Benefitting InfoVis with Visual Difficulties

Jessica Hullman, *Student Member, IEEE,* Eytan Adar, and Priti Shah

Abstract—Many well-cited theories for visualization design state that a visual representation should be optimized for quick and immediate interpretation by a user. Distracting elements like decorative "chartjunk" or extraneous information are avoided so as not to slow comprehension. Yet several recent studies in visualization research provide evidence that non-efficient visual elements may benefit comprehension and recall on the part of users. Similarly, findings from studies related to learning from visual displays in various subfields of psychology suggest that introducing cognitive difficulties to visualization interaction can improve a user's understanding of important information. In this paper, we synthesize empirical results from cross-disciplinary research on visual information representations, providing a counterpoint to efficiency-based design theory with guidelines that describe how *visual difficulties* can be introduced to benefit comprehension and recall. We identify conditions under which the application of visual difficulties is appropriate based on underlying factors in visualization interaction like active processing and engagement. We characterize effective graph design as a trade-off between efficiency and learning difficulties in order to provide Information Visualization (InfoVis) researchers and practitioners with a framework for organizing explorations of graphs for which comprehension and recall are crucial. We identify implications of this view for the design and evaluation of information visualizations.

Index Terms—desirable difficulites, cognitive efficiency, active processing, engagement, individual differences.

◆

INTRODUCTION

A commonly-held view in InfoVis is that visualizations should present data as clearly as possible, refraining from using distracting visual elements and irrelevant information and leveraging labelling and graphical formats that reduce cognitive processing by the user. It is thought that by optimizing the cognitive efficiency of the external representation a designer likewise optimizes the visualization's ability to support rapid, accurate communication of information as well as to act as an external memory aid. Edward Tufte is one well-cited source of design principles for the efficient design of graphs of abstract, quantitative information, including the guideline of avoiding extraneous elements (e.g., chartjunk) and maximizing the data-ink ratio. Other seminal works support such guidance with semiotic and psychologically-based theories of graph design and interpretation [20], [37], [51] or empirical evidence of the benefits to response time and accuracy that can result from cognitively efficient graphics [16]. The reasons proposed to explain these performance improvements tend to cite how visualizations augment cognition by providing external information storage (i.e. memory), which if optimized, can reduce visual search times. By offloading cognitive processing to the perceptual system, the natural abilities of humans to visually identify patterns can also be leveraged.

The cognitive efficiency view is also reflected in many evaluation models for Infovis, which frequently default to measurements such as response time and response accuracy. Recently, however, InfoVis researchers have voiced concerns that the design and evaluation guidance supplied by the cognitive efficiency view may not accurately capture the complex nature of visualizations as devices for communication and cognition [7].

With very few exceptions, evaluation measures tend to assess how a user interacts with a visualization and not how well they learn the important concepts or patterns represented. It has been noted that it is somewhat unrealistic to hold visualizations to an efficiency-based standard of minimal response times and high response accuracy as in practice these measures tend to be in a trade-off, with higher accuracy requiring more time examining the representation [34]. This observation raises the question of whether maximizing accurate comprehension might conversely involve increasing the visual difficulty of the representation or otherwise obstructing automatic processing.

This viewpoint might explain findings relevant to InfoVis that seemingly contradict cognitive efficiency models. Bateman et al. [5] found that subjects who were shown highly embellished, Nigel-Holmes-style charts had a significantly higher chance of accurately interpreting and recalling the chart's message weeks later compared to less embellished, cognitively-efficient alternatives (see Fig. 1). Participants also subjectively preferred embellished charts as being more attractive, more enjoyable and easier to remember. However, the particular cognitive mechanisms that underlie these findings remain ambiguous due to confounds in the experimental design.

Insights into the question of whether and why such efficiency contradictions are possible can be found in psychological and educational studies related to learning from graphs and other visual displays. For example, science education research on learning from graphs and diagrams suggests that techniques that rely on focused, engaged thinking about the graph in combination with constructive activities (e.g., generating self-explanations or manipulating internal visualizations) could benefit understanding and transfer of the patterns in a graph (e.g., [19], [45]). Mary Hegarty argued for the importance of internal cognitive processes of the end-user in a 2010 InfoVis keynote [30], in which she summarized experimental evidence suggesting the importance of understanding the role played by mental hypotheses and representations that a user engages with in visualization interaction. The new demands placed on users' cognitive abilities by these advances call for consideration of how to better leverage users' abilities to generate high-quality internal representations and inferences, without necessarily overwhelming them with possible functionality.

A focus on characterizing successful learning and cognition with a visualization holds promise for helping explain the observed

- *Jessica Hullman is with the University of Michigan, E-Mail: jhullman@umich.edu.*
- *Eytan Adar is with the University of Michigan, E-Mail: eadar@umich.edu.*
- *Priti Shah is with the University of Michigan, E-Mail: priti@umich.edu.*

limitations of conventional evaluation frameworks and design principles. Yet little research synthesizes the empirical evidence so as to derive the underlying forces that allow some effective graphs to contradict cognitive efficiency assumptions.

In this paper, we make three contributions to InfoVis design and theory. First, we provide a counterpoint to conventional InfoVis wisdom on effective graph design by summarizing a large body of psychological research on learning from graphs and diagrams that suggests that introducing visual difficulties to visualizations can be an effective way to stimulate important ingredients like active processing of information and engagement. We conceptualize visual difficulties as the set of beneficial learning and cognitive processing *obstructions* that may be applied to static and interactive information visualizations. This include but are not limited to visual techniques that manipulate the visual representation, also including manipulations of accompanying information and task and contextual characteristics of visualization interaction. In doing so, we propose a model of effective graph design as a trade-off between cognitive efficiency and obstructions that induce deeper learning strategies. Secondly, we use these alternative forms of guidance to derive design guidelines for InfoVis that can help designers better negotiate goals of both cognitive efficiency and *desirable difficulties*. This includes outlining moderating factors like individual differences. Lastly, we explain how evaluation frameworks for InfoVis could be updated to account for under-emphasized components of graph comprehension.
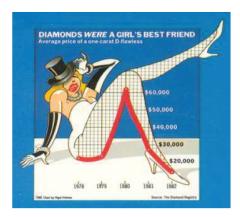

Figure 1: "Chartjunk" graph by Nigel Holmes used in Bateman et al [5].

1 COGNITIVE EFFICIENCY MODEL OF GRAPH DESIGN

Before elaborating how visual difficulties can be adapted to benefit a user's performance with visualization, we summarize theories on the importance of cognitive efficiency to InfoVis interpretation. We present common "talking points" or categories of concerns for which empirical efficiency research has proposed explicit design guidance. These include theories that focus predominantly on optimizing the efficiency of the end-user's cognitive operations, through the visual representation, forms of interacting like animation, or the use of labels.

1.1 Cognitive Operations

Inspired by information processing models of decision-making, cognitive psychologists theorized in the 1980's that graphic depictions are preferable over more conceptual representations as aids to cognition, based on their ability to visualize abstract relationships that otherwise remain implicit [40]. Larkin and Simon [46] provided empirically-supported guidelines for increasing the efficiency of graphic presentations as a means of quickly

conveying information. These include (1) allowing users to substitute perceptual inferences for more demanding logical inferences, and (2) reducing search for needed information. This is captured by a belief that distinguishing the "steps" taken to process a graph could predict response times [43]. More recent work [38] places more focus on interaction rather than the representation, but with the same 'cognitive cost' motivation.

These ideas are congruent with models of graph comprehension that break the process into multiple stages (e.g. [51]). Low-level parallel processing decomposes the visual scene into basic elements (e.g. form, color, texture), which are input to a pattern formation stage constrained by Gestalt laws of grouping and prior experience [51]). From this visual description, the user extract objects and patterns and generate propositions or conceptual messages about the variables. Such models informed principles for effective design (e.g. [37]) based on visual processing abilities. For example, graphs should maximize discriminability and manage complexity within perceptual and cognitive limits. Other early studies (e.g. [20]) provide evidence that some types of visual encoding (e.g., length judgments in bar charts versus arc judgments used in pie charts) are more accurate for representing quantitative information. More recent cognitive work echoes many of these assumptions [68].

Harnessing the low-level visual perceptive system also led to automatic methods for creating visualizations optimized for rapid, accurate visual processing (e.g. [16]). Systems like Tableau [64] represent comprehensive packages of visualization functions informed by literature on efficient design.

1.2 Visual Representation

Using the number of cognitive operations to operationalize visualization efficiency provided other researchers with a basis for testing and refining further design principles. We survey work related to optimizing data-ink ratios and information organization.

1.2.1 Data-ink

Tufte [66] proposed that effective visual displays for communicating information should optimize the data-ink ratio, calculated by dividing the ink used for displaying data (data-ink) by the total ink used in a graphic. Gillan and Richman [27] showed that high data-ink ratios correlated with faster response times and greater accuracy in experiments using bar and line graphs, while [10] found that high-data ink also increased awareness of variables. However, the effects of ink in the meaningful elements of a graph depended on the location and function of the element. For example, redundant ink in the objects in the graph indicating values had limited effects on performance, but pictorial backgrounds generally increased response time and decreased accuracy. Adding *x*- and *y*-axis lines decreased response time.

Studies on dimensionality, such as the addition of a gratuitous third dimension, have assessed the expectation that the third dimension would degrade performance. Carswell et al. [14] found that 3D line graphs resulted in less accurate recall and were associated with an overall performance decrement in a trend estimation task. Siegrist [61] similarly found that 3D bar and pie charts usually take more time to evaluate, although with accuracy on par with 2D versions.

An additional finding by Levy et al. [41] provided evidence that efficiency can be achieved by appropriately pairing data types with graph formats. Line graphs were more readily associated with trends, while bar graphs are associated with discrete points. Line graphs also appeared to be better for promoting memorability than for immediate use, whereas the opposite was true for bar graphs. Zacks et al. [70] manipulated rendering characteristics and relative heights of bars in bar graphs, finding that adding 3D depth cues decreased accuracy (see Fig. 2). The height of nearby bars and positions of other nearby elements also affected accuracy, supporting the general assumptions of efficiency principles that

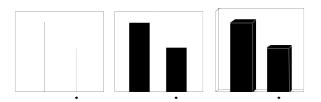propose avoiding unnecessarily distracting elements in graph design.

Figure 2: Recreation of graphs from Zacks et al.1998 [70] that vary rendering characters and depth cues.

1.2.2   Organization of Information

Psychological studies have investigated how a graph's ability to foreground important information using the visual organization can increase its effectiveness for depicting patterns such that end-users can spontaneously pick them out. This bears similarity to the 'immediate insight' goal of cognitively efficient visualizations. Carpenter and Shah [13] found that the time taken to interpret a graph is strongly related to the complexity of the graph, with three processes proposed to characterize graph comprehension: 1) detecting and encoding the visual patterns; 2) identifying the quantitative relationships that those visual features represent; and 3) relating those quantitative relations to the variables. Shah and Carpenter [58] show that the number of gazes on a line graph correlates strongly with the number of distinct line functions, each of which constitutes a visual chunk in a line graph. The model suggests that the cycle of processes scales up as the graph complexity grows, as a result of the additional chunks to encode and interrelate. Shah et al. [60] examined how graph formats that make trends appear more salient and require less mental computation better support viewers' abilities to spontaneously pick out patterns of interest. For example, reversing the variables on the x-axis can lead to a more (or less) salient presentation of a trend. Their *Organization Hypothesis* claims that if relevant trend information is perceptually grouped to form readily identifiable visual chunks (e.g., the relevant data points are grouped closely together in bar graphs), graph viewers are more likely to describe relevant trends [ibid].

1.3   Animation

While the question of whether animated visualizations are generally more effective than static representations may be debatable, many explorations of animation in InfoVis indicate an adherence to values of cognitive efficiency. Requiring a user to mentally animate or imagine transitions from a single static graph or set of graphs is often thought to be harder than animating the visualization so that the user might passively receive the information. Additionally, incorporating animation into visualization tools has been framed as beneficial to the user due to the speed with which it can convey the same information over multiple static graphs [54]. Animation can effectively draw attention to changes, retain context and help make occluded information visible (e.g., [29] [53]), and has been found to be preferable in some cases when accuracy is important [4]. A recent meta-analysis found an overall positive, although modest, effect of animation on learning [33]. This analysis also shed substantial light on moderators of the effect of animation, including whether animations are representational or decorational and whether they are used to communicate information about the time sequence of events, presumably because animation unfolds over time.

1.4   Labeling

Orienting the user to understand how visual elements in a graph represent data values is a vital consideration in graph design. Annotations help direct a user's attention and foreground particular insights. Generally, it has been proposed that direct labelling of important information supports the most efficient inferences. Kosslyn [37] noted that performance is affected by the considerable demands placed on short-term memory when using legends as a result of an individual being required to maintain a representation of the associations contained within the legend while they performed a task. This is supported by Lohse's [42] finding that adding a legend to a line graph can degrade performance measures like decision accuracy and the number of eye movements.

2   VISUAL DIFFICULTIES FOR INFOVIS DESIGN

Much of the work supporting the argument for cognitively efficient design is motivated by the scaffolding that increased visual efficiency can provide for several key roles a graph may play in cognition: acting as external memory or information storage and making abstract trends salient. In turning to assess graph design from an alternative view to the cognitive efficiency model, we suggest that interacting with an information visualization can also be considered a learning process. Decoding, assimilating, and drawing inferences from the visualized information represents a particular form of a more general process of learning in which new information is assimilated and integrated with existing knowledge. Constructive mental operations comprise the learner's active process of organizing and making sense of her emerging understanding. If it is represented at all, learning in the cognitive efficiency paradigm is reduced to whether a graph increases the speed and the accuracy of identification of patterns. Yet research in educational psychology and multimedia learning suggests that such measures may not capture the full range of factors that characterize effective learning, a concern echoed in calls for revisions to InfoVis evaluation models (e.g. [7]).

Motivated by the observation that predicting effective learning is more complex than merely the time spent studying or immediate recall of information, psychologists have applied the term desirable difficulties to a class of techniques that hinge on "going against one's intuitions" and "deviating from standard instructional practices" [9]. While some of this work focuses on adding difficulties to learning independently of study materials, such as by varying the conditions for study periods [8], various empirical studies have shown the potential benefits of introducing desirable difficulties to learning from visual media. By introducing obstructions, these studies provide a counterpoint to cognitive efficiency-motivated graph design research. In this section, we summarize research that finds that more difficult to interpret graphs and other visual materials yields better learning. We start with cognitive explanations for why difficulties might benefit comprehension. Specifically, we note that difficult displays require mentally manipulating internal representation, forming self-explanations, and appraising and noting disfluency, all of which promote *active* processing on the part of the user. We argue that this active processing leads to better long-term memory and comprehension. Related to the importance of active processing is the motivational-component of graph use, which we summarize as *engagement*. Often visual difficulties induce engagement with a graph design by manipulating novelty, tailoring and personalization, challenge and game-play, and aesthetic appeal. Engagement, in turn, increases the likelihood of active processing and thus may benefit learning.

2.1   Inducing Active Processing

All learning requires encoding information, yet the differing degrees to which learners can constructively process incoming information (as opposed to passively absorbing it through listening or reading) has led to definitions of *active processing* as active, self-directed searching for meaning on the part of the learner. At a basic cognitive level, active processing implies additional cognitive operations in the interest of reaching deeper understanding. When visual difficulties stimulate active processing, these additional

operations stand in contrast to the minimization goal of cognitive efficiency research. We describe particular forms of active processing including *self-explanations* and the manipulation of *internal visualizations*, and explain how a user's perceived processing effort can interact with active cognitive processing.

2.1.1 Self-explanation

One manifestation of active cognitive processing identified in multimedia learning research is *self-explaining* of information to oneself (e.g. [19]). In self-explaining, one attends to target material in a meaningful way such as by generating inferences to fill in missing information, integrating new information with prior knowledge, and monitoring and repairing faulty knowledge. Self-explanation is typically elicited by having learners spontaneously explain the meaning of sentences in a passage as they study a target domain, then coding responses into "low quality" statements based on simple rereading or paraphrasing and "high quality" statements that involve linking statements from the text using tacit knowledge or making inferences to fill gaps [19].

Several studies provide a precedent for how visual difficulties designed to induce self-explanations can be integrated into visualization interaction. In this sense, *visual difficulties* refers to means of stimulating more intense cognitive activity with a visualization, rather than specific manipulations to the visual representation. In studying risk communication, Natter and Berry [47] found in two experiments on the effects of active processing of graphs of risk information that completing a reflective task that involved portraying the size of risk on a bar chart in one experiment and answering a reflective question in the second (active conditions) increased participants' satisfaction with the information and led to more accurate judgments and estimates. Using textual or task prompts may be a more reliable way to induce self-explanation with a visualization. Evidence suggests that even highly relevant diagrams which learners might use as aids to self-explaining of concepts addressed in a text sometimes have no demonstrable effect on learning [21].

Designers of visual analytics tools often minimize cognitive operations as proposed by the efficiency model, by guiding users through execution of important functions or allowing them to automatically run complex processes after defining them. Yet just as some static graphs fail to support the "deep learning" facilitated by self-explaining, so too can automation of a complex process in an interactive visualization present risks to a user's understanding of a process in some cases. Automating processes central to solving a problem, for example, such as saving filters or views in a visual analytics tool, may reduce active self-explaining on the part of the user as she recreates the process a second time. While not ideal in all cases, the self-explanations that accompany early repetitions of the processes may provide her the opportunity to catch errors in her thinking. Evidence comes from results on "wizard" interfaces that break tasks into a sequence of steps rather than simply providing information on how to perform the task (which the user must integrate and translate to a sequence of actions themselves). Research suggests that users have trouble transferring knowledge gained from wizards to other, less guided environments [11]. A visualization-based example is the SocialAction tool [50], which uses the Systematic-Yet-Flexible model that, while ensuring completeness of analysis, assists users by "playing through" analyses through the use of shortcut buttons. Similarly, through automation the IVEE system [1] allows users to immediately visualize and explore data. While noting that a custom widget might lead to better specifications of query environments, the authors opt to reduce the cumbersome process of setting defaults to enhance efficiency. While both these systems represent innovative interface designs that may reliably benefit expert users, the aforementioned experimental findings tell a cautionary tale for the application of "low effort" designs in all situations.

2.1.2 Manipulating Internal Visualizations

Various cognitive psychologists studying the use of graphs and diagrams have made an argument for the importance of requiring viewers to manipulate *internal visualizations* to aid comprehension. Recent work in InfoVis by Liu and Stasko [44] argues for the importance of internal representations—mental representations overlaying text, visuals, and spatial relationships—in achieving interaction objectives like external anchoring, information foraging, and cognitive offloading. As such, prompting the manipulation of internal visualizations represents an alternative goal for visualization designers seeking to improve comprehension through visual difficulties. As in self-explaining, *visual difficulties* describes the addition of more cognitively intensive processing into a visualization task.

An example of beneficial manipulations of internal is provided by Trafton et al. [65]. They found that experts who formed and compared schematic internal representations to external visualizations drew the user's attention to gaps in her knowledge. Similarly, Mayer et al. [45], found that static displays lead to better comprehension than animations when users were appropriately cued to use the displays to identify and fill gaps in their knowledge. Like the reflective thinking prompt that has been used to induce self-explaining, requiring users to engage with internal visualizations can be accomplished by soliciting their predictions of how a visualized process works prior to examining a visualization of the process.

The benefits of visual difficulties that induce manipulation of an internal visualization can be contrasted with much of the dynamic graph drawing literature, which tends to be motivated by a desire to preserve mental models throughout visualization interaction by removing the need for internal visualization. Many algorithms for graph drawing rely on animations between time-slices, where the positions of nodes are interpolated between graph states and added or removed nodes fade in or out correspondingly. Design strategies aim to minimize the number of nodes that must be moved between time-slices and retain the overall graph shape, in order to preserve mental maps by reducing the need for internal visualizations of the changes. It is also generally assumed that minimizing movement and change improves graph readability [69]. However, preserving the mental map only brings benefits in certain tasks, such as in tasks that require identifying nodes by name [52].

In contrast to animated graphs that remove the need to manipulate an internal visualization to preserve one's mental map, small multiples graphs require some mental or internal animation to infer graph evolution. Thus using small multiples over animations represents one manifestation of visual difficulties based on internal visualization. Archambault et al. [4] compare small multiples versus animation with regard to how effective each is for comprehending how a graph changed over time. Despite the time saved in using animation to convey a complicated transition over using multiple static graphs, they find that small multiples graphs are generally more effective for quickly answering questions about particular states. For several types of questions, however, error rates were lower with animation. They also examine whether maintaining a mental map between graph time-slices, using either small multiples or animations, is effective for understanding graph evolution. Only small improvements in error rates and response times were observed when the mental map was preserved. These results suggest further exploring small multiples representations with an eye toward understanding how the internal visualization manipulation they promote can be more beneficial than supporting mental map preservation alone. Given the limitations of response times and error rates as performance measures, a natural next step would be to use alternative measures like transfer strength or long-term recall.

2.1.3 Perceptual Disfluency

The visual difficulties argument provides support that *disfluent* learning experiences—those that are perceived as more effortful—may often be more beneficial. Psychologists define fluency as a metacognitive judgment (e.g. [3]). As an example, consider the more effortful task of reading a passage of text in a degraded font like **Haettenschweiler** rather than one designed for clarity, such as Arial. While readers perceive degraded fonts as more effortful to read, they actually can lead to better comprehension and memory because disfluency cues the reader to avoid heuristics and defaults and reliance on peripheral cues (e.g. [2]).

Fluency research suggests the possibility that "easy-to-process" graphs may actually lead to more superficial interpretations. Fluency research problematizes the optimally "easy-to-process" graphs predicted to satisfy cognitive efficiency assumptions by suggesting that subjective assessments of processing ease can cause a learner to use representations that negatively affect comprehension and transfer compared to more difficult-to-process representations ([9]). This is partially the result of a learner interpreting the ability to recall information in the graph as storage strength (the degree to which the information is inter-associated with prior knowledge). In the case of minimalist graphics that strive for immediately clarity, the same misassumptions on the part of users may result.

Another benefit of visual difficulties in the form of perceptual disfluency is that a graph user's awareness that they are exerting a high level of effort can prompt processing that is more likely to be accurate for a task. One way such benefits have been explained is that perceived disfluency can lead a user of graphs or other learning materials to use systematic, deliberative analytical reasoning in place of less effortful, automatic, intuitive and heuristic reasoning processes [2]. This parallels the way that adding "difficulties" to a graph comprehension task may induce realizations of gaps in one's emerging mental model that stimulate a learner to work harder to understand the information. For example, recent research provides evidence that legends on graphs can increase understanding and memorability. Shah et al. [59] demonstrate that while labelling graph features improves response time as predicted by the efficiency model, using harder-to-process legends can increase data memorability and induce more explicit, thoughtful descriptions of the main effects and interactions depicted in the graph as compared to using labels.

2.2 The Role of Engagement

We define active processing as *user-directed* psychological activity in a learning context. Yet active processing need not imply a personal motivation to engage on the user's part. Rather, it is a neutrally defined cognitive activity that can be induced by experimental rewards or other extrinsic incentives. While the line between extrinsic and intrinsic motivation is often ambiguous, we present the following means of introducing visual difficulties as a second set of strategies for more explicitly engaging a user on an intrinsic motivational level. Doing so makes active processing of important content more likely.

We use case studies to present four potential strategies for increasing a user's engagement: novelty, aesthetic appeal, tailoring and personalization, and challenge and game-play (Fig. 3). We briefly summarize each of these relationships, before elaborating on their implications for design and evaluation.

2.2.1 Novelty

Engagement can be stimulated naturally by information presentation forms that are new to a user. Novel forms of media have been shown to enhance active information processing, resulting in better recall and recognition over non-novel stimuli (e.g., [25]). Novel stimuli can distract a user's self-focus away from an internal state to shift focus on the environment, as evidenced by studies in public health messaging such as findings

that periodically changing cigarette package warning messages is predictive of the warnings' salience and perceived effectiveness ([28]).

Novelty may help explain results in information visualization research suggesting that chartjunk, or ornamentation that is unnecessary to the data being visualized, may yield positive effects. Bateman et al. [5] found that subjects who were shown embellished Nigel Holmes-style charts had a significantly higher chance of comprehending the message of the chart as compared to non-embellished charts (see Fig. 2), as well as recalling relevant information several weeks later. Participants also found the Holmes-type charts to be more attractive, more enjoyable and easier to remember. These preferences are echoed by Inbar et al. [35].

It should be noted, however, that questions remained unanswered as to whether novelty was the primary driver of Bateman et al.'s results. The positive benefits may be the result of congruency of the form chosen and the message being visualized (such as the design of a graph depicting 'Monstrous Costs' that uses a drawing of a monster as background). In addition, the study used graphs with a small number of data values (~5) while the no-chartjunk rule was originally directed at displays depicting large numbers of data points.

2.2.2 Aesthetic Appeal

Graph aesthetics, referring to the experience and appreciation of a visualization's sensual qualities, have been avoided in many studies, and framed in opposition to graph readability [36]. This may suggest that an implicit belief in InfoVis asserts that aesthetically complex graphs are less efficient. Yet a system that is perceived as visually appealing as a result of aesthetic dimensions can also stimulate use (although in some cases such effects may be difficult to separate from novelty effects). The positive results from Bateman et al.'s chartjunk study, for example, may in part result from the attractiveness of the colored background graphics compared to the plain versions of the charts.

Sprague and Tory [63] provide some evidence to the motivational potential of aesthetic appeal in studying animation use in a music library visualization. Using an animated glyph correlated with clear performance decreases in a task, however, most subjects preferred a moving glyph, in part due to the aesthetic appeal of its 'dancing' appearance. The authors conclude that in designing casual visualizations, pragmatism can be emphasized for some users and aesthetics for others. We reflect further on the effects of individual differences on visualization effectiveness below.

While aesthetic considerations have been under-explored in many efficiency-motivated studies, an equally viable belief might be that aesthetically pleasing visualizations ease a user's ability to process depicted information. This highlights the proposal of this paper for a view of InfoVis design choices as a trade-off between cognitive efficiency and visual diffuclties, with the same techniques playing differing roles depending on other features of interaction.

2.2.3 Tailoring and Personalization

Various studies have documented the positive results of tailoring or personalizing graphical information in order to subtly induce interest in and active processing of the information content by users (e.g. [62]). Hence tailoring the display of visualization tools to users' prior knowledge or other attributes may similarly lead to increased motivation to attend to the data, spurring cognitive activity on the part of the user.

Personalization can also ease processing effort though, highlighting the trade-off between visual difficulties and cognitive efficiency. Mental models as conceptualized by [44] can help explain the potential efficiency of tailoring. By connecting either the framing or the nature of the information presentation with a user's prior knowledge and expectations, a message can have a

more pronounced effect on beliefs and behaviour. Transfer noise is minimized, allowing relevant information to be processed more easily.

Yet the benefits associated with the efficiency benefits of personalization are problematized by findings that show that applying graphical adaptation features to interface design can optimize performance on particular tasks, but often comes at the cost of deeper understanding of the tool's full functionality [26]. Hence some forms of personalization may become a crutch rather than a panacea for deeper understanding. Using tailoring to increase engagement and processing of visualized information is most likely to bring benefits where there are significant risks that users may not be sufficiently motivated to use the visualization in the first place.
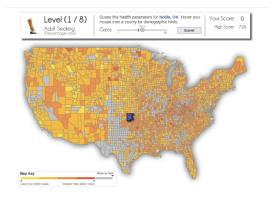

Figure 3: Guessing game version of an interactive health data visualization [22].

2.2.4 Challenge and Game-play

Interest in external stimuli like visualizations can also arise from a challenging environment that demands high attention and an intention to explore. Like novelty, stimulation and challenge are a prerequisite of personal development through the development of knowledge and skills, which in turn is a basic human need (e.g., [57]). Introducing visual difficulties offer a promising way to leverage a user's desire to be challenged in interaction.

One implementation of visual difficulties supported by work in social computing and InfoVis is through elements of game play, such as rules and goals. These have been used in social computing applications like games with a purpose (GWAP) (e.g., [24]), and online movie recommender systems, where they have been shown to increase user contributions [6]. Within InfoVis, incorporating challenges or problem solving tasks into visualization tools has been suggested to help turn visual data analytics into a game-like activity and motivate exploration of data [28]. More recently, Diakopoulos et al. [22] used game elements like rules and goals in an information visualization of health data. An online experiment showed that gamification led to demonstrable effects on exploration of the visualization, insights and learning, and enjoyment of the experience. In particular a guessing game led to significantly more exploratory behavior (e.g., unique health parameters visualized, hover activity, volume of interaction with slider feature), learning (e.g., number of insights, self-reported learning), and enjoyment (self-reported reliability ratings). The guessing game frame is one implementation of a prediction task such as that suggested by Hegarty et al. [32] as a way of deepening understanding by inducing internal visualization.

3 APPLYING VISUAL DIFFICULTIES IN INFOVIS

An important conclusion to be drawn from Sections 2 and 3 is that designing visualizations with lasting effects on knowledge

formation may not be achieved simply by easing users' effort and providing high levels of flexible interaction. In many cases, effective visualization practice requires navigating tradeoffs between the benefits of active processing through learning obstructions on the one hand, and the benefits of representation efficiency through more immediate pattern detection on the other. In this section, we demonstrate how the visual difficulties perspective can be applied to the "talking points" or common design concerns addressed by efficiency-based principles. A series of design implications intended to act as guiding goals for designers interested in using visual difficulties in their designs are then laid out. Additionally, we consider the implications for evaluation, and suggest ways in which InfoVis evaluation models can be adapted to account for benefits of visual difficulties.

3.1 Revisiting Common InfoVis Concerns

Table 1 summarizes how the visual difficulties findings on common InfoVis concerns compare to recommendations from the cognitive efficiency view.

	Cognitive efficiency	Visual difficulties
Cognitive operations	Minimize the cognitive steps required to process visualization	Induce constructive, self-directed cognitive activity on the part of the user
Data-ink ratio	Maximize the ratio of data to ink	Design representations that are most likely to engage a user to actively process the information
Organization	Choose the format which makes important information most visually salient	Choose the format that best stimulates deep cognitive reflection on the important data
Animation	Use animation to quickly and intuitively visualize important processes	Use static representations to induce interval visualization processing around causal mechanisms; consider animation in cases where mental animation lies beyond users' capacities
Labeling	Use labels rather than legends to optimize immediate clarity	Use legends to stimulate deeper reflection on data

Table 1: Efficiency versus difficulties recommendations in InfoVis.

3.1.1 Cognitive operations

Cognitive efficiency research proposes minimizing the number of cognitive steps required by a graph (see Section 2.1). The visual difficulties work suggests that cognitive steps do not accurately capture learning by a user (including comprehension and recall). Instead, researchers seek to increase the depth of cognitive operations. Classifying users' statements in response to a graph based on whether they represent low- or high-quality self-explanations [55] or spurring internal visualization manipulations [32] are several ways that this has been accomplished.

3.1.2 Visual representation – Data-Ink Ratio

Charts with higher data-ink ratios have conventionally been equated with 'embellishment' or 'decoration' (see Section 2.2.1). However, in a meta-analysis of 39 experiments, Carswell [15] found no support for the data-ink rule. The visual difficulties perspective provides evidence that low data-ink ratios may be functional in cases where the extra ink is used to personalize,

aestheticize, or otherwise make the visualization more enticing to end-users so as to drive more intrinsic desires to engage. Additionally, results like Levy et al. [41] stand in contrast to the studies on 3D that using immediate responses (e.g., reaction time, immediate recall) as performance measures by examining communication-based preferences instead. For example, a 3D graph can be preferred for group communication, for depicting details than trends, and for memorability than immediate use. Thus before opting for a more cognitively efficient 2D graph, designers might consider whether intrinsic motivation to engage might increase with the addition of the third dimension.

3.1.3   Visual representation – Information Organization

Cognitive efficiency dictates that designers should minimize the complexity of a visualization, by reducing the number of visual elements like line segments, and using formats that make trends readily salient (see Section 2.2.2). Yet under some conditions, using a complex visual organization may be preferable, such as when the complex visual more accurately represents the target relationship and is presented with a prediction task to insure that users will reflectively engage with the content. Doing so guards against the risk that users with little prior knowledge in a domain will rely on the abstraction presented in a simpler graph without considering the full implications of the depicted trends.

3.1.4   Animation

The visual difficulties view proposes evaluating the usefulness of animation on a case-by-case basis, rather than relying on a general rule. The main questions that a researcher should ask are 1) Does the animation direct user's attention to the information in the visualization that will be most helpful for the construction of accurate mental models?, and 2) Is the animation allowing the user to gain insight that is beyond the limits of their mental animation abilities?

3.1.5   Labeling

Visual difficulties can be beneficial to a graph comprehension task when legends are used in place of labels to stimulate deeper reflection on the part of users [59]. Using relatively disfluent fonts in labels may be preferable over use of fluent fonts when the presented information is crucial to a decision, such as in the design of graphs for score reporting that will inform college decisions, or graphs of an individual's health information.

4   IMPLICATIONS FOR DESIGN

The visual difficulties evidence highlights the importance of approaching design as a trade-off between efficiency and beneficial obstructions, one that is not easily captured by rules intended to apply to all visualization designs alike. Below, we present four considerations that designers can look to for guidance in navigating these trade-offs regardless of the specific combinations of cognitive efficiency and implementations of visual difficulties employed.

4.1.1   Individual Differences as Moderating Factor

Various visualization studies have noted that individual differences in prior knowledge can affect visualization interpretation and even bias perception. The quantity, specificity, and accuracy of a user's existing knowledge are important considerations in deciding whether to introduce visual difficulties. By improving an individual's sense of gaps in their mental model, prior domain-relevant knowledge can stimulate the conflict between internal and external representations that motivates revision of a developing mental model (e.g. [55]). Without the requisite prior knowledge to judge gaps in their understanding, novices can become overconfident [56]. In such cases, including visual difficulties that bring this to a user's awareness can offset negative effects on comprehension.

The benefits of introducing visual difficulties can also be constrained by cognitive abilities. For example, a technique to induce manipulation of internal visualizations should only be used when the required manipulations will be within the limits of the user's spatial abilities (e.g. [31]). Additionally, users can differ in their higher-level cognitive abilities in ways that influence the effects of using visual difficulties. Chi [19], for example, observed that learners differ in their natural tendencies to produce self-explanations while studying examples or reading text.

Cognitive needs represent a final individual difference that can interact with the effectiveness of visual difficulties. *Need for cognition* describes individuals' tendencies to differ in the degree to which they seek out and enjoy difficult mental work (see [12] for a review). Individuals low in need for cognition are more likely to ignore, avoid, or distort new information; rely on others' thinking; react negatively to a lack of structure; use cognitive heuristics; or use social comparisons in problem solving rather than independent reflection and reasoning. In applying visual difficulties to a visualization, an individual's need for cognition can affect how likely they are to meet the challenges presented versus to seek external help or give up. Low need-for-cognition users are also more likely to rely on superficial visual cues like the attractiveness of an information presentation, suggesting that increasing aesthetic appeal may be an effective way to motivate these users.

4.1.2   Managing Data and Task Complexity

Some of the strategies outlined above may not scale up appropriately as the complexity or dimensionality of the data set increases (or decreases). Yee et al. [80] find that preferences for polar versus linear coordinates in dynamic graph animation differed based on graph size and were lowest for a middle-ground condition, capturing an apparent trade-off between the benefits and weaknesses of the two approaches. The results of Bateman et al. [5] and Inbar et al. [41] on the validity of the chartjunk and data-ink principles, respectively, may be flawed due to over-simplification of the visualized data sets compared to the larger data sets that were presumably intended when the principles were presented. Similarly, introducing perceptual disfluency or substituting static diagrams for animations may not work if broadly applied within systems for complex analyses procedures that by nature impose high cognitive load. In fact, it may be the case that part of the effectiveness of these highly interactive environments comes from the natural difficulties imposed by having to vary parameters and explore large data sets with little guidance. While not identical to gamification, the success of these systems may derive from positive motivational influences that are similarly based on challenge. Future work might deepen understanding of the mechanisms driving successful interaction by examining more closely how the level of challenge presented to a given user by a system drives it's utility.

The above possibility again highlights the trade-off between difficulties and efficiency. When data and tasks are highly complex, difficulties may already be present and designers must thoughtfully **balance difficulties with efficiency techniques**. In addition to validating particular strategies using pilot studies, designers should **consider user experience levels**. As common operations in the context of a larger task become automatic for expert users, it may productive to pinpoint particular spots in the workflow where an analyst must transition to more focused thought, and introduce difficulties at these points in order insure highly active cognition. Hence a limitation of the visual difficulties view may be that many of the researched benefits may not produce the same benefits in expert visual analytics systems. An important direction for future work involves carefully evaluating the appropriateness of visual difficulties in highly complex scenarios where visualization is maximally useful. On the other hand, researchers should be aware that introducing too many obstructions to a novice user's interaction with an already unfamiliar visualization tool may induce unproductive levels of frustration.

4.1.3 Guiding Expectations and Inference Strategies

A user's sense of what is expected from them in using a visualization, including whether to use analytical or intuitive reasoning and the extent to which they should "let the graph do the work" can be a powerful influencer on how they structure their interaction. For example, improvements in learning can result when users of a graph are provided with the necessary information to use the self-explanation technique [55]. Visualization designers might put to use a variety of training procedures that have been used to induce self-explanation, ranging from simple **prompting** (e.g. [18]), to providing a **pre-question**, **prediction task,** or other **internal visualization activity** to guide learners' self-explanations in multimedia environments ([32]). In educational contexts, directly **training** graph users how to engage in high quality self-explanation and self-regulation strategies can be effective.

Depending on the characteristics of the intended users of a visualization, it may be necessary to provide higher levels of situational support to compel some users to engage in cognitively-demanding constructive activities like self-explaining [55]. An environment can **encourage learners to bring their prior knowledge to bear** on interpretations and allow them **to monitor and test** their developing understanding. In cases where prior knowledge is expected to facilitate accurate interpretation, it may be useful to use **multiple representational formats** to increase the chances of activating relevant prior knowledge [55].

The effectiveness of these strategies may be explained in part by how self-explaining and prediction or internal visualization tasks can help subjects to articulate and assess their intuitions about how target processes work in comparison to a visualization. Along these same lines, **providing dynamic and timely feedback** when possible can help a user to accurately assess her own learning and make adjustments to more deeply process content [49].

4.1.4 Cueing Self-Doubt and Representing Ambiguity

When presented with a difficult-to-read piece of text or otherwise disfluent visualization, the subject is cued to be more cautious, and potentially switch from using intuitive System 1 reasoning to more analytical and systematic System 2 reasoning [2]. For visualizations of problem domains where significantly better alternatives exist but may not be found using intuition (for instance, in a less personal task like deciding on nutritional supplements), this can lead to more accurate and informed decisions.

There are several ways that this goal can be accomplished. One is by giving the user opportunities to recognize holes in her learning during visualization interaction by **introducing learning events that cue uncertainty**, such as a prediction task prior to graph interaction (e.g. [32]). This can help curb the common yet unjustified confidence that often accompanies understanding of causal relationships [56].

In some situations an interactive visualization might cue uncertainty simply by **instructing a user to double-check her inferences**. This is best suited to visualizations where the user can actively represent the outcomes of their analyses, such as by highlighting or saving particular views or parts of a graph as evidence of some underlying mechanism. The system can use insurance measures like checklists at these points, such as through pop-ups, improving the analyst's accuracy.

A final strategy that is similarly motivated to induce doubt is to visually represent uncertainty around the values or patterns in a graph. The most common example of this may be the addition of error bars to represent confidence intervals in a box plot, or annotations describing the limits of inference (e.g. [67]).

4.2 Implications for Evaluation

The visual difficulties evidence problematizes InfoVis evaluation models that rely heavily on cognitive efficiency performance metrics by suggesting that these are only one view of how to assess a visualization. To capture the benefits of learning obstructions, evaluation models must allow for the role of active, constructive processing and the desire to engage with a visualization in the first place.

Prior work has pointed out that conventional evaluation models may not fully characterize the role of visualization in analysis [7]. Even evaluation methods based on an expanded view of the risks to validity that can occur in design tend to refer back to efficiency theories for validating visualizations (e.g. [46]). Yet measures like response time and accuracy may not be sensitive enough to catch differences between some visualizations, as the same levels of each can be achieved from different amounts of cognitive effort [34]. These conventional measures tend not to acknowledge the trade-off that often exists between comprehension and response times.

In recent years, approaches have been proposed which address particular considerations missed by a focus on response time and accuracy. Lam et al. [39] distinguish seven scenarios that entail alternative evaluation criteria. Examples include evaluating visual data analysis and reasoning (e.g., measuring eye movements, visual search times), human factors like user experience and performance (e.g., users' perceived effort, response times, and accuracy), work practice outcomes, and communication through visualization. We consider several alternative approaches to conventional evaluation as examples that encapsulate one or more of these scenarios, and discuss the potential for integrating visual difficulties into each. Our goal is not to provide a comprehensive treatment of proposed evaluation frameworks, but to demonstrate the questions that visual difficulties findings pose to evaluation.

One way in which conventional performance measures have been amended is through insight-based techniques that focus on productivity in knowledge gain (e.g. [48]). While introductions of visual difficulties are congruent with an emphasis on the quality of the new knowledge, it can be hard to define what constitutes an insight in a way that applies across users. Viewing graph design as a trade-off between cognitive efficiency and visual difficulties also may require expanding the ways in which insights are evaluated. For example, asking how active a role the user played in reaching the insights may help predict longer-term learning. This can be operationalized by capturing user's explanations of how they arrived at an insight using a graph and classifying these based on how many reflective thoughts were prompted (for example, 'I first considered [A], but then noticed [details] in the graph, and thought [B]'). Another way of asking this might be, who did more work, the system or the user? In cases where insights have been gained but there is no evidence of active effort on the part of the user, visual difficulties predict that learning may not last as it could.

Some models have proposed approaching evaluation from a human factors perspective where user experiences, including self-reported usability and effort with a tool, represent an important dimension in visualization effectiveness (e.g. [34]). The evidence of the potentially beneficial role of disfluent stimuli and conditions, often in the face of less favorable assessments of the learning experience by users, introduces a novel quandary to these techniques. What happens when a user prefers what is worse for them, in terms of understanding and retention of visualized content? Evidence on disfluency suggests using user-reported effort measures as indicators of the degree of disfluency different visualizations assume, and combining this data with more conventional metrics for better predicting which of several visualizations will more likely increase comprehension and retention.

An important observation is that many visual difficulties studies consider a visualization as one of multiple representations in a group of learning materials. In some cases, a visualization might most concisely represent a concept while texts or other instructional aids introduce beneficial obstructions. Research in educational contexts aims to optimize lasting, transferable knowledge overall rather than isolating features of the visual representation to evaluate. This outcome-focused approach is

paralleled by learning-based methods like Chang et al.'s [17] model for assessing how well a user transfers the knowledge gained with a visualization to a new task. A new tension posed by visual difficulties evidence, however, is that the level of constructive cognition entailed in one task (which tends to predict comprehension and retention) may not be cued in a second, similar task regardless of the visualization's effectiveness. If a user's prior knowledge is cued differentially across original and transfer tasks, the performance with new data may not be as effective. Similarly, if the user does not challenger herself to think actively in the transfer task based on lower intrinsic engagement with the data, the evaluation method is unlikely to generate reliable assessments. Evaluators could be better prepared to explain such situations by capturing users' prior knowledge and levels of interest.

Chang et al. frame their model against methods that measure *how* rather than *how well* a user learns with a visualization. The former learning efficiency concern is least congruent with visual difficulties evidence. Huang et al.'s [34] cognitive load approach, for example, suggests that cognitive load represents mental effort, or the amount of cognitive capacity that is allocated to accommodate the demands imposed by a task. The predictions of the model tend to support the conventional efficiency maxims: when memory demand is consistently lower than maximum memory capacity, task performance (e.g., accuracy, response time, perceived effort) is likely to be good assuming the load is not so low as to lead to boredom. When memory demand reaches maximum capacity, performance plummets. Yet the common discordance between a user's perceived mental effort and her learning gains challenges learning efficiency techniques, suggesting that some difficulty may be best when performance is most crucial.

In sum, while evaluation methods for InfoVis have expanded to consider characteristics beyond the visual encodings and sought new measures that capture the quality of learning, viewing graph design as a trade-off between difficulties and efficiency raises questions that existent models should address, including the importance of a user's depth active processing, the risks of using subjective effort appraisals as a sign of visualization quality, and her incentives to interact.

5 CONCLUSION

Cognitive effort appears to represent a natural guideline for designing effective information visualizations and has benefitted InfoVis as a foundation for practice. Yet contextualizing experimental support for cognitive efficiency against psychological and educational graph design literature suggests that visualization effectiveness is better characterized as a trade-off between efficient processing and desirable visual difficulties to stimulate learning. At the basis of this trade-off is active processing, which can guide designers considering diverse graph design strategies, including means of generating engagement. The visual difficulties evidence carries implications for design such as the importance of considering the expectations, level of certainty, and knowledge and abilities of the user. At the same time, the new questions that visual difficulties pose to evaluation approaches stand to increase the sophistication of those methods by integrating important psychological principles into the InfoVis evaluation toolbox.

REFERENCES

[1] C. Ahlberg and E. Wistrand, IVEE: an Information Visualization and Exploration Environment, in *Proc. of Vis.'95*, 1995.

[2] A.L. Alter, D.M. Oppenheimer, N. Epley, and R.N. Eyre, Overcoming intuition: Metacognitive difficulty activates analytic reasoning, *J. of Exp. Psych.-Gen.*, vol. 136, no. 4, 2007.

[3] A. L. Alter and D. M. Oppenheimer, Uniting the Tribes of Fluency to Form a Metacognitive Nation, *Personality and Social Psychology Review*, vol. 13, no. 3, 2009.

[4] D. Archambault, H. C. Purchase, and B. Pinaud, Animation, Small Multiples, and the Effect of Mental Map Preservation in Dynamic Graphs, *IEEE TVCG*, vol. 17, no. 4, 2011.

[5] S. Bateman, R. L. Mandryk, C. Gutwin, A. Genest, D. McDine, and C. Brooks, Useful junk?: the effects of visual embellishment on comprehension and memorability of charts, *CHI*, 2010.

[6] G. Beenen et al., Using social psychology to motivate contributions to online communities, in *CSCW*, 2004.

[7] E. Bertini, A. Perer, C. Plaisant, and G. Santucci, BELIV'08: Beyond time and errors: novel evaluation methods for information visualization (Workshop), *CHI*, 2008.

[8] R. A. Bjork. Memory and metamemory considerations in the training of human beings. In J. Metcalfe et al. (Eds.), *Metacognition: Knowing about knowing* Cambridge, MA: MIT Press, 1994.

[9] E.L. Bjork & R.A. Bjork, Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In M. A. Gernsbacher,et al. (Eds.), *Psychology and the real world: Essays illustrating fundamental contributions to society*. New York: Worth Publishers, 2011.

[10] A. J. Blasio and A. M. Bisantz, A comparison of the effects of data-ink ratio on performance with dynamic displays in a monitoring task, *Int.'l J. of Indust. Ergonomics*, vol. 30, no. 2, 2002.

[11] M. Burton, D. P. Wickham, L. Phelps, K. Spain, J. Crews, and N. Rich, Secondary navigation in software wizards, in *CHI '99 extended abstracts on Human factors in computing systems*, New York, NY, USA, 1999, p. 294–295.

[12] J. T. Cacioppo, R. E. Petty, J. A. Feinstein, and W. B. G. Jarvis, Dispositional Differences in Cognitive Motivation: The Life and Times of Individuals Varying in Need for Cognition, *Psych. Bulletin*, vol. 119, no. 2, 1996.

[13] P. A. Carpenter and P. Shah, A Model of the Perceptual and Conceptual Processes in Graph Comprehension, *J. of Exp. Psych.: Appl.*, vol. 4, no. 2, 1998.

[14] C.M. Carswell, S. Frankenberger, and D. Bernhard, Graphing in depth: perspectives on the use of three-dimensional graphs to represent lower-dimensional data, *Behav. & Info. Tech.*, vol. 10, no. 6, 1991.

[15] C.M. Carswell, Choosing specifiers: an evaluation of the basic tasks model of graphical perception, *Human Factors*, vol. 34, no. 5, pp. 535-554, Oct. 1992.

[16] S. Casner and J. H. Larkin, Cognitive Efficiency Considerations for Good Graphic Design. Technical Report, 1989.

[17] R. Chang, C. Ziemkiewicz, R. Pyzh, and J. Kielman, Learning-Based Evaluation of Visual Analytic Systems. *BELIV '08 (CHI)*, 2008.

[18] M.T.H. Chi, M. Bassok, M.W. Lewis, P. Reimann, and R. Glaser, "Self-explanations: How students study and use examples in learning to solve problems, *Cog. Sci.*, vol. 13, no. 2, 1989.

[19] M.T.H. Chi, Self-explaining expository texts: The dual processes of generating inferences and repairing mental models. In R. Glaser (Ed.), *Advances in Instructional Psychology*, Hillsdale, NJ: Lawrence Erlbaum Associates 2000.

[20] W. S. Cleveland and R. McGill, Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods, *J.of the Amer. Stat. Assoc.*, vol. 79, no. 387, 1984.

[21] J.L. Davenport, D. Klahr, & K. Koedinger, The influence of external representations on chemistry problem solving. Presented at *the Ann. Meeting of the Psychonomic Soc.,* 2006.

[22] N. Diakopoulos, F. Kivran-Swaine, and M. Naaman, Playable Data: Characterizing the Design Space of Game-y Infographics, *CHI* (forthcoming), 2011.

[23] C. Diemand-Yauman, D. M. Oppenheimer, and E. B. Vaughan, Fortune favors the **Bold** and the *Italicized*: Effects of disfluency on educational outcomes, *Cognition*, vol. 118, no. 1, 2011.

[24] ESP game. GWAP.com, 2009. http://www.espgame.org.

[25] F.L. Fahy, I. P. Riches, and M.W. Brown, Neuronal activity related to visual recognition memory: long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex, *Exp.. Brain Research.*, vol. 96, no. 3, 1993.

[26] L. Findlater, K McGrenere, and J. Dawson, Ephemeral adaptation: the use of gradual onset to improve menu selection performance, *CHI*, 2009.

[27] D.J. Gillan & E.H. Richman. Minimalism and the Syntax of Graphs. *Human Factors*, vol. 36, no. 4, 1994.

[28] D. Hammond, Effectiveness of cigarette warning labels in informing smokers about the risks of smoking: findings from the International Tobacco Control (ITC) Four Country Survey, *Tobacco Control*, vol. 15, no. 3, 2006.

[29] J. Heer and G. Robertson, "Animated Transitions in Statistical Data Graphics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, 2007.

[30] M. Hegarty, Representations i*n the Mind and in the World: How Cognitive Science Can I*nform the Design of Visualizations, Keynote presentation, *INFOVIS '10*, 2010.

[31] M. Hegarty and D.Waller. Individual Differences in Spatial Abilities. In P. Shah and A. Miyake (Eds), *The Cambridge handbook of visuospatial thinking*. Cambridge University Press, 2005.

[32] M. Hegarty, S. Kriz, & C. Cate, The roles of mental animations and external animations in understanding mechanical systems. *Cognition & Instruction, 21*, 2003.

[33] .T. Höffler & D. Leutner, Instructional Animation Versus Static Pictures: A Meta- Analysis. *Learning and Instruction*, vol. 17, 2007.

[34] W. Huang et al.. (Eds.), Measuring effectiveness of graph visualizations: A cognitive load perspective, *Information Visualization*, vol. 8, no. 3, 2009.

[35] O. Inbar, N. Tractinsky, and J. Meyer, Minimalism in information visualization: attitudes towards maximizing the data-ink ratio, in *Proc. of the Euro. Conf. on Cog. Ergonomics*, 2007,

[36] R. Kosara, Visualization Criticism – The Missing Link Between Information Visualization and Art, *IV '07*, 2007.

[37] S. M. Kosslyn, Understanding charts and graphs, *Applied Cognitive Psychology*, vol. 3, no. 3, 1989.

[38] H. Lam, A Framework of Interaction Costs in Information Visualization, *IEEE TVCG*, 2008.

[39] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, & S. Carpendale. Seven Guiding Scenarios for Information Visualization Evaluation. Technical Report, 2010.

[40] J. H. Larkin and H. A. Simon, Why a Diagram is (Sometimes) Worth Ten Thousand Words, *Cog. Sci.*, vol. 11, no. 1, 1987.

[41] E. Levy, J. Zacks, B. Tversky, and D. Schiano, Gratuitous graphics? Putting preferences in perspective, in *CHI'96*, 1996.

[42] G.L. Lohse, The role of working memory on graphical information processing, *Behaviour & Information Technology*, vol. 16, no. 6, 1997.

[43] G.L. Lohse, A cognitive model for the perception and understanding of graphs, *CHI'91*, 1991.

[44] Z. Liu and J. T. Stasko, Mental models, visual reasoning and interaction in information visualization: a top-down perspective, *IEEE TVCG*, vol. 16, no. 6, 2010.

[45] R.E. Mayer, M. Hegarty, S. Mayer, and J. Campbell, When static media promote active learning: annotated illustrations versus narrated animations in multimedia instruction, *J. of Exp. Psych.: Appl.*, vol. 11, no. 4, 2005.

[46] T. Munzner, A Nested Model for Visualization Design and Validation. *IEEE TVCG*, vol. 15, no. 6, 2009.

[47] H. M. Natter and D. C. Berry, Effects of active information processing on the understanding of risk information, *Appl. Cog. Psych.*, vol. 19, no. 1, 2005.

[48] C. North. Visualization viewpoints: Toward measuring visualization insight. *IEEE Comp. Graphics and Applic.*, vol. 26 no. 3, 2006.

[49] O. Pagliano and D. M. Bajzek, Improving Animation Tutorials by Integrating Simulation, Assessment, and Feedback to Promote Active Learning, in T. Bastiaens & S. Carliner (Eds.), *Proc. of World Conf. on E-Learning in Corp., Gov't, Healthcare, and Higher Educ.*, 2007.

[50] A. Perer and B. Shneiderman, Systematic yet flexible discovery: guiding domain experts through exploratory data analysis, in *Proc. on Intelligent User Interfaces*, 2008.

[51] S. Pinker. A theory of graph comprehension. In R. Friedle (Ed.), *Artificial intelligence and the future of testing*. Hillsdale, NJ: Erlbaum, 1990.

[52] H. Purchase, E. Hoggan, and C. Gorg. How important is the Mental map? an empirical investigation of a dynamic graph layout algorithm. In *Graph Drawing*, 2006.

[53] G. G. Robertson, S. K. Card, and J. D. Mackinlay, Information visualization using 3D interactive animation, *Comm. of the ACM*, vol. 36, 1993.

[54] G. Robertson, R. Fernandez, D. Fisher, B. Lee, and J. Stasko, Effectiveness of Animation in Trend Visualization, *IEEE TVCG*, vol. 14, Nov. 2008.

[55] M. Roy & M.T.H. Chi. The self-explanation principle. In R.E. Mayer (Ed.) *Cambridge Handbook of Multimedia Learning*, 2005.

[56] L. Rozenblit and F. Keil, The misunderstood limits of folk science: an illusion of explanatory depth, *Cog. Sci.*,vol. 26, 2002.

[57] S.H. Schwartz and W. Bilsky, Toward A Universal Psychological Structure of Human Values, *J. of Personality and Social Psych.* vol. 53, no. 3, 1987.

[58] P. Shah and P. A. Carpenter, Conceptual Limitations in Comprehending Line Graphs, *J. of Exp. Psych.: General*, vol. 124, no. 1, 1995.

[59] P. Shah & E. Freedman, Bar and line graph comprehension: An interaction of top-down and bottom-up processes. *Topics in Cog. Sci.*, 2009.

[60] P. Shah, R.E. Mayer, & M. Hegarty. Graphs as aids to knowledge construction: Signaling techniques for guiding the process of graph comprehension. *J. of Educ. Psych.,* vol 91, 1991.

[61] M. Siegrist, The use or misuse of three-dimensional graphs to represent lower-dimensional data, *Behav. & Info. Tech*, vol. 15, no. 2, 1996.

[62] C.S. Skinner, V.J. Strecher, and H. Hospers, Physicians' recommendations for mammography: do tailored messages make a difference? *Amer. J. of Pub. Health*, vol. 84, no. 1, 1994.

[63] D. Sprague & M. Tory. Motion in Casual InfoVis and the Interrelation between Personality, Performance, and Preference. Technical Report, 2009.

[64] Tableau, Tabeau Software Inc., http://www.tableausoftware.com/.

[65] J. G. Trafton, S. B. Trickett, and F. E. Mintz, Connecting Internal and External Representations: Spatial Transformations of Scientific Visualizations. *Foundations of Science*, vol. 10 no. 1, 2005.

[66] E. R. Tufte, *The Visual Display of Quantitative Information*, 2nd ed. Graphics Press, 2001.

[67] H. Wainer, *Picturing the Uncertain World: How to Understand, Communicate, and Control Uncertainty through Graphical Display*. Princeton University Press, 2009.

[68] C. Ware, *Visual Thinking: for Design*, First Edition. Morgan Kaufmann, 2008.

[69] K. Yee and D. Fisher, Animated Exploration of Dynamic Graphs with Radial Layout, *INFOVIS*, 2001.

[70] J. Zacks, E. Levy, B. Tversky, and D. J. Schiano, Reading Bar Graphs: Effects of Extraneous Depth Cues and Graphical Context, *J. of Exp. Psych.: Applied*, vol. 4, no. 2, 1998.