Value-Driven Evaluation of Visualizations

John Stasko
School of Interactive Computing
85 5th St., NW
Georgia Institute of Technology
1-404-894-5617
stasko@cc.gatech.edu

ABSTRACT
Existing evaluations of data visualizations often employ a series of low-level, detailed questions to be answered or benchmark tasks to be performed. While that methodology can be helpful to determine a visualization's usability, such evaluations overlook the key benefits that visualization uniquely provides over other data analysis methods. I propose a *value-driven evaluation* of visualizations in which a person illustrates a system's value through four important capabilities: minimizing the time to answer diverse questions, spurring the generation of insights and insightful questions, conveying the essence of the data, and generating confidence and knowledge about the data's domain and context. Additionally, I explain how interaction is instrumental in creating much of the value that can be found in visualizations.

Categories and Subject Descriptors
H.5.2 [**Information Interfaces and Presentation, (e.g., HCI)**]: User interfaces – *Evaluation/Methodology*.

General Terms
Measurement, Design, Human Factors, Theory.

Keywords
Data visualization, value, evaluation, interaction.

1. INTRODUCTION
The topic of evaluation (broadly considered) has risen in prominence in the data visualization research community over the past few years. To some degree, the novelty of creating new techniques and systems has worn off, or at least researchers have broadened their views beyond just creating new visualization techniques. The community has become more reflective and introspective, and thus evaluation seems to be a topic on everyone's minds currently.

What does "evaluation" mean in the context of visualization, however? What are we evaluating? For what purposes are we evaluating? I think these questions are more subtle than one would immediately surmise, and the answers are nuanced.

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

BELIV'14, November 10, 2014, Paris, France.
Copyright is held by the author. Publication rights licensed to ACM.
ACM 978-1-4503-3209-5/14/11 …$15.00.
http://dx.doi.org/10.1145/2669557.2669579

One potential angle of evaluation is to improve the techniques and systems one builds. That is, a developer of a new system should evaluate it and find the embedded problems and faults in order to help improve the system and make it better. This activity can be iterated repeatedly and is a fundamental component of formative evaluation that one encounters in the area of human-computer interaction.

A second style of evaluation is to compare two specific approaches to each other. Is technique A or technique B a better approach for a given problem? The specificity of this type of evaluation is appealing in many ways, but it is frequently quite difficult to conduct such an evaluation in the field of data visualization because very few techniques or systems are built for the exact same purpose, domain, and the same type of data. Examples of this type of evaluation do exist [14,29] but it is typically very difficult to "compare apples to apples."

Another angle of evaluation is more general than these first two. When a person develops a new visualization technique or system, there is a fundamental desire to determine whether it is any *good*. Simply put, is the technique useful and beneficial? Researchers want to show the value and utility of their new ideas, and they seek methods to answer those questions.

This notion of showing or identifying utility and value is one way to slightly recast the challenge of evaluation in data visualization. Rather than thinking about evaluating techniques and systems, one could focus on identifying the potential *value* of a system. Hence, the evaluation of a technique or system becomes a process where one identifies and illustrates its value.

I also believe that the notion of identifying value has a broader scope than evaluation does. Researchers in other fields probably will not be so interested in the evaluation of a particular visualization system. Conversely, showing the value of a visualization technique can be extremely important at a more general scale. Many approaches to analyzing and presenting data exist; visualization is just one of many approaches available to people and organizations today. Identifying the value of visualization better is vital to educating people not so familiar with the domain and convincing them of the potential impact that a visualization approach can achieve.

In the remainder of this article, I explore why the process of "evaluation" is more problematic in visualization research than in human-computer interaction research, on the whole. I recast the evaluation of a visualization as identifying its value, and I develop a qualitative formula for determining a visualization's value. Although the equation is more descriptive than prescriptive, it should help researchers to more accurately and objectively assess and identify the value of their new ideas, and it should help to more specifically communicate visualization's utility to those people less familiar with the area. Finally, I explain the fundamental importance of interaction in the value equation even though it is still the less understood half of data visualization.

2. BACKGROUND

After ten years of research focused on technique and system development, the field of information visualization began to consider issues involving evaluation more deeply around 2000. That year, a special issue of the journal *International Journal of Human-Computer Studies* [7] was devoted to empirical studies of information visualizations. Subsequently, the initiation of the BELIV workshop in 2004, whose focus is solely on evaluation in visualization, further stimulated research into this topic. Plaisant's keynote address there and accompanying article [21] outlined a number of challenges in evaluating information visualization including the need to match tools with users, tasks, and real problems, as well as difficulties in improving user testing. She called for repositories of data and tasks and a broader notion of evaluation including case studies and success stories. Scholtz [26] later argued for a metrics-driven evaluation methodology in visual analytics that moves beyond usability and includes evaluations of situation awareness, collaboration, interaction, creativity, and utility.

A number of different methodological approaches have been proposed for visualization evaluation. The insight-based evaluation methodology [25] attempts to count and document the insights gained during an analysis session. An insight is defined as an individual observation about the data by a participant, a unit of discovery. The MILC evaluation methodology [28] advocates performing long-term, multidimensional case studies of deployed systems in real use outside the laboratory. This approach follows a more ethnographic style and centers on observational evidence about the utility of a system.

Following a 2007 Dagstuhl seminar on information visualization, Carpendale authored a book chapter [5] focusing on many issues involving visualization evaluation. She described different methodological approaches that could be taken, along with the benefits and challenges inherent in each. Based on a meta-analysis of over 800 research articles, Lam et al. [15] describe seven canonical scenarios of empirical studies in information visualization. For each scenario, the authors describe goals and outputs, evaluation questions, and they provide methods and examples. A key notion of their work, also emphasized in this article, is understanding the purpose of an evaluation. Subsequently, Isenberg et al. [12] performed a similar study for scientific visualization articles.

Moving beyond evaluation, other researchers have carried out research explicating the value of visualization to cognition. Larkin and Simon [16] describe how visualization helps people think by acting as a temporary storage area (extra memory), by transforming cognitive challenges into perceptual actions, and by engaging people's strong pattern matching abilities, among other benefits. Norman [18] discusses the value of visualization in helping people to think and explains how external representations aid information access and computation. He specifically describes the importance of matching representations to the task at hand in order for visuals to provide the best cognitive value. Card, Mackinlay, and Shneiderman's book of collected information visualization readings [4] contains a first chapter with an extensive discussion on the benefits of visualization. The editors present a model of knowledge crystallization and list ways that visualization helps amplify cognition, such as by reducing searches for information, enhancing the recognition of patterns, and enabling perceptual inference operations.

With respect to analyses of visualization focusing on its value to people, van Wijk [33] presents a more formal, theoretical model of value. He first develops an equation for the central process in visualization, which describes time-varying images that are dependent upon the data and the visual specification. Knowledge is then defined by a person's initial knowledge plus the knowledge gained at each step as the person perceives the visualization image. Knowledge can be affected by interactive exploration as well. Next, he describes the costs of visualization which include initial development costs, initial costs per user and session, and perception and exploration costs. Finally, his economic model defines that the total (cognitive) profit is equal to the value minus those costs. Fekete et al.'s book chapter on the value of visualization [9] articulates multiple ways that visualization assists data understanding. The authors discuss cognitive and perceptual support, further describe van Wijk's economical model of value, and provide multiple "success story" examples that help illustrate visualization's utility.

3. VALUE-DRIVEN EVALUATION

In the field of human-computer interaction, evaluation often equates to assessing the usability of a system or interface. Through sessions with a series of benchmark tasks to be performed, evaluators determine whether potential users can employ a system to achieve desired results. That is, an evaluator seeks to learn whether people can effectively and efficiently use a system to successfully perform tasks. Such evaluation is a key aspect of user-centered design and it has become pervasive throughout software development, particularly for interactive systems.

Unfortunately, this type of benchmark task-focused evaluation is simply not sufficient for evaluating visualization systems, even though it is often what one observes when reading research papers within the discipline. This type of evaluation can help determine whether a visualization is learnable and comprehensible, and whether a person can use it to answer specific questions about a data set. However, this approach fails to examine some of the larger benefits of visualization that go beyond just answering specific data-centered questions. Visualization should ideally provide broader, more holistic benefits to a person about a data set, giving a "bigger picture" understanding of the data and spurring insights beyond specific data case values.

To help understand this premise, consider a data set of information about different car models and each car's attributes such as price, miles per gallon, horsepower, and so on. One could build a visualization of this data set and "evaluate" the visualization by conducting a user study where participants answer specific questions (i.e., perform benchmark tasks) about the data such as

- Which cars have the best miles-per-gallon?
- How much does a Ford Taurus cost?
- Is there a correlation between car weight and torque?
- How many different countries manufacture cars?
- What is the range of car's horsepower?

While it is important for the visualization system to support answering such questions thus illustrating that it is comprehensible and usable, a person could relatively easily gain those answers via a spreadsheet, query language, or statistics package as well.[1] Support for answering those types of questions is not a unique benefit of visualization. An evaluation that focuses only on these types of tasks simply fails to adequately assess the primary and unique benefits of visualization that go beyond simple data-driven queries.

What are the broader benefits that visualization provides? This question is at the heart of the approach I advocate for evaluating visualization systems. It involves a fundamental examination of the benefits that are more unique and specific to visualization, compared to other types of data analysis. The approach that I advocate focuses on identifying the *value* that a visualization provides. This value goes beyond the ability to support answering questions about data—it centers upon a visualization's ability to convey a true *understanding* of the data, a more holistic broad and deep innate sense of the context and importance of the data in "the big picture."

To be more specific, I have developed a simple equation that characterizes the value (**V**) of a visualization:

$$\mathbf{V = T + I + E + C}$$

Below, I elaborate on each of the four components of the value equation and then provide two examples to show how it applies to existing visualizations.

The **T** in the value equation represents:

> A visualization's ability to minimize the total **time** needed to answer a wide variety of questions about the data.

Effective visualizations allow a person to identify many values from a data set and thus answer different questions about the data simply by viewing the visualization or, at least, by interacting with the visualization and inspecting the resulting views. Rather than have to learn query languages and tools and issue syntactic requests, a person merely interacts with an interface using direct manipulation to select visual items or update the view to show the desired information. Effective visualizations also excel at presenting a set of heterogeneous data attributes in parallel. They make accessing both quantitative and nominal data take the same relatively low amount of effort.

The types of "low-level" questions about data that visualization can assist with have been described in earlier research [1,31,35]. They include tasks such as retrieving values, finding extrema, characterizing distributions, and identifying correlations, among others.

The **I** in the value equation represents:

> A visualization's ability to spur and discover **insights** and/or **insightful questions** about the data.

Effective visualizations allow a person to learn about and make inferences from a data set that would be much more difficult to achieve without the visualizations. This notion of knowledge gained goes beyond simply reading relevant data values. It often involves knowledge and inferences that are acquired by viewing combinations, alignments, aggregations, juxtapositions, and particular unique representations of the data.

The notion of insight has been a key component of information visualization for many years. The insight-based evaluation methodology [25] by Saraiya, North, and Duca, discussed earlier, provided an initial formal notion of insight. North [19] defined an insight as an individual observation about the data by the participant, a unit of discovery, and characterized each as being complex, deep, qualitative, relevant, and unexpected. This characterization, in particular the perception of insights being unexpected, seems to align with a notion of spontaneous insights, the "Aha!" moments when people finally see solutions to problems.

I do not require the notion of unexpectedness for the **I** component of the value equation. Instead, my view of insight aligns with that described by Chang et al. [6], and refers to the acts of knowledge-building and model-confirmation. Insight is like a substance that people acquire with the aid of a system. It is a by-product of exploration, not a task-driven result of analysis [38].

The **E** in the value equation represents:

> A visualization's ability to convey an overall **essence** or take-away sense of the data.

Effective visualizations allow a person to gain a broad, total sense of a potentially large data set, beyond what can be gained from learning about each individual data case and its attributes. In essence, to borrow from a well-known phrase, "the whole (understanding) should be greater than the sum of the parts". An effective visualization should convey information about the totality of data being presented, effectively the "big picture."

The notion of providing an overall sense of a data set is often implemented through visualization design principles such as "overview and detail" or "focus + context." Shneiderman's oft-repeated information visualization mantra [27], "Overview first, zoom and filter, details on demand", begins with an overview. The **E** component of the value equation goes beyond supplying an overview of the data set, however. A visualization with high utility should convey an overall essence of the data, its unique and most important characteristics, and the primary knowledge to be gained from the data.

The **C** in the value equation represents:

> A visualization's ability to generate **confidence**, knowledge, and trust about the data, its domain and context.

Effective visualizations allow a person to learn and understand more than just the raw information contained within the data. They help promote a broader understanding of the importance of the data within its domain context. Furthermore, effective visualization promotes a viewer's sense of confidence, and thus trust, of the data.

The knowledge and sense of confidence or trust can apply both to the visualization creators and the visualization consumers. Visualizations, and more specifically the visualization construction process, are a wonderful way to identify embedded problems in a data set such as missing, erroneous, or incomplete values. Similarly, viewing a data set with an effective visualization can highlight areas where more data may be needed

[1] Of course, each of these tools has a learning curve, but once understood they can be effective in this role.

or can signify the importance of adding and exploring other data sets.

This equation for identifying the value of a visualization is still just a qualitative metric. I have not developed quantitative measures of each of the four key components, nor the factors to determine how much each component contributes to the overall value. As such, the equation serves more as a *descriptive* aid than a *prescriptive* one.

One could imagine a variety of mechanisms for more precisely determining the benefits of a visualization toward each of these four components. The most basic method would be providing a thorough and clear explanation of a visualization, possibly including demonstrations of use through videos or direct hands-on interaction. A second method could be offering descriptions of detailed scenarios of use that illustrate how to apply the system in context and how it would be beneficial to potential relevant users. Going even further, the developers of the visualization could deploy it in the field and then report on its trial usage for an extended period of time. Furthermore, experiments and user studies that go beyond employing simple sets of benchmark tasks could be useful for helping to determine value.

In the remainder of this section, I examine two example visualization applications and illustrate how each provides value via the introduced equation. Here, I simply describe how each provides the four constituent components, but future work can take on the task of making these assessments even more descriptive and precise.

3.1 Map of the Market

The first example, shown in Figure 1, is the Map of the Market tool (http://www.marketwatch.com/tools/stockresearch/marketmap) that uses the treemap visualization technique to portray companies' performance on the stock market. It includes interface controls to search for a particular company and to adjust the period being shown, for example, showing year-to-date versus daily performance. The original version of this visualization was created by Wattenberg [34] for the smartmoney.com website. It used more saturated colors than the current version and provided more interactive controls for showing top gainers and losers.

T: The Map of the Market supports rapidly answering a variety of questions about stocks' performance simply by mousing over an item, clicking on it, or searching for a company. The viewer can learn how particular companies have performed, how bigger or smaller companies have performed, which companies have gained or lost significant amounts, how different sectors have performed, and so on. Answering these questions is quick and easy because it requires observation and simple interface actions. Note, however, that not every type of query can be answered easily. For instance, it may be difficult to determine which company has gained or lost the most because multiple companies will exhibit similar shades of a color.

I: The visualization supports the development of insights and insightful questions as the viewer notices patterns, trends, or outliers. For instance, one sector may have performed well during the period but another poorly, or perhaps large companies performed well while small did not. The viewer may wonder why such a result has occurred. By comparing the views of year-to-date and the prior 52 weeks performance, a person may notice different companies whose stock is trending upward or downward.

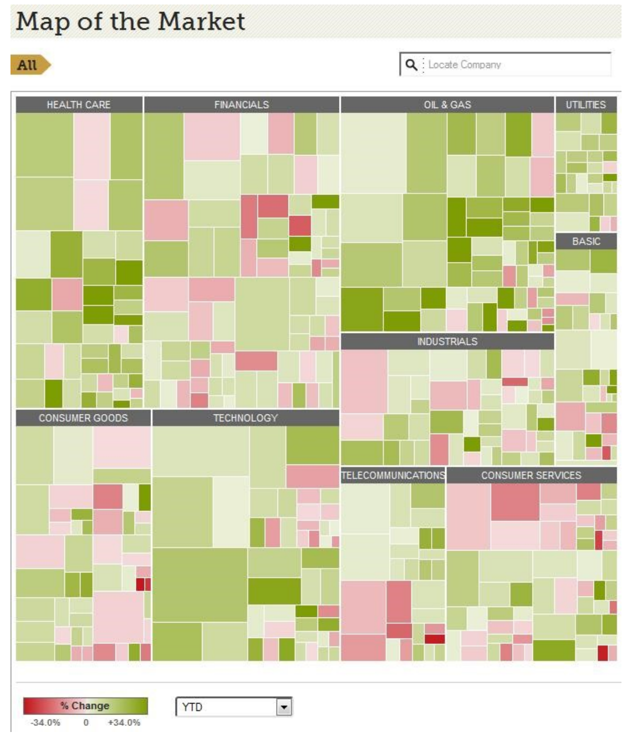

Figure 1. The Map of the Market showing stock performance (http://www.marketwatch.com/tools/stockresearch/marketmap)

E: Beyond the individual company performances, viewing the complete treemap gives a person a broad perspective on how the market has done as a whole. By observing the ratio of green to red cells, one can estimate the general performance. Similarly, on days when the market has a tremendous gain or crash, it communicates an overwhelming view of one bright color. Outliers also can emerge from these types of views. For example, during the 777 point drop on September 29, 2008, the former smartmoney.com site showed an overwhelming view of bright red except for one small green cell in the Basic Materials sector, a gold stock.

C: Viewing the market visualization from different periods gives a person a sense of the level of up and down swings that stocks typically make. Observing the sectors with the largest rectangles communicates where the largest companies reside, and thus informs the viewer about American companies on the whole. Effectively, this visualization goes beyond simply communicating stock performance data. It also begins to inform the viewer about companies, markets, and the economy.

3.2 CiteVis

The second example, shown in Figure 2, is the CiteVis application (http://www.cc.gatech.edu/gvu/ii/citevis/) that presents citation data about papers appearing at the IEEE Information Visualization Conference. The application was created by my students and me to help researchers learn about and understand the histories of citations to and among these articles [30]. It represents the papers for each year as a row of circles and uses interaction to show specific paper-to-paper citations rather than drawing edges between them. The application also allows a viewer to search for papers from specific authors or about specific topics.

determined. However, it shows how the visualization helped develop an insightful question. We speculate that many of these observations and insights would have been difficult to recognize simply by accessing the data through a query interface or tables and spreadsheets containing all the data values.

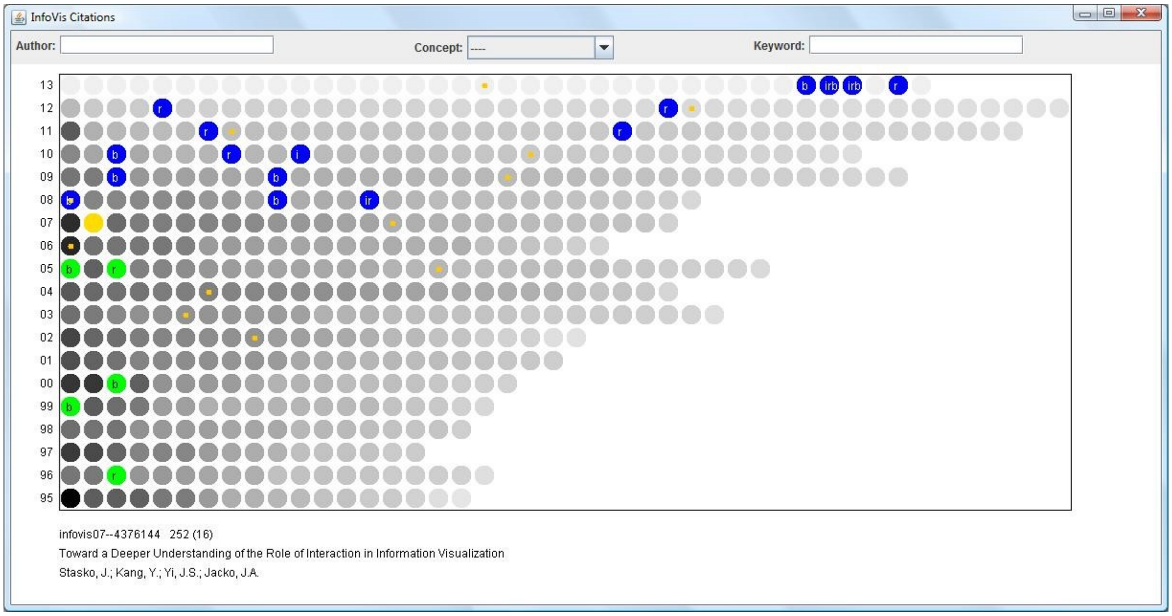

Figure 2. The CiteVis application showing InfoVis paper citations (http://www.cc.gatech.edu/gvu/ii/citevis/)

T: Just by viewing the CiteVis application, a person can quickly answer questions and learn about the number of papers each year and how the number of citations change going back in time. Furthermore, through simple interactions, many more questions can be answered and information acquired. By mousing over a paper's circle, the viewer learns its title, authors, number of total and internal (within the conference) citations that the paper has gained, and the internal papers it cites and is cited by. Through the menus and search fields at the top, viewers can observe the papers of a particular author or on a particular topic. Mousing over the darkest circles tells the viewer which papers have been cited the most. Observing the small yellow dots informs the viewer how the Best Paper Award winners have performed in terms of citations. Again, it is important to stress that no query language need be learned to acquire all this different information. Direct observation coupled with simple interface interactions answer the viewer's questions and communicate a wide variety of aspects of this data set very rapidly and easily.

I: The CiteVis application helps a person learn insights about the data. For example, just by examining the default view, one can notice the four highly cited (dark circles) papers from the year 2000 or the large set of more highly cited papers in 2004. Clearly, these were two good years for impactful, high quality articles. A viewer can notice how the Best Paper winners are not always among the most highly cited papers from a year. Issuing queries to spotlight papers mentioning "evaluation" or "user study" in their abstract highlights how the early years of the conference had relatively little work on that topic, but it has grown strongly since. When we explored the visualization one day, we noticed how many of the least cited papers to the right side of the view seemed to be the application/design study style of papers. We were curious if this is actually the case, but because the data set does not include the paper type, that answer cannot be conclusively determined. However, it shows how the visualization helped develop an insightful question. We speculate that many of these observations and insights would have been difficult to recognize simply by accessing the data through a query interface or tables and spreadsheets containing all the data values.

E: The CiteVis visualization conveys a global sense of how the conference has grown and how citation patterns change over time. There are fairly recent papers that have high citation counts, so it is not only older papers with that characteristic. The visualization conveys the essence of the data set, its overall trends, patterns, and peculiarities.

C: To create the CiteVis application, we gathered data from online digital libraries and we manually examined the papers from each year and logged relevant information. Upon attempting to run the visualization for the first time, we quickly learned of missing papers and faulty data values because the system would crash upon certain interactions. Similarly, we observed papers that supposedly cited a paper from the future, which clearly cannot happen and again pointed out errors in the data. Beyond inspiring more confidence in the data by highlighting errors, the visualization also informs the viewer about the field of information visualization itself, what subtopics are important, which authors have written high-impact papers, and how the field is growing. The visualization communicates knowledge about the domain and its context beyond simple citation counts on papers.

4. THE IMPORTANCE OF INTERACTION

The previous section explains how value is manifest in visualization through the identified four components. In this section, I dig deeper to uncover how that value is frequently generated. To begin, let us consider different design paradigms commonly employed in visualization today. Suppose that an organization has a data set with multiple variables (attributes) per data case. To create a visualization to help explore, analyze, and communicate that data, three fundamental visualization design paradigms could be used.

1. Create a rich visualization employing many graphical objects and a diverse set of visual properties that completely represents the entire data set. Effectively, this design packs all the data into one, usually complex, view or representation.
2. Provide multiple, coordinated views of the data set, where each view provides a different perspective on the data, perhaps just focusing on certain attributes of the data.
3. Represent only some of the data cases and/or attributes of the data set initially, but employ interaction to allow the viewer to change the cases, attributes, and representation being shown.

The first scenario is very common today and is exemplified by infographics so pervasive throughout print media and the web. These designs are static and often physically large, in order to include all aspects of the data set being shown.

As data sets become larger with more variables per case, it becomes increasingly difficult to rely on the first design scenario, however. There simply may be too much to show, or the representation becomes so complex that it is difficult for people to interpret. Thus, visualization designers often move to using the second or third design paradigms, or perhaps a hybrid of the two.

Both the Map of the Market and CiteVis applications discussed in the previous section are examples of the third paradigm. The key differentiator in this design is the use of interaction in the visualization. Interaction differentiates these types of visualizations from static infographics. Interaction, when applied well, allows a person using a visualization to engage in a rich dialog with the data, to effectively conduct a fruitful conversation including a variety of questions and answers. As a person considers aspects of the data and forms new schemas and models of the data's meaning, interaction allows the person to change perspective and address new inquiries that arise.

Representation and *interaction* are often seen as the two key components of visualization, but I believe it is fair to say that representation has received the majority of attention throughout the visualization research community. Interaction, while important, has not received as much focus. As discussed above, many infographics, which some people (erroneously) equate to data visualization, do not even include interactive capabilities.

What, therefore, is the role of interaction in visualization? Is it an equal peer with representation, playing just as important a role in providing value for analysis and presentation? Alternately, is interaction subordinate to representation in the visualization equation, serving merely as a facilitating mechanism to move between different representations? This is an interesting philosophical question. I tend to view interaction as a significant and vital component in the equation, playing an equally important role in providing effective, valuable understanding of data to people.

To explore interaction's role in the value equation more closely, it is helpful to examine the many different aspects of interaction. When a person interacts or seeks to interact with a visualization, s/he does so for a purpose and to gain further understanding. Thus, one important component of interaction is the user's *intent*, the reason why they interact. Yi et al. [37] studied over a hundred visualizations and examined how interaction functioned in each. They developed a list of seven types of intent that covered the vast majority of uses of interaction. They found that systems provide and people interact with visualizations to select, explore, reconfigure, encode, abstract/elaborate, filter, and connect.

Each of these intents facilitates knowledge acquisition from a visualization and thus plays a key role in the value equation. As shown by the two example systems in section 3, interaction is often crucial to simplified, fast answers to varieties of questions (**T**) that can be posed to a visualization. Interaction likewise enables the generation of insights and insightful questions (**I**) because new views of the data provide new mental perspectives for a person. Interaction may play a less important role in conveying the essence of a data set (**E**), but as data set sizes increase, a visualization simply may not be able to provide an adequate summary or overview, and hence may require interaction to communicate the "big picture." Finally, interaction facilitates confidence about the data, its context, and domain (**C**) by allowing a person to explore all dimensions and aspects of the data.

While intent is one aspect of interaction, how that intent is carried out operationally is yet another. That is, the set of interactive operations provided by a system also contributes fundamentally to its value. Prototypical lists of such operations are discussed in many articles and books [10,13,27] and include selecting items for identification and manipulation, navigation such as zooming, filtering items or values, and coordinating views, among others. Curious whether these operations are actually manifest in real systems, my doctoral student Charles Stolper recently conducted a survey of hundreds of visualization applications and systems found online. He logged the types of interactive operations each provides and found that the following small set of operations are pervasive and comprise the majority of interactions provided by the systems:

- Selection to access data details, including mouse-over "tooltip" or "balloon"-style details
- Navigation, including panning and zooming
- Brushing and linking across views or representations to show connections and relationships

We were surprised that these three interaction operations so dominate interactive visualization applications. We expected more diversity in systems because research has demonstrated many new types of interactive operations. Thus, we wondered if this lack of diversity means that either a) practitioners only feel that this small set of operations are truly useful or b) the visualization libraries and toolkits used to build the systems implicitly limit interaction to these types of operations.

Nonetheless, this observation shows that there is room to innovate on interaction within visualization. For example, the use of lenses or toolglasses [3] to provide different views of data is an old idea that is rarely seen in visualization applications. Other systems, such as Dust and Magnet [36] and Kinetica [22], represent data cases as physical objects that can be dragged and pushed to expose their underlying attributes. Furthermore, interaction can be used to implement semantic operations within a domain, such as the OnSet system's use of drag-and-drop style interaction to perform union and intersection operations on set-typed data [23].

Yet another aspect of interaction concerns the interface mechanisms facilitated by existing hardware used to implement interaction operations. The vast majority of visualization applications run on desktop computers with a keyboard, mouse, and monitor(s). They provide WIMP (window, icon, menu, pointer)-style interfaces with many controls and (often small) graphical objects representing the data. Hence, the use of a mouse is crucial and advantageous to manipulate these detailed controls and objects. Today, as more people transition to using tablets where finger or pen-touch are the only modes of interaction, visualization interfaces must transition to these modes as well [17]. A few research projects have explored visualization interaction via pen and finger touch [2,11,24], and they illustrate how new ideas are necessary to continue to provide value under this different hardware platform.

The goal of this section was to illustrate how vital interaction is to the value provided by visualization systems, but also how interaction is still the less understood half of the overall contribution. The research community has recently been making strides toward understanding interaction better. For instance, Elmqvist et al. [8] discuss the importance of interaction in visualization, focusing primarily on how fluid interaction can engage a person in a type of cognitive flow that enhances analysis and productivity. Furthermore, we need to move toward a science of interaction [20] including theories and testable premises that explain more fully how interaction contributes value to people's use of visualization systems. Fundamentally, interaction remains a relatively under-utilized component of visualization for providing greater analytical value to people.

5. CONCLUSION

In this article, I propose a new style of value-based evaluation of visualization techniques and systems. This approach hinges upon evaluators identifying the *value* that a visualization provides, beyond its ability to provide answers to basic queries about the data it represents. The technique defines a visualization's value through four constituent capabilities: T – minimizing the time to answer diverse questions about the data, I – spurring the generation of insights and insightful questions about the data, E – conveying the overall essence of the data, and C – generating confidence and knowledge about the data's domain and context.

This value-based evaluation methodology is presently still a qualitative, descriptive approach rather than a quantitative and more prescriptive technique. One might argue that for the approach and the value equation to become practical tools for evaluating visualizations' value, they need more of a quantitative characterization. While that is true at a pure sense, the equation as-is still provides a framework for visualization researchers to communicate the value of their creations, and it provides four dimensions that can serve as goals for researchers developing new evaluation methodologies. For instance, what type of evaluation could measure the ability of a visualization to convey the essence or take-away sense of a data set?

My approach in this article has largely been utilitarian, focusing on the analytic value of visualizations. Visualizations also can provide aesthetic impact and value [32]. Future work might explore bringing this notion to the value equation as well.

This article also advocates further study and use of interaction to provide even more value with visualization. Creative new methods of interaction from the research community have yet to be widely used in practice, and new platforms such as touch-based screens and tablets require new types of interactive operations.

Visualization goes beyond simply answering specific questions about a data set. It also helps to generate insights and insightful questions about the data that would be difficult to identify so easily using other methods of analysis. Visualization additionally conveys the "big picture" essence of a data set, and it facilitates generating confidence and learning about the domain of the data and its context. We need evaluation methodologies that focus on identifying the most beneficial, unique capabilities that visualization provides.

6. ACKNOWLEDGMENTS

This article was motivated by my EuroVis '14 Capstone address (http://vimeo.com/98986594). I thank my students for discussions helping in the development of these ideas. The research has been supported by the US Department of Homeland Security's VACCINE Center under Award Number 2009-ST-061-CI0001, DARPA's XDATA program, and the National Science Foundation via awards CCF-0808863, IIS-0915788, and IIS-1320537.

7. REFERENCES

[1] Amar, R., Eagan, J., and Stasko, J. 2005. Low Level Components of Analytic Activity in Information Visualization, In *Proceedings of the IEEE Symposium on Information Visualization* (Minneapolis, MN, October 2005). InfoVis '05. 111-117.

[2] Baur, D., Lee, B. and Carpendale, S. 2012. TouchWave: Kinetic Multi-touch Manipulation for Hierarchical Stacked Graphs. In *Proc. of ACM Interactive Tabletops & Surfaces* (Cambridge, MA, November 2012). ITS '12. 255-264.

[3] Bier, E.A., Stone, M.C., Pier, K., Buxton, W., and DeRose, T.D. 1993. Toolglass and magic lenses: the see-through interface. In *Proceedings of the ACM Conference on Interactive Graphics* (Anaheim, CA, August 1993). SIGGRAPH '93. 73-80

[4] Card, S., Mackinlay, J., and Shneiderman, B. 1999. *Readings in information visualization: Using vision to think*. Morgan Kaufmann, San Francisco, CA.

[5] Carpendale, S. 2008. Evaluating Information Visualizations, in *Information Visualization: Human-Centered Issues and Perspectives*, A. Kerren, J. Stasko, J.-D. Fekete, C. North, Eds. Springer LNCS, Berlin, Heidelberg, Germany, 19-45.

[6] Chang, R., Ziemkiewicz, C., Green, T.M., and Ribarsky, W. 2009. Defining Insight for Visual Analytics. *IEEE Computer Graphics & Applications*, 29, 2 (March/April 2009), 14–17.

[7] Chen, C. and Yu, Y. 2000. Empirical studies of information visualization: a meta-analysis. *International Journal of Human-Computer Studies*. 53, 5 (November 2000), 851-866.

[8] Elmqvist, N., Vande Moere, A., Jetter, H-C., Cernea, D., Reiterer, H., and Jankun-Kelly, TJ. 2011. Fluid interaction for information visualization. *Information Visualization*, 10, 4 (October 2011), 327-340.

[9] Fekete, J.-D., van Wijk, J., Stasko, J., and North, C. 2008. The value of information visualization. In *Information Visualization: Human-Centered Issues and Perspectives,* A. Kerren, J. Stasko, J.-D. Fekete, and C. North, Eds. Springer LNCS, Berlin, Heidelberg, Germany, 1-18.

[10] Heer, J. and Shneiderman, B. 2012. Interactive dynamics for visual analysis. *Communications of the ACM.* 55, 4 (April 2012), 45-54.

[11] Isenberg, P. and Isenberg, T. 2013. Visualization on Interactive Surfaces: A Research Overview. *i-com*, 12, 3 (November 2013). 10-17.

[12] Isenberg, T., Isenberg, P., Chen, J., Sedlmair, M., and Möller, T. 2013. A Systematic Review on the Practice of Evaluating Visualization. *IEEE Transactions on Visualization and Computer Graphics*. 19, 12 (December 2013), 2818-2827.

[13] Keim, D. 2002. Information Visualization and Visual Data Mining. *IEEE Transactions on Visualization and Computer Graphics*. 8, 1 (January-March 2002), 1-8.

[14] Kobsa, A. 2001. An Empirical Comparison of Three Commercial Information Visualization Systems. In *Proceedings of the IEEE Information Visualization Symposium* (October 2001). InfoVis ’01. 123-130.

[15] Lam, H., Bertini, E., Isenberg, P., Plaisant, C., and Carpendale, S. 2012. Empirical Studies in Information Visualization: Seven Scenarios. *IEEE Transactions on Visualization and Computer Graphics*. 18, 9 (September 2012), 1520–1536.

[16] Larkin, J. and Simon, H.A. 1987. Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science*. 11, 1 (January-March 1987), 65-99.

[17] Lee, B., Isenberg, P., Riche, N.H. and Carpendale, S. 2012. Beyond Mouse and Keyboard: Expanding Design Considerations for Information Visualization Interactions. *IEEE Transactions on Visualization and Computer Graphics*.18, 12 (December 2012), 2689–2698.

[18] Norman, D.A. 1993. *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA.

[19] North, C. 2006. Toward Measuring Visualization Insight. *IEEE Computer Graphics & Applications*. 26, 3 (May-June 2006), 6–9.

[20] Pike, W.A., Stasko, J., Chang, R., and O'Connell, T.A. 2009. The Science of Interaction. *Information Visualization*. 8, 4 (Winter 2009), 263-274.

[21] Plaisant, C. 2004. The challenge of information visualization evaluation, In *Proceedings of the International Conference on Advanced Visual Interfaces* (Gallipoli, Italy, May 2004). AVI ’04. 109-116.

[22] Rzeszotarski, J.M. and Kittur, A. 2014. Kinetica: naturalistic multi-touch data visualization. In *Proceedings of ACM Conference on Human Factors in Computing Systems* (Toronto, Canada, May 2014). CHI ’14. 897-906.

[23] Sadana, R., Major, T., Dove, A., and Stasko, J. 2014. OnSet: A Visualization Technique for Large-Scale Binary Set Data. *IEEE Transactions on Visualization and Computer Graphics*. 20, 12 (December 2014), to appear.

[24] Sadana, R. and Stasko, J. 2014. Designing and Implementing an Interactive Scatterplot Visualization for a Tablet Computer, In *Proceedings of the International Conference on Advanced Visual Interfaces* (Como, Italy, May 2014), AVI ’14. 265-272.

[25] Saraiya, P., North, C., and Duca, K. 2005. An Insight-Based Methodology for Evaluating Bioinformatics Visualizations. *IEEE Transactions on Visualization and Computer Graphics*. 11, 4 (July/August 2005), 443-456.

[26] Scholtz, J. 2006. Beyond usability: Evaluation aspects of visual analytic environments. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology.* (Baltimore, MD, October 2006). VAST ’06. 145-150.

[27] Shneiderman, B. 1996. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Proceedings of the IEEE Symposium on Visual Languages* (September 1996). VL ’96. 336-343.

[28] Shneiderman, B. and Plaisant, C. 2006. Strategies for evaluating information visualization tools: multi-dimensional in-depth long-term case studies, In *Proceedings of the BELIV Workshop* (May 2006). BELIV ’06. 1-7.

[29] Stasko, J., Catrambone, R., Guzdial, M., and McDonald, K. 2000. An Evaluation of Space-Filling Information Visualizations for Depicting Hierarchical Structures. *International Journal of Human-Computer Studies.* 53, 5 (November 2000), 663-694.

[30] Stasko, J., Choo, J., Han, Y., Hu, M., Pileggi, H., Sadana, R., Stolper, C.D. 2013. CiteVis: Exploring Conference Paper Citation Data Visually, (Poster). IEEE Information Visualization Conference (Atlanta, GA, October 2013).

[31] Valiati, E., Pimenta, M., Freitas, C. 2006. A Taxonomy of Tasks for Guiding the Evaluation of Multidimensional Visualizations, In *Proceedings of the BELIV Workshop* (Venice, Italy, May 2006). BELIV ’06. 1-6.

[32] Vande Moere, A. and Purchase, H. 2011. On the Role of Design in Information Visualization. *Information Visualization*. 10, 4 (October 2011), 356-371.

[33] van Wijk, J.J. 2005. The value of visualization. In *Proceedings of IEEE Visualization*, (Minneapolis, MN, October 2005). Vis ’05. 79–86.

[34] Wattenberg, M. 1999. Visualizing the Stock Market, In *Proceedings of ACM CHI Extended Abstracts* (May 1989). CHI ’99 EA. 188-189.

[35] Wehrend, S. and Lewis, C. 1990. A problem-oriented classification of visualization techniques, In *Proceedings of IEEE Visualization* (San Francisco, CA, October 1990), Vis ’90. 139-143.

[36] Yi, J.S., Melton, R., Stasko, J., and Jacko, J. 2005. Dust & Magnet: Multivariate Information Visualization using a Magnet Metaphor. *Information Visualization*. 4, 4 (Winter 2005), 239-256.

[37] Yi, J.S., Kang, Y., Stasko, J.T., and Jacko, J.A. 2007. Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*. 13, 6 (November/December 2007), 1224-1231.

[38] Yi, J.S., Kang, Y., Stasko, J.T., and Jacko, J.A. 2008. Understanding and Characterizing Insights: How Do People Gain Insights Using Information Visualization. In *Proceedings of the BELIV Workshop* (Florence, Italy, April 2008).