Implied Dynamics in Information Visualization

Caroline Ziemkiewicz
UNC Charlotte
9201 University City Blvd
Charlotte, NC 28262
caziemki@uncc.edu

Robert Kosara
UNC Charlotte
9201 University City Blvd
Charlotte, NC 28262
rkosara@uncc.edu

ABSTRACT

Information visualization is a powerful method for understanding and working with data. However, we still have an incomplete understanding of how people use visualization to think about information. We propose that people use visualization to support comprehension and reasoning by viewing abstract visual representations as physical scenes with a set of implied dynamics between objects. Inferences based on these implied dynamics are metaphorically extended to form inferences about the represented information. This view predicts that even seemingly meaningless properties of a visualization, including such minor design elements as borders, background areas, and the connectedness of parts, may affect how people perceive semantic aspects of data by suggesting different potential dynamics between data points. We present a study that supports this claim and discuss the design implications of this theory of information visualization.

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation**]: Miscellaneous

1. INTRODUCTION

The power of information visualization arises from its ability to apply perception and visual thinking to understanding complex data and solving difficult analytical problems. This use of visual thinking also raises many unanswered questions about how visualization works. How do images perceived by a user become information she can analyze? How do the surface properties of those images affect her understanding of that information?

Although visualization makes use of the idea that making information visual contributes somehow to the reasoning process, designing visualization in practice is often seen as simply a matter of maximizing the amount of visible information in the assumption that it will be transferred to the user's head in a straightforward manner. There is traditionally little consideration for the ways in which the design of this presentation of information can affect its comprehension. This lack of consideration can lead to unexpected and problematic results. There is evidence that informationally equivalent charts with different designs can lead to surprisingly large differences in user responses. Zacks and Tversky [18] showed that users viewing a two-point line graph will often interpret the data as a trend even when the graph depicts two separate groups, such as males and females. This can lead to odd descriptions of the data, such as, "The more male a person is, the taller he/she is." More seriously, Elting et al. [5] found significant differences in the number of errors produced by physicians viewing the same clinical study data in several charts of different types and presentation styles. The design of data presentation in this case affects not only users' immediate perceptions, but also the decisions they make.

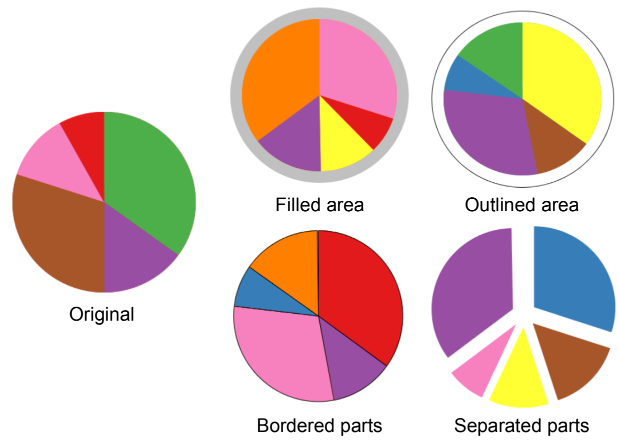

Figure 1: A pie chart in five different design configurations.

Part of the reason these effects are not accounted for in visualization theory is that it has historically focused on the perceptual qualities of individual visual objects. This has yielded a large body of knowledge on how perceptual elements such as color, size, and position can be read from a visualization, but this object-level theory has had less to say about how a user interprets the overall structure of a visualization scene. One problem of this narrow focus is that it places the emphasis of visualization theory on the syntax of information visualization, without considering its semantics. This reductionist viewpoint, while useful, needs to be supplemented by a holistic view of how visualization works. For all we have learned about how information is mapped to a visual form and then read by a user, we have very little sense of the impact of the visualization as a whole on the user's overall mental model of the data.

We argue that findings in visual cognition strongly suggest that design choices in a visualization do significantly affect how this mental model of data is constructed by a user. What is needed is a theoretical framework to explain how such design choices affect the interpretation and use of information visualization systems. We

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AVI '10, May 25-29, 2010, Rome, Italy
Copyright 2010 ACM 978-1-4503-0076-6/10/05 ...$10.00.

theorize that visualization use involves not just the passive decoding of information, but also an active simulation of visual dynamics that supports reasoning about that information. Elements that do not carry information about data points may therefore still carry semantic information about these implied visual dynamics. These design elements can thus affect how a user understands and reasons with visualized information.

In order to elaborate upon and test this hypothesis, we performed a study in which four design elements (Figure 1) were manipulated in various combinations across five simple visualization types (Figure 2). We measured whether participants rated the data visualized differently across each of these configurations on a number of semantic dimensions relating to the implied dynamics of the visual scene as well as the structure of the data. We contribute a review of previous work that places visualization use in the context of dynamic simulations of visual scenes, the quantitative and qualitative results of a study that provides preliminary support for this theory, and a set of design implications based on our findings.

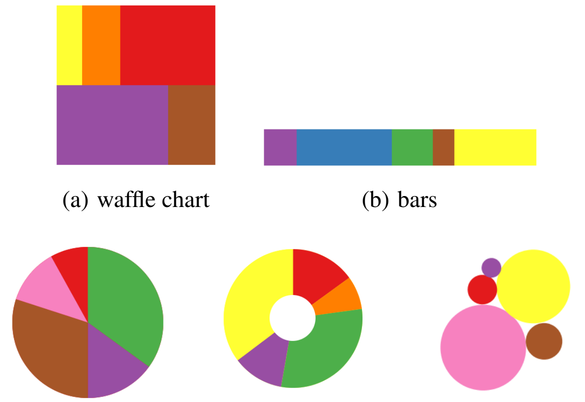

(a) waffle chart (b) bars (c) pie chart (d) donut (e) bubble chart

Figure 2: The five chart types used in the study.

2. RELATED WORK

The foundation of our work is the idea that the surface structure of a visualization can affect a user's understanding of the structure of the data it visualizes. Research in cognitive science and visualization do suggest a powerful connection between perceived visual dynamics and conceived information relationships. Zacks and Tversky [18] show that participants' interpretation of a simple two-point chart as representing either a trend or two separate groups is more influenced by the type of chart (i.e., bar chart or line graph) than by the type of data. Similarly, Richardson et al. [14] studied the types of simple diagrams generated by participants prompted to visually represent a series of verbs. The consistencies across participants found by the authors provides evidence for strong semantic associations with simple visual structures.

There has been work in visualization and human-computer interaction that studies the effects of visual structure by primarily focusing on how visual presentation affects the way a user accesses information [11]. However, there is also evidence from these fields that visual structure can affect mental models of data in a more profound way. Research on the effect of visual metaphors on visualization use suggests that visual structure can be thought of as analogous to conceptual metaphors in language, and that visual and verbal metaphors can interfere with each other in the course of reading a visualization [19]. An interesting aspect of verbal metaphors is that they usually map properties of abstract concepts such as emotions or ideas to properties of physical objects, and that inferences about the abstract concept are often drawn from physical simulations of the metaphor's source [10].

Similarly, our view of visualization is based on the idea that people reason with visual representations in part by viewing them as dynamic physical scenes and simulating the forces acting on them. This premise is strongly supported by findings in psychology and cognitive science. The foundation of this work comes from Freyd, who found that people recalling a variety of both photographic and diagrammatic scenes with implied motion tend to remember objects as being farther along in their path of motion than they actually were [6]. The same effect also appears in line drawings with no apparent motion but which suggest the presence of physical forces such as gravity or spring compression [7]. These and related findings provide evidence that perception involves active simulation of what might happen next in a scene, based on the apparent forces at work, such as gravity, friction, and momentum [9] and the perceived relationships between objects in a scene, such as whether they contact, support, or are attached to one another [15].

This work dovetails with that of the art theorist and psychologist Rudolf Arnheim, who suggests [2, 3] that a physical property such as gravity is naturally applied to our interpretation of the semantic content of shapes and their composition in a visual scene, explaining the emotional effects of abstract art such as that of Josef Albers or Piet Mondrian. These ideas also bear some relation to the idea of affordances, in which perceived objects are seen in terms of their potential for action by some agent [8]. Perceived affordances, in a broad interpretation, are often used as a framework for interface design [13]. Clearly, the idea that a scene's implied dynamics and physical properties can carry meaning is not unknown in design or human-computer interaction. However, we still lack the ability to reliably predict what meanings will be carried by a given design.

3. HYPOTHESIS

We propose that the meaning carried by visual design derives from the user's performing dynamic simulations with perceived visual objects and metaphorically interpreting the results of those simulations as inferences about information. As a first step in supporting this theory, we hypothesize that design elements that carry implicit physical information, such as borders, connectedness, and background shapes, can have significant and consistent effects on how a user interprets a visualization. Specifically, we believe that the implied dynamics of a visualization as derived from elements of its visual design can be shown to affect subjective responses to entities represented in visualized data in a reliable and systematic fashion. We further hypothesize that while these effects will be present and to some extent consistent across visualization types, they can also be generated by the visualization type itself, and design elements will affect different visualizations in different ways.

4. METHODS

In order to test our hypotheses, we designed a study to test the effects of visualization design elements on semantic judgements of data in a simple context. This study was meant to test whether a particular set of design elements indeed have a significant effect on semantic judgements, and to what extent these design elements affect various types of simple information visualizations.

4.1 Participants

We recruited 43 participants via Amazon's Mechanical Turk [1]. Participants performed the study online and were paid a base rate

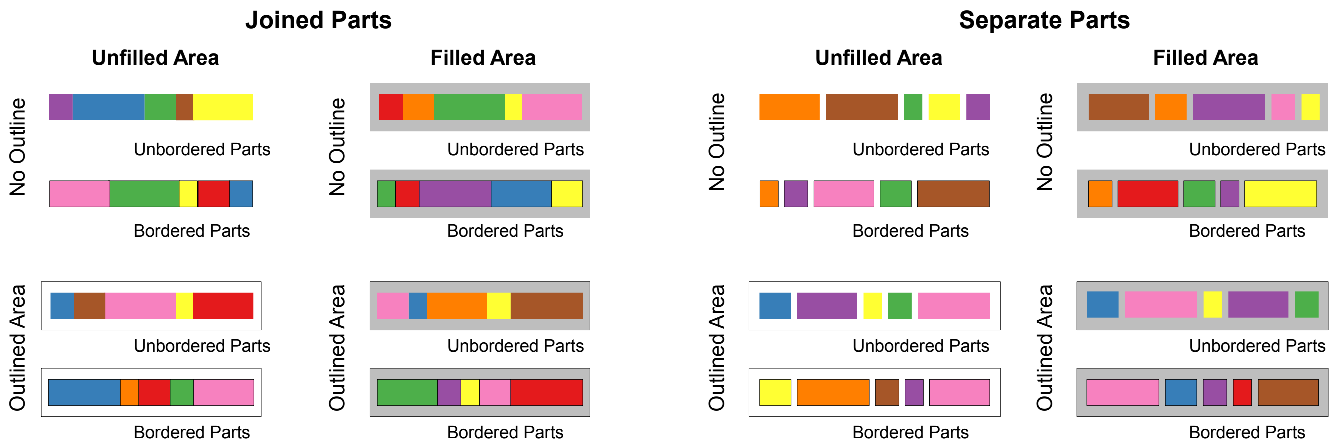

Figure 3: The stacked bar chart in each of the sixteen design configurations used in the study. These sixteen configurations represent every combination of the four binary design element variables we chose.

of $0.20 for their work, which took at most twenty minutes. They were told that giving especially helpful responses (defined as detailed, thoughtful, and clearly written) would yield a bonus payment of $1.00. Twelve participants ultimately received this bonus.

One participant was dropped from the study and not compensated due to a clear attempt to cheat the system by entering the same response for every question, leaving a total of 42 participants. Of the remaining participants, 25 were female (59.5%) and 17 were male (40.5%). Participant age ranged from 21 to 62, with an average age of 36.4.

4.2 Materials

Participants viewed a series of twenty charts which were described as representing fictional companies. Each company was divided into six departments, with the pieces in a chart representing departments sized according to the department's relative expenditures for a fiscal year. These proportions were the same for all twenty charts, but the order and coloring of the departments in each chart was randomized to conceal this fact. Although we cannot know for certain whether this was successful, only one participant's comments indicated that he realized that the proportions were all the same. The colors were derived from a categorical color scale from ColorBrewer [4]. This "company/department" description was used because it is a largely abstract concept that can lend itself to different conceptualizations, but is still familiar enough that participants could easily interpret data about it.

We used five types of charts to display the relationships of the departments to the company: a pie chart, a waffle chart (or one-level treemap), a horizontal stacked bar chart, a donut chart, and a bubble chart (Figure 2). These visualizations were chosen to represent a range of shapes and relative familiarity, while all being simple enough to evaluate quickly.

Having chosen this set of visualizations, we systematically altered them according to four design dimensions. These dimensions represented the presence or absence of some element we hypothesized to have semantic value. The elements were:

Filled area. In charts with a filled area, a gray background was visible behind the main chart which mimicked its overall shape but at a larger size. For the bubble chart, which is not space-filling, we used a circle as the filled area.

Outlined area. In charts with an outlined area, a black contour was drawn around the chart in the same shape as the main chart but at a larger size. For the bubble chart, we used a circular outline.

Bordered parts. In charts with bordered parts, black contours were drawn around each of the individual pieces.

Joined parts. In charts with joined parts, pieces were connected to one another in the manner most natural for a given visualization type. In separated parts conditions, the view was "exploded" so that a small amount of space was visible between parts.

With two possible states for each of these four variables, there were a total of sixteen design configurations for each of the chart types. Figure 3 shows these sixteen configurations applied to the bar chart.

These design elements were chosen because outlines, color areas, and connectivity are among the cues that visualization researchers such as Ware [17] and Tversky et al. [16] consider to be meaningful visual primitives. Additionally, we hypothesized each to have a unique semantic effect on the dynamics of the visual representation. We expected a filled area to suggest a stable foundation, the outlined area to suggest a limit or fence around the entire company, bordered parts to suggest limits on individual pieces, and the joining of parts to suggest connections between pieces.

To test these hypotheses, we developed a list of ten semantic variables which could describe a simple dataset of the kind we presented in structural terms. On a scale of one to five, we asked participants to rate how much the company presented in a chart was likely to be *stable, complete, controlled, inflexible, rigid, structured, isolated, unified, well-organized*, and *good place to work*.

While these variables were primarily chosen to represent a broad range of abstract structural qualities a company might have, they can also be interpreted as representing a range of dynamic physical properties. For example, *rigid* and *inflexible* may imply that a scene cannot be simulated as changing very much. *Stable* and *controlled* suggest that a figure is well supported and will not move without additional forces. *Unified* and *complete* may suggest that parts are seen as attached to one another and will move as a unit, while *isolated* may suggest that parts will move independently. *Structured* may imply that the scene represents a connected unit with dynamics in balance. The last two semantic variables, *good place to work*

Scale Item	Factor 1 (Balance)	Factor 2 (Strictness)
Complete	.795	.053
Controlled	.530	.559
Good Place to Work	.633	-.368
Inflexible	.141	.852
Isolated	-.250	.530
Rigid	.211	.852
Stable	.819	.057
Structured	.752	.331
Unified	.753	.073
Well-organized	.823	.062

Table 1: Factor loadings for the two factors we derived from the original ten semantic scale items, using a principal component analysis with Varimax rotation.

Factor	Group	Balance		Strictness	
		M	S.D.	M	S.D.
Filled Area	yes	-.032	.12	.004	.11
	no	-.067	.12	.106	.12
Bordered Parts	yes	-.071	.12	.125	.12
	no	.043	.11	-.086	.11
Joined Parts	yes	.160	.09	.073	.09
	no	-.188	.10	-.035	.09
Outlined Area	yes	-.030	.09	.022	.09
	no	.002	.08	.016	.09

Table 2: Means in the two derived semantic factors for each of the four design elements we studied. Only the variable of Joined Parts had a significant main effect on the factor of Balance.

and *well-organized*, are more subjective and may represent combinations of the other variables.

4.3 Procedure

We varied two of the four design elements between subjects (filled areas and bordered parts) and two of the elements within subjects (outlined areas and joined parts). The purpose of varying some elements between subjects was to reduce the number of charts a participant saw in order to avoid fatigue. We hypothesized that the design elements we chose as between-subjects variables were less likely to have an effect than the within-subject variables, although this did not entirely prove to be the case.

In the first part of the study, participants saw a series of twenty charts described as representing fictional companies. Participants were told that segments in a chart represented the departments of the company, and that the size of each segment represented the amount of spending by that department over a fiscal year. The ten semantic variables were presented below each chart, in an order which was randomized for each participant but kept consistent throughout a single participants' progress. Participants were told to rate each company on these semantic variables to the best of their ability based on the information in the chart. The charts each participant saw included four versions of each of the five chart types, varied on the two within-subjects variables of *outlined area* and *joined parts*. These charts were presented in random order. After the participant rated each chart on all ten variables, she clicked a button to continue to the next chart.

In addition to this main part of the study, we wished to give the participants an opportunity to explain their ratings in more detail so we could better understand how they viewed the charts semantically. We also wanted to make sure that participants felt comfortable rating the companies based on the charts and were able to give reasons for their ratings. After the first part of the study, we chose four pairs of charts to present to the participant. These pairs included two pairs in which the charts were visually similar (that is, were mostly the same on the visual dimensions) but were rated very differently by the participant, and two pairs in which the charts were visually different but rated similarly. We calculated the difference in ratings for every pair of charts by taking an average of the absolute difference between their ratings on each of the ten semantic dimensions. During the second phase of the study, we presented the two charts in each pair to the participant side by side and provided a text box in which she was asked to explain why she rated the charts either similarly or differently.

Once the participants finished these two phases of the study, they were asked to provide their gender and age in an additional form on the Mechanical Turk site.

5. RESULTS

Since our ten items in the scale we employed likely show a good deal of semantic overlap, the first step in analyzing these results was to reduce the number of dimensions using factor analysis. We employed a principle components analysis with varimax rotation and selected the resulting factors with an eigenvalue greater than one, which produced two factors. On examination of the factor loading (Table 1), the first factor seems to encompass the scale items that are more positive and suggest stability and good organization, while the second factor encompasses the items that are more negative and suggest rigidity and oppressiveness. However, since the written responses by participants show some variation between subjects as to which of the scale items were considered negatively or positively, I will refer to these factors more neutrally as "Balance" and "Strictness" in the subsequent analysis.

The effect of our overall model (a 2x2x2x2x5 design, with the four design elements and chart type) on these factors was first assessed with a multivariate analysis of variance (MANOVA) using Wilks' Lambda as the test statistic . This analysis found a significant main effect of chart type, $F(8, 31) = 3.04$, $p < .05$, $\eta^2 = .44$. There was also a significant effect of joined parts, $F(2, 37) = 7.11$, $p < .01$, $\eta^2 = .28$. In addition, there were significant interactions of filled area by bordered parts ($F(2, 37) = 3.82$, $p < .05$, $\eta^2 = .17$), outlined area by joined parts ($F(2, 37) = 3.94$, $p < .05$, $\eta^2 = .18$), and chart type by outlined area by filled area ($F(8, 31) = 2.51$, $p < .05$, $\eta^2 = .36$). There was also a marginally significant interaction between chart type and joined parts, $F(8, 31) = 2.21$, $p = .054$, $\eta^2 = .36$. Results that were significant in the multivariate analysis were further analyzed with univariate analyses of variance (ANOVA) for the two factors.

5.1 Main Effects

A summary of the main effects of the four design elements on the factors of Balance and Strictness is presented in Table 2. The only design element that produced a significant main effect in the MANOVA was joined parts, which a univariate repeated measures ANOVA found to be significant on the factor of Balance, $F(1, 39) = 12.18$, $p < .001$, $\eta^2 = .24$. That is, charts with joined parts ($M = .16$, $S.D. = .09$) were rated significantly higher on this factor than those with separated parts ($M = -.19$, $S.D. = .09$).

In addition, we found a significant main effect of chart type on both Balance, $F(2.73, 39) = 8.32$, $p < .001$, $\eta^2 = .18$, and Strictness, $F(2.43, 39) = 6.38$, $p < .001$, $\eta^2 = .14$. (Since the variable of chart type did not meet the assumption of sphericity for either factor, we employed a Greenhouse-Geisser correction on these ANOVAs.) The means for each chart type on these two

Chart Type	Balance		Strictness	
	M	S.D.	M	S.D.
waffle	.179	.11	.301	.11
bars	.085	.13	.204	.13
pie	.116	.08	-.045	.10
donut	-.002	.09	-.093	.12
bubble	-.447	.14	-.271	.12

Table 3: Means in the two derived semantic factors for each of the five chart types.

Chart Type	Joined Parts	Balance		Strictness	
		M	S.D.	M	S.D.
waffle	joined	.406	.13	.370	.11
	separated	-.049	.13	.231	.14
bars	joined	.140	.15	.202	.14
	separated	.029	.14	.206	.14
pie	joined	.340	.12	.121	.13
	separated	-.109	.12	-.211	.11
donut	joined	.295	.11	-.132	.12
	separated	-.299	.13	-.054	.14
bubbles	joined	-.382	.14	-.197	.13
	separated	-.512	.16	-.346	.13

Table 4: Means for each of the five chart types across the design variable of joined parts.

factors are summarized in Table 3. In general, these results show that the bubble chart is rated as much less Balanced than the other charts, meaning it was seen as less stable, unified, complete, and well-organized. Bubble charts were also rated as less Strict than the other chart types, and the rectangular charts–i.e., bars and waffles–were rated as more Strict than the predominantly circular charts, meaning they were rated as more rigid, inflexible, structured, and controlled.

5.2 Interactions Between Design Elements

In addition to these main effects, we also found several significant interactions in our overall model. Repeated measures ANOVAs found the interaction between filled area and bordered parts to be significant for both Balance ($F(1, 39) = 4.95$, $p < .05$, $\eta^2 = .12$) and Strictness ($F(1, 39) = 5.28$, $p < .05$, $\eta^2 = .12$). These interactions are illustrated in Figure 4. Either a filled area or bordered parts on their own seem to increase perceptions of Balance, but the presence of both elements or neither element creates the perception of less Balance. This suggests that the positive effects these elements have on perceptions of organization interfere with one another in some way. The effect of the interaction on Strictness adds some nuance to this interpretation. A chart with both a filled area and bordered parts is rated as highly Strict, as is a chart with neither element, which is complimentary to the Balance effect. However, while a chart with borders and no filled area is rated neutrally on the factor of Strictness, a chart with a filled area and no borders is rated as much less Strict. This suggests that a filled area produces a perception of flexibility which is somehow tempered by the presence of borders around parts.

The interaction of outlined area by joined parts was significant only for the factor of Balance, $F(1, 39) = 8.08$, $p < .01$, $\eta^2 = .175$. This interaction seems to arise from the fact that the difference in Balance between a chart with joined and separated parts is larger when there is no outline around the chart area. This offers the intriguing possibility that the perceived instability of the "exploded" charts is mitigated when there is a boundary limiting the perceived motion of the pieces.

In addition to interactions between design elements, we also found minor but significant interactions between chart type and joined parts for both the factor of Balance ($F(4, 39) = 3.56$, $p < .01$, $\eta^2 = .09$) and Strictness ($F(4, 39) = 2.51$, $p < .05$, $\eta^2 = .06$). The means for these conditions are summarized in Table 4. Generally, the perceived Balance of a bar chart or bubble chart is less affected by whether pieces are joined or not than that of other chart types. Additionally, pie charts show a greater loss of Strictness when parts are separated than do the other chart types. As with the factor of Balance, bars are largely unaffected by separation for the Strictness factor, and donut charts show a slight trend in the opposite direction, with separated charts perceived as more Strict on average than joined ones.

Finally, there is a three-way interaction of chart type, outlined area, and filled area for the factor of Balance, $F(4, 39) = 2.98$, $p < .05$, $\eta^2 = .07$. This difference seems to be mostly attributable to the fact that for waffle charts alone, an outline increases the perception of Balance only when there is no fill, and vice versa. This recalls to some extent the interaction of filled area and bordered parts, and suggests that outlines and fills, combined with an already quite rigid visual structure, can lead to a more negative impression.

6. GENERAL DISCUSSION

It is clear from the quantitative results of this study that design elements in a simple visualization context can have significant and to some extent consistent effects on a user's semantic evaluation of data. In addition to the simple semantic ratings, we also attempted to analyze the reasons for these ratings in the second part of our study, in which participants were asked to explain selected ratings. These explanations can shed some light on what design elements mean to users. There were no significant patterns to the types or configurations of charts automatically chosen for comparison, so participant comments covered a wide range of conditions. In general, these comments tended to reinforce our quantitative results and provided anecdotal evidence that a tendency to interpret charts as dynamic physical scenes leads to the patterns of semantic ratings that we found. They also demonstrate that our users were easily able to explain their ratings in almost all cases, suggesting that the task, while unusual, was understandable; only one of our 42 participants reported having difficulty rating the companies based on the charts. While many comments were minimal or did not address structural properties (for example, comments in which a participant preferred one chart to another because the colors were more pleasing), those which were more elaborate did tend to talk about dynamic and physical qualities of the chart.

6.1 Implied Dynamics of Chart Types

One common theme throughout these explanations was a treatment of both design elements and chart types as offering various potentials for movement or communication. For example, two separate participants explained that they considered the donut chart less stable because it seemed like it might "roll away." This kind of analysis also seemed to underly the evaluations of the bubble chart as unstable and uncontrolled, with participants describing this chart as "floating bubbles that were barely contained within the area" and "scattered." (It should be noted that apart from a mention of the pie chart in the initial instructions, we did not name or label the charts in the study, so the participant who described the bubble chart as "bubbles" did so spontaneously.) These comments suggest that at least some participants were implicitly applying gravity

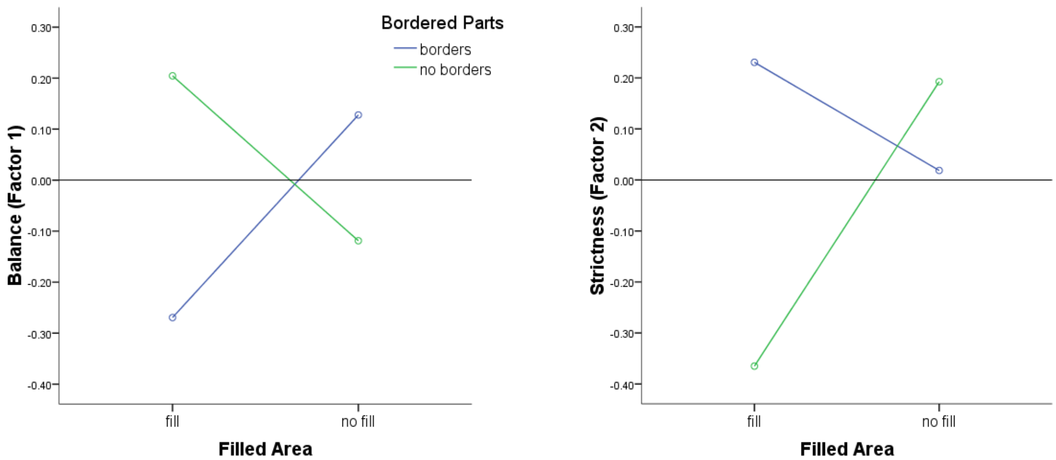

Figure 4: We found a significant interaction between the design factors of filled area and bordered parts. Balance increased when only one of the factors was present, but decreased when both or neither were present. Strictness was increased when both or neither were present, but was much lower when there was a fill but no part borders.

to the charts, and found the bubbles disconcerting because they so strongly seemed to violate the constraints of gravity.

This kind of description may also help to explain why why perceived Strictness was seemingly reduced with a filled area. One participant describing a bubble chart on a filled area said that it looked "as if the parts were placed randomly, with room to move them around," and another that it seemed "as if the company doesn't quite know how big it is." Of the five visualization types we presented, the bubble chart was the one with the least constraints on placement of parts; in the other four cases, the pieces are filling a predefined space of one shape or another. The arbitrary nature of the bubbles' placement, then, may be highlighted when there is a clearly defined space in which they can "move." In contrast, one participant described departments in the space-filling waffle chart as having "no room to move around."

6.2 Implied Dynamics of Design Elements

A similar sense of physical potential seemed to inform participants' ratings of the different design elements. For example, a common observation was that charts without joined pieces seem to be "flying apart" or "exploding." Since the pieces in this case are not supporting each other, the scene may be interpreted as being in a state of motion. It is therefore more unstable, complete, and flexible, since the pieces have not yet come to rest. Another possibility is that participants see a joined chart as a natural state, so that separation between pieces implies movement.

Aside from the expected physical simulations pointed to by these comments, another common theme was that different design configurations allowed for different amounts of communication between parts. This theme sheds light on the interaction between filled area and bordered parts. Describing a pie chart and a bubble chart that she rated similarly despite their differing on all dimensions except for filled area and bordered parts, one participant wrote that "they both represent difficulties in communication within the organization." Since the filled areas and bordered parts condition places two types of barrier between or around pieces, it may be seen as allowing less communication between parts and therefore a worse place to work.

A similar kind of analysis based on combined design features is hinted at by a participant who compared two charts with part borders, one with joined parts and one without. This participant described the company with separate parts as being "closed" and the one with joined parts as "not so closed," suggesting that the boundaries around parts are seen as less rigid when parts share borders. This participant went on to describe the company with joined parts as having "a more controlled flexibility."

Finally, the combined effect of filled areas and joined parts was referred to by two participants, although they seemed to offer two entirely different reasons for a similar assessment. One, comparing two similarly rated charts with filled areas and separate parts, wrote, "Both of these charts have so much gray area between the departments." This suggests that filled areas and separate parts are rated as poor places to work because of the presence of a visible barrier between parts. However, another participant who was comparing a chart with a filled area and separate parts (on the left) to one with a filled area and joined parts (on the right) wrote, "The company on the left is off its foundation (the gray circle), whereas the company on the right is centered on its foundation." (Both charts were, in fact, centered within their filled areas.) This offers the alternate possibility that the movement suggested by separated parts may create a sense of precariousness if the user views the filled area as a foundation on which pieces rest. Interestingly, these two sets of comments imply that a chart may lend itself to different interpretations depending on whether the user perceives gravity as moving downwards or moving into the screen; that is, whether she sees herself as looking at a side view or a top-down view of the scene.

The physical interpretations of chart elements provided by participants suggest the potential usefulness of implied dynamics as a framework on which to hang a theory of how design elements contribute to visual structure. However, the apparent differences in how participants interpreted these physical properties suggested in the last example make clear that reliably defining these mappings between visual elements and physical properties is by no means a trivial process.

6.3 The Extent of Elaborations

While users' comments suggested a strong influence of physical simulation on their semantic perceptions of the charts they saw, it could be argued that this influence does not go beyond simple visual organization to affect how participants actually think about the data. Most comments, such as those already quoted, tend to focus on properties that could apply equally well to either visual structure or the more abstract structure of a company; for example, balance, barriers, and movement are all common themes. This relatively direct metaphorical application of visual properties to abstract properties is largely what we expected to find in the user comments.

However, participants' comments often went beyond this simple metaphorical mapping to make elaborate inferences about the company's behavior and management style in terms that did not obviously map to visual properties of the chart. This was especially common in descriptions of companies visualized with a bubble chart, which various participants found "more organic and cooperative," "fun and open," "easy to get along with," and "open source." At the same time, one participant stated that they could not treat the company portrayed in a bubble chart seriously. The waffle chart also elicited a number of emotional responses in the opposite direction, and was described as "bulkier," "organized," and "regimented." These emotional responses to the waffle and bubble charts recall the fact that these were also the types that tended to receive the most extreme ratings on the semantic variables (Table 3). Since these are probably the two least familiar charts presented in our study, these results suggest that extensive elaboration is a greater factor in novel charts, whereas highly conventional representations such as pie charts may require less active interpretation on the user's part.

There were several cases in which these elaborate inferences went so far as to suggest that participants were able to imagine entire stories about the companies based on the simple charts they viewed. In describing a waffle chart, one participant wrote, "Going by the rules is the most important thing in this company and to violate them can get you in serious trouble." Another, describing the similarities between a donut and bubble chart, wrote, "They have some rules but they are mainly focused on encouraging people to do their best in terms of reaching a mutual goal. They don't want to stifle creativity, they want to encourage it." Eight of the 42 participants wrote at least one description that involved this kind of storytelling about the depicted companies. While these usually were based on the chart type, one participant elaborated upon the company based on the presence of a outlined area: "The graph on the left is encompassed by an extra circle—I took this to mean that there was some kind of higher up that kept all the smaller parts in line." That participants made these kind of imaginative elaborations at all, and that there seem to be some commonalities among them, is striking in itself, and suggests that visual structure can in some cases lead not only to minor semantic responses but to full-fledged inferences about the data.

6.4 Color and Weight Perception

A final trend in user comments suggested another factor that we did not directly consider in our initial study design. The colors of individual pieces were randomly generated at the point each chart was displayed, and we did not record these color combinations. However, several comments suggested that the arrangement of color affected how some users perceived the weight and balance of the charts. One participant said that a red segment in one pie chart looked larger than a blue segment in a second pie chart, even though she realized they were the same size. Another participant wrote that "the two largest sections, in brown and muddy green, recede a bit, making them seem more balanced by the red, vibrant blue and pink—that is, the two bigger sections have colors that make them seem less dominant."

These comments may reflect the established finding that darker colors are perceived as being heavier than bright ones [12], and the red color used in our study was in fact the color with the lowest luminance value. The different weights perceived in pieces with different colors may have affected how participants viewed the dynamics of the charts. For example, the comments about color balance may reflect the perceived center of gravity of an object. Also, while describing a donut chart with separated parts, a participant wrote that "the pieces appear to be flying apart, especially the dominating bright red of the largest slice." This implies that the heavier color is perceived as giving extra velocity to the implied movement of the piece. While these comments point towards a role for color in interpreting the dynamics of a visual representation, the fact that we did not record this information unfortunately makes it impossible to interpret this trend in the current study.

7. IMPLICATIONS FOR DESIGN

We have provided theoretical background and experimental evidence that the implied dynamics of a visualization design influence how users interpret the meaning of information. We further argue that these dynamic interpretations guide both inference and a user's perception of how a visual representation can be manipulated. This theory has a number of significant implications for the design of interfaces that use information visualization.

The first implication is that a better understanding of the semantics of design elements would make it possible to exploit their effects as is appropriate to the task. Rather than treating design as decoration or distraction, it could be used consciously to suggest global attributes of the dataset or to communicate interpretations of data in a collaborative context. While this may not be a novel idea to the design community at large, it is less intuitive in information visualization, where design elements such as those used in our study have traditionally been considered irrelevant at best. The fact that the presence or absence of these elements can actively influence interpretations of data, occasionally to a high degree of elaboration, suggest that design cannot be simply ignored or minimized in information visualization contexts. Every design choice is a choice about how the data will be interpreted.

A common theme in our results and user comments is that different design configurations suggested different levels of movement and freedom. The waffle charts and bar charts were seen as rigid and fixed; charts with filled background areas suggested more flexibility. These findings may have practical implications for designing interactive visualizations and other types of interfaces. Charts and visualizations with designs that imply flexibility and movement may suggest more potential points of interaction to the user. Rectangular, space-filling designs (such as treemaps) may be better suited to cases where limited interaction is needed, while a more open design with a clearly delineated space in which pieces can move may encourage users to interact more extensively.

The idea of implied dynamics may in general provide a framework for building intuitive interactions into novel user interfaces. How to make possible interactions in a novel system discoverable is a general problem in human-computer interaction, and is sometimes addressed with the idea of perceived affordances [13]. That an abstract scene such as a software interface may inherently present a set of implicit physical properties and forces suggests that designers could visually play up certain dynamics in order to guide a user's understanding of what can be done. For example, items that can be moved can be made to appear lighter, while items that cannot be changed can be made to look heavy and rigid.

Implied dynamics of visualizations may also constrain design in certain ways. Many of the negative comments made about the bubble chart and the various exploded charts (that is, those without joined parts) suggest that the lack of support and apparent violation of perceived gravity in these visualizations makes them seem chaotic and disorganized. This may or may not be a problem for a given task, but in the information visualization context it may have the undesirable effect of distracting from whatever organization or structure the data actually has. If implied dynamics are indeed an important part of visualization perception, then any visual representation that strongly implies motion, such as an exploded chart, should be used with care.

Another practical use for our findings is that they suggest certain visualization types which are more or less susceptible to the effects of design. In general, we found that the design elements we used had effects that were consistent across a variety of chart types, suggesting that they have basic effects on perception independent of their context. However, we also found a non-significant trend for ratings of the bubble chart to be more design-sensitive and those of the pie chart to be more robust to design changes. This may be an effect of the relative familiarity of the two chart types, or it may speak to a more essential difference between their structures. Studying the design robustness of a given visualization would provide valuable information in the evaluation process and help in making decisions about visualization use.

8. CONCLUSION AND FUTURE WORK

This work serves as a first step in the process of understanding how elements of visual structure can influence semantic interpretations of data. We have provided evidence that these influences can be reliable and to an extent predictable. The semantic impact of structural elements must be taken into account when designing and evaluating visualization methods. By focusing solely on a user's ability to read data points from a visualization, without considering how she views the dataset as a whole based on elements of the visual design, we may be missing the forest for the trees.

Ultimately, our goal should be to model these effects within a theoretical framework of how people perceive visual structure as information structure. Implied dynamics and the other physical properties frequently described by participants may be a useful basis for such a framework. We have begun studying the viability of implied dynamics as a model of visualization use by attempting to replicate the implied motion findings from cognitive science [7] in a visualization context, and preliminary results suggest that items in a visualization may indeed show an effect of gravity on the perceived motion of items in a scatterplot. However, much more work is needed to establish a set of correspondences between visual design, the dynamic properties it affords, and the semantic inferences these cause users to reach.

Finally, this work uncovered major gender effects in responses to visual structure. These gender differences are complex and difficult to interpret, going beyond differential spatial skills into a less well-understood area. There were no clear patterns in the explanations that shed light on these gender effects. Further study should be done in order to understand why different user groups may respond to visual structure differently, and how these differences should affect visualization research and design.

Interpretations of data in a visualization can go beyond the simple reading of data, and the semantic implications of design elements guide these interpretations in consistent ways. Although the effects of design elements may seem undesirable in a pure information visualization context, understanding their nature and how to predict them will ultimately give us the tools to control them.

9. REFERENCES

[1] Amazon Mechanical Turk, http://aws.amazon.com/mturk/. Retrieved 10-16-2008.

[2] R. Arnheim. *Art and Visual Perception: A Psychology of the Creative Eye*. UC Press, 1974.

[3] R. Arnheim. *The Power of the Center: A Study of Composition in the Visual Arts*. UC Press, 1988.

[4] C. Brewer and M. Harrower. ColorBrewer, http://www.colorbrewer.org/, 2002. Retrieved 12-4-2008.

[5] L. Elting, C. Martin, S. Cantor, and E. Rubenstein. Influence of data display formats on physician investigators' decisions to stop clinical trials: prospective trial with repeated measures. *BMJ*, 318:1527–1531, 1999.

[6] J. Freyd. Dynamic mental representations. *Psychological Review*, 94(4):427–438, 1987.

[7] J. Freyd, T. Pantzer, and J. Cheng. Representing statics as forces in equilibrium. *Journal of Experimental Psychology*, 117:395–407, 1988.

[8] J. J. Gibson. *The Ecological Approach to Visual Perception*. Houghton Mifflin, 1979.

[9] T. Hubbard. Environmental invariants in the representation of motion: Implied dynamics and representational momentum, gravity, friction, and centripetal force. *Psychonomic Bulletin and Review*, 2(3):322–338, 1995.

[10] G. Lakoff and M. Johnson. *Metaphors We Live By*. University of Chicago Press, 1980.

[11] G. Lohse and E. Johnson. A comparison of two process tracing methods for choice tasks. *Organizational Behavior and Human Decision Processes*, 68:28–43, 1996.

[12] M. Monroe. The apparent weight of color and correlated phenomena. *The American Journal of Psychology*, 36(2):192–206, 1925.

[13] D. Norman. Affordance, conventions, and design. *Interactions*, 6(3):38–43, 1999.

[14] D. Richardson, M. Spivey, L. Barsalou, and K. McRae. Spatial representations activated during real-time comprehension of verbs. *Cognitive Science*, 27:767–780, 2003.

[15] J. Siskind. Grounding the lexical semantics of verbs in visual perception using force dynamics and event logic. *Journal of Artificial Intelligence Research*, 15:31–90, 2001.

[16] B. Tversky, J. Zacks, P. Lee, and J. Heiser. Lines, blobs, crosses and arrows: Diagrammatic communication with schematic figures. In *Theory and Application of Diagrams*, pages 221–230. Springer-Verlag, 2000.

[17] C. Ware. *Visual Thinking for Design*. Morgan Kaufmann, 2008.

[18] J. Zacks and B. Tversky. Bars and lines: A study of graphic communication. *Memory & Cognition*, 27(6):1073–1079, 1999.

[19] C. Ziemkiewicz and R. Kosara. The shaping of information by visual metaphors. *IEEE TVCG*, 14(6):1269–1276, 2008.