A Framework of Interaction Costs in Information Visualization

Heidi Lam

Abstract—Interaction cost is an important but poorly understood factor in visualization design. We propose a framework of interaction costs inspired by Norman's Seven Stages of Action to facilitate study. From 484 papers, we collected 61 interaction-related usability problems reported in 32 user studies and placed them into our framework of seven costs: (1) Decision costs to form goals; (2) System-power costs to form system operations; (3) Multiple input mode costs to form physical sequences; (4) Physical-motion costs to execute sequences; (5) Visual-cluttering costs to perceive state; (6) View-change costs to interpret perception; (7) State-change costs to evaluate interpretation. We also suggested ways to narrow the gulfs of execution (2–4) and evaluation (5–7) based on collected reports. Our framework suggests a need to consider decision costs (1) as the gulf of goal formation.

Index Terms—Interaction, Information Visualization, Framework, Interface Evaluation

1 Introduction

Even though interaction is vital to interface success, the information visualization community has generally focused more on visual encoding than on interaction. Ideally, visualization designers should be able to weigh costs and benefits of interaction based on empirical results. Before we can effectively evaluate interaction, we need to first understand how it can contribute to visualization use and how designs can fall short in supporting these roles. While taxonomies of interaction techniques exist to study roles of interaction in interface use (see [69] for a survey), we still need a framework to study interaction costs.

Similar to understanding interaction techniques, one of the first steps in understanding interaction costs is to identify instances. To better design, we need to take an holistic approach and study interaction during interface use [3] instead of focusing on individual technique in isolation and with abstract tasks (e.g., [27]). Typical user studies seldom explicitly measure or even identify interaction costs, but reports on recorded usage patterns, participant strategies, and interface choice sometimes provide insights.

In this paper, we propose a framework of seven interaction costs based on cost reports gathered using a qualitative review. From 484 papers, we identified 32 that reported interaction-related usability issues. Reports collected were placed into a framework of action cycles in visualization use inspired by Norman's Seven Stages of Action [39]. We highlight interaction-design considerations to narrow Norman's gulfs of execution and evaluation, and propose adding a gulf of goal formation to study decision costs in establishing data-analysis focus.

After related work in Section 2, we summarize our framework in Section 3. Section 4 describes our review method and Section 5 elucidates our framework. Section 6 puts forth design considerations to mitigate some of the interaction costs.

2 Background and Related Work

For the purpose of this paper, we define **interaction** as actions from users that cause visible changes in the visualization, and **interaction techniques** as those actions. Since we gather our cost instances from papers, we only considered observable interaction in typical evaluations, ignoring non-action communications (e.g., eye-gazes) and unsolicited system actions (e.g., alerts). To us, an **interaction cost** is when the dialogue between users and system breaks down, or where users face enough difficulty accomplishing tasks to become aware of the user interfaces as obstacles to be overcome [67].

• *Heidi Lam is with the University of British Columbia, E-mail: hllam@cs.ubc.ca.*

Manuscript received 31 March 2008; accepted 1 August 2008; posted online 19 October 2008; mailed on 13 October 2008.
For information on obtaining reprints of this article, please send e-mailto:tvcg@computer.org.

One approach in framework development is to extend existing interaction taxonomies (see [69] for a survey) to cover costs. However, interaction costs can be implementation dependent. For example, navigation implemented in pan-and-zoom interfaces may fail to support object constancy, while navigation implemented in overview+detail interfaces may incur view-coordination problems.

Instead, we prefer a user-centric model depicting action steps. Researchers have developed models of user interaction with different foci and abstraction levels. Spence's navigation framework [61] and Card *et al.*'s knowledge crystallization task [8, p.10] are cognitive models with insufficient user-visualization interaction focus for our purpose. Chi and Riedl's operator framework classified operators into stages between raw data (value) and visual representation (view) [11], but our focus is on view only. Jankun-Kelly *et al.*'s P-Set Model describes the exploration process at the interaction-technique level (e.g., zoom, rotate), but does not cover low-level motion (e.g., mouse drag) or high-level cognitions (e.g., result interpretation) that can also incur costs [28]. We therefore based our framework on Norman's Seven Stages of Action [39] for its comprehensive coverage of user actions.

3 A Framework of Interaction Costs

We adapted Norman's Seven Stages of Action [39, p.46–53] to visualization use and classified interaction-related publication statements, or **reports**, based on the stage of interaction at which they occur. Our framework has seven costs (Fig 1):

1. *Decision costs to form goals*: When interfaces become more powerful and display more data points, users usually need to decide to focus on a subset of data (Section 5.1.1) and interface options (Section 5.1.2).
2. *System-power costs to form system operations*: Once users have a question in mind, they need to translate it into operations. Deciding on the correct operation sequences may be difficult especially for powerful systems (Section 5.2).
3. *Multiple input mode costs to form physical sequences*: When the input device offers multiple modes, translating system operations to device operations may be difficult due to inconsistent mode operations on multiple views (Section 5.3.1), mode change with inadequate visual feedback (Section 5.3.2), and overloaded input controls (Section 5.3.3).
4. *Physical-motion costs to execute sequences*: Even for young and healthy users over short time spans, motions such as mouse position (Section 5.4.1) and mouse drag (Section 5.4.2) can incur costs. Even low-cost motions can collectively cumulate into usability problems (Section 5.4.3).
5. *Visual-cluttering costs to perceive state*: Interaction such as mouse hovering can cause visual cluttering that makes state perception difficult (Section 5.5).

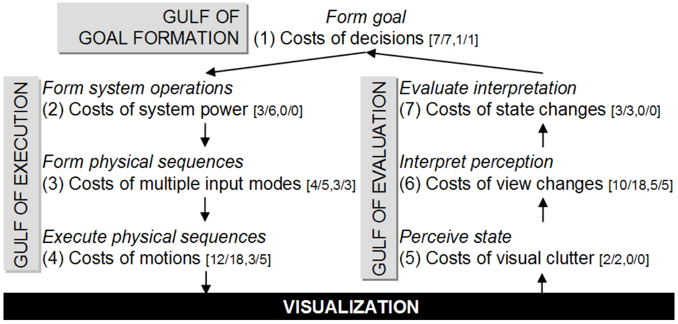

Fig. 1. A framework of interaction costs inspired by Norman's Seven Stages of Action (in italics) [39]. We added a gulf of goal formation to Norman's gulfs of execution and evaluation. Bracketed numbers are report and recommendation counts (included-in-this-paper/reviewed). Two general-observation reports are excluded in this figure.

6. *View-change costs to interpret perception*: Interaction usually result in view changes that requires re-interpretation based on expectations, which may not be met in automated systems (Section 5.6.1). Interpretation requires object association of: (1) temporal objects, as in zooming (Section 5.6.2); (2) spatial objects, as in view coordination (Section 5.6.3); and (3) local and global objects, as in navigation (Section 5.6.4).

7. *State-change costs to evaluate interpretation*: Data analysis often requires reflection on multiple data views or analysis states. Section 5.7 looks at the need for refinding in visualizations.

Norman's model has two gulfs. The **gulf of execution** is "the difference between the [user] intentions and the allowable [system] actions" [39, p.51], which covers costs (2–4) in our framework. The **gulf of evaluation** "reflects the amount of effort that the person must exert to interpret the physical state of the system and to determine how well the expectations and intentions have been met" [39, p.51], which covers costs (5–7). To emphasize data analysis and discovery supported in information visualization, we added a **gulf of formation** to cover cost (1), or the amount of effort required to formulate suitable intents. Section 6 discusses design considerations to narrow these gulfs.

Since Norman's model is an approximate model, behaviours are not required to involve all stages in sequence and each may take drastically different time durations [39, p.48]. For example, in information visualization, the steps may be carried out in rapid successions as in the case of dynamic queries [2]. In general, this paper treats cognitive-decision costs as gulf of goal formation, physical-motion costs as gulf of execution, and cognitive-interpretation costs as gulf of evaluation.

4 Methodology and Scope

Our work differs from most systematic reviews such as Chen & Yu's [10] as we did not start with specific questions. Instead, we took a bottom-up and qualitative approach to identify interaction costs from coded individual study reports. Our approach has three stages: venue selection; paper coding and selection; report and cost identification.

4.1 Venue selection

We reviewed papers from the IEEE Symposium on Information Visualization (InfoVis) proceedings, Palgrave's Journal of Information Visualization (IVS), Elsevier's International Journal of Human-Computer Studies (IJHCS), and ACM's Transactions of Computer-Human Interface (TOCHI). Our resource constraints limited our venue selection: we focused on journal publications as we believe researchers have more room to report observations.

4.2 Paper coding and selection

We coded all publications for InfoVis (1995–2007), IVS (2002–2007), and TOCHI (1994–2007). For IJHCS, we first filtered the papers with

	Category/Pub	InfoVis	IVS	IJHCS	TOCHI
	# papers	283	117	35	49
1.	No evaluation	178	50	8	12
2.	Case studies	48	32	1	4
3.	Informal studies (iSys|iReport)	18 (17|**2**)	5 (5|0)	3 (3|0)	1 (1|0)
4.	Qual study (iSys|iReport)	5 (5|**2**)	10 (10|**3**)	6 (6|**1**)	5 (4|0)
5.	Quant expt (iSys|iReport)	34 (20+2?|**8**)	20 (9+1?|**3**)	17 (11|**6**)	27 (20|**7**)

Table 1. Coded paper counts. For categories (3–5), we further classified if the study interface is interactive (iSys), and if the paper reported interaction use (iReport). '?' denotes cases where we were uncertain if the study interfaces were interactive. Qual study = qualitative evaluations with users and tasks; Quant expt = quantitative factorial experiments. Bolded numbers represent the 32 papers used in this review.

the query term "visualization" for title, abstract, or keywords and retrieved 49 papers between the years 1994 and 2007.[1]

For IJHCS and TOCHI, the two non-infovis specific venues, we included papers that focused on infovis subjects[2] and reported at least one visualization technique or system. We were then left with 26 IJHCS and 49 TOCHI papers, along with 283 InfoVis and 117 IVS papers. We coded the papers by their evaluation methods, as shown in Table 1. In short, we started with 484 infovis-related papers and ended up with 32 for our review.

4.3 Report and cost identification

From the 32 papers that reported interaction-related usability issues, we collected 61 reports of problems and 14 design recommendations. We took a bottom-up approach by grouping reports into costs and putting costs to action stages. Due to space constraints, we only included 42 reports and 12 recommendations from 27 papers here to illustrate our framework (Fig 1). The full list is at `http://www.cs.ubc.ca/~hllam/res_icost.htm`.

5 Interaction Costs

For each of the seven costs, we provide background literature and commonly-used solutions, followed by reports as examples and when appropriate, as test cases for these solutions. We included report authors' recommendations. Report-paper citations are denoted with '*'.

5.1 Decision costs to form goals

While command-line interfaces have been criticized for lack of cues to prompt user actions, modern flexible visualization systems may also share this problem. We found reports in two areas that require decision making: finding a data subset to explore (Section 5.1.1) and choosing amongst interface options (Section 5.1.2).

5.1.1 Choosing a data subset

When visualizations become more powerful and display more data points, users need to decide on data focus. Visualization can guide users by providing visual cues or information scent. Pirolli and Card developed the information foraging theory [44] that has been applied to designs of web-searching support tool designs (e.g., [43]) and to model web usage behaviours (e.g., [7, 12]). Their later Sensemaking process [45] incorporated Russell *et al.*'s sensemaking loop [52] with their foraging loop to depict intelligence analysis.

However, to provide information scents, the software needs to address the challenge in predicting user interests and intent (Section 5.4.1). One workaround is to simply provide more information

[1] Last query performed on December 12, 2007

[2] In our analysis, we excluded subjects such as theory (e.g., cognitive models, design architectures, and evaluation methods), human-behavioural studies (e.g., ethnographic studies), non-visual interfaces (e.g., audio), virtual or augmented reality, computer-assisted collaborative work focusing on human-human communications (e.g., designs of avators), and special issue introductory papers.

such as usage statistics (e.g., scented widgets [66]) and metadata to faciliate data reordering and selection (e.g., Table Lens [48]).

We found reports where users were lost in data due to insufficient information displayed. Kobsa reported that with Eureka,

> Since attributes in Eureka are vertically aligned, there is not very much room for attribute labels if data has more than, say, 20 dimensions. In this case, users have troubles making sense of the data and finding the attributes they need since the attribute labels on top of the columns are largely hidden. [*32, p.127]

Similarly, Abello *et al.* found that users did not know where to go when exploring a large unknown dataset using ASK-Graph:

> Users would be able to navigate around just fine, but had no idea where they should start to look for interesting features. As a result they sometimes stumbled upon something interesting, but spent most of their time randomly browsing the data. [*1, p.67]

In the context of menu design, Hornbæk and Hertzum recommended increasing information scent by expanding ahead:

> Perhaps hierarchical menus could also benefit from some of the ideas used for increasing the number of simultaneously visible menu items in nonhierarchical menus. Submenus could, for example, reveal more of their contents by dynamically expanding two levels of the menu structure in response to cursor movements. [*22, p.28]

5.1.2 Choosing amongst interface options

In some cases, users may form their goals based on available interface options. Even though human intuition may want more choices, decisions require mental and visual concentration [55]. Indeed, participants have expressed their preference to have less interface options, such as in Chung's evaluation of SBizPort:

> The subjects also liked the user-friendliness and ease of navigation of SBizPort. Nine subjects commented positively on it. For instance, a subject said that SBizPort's "categories make it easier to browse" and another subject said not having too much (many) options (in SBizPort) is helpful to the user because it makes it easy to look for the information. [...]
>
> In comparison, subjects were less satisfied with YahooES than SBizPort. They had difficulty obtaining precise information from YahooES. Ten subjects said that YahooES had too many options to choose from and hence distracted them from finding relevant information. [*13, p.821]

Participants may also express their preference by choosing less interactive but suboptimal interfaces:

> We believe this interesting choice [for single-level interfaces] was due to the cost of interface interaction complexity, which may also explain the lack of performance benefits over the optimal single-[level] interface for each task. Although seemingly tedious and laborious, using the high-[detail] plots has a low cognitive load: the only navigation available is scrolling, a relatively passive exercise, and the answer will usually be apparent sooner or later. In contrast, navigation in a multiple-[level] interface is complex, as it involves active selection of potential target candidates, an action that requires mental and visual concentration. [*34, p.1284–1285]

5.2 System-power costs to form system operations

Once users have a question in mind, they need to translate it into operations. Deciding on the correct sequences of operations may be difficult especially when the visualization offers a large set of operations. In evaluating Spotfire, Kobsa reported a cognitive set-up cost:

> While Spotfire offers several representations in parallel, in many cases not all of them are suitable for solving a given problem. It took users considerable time to decide on the right representation and to correctly set the coordinates and the parameters, particularly when the solutions required several steps. This seems to be caused both by the wealth of visualizations that the system offers, but also by the restrictions each of them imposes once it has been selected. [*32, p.127]

Without a clear set of available operations, users may have expectations based on standards, as found by Kang *et al.*'s NetLens study:

> [users observed an inconsistency that] while multiple histogram bars are selectable, using Shitft+Click is not supported [in NetLens]. [*30, p.29]

Users therefore assumed the burden to find the correct operations:

> Users must therefore plan in advance what variables should be used and how they should be represented. This planning must be performed without assistance from a visualization and takes up considerable time. [*32, p.129]

5.3 Multiple input mode costs to form physical sequences

Having multiple input modes or overloaded controls can lead to mode errors, where "the action appropriate for one mode has different meanings in other modes" [39, p.110].

The guideline to eliminate multiple input modes is impossible to follow when the input device needs to provide more actions than the number of available controls. The consequence may be costs in inconsistent mode use (Section 5.3.1), imperceptible mode changes (Section 5.3.2), or when overloading a single control (e.g., for parallel panning and zooming; Section 5.3.3).

5.3.1 Inconsistent mode use on multiple views

Users form habits with interface use. One interaction cost is due to inconsistent mode use on different views. For example, when studying zoomable user interfaces, Hornbæk *et al.* reported that,

> Subjects' habit formation highlighted some limitations in the interfaces. At least eight subjects tried to use a way of navigating from the overview window in the detail window or vice versa. Some subjects tried to click on the detail window, probably with the intention of jumping to the place where they clicked. This way of navigating seemed to be taken from the overview window, where clicking on a point centers the field-of-view box on that point. Similarly, some subjects tried to zoom in and out while they had the mouse over the overview window. This way of interacting seemed to be mimicked after the interaction with the detail view. [*23, p.381]

Hornbæk *et al.* suggested eliminating view-specific commands:

> We believe that interfaces with an overview should eliminate navigation commands that are specific only to the overview window or to the detail window, that is, they should aim at unifying navigation. All zoom and pan actions should therefore be similar across windows. [*23, p.384]

5.3.2 Imperceptible mode changes

Sometimes, mode change may not be intended or even noticed by users, especially with inadequate visual feedback. Hornbæk and Hertzum, in their study of Bederson's fisheye menu [4], reported that,

> [W]hen in focus-lock mode, many participants wanted to keep scrolling up or down toward an item they had seen before entering this mode. However, in the fisheye and multifocus menus, items visible in small font moved out of the menu as items in the transition region expanded. [...] Although visually indicated, these modes seemed to confuse participants. Especially when in focus-lock mode and accidentally crossing the center of the menu, at least seven participants expressed confusion when the focus area was consequently dynamically recentered to the mouse position. [*22, p.19]

Hornbæk and Hertzum suggested making the change continuous:

> Our data show that the binary nature of this mode caused participants problems. A simple idea would be to use a continuum instead. When the mouse is moved toward the righthand side of the menu, the selection height of menu items would increase toward a maximum of their visual height.[*22, p.28]

Or to make the mode change explicit by requiring user actions:

> [With the] quasimode [the user enters the] focus-lock mode when [he] presses the mouse button. This would lessen the possible confusion of modes by turning the focus lock on when users are about to complete their selection, thereby enabling any final adjustments of mouse position to be made at maximum selection height before users release the mouse button to select the target menu item. [*22, p.28]

5.3.3 Overloaded controls

Overloading input controls can lead to confusion and errors, even when users are aware of the operations. Hornbæk *et al.* found that with zoomable user interfaces,

> Many subjects experienced occasional problems with the combined zoom and pan button. Even though subjects practiced this combination button during the training tasks, 18 subjects zoomed at least one time when they verbally indicated that they were trying to pan. The delay before zooming began was sometimes too short. This appeared to happen when subjects began initiating a pan action without having made up their minds about which direction to pan. [*23, p.381]

Buering *et al.* also found similar problems with zoomable user interfaces in PDA-sized devices:

> Seven subjects mentioned that they had problems with the sliding technique of the ZUI. [...] As a result some subjects accidentally triggered a zoom operation when actually trying to slide. [*6, p.834–835]

How to best provide parallel zooming and panning remains unclear:

> Research is needed to find a method for interacting with zoomable user interfaces using a two-dimensional input device that is intuitive and supports habit formation. [...] Ideally, zooming and panning should be allowed to take place in parallel. [*23, p.384]

5.4 Physical-motion costs to execute sequences

Physical motion can be hard when the display is very large [14] or very small [49], or for specific user populations such as children [25] and the elderly [60]. Even with desktop-sized displays and physically-apt users tested over limited periods of time, we still found reports of user dissatisfaction with commonly-used mouse interaction such as positioning (Section 5.4.1) and dragging (Section 5.4.2). Furthermore, even low-cost simple motions can accumulate (Section 5.4.3).

5.4.1 Costs in Mouse Position

Fitts' Law models the ease of target selection [16]:

$$MT = a + b\log_2(A/W + 1) \quad (1)$$

where MT is average movement time, A is separation between the two targets, W is target width, and a and b are experiment constants. Fitts' Law implies that the less distance traveled and the bigger the target, the easier it is to reach. Despite its simplicity, applying Fitts' Law to interaction design often involves tradeoffs such as the need to predict user intents to optimize A; trading display capacity for larger W; and causing targets to move by improving A or W.

1. Reducing travel distance needs user intent
The challenges to minimize A are to predict user intent and preserve target locations, since moving targets are hard to select (see below). Predicting user intent can be difficult even for 1D lists. Sears and Shneiderman's Split menu, a menu added with a separate frequency ordering on top of the default alphabetical ordering, showed benefits over alphabetical menus only when frequently selected items were at the bottom of the list [*56].

2. Sacrificing target size for display capacity
Display capacity is a design priority especially for space-limited devices (as in mobile devices) or for large-data displays. We found reports when W was reduced too far and resulted in usability problems.

To accommodate up to six months' of calendar information on PDA-sized screens, Bederson *et al.* used a fisheye technique to create DateLens [*5]. With semantic zooming, users of DateLens can view each day in the calendar at various levels of detail. However, their study found small targets as the most serious usability issue:

> Experienced Pocket PC users often use their fingers to tap on targets in the user interface [...] Even with the stylus, users often invoked incorrect behaviors and actions accidentally when attempting to scroll or make UI selections in DateLens. [*5, p.115]

Indeed, we found similar reports from researchers that tested distortion-based desktop interfaces. Hornbæk and Hertzum examined the usability of Bederson's fisheye menu, where the menu-item font is reduced to increase display capability and reduce traveling distance [4]. Hornbæk and Hertzum found a cost with this optimization,

> the larger the selection height of menu items, the lower the selection time. [*22, p.27]

Li and North looked at dynamic query (DQ) sliders and brushing histograms and found that DQ sliders were more efficient than brushing histograms in range and criteria tasks, possibly because,

> The targets [in the brushing histograms] for clicking were narrower and smaller compared with DQ sliders. Thus it was easier for users to make incorrect or accidental selections [using the brushing histograms]. [*35, p.152]

In addition to ensuring selectable target sizes, Bederson *et al.* proposed another general guideline,

> allow the user to adjust the font size [...] in novel visualizations [*5, p.117]

3. Increasing target size may cause targets to move
Reducing travel distance A and target width W in Fitts' Law (Eqn 1) may cause targets to move. For single isolated targets, McGuffin and Balakrishnan showed that increasing W resulted in faster selection time [*36]. For multiple expanding targets, especially when the targets are closely packed together as in most visualization, there are tradeoffs. For example, distortion-based visualization technique such as fisheye can cause target expansion as the pointer moves closer:

> The problem that occurs when focus-targeting in fisheye views is that targets appear to move in the opposite direction to the motion of the magnifying lens. This means that a focus target will move towards an approaching pointer, and away from a retreating one, making it more difficult to precisely position the focus point relative to the underlying visualized data.
>
> Moving targets are always more difficult to hit—but to make matters worse, the gradually increasing magnification of a fisheye lens makes targets move faster and faster the closer the focus comes to them. In fact, targets move at their highest rate of apparent speed at the exact moment that the pointer nears the target, making it difficult for a user to precisely position the pointer over the target. [18, p.267–268]

Gutwin suggested speed-coupled flattering as a solution to ensure that "relatively static [view] during the acquisition phase of targeting, in order to simplify precise positioning" [18, p.269]. While Gutwin showed the effectiveness of speed-coupled flattering [18], focus lock in Bederson's fisheye menu [4] was found to be a substantial usability problem in Hornbæk and Hertzum's 2007 study [*22]. Focus lock is a mechanism where users,

> move the pointer to the right side of the menu, which locks the focus on the item the cursor is over. Then, when users move the pointer up and down, the focus stays fixed, but individual menu elements can still be selected. The focus region on the right side of the menu gets highlighted to indicate that the menu is in focus lock mode. [4, p.220]

In practice, Hornbæk and Hertzum observed difficulty in its use,

> One reason for the higher selection times [for the fisheye-based menus] appears to be an occasionally ineffective use of the focus lock. [...] This finding accords with the observation that participants sometimes had to leave the focus-lock mode because an item of interest was pushed out of view as a result of the expansion of menu items close to the mouse. Another reason is that participants may choose to enter the focus-lock mode only after having faced difficulties in acquiring the target. The time penalty associated with the focus lock suggests that the best time to make the shift to focus-lock mode was not obvious to participants. [*22, p.23]

In short, actual pixel space does not directly translate to selection space. Even though expanding targets may facilitate coarse navigation, actual target acquisition is more difficult with moving targets. Hornbæk and Hertzum thus concluded that,

> [Since i]tems are moving, the number of pixels in the motor space from which an item must be selected is lower than that of the hierarchical menu. [...] Stable position of menu items are central to the usability of the hierarchical menu.[*22, p.24]

Indeed, McGuffin and Balakrishnan also concluded that,

> The model [of expected benefit in multiple expanding targets] presented indicates that a net reduction in selection time with tiled expanding targets may be possible, however, in practice, the benefit may be negligibly small. [*36, p.419]

5.4.2 Costs in Mouse Drag

Mouse drags are almost ubiquitous on modern interfaces. Nonetheless, reports suggest limiting its use.

In a study to evaluate the PDQ Tree browser, Kumar *et al.* found that the panning action on the detail view, achieved by dragging on the field-of-view in the overview, was not universally welcomed:

> One subject emphasized the need to always fit the overview into one screen only, so that no scrolling of the overview is required. [...]
>
> Another subject suggested that users should be able to click anywhere in the overview and have the field-of-view jump to that position. This would enable fast coarse navigation [without mouse drag]. Fine-tuning could then be accomplished by dragging the field-of-view. [*33, p.119–120]

In PDQ's dynamic query panel, users filter by selecting tree-node attributes from a list using drag-and-drop. While some participants enjoyed the drag-and-drop mechanism, one obviously did not: "Drag-and-drop becomes a drag for experienced users, so drop it!" This sentiment was shared by others:

> Some other subjects also echoed the feeling that it might be easier and faster to just replace each drop area with a menu of attributes at that level. [*33, p.120]

Dragging is also used to specify bounding boxes for zooming. In a study of fisheye and zoomable interfaces on PDA-sized devices, Buering *et al.* reasoned that since,

> the fisheye interface required far fewer actions but, since task times were similar, it seems that they required more time to execute. Hence we assume that drawing a bounding box is cognitively more demanding than the more direct zooming of the ZUI. [*6, p.834]

5.4.3 Costs in accumulated motions

Even for simple movements such as mouse clicks, repeated actions can accumulate into measurable costs. In evaluating visualization systems to display microarray data, Saraiya *et al.* found that for GeneSpring,

> even though users tended to focus on a small number of basic visualization features, usability issues (such as the higher quantity of clicks required to accomplish tasks) reduced their overall insight performance. [*53, p.7]

We found a similar report from Hetzler *et al.* in their In-Spire study:

> The ability for a user to track a theme over time, while quite doable, involved too much manipulation and user intervention. This prompted the addition of the Keep Current capability [...]. [*20, p.94]

5.5 Visual-cluttering costs to perceive state

Interaction can cause visual distraction and occlusion that make perception difficult. For example, mouse hovering, despite providing tool-tip guidance, can cause unwanted visual distraction. Granitzer *et al.* noticed that in InfoSky when mouse hovering

> near the bottom of the hierarchy, where collections contained many documents, users were confused by the 'jumping around' of document titles. The prototype displayed the titles of those documents which were 'near' to the cursor. [*17, p.130]

In a study of Spotfire, Saraiya *et al.* found that,

> Spotfire's parallel coordinates view employs a poorly designed selection mechanism. Selected lines in its parallel coordinates results in an occluding visual highlight that made it very difficult for users to determine which genes were selected. [*53, p.7]

5.6 View-change costs to interpret perception

Interaction often causes change in the visual display. Users therefore need to associate objects in the old view to those in the new based on their expectations. Augmented interaction using machine intelligence may fail to meet these expectations and causes user confusion or even distrust (Section 5.6.1). Interpretation of changes is dependent on implementations. We found reports in three cases: costs in object association between between temporal frames (Section 5.6.2), between simultaneously displayed views (Section 5.6.3), and between local and global objects (Section 5.6.4).

5.6.1 Augmented interaction

To offload cognitive costs in dealing with large data, one design option is to offer automatic data processing. However, users have expectations as to what should happen after an interaction. Failure to meet expectations can lead to confusion or even distrust.

We found such reports in Siirtola and Mäkinen's study on the automatic reorderable matrix [*59]. Their participants rated the subjective satisfaction question, *Overall, the experimental application was easy to use and performed as expected, and did not do unexpected things*, one out of five, possibly because,

> When interviewed, four of the participants said that the reason for the rating was the uncomfortable mouse behavior—especially the proportional acceleration was different from what they were used to. The other reason was the heuristic nature of the reordering algorithm. For some participants, it was difficult to accept that the same setting of the slider would sometimes produce a new ordering. [*59, p.46]

Siirtola and Mäkinen also reported confusion in slider effects on matrix ordering based on their algorithm:

> Some of the participants commented during the experiment and in the interview that the continuous reordering feels a bit disturbing, although this does not show in the questionnaire results. They felt that a small change in the slider should not result a major change in the matrix ordering. This was the initial reaction, and most of the participants were able to accept this later in the experiment as a characteristic of the user interface. This behavior is due to the nature of the barycenter heuristic and cannot be avoided. [*59, p.47]

5.6.2 Temporal-frame association

When visual objects change with time, users need to keep objects in memory for association. Smooth animation has been proposed as a solution to connect different temporal views in zooming interfaces to preserve object constancy [51] and applied with success (e.g., [31]). Others solutions include minimizing visual changes [38] and providing visual cues such as background grids or landmarks [9, 70].

Our review indicates that animation alone may be inadequate. Siirtola and Mäkinen reported in a study of automatic reorderable matrix,

> Based on post-discussion, it was determined that subjects found the feedback to be inadequate during the row and column movements. There should be an outline of row or column visible during the move operation to indicate what is moving. The current implementation updates the matrix view as the mouse moves, but does not indicate the current selection. [*59, p.42–43]

McGuffin *et al.* suggested minimizing changes to allow landmarking:

> Features that allow the user to manually position or lock down. the relative placement of nodes would help alleviate the detrimental effects of rearrangement and allow for better landmarking and more consistent displays, thus reducing the time necessary to visually scan for nodes. [*37, p.125]

5.6.3 Multiple-view association

Interfaces may display different views of the same data, either at multiple visual levels as in overview+detail or focus+context interfaces [*34], or in different forms, as in multiform interfaces [50]. Traditionally, designers have used interaction techniques of brushing and linking to coordinate between the different views [64].

For example, in their Snap-together visualization study, North and Shneiderman found that coordination between overview and detail view enabled the use of the overview:

> If only the overview information is needed, then naturally coordination is not necessary. But for the important cases where access to details is needed, then coordination is critical. [*42, p.737]

However, we also found reports that suggest linking and brushing alone may be insufficient. In studying zoomable user interfaces, Hornbæk *et al.* reported problems in associating between the overview and the zoomable views, which resulted in slower task time for the interface with an overview because,

> switching between the detail and the overview window required mental effort and time moving the mouse. Our data modestly supports this explanation, since the number of transitions between overview and detail window were positively correlated with task completion time. [*23, p.382]

Another consideration is the amount of space needed for interaction:

> Coarse [navigation] and [...] resizing the field-of-view box could be difficult at low zoom factors. Subjects commented that the overview was hard to resize. In support of those comments, we note that the overview window used in the experiment occupied 256x192 pixels. When a zoom factor of 20 was reached, the field-of-view box was only 13x10 pixels, which was probably hard for most users to resize and move using the mouse. [*23, p.382]

These observations led to a design recommendation:

> To obtain the benefit of easy navigation provided by overviews, designers should use overviews at least one-sixteenth the size of the detail window (in area). For overviews coupled to a detail view less than the size of one screen or for screens on small devices, the overview might need to be larger to support navigation. For systems where much navigation is expected on the overview, for example, in support of monitoring tasks, a larger overview should be provided. For systems with zoom factors over 20 as used in our system, more usability problems will occur when using the overview, and consequently a larger overview will be necessary. [*23, p.384]

Another recommendation is landmarking as semantic linking between views in Crampes *et al.*'s KMap:

> There should be different KMaps according to the tasks, but that all KMaps should be semantically linked to maintain users' mental map. [*15, p.222]

5.6.4 Local-global association

When the display space contains little or no information for navigation, the problem of desert fog occurs [29]. Navigation becomes difficult once users have lost association between local and global objects.

In studying zoomable user interfaces, Hornbæk *et al.* observed that,

> At least six subjects repeatedly experienced what has been called desert fog, that is, they zoomed or panned into an area of the map that contained no map objects. [*23, p.381]

Kumar *et al.* observed in their PDQ Tree-browser study that,

> subjects get somewhat disoriented when the level of the tree was changed. This is because the layout algorithm generates a fresh layout whenever the tree structure changes, i.e. whenever more or less levels are requested to be seen. [*33, p.119]

In contrast, simpler navigation led to better interface use. Westerman *et al.*, in a document-browsing study with 2D and 3D display, hypothesized that,

> participants found the process of navigation less effortful in the two-dimensional condition, and therefore were prepared to adopt a more exploratory' strategy. [*65, p.731]

Simple navigation of the hierarchical menu in Hornbæk and Hertzum's study was also found to support better performance:

> Participants perform well with the hierarchical menu [as] it simplifies navigation. With fisheye and overview menus, participants made longer fixations, suggesting increased mental activity, compared with the hierarchical menu. Also, participants' scanpaths were longer with the multifocus and overview menus, indicating more visual search. Reasons for this could include: (a) the need with nonhierarchical menus to determine or remember which part of the menu structure one is currently in [...] [*22, p.25]

In addition to navigation simplicity, current focus can be used as landmarks to aid navigation in temporal view changes:

> It is felt that this [navigation] problem can be significantly alleviated by retaining the same current focus. For example, if the user asks to see the University level while the state Florida is near the center, the new view should be initialized to show universities within Florida. [*33, p.119]

Landmarks may also be useful aids in multiple spatial views:

> a detail-only interface could include cues about the current zoom factor, cues about the current position in the information space, and aids for avoiding desert fog. If such cues are integrated into the detail view, the mental and motor effort associated with shifting to the overview might be reduced, as would the screen real estate lost due to the presence of an overview. [*23, p.384]

5.7 State-change costs to evaluate interpretation

Data exploration often involves comparing between previously viewed data projections. Refinding has been studied in the context of web search (e.g., [63]). Lack of refinding support may inhibit exploration.

In our review, we found reports to suggest that in general, overview+detail interfaces are better than zooming or fisheye interfaces in supporting refinding, since participants explored more when using overview+detail interfaces. One such report is Hornbæk and Frokjær's document-reading study, where,

> The overview pane supports jumping directly to targets; it helps returning to previously visited parts of the document; and it invites and supports further explorations. Subjects using the fisheye interface depend extensively on the algorithm that determines which sections to collapse initially, even though subjects do not trust this algorithm. [*24, p.142–143]

Plumlee and Ware also found different interface-use strategies between the zooming and the multiple-window (WM) interfaces because of their differing support on refinding:

> The results show that subjects made dramatically more visits with the eye between windows than they made with the zooming interface. In addition, subjects made more eye-visits (in the multiple-window condition) than the model predicted would be necessary to achieve perfect performance.
>
> This suggests a kind of satisficing strategy with visual working memory as a limited-capacity, cognitively critical resource. When visits are cheap in time and cognitive effort, for example when they are made via eye movements, they are made frequently and people make a separate eye movement to check each component of the two patterns they are comparing. Thus their visual WM capacity relating to the task is effectively one. However, when visits are expensive in time and cognitive effort, for example when zooming is required, subjects attempt to load more information into visual WM and they also quit the task after fewer visits, which results in many more errors. [*47, p.205]

When the interface does not support refinding well, users may find it difficult to return to a state. Yi *et al.* reported this problem in their study of Dust & Magnet (DnM):

> Because DnM allows users a high degree of freedom in adjusting and manipulating dimensions, it is challenging to explicitly document certain clustering schemes. Yet, if the user keeps applying the 'Center Dust' feature, it is possible to regenerate similar clustering. However, the dimensions of DnM can move all over the main view, so DnM does not have the reproducibility that other visualization techniques have. [...] [T]he problem of reproducibility is really a trade-off with the heightened ability for users to freely explore and manipulate the data with a high degree of freedom. [*68, p.255]

6 Discussion: Narrowing the Gulfs

Interaction is vital to the success of modern visualization. Interaction has been found to benefit users by allowing multiple data views, for example, in controlling rotation [26] and object transformations [57]. Also, interaction can be fun for users, as seen in the use of Vizster social situations [19]. However, even though visualization evaluations seldom focus on interaction, interaction costs do impact usability. Indeed, in discussing their study results of testing five visualization tools for microarray analyses, Saraiya *et al.* commented that,

> The design of interaction mechanisms in visualization is critically important. Usability can outweigh the choice of visual representation. [*53, p.7]

In our review, we noticed a number of interaction costs that impact usability. In some cases, researchers attributed those interaction costs to inferior results obtained in their studies. We further discuss interaction design considerations based on our framework (Fig 1) as suggestions to narrow the gulfs of execution (Section 6.1) and evaluation (Section 6.2). We also raise the issue of whether visualizations should also address the gulf of goal formation (Section 6.3).

6.1 Narrowing the Gulf of Execution: Less is more

Similarly to visual cluttering, complex interaction is detrimental. Interfaces that provide too many choices, both in the form of input modes (Section 5.3) and as interface options (Section 5.1.2) may deter effective use by inducing unnecessary cognitive loads. In executing physical sequences, reports show users were annoyed by repetitive drag-n-drop movements in PDQ (Section 5.4.2) and in In-Spire (Section 5.4.3), were less effective in insight generation using GeneSpring (Section 5.4.3), made more errors with the Brushing histogram (Section 5.4.3), and produced worse time and accuracy performances with the Drill-Down method over the Distortion method in tree-node searches (Section 5.4.3).

It is unclear if training can mitigate these costs. Buering *et al.* believed so (Section 5.3.3):

> Seven subjects mentioned that they had problems with the sliding technique of the ZUI. We assume that this was mainly caused by the users' unfamiliarity with this kind of panning. [*6, p.834–835]

We also wondered if insufficient experience with novel and complex interfaces may have influenced interface choice where participants chose simpler but suboptimal interfaces (Section 5.1.2):

> Switching from a multiple-[level] mode to a single [level] can thus provide an easily perceived short-term benefit of lower cognitive load, despite potentially increasing the total time required to complete the task. Our study training for the users required them to demonstrate proficiency in the use of all four interfaces, as is usual in single-session laboratory settings. We conjecture that users trained to demonstrate proficiency in a multiple-[level] interface may still not have internalized confidence in its use: that is, may not have adequately understood the longer-term cost of these short-term choices. [*34, p.1285]

Nonetheless, we believe designs should aim for a small set of simple and predictable interaction, even at the cost of reduced user control. Simpler interaction can reduce errors, as seen in Li and North's study,

> simpler interactions of DQ [dynamic-query] sliders (only the slider thumbs were interactive) [as it] helped avoid mistakes. In contrast, all bars in the brushing histograms were interactive. [...] Thus it was easier for users to make incorrect or accidental selections [*35, p.152]

Simpler interaction can lead to better user performance, as reported by Shi *et al.* in their study with a tree-map like distortion technique for visual search tasks,

> The main reason that users perform better at the distortion techniques is due to the fact that in the drill-down approach the user has to drill-down and roll-up during several iteration[s] until the node is found. [*58, p.88]

Designers can replace sequences of actions by a single action to simplify interaction. The tradeoff is reduced interaction flexibility offered by intermediate steps. Our reviewed studies offer two examples: Hetzler *et al.* added a Keep Current capability in In-Spire to replace the more interactive approach to keep track of document themes over time (Section 5.4.3), and filtering with the DQ sliders instead of the more flexible brushing histograms. Designers can also use interaction techniques that are less physically demanding, such as using mouse clicks over drag-and-drops as suggested by participants in Kumar *et al.*'s PDQ Tree browser study (Section 5.4.2), or provide keyboard shortcuts such as Ctrl+S to avoid the more costly menu selections.

6.2 Narrowing the Gulf of Evaluation

To evaluate interaction outcomes, users need to first understand visual changes caused by interaction. While we believe interface intelligence can offer benefits especially in large-data exploration, we believe users should be kept in the control loop (Section 6.2.1). As for object associations, our reports suggest that commonly-used solutions are effective but may be inadequate (Section 6.2.2), and visualization should provide better support for reflection in analysis (Section 6.2.3).

6.2.1 Steerable and predicable augmentation

It is well known that interface intelligence can result in user confusion or even mistrust. Given the large quantity of data under visual analysis, users should be able to benefit from machine intelligence. Perhaps the solution is to emphasize augmentation rather than automation, allowing the human operator to understand and diagnose unexpected system behaviours, and to modify algorithm parameters [41].

6.2.2 Commonly-used solutions inadequate

Reports collected in this review indicate that while standard solutions for object associations such as animation, and linking and brushing are effective, users may need additional help to understand view changes. Lack of training may account for some of the troubles. For example, Saraiya *et al.* reasoned that the lack of multiple-view use was due to insufficient participant experience:

> Since the participants were novice users, they were also not experienced with performing data analysis on multiple views simultaneously.[...]
>
> Also, giving participants a longer training period on brushing and linking might have been helpful for them to better utilize the reverse brushing direction in which the parallel coordinate view is used to query the graph view. [*54, p.231]

Landmarking has been proposed as a visual aid (Sections 5.6.2–4). An open research question is to identify the number and type of visual or interaction aids required for object associations.

6.2.3 Support reflection in analysis

Our reports prompted us to consider the need for reflective cognition support in visualizations. While well-designed interaction supports experiential cognition as users effortlessly respond to incoming information without conscious reflection [40, p.22–31], reflective cognition is needed in analysis as users need to compare and contrast effects of different hypotheses [40, p.205]. Section 5.7 lists reports where users explored more with interfaces that supported refinding. Supporting refinding by allowing users to save visualization state is therefore not an implementation detail but an important design consideration.

6.3 Thinking about the Gulf of Goal Formation

Interface experimenters tend to provide study tasks and data for the participants. In real life, however, users have to come up with their own tasks and data focus with real goals. Even though the interface may be well designed at any point, users may not be able to use the visualization if they fail to form analysis questions. Large-data analysis is in itself a difficult task that required nontrivial amount of training and experience. Visualization perhaps should not be expected guide users who do not know what they are looking for, or are unable to articulate a question. Yi *et al.*, in evaluating the Dust-and-Magnet (DnM) interface, were ambivalent about adopting this role for visualization,

> One might question whether DnM can be effective for people who do not know how to articulate a question. This is a fair question because DnM does not supply any systematic approaches to producing this query initially. Therefore, the user should pose the question in advance in order to find answers using DnM effectively. Conversely, as shown in the user evaluation, the fact that DnM is easy and interesting to use might encourage users to explore data sets more vigorously, which smoothly lead them to pose proper questions. [*68, p.255]

On the other hand, users did get lost without guidance. In their evaluation of the Integrated Thesaurus-Results Browser, Sutcliffe *et al.* recommended system guidance:

> The visualization did appear to be comprehensible to users; however, it was hindered by lack of guidance on search strategies and possibly by the manipulations we provided for exploring the thesaurus and results browser visualizations. Basing visualization design on user tasks and data models has been advocated by others and demonstrated in successful products; however, in more complex tasks further research on visualization design methods that integrate active system guidance with visual browser and exploration tools is required. [*62, p.760]

Our framework therefore suggest a need to consider the gulf of goal formation, or if visualization should guide users in data explorations.

7 Limitations of Study

While we aimed to provide an objective and in-depth review to understand impacts of interaction costs on interface usability, our framework is limited by our venue scope and reported costs—we can only discuss techniques that had been studied and costs reported. Unlike study-statistics results, reporting on interface interactivity has not been standardized in publications. Our reports are therefore possibly subjective as they are observations and conjectures by paper authors.

Also, reasons behind the lack of reports are impossible to discern. Researcher may not report interaction costs for at least three reasons: first, study participants did not exhibit significant ill effects from interface interaction; second, interaction costs observed were considered to be due to study participants' inexperience with the interface and could be overcome with use; third, the study did not measure interaction costs, either with objective measurements or subjective observations.

Given these limitations, we took a qualitative route. A more quantitative approach, for example to assess the prevalence and severity of interactions costs, remains an open problem. Consequently, we can only conjecture rather than derive claims from our analysis, and our framework of seven costs is necessarily incomplete and perhaps non-representative. Also, our framework is only one interpretation of reports collected. With these limitations and the broad scope of our review, we can only provide broad general guidelines in Section 6.

8 Conclusion

We performed a systematic review on 484 InfoVis, IVS, IJHCS, and TOCHI publications and isolated 32 that reported on interaction uses. From these publications, we identified 61 interaction-related usability reports and grouped them into seven main costs inspired by Norman's Seven Stages of Action [39]. In addition to the Norman's gulfs of execution and evaluation, we proposed adding a gulf of goal formation to cover decision costs. Our framework is an initial step to study interaction costs.

Results of our review suggest a need to focus on interaction costs in visualization evaluations. We were surprised that even in journal publications where page constraints are less severe than in conference proceedings, less than 30% of user studies of interactive interface mentioned interaction. Even though cognitive costs seem to be more prevalent in modern visualization, we found that 30% of all reports are motion costs, perhaps since they are more observable. Our framework suggests a need to better identify cognitive costs in interaction. For example, our surveyed papers only identified costs in data and system option selections to form goals (Section 5.1), which are insufficient to capture the essence of sensemaking in visual analysis. Our framework therefore suggests a need to diversify from our traditional focus on visual encoding (cost 5) to cover the entire action cycle.

In terms of study methodologies, given the challenges in evaluating visualizations (e.g., [46, 21]), we believe recording observations in laboratory or field studies is a good starting point, as by collecting statements on interaction during interface use, we can begin to identify possible factors in interaction before we can develop more objective usability metrics, and perhaps new technologies, to capture and quantify these costs. Only then can we truly quantify the prevalence and severity of interaction costs on interface usability.

REFERENCES

[1] *J. Abello *et al.* ASK-Graph View: A Large Scale Graph Visualization System. *IEEE TVCG*, 12(5):669–676, 2006.

[2] C. Ahlberg and B. Shneiderman. Visual Information Seeking: Tight Coupling for Dynamic Query Filters with Starfield Displays. In *Proc ACM CHI*, pages 313–317, 1994.

[3] M. Beaudouin-Lafon. Designing Interaction, Not Interfaces. In *Proc ACM AVI*, pages 15–22, 2004.

[4] B. Bederson. Fisheye Menus. In *Proc ACM UIST'00*, pages 217–226, 2000.

[5] *B. Bederson *et al.* DateLens: A Fisheye Calendar Interface for PDAs. *ACM TOCHI*, 11(1):90–119, March 2004.

[6] *T. Buring *et al.* User Interaction With Scatterplots on Small Screens—A Comparative Evaluation of Geometric-Semantic Zoom and Fisheye Distortion. *IEEE TVCG*, 12(5):829–836, 2006.

[7] S.K. Card *et al.* A Morphological Analysis of the Design Space of Input Devices. *ACM TOIS*, 9(2):99–122, 1991.

[8] S.K. Card *et al. Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, 1999.

[9] M.S.T. Carpendale *et al.* Three-dimensional pliable surfaces: for the Effective Presentation of Visual Information. In *ACM UIST)*, pages 217–226, 1995.

[10] C. Chen and Y. Yu. Empirical Studies of Information Visualization: A Meta-Analysis. *Intl. J. of Human-Computer Studies (IJHCS)*, 53:851–866, 2000.

[11] E.H. Chi and J.T. Riedl. An Operator Interaction Framework for Visualization Systems. In *Proc IEEE InfoVis*, pages 63–70, 1998.

[12] E.H. Chi *et al.* Using Information Scent to Model User Information Needs and Actions and the Web. In *Proc ACM CHI*, pages 490–497, 2001.

[13] *W. Chung. Studying Information Seeking on the Non-English Web: An Experiment on a Spanish Business Web Portal. *IJHCS*, 64(9):811–829, 2006.

[14] M. Collomb *et al.* Improving Drag-and-Drop on Wall-Size Display. In *Proc GI*, pages 25–32, 2005.

[15] *M. Crampes *et al.* Concept Maps for Designing Adaptive Knowledge Maps. *Information Visualization*, 5(3):192–224, 2006.

[16] P. Fitts. The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. *J. Expt. Psych.*, 47:381–391, 1954.

[17] *M. Granitzer *et al.* Evaluating a System for Interactive Exploration of Large, Hierarchically Structured Document Repositories. In *Proc IEEE InfoVis*, pages 127–134, 2004.

[18] C. Gutwin. Improving Focus Targeting in Interactive Fisheye Views. In *Proc ACM CHI*, pages 267–274, 2002.

[19] J. Heer and D. Boyd. Vizster: Visualizing Online Social Networks. In *Proc IEEE InfoVis*, pages 33–40, 2005.

[20] *E.G. Hetzler *et al.* Turning the Bucket of Text into a Pipe. In *Proc IEEE InfoVis*, pages 89–94, 2005.

[21] K. Hornbæk. Current Practice in Measuring Usability: Challenges to Usability Studies and Research. *IJHCS*, 64:79–102, 2006.

[22] *K. Hornbæk and M. Hertzum. Untangling the Usability of Fisheye Menus. *ACM TOCHI*, 14(2), 2007.

[23] *K. Hornbæk *et al.* Navigation Patterns and Usability of Zoomable User Interfaces with and without an Overview. *ACM TOCHI*, 9(4):362–389, 2002.

[24] *K. Hornbæk *et al.* Reading Patterns and Usability in Visualization of Electronic Documents. *ACM TOCHI*, 10(2):119–149, 2003.

[25] J.P. Hourcade *et al.* Differences in Pointing Task Performance between Preschool Children and Adults using Mice. *ACM TOCHI*, 11(4):357–386, 2004.

[26] G.S. Hubona *et al.* The Effects of Motion and Stereopsis on Three-dimensional Visualization. *IJHCS*, 47(5):609–627, 1997.

[27] R.J.K. Jacob *et al.* Integrality and Separability of Input Devices. *ACM TOCHI*, 1(1):3–26, 1994.

[28] T.J. Jankun-Kelly *et al.* A Model and Framework for Visualization Exploration. *IEEE TVCG*, 13(2):357–369, 2007.

[29] S. Jul and G.W. Furnas. Critical Zones in Desert Fog: Aids to Multiscale Navigation. In *Proc ACM UIST*, pages 97–106, 1998.

[30] *H. Kang *et al.* NetLens: Iterative Exploration of Content-Action Network Data. *Information Visualization*, 6(1):18–31, 2007.

[31] C. Klein and B.B. Bederson. Benefits of Animated Scrolling. In *Extended Abstracts of CHI*, pages 1965–1968, 2005.

[32] *A. Kobsa. An Empirical Comparison of Three Commerical Information Visualization Systems. In *Proc IEEE InfoVis*, pages 123–130, 2001.

[33] *H.P. Kumar *et al.* Browsing Hierarchical Data with Multi-level Dynamic Queries and Pruning. *IJHCS*, 46(1):103–124, 1997.

[34] *H. Lam *et al.* Overview Use in Multiple Visual Information Resolution Interfaces. *IEEE TVCG*, 13(6):1278–1283, 2007.

[35] *Q. Li and C. North. Empirical Comparison of Dynamic Query Sliders And Brushing Histograms. In *Proc IEEE InfoVis*, pages 147–153, 2003.

[36] *M.J. McGuffin and R. Balakrishnan. Fitts' Law and Expanding Targets: Experimental Studies and Designs for User Interfaces. *ACM TOCHI*, 12(4):388–422, 2005.

[37] *M.J. McGuffin *et al.* Expand-ahead: A Space-Filling Strategy for Browsing Trees. In *Proc IEEE InfoVis*, pages 119–126, 2004.

[38] K. Misue *et al.* Layout Adjustment and the Mental Map. *J. of Visual Languages and Computing*, 6(2):183–210, 1995.

[39] D.A. Norman. *The Design of Everyday Things*. Doubleday/Currency, 1990.

[40] D.A. Norman. *Things that Make Us Smart*. Basic Books, 1993.

[41] D.A. Norman. *The Design of Future Things*. Basic Books, 2007.

[42] *C. North and B. Shneiderman. Snap-Together Visualization: Can Users Construct and Operate Coordinated Visualizations. *IJHCS*, 53(5):715–739, 2000.

[43] C. Olson and E.H. Chi. ScentTrails: Integrating Browsing and Search on the Web. *ACM TOCHI*, 10(3):177–197, 2003.

[44] P. Pirolli and S.K. Card. Information Foraging. *Psych. Review*, 106:643–675, 1999.

[45] P. Pirolli and S.K. Card. Sensemaking Processes of Intelligence Analysis and Possible Leverage Points as Identified Through Cognitive Task Analysis. In *Proc Intl. Conf. on Intelligence Analysis*, 2005.

[46] C. Plaisant. The Challenge of Information Visualization Evaluation. In *Proc ACM AVI*, pages 109–116, 2004.

[47] *M. Plumlee and C. Ware. Zooming versus Multiple Window Interfaces: Cognitive Costs of Visual Comparisons. *TOCHI*, 13(2):179–209, 2006.

[48] R. Rao and S.K. Card. Table Lens: Merging Graphical and Symbolic Representations in An Interactive Focus Plus Context Visualization for Tabular Information. In *Proc ACM CHI*, pages 318–322, 1994.

[49] X. Ren and S. Moriya. Improving Selection Performance on Pen-Based Systems: A Study of Pen-Based Interaction for Selection Tasks. *ACM TOCHI*, 7(3):384–416, 2000.

[50] J.C. Roberts. Multiple-View and Multiform Visualization. In *Visual Data Exploration and Analysis VII. Proc SPIE 3960*, pages 176–185, 2000.

[51] G.G. Robertson *et al.* The Cognitive Coprocessor Architecture for Interactive User Interfaces. In *Proc ACM UIST*, pages 10–18, 1989.

[52] D.M. Russell *et al.* The Cost Structure of Sensemaking. In *Proc ACM CHI*, pages 269–276, 1993.

[53] *P. Saraiya *et al.* An Evaluation of Microarray Visualization Tools for Biological Insight. In *Proc IEEE InfoVis*, pages 1–8, 2004.

[54] *P. Saraiya *et al.* Visualization of Graphs with Associated Timeseries Data. In *Proc IEEE InfoVis*, pages 225–232, 2005.

[55] B. Schwartz. *The Paradox of Choice: Why More Is Less*. HarperCollins, 2004.

[56] *A. Sears and B. Shneiderman. Split Menus: Effectively Using Selection Frequency to Organize Menus. *ACM TOCHI*, 1(1):27–51, 1994.

[57] K. Sedig *et al.* Application of Information Visualization Techniques to the Design of a Mathematical Mindtool: A Usability Study. *Information Visualization*, 2(3):142–159, 2003.

[58] *K. Shi *et al.* An Evaluation of Content Browsing Techniques for Hierarchical Space-Filling Visualizations. In *Proc IEEE InfoVis*, pages 81–88, 2005.

[59] *H. Siirtola and E. Mäkinen. Constructing and Reconstructing the Reorderable Matrix. *Information Visualization*, 4(1):32–48, 2005.

[60] M.W. Smith *et al.* Aging, Motor Control, and the Performance of Computer Mouse Tasks. *Human Factors*, 41(3):589–596, 1999.

[61] R. Spence. A framework for Navigation. *IJHCS*, 51(5):919–945, 1999.

[62] *A.G. Sutcliffe *et al.* Evaluating the Effectiveness of Visual User Interfaces for Information Retrieval. *IJHCS*, 53(5):741–763, 2000.

[63] J. Teevan. The Re:Search Engine: Simultaneous Support for Finding and Re-Finding. In *Proc ACM UIST*, pages 23–32, 2007.

[64] M.O. Ward. XmdvTool: Integrating Multiple Methods for Visualizing Multivariate Data. In *Proc IEEE Vis*, pages 326–333, 1994.

[65] *S.J. Westerman *et al.* Browsing a Document Collection Represented in Two- and Three-Dimensional Virtual Information Space. *IJHCS*, 62(6):713–736, 2005.

[66] W. Willett *et al.* Scented Widgets: Improving Navigation Cues with Embedded Visualizations. *IEEE TVCG*, 13(6):1129–1136, 2007.

[67] T. Winograd and C.F. Flores. *Understanding Computers and Cognition*. Ablex Publishing Corp., 1986.

[68] *J.S. Yi *et al.* Dust & Magnet: Multivariate Information Visualization Using a Magnet Metaphor. *Information Visualization*, 4(4):239–256, 2005.

[69] J.S. Yi *et al.* Toward a Deeper Understanding of the Role of Interaction in Information Visualization. In *Proc IEEE InfoVis*, pages 1224–1231, 2007.

[70] A. Zanella *et al.* On the Effects of Viewing Cues in Comprehending Distortions. In *ACM CHI*, pages 119–128, 2002.

* denotes publications with listed reports