit quantitative analysis.

Two primary research goals were addressed in the following ways:

1. **Perceptions of utility, aesthetics and entertainment relating to system**

use (Goal 3): Performance (utility) and aesthetic trade-offs with regard to system use are at the heart of the motion encoding study.

2. **Casual InfoVis Methods (Goal 4):** The motion encoding study was an initial attempt at performing a controlled empirical evaluation of motivations, use patterns, and opinions regarding casual InfoVis. As such, this experience helped mold the quantitative methodology suggestions found in Chapter 3

5.2 Previous Literature

We define *motion* in the current context as the perceived sequence of changes in a visualization’s image over time. Motion may include icon movements, transitions between system states, and colour changes [127]. We further define *extraneous motions* in information visualizations as any motion that does not encode data directly used in a task. This includes icons moving randomly [34], extra movement during an animated transition, and icons that change colour or shape for aesthetic purposes. *Motion based data encoding* refers to the use of temporal changes to encode relevant information [138].

5.2.1 Benefits and Drawbacks of Motion

Previous research has shown the benefits of motion for some contexts. Ware and Bobrow [138] demonstrated that repetitive sinusoidal motions around a glyphs can be used to encode data. Furthermore, the movement frequency, amplitude, and phases could all be used to encode information, with subjects being most sensitive to relative phase patterns and least sensitive to oscillation frequency differences. Motion is perceived preattentively, meaning that identification of a moving object within a person’s field of view does not rely on conscious recognition. Motion attracts attention and requires conscious effort to be ignored [10]. Bartram *et al.*[11] suggests that motion should be used sparingly as it may distract users from important information.

Motion may also indirectly affect performance through aesthetics. Motion can be visually appealing to many users resulting in potential increases in system use. The halo effect suggests that aesthetically pleasing interfaces are described in more positive terms, even if the positive trait has nothing to do with the interface [5]. Norman’s [93] three process theory of human judgment suggests that behavioral evaluations involving logic and utility (behavioral) are influenced by previous experience (reflective)

and aesthetics (visceral). We therefore expect motion’s appeal and evocative nature to provide aesthetic and visceral benefits that will improve perceived usability. For casual use, we thought motion’s appeal may be more beneficial than the distraction it causes.

5.2.2 Uses of Motion

Motion is traditionally used in visualizations for two main purposes: to alert users [11, 46, 138], and to provide animated data transitions [15, 115]. For example, Gluck *et al.* [46] used animation to provide strong alert signals to subjects performing a concentration task. They found that fast motion was noticed more quickly than oscillating colours or slow motions. Animated transitions were shown by Bederson and Boltman [15] to help users map out or understand a family tree relationship viewed one family member at a time. Research on smooth transitions showed that subjects made fewer errors in judging connectedness in node-link diagrams when seeing smooth animations compared to discrete image changes [115].

Motion data encoding research has been limited due to well documented distraction concerns [10]. Instead, motion based data encoding is normally used temporarily, often as a means of data highlighting. The preattentive perception of motion makes it effective for filtering and brushing subgraphs in node-link diagrams, thus encoding subgraph relationships [137]. Ware and Bobrow [138] suggest that static data encoding may be more efficient for displaying information in the users foci since there is no need to temporally track a glyphs state. More recently, Robertson *et al.* [111] have shown motion in trend visualizations to be effective and highly appealing for demonstrations but ineffective for analysis tasks. For a brief review of motion based encoding methods see work by Ware and Bobrow [138]. Our current work extends motion based data encoding by examining its use in casual InfoVis and to what we believe to be a rare, semantically appropriate application domain for its use: music visualizations.

Extraneous motion has been used stylistically by both academics and industry. For example, the Vizster Friendster visualization uses randomized node movements to increase system appeal and reinforce the system’s informal nature [56]. The InfoLotus email visualization uses a slowly growing flower metaphor to represent the emails received, and distinguishes urgent emails using an continuously animating bee [145]. Song icons in the Burst Labs visualization [22] drift constantly within an enclo-

sure seemingly for stylistic reasons alone. Similarly, the Fizz visualization represents Facebook posts using bubbles that drift and move in 2D space, yet bubble movements are presumably for aesthetic purposes only [18]. The YouTube bubble visualization [42] uses long animated transitions to make video nodes appear to drift relative to one another. Long animated transitions and extraneous planet motion used in tag galaxy’s visualization of Flickr tags help support the solar system metaphor used but not the user’s primary tasks [143]. Finally, the digg visualizations [34] use motion extensively when representing recent digg messages and their popularity. These movements attract attention, indicate changes in content popularity, and may add a playfulness and style to the information. Motion may also help people monitor new articles, which appear suddenly and move the most. This motion may help users gain a more general awareness of what is available rather than concentrating on a subset of posts. However, it is currently unknown whether motion improves system appeal, how substantially motion affects data readability and how motion based distractions affect user opinions. Our studies attempt to address these unknowns.

5.3 Research Overview

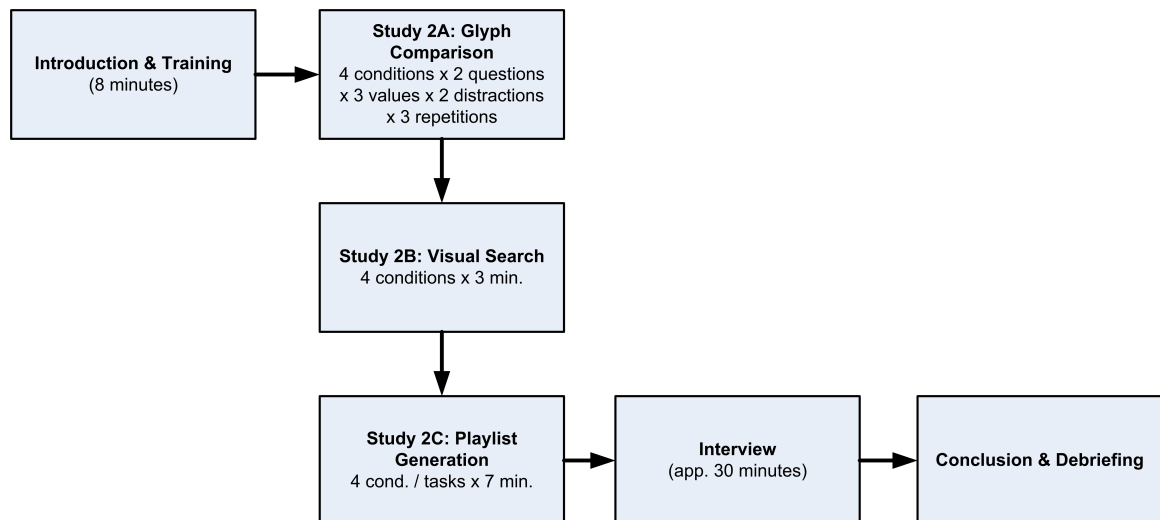


Figure 5.1: The ordering of motion encoding sub-studies for each research session. The total session took approximately 2 hours.

Most of the previous motion research has been performed in highly controlled environments, independent of the final proposed application domain. This approach

is effective for determining preattentive cognitive abilities and the task independent effects of motion. However, we wished to examine complex user-system interactions with potential cost / benefit trade-offs. We felt that a more application specific approach was necessary but we still wanted a degree of control to permit quantitative analysis.

In an attempt to tease apart relatively objective performance measures from user opinions, research was conducted using three complementary sub-studies (see Figure 5.1). First, we measured the ability for participants to read and interpret motion based data encodings. Study 2B tested how distracting motion was during an orthogonal visual search task. Finally, study 2C required participants to create a musical playlist with beat strength and tempo information (the motion-encoded data) being important to the task. Subjects were then interviewed to examine their opinions about motion-based data encodings. Subjects participated in all three studies during a single two hour session. Hence, we expected subjects to be completely aware of the performance costs associated with the motion conditions (from study 2A) during studies 2B and 2C. This allowed us to investigate how people interpreted potentially appealing stylistic features with a known utility cost.

5.3.1 Instructions, Training, and Consent

I met subjects in ECS 555 (the Usability Lab) and briefly explained the purpose of the study. Subjects were asked to read and sign a consent form (see Appendix B) before they were received a brief tutorial on the system used in studies 2B and 2C. The tutorial consisted of a brief description of system functionality and a demonstration of the four beat encoding techniques. Subjects were asked to experiment with all system functionality, explore the music library, and locate familiar songs for eight minutes. Built in speakers were used for music playback. Subjects were reminded about unused system functions in the last 2-3 minutes of practice. Subjects were financially reimbursed for their time at the end of the study or whenever they chose to withdraw, provided the consent form was signed.

Subjects

Sixteen subjects aged 19-39 (*Mean* : 25.33, *S.D.* = 7.9) were recruited for this set of studies. People were recruited using online postings, group emails, and paper posters distributed throughout the city of Victoria (see Appendix B for recruitment

materials). One working engineer, nine undergraduate and six graduate students participated. Six of the subjects were female and four were not engineers or computer scientists. The subject sixteen session was conducted twice with two different participants because the first subject did not adequately follow instructions (his data was removed). Condition orders were counter-balanced for all tasks using a Latin squares design, which repeated every four subjects.

5.4 Study 2A: Glyph Comparison Task

5.4.1 Method

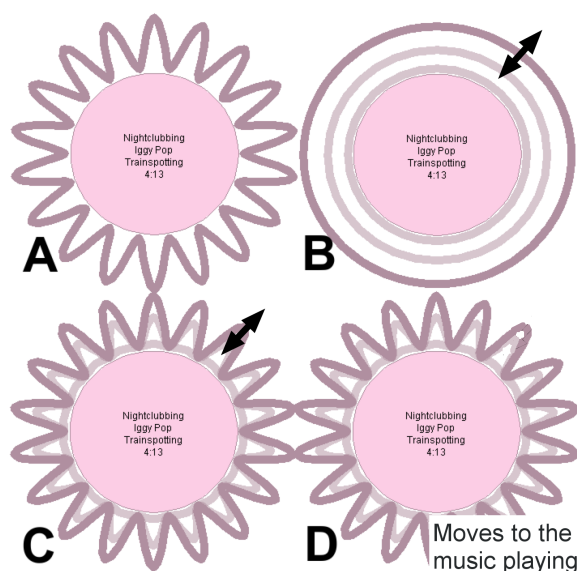


Figure 5.2: Song glyph representing the song ‘Night Clubbing’ by Iggy Pop. Beats per minute and beat strength encodings differ between conditions: static (A), motion (B), redundant (C), and extraneous (D). Faded lines indicate movement paths.

We conducted a traditional forced decision task using a within subjects design with four beat encoding conditions. For consistency with the playlist generation and visual search studies (both described later), each glyph represented a song with song tempo (**beats per minute** or **bpm**) and intensity (**beat strength** or **BStr**) being mapped to an animating pattern or a sinusoidal line pattern (see Figure 5.2). Conditions were:

- **Static Encoding (S):** A sinusoidal line encompassing the glyph represented tempo using the number of line oscillations and beat strength using amplitude

- **Motion Encoding (M):** an expand/contract animation of a circle encompassing the glyph represented tempo using the animation frequency and beat strength using movement size or amplitude. All animations oscillated at the same frequency as the beat they represented and followed a sinusoidal movement rate. Thus the motion encoding was a temporal equivalent of the sinusoidal line curve.
- **Redundant Encoding (R):** Tempo information was encoded redundantly using a sinusoidal line that expands and contracts the same as the motion encoding.
- **Extraneous Motion Encoding (E):** The static encoding method was used to represent the glyph’s beat information with the line animating with the beat of the currently playing song. Thus, when a song was playing, the glyphs all appeared to ‘dance’ to the music.

Using a sinusoidal line enabled us to keep conditions almost identical except for the trait being manipulated. The movement’s repetition also addressed the apprehension principle [127]: sound has traditionally been mapped to sinusoidal wave patterns and to repeated movement. The motion encoding should have also provided an appropriate cognitive mapping to a beat’s regular pattern (particularly when motion and beat were synchronized)[92, 6]. Thus we expected the encoding methods to be more intuitive than changes in colour or luminosity. Bartram et al’s [11] work on motion and distraction suggests that growing and shrinking motions will be less distracting to users than translation, rotation, or explosive motions.

Procedure

The task consisted of subjects performing a series of timed pair-wise glyph comparisons to evaluate performance differences between conditions. The test application was written in Java and had a refresh rate of 40Hz for all conditions. Tasks were conducted on a 1.66 GHz Intel Centrino Duo Core laptop with Windows XP and 1GB of RAM.

Trials followed a 4 (encoding condition) \times 2 (question type) \times 3 (value) \times 2 (distraction) design. Each configuration was repeated three times. Conditions orders were counter-balanced for all tasks using a Latin squares design, which repeated every four subjects. For each trial, a question was provided, users pressed the space bar to

start the trial, and after five seconds without a decision, the beat encoding line was removed to force a decision. The keyboard arrow keys were used to choose between the left and right glyphs. Subjects were asked to answer as quickly as possible while maintaining a greater than 70% response correctness (displayed at the end of each condition based set of trials). Glyphs had no text, were colour coded by question type, and glyph bodies were approximately 5cm in diameter. Question data was synthetically generated a priori based on beat data ranges from the library used in studies 2B and 2C. Short breaks were provided between sets. Experimental factors were:

- **Encoding:** Static, motion, redundant, or extraneous beat encodings. Extraneous motion rates were constant but did not match the current glyph’s bpm
- **Question Type:** Two questions were asked: “Which song has the fastest tempo (beats per minute)?” and “Which song has the stronger beat?”
- **Value:** The difference in beat data between glyph pairs was held constant at 31 for bpm and 3.37 for BStr questions. Beat values between trials varied between high (BPM : 184/153, BStr : 16.06/12.69), medium (BPM : 113/144, BStr : 11.85/8.48), and low (BPM : 73/104, BStr : 7.63/4.26)
- **Distraction:** For each question type, the other beat information dimension was either identical between glyphs or differed by the pair-wise beat difference (31 or 3.37)

Trials were grouped by condition and question type with bpm and BStr questions randomly chosen to start. Value and distraction factors were randomized within a subset. Beat differences were made sufficiently large (31 bpm) to avoid minor frequency differentiations (which are difficult to distinguish [22]). Glyph readability was measured using task correctness and task time. A training period was provided at the start of the task, where subjects were required to correctly answer four beat strength questions and four bpm questions with each encoding method.

5.4.2 Hypothesis

Beat data will be read most effectively using static encodings, followed by redundant and extraneous encodings. Static curves can be read immediately and do not require

monitoring over a period of time like motion encoding. Movement phase issues are also irrelevant with the curve encoding.

5.4.3 Results

Table 5.1: The effects of question type and distractions on mean trial time and accuracy.

Question Type	Trial Time	Accuracy
BPM Distraction	1769ms	90%
BPM No Distraction	1637ms	95%
BStr Distraction	1918ms	92%
BStr No Distraction	1705ms	98%

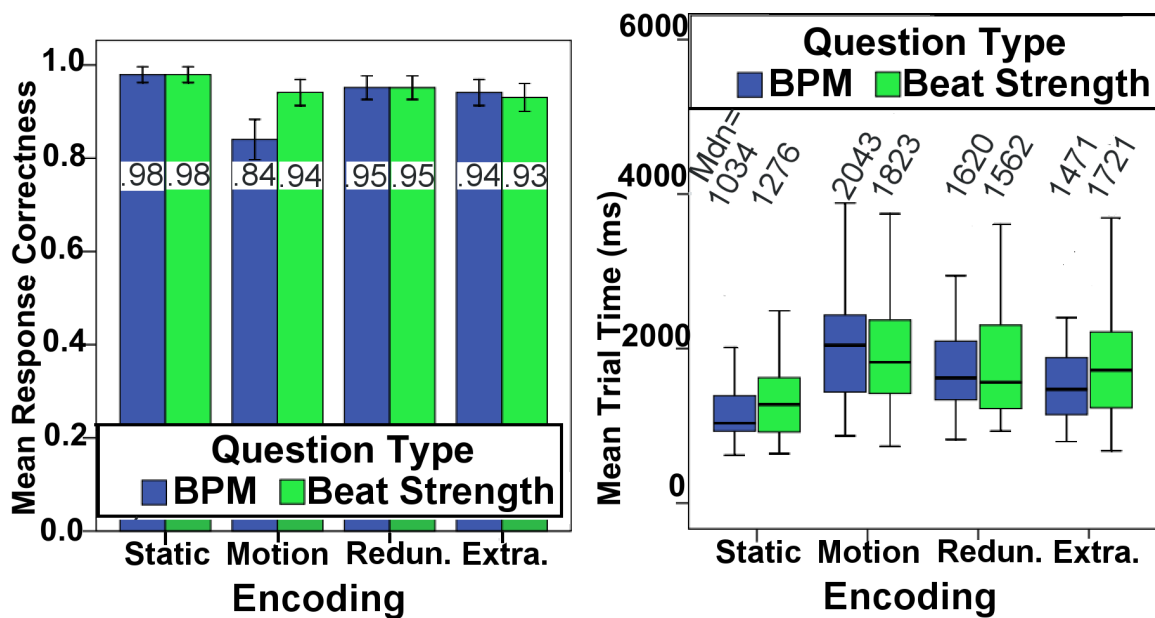


Figure 5.3: Glyph comparison task results showing the interaction between question type and conditions on mean trial time (milliseconds) and correctness (%/100). Error bars indicate 2 standard errors.

Experimental results confirm our hypothesis. Significant differences in trial time ($F_{TimeBPMCond}(3, 39) = 13.282, p < .001$, $F_{TimeBStrCond}(1.603, 6.451) = 20.844, p < .001$) and task correctness ($F_{IsCorrectBPMCond}(3, 39) = 15.713, p < .001$) were identified between encoding conditions, with post-hoc analyses showing static data encoding being significantly faster than the motion ($p < .001$), redundant ($p < .001$)

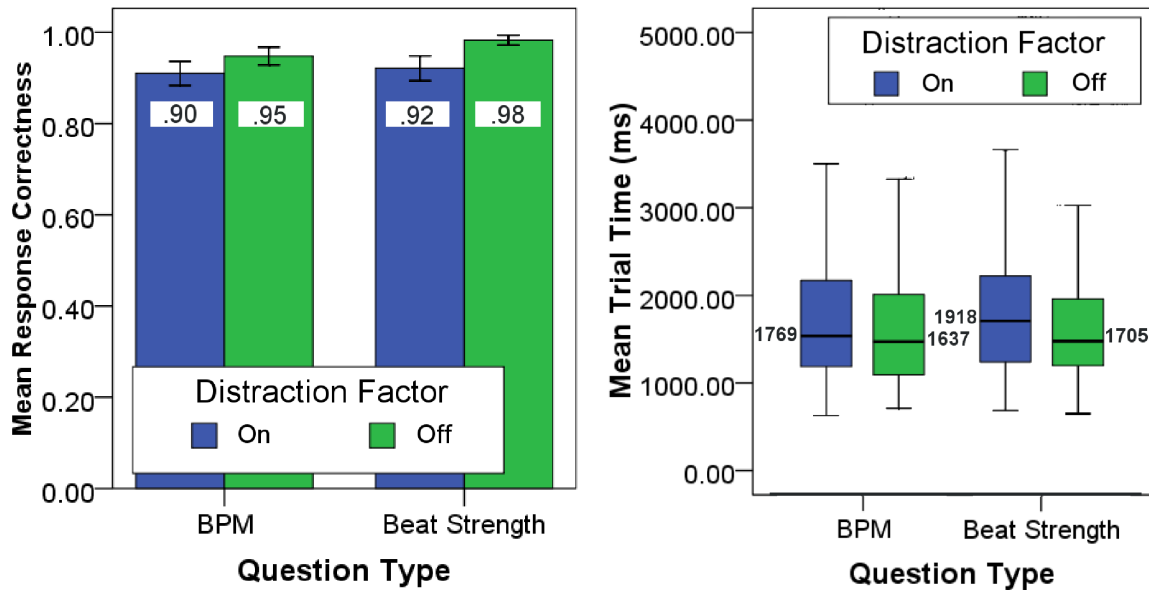


Figure 5.4: Glyph comparison task results showing the interaction between the presence of a distraction factor and the question type on mean trial time (milliseconds) and correctness (%/100). Error bars indicate 2 standard errors.

and extraneous ($p < .018$) conditions (see Figure 5.3). Static encoding was significantly more accurate than the moving conditions ($p < .001$, $p < .023$, $p < .012$ for motion, redundant and extraneous respectively). Motion encoding was also significantly slower than the extraneous condition ($p < .031$) and less accurate than both the extraneous ($p < .002$) and redundant conditions ($p < .001$). Trials with a distraction were slower ($F_{TimeBPMDis}(1, 13) = 15.344$, $p < .001$, $F_{TimeBStrDis}(1, 13) = 13.676$, $p < .002$) and less accurate ($F_{CorrBStrDis}(1, 13) = 5.194$, $p < .038$) (see Table 5.1 and Figure 5.4).

5.4.4 Discussion

The glyph comparison task results clearly show that static encoding was significantly faster and more accurate than the three motion conditions, particularly when distraction factors were involved. We believe that the stability of static encodings simplified data reading, thus improving task scores. Alternately, the motion in the extraneous and redundant conditions may have distracted subjects from reading the sinusoidal line. Subjects during the redundant trials may have chosen to read the data using the motion encoding much of the time rather than use the sinusoidal curve. The

ephemeral nature of motion means that glyph evaluations must be timed and likely requires working memory [138]. Task performance using the motion encoding was significantly worse than redundant and extraneous encodings, suggesting that reading data encoded in motion is more difficult than reading the curve encodings (even when animating). Our results suggest that it is better to have motion encode nothing and act as a distraction than to have it solely encode information.

The significant effect of distraction suggests that encoding multiple dimensions of information into motion can lead to problems. The distraction effect on BStr reading is likely caused by the phase discrepancies between glyphs such that the peak movement positions are not simultaneously reached for the two glyphs and visual memory is required to respond. Beat strength distractions may affect bpm questions since pixel movement rates increase with beat strength. These data demonstrate that encoding more than one piece of information using motion can hinder data reading, even when the motions are considered conceptually orthogonal.

Study 2A not only provided us with a quantifiable ‘cost’ for our motion encoding method, but it also made subjects aware of the performance differences between conditions during studies 2B and 2C (rather than being told). At the end of each set of trials, the subject’s accuracy was displayed and the time required to perform a trial was regularly commented on by subjects. Study 2A enabled us to examine if participants preferred motion encodings despite knowing the draw-backs.

5.5 Study 2B: Visual Search Task

The visual search task was intended to test the effects of motion on user distraction in what we thought was a somewhat realistic task (for example, searching through a friend’s CD collection for new music you may enjoy). The visual search task required subjects to visually explore the music map and locate songs from a particular album. Study 2B did not require the use of beat information and therefore tests for condition based distractions, ignoring the ability to read motion based encodings. We expected this approach to identify possible extreme distraction effects. Given that results were not significant, the study will be discussed briefly.

5.5.1 Method

Study 2B was conducted immediately following study 2A and involved the same four glyph conditions. Characteristics unique to this study are described below.

Subjects were given three minutes to search for all songs in a given album. Four different albums from a single musical artist were used: Elvis Costello, Loretta Lynn, Counting Crows, and Neil Young. Artist order was kept consistent across subjects but condition order varied (but was the same as study 2A). Each album had 13 songs and each artist only had one album in the library. Target songs were marked using a double click. No subject found all 13 target songs. The extraneous condition glyphs animated to randomly chosen songs. The experimenter monitored subjects from another room using a video camera, a feed of the laptop screen and an area microphone; all of which were recorded. The training stage, described previously, also got participants accustomed to this study’s interface, however, no practice searches were given to participants.

Software and Equipment

The music library visualization was written in Java and displayed music glyphs on a zoomable two dimensional (2D) music layout which we call a *music map* (see Figure 5.5). The system provided an editable playlist and traditional ‘iTunes like’ navigation (which was disabled during trials). Music came from my personal music collection. The text based navigation was only enabled during training to ensure the music map was used during testing. The applications updated at 30Hz to ensure consistency between conditions but update rates occasionally decreased when zooming occurred. A USB optical wheel mouse was used to pan (click and drag) and zoom (mouse wheel).

Music beat analysis and the music map layout were performed in advance. The final 1432 songs were chosen to be as diverse as possible while emphasizing recognizable music like the Beatles. Human and computer beat identification errors are usually double or half the actual value [129]. We filtered out songs not in 2/4 or 4/4 time so disparities between the perceived beat and computed beat would remain multiples of each other, limiting confusion. Euclidean distances between each pair of songs were calculated, with both song metadata and sound similarity metrics, using the same formulation as used during the PartyVote study (see Chapter 4). MDSteer [140] then performed multidimensional scaling to generate a two dimensional layout for these vectors. Glyph positions were programmatically scaled and shifted with

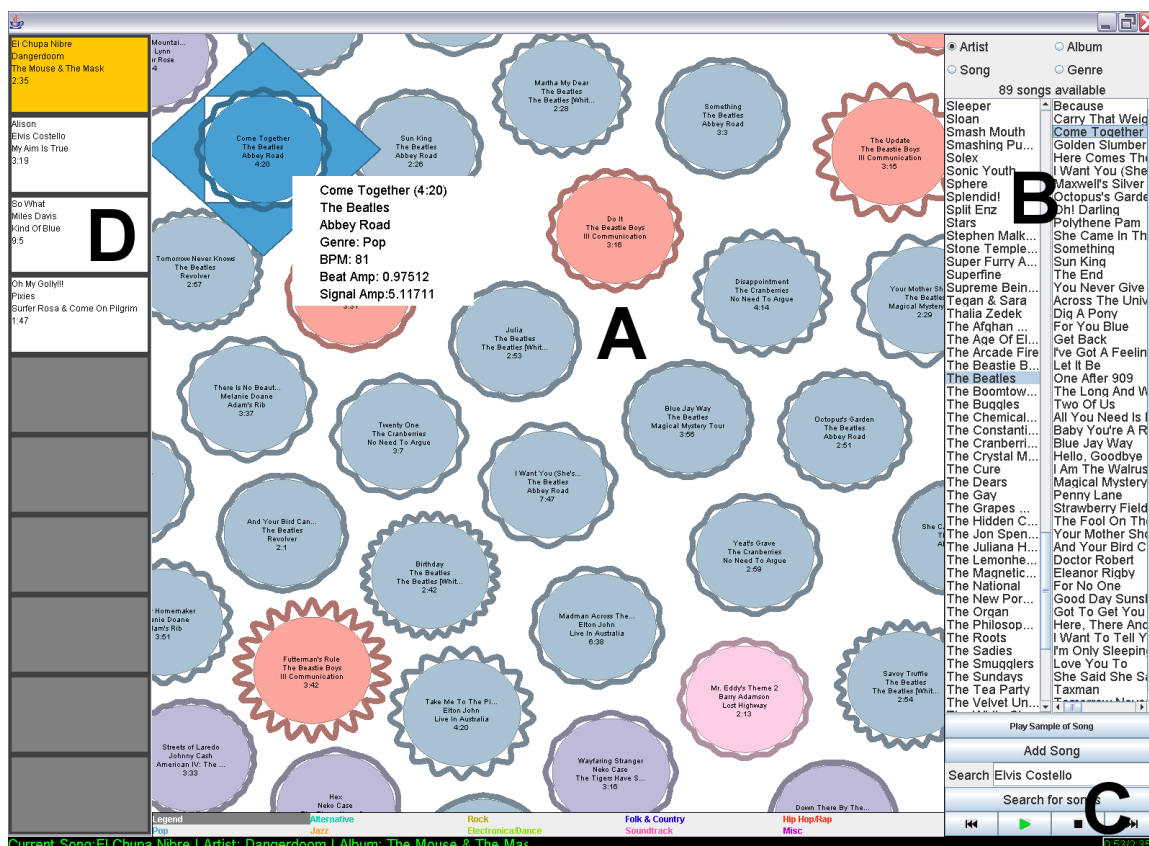


Figure 5.5: The visualization used during the visual search tasks and playlist creation tasks (study 2C). Features include A) a zoomable genre coloured music map, B) navigation window, C) music controls, and D) the current playlist.

experimenter guidance to ensure that they did not overlap. Marsyas' bextract application [128] calculated song beat information for each song in the music library. For each song, I validated beat measures and manually adjusted the values when the computed result was clearly incorrect.

5.5.2 Hypothesis

We believed motion would be modestly distracting, drawing user attention away from a concentrated search task. Thus, during the visual search task, more songs would be located using the static condition than the moving conditions.

5.5.3 Results and Discussion

Our hypothesis was not confirmed from the visual search task data. There were no significant differences in the number of songs found between conditions ($p < .533$). Furthermore, the Partial Eta Squared effect size was .086, suggesting a very small effect size or no effect. Eight subjects found the motion distracting but most believed that they could ignore it during the search task and the playlist creation task. Since the distribution of songs was not held consistent between tasks, this lack of effect could be due to experimental noise. However, there is no evidence suggesting that motion based distractions would reduce performance in study 2C.

5.6 Study 2C: Playlist Generation + Interviews

Our third study was a qualitative examination of the trade-offs between aesthetics and performance, and how these affect user preferences. A controlled quantitative experiment would not adequately address the complexity of our research concerns. We also wanted to examine casual visualization use under normal circumstances. Subjects performed a music playlist generation task using a custom made music library visualization (discussed in the previous section), as this seemed sufficiently complex, realistic, and open-ended to examine user motivation and thought processes during the follow up interviews. We explored behavioral patterns using semi-structured interviews, observations, and recorded user actions, using a method similar to grounded theory. See work by Isenberg *et al.* [65] and Creswell [26] for a more thorough discussion of grounded theory.

Procedure

The playlist creation tasks required subjects to find music that they felt was appropriate for a given situation like a romantic evening or a party. I monitored subjects from another room using a video camera, a feed of the laptop screen and an area microphone; all of which were recorded. Condition orders were the same as study 2A. The semi-structured interview, immediately following study 2C, was conducted for approximately a half hour. Subjects gave background information, their opinions about the playlists they created, and their glyph preferences.

The interface used was almost identical to the one used for the visual search task (discussed above). However, the music was no longer muted. Double clicking a music

map's song icon added it to the displayed playlist (see Figure 5.5). Playlist entries could also be re-arranged and removed during study 2C.

Playlist Generation Task: Subjects were asked to imagine they were looking through a friend's music collection to create a list of songs for four different occasions: a romantic evening, a party with 20 attendees and perhaps some dancing, studying with a personal music player, and exercising with a personal music player. Subjects were allowed to modify these tasks enough for the activity to be realistic for them. For example, exercise music was chosen to be appropriate to the exercises subjects performed, such as running or lifting weights. Occasion task orderings were consistent across subjects but condition order varied. Tasks were intended to explore how beat information was used for finding slow music (studying & romantic) versus fast music (party & exercise) and group music (romantic music & party) versus personal music (exercising & studying). Each subject was given seven minutes to create a playlist, however, it was stressed that the number of songs found was unimportant and subjects were to approach the problem the way they normally would at a friend's house. We chose this approach to ground the experimental task in a more realistic context.

Participants were informed that playlist music would start automatically as an added incentive to choose music they liked. Subjects were also encouraged to play song samples before adding music to the playlist to reduce the frequency of add/remove operations and to improve music appropriateness. Subjects could control the music volume as long as it remained audible. Beyond this, no formal instructions about how to perform the task were provided, as subject approaches were deemed as important as how conditions affected performance. We also wanted behaviors to be as natural as possible.

Interviews were analyzed by repeatedly listening to the audio recordings, noting interesting or important quotes, identifying subject biographical information, and iteratively tagging subject comments. Tags were generated ad hoc using open coding and then iteratively defined and categorized using theoretical coding [27, 65]. I grouped subjects based on commonalities of their interview responses such as their desire to discover new music, how frequently beat information was used, and their educational background. Subject log files were also analyzed for patterns and commonalities between subjects. I then attempted to group and cluster subjects and behaviors based on the interviews and log data. My notes identifying group commonalities are provided in Appendix B. Our grounded theory clustering approach

inherently includes some subjectivity, but enabled me to discover emergent patterns in the inherently noisy results.

Despite using a predominately qualitative approach, some metrics were collected to explore the relationship between user action and opinion. Subjects ranked the four glyph encodings and indicated their music collection size. We also recorded the frequency at which a subject sampled songs before adding them to the playlist. We use this as an indirect measure of music novelty as familiar music would not need to be sampled first. Subjects were asked about their search criteria for each task to identify when beat was important. They were also asked whether they used the beat encodings to find music. We consider cases where beat was important but the beat encodings were unused to be a *beat contradiction*. Beat contradictions show that the subject did not use the tools provided when it may have helped.

5.6.1 Guiding Questions

Grounded theory emphasizes the process of pattern discovery rather than hypothesis testing. As such, study 2C’s design was built around two guiding questions:

- **Q1: How do productivity and visual appeal affect casual InfoVis system expectations and opinions?** It is not clear when motion is preferred by casual InfoVis users if there is a clear trade-off between visual appeal and practical utility.
- **Q2: What factors outside of pragmatics and hedonics (e.g. experience, context, and user goals) influence user opinions and approaches to casual InfoVis use?** Q2 was a secondary question, but determining what pre-existing and task specific factors most strongly affect opinions and approaches may aid future casual InfoVis designers.

5.6.2 Findings

As expected, user reactions to the various visualization conditions were varied. However, user evaluations and comments appeared to fit into three main categories, suggesting that system appreciation is highly dependent on previous experiences, fun/practicality evaluations, and personality. These groups are discussed in the ‘Subject Groups’ section. We also found that subject strategies and approaches to music

selection were complex and varied, despite most subjects pursuing similar experimenter defined goals (discussed below).

Every condition ranked highest for at least two subjects and lowest for at least two others. The static encoding was generally preferred; ranked first by six subjects. However, six other subjects rated static encoding as one of their least favorite, despite knowing they performed better using it during the glyph comparison task. This result addresses Q1: productivity and utility were not the only factors determining preference for all subjects, however, it was considered very important for half of them. Beat use was associated with the playlist type, however, it was not associated with encoding condition. If motion was substantially more appealing than the static encodings, we expected beat information to be utilized more during the three motion conditions (Q1). We found that the answers to our research questions corresponded greatly with the subject groups, discussed below. Music playlists and songs sampled did not correspond to subject groups and revealed few usage patterns as these data were extremely noisy. Furthermore, no correlation or relationship was found between task and condition preferences.

Log data and interviews demonstrate that song selection approaches were strongly influenced by previous experiences, the perceived usage context, and user goals. I will limit my discussion to several illustrative examples. Beat information was more frequently used during the exercise and party tasks (fast music) while interviews revealed that subjects sought unfamiliar music when accommodating group tastes, but infrequently for personal music. Romantic evening playlists predominately had people searching for songs about love, jazz, or slow music. Party music was typically hip-hop or rock and chosen to appease the subject's friends. Study playlists were perhaps the most musically diverse between subjects, but most participants sought non-distracting music. However, what distracted participants varied greatly. Strategies included searching for music without lyrics, familiar music, unfamiliar music, and music the subject could not easily sing along with. Exercise music choices depended on the exercise chosen but subjects predominately chose faster music or music with a stronger beat. As such, several subjects enjoyed using the beat encoding more for this task.

Previous experiences appeared to affect how people approached playlist creation and how they perceived the data encodings. Subject 1 designed the party playlist with musical ebbs and flows due to his music background, while subject 10 simply looked for more popular or recognizable music. Subjects 4, 11 and 12 believed that music

choice could not be simplified to a couple key dimensions like beat based on their experiences and thus did not use the beat encodings. Subject 2 frequently noted his dislike of blinking html to explain his dislike of extraneous motion. Similarly, higher level cognition and anthropomorphic image interpretations seemed to positively affect motion appeal. Subject 15 stated that the moving glyphs “seem happy that you’re using it” while subject 3 liked the redundant condition because it “looks like it’s alive.”

Subject Groups

Consistent patterns of behavior, opinion, and background were found between subject subsets. From audio recordings, we found that subjects fell into three major groupings. These groupings corresponded to user comments about pragmatics, hedonics, and how well conditions mapped to the user’s cognitive model about music. Thus, it seems that subjects split along practicality and visual appeal attitudes. Interestingly, these groupings also corresponded with music interests, desire to find new music, gender, and beat contradiction frequency. Each group will be discussed in turn. See Table 5.2 for a summary.

Pragmatics: “If [motion] was an option, I would turn it off. It’s something I don’t need.” (subj. 2)

The eight subjects concerned about task performance, functionality, and utility were exclusively male, the vast majority of subjects had a technical background (eg. studying computer science or engineering), all but one made explicit comments about the importance of productivity and system utility, and all but three had the static encoding as their favorite condition (two of the which thought redundant encoding improved their performance). Comments like “static is fine” (subject 4), “if animating then [I want] the useful animation” (subject 2), and “[motion was] more distracting than anything else” (subject 11) illustrate the general attitude about motion. Only subjects 2, 4, and 12 did not chose the motion only condition as their least favorite condition. Pragmatic subjects were substantially less likely to use the beat encodings during playlist creation (only used in 11 of 32 trials). Beat was considered an important criterion in 11 of 21 trials where beat encodings were ignored (beat contradictions). Seven of the eight subjects claimed to search for familiar music rather than discovering new music. For example, subject 11 said it was “easier to find things [he] recognized.” Curiously, seven of the nine alternative music fans were pragmatics.

Hedonics: “Although [motion] was the most distracting, it was certainly the most fun.” (subj. 16)

Hedonic participants typically mentioned the importance and appeal of beauty, fun, or aesthetics relating to the beat encodings. Subject 14 thought the static condition was “weird that it just stands still.” Subject 13 similarly said “if [the glyphs are] moving, then I think it’s more pretty.” The static condition was not the favorite condition for any hedonic group members. Three of the five subjects were women, and the majority had a non-technical background. According to playlist and sampling data, hedonics were more likely to play a sample without adding it to the playlist than other subjects (Samples added: Hedonics = 18.75% of 672, Others = 25.68% of 693), suggesting they might be finding more unfamiliar music. Subject interviews further support this, where 4 of the 5 subjects who actively sought new music were in this group. Beat contradictions were rare with only subject 14 claiming that tempo was important while ignoring beat information. Four of the six pop music fans were members of this group.

Intuitives: “[Motion] gives you the information more directly.” (subj. 3)

The intuitives group found that motion was a more direct, easier to read, or intuitive way to encode beat information than the curved lines. All three subjects were women. Subject 9 stated that motion encoding “kind of clicked” while subject 15 stated “if it’s not moving, then I didn’t pay attention to it.” The redundant condition was preferred by subjects 3 and 9, potentially suggesting that users thought the motion was intuitive but they still pragmatically wanted the static encoding to fall back on. Intuitives used beat encodings the most frequently, utilizing it in ten of their twelve trials. Only one beat contradiction was identified.

5.6.3 Discussion

Study 2C results show that motion only improved system appeal for a subset of the subjects, while others were acutely concerned about system utility and efficiency. Furthermore, users can frequently have conflicting desires and goals. Even within the context of a constrained music selection task with clearly formulated instructions, individual freedoms associated with casual InfoVis resulted in extremely divergent results, opinions, and perceptions about system effectiveness.

Subject explanations for their behavior tended to relate to their existing beliefs and goal directed actions. For example, subjects 4 and 11 believed music choice could not be condensed to several key measures like beats per minute. Subject 6 avoided zooming to enable music browsing by genre, based on that measure's past effectiveness. To address individual differences, visualization customizations may be beneficial but time consuming. Provided subject groupings are possible, high level customizations based on group memberships could greatly improve user experiences.

Effects of productivity and Visual Appeal(Q1)

As expected, productivity and functionality were not the exclusively important factors for system appeal. However, the factors affecting appeal were far more complex than a simple fun/functionality trade-off. During the interviews, subjects 1 and 5 believed they performed better using the redundant encoding. Incorrect performance perceptions suggest either 'the more data the better' preconceptions or motion's appeal resulted in a discrepancy between objective metrics and perceived performance. Norman [93] suggests positive aesthetic judgments can influence perceived performance.

Our results suggest that motion was not as appealing to participants as expected, although this was group dependent. Although many subjects stated they enjoyed one of the moving conditions, subjects were equally likely to mention negatives like decreased productivity and distractions. Negatives may counter motion's potential aesthetic benefits. Six hedonic and intuitive group members reported enjoying the motion despite the distractions while pragmatics tended to concentrate on the negatives. It is unlikely that participants did not notice the performance differences during study 2A given their reactions during testing. However, not noticing performance costs could have led some subjects to prefer moving conditions. We were also surprised at what subjects considered 'useless' or distracting and how subject specific this was. For example, subjects 5 and 16 found redundant encoding problematic but could "tune out" extraneous encodings, while subjects 2 and 7 found the opposite.

For hedonics, the emphasis on fun or style seems to be reflected in their approaches to the task. Hedonics more frequently sampled songs without adding them to their playlists and, as a group, they more actively sought out new music rather than familiar songs. Users wishing to be efficient would conceivably tend to look for familiar music and only choose to sample music if there was a good possibility it would be added

to the playlist. Listening to thirty second samples of new music would decrease the number of songs chosen. Pragmatic users' large personal music libraries may also have led them to a learned emphasis on utility. Searching for familiar music rather than sampling unfamiliar music may be an acquired approach due to their music collection size. Alternately, pragmatics might be more interested in owning music (and presumably listening to the same music repeatedly) than other subjects.

Previous Experience, Context, and Task (Q2)

We found clear indications that subjects approached the playlist creation task differently depending on the task particularities, previous experiences, and the context. This may not be a surprising result by itself. However, given the importance of practicality and pleasure on user opinions, the diversity of interpretations as to what was useful and appealing is intriguing. We note subjects approached the playlist creation task differently depending on task specifics with group music leading to more music exploration and fast music tasks leading to increased beat use. Interviews revealed that perceived distraction and the utility of data presented was influenced by training and previous experience. Thus utility and enjoyment appear to be important but mitigated by factors such as user goals and training. These factors may not have been identified in a more controlled experiment.

Overall, it appears that the perceived ease of use and perceived utility of the system greatly influenced user preferences and opinions. Our results seem congruent with the Technology Acceptance Model (TAM), which states that perceived ease of use and perceived utility are the primary factors that predict a person's desire to use a technology or system [16]. However, our results suggest that utility and ease of use may be extremely subject and context dependent. We note that user goals, as dictated by the task, tool, and the person involved seem to be key factors determining participant opinions. Latham and Locke's [81] Goal-Setting Theory similarly predicts that user goals, developed through mechanisms like choice and required effort and moderated by factors such as feedback and task complexity, directly predict motivation to use a visualization. Explanations using Goal-Setting Theory, however, may just shift the research problem upstream to studying how visualization factors affect goals. The current study's findings directly motivate our subsequent focus on goals and motivating factors. Goal-Setting Theory and user goals are explored in more detail in Chapter 6.

Subject Groupings

The fact that subjects were clusterable along pragmatic / hedonic lines both implies that fun / functionality trade-offs are potentially being evaluated by participants and that there may be circumstances where users are willing to sacrifice utility for hedonics. Motion may help attract ‘eyes on the screen’, even if people are aware of a functionality cost. System preferences might be addressed using group-based visualization customizations. Furthermore, our results suggest that with future research, categories of similar thinking users might be identifiable based on markers such as music collection size, background training, and system preferences. The ideal markers for identifying user groups is an open research question. If systems could account for an enumerable number of groups in a target population, this may be sufficient to address most individual user needs. The set of all possible customizations and permutations would not need to be supported. Similarly, less user information may need to be collected to identify user opinions. For example, knowing a user is a philosophy major who finds motion appealing may mean she wants to discover new music, she is unconcerned about task efficiency, and a more stylized visualization would be appreciated. An engineer with a large personal music collection may want a more pragmatic visualization. More research is needed to establish whether consistent subject groups can be identified and whether these groupings exist in other application domains. We are not aware of any such research.

The underlying mechanism causing the observed clustering is unknown, however, numerous explanations are possible. Personality characteristics may dictate how a user approaches system use. Users concerned about utility may be unwilling to try new features or new music because it may impact their perceived performance, while hedonics may prefer novelty and thus the moving glyph encodings. Caution may explain the reluctance of pragmatic group members to use beat encodings and their emphasis on functionality. Perhaps the lack of beat encoding precision also frustrated pragmatic subjects. Adventurous users may want to find new music and hence sample far more music than they add to their playlists (e.g. the hedonics). Subject backgrounds, training/education, and gender would logically affect opinions before the study, and these experiences may affect goals, responses, and performances more than in-study training. These goal differences could then affect which system features are important. This concept is discussed further in Chapter 6. Previous experience differences would explain the group differentiations along music genre, background,

opinion, and education lines. It is currently unclear why music genre co-clustered with our groups and this remains an open research question. There do seem to be more factors influencing motion opinions than we originally suspected and one or more of these factors may be the underlying reason for the group clusterings. However, for effective group based customizations, the reason for the grouping is not required, so long as a membership predictor can be identified.

Gender and occupational differences between the subject groups seem to be the most alarming result from our grouping scheme. Pragmatics and intuitives broke perfectly along gender lines but this does not necessarily mean gender affects preferences. First, there were significant age and occupational differences between men and women which may explain this result. Men were older and were predominately from technical fields like engineering and computer science. Women mostly had non-technical backgrounds. Second, groupings were based on user comments so this may indicate a gender/culture based language difference. Substantial differences between small groups may have caused an inadvertent sampling artifact. Recruitment order and possible unintended changes to the study's methods over time offer another possible explanation, as most engineers were tested first.

5.7 Study Limitations

Examining what aspects of the study worked and did not work can be crucial to improving future research approaches. One limitation of the study's design was the use of the similarity based music map. The music map was employed since we needed an environment where glyphs could be positioned in space without the layout being based on beat information or genre (since that data was encoded on the glyph). The axis independent layout was a frequent source of confusion and contention. Subjects were informed that the song layout was based on song similarity but many subjects still assumed there was a one or two dimensional criteria for the layout and wanted to know what it was. Furthermore, subjects had difficulties re-discovering music located earlier with S15 describing the task as a 'treasure hunt.' Our results suggest that visualizations using dimensionality reduction may be a poor choice for casual InfoVis. Abstract concepts that users are unfamiliar with appear to frustrate many people.

In retrospect, our choice to use motion frequency differences to encode beat per minute information may have put the motion conditions at an unfair disadvantage as

these were difficult to differentiate when beats were weak. However, given that movement was synchronized to the music, this continues to seem the most natural and appropriate mapping between bpm and motion. Another motion encoding method seems unlikely to be more effective and the music library size meant we could ensure each song had sufficiently different beat information from the others.

5.8 Conclusions & Future Work

Results of these studies begin to examine the cost/benefit trade-offs between designing for appeal and designing for functionality. Our first experiment clearly identified performance costs associated with our chosen motion based data encoding. Subjects were aware of such performance issues in study 2C, where we clearly demonstrated that some participants were willing to sacrifice performance for visual appeal. Furthermore, subject opinions were strongly task and experience dependent. Three user categorizations based on glyph opinions were identified: pragmatics, hedonics, and intuitives. Differences in subject background, judgment, and actions corresponded with this grouping. Our findings suggest that a group based customization approach to casual InfoVis is a promising new research direction. Furthermore, knowledge about a user group may be critical for effective casual InfoVis design. Motion's appeal to some people may explain its use in some current systems. Motion may provide an effective means for attracting a particular subset of users, provided it is designed intelligently. However, many of our subjects expressed motion based distraction concerns. Care should be taken before introducing it into a visualization, particularly for entertainment / aesthetic purposes only. However, identifying why motion and other stylistic factors are only appealing during some circumstances or for some people is an obvious next step. This research direction is explored further in the subsequent chapter. However, for Chapter 6's investigation we wished to avoid constraining our exploration to one visualization or one set of factors; hence an open exploratory study of how people see data in their daily lives seemed appropriate.

Table 5.2: Subject groups and their distinguishing traits (in bold). Factors are (in order): gender, favorite condition, least favorite condition, the frequency a subject sampled a song but did not add it to the playlist, the number of times beat information was considered important but not used (beat contradictions), explicit claims a subject looked for familiar or new music, and the approximate size of the subject’s personal music collection.

Subject	Group	Gender	Fav. Cond.	Least Fav. Cond.	Sample Only Freq.	Beat Contra. (/4)	Find Familiar (/4)	Find New (/4)	Personal Music (Songs)
1	Prag.	Male	Redun.	Motion	Med.	1	1	1	4000
2	Prag.	Male	Static	Extra.	Med.	0	1	3	400
4	Prag.	Male	Static	Extra.	High	1	2	1	10000
5	Prag.	Male	Extra.	Motion	Med.	1	1	0	20000
6	Prag.	Male	Static	Motion	Med.	3	0	0	12000
7	Prag.	Male	Static	Motion	Med.	0	4	0	2000
11	Prag.	Male	Static	Motion	Med.	3	4	0	500
12	Prag.	Male	Motion	Redun.	Low	1	1	0	5000
8	Hedonic	Male	Motion	Static	High	1	0	0	1500
10	Hedonic	Female	Redun.	Static	High	0	1	4	1500
13	Hedonic	Female	Motion	Extra.	High	0	0	4	500
14	Hedonic	Female	Redun.	Static	High	3	0	4	1400
16	Hedonic	Male	Extra.	Redun.	Med.	0	2	1	6000
3	Intuit	Female	Redun.	Static	High	0	0	0	1000
9	Intuit	Female	Redun.	Motion	Low	1	4	0	2000
15	Intuit	Female	Static	Motion	Med.	0	1	0	2500

Chapter 6

Study 3: Visual Representations of Data During Non-Work Time - Modeling User Goals and Regulated Motivations

6.1 Introduction

The studies reported in Chapters 4 and 5 provided an initial and very specific exploration of casual InfoVis use. However, the motion study revealed that opinions and motivations surrounding casual InfoVis are more complex and involve more factors than we realized. Controlled testing of individual system features is ineffective provided unknown confounding factors influence the results. For study 3 we used a less controlled, more exploratory research approach to accurately identify how people use visual representations of data in their daily lives. We hoped this would reveal how various factors influence casual InfoVis use. A more general understanding of user motivations also required expanding the research scope to include numerous different visualizations and infographics rather than a single system. Understanding motivating factors should help designers to create visualizations that are more likely to be adopted by casual users.

Three primary research goals were addressed in the following ways by the Artifact study:

1. **Explore why people use visualizations during leisure time (Goal 1):**

We examined how people currently use visual representations of data during their leisure time and then extrapolated why visualizations may be used.

2. **Design a casual InfoVis model (Goal 2):** The main contribution of the artifact study is the Promoter/Inhibitor Motivation Model (discussed later), derived directly from the artifact study findings.
3. **Perceptions of utility, aesthetics and entertainment relating to system use (Goal 3):** Interviews and journal entries were analyzed to explore how participant opinions about utility, usefulness, aesthetics and entertainment related to why they examined an artifact. I also identified an emergent taxonomy of user goals based on participant interviews and related these to the duration and frequency of artifact use.

Our goal was to address why people explore data in their free time by observing casual use of visualizations ‘in the wild.’ However, most visualizations are not yet used by the general public. As such, observing existing visualization use may focus too much on early adopters’ habits and beliefs, which may differ substantially from general use. We instead examined how people use a superset of visualizations: visual representations of data (termed *artifacts* in this study). This superset includes non-interactive infographics, icons, and charts.

Results from our study revealed the importance of personally relevant data, designing for ‘productive relaxation’, and how opinions about system features changed depending on user goals opinions and motivations. The primary research contribution in this chapter is a proposed model of motivation, *the Promoter / Inhibitor Motivation Model (PIMM)*, which provides an initial framework to guide future casual InfoVis design and research evaluations. We use the model and our results to infer appropriate design considerations for casual visualizations.

6.2 Artifacts vs. Visualizations

Focusing on artifacts enabled us to identify user goals and explore usage patterns independent of a particular system. An assumption underlying our work is that motivations to use visual representations are very similar to motivations for using interactive visualization tools, even if aspects of the interaction differ. We argue that there is substantial overlap:

1. Data artifacts can be non-interactive and more simplistic than typical visualizations. However, data analysis and interpretation are likely required for the artifacts as they are for conventional visualizations.
2. There is a well-established research precedent (e.g. in grounded evaluations, cultural probes, and computer-supported collaborative work) for initially researching non-digital settings to help clarify what users may do in digital settings. My approach can help identify user needs and practices in the pre-design research stage [44, 65] and identify domain problems independent of system designs[90].
3. ‘Target’ Artifacts (defined later) were categorized as interactive (8) or static (78). Using a Spearman (non-parametric) correlation, we found no significant interactions between interactivity and user goals, data type (abstract / concrete / geospatial), or circumstances of use (necessary to use, scheduled use, or randomly encountered). Artifact use duration did correlate with interactivity and this is discussed in the Stages of Artifact Use section. Goals and use contexts categories were based on journal entries.

6.3 Casual Information Visualization Studies

Evaluations of visualizations specifically designed for casual use have been limited, and have typically used controlled laboratory study and survey approaches [53, 132]. More holistic evaluations have predominately examined early adopters and heavy users of particular systems, most notably Many Eyes [139, 135, 134]. For example, Danis *et al.* [30] noted that the participants they interviewed were not typical users: half were regular bloggers, most worked with data, and they all used Many Eyes far more than average. In addition, system evaluations may not sufficiently isolate casual users in particular. For instance, a qualitative examination of Many Eyes users revealed it was frequently used for a person’s job as a free tool [30]. Despite these limitations, however, we can still gain many insights about user motivation from these studies.

6.3.1 Factors Influencing Visualization Use

Previous studies have identified the following factors that influence casual InfoVis use:

Usefulness

Based on blog posts, Many Eyes was repeatedly used to serve some task based use, such as data analysis and research (Viégas *et al.* [135]). Even systems designed for reflection, identity construction / expression, and aesthetic appeal may be evaluated in terms of perceived usefulness (defined in Table 6.1). For example, Neumann *et al.*'s [91] visualization of keyboard typing style had users express a desire to read back a message from the graphic despite utility not being a consideration in the system's design.

Chart Junk, Aesthetics, and Self-Reflection

Ghahary [45] suggested finding the appropriate balance between aesthetics and utility. He noted that more abstract visualizations are often harder to remember and understand. Bateman *et al.* [14] similarly found that so-called 'chart junk' aids memory recall compared to minimalist and 'clean' data representations.

Numerous researchers have identified or advocated the importance of aesthetics [24, 54, 91] and self-reflection [106] as reasons why people may choose to use a casual visualization. For example, a visualization system for email archives called TheMail [132] and the 'Imprint' system [104] both encouraged reflection and data exploration where productivity is irrelevant. Informative art also stresses the importance of self-reflection or what Viégas *et al.* [134] term 'data mirrors' [45]. However, in the above examples, aesthetics and self-reflection were examined either in an atypical environment (art galleries and laboratory settings) or by studying people already using a system. We were unsure how aesthetics and self-reflection would be viewed during the daily lives of our participants.

Reduced Learning Costs

Grammel *et al.* [50] examined barriers non-experts encounter when constructing and interpreting simple visualizations. Although their study focused on a work-related context, we believe similar learning costs would be encountered by many casual visualization users. Similar to Heer *et al.* [57] and Kang *et al.* [67], Grammel *et al.* [50] suggested several ways to reduce learning costs. They also found that participants performed tasks based on goal-directed hypothesis testing, which our model also proposes.

Personal Interest

Pousman *et al.* [106] claim that data needs to be personally meaningful for it to be explored in a casual context. Supporting this claim, Danis *et al.* [30] showed that personal interest can direct non-work related system use. Wattenberg's [139] Baby Name Voyager demonstrated the ego-centric nature of casual visualization use as many users initially searched for their own name. Research involving a job trend visualization revealed a similar pattern [44]. However, upon repeat use of the Baby Name Voyager, a subset of users examined names unrelated to themselves, presumably out of curiosity or due to social incentives. Similarly, other researchers claim that citizen scientists and so-called 'data junkies' may be more interested in the analysis process than in the dataset itself [30, 82]. Nonetheless, citizen scientists may constitute a small, atypical group of users. It is currently unclear how important personal interest in the data is for most users; a factor examined during our study.

Communities and Socializing

Wattenberg [139] discovered that shared data, collaborations, and community discussions were critical to system use. People shared data patterns with others and proposed data exploration challenges, suggesting that extrinsic social pressures may have helped drive repeat use. Viégas *et al.* [135] confirmed that group reflection and sociability (along with utility, personal reflection, creativity, and play) were important motivators for visualization creation. For example, interest-based communities (such as Christian groups) used Many Eyes visualizations even when the patterns revealed were not particularly informative (such as Jesus being a central New Testament figure). The Many Eyes result suggests that the topic and community may influence visualization use more than the magnitude of data discovery.

Previous User Goal Categories

Work by Wattenberg [139] and Bartle [9] previously identified four roles (*achievers*, *explorers*, *socializers*, and *killers*) that online community users exhibit which seem highly appropriate to casual InfoVis motivations. Achievers attempt to accomplish a task, whether achieving a high score in a game or finding a name for their expectant baby. Visualization explorers typically find new patterns or unexpected patterns in a data set presumably out of curiosity, while the main concern of socializers was interactions with other users. Finally, killers are people who enjoy imposing themselves on

others and causing distress. For the Baby Name Voyager, killers ridiculed the names of others. Killers and socializers can both be considered motivated by social factors. Achievers appear motivated by perceived usefulness and explorers appear driven by curiosity and learning goals. Bartle's taxonomy of users might roughly describe casual InfoVis users, however, we did not wish to constrain our exploration to these categories. Instead, Bartle's four user categories illustrate some of the ways that user motivations are determined by goals and system affordances.

6.3.2 Research Approach

Our research was guided by a series of questions related to my primary research goals:

1. What factors motivate people to examine data sets in their non-work time?
2. What factors influence the continued use of systems?
3. What goals do users have when examining visualizations and artifacts?
4. Can we model such behavior to enable hypothesis testing and design recommendations?

We conducted our study in three phases: recruitment, artifact collection, and a final interview. We used a cultural probes approach (described below) for artifact collection; participants carried a digital camera (or camera phone) during the artifact collection phase and submitted daily journal entries. Participants submitted photographed artifacts with the journal entries whenever possible. I conducted individual one-hour semi-structured interviews within nine days after each artifact collection period ended. I read electronic journal entries before the interviews to help guide the semi-structured discussions. Interviews were audio and video recorded.

6.3.3 Participants

Twenty-two Victoria, British Columbia residents (fourteen women), between 19 and 30 years old ($Mean = 22.9$, $S.D. = 1.80$), without children participated in the study. Our selection criteria were chosen to promote results convergence and because we thought this demographic would have more free time and a less structured schedule than older people with children. We hoped this would lead to more observed

artifacts and more curiosity-based searches. Participants were recruited using on-line job postings on websites such as Craigslist Victoria¹ and using paper posters distributed throughout the city of Victoria (see Appendix C for recruitment materials). Each person was required to have access to a digital camera for the two-week research period. Sixteen participants were college and university students. Of the remaining people, two were unemployed, one was on medical leave, one worked in construction, one worked in an office, and one was a long-stay tourist. Each participant chose a pseudonym for online journal anonymity. Participant names were **ABC**, **Benji**, **Broccoli**, **Dora**, **Flamingo**, **Galoobglaz**, **Gorffy**, **Grahame**, **Gravity**, **Leon**, **Meerkat**, **NZGinger**, **Poppy**, **Que_sera**, **Rubix**, **S2**, **Toast**, **Toothpaste**, **Willie**, **YN**, **116958**, and **Z10**.

Participants were asked to sign a consent form before they could participate (see Appendix C). The consent form was provided when I initially met them. Ongoing consent was inferred with continued completion of paper journal entries or explicitly indicated (via a check box) for online journal entries. Participants who did not reply to emails and submitted fewer than 4 online journal entries were assumed to have dropped out of the study. Participants who completed the minimum number of journal entries were paid after the interview and debriefed. Only one person (not listed) did not complete the study and he only submitted one journal entry.

6.3.4 Cultural Probes

We believe that usage context is a key component to understanding how and why people use visual representations of data. We also wished to provide a more holistic picture of how people use visual representations of data in their daily lives than, for example, a laboratory study could provide. We considered an interview-only approach, but felt that recall biases would be too pronounced. Cultural probe studies examine a user activity by supplying each participant with a kit (cultural probe) with which to record their activities and opinions [44, 49]. Records can include photos, journal entries, and drawings. Cultural probes, unlike journal studies, are intentionally playful or evocative to encourage use. Traditionally, cultural probes also intentionally bias participant reporting by emphasizing creativity, under the assertion that self-reporting is always biased [44].

Our approach is a variant on this method as cameras were not provided, and ma-

¹<http://victoria.en.craigslist.ca/>

materials were not overly playful or exotic. Participants were required to have access to a digital camera to join the study. Unlike a journal study, photos were highly emphasized while the journal simply recorded a brief summary of initial thoughts. Each participant was given an artifact kit with printed journal entry forms, instructions, a business card sized summary sheet, and a pen (paper probe materials are provided in Appendix C). Journal entries could be completed electronically or on paper, but electronic submissions were encouraged. Reminder emails were sent if the participant missed more than one consecutive electronic journal entry.

Using cultural probes had several advantages: the daily journal would remind users to remain aware, and the photos and short journal entries would be simple enough to not discourage artifact finding but sufficient to get immediate impressions about an artifact before long term recall errors became too pronounced. Interviews could flush out any additional details with the photographs providing a reminder. My cultural probes approach also provides multiple data sources, enabling the triangulation of results.

An open-ended approach was necessary since we expected substantial artifact diversity and we wished to not constrain or direct user actions. Our results about data artifact use cannot be generalized away from the participant group and the time of the study, as is the case with all cultural probe studies. However, trends and patterns observed through qualitative analysis can aid hypothesis and model generation, our primary concern at this research stage.

6.3.5 Artifact Definition and Instructions to Participants

For participants, we defined a *data artifact* as any visual representation of data that encodes more than one *unit of information* (the most simplistic graphical image that conveys meaning [38]) in a non-alphanumeric way. Thus, an unordered table of values only presenting information using alphanumeric symbols and row/column intersection information is invalid. However, an ordered table of values with colour coded cells would be a valid artifact. Artifacts could be physical or electronic. Artifacts could not be used for job or school-work related purposes, but could still be found during work breaks. Each artifact could be reported only once.

I met each participant separately at the start of the study, provided the cultural probe, gave an experimental overview, and answered any questions they had. People were instructed that if they were unsure if something was an artifact, to assume it

was. After two electronic artifacts were submitted, feedback was given to participants, encouraging them and indicating which artifacts were on the right path. We emphasized that participants should go about their daily lives and not actively search for artifacts, but they should still be actively aware of their surroundings. This emphasis was done to collect an accurate representation of each person’s daily life patterns.

6.3.6 Data Analysis

We analyzed interviews, journal entries, and artifact photos for patterns and commonalities using a process inspired by grounded theory [27] and following grounded evaluation guidelines [65]. We triangulated our results using more than one participant and multiple data sources (e.g. journal entries, interviews, and images). Our results subsequently formed the foundation of our proposed **Promoter / Inhibitor Motivation Model (PIMM)**. As an initial approach to the data analysis, we grouped artifacts according to their characteristics, the context of their use, and the explicated user goals associated with that context. Journal entries form the basis of our descriptive statistics.

We first transcribed the interviews and removed non-relevant conversations. We then iteratively tagged the data using the qualitative analysis tool Weft QDA. We used open coding to generate ad hoc tag names [27, 65], and then hierarchically grouped, organized, and merged tags using the process of theoretical coding to help identify commonalities between topics [27]. Our categories eventually formed the basis of our analysis of promoting and inhibiting factors, system feature clusters, user goals, and circumstances of use.

Artifact photos were also analyzed for patterns. We printed each artifact photo submitted. Artifact submissions without photos were roughly sketched by hand or found online, based on participant descriptions (see Appendix C for all artifact images). All journal entries without images described well-known objects. We organized these images by various characteristics, including by format (e.g. geospatial), purpose of use, context of use, number of dimensions, complexity, and use factors. Use factors included the degree of user focus, time used, trust, importance of the data, and user interest. We also noticed that duration and use frequency seemed related to the artifacts used. Journal data led us to order photos by use patterns, providing the basis for our stages of use, described later.

6.3.7 Data Cleaning

As expected, numerous submissions did not fit our definition of an artifact, nor would all valid artifacts be considered conventional visualizations. We chose a group of approximately ten artifacts that matched traditional visualization criteria, and used commonalities between them as the criteria for target artifacts. We defined *target artifacts* as visual abstractions of quantifiable data that demonstrated at least one of the following:

1. Multiple uses/interpretations per data object, while remaining precise and unambiguous. For example, a weather forecast can help someone choose their clothes that day or to explore how weather patterns have changed in the past year.
2. Meta-patterns and trends between data objects. Pattern tasks may include determining the family of plants that grows best in Victoria, based on a planting chart.
3. High information density. High density may include artifacts representing a lot of data objects or each data object having a high dimensionality.

Artifacts with none of these characteristics but which clearly fit the definition given to participants were designated as *borderline* artifacts. Borderline artifacts included cell phone icons, an electric guitar tuner, totem poles, an iPod progress bar, and video game icons. Finally, submissions that did not meet the definition were designated *inappropriate*, and were not analyzed. Borderline artifacts were analyzed for the sake of completeness. However, all results exclusively refer to target artifacts, unless otherwise noted.

6.3.8 Model Generation

Through our coding of interview data, we identified factors that influenced a participant's motivation and enjoyment when using an artifact; these form the basis for our regulating factors (see below). Similarly, we clustered reported user goals, forming the basis for the proposed user goal categories. Finally, we used interviews, journal entries, and artifact images to identify which elements affected goals and regulating factors. Context, user characteristics, visual artifact characteristics, and data represented form the basis of our source factors, which we believe can exhaustively describe

the state of artifact use at any time. Connections between regulating factors, user goals, and source factors are speculative with inspiration taken from *Goal-Setting Theory* [81] and some support derived from interviews.

6.4 Findings

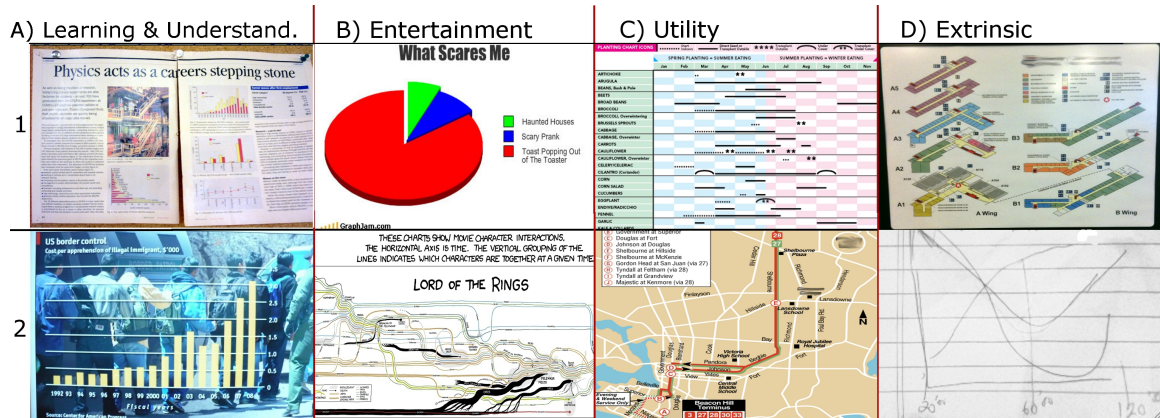


Figure 6.1: An assortment of submitted target artifacts and their primary goals (A to D). Artifacts are: 1A) a newspaper article posted on a wall showing physics degree career possibilities, 2A) a bar graph showing the cost to apprehend and deport individual illegal immigrants, 1B) a joke pie chart found online, 2B) an online flow diagram showing how Lord of the Rings characters interacted over time, 1C) a planting chart found in a book, 2C) a local bus map found online, 1D) a building map found outside a classroom, used to waste time, and 2D) a hand-drawn graph of forgotten content generated by a friend’s mother to explain a topic.

6.4.1 Artifacts Reported

Types of Artifacts

One hundred and seventy-two artifacts were reported by our twenty-two participants (86 target, 49 borderline, and 37 inappropriate). Examples are shown in Figure 6.1. Each participant submitted a mean of 1.7 inappropriate ($SD = 1.9$), 2.2 borderline ($SD = 2.0$) and 3.9 target artifacts ($SD = 3.0$) during the two weeks. Two participants submitted no target artifacts. A total of 230 journal entries (including ‘no artifact found’) were reported.

Forty-two of the valid (target + borderline) artifacts could be considered charts or graphs, the vast majority of which were pie charts and bar charts. Twenty-seven maps and floor plans were reported, both electronic and physical. Although nineteen icon or status monitoring artifacts (such as cell phone screens) were reported, all but one were borderline. Finally, we noted ten game artifacts, of which four qualified as target artifacts.

Context of Artifact Use

Journal entries and interviews revealed that participants generally encountered artifacts during scheduled relaxed leisure time such as reading, surfing the Internet, and watching TV. Journal entries revealed that fifty-four target artifacts were encountered during relaxed leisure time, followed by artifacts found doing chores and necessary tasks (11), taking a break (9), social time (7), and just ‘wasting time’ (avoiding boredom) (5). Most encounters with artifacts were not random. The 86 target artifact journal responses indicated that participants were relaxed 36 times, going at their typical pace 31 times, very relaxed 12 times, and in a rush 8 times.

Few participants encountered artifacts while procrastinating (see definition in Table 6.1); instead most people allocated free time to themselves. As an example, participant 116958 explored the ‘Wolfram Alpha’ computational knowledge engine [142] for over a half hour, based on a friend’s suggestion. She allocated this time to not doing course work and thus did not feel pressure to get back to work. Similarly, YN, Gorffy, Que_Sera16, Grahame, NZGinger, and 116958 all made Internet use an accepted part of their daily routine. YN looked at statistics about her blog readers to know which friends and family members read her posts. Despite considering herself a highly efficient person, she chose to explore this data for however long it required (perhaps due to the strong social motivators). YN stated “I guess, this is quite fast because sometimes you don’t get anything...But if, let’s say I found something interesting...I found ‘Oh yeah, this is not the friend that actually frequently view my blog’ and I would try to view their blog and see why are they interested in my blog.”

At least four participants reported artifacts found using ‘Stumble Upon’²; a program that directs users to random but personally relevant web pages. Even this application, which seems geared toward killing time and avoiding work, was used for planned breaks and as a reward for work completed; although some participants

²www.stumbleupon.com

reported they have used it to procrastinate (but did not find an artifact).

It is possible that participants were too distracted to notice artifacts they used when not relaxed, however, we found no evidence for this. Additionally, reflection biases may have meant that artifact use was retrospectively considered planned when in fact it was just procrastination.

NZGinger identified that for him, procrastination is work avoidance and therefore must be effectively effortless. He stated “...my usual pace is pretty relaxed...Say I was at work procrastinating from work, I don’t know if I would [install a flash update] because at that point I would be like, ‘I should really get back to work’... something came up that sort of made me have to work instead of muck around.” The lack of procrastination artifacts might reflect the learning effort required and people avoiding complexity when procrastinating.

6.4.2 The Promoter / Inhibitor Motivation Model (PIMM)

Interview questions predominately focused on the circumstances surrounding an artifact’s use and why the participant decided to use it. Subsequently, interview tags about goals, factors influencing artifact use and circumstances were the most common. We grouped tags within each of these categories (using affinity diagramming) and then searched for between category correspondences. This approach led to our proposed model. We summarize our findings about casual InfoVis motivation as the Promoter / Inhibitor Motivation Model (PIMM)), a model which describes a theoretical underpinning behind casual InfoVis use. The model consists of two parts: the primary PIMM framework (see Figure 6.2) and the stages of use (see Figure 6.4). We first provide an overview of our cost/benefit model and then describe in detail the goals, source factors, and regulating factors that we identified in our study, with specific examples.

Model Overview

PIMM operates under the theory that a user will continue to use a visualization provided perceived benefits exceed perceived costs. Thus, a subjective weighing of factors promoting use and factors inhibiting use in relation to constructed user goals determines a user’s motivation and subsequent actions. Cost / benefit evaluations likely have a large subconscious component and are not likely directly summative in nature.

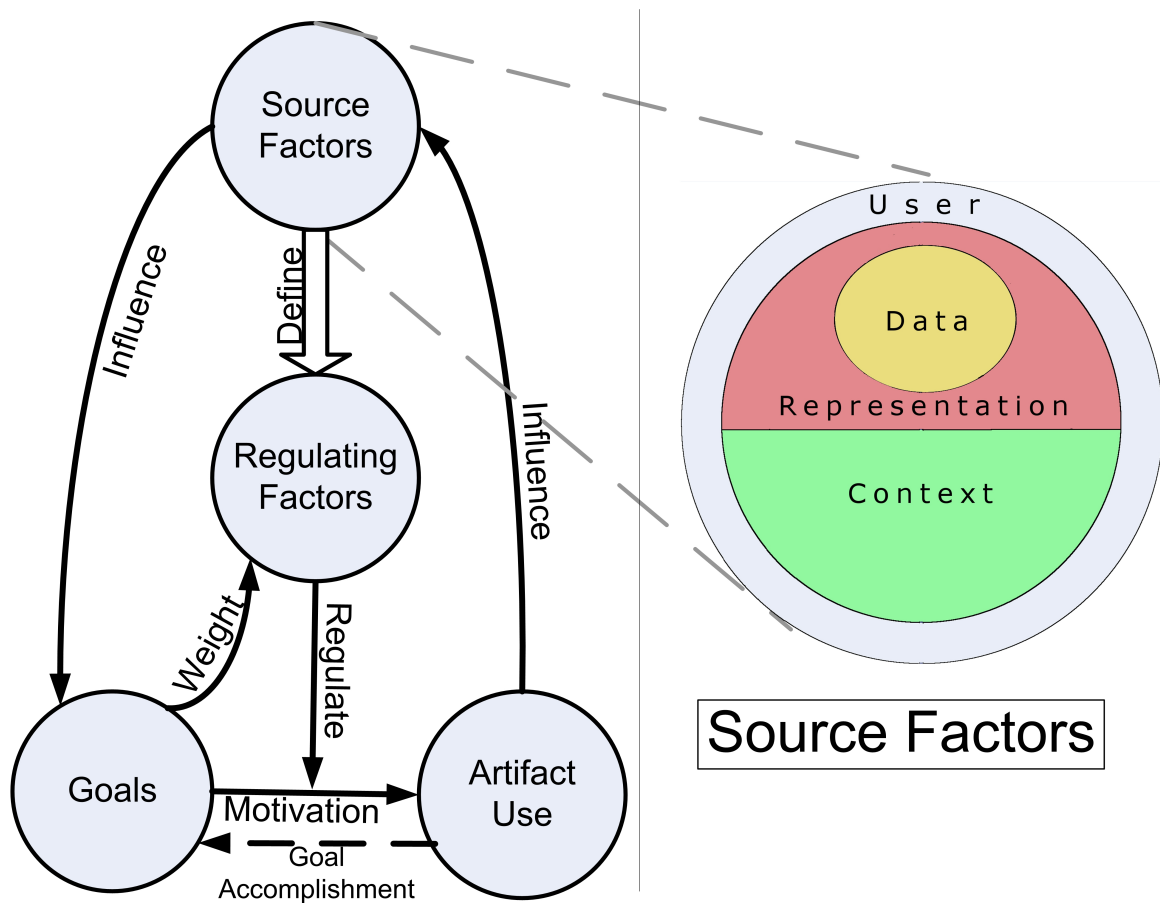


Figure 6.2: The proposed relationship between source factors, regulating factors, goals and artifact use. Regulating factors are determined via source factors interpreted through user perception and cognition. Source factors influence goal production while goals affect the weights of regulating factors. Regulating factors subsequently modify user motivations to accomplish their goals. Artifact use alters source factors and potentially leads to goal accomplishment. The source factor diagram (right side) illustrates that data are influenced by the system representation or visualization and that all source factors are interpreted by the user's thought processes and experience.

We expect construction of the user's goal to be influenced by '*source*' factors of context, system design, data, and user characteristics (see Figure 6.2). '*Regulating factors*', such as learning effort, then promote or inhibit the motivation to accomplish these goals. Users weigh the importance of each regulating factor based on their goals. Finally, visualization use affects the context surrounding subsequent visualization use, providing a feedback mechanism. Promoter and inhibitory factors are also expected to change in magnitude and importance based on the use duration and whether the visualization had been used before. Thus goals promote different use patterns over time (see Figure 6.4).

Gorffy's planting chart (see Image 1C: Figure 6.1) will provide an illustrative example of PIMM. Thought processes in this example are inferred based on the interview. Sections of Figure 6.2 being discussed will be identified in bold font.

Gorffy originally sought out the planting chart based on an acquaintance's recommendation (**Source Factor: Context**) to aid Gorffy's new gardening hobby (**Source Factor: User**). This chart identified vegetables that can be grown locally, the planting schedule, and the required growing conditions (**Source Factor: Data**). The artifact was paper based and static, simplifying repeat referential use (**Source Factor: Representation**).

Gorffy's primary goal was to learn when to plant different seeds. Usefulness and correctness regulating factors weighed heavily based on these goals (**Goal to Regulating Factors arrow**). Gorffy also found the chart aesthetically pleasing (**promoter**). The fact that the artifact was useful and correct (**promoters**) and **inhibitory** 'system learning costs' were modest meant that Gorffy was motivated to use the artifact (**Regulating Factors to Motivation arrow**).

When Gorffy acquired the information he wanted, his primary goal was accomplished (**Artifact Use to Goal arrow**) and motivation surrounding the goal ceased. However, the artifact also piqued his curiosity and he articulated a second goal of 'learning what can be grown locally'. Regulating factors would be weighted differently with this learning goal. Gorffy would continue using the artifact (**motivation was non-zero**), provided inhibitory factors such as 'limited time' did not outweigh the benefits of satisfying his curiosity. When inhibitors exceeded promoters for all goals, artifact use ceased. Artifact use altered source factors and a positive experience increased the perceived usefulness of the artifact (**Artifact Use to Source Factors arrow**).

Source Factors

We found that four main source factors were sufficient to describe any instance of artifact usage: context, user characteristics, the artifact’s representation, and the data. Source factor categories were decided based on the traditional data / representation / user distinctions, and with context subsuming all additional factors. Categories were subsequently verified based on the interview tags. Proposed source factors are hypothesized to influence goal generation. Source factors, when viewed through the lens of user perception and cognition, can serve to regulate motivation (regulating factors; discussed later), and we hypothesize they influence goal generation.

Context represents circumstances outside of the user or artifact. This could include the physical environment (e.g. riding in an elevator or sitting at a desk at home), time of day (e.g. planning an outfit based on the weather forecast), and the user’s activities (e.g. playing a video game, waiting for a lecture to start, or running to class).

Data is the abstract information being displayed. Relevant characteristics of data included accuracy, correctness, precision, the type of data (e.g. geospatial), and the subject matter. We hypothesize that system features lead directly to how data are interpreted, making representation and data highly intertwined.

Representation is the artifact itself and how data was encoded within it. Relevant characteristics included novelty, interactivity, data clarity, data complexity, chart junk, and information density. We were surprised to find only four target (Wolfram Alpha, an unorthodox re-forestation chart, NZGinger’s ‘Maturity Through Goats’ diagram, and Gorffy’s planting chart) and two borderline artifacts (totem poles and a Chakra chart) were considered novel. Only the planting chart required the participant to learn how to interpret the data. Participants did not seek out novelty and novelty was never given as a reason why someone used an artifact. **User characteristics** included hobbies, ability to withstand boredom, need for efficiency, self-confidence, self-efficacy (the belief that one can do a task), and learning style.

User Goals

Based on existing motivational theories [32, 81] and our interviews, we separated goals into intrinsic and extrinsic categories. Although intrinsic and extrinsic categories are usually applied to motivation, we consider all motivations to be related to goal accomplishments.

We defined an intrinsic goal as the plan to use an artifact based on factors directly relating to the artifact's information. Conversely, a plan based on factors not related to the artifact is an extrinsic goal. Extrinsic goals may include alleviating social pressures, avoiding boredom, and acquiring monetary benefits. Intrinsic motivators were far more common and were sub-categorized into utility, learning & understanding, and entertainment. The goal categorization and examples are shown in Figure 6.3, and each category will be discussed below.

We do not provide an exhaustive taxonomy for all possible visualization goals, but our categories do encompass all the reported artifacts' goals. Information visualization research typically focuses on tasks such as gaining insights about the data, communicating information, and externalizing cognition (visually working through a problem). We did not categorize goals in this way because these categories did not exhaustively describe our results. Instead we followed a grounded process, allowing categories to emerge from the data.

Each artifact may have more than one goal. For example, a participant initially used the Lord of the Rings flow diagram (see image 2B, Figure 6.1) based on peer recommendations but then also used it out of curiosity and for entertainment. NZGinger's 'Maturity Through Goats' artifact (see image E, Figure 6.7) was found using the application 'stumble upon', initially examined for entertainment, and then more thoroughly examined to see where he fit into the maturity hierarchy (learning & understanding). Furthermore, goals can often overlap: z10's Kabbalah Tree of Life artifact (see image E, Figure 6.9) enabled him to contemplate the religion (learning and understanding), which he found enjoyable (entertainment) and was a visual aid for text-based discussions he read (Utility).

We designated a primary goal for each target and borderline artifact and found that utility was the largest cluster with fifty-five artifacts. We were able to classify 113 of the artifacts (see Figure 6.3); the remaining twenty-three were ignored or no clear goal was indicated.

Intrinsic Goals:

Learning & understanding goals involved a desire to use an artifact to acquire new information (learning) but not to complete another task. This goal set includes alleviating curiosity, acquiring information, and reflection (learning more about one's self or contemplating a complex and subjective topic such as religion).

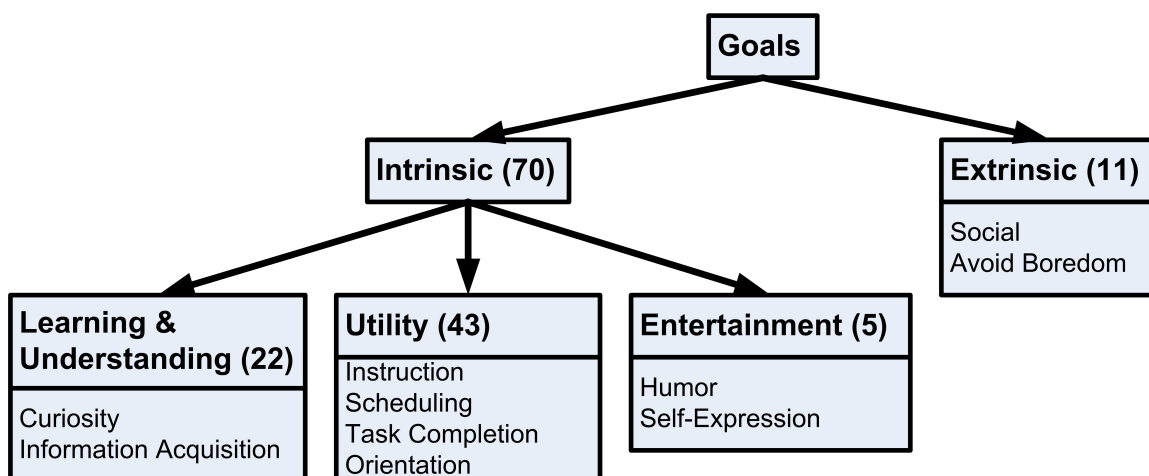


Figure 6.3: Hierarchy of participant goals. Example goals, such as humor, are provided to illustrate the four major categories. Numbers indicate the number of artifacts whose primary goal belonged to each category

Curiosity and information acquisition were by far the most commonly articulated learning goals. Occasionally, participants encountered artifacts at random, but most curiosity based data examinations were done during scheduled relaxed leisure time (twenty-seven of thirty-eight learning artifacts). Curiosity based goals were often regulated by higher priority tasks, such as getting to a class, and the visual appeal of the representation.

Utility goals involved well-defined, close-ended activities and were often associated with routine or well practiced tasks. Thus the steps required to reach a clearly articulated utility goal were well understood. Artifacts used for utility purposes could be tools, instructions, scheduling aids, maps, status monitors, and means to organize and filter data. Four of the five participants asked to choose their favorite artifact chose their favorite based on its perceived utility. For example, Flamingo said “probably Google maps. It’s just so damn useful you know?”

When an artifact was used primarily for utility purposes, features that reduced productivity were strongly disliked. This effect appeared less pronounced for other goals like entertainment. Aesthetics at the expense of utility was almost universally disliked. Toast summed up this opinion succinctly, stating “The point of the object is that it works, not that it’s nice [looking].” Similarly, NZGinger said comments like “Information is a tool” and “I like my information quick.”

Entertainment goals encompassed the user’s desire to be amused. Artifacts used for entertainment included comedic artifacts such as those at Graph Jam [63], playful artifacts, board games, graphs within video games, unintentionally humorous artifacts, and artifacts used for personal expression / identity construction (such as a hand made chore wheel with each house-mate’s picture pointing to that week’s tasks: not shown).

Artifacts designed for entertainment are unconventional from a traditional visualization perspective; however, entertainment seemed to be a strong motivator for use. Most comedic artifacts were quite simplistic and used once, however, dynamic in-game artifacts, such as Gorffy’s Civilization IV line graph (see image D, Figure 6.9), were often used the most frequently and for the longest periods of time out of all the reported artifacts. NZGinger examined his humorous yet complex ‘Maturity Through Goats’ diagram longer than any of his other artifacts (see Image E, Figure 6.7).

Extrinsic Goals: Extrinsic goals, or goals independent of the data or system, predominately involved social pressures and avoiding boredom. As an example of social pressures, 116958 claimed she felt the need to pay attention when a friend’s mother created a series of hand-drawn charts to explain a complex topic (see image 2D, Figure 6.1). The charts did not make much sense to her and she stated, “It kind of made me pay attention more but it didn’t really explain her point.” Other extrinsic pressures included procrastination and wasting time / avoiding boredom. Wasting time was common and explained why some participants examined artifacts they did not have a strong personal interest in; these were often objects in the nearby environment. Gravity reported a library floor map, explaining, “I don’t normally look at these things for fun. It’s definitely for information or sheer boredom in the elevator.” Despite this, some level of personal interest was still present most times. For example, Benjii studied a map of the local area’s tectonic stability to see if her home was in a safe area, while waiting for a lecture to begin (see Image A in Figure 6.7).

6.4.3 Regulating Factors

Goal setting theory suggests that the ability to achieve one’s goals is an important driving force behind any activity [32]. However, participants reported numerous other factors not exclusively related to their goals that they claimed influenced their mo-

tivation to use an artifact. We characterize regulatory factors as *promoters* (which encourage artifact use) and *inhibitors* (which discourage use). A subset of the most prominent or interesting regulating factors will be discussed below. A complete list is available in Table 6.2.

Promoters

Personal Interest was one of the most evident promoters, noted for almost all reported artifacts. For example, Gorffy’s planting chart directly aided his gardening hobby and Gravity’s chart on the costs to handle illegal immigration was meaningful because he had lived in Mexico (see image 2A, Figure 6.1). YN summarized this idea, saying “If it’s complicated and meaningful to me then I may take the time to learn about the data.” Personal interest appeared to be associated with all four types of goals. It was a key factor in artifact use for 95 target and borderline artifacts.

Usefulness was important for utility goals (ex. email filtering by colour to improve email monitoring) and learning & understanding goals (ex. pie charts in Poppy’s bank statements and information about who read YN’s blog). Useful visualizations did not necessarily improve perceived task efficiency. For example, a user with a learning goal would find a visualization that improves understanding useful, even if the system took longer to use. Participants primarily judged usefulness based on what the data could be used for, but context, personal priorities, and representation also influenced this evaluation. For example, a campus map with a strong emphasis on parking areas still helped Gorffy find a building and was thus useful.

Usefulness also depended on whether the level of detail was appropriate. Numerous participants reported that approximations were ‘good enough.’ For instance, when discussing an interactive chart in an airport, YN said, “Yeah, of course your data can give me a lot of information, but as a tourist actually I don’t want to know the exact statistics...an overview...is more than enough.” Conversely, however, Flamingo took extra time and effort to check online weather information rather than use a desktop widget, stating, “I would usually just go to the website for the weather too because [the widget] doesn’t give you very much information.”

Curiosity was a major promoter for initial use, but was highly dependent on personal meaning of the data and individual differences. For example, Gorffy claimed to actively seek out new data and responded to surveys out of curiosity, while YN only explored non-utilitarian artifacts when she was wasting time. Curiosity appeared

to be a relevant regulating factor for all four goal categories but it predominately appeared to affect first time use.

Data Correctness/Trust was a more complex promoter than we expected. Participants articulated that trust in the data's accuracy only mattered when the data was personally meaningful or there was a high cost to receiving incorrect information. For example, Gorffy wanted to know the data source of a global mortality rate visualization because he believed the data was too important not to note this. Trust affected all goal categories.

Cost of a Misinterpretation: Usage time increased when the cost of misinterpretation increased. Gravity, S2, Leon, NZGinger, Broccoli, Poppy, and Toothpaste all stated that they did not care about interpretation or data errors as much when the data was less important. For example, Broccoli was not concerned reading the various gages when playing a golf video game because she was not very engaged in the game and errors were not a major concern. It is conceivable that some artifacts may be avoided altogether if the perceived error costs are too high. This factor was not considered a pure promoter for this reason (see Table 6.2).

Aesthetics was a relatively minor promoter and no participants claimed it was a primary factor. Furthermore, aesthetics at the expense of utility was often highly disliked. Flamingo, NZGinger, Gorffy, and YN were the only participants who stated that visual beauty was a reason they enjoyed an artifact and in all cases, they preferred 'clean' and simplistic visualizations. NZGinger did state "Eye appeal is always good" but later noted that it was secondary to efficiency. Gorffy stated, "I liked it because it's pretty simple. When I'm just like relaxing, browsing through news articles I don't want to see super complex charts; that kind of bling." Context of use and data may play a key role in how aesthetics and complexity are perceived.

Inhibitors

Time Constraints & Higher Priority Tasks were perhaps the strongest and most common inhibitors. Frequently, participants stopped using an artifact because they had another task to perform (ex. Leon examined a distribution of international students while searching for information for an essay) or when time was limited (ex. a lecture was about to begin when Benjii examined the earthquake risk map: image A, Figure 6.7). Time constraints affected all goal categories. Broccoli articulated her enjoyment for seeking out new random information but said, "If I had time, I would

do [random data searching] more.” Higher priority tasks were not noted for utility goals; we suspect that these tasks had a high priority already.

Artifact Learning Effort tended to inhibit artifact use, particularly for lower priority artifacts used only once, such as most of the reported charts. Broccoli explicitly stated that she avoided “hard-core sciences” artifacts, stating, “I guess I’m intimidated by things that I don’t understand.....” Learning effort seems inversely related to novelty; familiar systems lack novelty, but have reduced barriers to use and are therefore more likely to be examined, as also reported by Grammel *et al.*[50]. Interestingly, format familiarity or learning effort regulated motivation for all goal categories, suggesting that this regulating factor is pervasive.

Insufficient Data Context: We were somewhat surprised that participants never sought out data; instead they sought a topic of interest and used artifacts only when encountered in their search. Participants were therefore interested in commentaries and summaries. Rubix said, “Usually with the text you have a full explanation, so then you can make your own decision if you believe it or not [...] if they don’t give a very good explanation you usually don’t trust it.” Similarly, Toothpaste viewed charts as supplementary to a persuasive argument, explaining, “No I think I would rather have kind of a holistic [approach]. You know, here’s some data to back up this article.” Thus, visualizations seem to be better served via linkages to text information about a topic rather than a repository of visualizations being grouped together. Earlier suggestions of a ‘martini glass’ design approach also seem justified given these results [113]. Martini glass approaches initially guide a user’s data exploration but progressively open up to more unstructured or open-ended data explorations. Galoobglaz took this approach when reading about whale migration patterns. She initially examined a whale range map to see if Bryde’s whales were nearby and then examined all their ranges (also illustrating personal interest promoters).

6.4.4 Stages of Artifact Use

If artifact use is simply a matter of cost / benefit trade-offs, what causes people to keep using an artifact or to suddenly stop? Our results indicated that regulating factors and goals change over time as experience with the artifact changes. Artifacts grouped by frequency and duration of use had substantially different promoters, inhibitors, and use goals. Perhaps an artifact’s characteristics afford a certain type of use, and / or

because people actively choose artifacts that support an intended use. In either case, by identifying factors associated with different stages of use, we hope to better predict how systems can be designed to support the expected time frame of use.

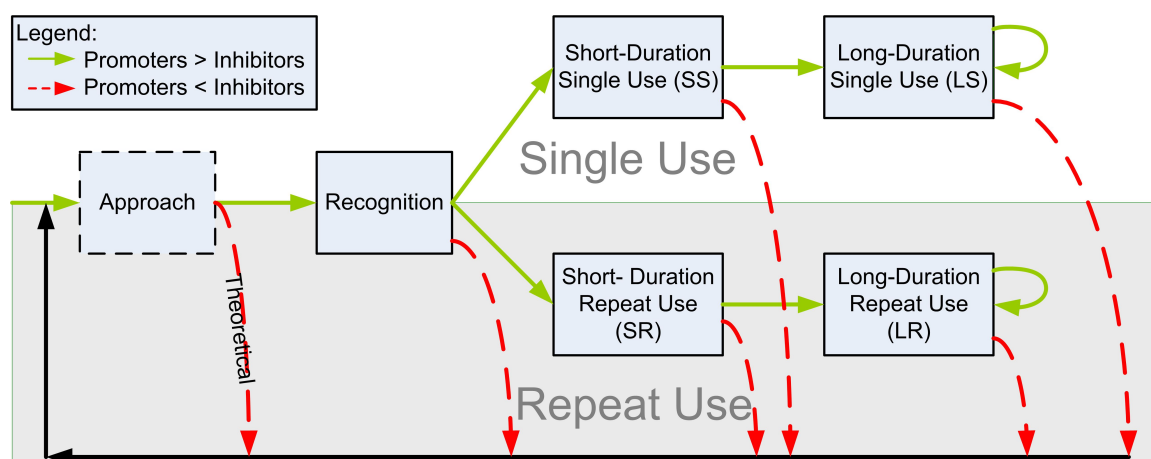


Figure 6.4: Stages of artifact use. Artifacts continue to be used until inhibitors exceed promoters. With each use stage, promoters and inhibitors can change. Repeat looping occurs when long-duration stages are reached. Experience with the artifact subsequently influences whether it will be used again (black arrow).

We grouped artifacts into five clusters based on reported use patterns: **recognition**, **short-duration single use (SS)**, **long-duration single use (LS)**, **short-duration repeat use (SR)**, and **long-duration repeat use (LR)** (see Figure 6.4). Although a different grouping could have been used, we felt the categories chosen were easy to identify, artifacts were well distributed, and the granularity was sufficient to find commonalities between co-clustered artifacts. Participants were often unsure about the duration of longer term artifact use so *long-duration* was defined as artifact use at least once for longer than one minute. *Repeat use* was an artifact used more than once by the participant. We also propose an ‘approach’ group, but since ‘seen but not recognized’ artifacts were not possible to report, this category is theoretical.

Long duration use is hypothesized to require promoters consistently exceeding inhibitors throughout artifact use. For repeat use, novelty, learning costs, and curiosity regulating factors are minimal and previous experience becomes important. The eighty-eight target artifacts consisted of four recognition, twenty-eight SS, eighteen LS, twenty-one SR, and seventeen LR stage artifacts. Examples of artifacts from each stage will be provided to help illustrate our findings. See Figures 6.5 to 6.9 and Table

6.3 for examples and Table 6.2 for an overview of the stages of use in which each regulating factor was observed.

Approach

Approach is a hypothesized use stage that involves observing that the artifact exists, without necessarily identifying what data is being represented. Holmquist *et al.* [60] propose a similar ‘awareness’ stage for ambient InfoVis comprehension. Although we do not have direct examples of this stage, YN’s reported ‘periodic table of mixology’ artifact provides some support (image C, Figure 6.5). This poster was incorrectly read by her to be the periodic table of elements and subsequently ignored, suggesting that it is possible to see an artifact and not exert the effort to identify the data being shown.




Figure 6.5: Examples of recognition stage artifacts ($N = 4$). Identified but unused recognition stage artifacts were a globe (A), a building map emergency evacuation plan (B), a ‘Periodic Table of Mixology’ comedic poster (C), and another emergency evacuation plan (D).

Recognition: Examples of this stage were relatively rare (four target artifacts), each associated with a distinct lack of user interest about the data. Participants had no explicit goals and frequently had higher priority tasks or time constraints. This lack of interest suggests that when an artifact is recognized, a goal for its use is identified or the artifact is ignored.

Short-duration Single (SS) Use: SS artifacts were typically simplistic, personally meaningful, cognitively interesting, and often interspersed with a task currently being performed (such as Internet surfing). Most SS artifacts used a familiar format, although this may be a choice of the artifact designers. The consequence for misinterpreting the data was particularly low for this group. We also hypothesize, based on




Figure 6.6: Examples of Short-duration Single Use Stage artifacts ($N = 28$). SS stage artifacts were used once and for less than a minute. Artifacts include the discussed immigration chart (A), a comedic pie chart (B), a chart showing peak oil production in the U.S.S.R./Russia (C), and a pie chart used during a Japanese television show (D).

participant reports and previous research [139], that users tend to look at personally meaningful data first. To progress to longer-term analysis, users would subsequently need to find other data of interest.

A subset of SS artifacts were not personally meaningful and thus examined only for a rudimentary understanding. For example, Grahame examined two stacked bar charts during a recorded online presentation, one that was relevant to her hobbies and studied extensively (LS) and another that was only examined until the slide automatically progressed (SS).

Some artifacts, such as Benjii's physics charts, were partially examined but higher priority tasks cut artifact use short (image 1A, Figure 6.1). However, we note that the simplicity of most SS artifacts (such as pie charts) may have enabled them to be fully explored quickly, removing curiosity, entertainment, and utility promoters.




Figure 6.7: Examples of Long-duration Single Use Stage artifacts ($N = 18$). LS artifacts were used once for more than a minute and include a map showing earthquake risk in Victoria, BC (A), graphs displayed when learning to use Wolfram Alpha (B), a 'stumble upon' infographic describing which countries are the primary producers for different resources (C), drywall mounting screw instructions (D), and a humorous diagram describing characteristics of maturity using goat illustrations (E).

Long-duration Single (LS) Use: Eighteen LS artifacts were reported and these only appear to be used when daily productivity pressures are extremely decreased. LS artifacts were personally meaningful (such as ABC’s map of where friends traveled and z10’s world resources chart), they were used for a long-duration task (116958’s instructions on how to put up drywall) and/or provided an outlet to avoid boredom (Benji’s earthquake risk map, image A, Figure 6.7). Of all reported artifacts, LS artifacts required the greatest (amortized) learning time, had high complexity, and participants gave total focus to using them. Time constraints and simultaneous higher priority tasks were extremely rare for LS artifacts. The longer duration suggests that user goals may have required more time to satisfy or that the user defined subsequent goals (such as NZGinger’s ‘Maturity Through Goats’ artifact discussed previously. See image E, Figure 6.7). Interestingly, goals or regulators often had an extrinsic component, which may also encourage longer-term use.




Figure 6.8: Examples of Short-duration Repeat Use Stage artifacts ($N = 21$). SR stage artifacts were used repeatedly but for less than a minute each time. Artifacts include a weather forecast widget (A), an online two-week weather forecast (B), a local bus map (C), a pie chart showing the country of origin for readers of YN’s blog (D), and customized colour-coded emails and symbols in Gmail (E).

Short-duration Repeat (SR) Use: SR artifacts were predominately simple tools, orientation aids, and state monitoring visualizations such as cell phone battery icons. In fact, many of the SR artifacts were borderline artifacts because they were very simplistic. Still, twenty-one target artifacts were SR. Task times were well known and the artifacts were not novel. Curiosity played no role and utility seemed paramount. SR artifacts were often dynamic, thus warranting return use. Efficiency was a common promoter, as was usefulness and personal interest. For example, 116958’s Gmail (see image E Figure 6.8) was set up for optimal efficiency and was monitored throughout the day.



Figure 6.9: Examples of Long-duration Repeat Use Stage artifacts ($N = 17$). LR artifacts were used repeatedly and used for more than a minute at least once. Artifacts included a planting chart (A), a map of local bike trails including trail difficulty and elevation information (B), a building map for a regularly used but confusing building (C), a culture progression line graph in the video game Civilization IV (D), and a Kabbalah Tree of Life diagram used to reflect on the philosophy surrounding the faith (E).

Long-duration Repeat (LR) Use: LR artifacts are perhaps the best examples of traditional visualizations for casual users. Sixteen of the seventeen target artifacts were sufficiently complex for long term use, were predominately used for planning or strategy, and provided a clear utility benefit. The sole exception was the visually simplistic Kabbalah Tree of Life diagram (image E, Figure 6.9) which z10 used as a framework to reflect on the religious philosophy. LR artifacts were frequently related to strong personal meaning such as a hobby, and were rarely examined under time constraints or when a higher priority task was to be performed. Interestingly, most LR artifacts had the highest consequences for misinterpretation, which may suggest the longer use duration was necessary to verify the data. SR artifacts had the next highest cost for misinterpretation, followed by LS. LR artifacts continued to be useful, often because the data was dynamic, participants needed to repeatedly reference the same information, or the artifact served multiple purposes.

6.5 Discussion

The Promoter/Inhibitor Motivation Model is a simplistic but effective tool for predictive empirical testing and provides a lens through which to view visualization use, potentially guiding future design. Here we re-examine the various aspects of the model, extracting tentative design and evaluation considerations. Our suggestions need subsequent controlled verification, which is outside the scope of this exploratory work.

6.5.1 Design Considerations

Guiding Use Stages

Casual visualizations may not always be designed for long-duration repeat use. In many cases, such as ambient visualizations, productivity dashboards, and dynamic real-time charts, short system monitoring events may be desired. For informative art, a single session long-duration use seems ideal.

Recognition: Factors such as curiosity, aesthetics, social pressures, and familiarity can all theoretically improve the rate of entering this stage. Familiar visualizations should presumably require less effort to recognize.

Designing for Long-Duration vs. Short-Duration Use: Learning challenges are unique for first time use and thus one-use visualizations need to be easy to learn or familiar, especially when short-duration use is expected. In addition, interaction efficiency seems to be important for short-duration use artifacts, particularly those used repeatedly.

Complexity appears to be a key difference between short-duration and long-duration use, probably because learning and curiosity promoters would quickly disappear with a simple artifact. A complex artifact may initially discourage use but the extra information may provide more learning and utility opportunities if the participant is willing (see Table 6.2). Our results also indicate that artifacts providing personally relevant information may offset system learning and time costs, enabling long-duration use.

Our results suggest that participant willingness was increased by extrinsic goals and pressures; another commonality between long duration (LS and LR) stages. Furthermore, ‘time constraints’ and ‘higher priority task’ inhibitors were less common for LS and LR stages.

Promoting Repeat Use: Single use artifacts reported here are not precluded from future repeat use, however, participants frequently stated they had no desire to examine them again, and we noticed clear differences from repeat use artifacts. Factors such as boredom avoidance, random encounters, curiosity, humor, aesthetics, novelty, social pressures, and learning effort all influenced initial artifact use but these factors became less relevant or disappeared when repeat use occurred. Generally, repeat use

relied on positive previous experiences, personal relevance, ongoing utility benefits, and changing data. For short-term repeat use, the ability to anticipate use duration may also promote repeat use, enabling users to fit the activity into their regular routines. System learning was a less pronounced inhibitor with repeat use; however, bad system design or bad experiences seemed to frustrate people more upon repeat system use.

Other Design Implications

Further design implications can be suggested based on our study findings and implications from the proposed model.

1. **Personal Relevance is Critical:** Personal relevance or interest played a major role for almost all artifacts. Even when the artifact topic was not highly interesting, participants sought out how that data related to them. Providing an initial user-centric visualization view (when possible) should encourage system use.
2. **Design for Productive Relaxation:** We were surprised by two common occurrences: participants scheduled relaxed leisure time despite their generally flexible schedules, and even when relaxing they wanted to be productive. Hence, artifacts that served no purpose, were difficult to use, or were not useful were ignored or barely used. These findings are reflected in PIMM: relaxed leisure time minimizes inhibitory factors but promoters are still required. Relaxed leisure time may be the ideal context in which to use a casual InfoVis system as these times are associated with a non-stressed attitude and openness to curiosity based data analysis, and participants tend to use an artifact for an extended time period.
3. **Reduce Costs to Initial Use:** Using a familiar display format and generally reducing the learning effort should increase the rate of initial visualization use, at least to the SS stage. Novel formats need to be sufficiently clear to minimize learning costs and intimidation. Minimizing costs seems especially useful for chance visualization encounters where users have not determined system benefits.
4. **Data Context is Key:** Participants wanted data sources available, many wanted textual discussions, and people generally enjoyed when key patterns

were identified for them. My results contradict Danziger’s [31] notion that people may prefer open-ended explorations using a visualization. Identifying key patterns may reduce the cost to initial artifact use. Participants also preferred information to be organized by topic (with visualizations and text together). This result corresponds with Danis *et al.*’s [30] findings that Many Eyes visualizations were often discussed on other topical sites.

5. **Appeal Changes Over Time:** Curiosity, entertainment, novelty, and aesthetics may all increase short-duration initial use but our results indicate that a participant’s personal interests and expectations of utility are the best ways to encourage repeat use. Providing extra details that do not distract from initial understanding and contextual information may encourage longer-duration use by piquing curiosity after the primary motivation has ceased. Systems should be simple, predictable, and efficient to read for quick look-up tasks (SR artifacts). This suggestion is consistent with Hinrichs *et al.*’s [59] museum visualization results, where curiosity and other users promoted initial system use and personal interest guided longer term interactions.
6. **Design for Goal Type:** For utility goals, usefulness should be clearly identifiable and should never be sacrificed for aesthetics. For learning goals, context and related information (e.g. references) should be readily available.
7. **Constrain User Goals:** Systems used for a wide variety of tasks may insufficiently satisfy users for any given goal. Focusing the system design on a few tasks may therefore reduce user frustration. Developers may wish to support a wide variety of goals, but PIMM indicates this should be done with caution.
8. **Close-Ended Analysis:** Perhaps the most glaring characteristic of our findings was that none of the submitted artifacts were used for the open-ended and unstructured data analysis and exploration tasks proposed by many casual InfoVis promoters. Only the ‘Maturity Through Goats’ artifact (image E, Figure 6.7) was used for self-reflection and only z10 reflected on an artifact (image E, Figure 6.9). We found no examples of artifact use for aesthetic factors, or ‘data junkies’ exploring random data such as reported by Wattenberg [139] or proposed by Pousman *et al.* [106]. We always observed a clear purpose for artifact use, even if it was to avoid boredom. Some participants did explore artifacts out of curiosity, but data was always personally meaningful in some way. Purpose

driven artifact use seems to contradict earlier claims that casual visualizations are used for aesthetic appeal and self-reflection. We do not doubt that some people may use visualizations for these purposes, but these people may be rare or may have training in data analysis (hence all data is personally meaningful). Wattenberg similarly reported that most people explored personally meaningful naming trends. Perhaps open-ended data analysis and reflection only occur in rare circumstances not detected within a two-week interval, or by a different population than the young adults we studied.

6.5.2 Evaluation Implications

PIMM suggests that the expected or desired use patterns (duration and frequency) need to be clearly identified for the appropriate evaluation of a casual InfoVis system. For example, if a system is designed for repeat long-duration use, immediate usability tests are less critical than observing long-term use. In contrast, SR visualizations require immediate usability and are more appropriate for experimentally controlled testing with no training. Furthermore, system adoption depends on user goals and regulating factors. The effects of different regulating factors can be methodically tested, provided user goals are constrained, a key consideration for controlled studies.

6.5.3 Caveats

We note a few caveats that should be considered when interpreting our results. A direct causal relationship between source factors, goals, and regulating factors is not possible using participant recall, and self-reflection biases may distort results. Recall issues were minimized by having participants note artifacts immediately and write daily journal entries. Our results provide foundational evidence for our proposed model of casual visualization use but cannot generalize beyond the population studied. For example, one might reasonably expect motivational factors to differ for older adults or children. Our twenty-two participants constitute a relatively small, and somewhat homogeneous subset of the population. Few of our participants had full time jobs, meaning their behaviors and time constraints may be different from a wider user population. Furthermore, there are some clear differences between our artifacts and what would typically be called casual information visualizations. For example, the non-interactive nature of many artifacts means data had to be simplistic to avoid clutter, which may limit open data exploration possibilities. Nonetheless, we believe

the current research provides useful predictions about casual InfoVis motivations and is able to avoid the serious problems associated with focusing on early adopters.

6.5.4 Study Limitations

Difficulties were encountered during the current study, and like all research, there was retrospectively some room for improvements. I was primarily disappointed with the lack of participant diversity during the initial sixteen participant phase of the study. We did not constrain our participants to university students, but few non-students in the same age range volunteered to participate. Six additional non-student participants were recruited to help diversify our results, but we still only had one participant with a typical work-day job. Time limitations forced us to accept this issue, but a wider user population would have greatly improved how we interpreted the results.

The artifact study results may have also been improved with an improved artifact definition. Several participants had difficulties understanding what we were looking for, even despite the feedback. Although I did not want to constrain the search space too severely, I believe that more ideal artifacts would have been reported if participants did not spend their time reporting as many borderline and inappropriate artifacts.

Finally, in terms of researcher errors, the questions posed to participants during the initial set of interviews retrospectively appear too open and unfocused. Furthermore, in an attempt to improve participant comfort and dialog flow during the interviews, I felt I talked too much. My excess talking meant participants occasionally agreed with statements rather than generating them themselves; thus fewer than expected succinct interview quotes were available. Balancing interviewee comfort and interview direction is a common difficulty for unstructured interviews [27]. All interviews also pose a potential danger that the interviewer will lead an interviewees thoughts and responses. This danger was also possible for the artifact study, but there were no obvious examples of this during the study.

6.6 Related Computer Use and Motivation Theories

Here we will discuss pertinent technology use and motivation theories and how each relates to our proposed model. However, we note that these theories either provide

insufficient predictive abilities, a lack of theoretical underpinnings, or a disregard of how factors change according to context, task, and user traits. Most motivation and system adoption theories also do not address how use patterns change over time. Hence the reason we propose PIMM. PIMM is also grounded in the empirical data itself and explains our observations better than previous models. The grounded theory approach also requires generated models and theories to derive from the data rather than extending previous work. Hence we did not use other theories to explain our results.

6.6.1 Motivation Theories

Herzberg's motivation theory [43] suggests two independent factors for driving actions: motivating factors and hygiene factors. Motivating factors increase satisfaction, thereby promoting an activity. Hygiene factors influence 'dissatisfaction', inhibiting an activity. Our inhibitors can be considered hygiene factors while promoters are analogous to motivation factors. Motivation is also typically categorized as intrinsic or extrinsic [110], and our model continues this tradition. We suspect that intrinsic motivation may have greater relevance for casual visualizations compared to traditional visualizations, because fewer extrinsic motivations such as job performance are present.

Self-Determination Theory (SDT) dictates that humans are more motivated to perform an action if there is a perception of autonomy, competence, and relatedness associated with the task [110]. Autonomy refers to the urge for people to be causal agents in their own life, competence involves the desire to effectively deal with one's environment, and relatedness refers to the need for interactions and connections with others. Although SDT is well researched, this macro theory of human motivation did not sufficiently reflect the artifact study findings. Competence and autonomy could help describe curiosity, utility, and personal interest promoter, and relatedness explains social promoters, but SDT seemed too theoretical for the purpose of design recommendations. Hence SDT was not a major influence on PIMM.

A recent informal usability model proposed by Anderson [4] is similar to PIMM in that context and the users's activities / goals are considered paramount to an effective user experience. However, this model appears relatively vague, provides minimal mechanisms for empirical testing, and is not directly supported by research.

6.6.2 Goal Setting Theory

Perhaps the most relevant theory to PIMM is Locke and Latham's [81] Goal Setting Theory. Goal Setting Theory proposes that the task difficulty, clarity, goal commitment, probability of success, and specificity of a user's goals directly influence that person's motivation to perform a task. Thus, contextual, user, system and data factors can indirectly affect motivation like our model proposes. Goal commitment in particular, is directly related to the importance of the task and self-efficacy, making personal interest an important factor in goal pursuit. Goal setting theory also suggests that more complex but understandable artifacts that challenge a user will be used more frequently or for longer, which our stages of use support. Furthermore, Goal Setting Theory proposes that personal goal accomplishment leads to user satisfaction and is a primary motivating factor for artifact reuse, a central tenet of our model.

Goal Setting Theory, however, does not explain our results as well as PIMM. Goal Setting Theory is more abstract and theoretical, it does not provide specific regulating factor details (termed moderators and mechanisms), the model does not suggest that goals influence regulating factors, and temporal use patterns are generally ignored. Goal Setting Theory's primary similarity to PIMM is that goals are influenced by system factors, goal accomplishment is the primary motivation for system use, and goal accomplishment affects regulating factors using a feedback cycle.

Technology Acceptance Model (TAM)

The highly cited Technology Acceptance Model (TAM) [32, 36] proposes that perceived usefulness (PU) and perceived ease of use (PEOU) are central factors influencing whether an information system will be adopted. In relation to PIMM, we believe that PU can be considered a set of promoters such as usefulness, potential usefulness, and the ability to accomplish user goals, while PEOU is essentially represented by inhibiting factors such as system learning effort. Despite commonalities between TAM and our model, however, the simplicity and the generalized nature of TAM means that it provides little predictive power for casual InfoVis use. Our proposed model also offers numerous additional regulating factors unrelated to PU or PEOU. Expansions to TAM typically just add additional factors; these factors predict initial system use and eventual regular adoption rather than exploring how technology is used over time [16].

Task Technology Fit (TTF)

Other technology use theories are the Task Technology Fit (TTF) model and an extension, the Technology-to-Performance Chain (TPC) [48]. TTF and TPC contend that a technology is more likely to be perceived as having a positive effect on performance if the system's characteristics match the user's task. A positive fit subsequently leads to improved performance and potential system use, as user frustrations appeared most pronounced when participants used an artifact for utility purposes but system efficiency was low. However, these models seem insufficiently precise for our purposes and predict technology use rather than exploring how it is used over time. We feel PIMM is more nuanced than TPC and TTF, as context and data source factors are considered. TPC and TTF also stress task performance improvements, which are typically not a casual activity focus.

6.7 Future Work

PIMM provides a theory to describe user motivations surrounding casual visualization use. Future experiments should be conducted to more precisely determine the interrelation of model factors, such as how motivation changes when regulating factors change. A survey of existing casual InfoVis systems and how they map to PIMM is a simple yet interesting follow up study. We are also interested in better understanding how source factors influence goal formulation. We would like to empirically test our design considerations to determine if they can be used to predict length and duration of use. Although PIMM was limited to visualizations and visual representations of data, we suspect that motivation to use other systems may be similar in many ways. Extending the model to explain more general computer system use may therefore be valuable.

6.7.1 PIMM in Other Contexts

PIMM is potentially applicable to other computer interactions and even different research disciplines. For example, it potentially can explain why an operating system or web page was not adopted by users. Numerous users have contested *FaceBook's* privacy concerns and the less than intuitive interface, yet it continues to be used rather than other social media networks. The PIMM model predicts that FaceBook's use is dependent on what people wish to use it for (games, keeping in contact, or sharing

information with friends or finding high school classmates) and whether the perceived inhibiting factors (such as privacy concerns) overshadow the perceived benefits (such as the increased probability that former classmates will also be users). Provided competing sites do not have the same number of users and thus the same promoters, Facebook's continued use makes sense.

The PIMM model is also not strictly limited to software applications. Technically PIMM could help explain day to day choices that people make such as whether to wash the dishes or have a coffee with friends instead. However, the PIMM model's primary purpose remains to explain why people look at data in their free time. I emphasize this nuance because user goals can differ substantially with the same system. For example, the Baby Name Voyager [139] can be explored out of curiosity or to help determine an upcoming baby's name. Provided different users consistently have the same primary goals (e.g. playing video games for entertainment), the PIMM model reduces down to a simple cost / benefit analysis (see Figure 6.10). Potential model generalizability can be considered a strength of PIMM, however, I also note that the empirical evidence all relates to visual data use rather than other domains. Hence, factors specific to artifacts and visualizations could make the results unique to the domain. Most of the empirical data is also exclusively relevant to the studied domain; only the general PIMM framework has other potential applications. PIMM could be relevant to other domains but the foundational results are firmly situated within our results, so I consider the model is best suited for casual InfoVis design.

6.8 Conclusion

Our cultural probe study examined how people use visual representations of data in their daily lives. We inferred that casual users continuously perform a perceived cost/benefit evaluation when using a visual representation of data. We then identified and categorized factors influencing artifact use as either promoters (encouraging system use) or inhibitors (discouraging system use). We propose PIMM, a model of information visualization use in casual contexts based on regulating factors influencing motivation, and derive casual InfoVis design considerations based on PIMM.

Our results illustrated that artifacts were always used to achieve some goal, even if it was just to alleviate curiosity. We identified four major categories of goals: extrinsic, entertainment, utility, and learning. Based on our observations, we suggest that designers consider the goal type and intended duration / number of uses, reduce

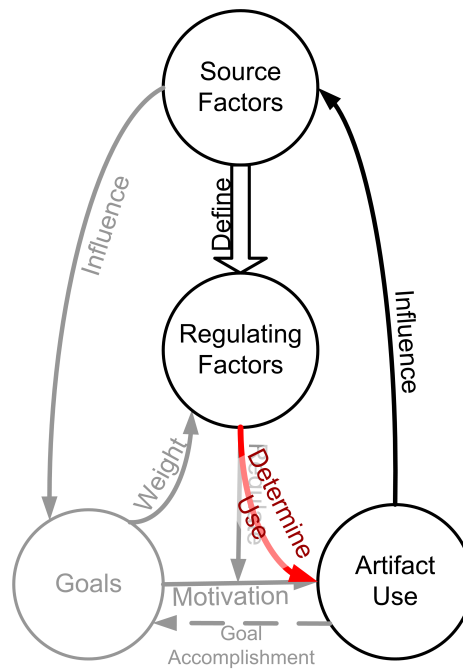


Figure 6.10: PIMM when the primary goals surrounding visualization /system use are uniform across users. Source factors determine regulating factors which can be conceptualized as directly influencing artifact use.

costs to initial use (especially learning effort), increase initial personal relevance of the data when possible, provide clear utility, organize information topically, and provide context and related information, perhaps in the form of text summaries. Our findings open exciting new motivation based research avenues, identifying goals and regulating factors associated with system use, which we hope to potentially explore in the future.

Table 6.1: Definitions provided to clarify points of discussion in Chapter 6. In some cases, terms can be defined numerous ways and an explicit defining of a term can frame the scope of the discussion.

Term	Definition
Visual Representations of Data (Artifacts)	Any means to visually represent or symbolize quantifiable information without solely using alpha-numeric representations. Data encoding methods may include using size, color, shape, position, orientation, and motion. For the study, artifacts required two or more non-alphanumeric data encodings to be valid.
Motivation	The internalized desire or drive to pursue and accomplish a task or activity.
Usefulness	The perceived or expected benefit from using a visualization relating to task accomplishment. Useful visualizations do not necessarily improve task performance measures. For example, a user with a learning goal would find a visualization that improves understanding useful, even if the system took longer to use (efficiency).
Utility	Objective task performance improvements (such as efficiency) associated with using a visualization.
Curiosity	The internalized drive to discover or understand previously unknown information.
Information Acquisition / Learning	A change in cognitive functions due to an external stimuli that can potentially be used at a later date. These changes typically include the appropriation or modification of a skill, memory, or strategy.
Data Exploration	The open-ended act of examining a data set to understand individual set elements and potentially the relationships between them.
Self-Reflection	A person's contemplation about personal actions, activities, and thought processes in an attempt to better understand his or her conscious self and the factors influencing this.
Procrastination	Activities performed to avoid a disliked primary task. E.g. Internet use instead of writing a doctoral dissertation.
Wasting Time	Activities performed when no alternate interesting or necessary tasks are available. Activities are presumably done to avoid boredom.
Personal Interest	Information that aligns with a person's interests, hobbies, cultural history, or personal experiences.
Social Pressure	The motivation to perform a task in order to appease or not frustrate / offend other people. Examples included examining a web page based on peer suggestions and looking at a hand drawn chart despite not understanding the topic of conversation.
Higher Priority Task	Any task considered more important to perform than using the artifact. Choosing to use the artifact reduces the possibility the important task can be completed to the user's satisfaction. For example, going to class was a higher priority task than examining a randomly encountered bar graph.

Table 6.2: Regulating factors identified via interviews and journal entries. Regulating factors are grouped based on the source factor that influenced them the most. The number of participants reporting each factor is indicated in parentheses. Promoter cells are coloured green, inhibitor cells are coloured red, and factors that can promote and inhibit are in white cells. Letters P and I indicate expected state influences based on PIMM; these are in gray cells if the effect is predicted but not verified. Commonly reported factors are underlined. A matrix (right) identifies which goals and stages of artifact use regulating factors were associated with.

	Regulating Factor	Examples	Goals				Stages				
			Utl	Lrn	Ent	Ext	Recog.	SS	IS	SR	LR
Context	Efficiency (11)	Blog charts, ticket purchase system								P	P
	Higher Priority Tasks (8)	Physics charts, earthquake map					I		I	I	
	Social Pressures (8)	Wolfram Alpha, World map showing where friends travelled						P	P	P	P
	<u>Time Constraints (6)</u>	Floor plans, campus map, Slashdot survey					I		I	I	
	<u>Usefulness (14)</u>	Campus map, weather forecast							P	P	P
	User Focus on Artifact (6)	Planting chart, dry-wall instructions							P		P
Data	Completeness(4)	Planting chart, weather forecasts						P	P	P	P
	Complexity (5)	Bike trail map, complex building map					I	I	P	I	P
	Correctness/Trust (10)	Mortality rate vis., GPS, battery icon							P	P	P
	<u>Curiosity (20)</u>	Planting chart, Slashdot survey					P	P	P		P
	Data Set Comfort (8)	Scrabble board, campus maps, Civilization IV line graphs					P	P	P		
	<u>Perceived Cost of Misuse / Misinterpretation (9)</u>	Bus map, Google maps, GPS, weather forecast					I		P		P
Representation	Aesthetic Appeal (10)	Planting chart, blog charts					P	P	P	P	P
	Appropriate Detail (8)	Downtown Vancouver map, GPS						P	P	P	P
	<u>Artifact Learning Effort (9)</u>	Earthquake map, planting chart					I	I	I		
	Chart Junk (5)	Complex emergency floor plan								I	I
	Data Clarity (6)	Planting chart, dry-wall screw instructions						P	P		
	Insufficient Data Context (4)	Newspaper article, peak oil in USSR chart						I	I		
User	<u>Personal Interest (18)</u>	Planting chart, blog charts, physics chart, Tree of Life diagram					P	P	P	P	P
	Potential Use (3)	Bus map, Wolfram Alpha						P	P		

Table 6.3: Artifact examples from of the PIMM stages of use and an associated description of activities.

Example	Stage	Description
Globe	Recog.	Broccoli reported a globe that she randomly encountered and ignored, but later recalled being an artifact when filling out her journal (see image A, Figure 6.5). When asked, she stated that although she liked the colours (aesthetic promoter), she was doing something else at the time and had no purpose for the object.
Immigration Chart	SS	Gravity briefly examined a magazine chart showing how the cost per apprehension of illegal immigrants in the United States has changed over ten years (see image A in Figure 6.6). The data was personally meaningful as Gravity spent years living in Mexico. The background image attracted Gravity’s attention, and he felt that it gave context and improved aesthetics. Gravity misread the chart as he interpreted the data as the cost to handle illegal immigrants (and not per immigrant), but he was not upset about his misinterpretation because the data was not critical to him. Gravity presumably examined the chart based on curiosity but his curiosity was probably satisfied quickly given he missed some graph details. Other tasks to perform (reading the magazine) may have also led to the short artifact use.
Earthquake Risk Map	LS	Benjii encountered an earthquake risk map of Victoria while waiting for a lecture (see image A, Figure 6.7). She explored the map, which was near the lecture hall, to fill the time. Her first priority was to locate her own home on the map, so personally meaningful information was important. She was not in a rush and continued looking at the map until the lecture began. The map was also sufficiently novel that it may have taken her more time to learn how to read it and was complex enough to allow her to continue examining it after her initial goal was achieved.
Wolfram Alpha	LS	116958 examined Wolfram Alpha based on a peer recommendation, claiming in her journal “Friend described search site for charts. Used out of curiosity about the site itself. May not go back.” Wolfram Alpha is an information search engine that directly provides topical information, including applicable graphs and charts, based on the search terms (see image B, Figure 6.7). During the interviews, 116958 claimed she examined the site for approximately ten minutes until she felt she understood what it did, which was her primary goal. 116958’s exploration was conducted during relaxed leisure time (no time pressure), and on the advice of a friend (secondary extrinsic goal or social promoter).
Weather Forecast	SR	Broccoli reported the weather forecast widget on her computer that she checks almost every day (for less than a minute) to determine what to wear (see image A, Figure 6.8). Since Broccoli was very familiar with the artifact, there was no learning curve. She examined the widget just before leaving the house, suggesting a known time constraint.
Planting Chart	LR	Gorffy regularly examined a text-based chart of plant growing times and conditions to aid his gardening hobby (image 1C, Figure 6.1). A friend recommended the chart but he used it again due to its functionality. The chart also piqued his curiosity so he examined the artifact for what could be grown locally. A more detailed description of the artifact is in the main text.

Chapter 7

Revisiting the Earlier Studies

PIMM not only helps to explain the actions and opinions surrounding artifact use in Chapter 6 but may also help explain the results in Chapters 4 and 5. I will re-examine these studies to discuss the results in light of the model and how we might have improved the research in retrospect.

7.1 Revisiting PartyVote

7.1.1 PartyVote According to PIMM

The results of the PartyVote study seem quite comprehensible when examined via PIMM. First, minimizing effort and learning costs (inhibitors) should greatly improve initial system use, eliminating the need for a brief introductory tutorial. Minimizing learning costs should also reduce the mild sense of intimidation felt by two participants in the second PartyVote study. Replacing the multi-dimensional scaling with axes that are meaningful to users would also seem to help reduce visualization learning costs. The interface and visualization could also be simplified such as removing extraneous animations and simplifying how voting worked.

Second, maximizing initial system appeal by quickly revealing potential benefits and encouraging curiosity about the system should attract more use. Both suggestions act as temporary promoters but following this, actual system usefulness would then promote repeat use. Entertainment (play), learning how the system worked, the music selection utility of using PartyVote and social pressures all affected PartyVote use. However, the playful actions of Pete negatively affected the utility based goals of Scott. Hence, the ability to remove a participant's influence once he or she left the

party and voting preview functionality may allow people to play with or learn about the system without negatively influencing others.

As previously stated, social pressures did not constrain PartyVote use as expected. PIMM can explain this finding based on user goals, the strength of the social inhibitor, and alternate promoters. Amy predominately had utility based goals as she attempted to influence the music played. Scott also had utility goals, but the benefits of using PartyVote extensively did not outweigh the time costs (inhibitor); hence his minimal interactions each time he voted. Pete attempted to diversify the music played and so the benefits of voting multiple times (promoter) outweighed perceived inhibitory factors (such as time and social costs). Bill's goal was to promote music he wanted to hear and learn how PartyVote operated. Thus choosing a song he enjoyed was a strong promoter and did not constrain himself based on group norms since he believed there would be minimal social repercussions with his friends (minimal inhibitor). The utility goals of the users also suggest that the playful space metaphor should not be used as entertainment. The playfulness would be a weak promoter and motion based distractions would strongly inhibit use.

Expected social pressures were minimal for both studies and thus social norms did not sufficiently inhibit and control user behaviors. For social norms to be effective, the perceived costs for breaking a rule must exceed the perceived benefits of the action. Thus a music moderator with the ability to delete a person's vote could effectively strengthen social pressures. Multiple votes could be met with group scorn and the benefit of the multiple votes can be eliminated (by removing these votes).

There was also a clear difference between utility based use, play (entertainment), curiosity (learning), and extrinsic pressures. Pete voted repeatedly out of curiosity and to learn how the system operated. Thus, he spent more time examining the system and music collection than actively voting. Amy, however, had game-like goals, strategically voting for music in an attempt to cancel Bill's selections. This led to extensive use of the music map and long term system use. Scott selected music from the purely utilitarian perspective, ignoring the music map and assuming that his votes would influence the music while minimizing the work he needed to do. Extrinsic social pressures were evident based on the clustered voting behaviors of the group. I hypothesize that when one person voted, others were reminded about the system, they wanted to ensure they maintained their influence over the music, or they were curious as to how the music space may have changed. Alternately, the act of voting by another person may have coerced participants into using the system again.

As an unfortunate side-effect of the research method, participants during the house party could have felt some inadvertent extrinsic pressure from me to use the system. However, interviewed participants claimed this was not the case. Bill's music choices led to a surprising lack of social pressure. Both Amy and Scott overtly complained about his music choices, but he felt they were good enough friends to comfortably ignore their protests. However, Bill's votes may have provided an extrinsic pressure for Amy, promoting longer term use for her.

Entertainment goals can be mutually exclusive to utility and learning goals. Using PartyVote as a game (such as Amy's strategic voting) annoyed users who did not wish to use the system this way (such as Bill). Similarly, when Pete was learning to use PartyVote, his votes unfairly provided him more influence over the music, potentially annoying people with strong utility goals. Constraining PartyVote usage patterns when possible may aid emphasizing utility based goals and discouraging entertainment goals. Alternately, indicating system rules may help reduce conflicting user goals. Otherwise user goals need to not negatively influence other participants.

7.1.2 Potential Improvements

When designing the PartyVote system, foundational casual InfoVis research was quite limited. It was not immediately clear how to design both for casual environments and for multiple simultaneous users. Design guidelines for both of these circumstances are still relatively unknown, but it is clear that many of our assumptions about appeal, social norms, and immediate usability were not justified. The social environment and user goals surrounding PartyVote use were far more complex than expected. People used the system as a game, out of curiosity, to influence the music in general, and to choose a particular song. Social norms also did not influence behaviors as much as expected.

A more iterative development process would likely have aided system design (and resulted in the design changes discussed below). For example, in retrospect, we would likely not have used multi-dimensional scaling, a convex hull, and the music map as users repeatedly had difficulties understanding how music could be represented in two dimensions, particularly when the dimensional axes had no inherent meaning. In retrospect, I think a different visualization from the music map may lead to less confusion and increased visualization use. A music map with clearly identified axes may also be less confusing (although more difficult to choose).

A node-link diagram with edges representing songs similarities above a threshold level may have also alleviated the confusion surround multi-dimensional scaling. Instead of a convex hull, votes and weights would determine how many hops from the voted for song are permitted. Songs within that distance would then be playable. Of course there are technical issues with node-link diagrams such as how many nodes to display at once and if / how to dynamically change the layout.

Immediate usability seems to be a critical factor based on the artifact study, given that potential embarrassments could inhibit use [59, 73], learning costs may be considered inhibitive, and choosing a party's music is likely a modest promoter. Provided choosing music was a stronger promoter, I would expect more self-designated DJs during regular social gatherings. Hence, for a modest promoter to encourage system use, inhibiting factors (such as system learning costs) must be even milder. The playful space metaphor and extraneous motion also do not appear justified.

Perhaps our system assumption that was least justified was the space travel analogy and the extraneous motion. The motion encoding study clearly identified that these stylistic additions may have annoyed or distracted some users. Furthermore, the PIMM model proposes that such additions may only be considered beneficial in some limited contexts such as system demonstrations. Most participants in the house party study concentrated on using the system to alter the party's music (utility goal) and entertainment/aesthetics were therefore probably less important. Based on user interviews, the space travel analogy did not improve the system's fun factor and was most appreciated by people playing games with the data or exploring out of curiosity. People who wanted to vote quickly ignored the music map completely, and people using the system to choose music appeared to use PartyVote as a tool rather than as a source of entertainment itself.

PIMM suggests that the context of system use affects regulating factors such as social norms, system learning costs, curiosity and social pressures. The importance of context was distinctly noted during the house party. Voting occurred in bursts, presumably with one participant's actions increasing curiosity about the system or adding social pressures (encouraging people to vote again so that their influence is not diminished). Other people using the system may have also reduced perceived learning costs as they could help if technical issues arose. However, embarrassment in such a social situation with a publicly viewable system may have hindered trial and error experimentation of the system. Potential embarrassment was most noted by Judy and Mary, who minimally used the system and stated (during the party)

that they did not feel comfortable using the system after arriving so late.

The Coffee House study, despite not revealing much about how well PartyVote was designed, clearly illustrated how novelty and duration of system use influenced user goals and regulating factors. Many participants treated PartyVote as a demonstration, choosing music to increase the convex hull size or to see what would happen rather than to vote for music. In such situations, I believe the space travel analogy may have been more successful. Votes cast during the house party appeared to be more thought out and personally meaningful to the user. This may have been due to the party's longer duration (and higher cost for choosing an unliked song) or the fact that demos were expected at the coffee house. In retrospect, a multiple session use of the PartyVote system would likely be much more indicative of regular system use, with system novelty and curiosity not driving use. System intimidation would also not inhibit use. Social norms associated with PartyVote use may also begin to be established with repeat sessions.

7.1.3 Redesigning the PartyVote Study

Lessons learned during this thesis work have led me to contemplate how the PartyVote research could have been conducted more ideally. A controlled experiment would still likely not have identified important design considerations about the system. However, repeat single user pilot testing could help system refinement.

A more ideal context to test PartyVote would seem to be over a span of several parties with predominately the same participants or by releasing PartyVote to the general population. A multi-session study would help eliminate novelty and system curiosity promoters plus minimize the effect of the researcher's presence. We noted that novelty was a strong motivator for initial artifact use but this quickly ceased. A multi-session study would also help identify how system use practices changed over time and how social norms might develop around the system. Unfortunately, multi-session studies can be difficult to design, take considerably more time to run, and many participants may not return. Furthermore, since the music was from my collection, participants may not want to use the system again, instead opting to bring their own music. This was noted when PartyVote was informally tested in France. Attendees initially used the system but decided to instead plug in their iPods to play their own music later that night. Despite these issues, however, multiple sessions do seem like the best way to realistically investigate PartyVote and how system use

patterns change over time.

Ideally, releasing PartyVote to the general public could dramatically increase the population size examined. In such a scenario, data logs above a minimal duration and number of voters (to ensure actual parties are being examined) could be analyzed to identify how the system is being used. A follow up questionnaire or interviews of willing users could then clarify results. However, the effort required to refine PartyVote to be a deployable product seems prohibitively high, given there is no guarantee of user adoption. Thus deploying PartyVote could lead to minimal results. For example, sound similarity analysis with bextract was an arduous process to set up and would not be easily performed by end users. Furthermore, multidimensional scaling times do not scale linearly. The music library used for the studies required more than ten minutes to run each time and far more time with just a modest increase in library size. PartyVote was also only extensively tested on one computer. These implementation issues could effectively be ignored for the PartyVote studies, but the setup process would have to be greatly streamlined before deployment is possible.

7.2 Revisiting the Motion Encoding Studies

When interpreted via PIMM, the motion encoding studies described in Chapter 5 seem to make more sense. Each of the three user groups described tended to have distinctly different goals when it came to the playlist creation task.

- **Pragmatics:** The eight subjects primarily concerned about task performance had predominately technical backgrounds. As suggested by PIMM, user experiences can dictate how goals are framed. Furthermore, given their performance based goals, any system feature that reduced productivity would be considered inhibitory. This was exemplified by the pragmatics' dislike of motion. A productivity emphasis also meant that searching for familiar music was probably considered more efficient than finding new music. The time costs associated with finding or identifying an appropriate unfamiliar song would be inhibitory to visualization use. Beat contradictions would therefore be more frequent since the beat information was ignored for efficiency's sake. The fact most pragmatics were male and alternative music fans may be a side effect of their educational background/demographic.

- **Hedonics:** Hedonic participants typically mentioned the importance and appeal of beauty, fun, or aesthetics relating to the beat encodings. As such, enjoyment or entertainment goals when performing the playlist generation task may have been more common. Provided the playlist generation rate was not a primary concern, then the performance decreasing effects of motion may not constitute much of an inhibitory factor. Aesthetic appeal, conversely, would promote system use more, provided entertainment is the user's goal. Time costs to sampling songs would also be less inhibitory and the enjoyment of discovering new music may promote listening to new music. The lack of beat contradictions in the hedonics group suggests that finding appropriate music irrespective of time was a primary concern. Again, the non-technical background associated with hedonic participants may have primed or influenced the goals they set. A preference for pop music may be due to the associated demographic.
- **Intuitives:** The intuitive group found motion was a more direct, easier to read, or intuitive way to encode beat information than the curved lines. As such, goals associated with intuitives are harder to theorize. I presume that the more natural encoding would make moving glyphs more pleasant to read, despite any performance costs. As such, inhibitory performance costs associated with motion encoding would be minimal. The lack of motion based costs is reflected in the beat encoding's high rate of use in this group and the lack of beat contradictions. The results also imply that the intuitive nature of the motion encoding provided had a promoting effect on beat use. The redundant condition, however, was most preferred by two of the three group members, perhaps indicating that motion was natural but the static encoding was perceived as more efficient. The gender based clustering suggests that either women find motion based encoding more natural or that the majority of non-technical background participants (who did not concentrate on productivity) were women. With only three participants, the fact only women constituted this group could also be due to random chance. In fact the intuitives group could simply be a miscellaneous collection of participants that did not fit into the other groups. Hence care should be taken when considering this group.

Overall, the clustering associated with the playlist generation study can best be explained by the disparate goals each group had. Since user goals were not uniform, the promoting and inhibiting factor weights differed between groups, resulting in

the identified clusters. The identified importance of utility (promoting factor) and perceived ease of use (an inhibiting factor when lacking) also supports the explanatory power of PIMM. More research is required, but provided the user's primary goal can be identified or dictated, it appears that user clusterings could be minimized.

7.2.1 Redesigning the Motion Studies

The clustering associated with the motion encoding studies was an accidental but interesting result which helped lead to PIMM. However, with the PIMM model now being available, changes could be made to the playlist generation task to potentially improve results. We attempted to make the playlist generation task as realistic as possible by avoiding task constraints. This appears to lead participants to frame their own goals based on their previous experiences, training, and how the task is framed. Actively controlling and / or manipulating user goals seems necessary to ensure consistent results. For example, although we asked participants to find music for exercising, we perhaps should have asked people to find music above a minimum beat strength or frequency. Although this change might seem to reduce realism, the underlying tasks (some) participants performed would be the same. The key caveat is identifying the underlying heuristics most participants used for finding music.

7.3 Conclusion

Overall, the studies conducted during this doctoral research effectively built upon each other and examined casual InfoVis system design from three complementary perspectives. Although in retrospect the PartyVote and motion encoding studies did not ideally provide revolutionary results, they were appropriately structured and conducted based on our knowledge at the time. The study revisions proposed above therefore do not demonstrate how the studies failed, but instead how they would be conducted now given what was learned. I have identified the difficulties in conducting mixed-method controlled / exploratory studies and I suggest that system novelty and divergent user goals can pose major problems to identifying research conclusions. For short term studies such as the motion encoding studies, extracting actions performed while performing realistic tasks can help researchers introduce more structure to tasks while maintaining some realism. System novelty can also be reduced using a multi-session in-situ approach.

Chapter 8

Conclusions

My doctoral research has used a variety of casual InfoVis research methods to identify what motivates people to use such visualizations. My research progressed from a specific system design based on assumptions about casual use (the PartyVote system: Chapter 4) to a semi-controlled analysis of utility / appeal trade-offs surrounding a single design feature (the motion study motion: Chapter 5). The motion study then directly led to the artifact study (see Chapter 6) which explored how a subset of the casual InfoVis population examined data in their non-work time. Artifact study results directly led to my postulation of the Promoter / Inhibitor Motivation Model (PIMM) of casual InfoVis use. Not only did each research step progress based on the previous research, but the three main research topics used disparate approaches. As a corollary, the studies enabled me to suggest casual InfoVis methods based on the current state of the field.

8.1 Extending the Research

The current research, framed within Munzner's [90] nested model of visualization design and validation, has explored the domain problem, data abstraction, and encoding technique levels of casual InfoVis use. The PartyVote study attempted to validate the system's final design (data abstraction level) and the motion encoding study explored user opinions surrounding motion based data encoding techniques (the encoding technique level). I argue that the motion encoding study also explored the domain problem, examining factors that influence casual InfoVis use. Finally, the artifact study provided an exceptional case of problem domain characterization using

more than just interviews [90]. The logical next step to follow up on this research is the downstream evaluation of an implemented casual InfoVis system to re-examine the domain characterization. Downstream evaluations can be done by analyzing log data from current ‘InfoVis for the Masses’ systems such as Many Eyes or Tableau Public, however, gaining access to these systems can be difficult. Alternately, designing and deploying a visualization could also address this issue. I contemplated a deployed version of PartyVote, but the time costs associated with this seemed prohibitive.

Finally, PIMM has not been extensively tested or validated. As such, numerous possible research directions derive directly from searching for model support. For example, a controlled study examining how user priorities and system satisfaction change with user goals seems promising. Participants could be asked to conduct a relatively simple visualization search task but be given different goals to accomplish such as “do the task as fast as possible”, “beat the current fastest time” and “have fun.” By varying the context of use (and hopefully the participant’s internalized goals as well), I could then examine how opinions about the visualization differed. In particular, what system factors promoted system use and which inhibited them. Although a further examination of data abstractions and encodings techniques could be beneficial, the scale of such research appears to not address my high level research interests: examining why people use casual InfoVis.

Overall, I hope future PIMM related studies will examine both the details surrounding individual model components (source factors, goals, etc.) and the applicability of PIMM to other domains. An exploration of long-term visualization adoption in relation to PIMM would also be extremely useful. Furthermore, provided future evidence supports PIMM, a further refinement and expansion of PIMM based design suggestions seems appropriate. I sincerely hope that PIMM will be a useful framework to evaluate and design casual visualizations around and I feel my research constitutes an important first step.

8.2 Contributions

Motivation or reason to use a system is not always ingrained in casual systems. This results in potentially different use behaviors where utility and productivity are not the only critical metrics. Considering what factors motivate people to use casual systems is also important. Although traditional discussions about casual InfoVis have been framed around an implicit productivity / appeal trade-off [31, 24, 76], our results and

the proposed Promoter / Inhibitor Motivation Model indicate such a framing may be counter-productive. Instead, the PIMM model implies that previous experiences, use context and system design direct user goals, which in turn determine what effect system factors have on visualization use. Thus, instead of the productivity / appeal evaluations we initially examined, the fit between the system and a user's goals is the critical test (an idea also proposed by the task-technology fit model).

The proposed PIMM model is a primary contribution of my research. PIMM provides both a testable framework for determining user satisfaction and visualization adoption and a mental framework for designers. The PIMM model also opens research avenues for validation. As such, this is a potential first step into exploring how casual visualizations are used with this dissertation's three studies acting as initial evidence.

System motivation over time was another major topic of my dissertation. The proposed PIMM model provides a framework to test whether a system will be used initially and repeatedly. The model's six stages of artifact use (the time component) identify how user motivations change with experience and time. The proposed stages clearly identified that curiosity, learning, and fun tended to be effective but temporary system promoters, while visualizations related to personal hobbies and personally meaningful information motivated repeat long-term use. Thus utility, aesthetics, and entertainment all promote system use, however, known utility benefits promote use in the longer term.

Finally, I explored casual InfoVis design methods when conducting this research. Numerous methodological approaches are proposed in Chapter 3.

Several minor research contributions not directly relating to the dissertation's research question are also identifiable. First, PartyVote was a novel, light weight group behavior mediating visualization (see Chapter 4). PartyVote also relied on group norms and participant visibility to ensure fair system use; a novelty for music visualizations. The system required no user profiles and minimal user interaction but still ensured each person could contribute to the music played.

The motion encoding study explored the benefits and costs of using motion to encode data for a music library visualization (a presumed optimal place to use motion encoding. See Chapter 5). The motion study demonstrated that subject opinions were strongly task and experience dependent. Furthermore, participants were clusterable into distinct groups. Future research may enable quick visualization customizations based on distinct user groups. The motion study also provided evidence that some participants are willing to sacrifice system utility for visual appeal and aesthetics.

Design Implications

Numerous design implications derive directly from my studies. First, clearly identifying user goals and desires with the system can greatly aid system adoption. Unlike PartyVote’s open design, the PIMM model indicates that constraining user goals to a subset of activities can help system features match user tasks. The diversity of PartyVote’s potential uses (social games, minimal interactions, music exploration, strategic voting, etc) meant system features designed to promote use (such as the ability to strategically vote) actually were sometimes inhibitory. Constraining use options, either via customizations or permanently, may promote system use for specific tasks. Furthermore, ‘jack of all trades’ systems may not adequately appease users for any particular task. Hence, casual InfoVis research should attempt to identify goals and commercial systems should identify what the system is being used for (the problem domain [90]) to contextualize any evaluation metrics. Constraining priorities was a problem identified during the motion encoding study; different user groups had differing task goals. Thus, opinions about system features diverged despite all participants experiencing the same conditions. Constraining use tasks is easier said than done, but promoting diverse use patterns should at least be discouraged if the variety of possible user goals cannot be sufficiently supported.

The time component of PIMM indicates that long-term repeat use visualizations tended to have a known utilitarian component while entertainment, aesthetics, and curiosity based goals were predominately found with single use or short-duration use visualizations. Visualization re-use can be promoted by encouraging utility or by identifying personally relevant information to the user. Knowing the desired visualization use duration can also help guide design decisions. For example, a recipe visualization designed for repeat long duration use might clearly identify recipes similar to previously found recipes or related to known user interests (e.g. Star Wars cakes) and the visualization design should concentrate on improved user efficiency rather than discovering inter-recipe relationships (such as how many cheesecake recipes use ricotta cheese).

8.3 Final thoughts

Overall, the mix of exploratory and confirmatory research conducted for my dissertation evaluated casual InfoVis use patterns and user motivations to gain a better

understanding of why people might use visualizations in their leisure time. Casual visualizations are currently only beginning to be developed and this research provides an opportune lens through which to design and interpret new casual InfoVis systems. My dissertation also describes a model and methodology framework for other researchers to build upon. I hope that future casual InfoVis development and studies will be greatly facilitated as a result of my research.

Bibliography

- [1] AGIV. Thermografische kaart, 2010. Available at <http://zoominopuwdak.antwerpen.be/>, accessed June 16, 2011.
- [2] Christopher Ahlberg and Ben Shneiderman. Visual information seeking using the filmfinder. In *Conference companion on Human Factors in Computing Systems (CHI '94)*, pages 433–434, 1994.
- [3] Christopher Ahlberg, Christopher Williamson, and Ben Shneiderman. Dynamic queries for information exploration: an implementation and evaluation. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '92)*, pages 619–626, 1992.
- [4] Stephen P. Anderson. The fundamentals of experience design, 2009. Available at <http://www.poetpainter.com/thoughts/article/ia-summit-2009-the-fundamentals-of-experience-design->, accessed March 25, 2011.
- [5] Antonella De Angeli, Alistair Sutcliffe, and Jan Hartmann. Interaction, usability and aesthetics: what influences users' preferences? In *Proc. Conf. Designing Interactive Sys.(DIS '06)*, pages 271–280, 2006.
- [6] M. Anshel and D. Marisi. Effects of music and rhythm on physical performance. *Research Quarterly*, 49:109–112, 1979.
- [7] Yannick Assogba, Irene Ros, Joan DiMicco, and Matt McKeon. Many bills: engaging citizens through visualizations of congressional legislation. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*, pages 433–442, 2011.

- [8] The Take Away. The take away's electoral college prediction tracker, 2008. Available at <http://vote2008.thetakeaway.org/2008/09/20/track-the-electoral-college-vote-predictions/>, accessed June 16, 2011.
- [9] R. Bartle. Hearts, clubs, diamonds, spades: players who suit MUDs. *The Journal of Journal of MUD Research*, 1(1), 1996. Available at <http://www.mud.co.uk/richard/hclds.htm>, accessed June 16, 2011.
- [10] Lyn Bartram. Can motion increase user interface bandwidth in complex systems? In *Intl. Conf. Systems, Man, and Cybernetics*, volume 2, pages 1686–1692, 1997.
- [11] Lyn Bartram, Colin Ware, and Tom Calvert. Moticons: detection, distraction and task. *Int. J. Hum.-Comput. Stud.*, 58(5):515–545, 2003.
- [12] Lynn Bartram, Johnny Rodgers, and Rob Woodbury. Supporting sustainable living: Aware homes and smart occupants. In *UbiComp 10 Workshop on Ubiquitous Computing for Sustainable Energy*, 2010. Available at <http://www.hcilab.org/events/ucse2010/images/a/a8/UCSEpaper02.pdf>, accessed June 16, 2011.
- [13] Arianna Bassoli, Julian Moore, and Stefan Agamanolis. tunA: synchronized music-sharing on handheld devices. In *UbiComp '04 Adjunct Proceedings*, pages 171–172, September 2004.
- [14] Scott Bateman, Regan L. Mandryk, Carl Gutwin, Aaron Genest, David McDine, and Christopher Brooks. Useful junk?: the effects of visual embellishment on comprehension and memorability of charts. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*, pages 2573–2582, New York, NY, USA, 2010. ACM.
- [15] Benjamin B. Bederson and Angela Boltman. Does animation help users build mental maps of spatial information? In *Proc. IEEE Symp. Information Visualization (InfoVis '99)*, pages 28–35, 1999.
- [16] Izak Benbasat and Henri Barki. Quo vadis, TAM? *J. Assoc. Information Systems*, 8(4):211–218, 2007.

- [17] Mathew Bloch, Lee Byron, Shan Carter, and Amanda Cox. The ebb and flow of movies: Box office receipts 1986-2008, 2008. Available at http://www.nytimes.com/interactive/2008/02/23/movies/20080223_REVENUE_GRAPHIC.html, accessed June 15, 2011.
- [18] Bloom.io. Introducing fizz, 2011. Available at <http://blog.bloom.io/2011/02/25/introducing-fizz/>, accessed March 3, 2011.
- [19] Richard Borovoy, Fred Martin, Sunil Vemuri, Mitchel Resnick, Brian Silverman, and Chris Hancock. Meme tags and community mirrors: moving from conferences to collaboration. In *Proc. Computer Supported Cooperative Work (CSCW '98)*, pages 159–168, 1998.
- [20] John S. Breese, David Heckerman, and Carl Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proc. Conf. Uncertainty in Artificial Intelligence (UAI '98)*, pages 43–52, 1998.
- [21] Harry Brignull and Yvonne Rogers. Enticing people to interact with large public displays in public spaces. In *Proc. IFIP Conf. on Human-Computer Interaction (INTERACT 2003)*, pages 17–24, 2003.
- [22] burst Labs. burst labs home page, 2008. Available at <http://www.burstlabs.com/>, accessed May 16, 2008.
- [23] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman, editors. *Readings in information visualization: using vision to think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.
- [24] N. Cawthon and A. Vande Moere. The effect of aesthetic on the usability of data visualization. In *IV '07: Proc. 11th Intl. Conf. Information Visualization*, pages 637–648, 2007.
- [25] Dennis L. Chao, Justin Balthrop, and Stephanie Forrest. Adaptive radio: achieving consensus using negative preferences. In *Proc. GROUP '05*, pages 120–123, 2005.
- [26] John Creswell. *Qualitative inquiry and research design: choosing among five traditions*. Sage Publications, London, 1998.

- [27] John W. Creswell. *Research design : qualitative & quantitative approaches*. Sage Publications, Thousand Oaks, Calif., 1994.
- [28] Andrew Crossen, Jay Budzik, and Kristian J. Hammond. Flytrap: intelligent group music recommendation. In *Proc. Conf. Intelligence User Interfaces (IUI '02)*, pages 184–185, 2002.
- [29] Catalina Danis and Alison Lee. Evolution of norms in a newly forming group. In *Proc. IFIP Conf. on Human-Computer Interaction (INTERACT '05)*, 2005.
- [30] Catalina M. Danis, Fernanda B. Viégas, Martin Wattenberg, and Jesse Kriss. Your place or mine?: visualization as a community component. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, pages 275–284, 2008.
- [31] Mike Danziger. Information visualization for the people. MSc thesis, Massachusetts Institute of Technology, 2008, Available at <http://cms.mit.edu/research/theses/MichaelDanziger2008.pdf>, accessed June 16, 2011.
- [32] Fred D. Davis, Richard P. Bagozzi, and Paul R. Warshaw. User acceptance of computer technology: a comparison of two theoretical models. *Manage.Sci.*, 35(8):982–1003, 1989.
- [33] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer, Berlin, Germany, 1997.
- [34] digg labs. The dig labs swarm visualization, 2008. Available at <http://labs.digg.com/>, accessed May 16, 2008.
- [35] Marian Dörk, Sheelagh Carpendale, Christopher Collins, and Carey Williamson. Visgets: Coordinated visualizations for web-based information exploration and discovery. *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 14(6):1205–1212, 2008.
- [36] Yogesh Kumar Dwivedi, Michael D. Williams, Banita Lal, and Andrew Schwarz. Profiling adoption, acceptance and diffusion research in the information systems discipline. In Willie Golden, Tom Acton, Kieran Conboy, Hans van der Heijden, and Virpi Kristiina Tuunainen, editors, *16th European Conference on Information Systems*, pages 1204–1215, Galway, Ireland, 2008.

- [37] General Electric. Our aging world, 2010. Available at <http://www.ge.com/visualization/aging/>, accessed June 16, 2011.
- [38] Yuri Engelhardt. The language of graphics: A framework for the analysis of syntax and meaning, 2002. Doctoral thesis, the University of Amsterdam.
- [39] Stephen Few. Infovis as seen by the world out there: 2007 in review. In *IEEE Symp. Visual Analytics Science and Technology (VAST '07)*, 2007. Capstone Address. Available at <http://www.perceptualedge.com/articles/11-13-07.pdf>, accessed June 16, 2011.
- [40] Stephen Few. What ordinary people need most from information visualization today, 2008. Available at http://www.perceptualedge.com/articles/visual_business_intelligence/what_people_need_from_infovis.pdf, accessed June 16, 2011.
- [41] ffunctions Inc. Expense visualizer, 2010. Available at <http://ffctn.com/expensevisualizer>, accessed June 16, 2011.
- [42] Flowing Data. Youtube releases visualization for related videos, 2008. Available at <http://flowingdata.com/2007/12/18/youtube-releases-visualization-for-related-videos/>, accessed March 11, 2008.
- [43] Godfrey Gardner. Is there a valid test of herzberg's two-factor theory? *Journal of Occupational Psychology*, 50(3):197–204, 09 1977.
- [44] Bill Gaver, Tony Dunne, and Elena Pacenti. Design: Cultural probes. *Interactions*, 6(1):21–29, 1999.
- [45] Amir Ghahary. Towards an ideal in informative art, 2009. Simon Frasier University Master's Thesis. Document provided by Dr. Lynn Bartram.
- [46] Jennifer Gluck, Andrea Bunt, and Joanna McGrenere. Matching attentional draw with utility in interruption. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, pages 41–50, 2007.
- [47] Nathaniel S. Good and Aaron Krekelberg. Usability and privacy: a study of Kazaa P2P file-sharing. In *CHI '03 Conference Proceedings*, pages 137–144, 2003.

- [48] Dale L. Goodhue and Ronald L. Thompson. Task-technology fit and individual performance. *MIS Q.*, 19(2):213–236, 1995.
- [49] Connor Graham, Mark Rouncefield, Martin Gibbs, Frank Vetere, and Keith Cheverst. How probes work. In *Proc. Australasian conference on Computer-Human Interaction (OZCHI '07)*, pages 29–37, 2007.
- [50] Lars Grammel, Melanie Tory, and Margaret-Anne Storey. How information visualization novices construct visualizations. *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 16:943–952, October, 2010 2010.
- [51] Saul Greenberg and Michael Rounding. The notification collage: posting information to public and personal displays. In *Proc. SIGCHI Conference on Human Factors in Computer Systems (CHI '01)*, pages 514–521, 2001.
- [52] Jonathan Harris and Sep Kamvar. wefeelfine, 2006. Available at http://www.wefeelfine.org/wefeelfine_mac.html.
- [53] Marc Hassenzahl and Daniel Ullrich. To do or not to do: Differences in user experience and retrospective judgments depending on the presence or absence of instrumental goals. *Interact. Comput.*, 19(4):429–437, 2007.
- [54] Marti A. Hearst and Daniela Rosner. Tag clouds: Data analysis tool or social signaller? In *Proc. Hawaii International Conference on System Sciences (HICSS '08)*, page 160, 2008.
- [55] Jeffrey Heer. Visualization evaluation of the masses, by the masses, and for the masses, 2010. BELIV workshop position paper. Available at <http://www.beliv.org/papers/beliv2010/Position\%20Papers/25-beliv-jeff-heer.pdf>, accessed June 16, 2011.
- [56] Jeffrey Heer and Danah Boyd. Vizster: Visualizing online social networks. In *Proc. IEEE Symp. Information Visualization (InfoVis '05)*, pages 33–40, 2005.
- [57] Jeffrey Heer, Frank Ham, Sheelagh Carpendale, Chris Weaver, and Petra Isenberg. Information visualization. In Andreas Kerren, John T. Stasko, Jean-Daniel Fekete, and Chris North, editors, *Information Visualization*, chapter Creation and Collaboration: Engaging New Audiences for Information Visualization, pages 92–133. Springer-Verlag, Berlin, Heidelberg, 2008.

- [58] Jeffrey Heer, Fernanda B. Viégas, and Martin Wattenberg. Voyagers and voyeurs: supporting asynchronous collaborative information visualization. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, pages 1029–1038, 2007.
- [59] Uta Hinrichs, Holly Schmidt, and Sheelagh Carpendale. EMDialog: Bringing Information Visualization into the museum. *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 14(6):1181–1188, 2008.
- [60] Lars Erik Holmquist and Tobias Skog. Informative art: information visualization in everyday environments. In *GRAPHITE '03: Proc. 1st Intl. Conf. on computer graphics and interactive techniques in Australasia and South East Asia*, pages 229–235, 2003.
- [61] Elaine M. Huang and Elizabeth D. Mynatt. Semi-public displays for small, co-located groups. In *Proc. SIGCHI Conference on Human Factors in Computer Systems (CHI '03)*, pages 49–56, 2003.
- [62] Nike Inc. What is nike+?, 2010. Available at http://nikerunning.nike.com/nikeos/p/nikeplus/en_US/plus/#//dashboard/, accessed June 16, 2011.
- [63] Pet Holdings Inc. Graph jam: Life and pop culture graphed for your inner geek, 2010.
- [64] Information Architects. Web Trend Map 2008, 2008. Available at <http://informationarchitects.jp/start/>, accessed June 16, 2011.
- [65] Petra Isenberg, Torre Zuk, Christopher Collins, and Sheelagh Carpendale. Grounded evaluation of information visualizations. In *Proc. Conference on Beyond time and errors: novel evaluation methods for Information Visualization (BELIV '08)*, pages 1–8, 2008.
- [66] Anthony Jameson and Barry Smyth. Recommendation to groups. In *The Adaptive Web*, pages 596–627, 2007.
- [67] Hyunmo Kang, Catherine Plaisant, and Ben Shneiderman. New approaches to help users get started with visual interfaces: multi-layered interfaces and integrated initial guidance. In *Proc. Annual National Conference on Digital Government research (dgo '03)*, pages 1–6, 2003.

- [68] Robert Kosara. Rethinking the user, 2007. Available at EagerEyes.org: <http://eagereyes.org/blog/rethinking-the-user.html>, accessed June 16, 2011.
- [69] Robert Kosara. Visualization criticism - the missing link between information visualization and art. In *IV '07: Proc. 11th Intl. Conference Information Visualization*, pages 631–636, 2007.
- [70] Robert Kosara. Trivialization for the masses, 2010. Available at EagerEyes.org: <http://eagereyes.org/blog/2010/trivialization-for-the-masses>, accessed June 12, 2011.
- [71] Robert Kosara. Visualization can never be art, 2010. Available at EagerEyes.org: <http://eagereyes.org/criticism/visualization-can-never-be-art>, accessed June 12, 2011.
- [72] Robert Kosara. The visualization cargo cult, 2010. Available at EagerEyes.org: <http://eagereyes.org/criticism/the-visualization-cargo-cult>, accessed June 16, 2011.
- [73] Bill Kules, Hyunmo Kang, Catherine Plaisant, Anne Rose, and Ben Shneiderman. Immediate usability: a case study of public access design for a community photo library. *Interacting with Computers*, 16(6):1171–1193, 2004.
- [74] Florian Lambers. Party jukebox: Supporting co-located music consumption in a public environment, 2010. Research talk. Available at http://www.medien.ifi.lmu.de/lehre/ss10/os/material/DA_Abschlussvortrag_FlorianLambers.pdf, accessed April 10, 2011.
- [75] Last.fm. The social music revolution. Available at <http://www.last.fm/>, accessed June 16, 2011.
- [76] A. Lau and A. Vande Moere. Towards a model of information aesthetics in information visualization. In *Proc. International Conf. of Information Visualization*, pages 87–92, 2007.
- [77] Paul D. Leedy and J. E. Ormrod. *Practical research: planning and design*. Merrill Prentice Hall, Saddle River, NJ, USA, 2001.

- [78] Ralph Lengler and Martin J. Eppler. Towards a periodic table of visualization methods for management. In *Proc. Conference on Graphics and Vis. in Engineering (GVE 2007)*, 2007. Available at http://www.visual-literacy.org/periodic_table/periodic_table.pdf, accessed May 16, 2009.
- [79] Golan Levin, Kamal Nigam, and Jonathan Feinberg. The dumpster, 2006. Available at <http://artport.whitney.org/commissions/thedumpster/index.html>, accessed June 11, 2008.
- [80] Manuel Lima. Information visualization manifesto, 2009. Available at <http://www.visualcomplexity.com/vc/blog/?p=644>, accessed March 3, 2011.
- [81] E. A. Locke and G. P. Latham. Building a practically useful theory of goal setting and task motivation. a 35-year odyssey. *The American Psychologist*, 57(9):705–717, September 2002.
- [82] Kurt Luther, Scott Counts, Kristin B. Stecher, Aaron Hoff, and Paul Johns. Pathfinder: an online collaboration environment for citizen scientists. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*, pages 239–248, 2009.
- [83] Gloria Mark. Conventions and commitments in distributed CSCW groups. *Proc. Computer Supported Cooperative Work (CSCW '02)*, 11(3):349–387, 2002.
- [84] Joseph E. McCarthy and Theodore D. Anagnost. MUSICFX: an arbiter of group preferences for computer supported collaborative workouts. In *Proc. Computer Supported Cooperative Work (CSCW '98)*, pages 363–372, 1998.
- [85] Joseph F. McCarthy, David W. McDonald, Suzanne Soroczak, David H. Nguyen, and Al M. Rashid. Augmenting the social space of an academic conference. In *Proc. Computer Supported Cooperative Work (CSCW '04)*, pages 39–48, 2004.
- [86] Patients Like Me. Patients like me, 2010. Available at <http://www.patientslikeme.com/>, accessed June 16, 2011.
- [87] Meredith Ringel Morris, Anthony Cassanego, Andreas Paepcke, Terry Winograd, Ann Marie Piper, and Anqi Huang. Mediating group dynamics through tabletop interface design. *IEEE Comput. Graph. Appl.*, 26(5):65–73, 2006.

- [88] Ann J. Morrison, Peta Mitchell, and Margot Brereton. The lens of ludic engagement: evaluating participation in interactive art installations. In *Proc. ACM International Conference on Multimedia (MULTIMEDIA '07)*, pages 509–512, 2007.
- [89] Fabian Mrchen, Alfred Ultsch, Mario Ncker, and Christian Stamm. Databionic visualization of music collections according to perceptual distance. In *Proc. International Conference on Music Information Retrieval (ISMIR '05)*, pages 396–403, 2005.
- [90] Tamara Munzner. A nested process model for visualization design and validation. *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 15:921–928, 2009.
- [91] Petra Neumann, Annie Tat, Torre Zuk, and M. Sheelagh T. Carpendale. Keystrokes: Personalizing typed text with visualization. In *EuroVis*, pages 43–50, 2007.
- [92] Donald A. Norman. *The Psychology of Everyday Things*. Basic Books, 1988.
- [93] Donald A. Norman. *Emotional Design: Why We Love (or Hate) Everyday Things*. Basic Books, 2004.
- [94] Kenton O’Hara, Matthew Lipson, Marcel Jansen, Axel Unger, Huw Jeffries, and Peter Macer. Jukola: democratic music choice in a public space. In *Proc. DIS '04*, pages 145–154, 2004.
- [95] Yee Von Ooi. Can CASSM bridge the gap between InfoVis and the user?, 2009. Master’s thesis, available at http://www.ucl.ac.uk/uclic/taught_courses/distinction/2008_theses/2009-Ooi, accessed June 16, 2011.
- [96] Elias Pampalk, Andreas Rauber, and Dieter Merkl. Content-based organization and visualization of music archives. In *Proc. ACM International Conference on Multimedia (MULTIMEDIA '02)*, pages 570–579, 2002.
- [97] Pandora.com. Radio from the music genome project. <http://www.pandora.com/>, accessed June 16, 2011.

- [98] Eric Paulos and Chris Beckmann. Sashay: designing for wonderment. In *Proc. SIGCHI Conference on Human Factors in computing systems (CHI '06)*, pages 881–884, 2006.
- [99] Sarah Perez. The best tools for visualization, 2008. Available at http://www.readwriteweb.com/archives/the_best_tools_for_visualization.php, accessed July 16, 2009.
- [100] Catherine Plaisant. The challenge of information visualization evaluation. In *Proc. Working Conf. Advanced Visual Interfaces (AVI '04)*, pages 109–116, 2004.
- [101] The Washington Post. Top secret america, 2010. Available at <http://projects.washingtonpost.com/top-secret-america/>, accessed June 16, 2011.
- [102] Zachary Pousman. Thinky.org, 2008. Available at <http://thinky.org/research-casual-information-visualization/>, accessed March 12, 2011.
- [103] Zachary Pousman. Casual infovis in the workplace: Supporting worker reflections and conversations, 2010. Workshop paper at SIGCHI Conf. on Human Factors in Computing Systems (CHI 2010). Available at http://www.personalinformatics.org/docs/chi2010/pousman_casual_infovis.pdf, accessed June 16, 2011.
- [104] Zachary Pousman, Hafez Rouzati, and John Stasko. Imprint, a community visualization of printer data: designing for open-ended engagement on sustainability. In *Proc. Computer Supported Cooperative Work (CSCW '08)*, pages 13–16, 2008.
- [105] Zachary Pousman and John Stasko. A taxonomy of ambient information systems: four patterns of design. In *Proc. Working Conf. Advanced Visual Interfaces (AVI '06)*, pages 67–74, 2006.
- [106] Zachary Pousman, John Stasko, and Michael Mateas. Casual information visualization: Depictions of data in everyday life. *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 13(6):1145–1152, 2007.
- [107] Derek F. Reilly and Kori M. Inkpen. White rooms and morphing don't mix: setting and the evaluation of visualization techniques. In *Proc. SIGCHI Con-*

- ference on Human Factors in Computer Systems (CHI '07)*, pages 111–120, 2007.
- [108] Magnus Rembold and Jrgen Spth. `munterbund.de`, 2004. Available at http://www.munterbund.de/visualisierung_textaehnlichkeiten/essay.html.
- [109] Theresa-Marie Rhyne, Melanie Tory, Tamara Munzner, Matt Ward, Chris Johnson, and David H. Laidlaw. Information and scientific visualization: Separate but equal or happy together at last. In *Proc. IEEE Conf. Visualization 2003 (VIS'03)*, pages 115–118, 2003.
- [110] Jeffrey A. Roberts, Il-Horn Hann, and Sandra A. Slaughter. Understanding the motivations, participation, and performance of open source software developers: A longitudinal study of the apache projects. *Manage. Sci.*, 52(7):984–999, 2006.
- [111] George Robertson, Roland Fernandez, Danyel Fisher, Bongshin Lee, and John Stasko. Effectiveness of animation in trend visualization. *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 14(6):1325–1332, 2008.
- [112] Hans Rosling. The best stats you’ve ever seen, 2006. Available at http://www.ted.com/talks/hans_rosling_shows_the_best_stats_you_ve_ever_seen.html, accessed January 31, 2011.
- [113] Edward Segel and Jeffrey Heer. Narrative visualization: Telling stories with data. *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 16:1139–1148, 2010.
- [114] Bernhard Seifert and Oliver Schneider. Windirstat, 2010. Available at <http://windirstat.info/>, accessed June 16, 2011.
- [115] Maruthappan Shanmugasundaram, Pourang Irani, and Carl Gutwin. Can smooth view transitions facilitate perceptual constancy in node-link diagrams? In *Proc. Graphics Interface 2007*, pages 71–78, 2007.
- [116] Ben Shneiderman. Promoting universal usability with multi-layer interface design. In *Proc. Conference on Universal Usability (CUU '03)*, pages 1–8, 2003.
- [117] Ben Shneiderman and Catherine Plaisant. Strategies for evaluating information visualization tools: multi-dimensional in-depth long-term case studies. In

Proc. AVI workshop on BEyond time and errors: novel evaLuation methods for Information Visualization (BELIV '06), pages 1–7, 2006.

- [118] Martin Shuttleworth, Mark McCormick, Christine Oliver, Lisa Villani, and Alastair Dan. World cup 2010 twitter replay, 2010. Available at <http://www.guardian.co.uk/football/world-cup-match-replay>, accessed January 11, 2011.
- [119] Nelson Siu, Lee Iverson, and Anthony Tang. Going with the flow: email awareness and task management. In *Proc. Conf. Computer Supported Cooperative Work (CSCW '06)*, pages 441–450, 2006.
- [120] James Slack, Kristian Hildebrand, Tamara Munzner, and Katherine St. John. Sequencejuxtaposer: Fluid navigation for large-scale sequence comparison in context. In *German Conference on Bioinformatics*, pages 37–42, 2004.
- [121] Slife Labs. Slife labs home page, 2007. Available at <http://www.slifeweb.com/>, accessed June 16, 2011.
- [122] Smashing Magazine. Data visualization: Modern approaches, 2008. Available at <http://www.smashingmagazine.com/2007/08/02/data-visualization-modern-approaches/>, accessed July 16, 2009.
- [123] David Sprague, Fuqu Wu, and Melanie Tory. Music selection using the PartyVote democratic jukebox. In *Proc. Working Conf. Advanced Visual Interfaces (AVI '08)*, pages 433–436, 2008.
- [124] Ian Stavness, Jennifer Gluck, Leah Vilhan, and Sidney Fels. The MUSICtable: a map-based ubiquitous system for social interaction with a digital music collection. In *Proc. International Conference on Entertainment Computing (ICEC '05)*, pages 291–302, 2005.
- [125] Melanie Tory, Sheryl Staub-French, Barry A. Po, and Fuqu Wu. Physical and digital artifact-mediated coordination in building design. *Comput. Supported Coop. Work*, 17:311–351, August 2008.
- [126] Edward Tufte. *Envisioning information*. Graphics Press, Cheshire, CT, USA, 1990.

- [127] Barbara Tversky, Julie Bauer Morrison, and Mireille Betrancourt. Animation: Can it facilitate? *Intl. J. Hum.-Comput. Stud.*, 57(4):247–262, 2002.
- [128] George Tzanetakis. Musical genre classification of audio signals. In *Proc. IEEE Trans. Speech and Audio Processing*, pages 293–302, 2002.
- [129] George Tzanetakis and Perry Cook. Musical genre classification of audio signals. *IEEE Trans. Speech and Audio Processing*, 10(5):293–302, 2002.
- [130] A. Vande Moere. Form follows data: The symbiosis between design and information visualization. In *Proc. Intl. Conf. Comp.-Aided Architectural Design (CAADfutures '05)*, pages 31–40, 2005.
- [131] A. Vande Moere. Information aesthetics - data visualization and visual design, 2006. Available at <http://infosthetics.com/>.
- [132] Fernanda B. Viégas, Scott Golder, and Judith Donath. Visualizing email content: portraying relationships from conversational histories. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*, pages 979–988, 2006.
- [133] Fernanda B. Viégas and Martin Wattenberg. Timelines: Tag clouds and the case for vernacular visualization. *Interactions*, 15:49–52, July 2008.
- [134] Fernanda B. Viégas, Martin Wattenberg, Matt McKeon, Frank van Ham, and Jesse Kriss. Harry Potter and the meat-filled freezer: A case study of spontaneous usage of visualization tools. In *Proc. Hawaii International Conference on System Sciences (HICSS '08)*, pages 159–169, 2008.
- [135] Fernanda B. Viégas, Martin Wattenberg, Frank van Ham, Jesse Kriss, and Matt McKeon. ManyEyes: A site for visualization at Internet scale. *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 13(6):1121–1128, 2007.
- [136] Amy Vaida, Rebecca E. Grinter, Nicolas Ducheneaut, W. Keith Edwards, and Mark W. Newman. Listening in: practices surrounding iTunes music sharing. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*, pages 191–200, 2005.

- [137] Colin Ware and Robert Bobrow. Supporting visual queries on medium-sized node-link diagrams. *Proc. IEEE Symp. Information Visualization(InfoVis '05)*, 4(1):49–58, 2005.
- [138] Colin Ware and Rusty Bobrow. Motion coding for pattern detection. In *Proc. Symp. Applied Perception in Graphics and Visualization (APGV '06)*, pages 107–110, 2006.
- [139] Martin Wattenberg. Baby names, visualization, and social data analysis. In *Proc. IEEE Symp. Information Visualization (InfoVis '05)*, pages 1–10, 2005.
- [140] Matt Williams and Tamara Munzner. Steerable, progressive multidimensional scaling. In *Proc. IEEE Symp. Information Visualization (InfoVis '04)*, pages 57–64, 2004.
- [141] T.D. Wilson. Human information behaviour. *Informing Science*, 3(2):49–56, 2000.
- [142] Stephen Wolfram. Wolfram alpha computational knowledge engine, 2010. Available at <http://www.wolframalpha.com/>, accessed June 16, 2011.
- [143] Steven Wood. Tag galaxy, 2008. Available at <http://taggalaxy.de/>, accessed November 25, 2008.
- [144] Bill Zeller. Graph your inbox, 2010. Available at <http://graphyourinbox.com>, accessed June 16, 2011.
- [145] Leizhong Zhang, Nan Tu, and Dave Vronay. Info-lotus: a peripheral visualization for email notification. In *Extended Abstracts on Human Factors in Computer Systems (CHI '05)*, pages 1901–1904, 2005.

Appendix A

Additional PartyVote Information

The following files were used for the PartyVote study described in Chapter 4. Song meta-data, song positions, and user log data were excluded from the appendix due to their size and the lack of human readability. These files should be included with the corresponding data disk/hard drive provided to Dr. Melanie Tory. Program code will also be provided in this format.

Documents included below are in the following order:

- A list of PartyVote music selected (both studies)
- A list of PartyVote house party music played
- The study questionnaire
- The consent form (ethics protocol number 08-06-128-c)

PrevVotes_safe_combined.txt

Cofee House Votes

David S	2049	0	2049	2	
Thanh	1491	1	1491	2	
lmartins		1328	2	1328	2
Del	2299	3	2299	2	
Sean	2525	4	2525	2	
Bob	439	5	439	2	
Oana	2698	6	2698	2	
Anguel Tzavkov		2753	7	0	3
Deer Li	1556	8	1556	2	
dandan	756	9	756	2	
ting	1626	10	178	0	
Kedar	598	11	598	2	
Lindsey	1957	12	1957	2	
David	936	13	936	2	
Brittany		2292	14	2292	2
dale	966	15	966	2	
Lior	2734	16	2734	2	

House Party Votes

Melanie	2266	0	2266	2	
Davids	2049	1	2049	2	
sean	1067	2	1067	2	
fidler	914	3	914	2	
terryf	614	4	614	2	
Melanie2		2457	5	2457	2
tjf	936	6	936	2	
c	1826	7	1826	2	
sean2	1694	8	1694	2	
sean3	338	9	18	0	
tjfi	2812	10	2812	2	
d	1983	11	1983	2	
tjfid	326	12	326	2	
viv1	1126	13	1126	2	
nineke	574	14	574	2	

SongInfo	Weight	User1	User1	Weight	User2	User2	Weight	User3	User3	Weight
Rockin' The Suburbs Artist: Ben Folds Album: Rockin' The Suburbs Genre: Alternative & Punk	1.05	Melanie	1.0	0.0	0.0					
Cool It Down Artist: The Velvet Underground Album: Loaded (Disc 1) Genre: Alternative & Punk	0.54	Melanie	0.47	DavidS	0.03	0.0				
Gone Artist: Ben Folds Album: Rockin' The Suburbs Genre: Alternative & Punk	0.32	Melanie	0.25	DavidS	0.03	0.0				
Take Me Out Artist: Franz Ferdinand Album: Franz Ferdinand Genre: Alternative & Punk	0.6	Melanie	0.03	DavidS	0.53	0.0				
Caribou Artist: Pixies Album: Surfer Rosa & Come On Pilgrim Genre: Alternative & Punk	0.6	Melanie	0.03	DavidS	0.53	0.0				
Sweet Thing (Reprise) Artist: David Bowie Album: Diamond Dogs Genre: Alternative & Punk	0.6	Melanie	0.03	DavidS	0.53	0.0				
Enid Artist: Barenaked Ladies Album: Gordon Genre: Alternative & Punk	0.59	Melanie	0.52	DavidS	0.03	0.0				
This Is The Sound Artist: The Juliana Hatfield Three Album: Become What You Are Genre: Alternative & Punk	0.59	Melanie	0.52	DavidS	0.03	0.0				
Karma Police Artist: Radiohead Album: OK Computer Genre: Alternative & Punk	0.6	Melanie	0.03	DavidS	0.53	0.0				
Rock & Roll (Alternate Mix) Artist: The Velvet Underground Album: Loaded (Disc 1) Genre: Alternative & Punk	0.15	Melanie	0.06	DavidS	0.05	0.0				
Only Dreaming Artist: K's Choice Album: Paradise In Me Genre: Alternative & Punk	0.17	Melanie	0.08	DavidS	0.05	sean	0.02			
Not Even Jail Artist: Interpol Album: Antics Genre: Alternative & Punk	0.61	Melanie	0.03	DavidS	0.53	sean	0.01			
Crucify (Remix) Artist: Tori Amos Album: Crucify [EP] Genre: Alternative & Punk	1.08	Melanie	0.02	DavidS	0.02	fidler	1.0			
King Of Carrot Flowers Part 1 Artist: Neutral Milk Hotel Album: In The Aeroplane Over The Sea Genre: Alternative & Punk	0.24	Melanie	0.04	DavidS	0.14	fidler	0.02			
If Trouble Was Money Artist: Albert Collins Album: Festival International de Jazz de Montreal - 20 Years Of Music - The Blues and Beyond (Disc 1) Genre: Jazz	0.73	Melanie	0.02	terryf	0.16	sean	0.54			
I Only Eat Candy Artist: Nerf Herder Album: Nerf Herder Genre: Alternative & Punk	0.68	terryf	0.04	Melanie2	0.56	sean	0.04			
This Is Just A Modern Rock Song Artist: Belle & Sebastian Album: This Is Just A Modern Rock Song Genre: Alternative & Punk	0.68	terryf	0.04	Melanie2	0.56	sean	0.04			
Deeper Into Movies Artist: Yo La Tengo Album: MATADOR AT FIFTEEN (Disc 2) Genre: Alternative & Punk	0.67	Melanie2	0.56	fidler	0.03	terryf	0.04			
5/4 Artist: Gorillaz Album: Gorillaz [UK] Genre: Alternative & Punk	0.69	sean2	0.06	DavidS	0.03	c	0.56			
High 5 (Rock The Catskills) Artist: Beck Album: Odelay Genre: Alternative & Punk	0.69	sean2	0.07	DavidS	0.03	c	0.55			
Icicle Melts Artist: The Cranberries Album: No Need To Argue Genre: Pop	0.1	Melanie	0.01	terryf	0.02	Melanie2	0.01			
Apparitions [Edit] Artist: Matthew Good Band Album: Big Shiny Tunes 3 Genre: Alternative & Punk	1.11	Melanie	0.03	terryf	1.0	Melanie2	0.04			
Not The Same Artist: Dinosaur Jr Album: Where You Been Genre: Alternative & Punk	0.17	Melanie	0.05	DavidS	0.03	sean2	0.03			
A Dame With A Rod Artist: The Juliana Hatfield Three Album: Become What You Are Genre: Alternative & Punk	0.28	Melanie	0.16	DavidS	0.04	terryf	0.03			
Insensatez Artist: Antônio Carlos Jobim Album: Lost Highway Genre: Soundtrack	1.13	sean2	1.0	DavidS	0.02	c	0.07			
Bullet With Butterfly Wings Artist: Smashing Pumpkins Album: Mellon Collie And The Infinite Sadness - Dawn To Dusk [Disc 1] Genre: Alternative & Punk	1.13	sean2	0.07	DavidS	0.03	c	1.0			

Airbag | Artist: Radiohead | Album: OK Computer | Genre: Alternative & Punk | 1.11 Melanie 0.03 DavidS 1.0 c 0.03

Get Your Hands Off My Woman | Artist: Ben Folds | Album: Supersunnyspeedgraphic: The LP | Genre: Alternative & Punk | 0.32 Melanie 0.2 DavidS 0.03 terryf 0.04

Army | Artist: Ben Folds Five | Album: The Unauthorized Biography Of Reinhold Messner | Genre: Alternative & Punk | 0.66 Melanie 0.55 DavidS 0.03 terryf 0.03

Soul In Lust | Artist: The High Dials | Album: War Of The Wakening Phantoms | Genre: Alternative & Punk | 0.66 Melanie 0.55 DavidS 0.03 terryf 0.03

Un Canadien Errant | Artist: Emile Campagne | Album: Emile Campagne | Genre: Folk | 1.07 terryf 0.01 sean 1.0 Melanie2 0.02

Dragula (The Hod Rod Herman Remix) | Artist: Rob Zombie | Album: Big Shiny Tunes 3 | Genre: Alternative & Punk | 1.11 Melanie2 1.0 fidler 0.03 terryf 0.04

Mary | Artist: Sarah McLachlan | Album: Fumbling Towards Ecstasy | Genre: Rock | 0.91 d 0.62 tjfi 0.02 tfj 0.3

One | Artist: U2 | Album: Achtung Baby | Genre: Rock | 0.93 d 0.63 tjfi 0.02 tfj 0.31

Five Years | Artist: David Bowie | Album: Ziggy Stardust | Genre: Rock | 0.92 d 0.3 tfj 0.63 tjfi 0.02

Money | Artist: Pink Floyd | Album: Dark Side Of The Moon | Genre: Rock | 1.26 d 0.27 tfj 1.0 tjfi 0.02

Leash | Artist: Pearl Jam | Album: Vs. | Genre: Rock | 0.88 d 0.24 tfj 0.63 tjfi 0.02

Nightclubbing | Artist: Iggy Pop | Album: Trainspotting | Genre: Soundtrack | 0.87 d 0.23 tfj 0.63 tjfi 0.02

Something In The Way | Artist: Nirvana | Album: Nevermind | Genre: Rock | 0.85 d 0.2 tfj 0.63 tjfi 0.02

Indifference | Artist: Pearl Jam | Album: Vs. | Genre: Rock | 0.74 d 0.21 tfj 0.51 tjfi 0.02

The Lines You Amend | Artist: Sloan | Album: 4 Nights At The Palais Royale (Disc 1) | Genre: Alternative & Punk | 0.72 Melanie 0.07 Melanie2 0.57 terryf 0.03

Everything You've Done Wrong | Artist: Sloan | Album: 4 Nights At The Palais Royale (Disc 1) | Genre: Alternative & Punk | 0.72 Melanie 0.07 Melanie2 0.57 terryf 0.04

Johnny Go | Artist: Jean Leloup | Album: Le Dôme | Genre: World | 0.13 d 0.02 tfj 0.02 tjfi 0.04

The Piccolo Snare | Artist: Super Furry Animals | Album: Phantom Power | Genre: Alternative & Punk | 0.72 Melanie 0.07 DavidS 0.03 Melanie2 0.57

Hey Ya! | Artist: Andre 3000 | Album: The Love Below | Genre: Hip Hop/Rap | 1.09 d 0.02 tfj 0.03 tjfi 1.0

On A Plain | Artist: Nirvana | Album: MTV Unplugged in New York | Genre: Rock | 0.14 Melanie 0.03 d 0.02 tfj 0.03

All The Wine | Artist: The National | Album: Alligator | Genre: Indie | 0.17 sean2 0.05 d 0.03 tfj 0.03

Desperado | Artist: Johnny Cash | Album: American IV: The Man Comes Around | Genre: Country | 1.24 sean3 0.03 viv1 0.21 tjfid 1.0

Bird On A Wire | Artist: Johnny Cash | Album: American Recordings | Genre: Country | 0.16 sean3 0.03 tjfid 0.04 viv1 0.04

Rolex By Solex | Artist: Solex | Album: Solex Vs. The Hitmeister | Genre: Electronica/Dance | 0.84 tjfid 0.55 terryf 0.03 viv1 0.29

Fido, Your Leash Is Too Long | Artist: The Magnetic Fields | Album: 69 Love Songs, Vol. 1 | Genre: Indie | 0.15 tjfid 0.03 sean 0.03 viv1 0.04

Ride Into The Sun (Demo) | Artist: The Velvet Underground | Album: Loaded (Disc 1) | Genre: Alternative & Punk | 0.19 Melanie 0.06 DavidS 0.04 terryf 0.03

Tether | Artist: Sarah Harmer | Album: All Of Our Names | Genre: Folk | 0.84 tjfid 0.21 terryf 0.02 viv1 0.61

Sleepwalking Zombies | Artist: Leah Abramson | Album: 8 Songs From The Attic | Genre: Folk | 0.84 tjfid 0.2 terryf 0.03 viv1 0.61

Andromeda | Artist: Leah Abramson | Album: 8 Songs From The Attic | Genre: Folk | 0.81 viv1 0.16 terryf 0.03 tjfid 0.62

Mountain On Mountain | Artist: Sparrow | Album: Sparrow | Genre: Pop | 0.69 viv1 0.15 terryf 0.03 tjfid 0.5

Claudia Acuna | Artist: Gracias A La Vida | Album: Festival International De Jazz De Montreal 21e Edition | Genre: Jazz | 1.23 tjfid 0.21 terryf 0.02 viv1 1.0

To This Day | Artist: K's Choice | Album: Paradise In Me | Genre: Alternative & Punk | 0.19 Melanie 0.06 DavidS 0.06 c 0.03

All Mine | Artist: Portishead | Album: Roseland NYC [Live] | Genre: Electronica/Dance | 0.12 tjfi 0.02 viv1 0.02 tjfid 0.02

The Last Polka | Artist: Ben Folds/Anna Goodman | Album: Ben Folds Five | Genre: Alternative & Punk | 0.68
 Melanie 0.57 DavidS 0.03 terryf 0.03

My Hero | Artist: Foo Fighters | Album: The Colour And The Shape | Genre: Rock | 0.14 Melanie 0.03 sean2 0.03 tjf
 0.03

Territorial Pissings | Artist: Nirvana | Album: Nevermind | Genre: Rock | 1.27 d 1.0 tjfi 0.02 tfj 0.27

Have A Cigar | Artist: Pink Floyd | Album: Wish You Were Here | Genre: Rock | 0.27 nineke 0.06 d 0.09 tfj 0.1

High On Cruel | Artist: The Corn Sisters | Album: The Other Women | Genre: Country | 0.87 tjfid 0.55 nineke 0.04
 viv1 0.32

Music Is My Boyfriend (Live On CBC Radio 3) | Artist: The Hidden Cameras | Album: CBC Radio 3 Sessions
 Volume 1 | Genre: Easy Listening | 0.15 Melanie 0.04 tjf
 0.02 sean2 0.03

The Rising Tide | Artist: Sunny Day Real Estate | Album: Rising Tide | Genre: Alternative & Punk | 0.24 Melanie
 0.04 DavidS 0.11 Melanie2 0.04

Until Morale Improves, the Beatings Will Continue | Artist: Murder By Death | Album: Who Will Survive, and
 What Will Be Left of Them? | Genre: Alternative & Punk | 0.31
 Melanie 0.09 Melanie2 0.14 terryf 0.04

Sentimental Guy | Artist: Ben Folds | Album: Songs For Silverman | Genre: Alternative & Punk | 0.72 Melanie 0.06
 Melanie2 0.57 terryf 0.04

Judy Is A Dick Slap | Artist: Belle & Sebastian | Album: Legal Man | Genre: Alternative & Punk | 0.72 Melanie 0.06
 Melanie2 0.57 terryf 0.04

Swan Song | Artist: Juliana Hatfield | Album: Bed | Genre: Alternative & Punk | 0.2 Melanie 0.06 DavidS 0.05
 Melanie2 0.04

It's All Gonna Break | Artist: Broken Social Scene | Album: Broken Social Scene | Genre: Alternative & Punk | 0.71
 Melanie 0.57 Melanie2 0.06 terryf 0.03

The Old Apartment | Artist: Barenaked Ladies | Album: Rock Spectacle (Live) | Genre: Alternative & Punk | 0.71
 Melanie 0.57 Melanie2 0.06 terryf 0.03

Jivin' With Dal | Artist: Michael Kaeshammer | Album: Best Of The Beaches International Jazz Festival | Genre:
 Jazz | 1.13 nineke 1.0 tjfid 0.04 viv1 0.04

Jeez Louise | Artist: Zumpano | Album: Look What The Rookie Did | Genre: Alternative & Punk | 0.19 Melanie2
 0.03 terryf 0.03 fidler 0.07

Gracie | Artist: Ben Folds | Album: Songs For Silverman | Genre: Alternative & Punk | 1.15 Melanie 0.07 Melanie2
 1.0 terryf 0.03

Thrice All American | Artist: Neko Case | Album: Furnace Room Lullaby | Genre: Country | 0.18 tjfid 0.05 viv1
 0.06 nineke 0.02

Questionnaire

1. Did you interact with or examine the PartyVote system during the party?

Yes No

Why or why not?

If you didn't interact or look at the PartyVote system, go to question 11.

2. How many times did you vote using the PartyVote system?

a) 0 times b) 1 time c) 2 times d) more than 2 times

Why did you vote or not vote?

If you voted 0 times, jump to question 5

For the following statements, please circle a number indicating how much you agree or disagree with the sentence.

<i>Question</i>	<i>Strongly Disagree</i>	<i>Disagree</i>	<i>Neutral</i>	<i>Agree</i>	<i>Strongly Agree</i>	<i>Not Applicable</i>
3. I felt my vote had no influence on what songs were played	1	2	3	4	5	N/A
4. The ability to vote for music helped me feel empowered	1	2	3	4	5	N/A
⇒ 5. I clearly understood how votes affected which songs were played	1	2	3	4	5	N/A
6. I understood how songs/planets were organized on the space themed visualization	1	2	3	4	5	N/A
7. The space/planet themed visualization was childish or inappropriate	1	2	3	4	5	N/A
8. There were songs I wanted to hear in the music library	1	2	3	4	5	N/A
9. Finding and selecting music was difficult	1	2	3	4	5	N/A
10. Planet size clearly illustrated the probability of a song being played	1	2	3	4	5	N/A
⇒ 11. Today's music was diverse and varied	1	2	3	4	5	N/A
12. I enjoyed most of the music played today	1	2	3	4	5	N/A
13. I recognized most of the music played today	1	2	3	4	5	N/A
14. I normally take an active role in selecting music during parties and gatherings	1	2	3	4	5	N/A

What percentage of the music did you enjoy?

What did you like the most about the PartyVote system?

What did you like the least about the PartyVote system?

Additional Comments:



Department of Computer Science

Participant Consent Form

Human Factors in Visualization

You are being invited to participate in a study entitled *Mediating Collaborative Music Selections Using Voting* that is being conducted by Melanie Tory, David Sprague, and Fuqu Wu.

Melanie Tory is a faculty member in the department of Computer Science at the University of Victoria and you may contact her if you have further questions by email at mtory@cs.uvic.ca or by phone at (250) 472-5798.

This research is being funded by NSERC.

The purpose of this research project is to investigate how people use information, and how different visual representations of data affect how people perform tasks such as data analysis and decision making. Research of this type is important because it allows us to design better data displays to allow more effective, efficient, and enjoyable analysis of data in a variety of applications.

If you agree to voluntarily participate in this research, your participation will include

- Completing computer-based tasks.
- Filling out a questionnaire about the computer-based tasks and tools you experienced.
- Being video and audio-taped.
- Being watched by live observers.

The research session is expected to take approximately 1 hour and will take place at the University of Victoria Engineering/Computer Science Building in room 668.

There are no known or anticipated risks to you by participating in this research.

Your participation in this research must be completely voluntary. If you do decide to participate, you may withdraw at any time without any consequences or any explanation. If you do withdraw, we will ask whether we may use your data for data analysis. If you decline, your data will be destroyed.

Your confidentiality and the confidentiality of the data will be protected by identifying data only with a participant number rather than your name, password-protecting computer files, and storing video and audio tapes in a locked office. Because this is a group study, confidentiality cannot be fully guaranteed since other participants in your group may know your identity.

It is anticipated that the results of this study will be shared with others in the following ways:

- Published articles
- Conference presentations
- Video publications
- Theses
- Internet project descriptions

Data from this study will be disposed of within 5 years. Electronic data will be erased, paper copies will be shredded, and video/audio tapes will be recorded over or physically destroyed.

In addition to being able to contact the researcher at the above phone numbers, you may verify the ethical approval of this study, or raise any concerns you might have, by contacting the Associate Vice-President, Research at the University of Victoria (250-472-4545).

Your signature below indicates that you understand the above conditions of participation in this study and that you have had the opportunity to have your questions answered by the researchers.

Name of Participant

Signature

Date

A copy of this consent will be left with you, and a copy will be taken by the researcher.

Appendix B

Motion Study Additional Information

The following files were used for the motion study described in Chapter 5. Song meta-data, song positions, and user log data were excluded from the appendix due to their size and the lack of human readability. These files should be included with the corresponding data disk / hard drive provided to Dr. Melanie Tory. Program code will also be provided in this format.

Appendix documents below consist of the following:

- Study instructions and guiding questions
- A system functionality reminder document
- Experiment recruitment text
- The consent form (ethics protocol number 08-06-128-c)

Experimental Dialog:

Thank you for participating in this experiment. Before I ask you to sign a consent form, I want to give you some brief information about what you will be doing during the next two hours. You will be asked to look through a music library that you are unfamiliar with and create playlists for certain activities or occasions. You will also be asked to search the library for songs with particular characteristics like finding all songs with the letter y in the title. Regular breaks will be provided so please don't stop in the middle of a task. After a brief training you will see four different versions of the system. Finally, I will interview you to get your opinions and thoughts on the system.

Please sign the following consent form if you are willing to participate. Remember that your participation is voluntary and you can stop at any time. Do you have any questions?

Here is a paper printout of how one would add the song "Come Together" to the music playlist and play it. Please go through the instructions and ask any questions you may have. Let's go through some common activities you will be performing:

Search: *Type words into the search text box and press the search button. This will show all words with that text in the title, artist, or album fields.*

Music Browsing: *The four radio buttons at the top right allow you to search the music collection by artist name, album, genre, or song title. Notice that by clicking items in the left list, songs associated with that item appear in the right list.*

Panning & Zooming: *If you want to see song icons (the circles in the center view) in more detail you can zoom into the visualization using the mouse wheel. The zoom is centered on where your mouse is positioned. When zoomed in, you can pan the view by clicking the white background using the left mouse button and dragging the mouse while the button is down. This is similar to Google maps if you've used that.*

Music Icons & Beat Information: *Notice the song icons. Each circle displays the song title, artist, and album information. The icon colour represents the song's genre. Song beat strength and tempo (or beats per minute) information will vary between conditions. In the current example, the number of bumps around the circle indicate the number of beats per minute and the height of the bumps indicate beat strength. You would expect dance music to have a lot of high bumps or high beat strength and many beats per minute. Slow songs would be expected to have very few shallow bumps. Please note that we used automatically extracted song metadata for this information so please try and ignore any incorrect information if you disagree with the data being shown.*

Popup Song Information: *Hovering the mouse over a given song icon causes a pop-up box of song information to be displayed.*

Selecting Songs: *Left clicking a song icon with the mouse selects an icon.*

Adding songs (playlist or choosing): Selected icons can then be added to the song playlist or chosen for the search task by clicking the “add song” button. Alternately, double clicking the icon will do the same thing.

Playlist: Full songs are only played if they are on the playlist. The current song on the playlist is indicated either in orange (playing) or dark orange/brown (next to be played). Double clicking an item in the playlist makes it the new current song.

Moving songs in the playlist: The song playlist on the far left of the screen permits drag and drop functionality. Left click on a song and then drag the mouse to a new position while holding down the button. Notice that the song being dragged is outlined in gray and the destination position is indicated by a gray line. To drop the song in the new position, release the mouse button when at the right location.

Sound control: Pressing the play button on the bottom right of the screen plays the next song to play. Pressing the same button (now showing a pause symbol) pauses it. Stop, skip forward, and skip back functionality are also provided.

Sampling Songs: However, if you wish to quickly see what a song sounds like, you can select the music icon and click the “play sample” button to play a 30 second sample. Playing a sample will temporarily pause any currently playing song.

Glyph Encoding: You will be doing each experimental task four times and the beat information about each song will change. The number of beats per minute (or tempo) of the song and the song’s beat strength will be shown in one of four ways (cycle through encoding types using the “n” key).

Static Condition: In the current example, the number of bumps around the circle indicate the number of beats per minute and the height of the bumps indicate beat strength.

Animation Only Condition: In the current task, the frequency of the line’s movement indicates the number of beats per minute and the height of the animation indicates the beat strength.

Redundant Condition: In the current task, the number of beats around the circle and the frequency of the line’s movement both indicate the number of beats per minute. The height of the animation indicates the beat strength.

Extraneous Animation Condition: In the current example, the number of bumps around the circle indicate the number of beats per minute and the height of the bumps indicate beat strength. The line around each circle will animate in time with the currently playing song.

For practical purposes you should expect quiet and slow music like Simon and Garfunkel songs to have few bumps/move slowly and not move much. You should expect electronic dance music to be the opposite with many bumps on the line and/or big and fast movements.

Now you will be given five minute to explore the music library so you get a better feel for the songs available. Feel free to use any system functions you want. Do you have any questions?

Glyph Comparison Task: *For this next set of tasks you will be asked to compare two song glyphs at a time and identify the one with the larger value. Half of the questions will ask you to find the song with the most beats per minute and the other half will ask you to find songs with the strongest beat. You will notice that the colour of the glyphs will change depending on the question to hopefully make it easier for you to keep track of the question. Between each question you will see the next question being asked. You will need to press a key to begin the next trial (demonstrate). To select the left glyph, press the left arrow and to select the right glyph, press the right arrow. Please try and get less than 10% of the questions wrong and make your selection as quickly and accurately as possible. The first four trials will be practice trials. If you get a trial wrong, don't be too concerned and continue to try your best. At the end of each session, you will be informed of your score. Now, try to answer a few questions for practice. You will get practice with each condition and then we will begin the actual trials. Please try your best.*

Playlist Generation Task: *For this task please read the following instructions and click ok whenever you are ready to begin. Please remember to just try your best. There is no right answer to this task, just the music that you think it appropriate.*

Two-Dimensional Search Task: *For this task, please find all the items in the music library based on the criteria written. To choose a song, double click the song icon. If you add a song by accident, just continue. Please work as fast as you can while still being accurate.*

Break: *You can now take up to a 5 minute break to stretch, relax or ask any questions you may have. Please let me know when you are ready to continue.*

Interview: *I would now like to interview you to see what your thoughts were on the various systems. Please feel as much detail as you like. Remember, we are looking for your option. There is no right answer. We would also like to audio tape you to make sure we transcribe your comments accurately. This recording will not be used for anything other than transcription purposes. Are you ok with me taping you?*

Experiment End: *The experiment is now finished. Please sign the payment receipt form. We can provide you with a debriefing as to the purpose of this experiment if you want. Do you have any final questions or do you want to hear the purpose of the experiment?*

Debriefing: *This experiment is designed to examine how animation can be used to motivate people to use a system and to encode data. You experienced four conditions: a static or non-moving condition where beat information was encoded using a oscillating line, an animation only condition where beat strength and beats per second were represented by the line's animation, a redundant condition where data was represented both using the line and animation, and an extraneous animation condition where beat data was represented by the line but songs animated to the currently playing music (not*

representing any relevant data). We wanted to see what your preferences were, if the animations distracted you (during the search task), and if the animation helped you find music to add on your playlist. We also wanted to explore the trade-offs between the distractions caused by animations and aesthetic judgments about each system.

Playlist Creation Tasks:

1. Imagine you are at a friend's house and you need music for a romantic evening. You don't have time to go home for your own music so you have to use his or her music collection. In a ten minute period, create a playlist of music that YOU would like to listen to during a romantic evening.
2. Imagine you are at a friend's house and you need music for a party you are throwing there that night. You are expecting some dancing and more than 20 people attending. In a ten minute period, create a playlist of music that YOU would like to listen to during a **party**.
3. Imagine you are at a friend's house and you are about to run off to study on-campus for the day. If you don't study to music, pretend that the music is being used to reduce distractions in a noisy room and choose the music accordingly. In a ten minute period, create a playlist of music that YOU would like to listen to while **studying**.
4. Imagine you are at a friend's house and you want to make an exercise playlist for use while working out. In a ten minute period, create a music playlist that YOU would like to listen to while exercising.

Library Search Tasks (randomly assigned without replacement to the 4 conditions):

In a three minute period, find all the songs in the library from the album "Recovering the Satellites" by Counting Crows. You can select a song by double clicking or selecting the song and pressing the "Add Song" button.

In a three minute period, find all the songs in the library from the album "Van Lear Rose" by Loretta Lynn. You can select a song by double clicking or selecting the song and pressing the "Add Song" button.

In a three minute period, find all the songs in the library from the album "Decade (disk 1)" by Neil Young. You can select a song by double clicking or selecting the song and pressing the "Add Song" button.

In a three minute period, find all the songs in the library from the album "My Aim is True" by Elvis Costello. You can select a song by double clicking or selecting the song and pressing the "Add Song" button.

Interview Questions for Animation Experiment

Reminder about 4 conditions: no anim, anim only, redundant anim, and extraneous anim

WE WOULD GREATLY APPRECIATE AS MUCH DETAIL AND REASONING FOR YOUR ANSWERS AS YOU CAN GIVE. THE INTERFACE, MUSIC LIBRARY, AND MUSIC SELECTION METHOD ITSELF ARE NOT BEING STUDIED SO FEEL FREE TO CRITIQUE ANY ASPECT OF THIS STUDY WITHOUT UPSETTING US.

BACKGROUND:

What is the size of your personal music collection (# of songs)? Is this collection digital?

Do you own an iPod or other portable mp3 player?

How often do you create a music playlists or mixed CDs? (# of times a year for example) For what occasions do you create these playlists?

How frequently do you explore another person's music library (online using lastFM or glancing through a friend's CD collection for example)?

Do you have any music education or experience? For example, can you play any instruments?

Playlist Generation

What do you look for in terms of romantic music? How did you search for these criteria during the playlist creation task?

What did you look for in terms of dancing or party music? How did you search for these criteria during the playlist creation task?

What did you look for in terms of music to study to? How did you search for these criteria during the playlist creation task?

What did you look for in terms of music to exercise to? How did you search for these criteria during the playlist creation task? What type of exercise did you imagine the music would be used for?

Interface Opinions:

Were you familiar with the music in the given library? Can you give a rough percentage for the songs you knew? What percentage of the songs you knew did you like? What percentage did you dislike?

How satisfied were you with each of the playlists you created? Please explain your answer.

Did you use any of the beat per minute or beat strength information to perform either the playlist creation task or the library search task? How useful was this information?

Which condition did you like performing the playlist creation task with the most? Why? (feel free to rank conditions)

What condition did you like performing the playlist creation task with the LEAST?

Which condition did you like performing the playlist search task with the most?

The least?

What did you think of each of the conditions for the glyph comparison task?

Were you able to ignore the movement and animation in the search task?

Which condition did you find to be the most pleasing condition to use (in terms of aesthetics)?

Generally, what are your thoughts on using animations or movement for the given tasks? Did you ever find it useful or beneficial?

Which condition did you feel like you were the most productive in? For the playlist task? For the library search task? If you felt you were equally productive using all conditions, please indicate this.

For the glyph comparison task, which conditions did you think was the easiest to do the task using? Why?

Were there any times or questions that were particularly difficult during the glyph comparison task?

Was there anything about the tasks you were asked to perform that you found annoying or frustrating?

Did you find any of the conditions distracting? Why?

Did you notice any errors or problems with the beat per minute or beat strength information?

You may have noticed that songs were not laid out on the computer screen in a grid pattern. How do you think songs were positioned on the screen? How do you think songs should have been positioned on the screen?

If you were able to explore your own music library using one of the four experimental conditions, would you? Why or why not?

Which condition would you use at home and why?

Music Visualization Instructions


The following instructions on how to add “Come Together” by the Beatles to top of the playlist provide an example of how to use the system.

1. Music Browsing: The four radio buttons at the top right allow you to search the music collection by artist name, album, genre, or song title. Clicking on artist and selecting “The Beatles” makes all the library’s Beatles songs show appear in the right column. Selecting “Come Together” make selects the appropriate icon in the main window.
2. Search: Typing in “Come Together” in the search box, clicking the “Search” button, and selecting the song in the list also selects the desired song. Searching for “The Beatles” will make all Beatles songs appear in the left list.



3. Selecting songs: Song icons can also be selected by clicking them with the left mouse button.

4. Panning & Zooming: To take a closer look at the selected song, turn the mouse wheel towards you. To zoom back out, rotate the wheel away from you. You can move the view by clicking and dragging.



5. Music Icons & Beat Information: Each song icon displays the song title, artist, and album information. “Come Together” is in blue indicating it is a pop song (see the legend at the bottom of the main window for genre colours).
6. Popup Song Information: Hovering the mouse over a song displays pop-up song information.
7. Playing a Sample: After selecting “Her Majesty”, you can play a short sample of the song by clicking the “Play Song Sample” button or clicking the song icon and pressing the space bar.

8. Adding a Song: Selected songs are chosen by clicking the “Add Song” button or by double clicking the song icon.
- Playlist creation task: Chosen songs can be added to the bottom of the playlist. (image)
 - Search task: Chosen songs are indicated by a square line around the icon.
9. Moving Songs in the playlist (playlist creation task only): Drag and drop functionality is permitted. The song being dragged is outlined in gray. A gray line indicates the position the song will be inserted.
10. Playing Music: Provided music is in the playlist, songs can be played by clicking the play button. The next song to be played is indicated by a dark orange colour. The song currently playing has an orange background. You can select the current song by double clicking its playlist icon or by using the previous and next buttons.

Participant Recruitment Documents

Subject: Needed: Participants in music library visualization experiment (grads) / Needed: Participants in music library exploration experiment (Craig's List)

Craig's List:

From October 6-17 and October 25-31, we will be conducting an experiment looking into how animation can be used to explore digital music libraries. As such, we are looking for participants under 35 years old and preferably with some interest in music. Participants will come into the University of Victoria computer science department and use a new system to search a digital music library and create playlists. Afterwards you will be interviewed to get your thoughts and opinions. The experiment is expected to last two hours and you will be paid \$20 for your time.

If you are interested in helping us, please contact David Sprague at dsprague@cs.uvic.ca or phone (250) 472-5749 (email is preferred)


Grads:

From October 6-17, the VisID research group will be conducting an experiment looking into how animation can be used to explore digital music libraries. As such, we are looking for participants (preferably with some interest in music). Participants will use a new visualization to search a digital music library and create playlists. Afterwards you will be interviewed to get your thoughts and opinions. The experiment is expected to last two hours and you will be paid \$20 for your time.

If you are interested in helping us, please contact David Sprague at dsprague@cs.uvic.ca or phone (250) 472-5749 (email is preferred)

Interested in Digital Music?

The VisID research group at U.Vic. is exploring how animations can be used to explore music libraries and we are looking for participants



If you are < 35 and can spare 2 hours, we want your help!

You will be financially compensated for your time.

If you are interested in helping us, please contact David Sprague at dsprague@cs.uvic.ca or phone (250)472-5749 (email is preferred)



UVic

Department of Computer Sciences
University of Victoria
Engineering/Computer Science Building
(ECS), Room 504
PO Box 3055, STN CSC
Victoria, BC
Canada V8W 3P6

Participant Consent Form¹⁸⁰

Human Factors in Visualization

You are being invited to participate in a study entitled *The Effects of Beat Encoding Methods on Music Library Use, Preference and Performance* that is being conducted by Melanie Tory and David Sprague.

Melanie Tory is a faculty member in the department of Computer Science at the University of Victoria and you may contact her if you have further questions by email at mtory@cs.uvic.ca or by phone at (250) 472-5798.

This research is being funded by NSERC.

The purpose of this research project is to investigate how people use information, and how different visual representations of data affect how people perform tasks such as data analysis and decision making. Research of this type is important because it allows us to design better data displays to allow more effective, efficient, and enjoyable analysis of data in a variety of applications.

If you agree to voluntarily participate in this research, your participation will include answering a background questionnaire that asks about your experience with computer technology and data analysis applications, as well as personal characteristics such as age and gender.

- Completing computer-based tasks.
- Participating in a verbal interview.
- Being video and audio-taped.
- Being watched by live observers.

The research session is expected to take approximately 120 minutes and will take place at The University of Victoria, Engineering and Computer Sciences Building Room 555.

There are no known or anticipated risks to you by participating in this research.

Your participation in this research must be completely voluntary. If you do decide to participate, you may withdraw at any time without any consequences or any explanation. If you do withdraw, we will ask whether we may use your data for data analysis. If you decline, your data will be destroyed.

Your confidentiality and the confidentiality of the data will be protected by identifying data only with a participant number rather than your name, password-protecting computer files, and storing video and audio tapes in a locked office.

It is anticipated that the results of this study will be shared with others in the following ways:

- Published articles
- Conference presentations
- Video publications

- Theses
- Internet project descriptions

Data from this study will be disposed of within 5 years. Electronic data will be erased, paper copies will be shredded, and video/audio tapes will be recorded over or physically destroyed.

In addition to being able to contact the researcher at the above phone numbers, you may verify the ethical approval of this study, or raise any concerns you might have, by contacting the Associate Vice-President, Research at the University of Victoria (250-472-4545).

Your signature below indicates that you understand the above conditions of participation in this study and that you have had the opportunity to have your questions answered by the researchers.

Name of Participant

Signature

Date

A copy of this consent will be left with you, and a copy will be taken by the researcher.

Appendix C

Artifact Study Additional Information

The first five documents listed below were included with each cultural probe or artifact kit given to participants. The online survey was slightly different than the paper journal entries as it spanned more than one screen and required an explicit indication of consent. The list of journal articles submitted does not include all details due to the file size and the lack of human readability.

Appendix documents consist of the following (in order):

- The participant consent form (ethics protocol number 09-417)
- Instructions on the study and what is an artifact
- A Paper journal entry
- A matrix of pocket reference cards (front and back)
- A study “Withdrawal with continued data use” form
- A list of artifact journal entries reported (inappropriate entries are in gray)
- A grid of artifact photos submitted

Participant Consent Form

Visual Encoding of Data Away from Work: Exploring Casual Data Analysis and Key Contextual Factors

You are being invited to participate in a study entitled Visual Encoding of Data Away from Work: Exploring Casual Data Analysis and Key Contextual Factors that is being conducted by David Sprague and Melanie Tory.

As a graduate student, I (David Sprague) am required to conduct research as part of the requirements for a degree in Computer Science. It is being conducted under the supervision of Dr. Melanie Tory. You may contact me if you have any further questions by email at dsprague@cs.uvic.ca or by phone at (250) 472-5749. You may contact my supervisor by email at mtory@cs.uvic.ca or by phone at (250) 472-5798.

This research is being funded by NSERC.

The purpose of this research project is to investigate how people use information, and how different visual representations of data affect how people perform tasks such as data analysis and decision making. Research of this type is important because it allows us to design better data displays to allow more effective, efficient, and enjoyable analysis of data in a variety of applications.

If you agree to voluntarily participate in this research, your participation will include:

- Recording a daily journal about your experiences with visual representations of data you have encountered in your free (non-work) time. The journal can be filled out online or on paper. We expect this to take less than ½ hour each day.
- Participating in a verbal interview, expected to take one hour.
- Being video and audio-taped during the interview.

The research session is expected to take approximately 2 weeks with less than a half hour spent writing journal entries each day. A 1 hour interview at the end of the two weeks will take place in room 555 of the Engineering and Computer science building at the University of Victoria. Interview times will be held at your convenience.

There are no known or anticipated risks to you by participating in this research.

Participation in this study may cause some minor inconvenience to you, including interruptions to your daily routine in order to note artifacts and write journal entries. This disruption is expected to be minimal.

As a way to compensate you for any inconvenience related to your participation, you will be offered \$60. For each journal entry you submit, your name will be entered into a series of three \$10 cash draws. To qualify for these draws and financial compensation, you must complete at least 5 journal entries (including “no artifact found” entries) and attend the final interview. We assume that participants who filled in fewer than 5 journal entries have dropped out of the study. It is important for you to know that it is unethical to provide undue compensation or inducements to research participants and, if you agree to be a participant in this study, this form of compensation to you must not be coercive. If you would not otherwise choose to participate if the compensation was not offered, then you should decline. If you attend the interview and wish to withdraw from the study, you will be fully compensated. If you withdraw before the interview, compensation will not be offered.

Your participation in this research must be completely voluntary. If you do decide to participate, you may withdraw at any time without any consequences or any explanation. If you do withdraw, we will ask whether we may use your data for data analysis. If you decline, your data will be destroyed.

A pseudonym will be assigned to your information to provide you anonymity to anyone outside of the research group. All published results will use this pseudonym. Your anonymity to the research group is not possible given the need to connect interview responses with online journal entries.

Your confidentiality and the confidentiality of the data will be protected by identifying data only with a participant number and user name (that you decide) rather than your name. Protection is also provided using password-protecting computer files, and storing study video and audio tapes in a locked office. Please try to avoid revealing personally identifying information (such as providing photos of yourself) during your journal entries as this will compromise your anonymity and confidentiality to people outside the research team. If you wish to examine the data relating to your participation, please ask.

It is anticipated that the results of this study will be shared with others in the following ways:

- Published articles
- Conference presentations
- Video publications
- Theses
- Internet project descriptions

Data from this study will be disposed of within 5 years. Electronic data, excluding online records, will be erased, paper documents will be shredded, and video/audio tapes will be recorded over or physically destroyed. We will delete online data but since we do not control this data, data backups may continue to exist. However, the only information connecting your online responses to you will be destroyed.

In addition to being able to contact the researcher at the above phone numbers, you may verify the ethical approval of this study, or raise any concerns you might have, by contacting the Associate Vice-President, Research at the University of Victoria (250-472-4545).

Your signature below indicates that you understand the above conditions of participation in this study and that you have had the opportunity to have your questions answered by the researchers.

Name of Participant

Signature

Date

A copy of this consent will be left with you, and a copy will be taken by the researcher.

Research Instructions

Dear potential participant

Thank you for agreeing to participate in the research study titled Visual Encodings of Data Away from Work: Exploring Casual Data Analysis and Key Contextual Factors. This study is expected to run over a two week period. It will require you to report “data artifacts” (defined below) and participate in an interview to be held at the University of Victoria.

Please CAREFULLY read the description of the experiment below before joining the study. The two week period will start once you fill in your first journal entry, but this must start within 2 days of receiving these instructions. If you have any questions please do not hesitate to contact David Sprague at (250) 472-5749 or dsprague@cs.uvic.ca.

Qualifications for Participation

Participants must be 19 or older. We require that you own or have access to a digital camera for the two week research period. Owning a USB flash drive is helpful, however, one can be lent to you. This study does not require any special skills, training, or computer proficiency; we are investigating how regular people use data artifacts.

Instructions

Participating in this study will require you to report encounters you have with visual representations of data during your daily life. We are particularly interested in knowing how you use and interpret data in situations outside of work, where job related pressures aren’t affecting your decisions. Please try to constantly pay attention to the data you see around you but do not change your daily behaviors more than necessary for this study. We definitely do not want you to actively search for data artifacts to fill a quota; if you happen to see or use an artifact, report it. Otherwise, it is perfectly acceptable for your journal to indicate you found nothing on a particular day. Furthermore, we are only interested in how you encounter information during your non-work time (defined below) so this study should not disrupt your studies or your job.

Artifact Collection

Please record or note your findings by storing data and screen shots on a flash drive, taking pictures of the artifact, or by providing a detailed description of it. Report an artifact after you are no longer using or examining it. We ask that you carry your flash drive and camera almost everywhere you go, therefore, small digital cameras and camera-phones are ideal. Use of the flash drive is more suited for computer based data sets that you can save and may not be easy to photograph. Screen shots can be saved to a USB drive (instructions for taking a screen shot: PC - <http://www.microsoft.com/windowsxp/using/setup/tips/screenshot.msp> Mac - <http://docs.info.apple.com/article.html?path=Mac/10.4/en/mh364.html>). Screen shots, photos, and saved data sets should *ideally* help your journal to illustrate what the data was and how it was presented. Photos can also show the context surrounding the artifact’s use (where were you and what were you doing?) and how you used the data.

Journal Entries

<http://www.surveygizmo.com/s/188241/artifact-journal>

At your convenience each day, please fill in an online journal entry (or a printed copy). Even if you have nothing to report, we expect you to fill out a brief entry. If you miss an entry, you may be reminded via email. We expect these entries to take between 3 minutes (no artifact) and a half hour (extensive details) with an typical artifact report taking 10-15 minutes. There is no minimum number of artifacts to report but we believe everyone will probably find more than 1 in two weeks. We should stress that you do not have to report every potential artifact you encounter, especially if this will embarrass you or you would consider this an invasion of privacy. We will occasionally monitor the online journal entries to both keep track of everyone’s progress and to ensure that you are performing the task as expected. Participants that fill in fewer than five (5) entries (including “nothing to report” entries) will be dismissed from the study. For each journal entry you submit, your name will be entered into a draw for three \$10 cash prizes. Please make sure you take a copy of the Implied Consent agreement found on the first page of the daily journal. If you choose to use the printed journal forms, please contact David. Finding a consistent time each day to write your journal may help establish a journal entry routine.

To ensure your privacy, your journal entries will only be associated with a user name you will email us. Please do not use your real name or personally identifying information as this information is the only way to connect your journal responses to you. Your user name will only be known to the experimenters and you. Therefore your online responses are anonymous to everyone else.

Interviews

Finally, we will arrange a meeting with you. As such, please provide us with a telephone number or email address that you can be readily contacted by. Please also let us know what days and times work best for your schedule.

Turn Over ->

Research Instructions

During the interview, we will discuss what you reported and your opinions. This interview is not expected to take more than an hour and it will be conducted in ECS 555 (the Engineering and Computer Science building) at the University of Victoria. [Click here for a map](#). At the end of the interview you will be paid (\$60) and be debriefed.

Definition of a Data Artifact

A data artifact is any sufficiently complex visual representation of data that fits the following criteria:

1. **Not All Alphanumeric Information** - Artifacts must represent some information without exclusively using letters or numbers (alphanumeric symbols). This information can be “encoded” in a wide variety of ways including position, order, shape, size, and colour. For example, a bar graph used for marathon training can show how many kilometers you and your friends have run each day. The bar order can indicate chronological time, bar size can indicate the distance run, and bar colours can differentiate people. Similarly, the board game Risk uses the colour, shape, and position of board pieces (military units) to convey who controls the unit, what that unit represents, and where the unit is located.

Minimum Requirements - Data can be considered almost anything from an alphabetical listing of books to a representation of the human genome. Therefore reported data artifacts should have multiple (>1) pieces of non-alphanumeric information displayed. An analogue watch is not a valid artifact since it only provides the current time (one piece of information). A student list sorted by an exam grade is not valid either, however, if this sorted list had names coloured according to academic major, it would be acceptable.

2. **Physical or Electronic** – Despite this being Computer Science research, we are equally interested in both physical and electronic artifacts.
3. **All in One Image**: Although a parking lot could be considered a data artifact (shape, position and colours represent parking permissions, locations, and availability), for this study the entire artifact must be (potentially) visible within a single photograph or screen shot to be acceptable (although you can still save multiple images).
4. **No repeats** - If you encounter an artifact multiple times, you should report it only once unless the context of its use has changed. For example, the marathon training bar chart could be used to explore your progress over time or to compare yourself against your friends (for a second report).

We are intentionally leaving the definition of an artifact somewhat vague because we want to see how you approach this task and what you come up with. A good question to ask yourself may be “If you take away all the text on the potential artifact (or the words are replaced with an unknown language), does it still convey multiple pieces of information?” (even if you would not understand what is represented). Just try your best.

Examples of data artifacts include but are not limited to...

Computer detected patterns in your daily life	Elaborate ordering of objects on a shelf	Node link diagrams	Color coded schedules
Charts & graphs	Stock market trends	Scientific diagrams	Video game “status bars”
Product comparisons	Dashboard displays	Floor plans	Maps

Remember: The more information presented with the artifact, the better. If you discover more than 5 artifacts per day and you don’t want to report them all, choose your favorites or the most complex ones. If you are unsure if something is an artifact, assume it is.

What is Non-Work Time?

We are only interested in seeing how you deal with data artifacts during your “non-work” time. This is essentially whenever you are not working at a job or studying / attending school. Non-work tasks can include: procrastinating (including work breaks), hobbies, house chores, running errands, socializing with family or friends, high energy leisure activities (jogging, sky diving, going to a club.....), and low energy activities (meditation, reading a book.....).

Needed Information (if you haven’t already given this):

Please send the following information to dsprague@cs.uvic.ca: 1) your name, 2) contact information (an email you regularly check and optional phone number), 3) your preferred interview time (two weeks after your first journal entry) and 4) the user name you want to use for the journal. Only the research team will have access to this information.

Data Artifact Experiment: Printable Journal Entry

=====
 Visual Encodings of Data Away from Work: Exploring
 Casual Data Analysis and Key Contextual Factors
 =====

This is a copy of the daily journal form to print out for use when you do not have Internet access. We encourage you to complete your daily journal online whenever possible. This can be found at: <http://www.surveygizmo.com/s/188241/artifact-journal>

=====
 Instructions
 =====

Everyday Visual Representations of Data:

This study examines how people use visual representations of data in their non-work related lives (leisure time, family time, and during chores). We ask you to notice how data sets are presented to you (a data artifact) and summarize your experiences via daily journal entries.

Data Artifacts: A visual representation of data which conveys multiple pieces of non-alphanumeric information. A simple question to ask yourself might be "if all the text and numbers were removed, how much information would I still see?" See the experimental instructions for more details. Examples include, but are not limited to: charts, graphs, floor plans, colour coded schedules, and node-link diagrams. If in doubt, assume something is a data artifact.

- Feel free to leave out any information/details that you feel uncomfortable reporting or are embarrassed about.
- If you collected supplementary information about the artifact (like photos or information saved on your flash drive) to accompany your paper journal entry, please remember to bring these materials to your interview. These supplementary materials are very important for our understanding. You can also email these files to David Sprague (dsprague@cs.uvic.ca).
- To qualify as an active participant, you are required to fill out a minimum of five (5) journal entries (including entries where you report 0 artifacts) and attend an end of study interview. Only active participants will be financially compensated for their time (\$60).
- Each journal entry made by an active participant will be entered into a series of three \$10 cash draws at the end of the study.

 Please make sure you have read and agree with the implied consent form (provided). Continuing with the survey implies your consent. You may pick up printed copies of all research documents (including journal entry forms) in room 654 of the ECS building (University of Victoria).

Data Artifact Experiment: Printable Journal Entry

=====

Data Artifact Report

Fill out 1 form per artifact OR one per day if you found no artifacts

=====

1. Please write your unique user name you have chosen. This will uniquely identify your reports.

User Name: _____ Date: _____

2. What non-work tasks did you perform today (check all that apply)?
- | | |
|---|---|
| <input type="checkbox"/> Surfing the Internet | <input type="checkbox"/> Shopped (food, clothes, etc) |
| <input type="checkbox"/> Sent Emails | <input type="checkbox"/> Did Banking |
| <input type="checkbox"/> Read for leisure | <input type="checkbox"/> Procrastinated at Work |
| <input type="checkbox"/> Watched TV/movie | <input type="checkbox"/> Pursued Hobby |
| <input type="checkbox"/> Played Video Games | <input type="checkbox"/> Calculated Personal Finances |
| <input type="checkbox"/> Exercised | <input type="checkbox"/> Did Household Chores |
| <input type="checkbox"/> Socialized (Friends/ Family) | <input type="checkbox"/> Read Instructions |
| <input type="checkbox"/> Took a Bus, Car or Bike | <input type="checkbox"/> Played a board/card game |
| <input type="checkbox"/> Other (Please specify) _____ | |

 *****If you are not reporting an artifact today, you may stop here*****

Artifact Being Described: _____

3. In how much of a rush were you (not necessarily related to the data artifact)?

- Very rushed
 Rushed
 Going at my usual pace
 Relaxed
 Very relaxed

4. What kind of "free time" were you experiencing when you noticed the data artifact (check all that apply)

- Relaxed leisure time (watching TV, meditating...)
 Active leisure time (jogging, dancing at a club...)
 Social time (spending time with family or friends)
 Chores & Tasks (grocery shopping, eating at home...)
 Taking a break from work
 Other

5. Please tell us the circumstances of your experience. Where were you and what were you doing?

Data Artifact Experiment: Printable Journal Entry

6. Why did you examine the data artifact (check all appropriate responses)?

- | | |
|--|--|
| <input type="checkbox"/> Wasting time | <input type="checkbox"/> To learn something new |
| <input type="checkbox"/> Curious about the data | <input type="checkbox"/> Data had personal meaning |
| <input type="checkbox"/> Curious about the system | <input type="checkbox"/> Self-reflection |
| <input type="checkbox"/> Data helped me perform a task | <input type="checkbox"/> I barely looked at the data |
| <input type="checkbox"/> Other (please specify) _____ | |

7. What plans/goals (if any) did you have in mind when using the artifact? For example, you may have been trying to figure out when your bus would arrive using a bus schedule.

8. What about the artifact caught your attention?

9. What did you like about the artifact?

10. What did you dislike about the artifact?

11. Additional comments:

Thank you for filling in your daily journal. We appreciate your help and feedback. If you have any questions or problems, you can contact the lead researcher David Sprague at dsprague@cs.uvic.ca.

<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)	<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																
<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)	<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																
<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)	<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																
<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)	<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																
<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)	<table border="1"> <thead> <tr> <th colspan="2">Journal Entry Website</th> </tr> </thead> <tbody> <tr> <td colspan="2">http://www.surveygizmo.com/s/188241/artifact-journal</td> </tr> <tr> <th>User Name</th> <th>Interview Time/Place</th> </tr> <tr> <td></td> <td>Time: ECS 555 (University of Victoria)</td> </tr> </tbody> </table> <p>Contact: David Sprague, (250) 472-5749, dsprague@cs.uvic.ca</p>	Journal Entry Website		http://www.surveygizmo.com/s/188241/artifact-journal		User Name	Interview Time/Place		Time: ECS 555 (University of Victoria)
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																
Journal Entry Website																	
http://www.surveygizmo.com/s/188241/artifact-journal																	
User Name	Interview Time/Place																
	Time: ECS 555 (University of Victoria)																

Permission to Use Data After Withdrawal

I hereby permit the researchers conducting the research study Visual Encodings of Data Away from Work: Exploring Casual Data Analysis and Key Contextual Factors to continue using the data I have provided, despite having withdrawn from the study.

Name: _____ Date: _____

I hereby permit the researchers conducting the research study Visual Encodings of Data Away from Work: Exploring Casual Data Analysis and Key Contextual Factors to continue using the data I have provided, despite having withdrawn from the study.

Name: _____ Date: _____

I hereby permit the researchers conducting the research study Visual Encodings of Data Away from Work: Exploring Casual Data Analysis and Key Contextual Factors to continue using the data I have provided, despite having withdrawn from the study.

Name: _____ Date: _____

User	File Name	Artifact Name
Broccoli	file_188241_60370983_0_dsc_1035.jpg	United Way Chart
Flamingo	NP_002	Google Map
Flamingo	file_188241_60437182_1_coins.jpg	Coins
Broccoli	file_188241_60525175_0_dsc_1335.jpg	Cell Phone Battery
Broccoli	IMG_06351 1 copy.jpg	Joke Pie Chart
Gravity	file_188241_60624002_0_nosigns.jpg	Private Parking Sign
Broccoli	file_188241_60768403_0_dsc_1341.jpg	Metro Map
Broccoli	file_188241_60913872_0_dsc_1344.jpg	Nutrition Info
Gorffy	file_188241_60915845_0_picture1.png	Slashdot Survey Results
Poppy	file_188241_60918757_0_img_3382.jpg	Psych Textbook
Meerkat	IMG_0557.JPG	Microwave
Meerkat	IMG_0556.JPG	Cornett Map
Broccoli	file_188241_61152227_0_dsc_1348.jpg	Chore Wheel Chart
Leon	file_188241_61242750_0_1visualdata.jpg	Housing Costs: Bar graph and line diagram
Leon	file_188241_61242750_1_2visualdata.jpg	Blog Bar graph
Broccoli	file_188241_61262088_0_dsc_1352.jpg	Globe
Dora	IMG_2498.JPG	Blackberry Screen
Gorffy	file_188241_61309046_0_edres-jan09.pdf	Uvic map
Benji	NP_009	Pie Chart (physics building)
	116958 file_188241_61422531_0_gmail.jpg	Gmail (each email is object)
	116958 file_188241_61422531_1_imdb.jpg	IMDB
	116958 file_188241_61422531_2_firefox.jpg	Firefox
Broccoli	file_188241_61426919_0_dsc03737.jpg	Organ
Dora	IMG_2501.JPG	License plate
Broccoli	file_188241_61460302_0_dsc_1372.jpg	Wii Golf Gauge
Benji	NP_007	Gulf Islands Map
Poppy	file_188241_61481918_0_img_3383.jpg	CBC News Poll
Dora	IMG_2503.JPG	Traffic lights
	116958 file_188241_61497816_0_faucet.jpg	Shower Faucet
Broccoli	file_188241_61578169_0_dsc_1384.jpg	Toaster
	116958 file_188241_61578656_0_uvic.jpg	Uvic Emblem
	116958 Mix.jpg	Electric Mixer
Poppy	file_188241_61580102_0_img_3387.jpg	Costco Tanning Poll
Poppy	file_188241_61580102_1_img_3388.jpg	Japanese TV Poll
Gorffy	file_188241_61581953_0_antingchart.png	Planting Chart
Meerkat	NP_006	Cribbage Board
Leon	file_188241_61653631_0_img_1566.jpg	Clock
Poppy	file_188241_61686662_0_img_3390.jpg	CIBC Graph
	116958 file_188241_61698475_0_afficlights.jpg	Traffic Lights
	116958 file_188241_61698475_1_walk.jpg	Crosswalk Button
S2	file_188241_61700558_0_p1210054.jpg	Alarm Clock Warnings
S2	file_188241_61700558_1_p1210056.jpg	Alarm Clock Battery Instructions
Flamingo	file_188241_60437182_1_coins.jpg	Coins
Flamingo	file_188241_61701456_0_at6.25.17pm.png	YouTube Video
Flamingo	file_188241_61701456_1_ner-diagram.jpg	Guitare Tuner
Flamingo	NP_002	Google Map
Flamingo	NP_004	Traffic lights
Rubix	file_188241_61707123_0_bus_map.gif	Bus Map
Benji	P1000108.JPG	Physics Building Charts (bar charts)
Broccoli	file_188241_61757921_0_v.jpg	Hotmail Folder
Leon	file_188241_61776044_0_day6profile.jpg	World Income Distribution
Gorffy	IMG_0178.jpg	Building Map
	116958 file_188241_61847059_0_soap.jpg	Soap Dispenser
Dora	IMG_2505.JPG	Street signs
Broccoli	file_188241_61918235_0_x.jpg	Vista Weather Widget



DSCN6144.JPG




CIMG1804.JPG



CIMG1805.JPG



CIMG1809.JPG



CIMG1810.JPG



DSCN6156.JPG



file_188241_140363506_0_inception.png



file_188241_140446870_0_weather.png



file_188241_140336551_0_cippers.png



file_188241_139198604_0_aintraining.png




gotes.jpg




DSCN6149.JPG




DSCN6155.JPG



file_188241_140183412_0_eatlossgain.jpg



file_188241_140183412_1_eneutrality.png



file_188241_140183412_2_energygraph.png




file_188241_137701315_0_dsc04472.jpg



file_188241_138987657_0_img_0018 (1).jpg



file_188241_138987657_1_img_0019.jpg



file_188241_140845133_0_nation.png.jpg



file_188241_140674760_0_eatingchart.png



file_188241_140531121_0_img_0012.jpg



file_188241_139614531_0_img_0001.jpg



file_188241_139174099_0_img_0064.jpg



IMG_0011.JPG



Artifact 4 (1).jpg




Artifact 1.jpg



Artifact 2.jpg



Artifact 3.bmp




file_188241_142130301_0_jesuschrist.jpg




file_188241_141774841_0_e-princess.jpg



file_188241_141549192_0_kabballah.jpg




file_188241_141015080_0_planetsizes.jpg



file_188241_140684473_0_rces-map-r2.gif



file_188241_140366146_0_mainmap.jpg




1 Brain history.JPG




1 instruction.JPG



1 weather forecast in Mt. Washington.bmp



2 Brain Profile.JPG




2 poster.JPG



2 snowfall in Mt. Washington.bmp



4 google reader (analysis).bmp



analysis.bmp



Bag.jpg



Cup.jpg



downtown vancouver.gif



driversconsole.JPG




Emergency Evacuation plan.JPG



Evony.png



file_188241_60370983_0_dsc_1035.jpg



file_188241_60437182_0_ner-diagram.jpg



file_188241_60437182_1_coins.jpg




file_188241_60525175_0_dsc_1335.jpg



file_188241_60525175_1_unintitled.jpg



file_188241_60624002_0_nosigns.jpg



file_188241_60768403_0_dsc_1341.jpg



file_188241_60913872_0_dsc_1344.jpg



file_188241_60915845_0_picture1.png



file_188241_60918757_0_image_3382.jpg



file_188241_61152227_0_dsc_1348.jpg



file_188241_61242750_0_visualdata.jpg



file_188241_61242750_1_visualdata.jpg



file_188241_61422531_0_gmail.jpg




file_188241_61422531_1_image_db.jpg




file_188241_61422531_2_firefox.jpg



file_188241_61426919_0_dsc03737.jpg



file_188241_61460302_0_dsc_1372.jpg



file_188241_61481918_0_image_3383.jpg



file_188241_61497816_0_facet.jpg



file_188241_61578169_0_dsc_1384.jpg



file_188241_61578656_0_uv
ic.jpg



file_188241_61580102_0_im
g_3387.jpg




file_188241_61580102_1_im
g_3388.jpg



file_188241_61581953_0_a
ntingchart.png



file_188241_61653631_0_im
g_1566.jpg



file_188241_61686662_0_im
g_3390.jpg



file_188241_61698475_0_af
fights.jpg



file_188241_61698475_1_w
alk.jpg




file_188241_61700558_0_p
1210054.jpg



file_188241_61700558_1_p
1210056.jpg



file_188241_61701456_0_at
6.25.17pm.png



file_188241_61701456_1_n
er-diagram.jpg



file_188241_61707123_0_b
us_map.gif

file_188241_61757921_0_v.j
pg



file_188241_61776044_0_d
ay6profile.jpg



file_188241_61847059_0_so
ap.jpg



file_188241_61918235_0_x.j
pg



file_188241_61993960_0_at
artifact.png



file_188241_62041950_0_ce
ll.jpg



file_188241_62101678_0_wi
dgest.png



file_188241_62131163_0_b
usmap.jpg



file_188241_62133220_0_ja
nuary25.jpg



file_188241_62133477_0_ds
c02728.jpg



file_188241_62133657_0_p
1210059.jpg



file_188241_62133657_1_p
1210063.jpg



file_188241_62133657_2_p
1210065.jpg



file_188241_62134122_0_im
g00001[1].jpg



file_188241_62163292_0_d
ay8profile.jpg



file_188241_62177738_0_e
psimachine.jpg



file_188241_62246630_0_im
g_0011.jpg



file_188241_62246630_1_im
g_0009.jpg



file_188241_62252351_0_cu
p.jpg



file_188241_62258798_0_ds
c02732.jpg



file_188241_62262941_0_w
eather.jpg



file_188241_62301537_0_pi
cture1.png




file_188241_62350909_0_w offram.jpg




file_188241_62433414_0_m acintosh.jpg




file_188241_62437141_0_im g_0011.jpg



file_188241_62557275_0_ay 11profile.jpg



file_188241_62557275_1_ay 12profile.jpg



file_188241_62559014_0_ti me_chart2.png



file_188241_62563216_0_dr ywall.jpg



file_188241_62563216_1_st reet.jpg



file_188241_62653862_0_p 1210066.jpg



file_188241_62693964_0_re creation.jpg



file_188241_62837287_0_ds c02738.jpg




file_188241_62837287_1_ds c02739.jpg



file_188241_62840211_0_gr aph.jpg



file_188241_63067979_0_sc reenshot.png



file_188241_63098340_0_fu nnygraph.jpg



file_188241_63098340_1_sh ampool.jpg



file_188241_63343148_0_im g_0009.jpg



file_188241_63431872_0_im g_0017.jpg



file_188241_63520597_0_ds c02749.jpg



Frames.jpg



GulfIslandsMap1.gif



GulfIslandsMap2_serviceAre as.gif



imagejpeg_2.jpg



imagejpeg_map.jpg



IMG0001[1].jpg



IMG_0178.jpg



IMG_0179.jpg



IMG_0180.jpg




IMG_0556.JPG



IMG_0557.JPG



IMG_0563.JPG



IMG_06351 1 copy.jpg



IMG_2498.JPG



IMG_2501.JPG



IMG_2503.JPG



IMG_2505.JPG




Mafia Wars.png



Mix.jpg




Mobsters 2.png



movie_narrative_charts_large.png




NN 1 (number).bmp



NN 2 (origin).bmp



NN 3 (key word).bmp



NN 4 (referral).bmp



P1000107.JPG




P1000108.JPG



P1000123.JPG



P1010486.JPG



P1010492.JPG



P1010493.JPG




P1010498.JPG




P1010500.JPG




P1010501.JPG




P1010508.JPG




P1010511.JPG



P1010512.JPG



P1010516.JPG



P1010519.JPG



P1010523.JPG



P2010640.JPG



P2050639.JPG



weather forecast in Mt. Washington.bmp