

# Information Visualization

<http://ivi.sagepub.com/>

---

## **Fluid interaction for information visualization**

Niklas Elmqvist, Andrew Vande Moere, Hans-Christian Jetter, Daniel Cernea, Harald Reiterer and TJ Jankun-Kelly  
*Information Visualization* 2011 10: 327 originally published online 22 August 2011  
DOI: 10.1177/1473871611413180

The online version of this article can be found at:  
<http://ivi.sagepub.com/content/10/4/327>

---

Published by:



<http://www.sagepublications.com>

**Additional services and information for *Information Visualization* can be found at:**

**Email Alerts:** <http://ivi.sagepub.com/cgi/alerts>

**Subscriptions:** <http://ivi.sagepub.com/subscriptions>

**Reprints:** <http://www.sagepub.com/journalsReprints.nav>

**Permissions:** <http://www.sagepub.com/journalsPermissions.nav>

**Citations:** <http://ivi.sagepub.com/content/10/4/327.refs.html>

>> [Version of Record](#) - Oct 14, 2011

[OnlineFirst Version of Record](#) - Aug 22, 2011

[What is This?](#)

# Fluid interaction for information visualization

Niklas Elmqvist<sup>1</sup>, Andrew Vande Moere<sup>2</sup>, Hans-Christian Jetter<sup>3</sup>, Daniel Cernea<sup>4</sup>, Harald Reiterer<sup>3</sup> and TJ Jankun-Kelly<sup>5</sup>

Information Visualization  
10(4) 327–340  
© The Author(s) 2011  
Reprints and permissions:  
sagepub.co.uk/journalsPermissions.nav  
DOI: 10.1177/1473871611413180  
ivi.sagepub.com



## Abstract

Despite typically receiving little emphasis in visualization research, interaction in visualization is the catalyst for the user's dialogue with the data, and, ultimately, the user's actual understanding and insight into these data. There are many possible reasons for this skewed balance between the visual and interactive aspects of a visualization. One reason is that interaction is an intangible concept that is difficult to design, quantify, and evaluate. Unlike for visual design, there are few examples that show visualization practitioners and researchers how to design the interaction for a new visualization in the best manner. In this article, we attempt to address this issue by collecting examples of visualizations with 'best-in-class' interaction and using them to extract practical design guidelines for future designers and researchers. We call this concept *fluid interaction*, and we propose an operational definition in terms of the direct manipulation and embodied interaction paradigms, the psychological concept of 'flow', and Norman's gulfs of execution and evaluation.

## Keywords

fluidity, flow, embodiment, design, information visualization, human–computer interaction

## Introduction

Not all information visualization (InfoVis) tools are created equal. As all users of such tools know, while we can expect any competent tool to be able to represent abstract data in graphical form, there is a certain class of InfoVis tools that take this a step further through engaging, compelling, and even absorbing user experiences that turn the analytical sensemaking<sup>1</sup> process into a pleasurable task. However, the academic portion of the InfoVis field, in general, puts little emphasis on design, aesthetic, and user experience aspects of information visualization tools, and so far there has been virtually no effort towards characterizing this class of InfoVis tool in the research community. In contrast, much of information visualization in the real world is directly concerned with creating compelling—even playful—and beautiful tools that are capable of capturing the attention of general users on the Internet and in public spaces such as museums, exhibition halls, and corporate lobbies.

In this article, we try to remedy this state of affairs by proposing a unifying concept for both researchers and practitioners that captures this class of *InfoVis exemplars* in a single definition: *fluid interaction*. Fluidity in information visualization is an elusive and intangible

concept characterized by smooth, seamless, and powerful interaction; responsive, interactive and rapidly updated graphics; and careful, conscientious, and comprehensive user experiences. Our hypothesis is that an InfoVis tool that exhibits this fluidity in all aspects will transform the sensemaking process into an efficient, illuminating, and even enjoyable experience because it helps the user stay in the flow<sup>2</sup> of the work process. However, creating a fluid design is far from trivial, mainly because of this intangibility and elusiveness.

To better illustrate our definition of fluid interaction for information visualization, we collect and describe a subset of these exemplar InfoVis tools from the scientific community, including BabyNameVoyager,<sup>3</sup> FacetStreams,<sup>4</sup> and Scatter/GraphDice.<sup>5,6</sup> Unlike most

<sup>1</sup>Purdue University, USA.

<sup>2</sup>K. U. Leuven University, Belgium.

<sup>3</sup>University of Konstanz, Germany.

<sup>4</sup>University of Kaiserslautern, Germany.

<sup>5</sup>Mississippi State University, USA.

## Corresponding author:

Niklas Elmqvist, School of Electrical and Computer Engineering, Purdue University, 465 Northwestern Ave., Electrical Engineering Building, West Lafayette, IN 47907-2035, USA  
Email: elm@purdue.edu

existing research articles in information visualization, we also collect exemplars from outside the scientific community, such as from the design, aesthetics, and infographics communities. These ‘real-world’ exemplars include *mæve*,<sup>7–9</sup> *We Feel Fine*,<sup>10</sup> and the interactive holographics from the film *Iron Man 2*.<sup>11</sup>

Using these exemplars as a starting point, we derive tips, guidelines, and principles for how to achieve fluidity in InfoVis in terms of both interaction and visual representations. We hope that these snippets of practical experience, elevated almost to the level of being design idioms and patterns,<sup>12</sup> will help developers of InfoVis tools—academics and practitioners alike—to build better, more rewarding, more captivating, and ultimately more efficient information visualization tools that will propel our field to make an even bigger impact in both our own scientific community as well as the real world.

In the rest of this article, we first survey the state of the art in fluid interaction, design, and user experience for information visualization. We then attempt to define the concept of fluid interaction through a detailed list of properties that we think characterize fluid InfoVis tools. These properties are all exhibited in the context of our InfoVis exemplars, that we describe next. Starting from these exemplars, we derive design guidelines for fluid information visualization, and conclude the paper with our vision for future research directions in design aspects of information visualization.

## Related work

Unlike its sister field of human–computer interaction, where design and user experience are major components, InfoVis research articles generally do not place much emphasis on interaction design aspects of information visualization. The few papers that discuss these topics tend to do so from a purely scientific, engineering, or implementation viewpoint. For example, although Amar et al.<sup>13</sup> present a meta-level review of typical analysis tasks and Yi et al.<sup>14</sup> collect seven archetypes of interaction, these are all still descriptive classifications, which makes them difficult to use in a generative and design purpose. They also do not describe ‘softer’ and less goal-driven aspects of interaction design such as aesthetics, user experience, and rewarding interaction.

Below we discuss both the visual and the interactive aspects of visualization design.

## Visual aspects

Textbooks on graphic design are an excellent source of information on maximizing clarity and expressiveness

in static (often printed) visual representations. Bertin<sup>15</sup> discuss retinal variables for visual marks in data visualization (later improved and extended by Mackinlay<sup>16</sup> and again by Card et al.<sup>17</sup>). Furthermore, Edward Tufte’s books<sup>18–21</sup> are key resources for visual design of data displays, and have influenced many InfoVis systems and papers through the years. Following this tradition, Stephen Few’s books give clear and concrete design guidelines for how to design graphs and tables<sup>22</sup> as well as information dashboards.<sup>23</sup>

Visualization textbooks are typically even more relevant because they focus on interactive visualization applications. Colin Ware’s books<sup>24,25</sup> on perception and cognitive aspects of visualization provide valuable background on the psychology of visual thinking for interactive visualizations. Robert Spence takes a design and interaction-oriented approach to information visualization in his textbook.<sup>26</sup> A recent book, *Beautiful Data*,<sup>27</sup> gives a hands-on approach to visualization design through a set of case studies involving real datasets collected from the web. Most recently, Ward et al.<sup>28</sup> devote a full chapter to step-by-step guidelines on how to design effective visualizations.

## Interactive aspects

While visual aspects are important for fluid interaction, it is clearly the interactive aspects that are central for effective visualization interaction design. Unfortunately, interaction is not discussed at all in graphic design, and even visualization textbooks tend to downplay this angle.

One exception is Robert Spence’s book,<sup>26</sup> which takes interaction design as its starting point for the study of information visualization. Many of the visualization techniques presented in the book include a discussion on key interaction design aspects associated with the visual representation. Another exception is Stephen Few’s newest book,<sup>29</sup> which incorporates a chapter on analytical interaction where Few gives recommendations for how visualization software should best support the analytical discourse. Few even uses the word ‘fluid’ to describe a desirable feature of the interactive exploration process.<sup>29</sup> Our work in this paper builds on these existing efforts, but we formalize the concept of flow and fluidity further.

The recent paper by Pike et al.<sup>30</sup> on the ‘science of interaction’ is highly relevant to our work. Pike makes a case for the role of interaction in visualization and visual analytics, and emphasizes themes such as interaction design, user experiences, and best practices for interactive tools—the very same themes we promote here. This article ends with seven broad areas for future research, but does not take the practical design approach that our paper does.

Given the dearth of interaction design in InfoVis research, it is fortunate that there exists much HCI research that is highly relevant to the concepts of fluidity and flow in visualization. We draw on a wide variety of disciplines—HCI and interaction design included—when we define fluid interaction in the next section and we will thus cite and discuss these sources below. However, the most influential work includes the direct manipulation paradigm,<sup>31</sup> Norman's gulfs of execution and evaluation,<sup>32</sup> the instrumental interaction model,<sup>33,34</sup> tangible<sup>35</sup> and embodied interaction,<sup>36</sup> and, most recently, the reality-based interaction framework.<sup>37</sup>

In the next section, we show how this rich collection of related work, theories, and frameworks can be tied together into the concept of fluid interaction and how it can be applied to information visualization.

## Fluid interaction for information visualization

As can be seen from the above literature survey, there exists very little work on the topic of *fluid interaction* for information visualization. Part of the reason for this is that notions of flow<sup>2</sup> and fluidity are very elusive and difficult to pin down.<sup>38</sup> In this section, we will make inroads towards an operational definition of this concept. In the next section, we will describe a set of existing applications that fulfil this definition: we call them InfoVis exemplars. Finally, in the section following that, we synthesize the characteristic properties collected from these exemplars into general guidelines for how to design fluid InfoVis applications.

### Operational definition

As mentioned above, fluid interaction is an elusive concept that is not easily amenable to a theoretical definition. Instead, we present an *operational definition* of the properties of fluidity that draws from a plethora of sources, including HCI design, the concept of flow,<sup>2</sup> embodiment,<sup>36</sup> immersion, and natural interaction. As a starting point, Merriam-Webster defines fluidity as follows:

**fluidity, n.:** 1. the quality or state of being fluid; 2. the physical property of a substance that enables it to flow.

Following this general definition, a *fluid interface* for information visualization is characterized by one or several of the following properties:

- **Promotes flow:** The interaction should be designed to promote staying in the flow. 'Flow'<sup>2</sup> is defined as a mental state of total immersion in an activity where

the challenges of the activity and the skills of the participant are perfectly balanced, leading to high focus, involvement, and rewarding outcomes. Bederson<sup>38</sup> previously proposed an interaction design philosophy based on helping the user stay in the flow, emphasizing five characteristics for user interface design. In the below list, we take a somewhat broader view of Csikszentmihalyi's factors for flow that may influence the concept of fluid interaction:

- *Balanced challenge:* the skill required by the activity and the user's skill level should be matched;
- *Concentration:* the activity should allow for a high degree of focus on a limited field of attention;
- *Loss of self-consciousness:* the user should be able to merge action and awareness;
- *Transformation of time:* enable users to 'lose themselves' in the activity, essentially losing track of time;
- *Prompt feedback:* users should be immediately informed of progress towards their goals;
- *Sense of control:* ensure users feel in control over the activity so that they can truly affect the outcome; and
- *Intrinsically rewarding:* the activity should have a tangible reward in and of itself.
- **Supports direct manipulation:** The direct manipulation paradigm<sup>31</sup> (further extended by the instrumental interaction model<sup>33,34</sup>) promotes an explicit method of interacting with computers by directly interacting with the domain objects themselves, thereby minimizing the indirection in the interface. The paradigm is based on four main principles:
  - continuous representation of the object of interest;
  - physical actions instead of complex syntax;
  - rapid, incremental, and reversible operations whose impact on the object of interest is immediately visible; and
  - layered or spiral approach to learning that permits usage with minimal knowledge.
- **Minimizes the gulfs of action:** According to usability expert Donald Norman, the challenge of interacting with any system, physical or virtual, can be described in terms of two *gulfs*<sup>32</sup>:
  - *Gulf of Evaluation:* The difference between the system's state and the user's perception of that state.
  - *Gulf of Execution:* The difference between the allowable actions of a system and the user's intentions for using the system.

### Towards a cognitive account of fluid interaction

From a cognitive perspective, the aforementioned desired properties of a *fluid interface* for InfoVis share

some important commonalities. Understanding these commonalities and identifying an underlying common concept could help to better understand what creates and what hampers fluidity in InfoVis. In the following, we suggest such a common concept which is deduced from existing cognitive models. This first step towards a cognitive account of fluidity is still far from being a comprehensive theoretical model. However, viewing visualization design through this new lens could help to analyse existing information visualizations and to inform and refine future models.

In our view, a basic requirement for fluidity and the underlying concept behind it is the users' feeling of direct participation and embodiment in the interface. Fluid interfaces for InfoVis must make the users feel that are able to directly 'touch' and manipulate the visualization instead of indirectly conversing with a user interface. Users should get a feeling of immersion, first-personness and direct engagement with the objects and the visualizations that concern them.

This phenomenon was first described by Hutchins et al.<sup>39</sup> in their cognitive account of direct manipulation. They differentiated between two major metaphors for the nature of human–computer interaction, a *conversation metaphor* and a *model-world metaphor*: in the former, the interface serves as a language for interacting with the world, whereas in the latter, the interface itself is the world which the user can manipulate. For Hutchins et al., model–world interfaces create a feeling of directness and direct manipulation by minimizing the *gulfs of evaluation* and *execution*<sup>32</sup> and thereby using much less of the users' cognitive resources.

We believe, that InfoVis—and in particular *fluid* InfoVis—should follow this model–world metaphor. This is also in line with the growing importance of theories of *embodied cognition* in cognitive science and *embodied interaction* in human–computer interaction.<sup>36</sup> These embodied views emphasize that our cognitive abilities are specifically designed to reason, act, and move in our natural physical and social world. In other words, our entire way of thinking and our perception are optimized for these real-world tasks and cannot be separated from our physical and social existence. These physical and social skills are far more constituent of what we consider as human cognition than the disembodied and formal processing of symbols when conversing with a computer by clicking through labelled buttons, menus, hyperlinks, or forms. Thus, we are wasting a great deal of our true skills when using computing technology without directly acting in metaphorical (model-)worlds.

This insight has led the field of HCI to a new generation of user interfaces based on *reality-based interaction*.<sup>37</sup> These interfaces use modalities such as body

tracking, bi-manual multi-touch interaction, or tangible objects to further reduce the users' *gulf of execution* and to draw strength by employing themes of reality such as *body*, *environment* and *social skills & awareness*.<sup>37</sup> Thus, they transfer an even stronger form of direct manipulation from desktop to post-desktop computing.

In conclusion, given our innate cognitive abilities, creating model–world interfaces into which users can immerse themselves and in which objects of concern become virtually or even physically tangible can help to achieve Csikszentmihalyi's factors of flow. Directly manipulating the objects in the model–world with a greater set of motor skills (e.g. directly dragging an object with multi-touch or mouse instead of pushing keys or buttons to that effect) mediates a *sense of control* and the model–worlds provide the desired *prompt feedback*. They use less cognitive resources and thus enable *concentration* on the task instead of concentration on handling the user interface. Less usage of cognitive resources can also help to achieve a greater design space for a more *balanced challenge*. Furthermore, good direct manipulation interfaces are in many respects similar to computer games,<sup>31</sup> and thus could lead to the *loss of self-consciousness* and *transformation of time*.

### Utility of fluid interaction

Our argument so far has assumed that fluid interaction is a desirable attribute in visualization design. However, it is true that fluidity is not a necessary condition for any given visualization application, and we can even go so far as to say that there probably exist numerous very successful visualization and analysis applications that score low on what we would call 'core' fluidity properties. For example, the statistical package R (<http://www.r-project.org/>) has a command-based user interface with few visual components (i.e. a conversation interface), but is widespread and highly successful in many communities (we should note that R certainly does promote flow for expert users, while other fluidity properties are somewhat neglected).

On the other hand, interaction is the catalyst for the interplay between the data and the user, and is an essential part of visual exploration.<sup>14,40,41</sup> Thus, it follows that achieving the 'optimal experience' (characterized as flow<sup>2</sup>) while interacting with a visualization application will cause the user to perform better.<sup>38</sup> In other words, if we improve the fluidity of a visualization application, user performance would directly benefit. This is also the message of this article: showing how interaction design can be used to make visualization applications, existing and novel ones alike, more effective. For example, RStudio (<http://www.rstudio.org/>) is a new



graphical IDE for R that likely will help novice users overcome the steep learning curve of the R system

Many authors argue in favour of streamlining the interactive process in this way. Pike et al.<sup>30</sup> emphasize the need for natural and seamless interaction methods in support of discovery. Although they note that disruptions in the analytical discourse are inevitable and, in fact, often useful, their use of the word ‘fluid’ is different than ours, and their argument still seems to advocate the user staying in the flow of a visualization tool during analysis. Bederson<sup>38</sup> use interruptions as a reasoning tool when arguing for maintaining flow in interactive applications, noting that literal or conceptual interruptions can have large impact on user productivity. While sensemaking literature<sup>1</sup> typically discusses interaction with information at a conceptual level, many of Pirolli and Card’s leverage points<sup>42</sup> directly or indirectly involve improving the cost structure of analysis through tool innovation.

Given the above operational definition, cognitive account, and utility for fluidity, we are now ready to study concrete examples that embody these concepts.

## InfoVis exemplars

In this section, we review six InfoVis exemplars that we feel exhibit the fluid interaction properties discussed earlier in this article. Table 1 gives a summary of the exemplars based on domain (the primary community the system targets), audience (intended users), task (the primary task and intention with the system), and properties (the primary properties of fluidity that the system exhibits). We describe each of these systems in detail below.

Of course, any choice of a mere six examples from the rich array of excellent visualization systems is highly subjective and could even be seen as somewhat arbitrary. We used the following criteria when identifying the exemplars:

- **Diversity:** We wanted our selection to reflect the practice of information visualization in both

academic and design communities, as well as on the web, in movies, and in physical locations such as museums.

- **Illustrative:** Our choice was influenced by a desire to illustrate *different*—seemingly disparate—aspects of each system that the concept of fluidity could help explain.
- **Availability:** We chose only work that that was directly available (through demonstrations or videos) and amply described in the literature (for academic work) or in the blogosphere or popular press (for non-academic work). In other words, we relied on peer-review (academic work) or public acceptance (non-academic work).
- **Depth instead of breadth:** This paper is not a comprehensive survey of best-in-class visualizations, and thus our focus was on deep analysis of each exemplar instead of a large and broad enumeration of **all** InfoVis systems that incorporate fluid interaction.

Other factors beyond our control naturally also played a role in this selection process, including the authors’ own biases, personal tastes, and familiarity with the literature. However, regardless of the sparse sampling of the design space, we are convinced that the characteristic properties that we extract from each exemplar together are broadly applicable for fluid interaction in general.

### Facet-streams

Facet-Streams<sup>4</sup> is a system for co-located collaborative product search by multiple users around a tabletop (Figure 1). It supports small groups during decision-making and negotiation by enabling a faceted exploration of a product catalogue, e.g. a catalogue of hotels for a family’s vacation. Queries can be formulated by each participant by putting small glass discs as query tokens on the tabletop and assigning the desired data field and value ranges to them (Figure 2). These tokens can then be visually linked to form a directed

**Table 1.** Overview of the six InfoVis exemplars reviewed in this article with informal classifications on the domain, intended audience, main task, and primary fluid properties of each system

Exemplar	Domain	Audience	Task	Properties
Facet-Streams <sup>4</sup>	Academic	Groups	Collaborative search	Direct, embodied, minimizes gulfs
BabyNameVoyager <sup>3</sup>	Academic/design	Web	Exploration	Direct, prompt feedback, rewarding, control
Scatter/GraphDice <sup>5,6</sup>	Academic	Analyst	Exploration	Direct, prompt feedback, minimizes gulfs
Mæve <sup>8,9</sup>	Museum	Casual	Browsing	Embodied, rewarding, aesthetic
We Feel Fine <sup>49</sup>	Design	Web	Social navigation	Rewarding, aesthetic, minimal knowledge
Iron Man 2 <sup>11</sup>	Movie	Casual	Immersion	Embodied, feedback, aesthetic



**Figure 1.** The Facet-Streams system for co-located collaborative search on tabletops (left). Glass tokens form a tangible filter/flow representation of faceted Boolean search (right).



**Figure 2.** Assigning the desired data field (left) and value range (right) to a query token using touch input.

graph that serves as a visual filter/flow representation<sup>46</sup> of faceted Boolean search. All products from the catalogue flow along the edges of the network and are filtered by the nodes they pass. Logical AND and OR can be expressed without symbolic notations or query languages, simply by connecting nodes or letting edges flow together. Query results can be inspected by the users by touching an edge to reveal the products that flow therein.

Through tangible and touch input, Facet-Streams exploits a greater range of the users' real-world motor skills than normal desktop interfaces. The number and spatial layout of nodes can be altered by familiar physical manipulations, similar to placing, lifting or moving the pieces of a board game like checkers. The topology of the network can be changed by touch interaction, e.g. by dragging new connections between nodes with the fingers or by cutting them with a crossing out gesture. The selected data fields and value ranges of a node can be changed by touching and sliding the finger over the token's field and value dials (Figure 2). This also enables users to develop more advanced techniques, e.g. bi-manual selection of value ranges

during which one hand rotates the glass token and its attached dial while a finger of the other hand selects the segments of the dial that is rotating below. Thereby, every kind of tangible and touch input into the system leads to immediate visual feedback. For the users, this creates the illusion of direct physical interaction with the visual representation.

The benefit of this design is a low viscosity of the query's visual representation, i.e. a 'low resistance to change' in the interface.<sup>4</sup> It enables users to rapidly modify it according to their individual or shared goals. During initial search phases, each group member can formulate and explore their own criteria individually during phases of loosely coupled parallel work. These personal query networks can then be effortlessly combined into a larger group network for collective reviewing during phases of tightly coupled collaboration. However, the query's network can easily be dissolved into smaller parts again, e.g. for returning to parallel work or to separate the satisfactory parts from those that need further refinement. Therefore, low viscosity also gives the necessary flexibility to support different working styles or different collaborative phases.

In summary, we can extract the following characteristic properties of fluidity from Facet-Streams:

- *Tangible and touch interaction* with immediate visual feedback creates the illusion of physically interacting with the visual representation;
- This results in a *low viscosity* of the visual representation—i.e. a low resistance to change—that enables users to rapidly modify it according to their individual or shared goals; and
- *Flexibility in working styles* for the group of users by enabling smooth changes between loosely coupled parallel work and tightly-coupled collaboration.

### BabyNameVoyager

The BabyNameVoyager<sup>3</sup> is a web-based visualization tool (accessible from <http://www.babynamewizard.com/name-voyager>) for interactive representation and analysis of historical trends in baby naming. Figure 3 shows a screenshot of the application where a subset of names have been filtered out. The data are represented as a set of stacked graphs,<sup>47</sup> where the horizontal axis represents the time, and the vertical axis the amount of children that were given a certain name in the corresponding period. Initially, all names in the dataset are displayed as narrow coloured overlaying threads that are sorted alphabetically. Filtering—the most important interaction in this tool—can be achieved in two different ways that both exhibit a high degree of

fluidity: by interactively browsing the data, or by issuing a partial textual query.

The user can browse the data by moving the mouse over the name segments. Clicking a segment will select the corresponding name and expand it, while filtering out the other stacks in a gentle animation. The smooth transitions between states provide continuity between the information presented in different states. Moreover, the presence of a direct interface with the visualization suggests similarities to other real-world exploration tasks (e.g. browsing a drawer of ordered folders; clicking can be viewed as the selection and smooth opening of a folder to reach the actual information).

Textual queries are issued by typing a name, or parts of a name, into a text box. The filtering is immediate, as each key stroke generates a new intermediary name and a stricter filtering rule, causing the visual display to update as the user is typing. Again, the interaction mimics widely known and used actions from our everyday life, which suggests a coupling between fluidity and real-world metaphors.

In summary, we can extract the following characteristic properties of fluid interaction from the BabyNameVoyager:

- *Smooth animated transitions* between visualization states;
- *Minimalistic interface* using direct manipulation<sup>31</sup> (click-to-query) or an integrated query box (textual queries);

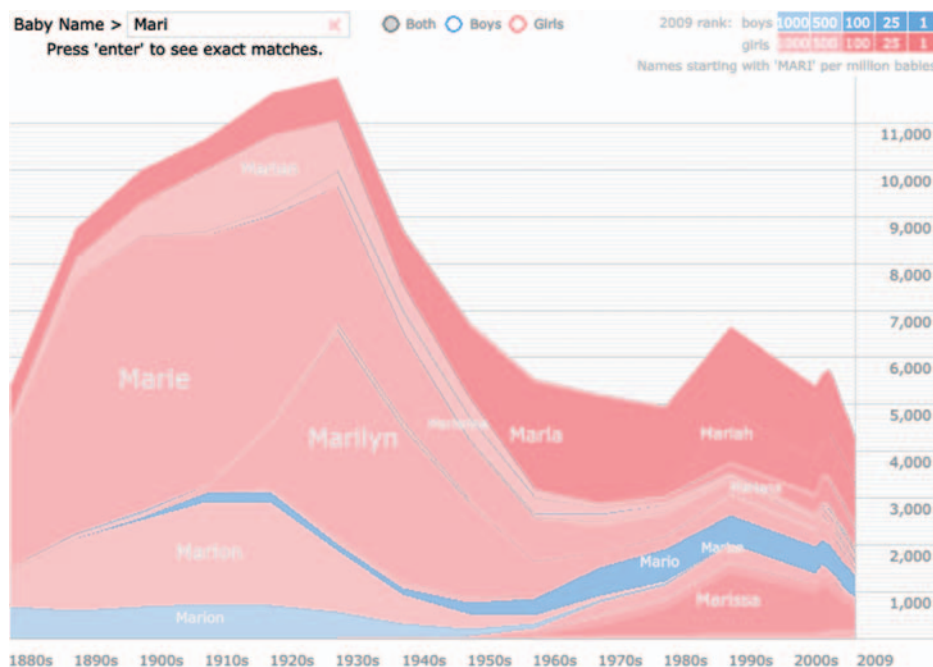


Figure 3. A screenshot of Baby Name Voyager showing names starting with 'Mari'.



- *Immediate visual feedback* for not only final queries, but also intermediate ones; and
- *Aesthetic visual design* that does not sacrifice correctness.

*ScatterDice and GraphDice*

ScatterDice<sup>6</sup> and GraphDice<sup>5</sup> are visualization tools for interactive visual exploration of multidimensional tables and multivariate graphs, respectively. Figure 4 shows a screenshot of the GraphDice interface (the ScatterDice interface is similar, but uses scatterplots instead of node-link diagrams). As can be seen from the image, the interface is dominated by the current plot (a 2D scatterplot for ScatterDice and a node-link graph with attribute-based layout for GraphDice), as well as a plot matrix that shows all the dimensions in the dataset. Users can change the current plot by navigating in the plot matrix; the transition from one position to another is communicated using an animated 3D rotation.

This navigation method—called *scatterplot matrix navigation*—is an example of fluid interaction because (1) the transition is smoothly animated with no discontinuities, and (2) it builds on a cohesive and consistent model for the visual representation that is closely integrated with an interaction technique for navigation in data (and not just geometric) space. In this way, scatterplot matrix navigation minimizes both the gulfs of evaluation and execution.<sup>32</sup> The method is also integrated with another fluid interaction technique called *query sculpting* where users select data items in one plot using a lasso or selection box, and

can refine their queries in other plots as they navigate the data space. As the plots change, the *query hulls* representing each selection animate smoothly as well to maintain the user’s mental model of the exploration.

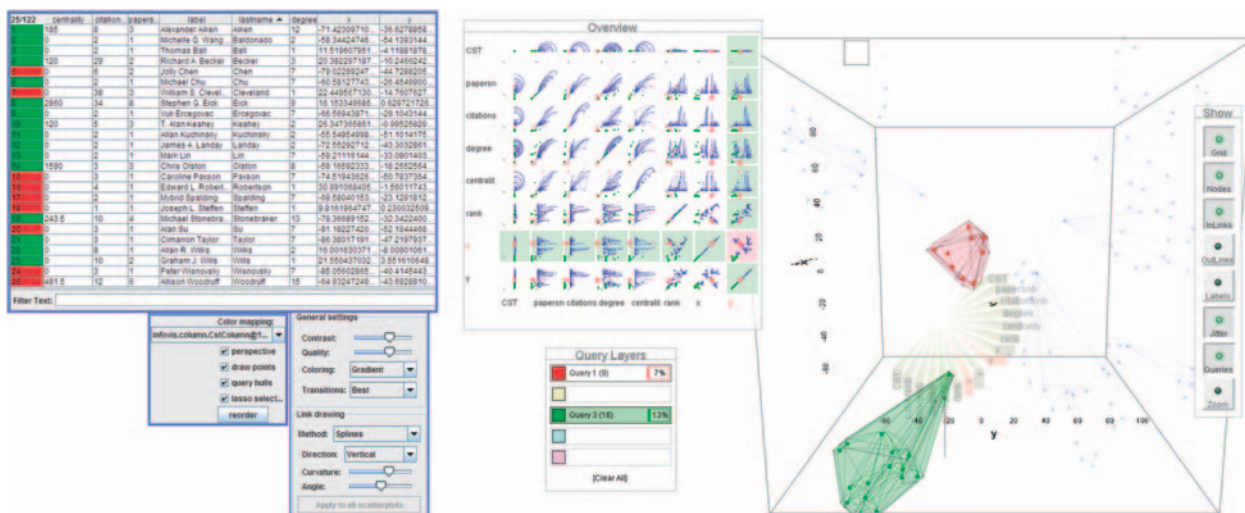
In summary, we can extract the following characteristic properties of fluidity from the ScatterDice and GraphDice tools:

- *Smooth animated transitions* between visualization states;
- *Minimalistic interface* using direct manipulation<sup>31</sup> (query sculpting);
- *Immediate visual feedback* for both query sculpting and navigation operations; and
- *Coherent conceptual model* that allows users to think and reason about the visual representation.


*Mæve*

Mæve<sup>8,9</sup> is an interactive multi-touch tabletop application for a museum installation (Figure 5) that employs a tangible interaction metaphor to enable the manipulation of a network graph, which itself represents the relationships between various architectural projects. When a paper card of a specific project is placed on the tabletop, its associated information structure appears, including media files, keywords, and related projects. The new card is then also visually connected to other cards that might be present in the table, to highlight the various similarities between the projects.

According to Bardzell et al.,<sup>7</sup> ‘It is easy to imagine a traditional browser-based presentation [...] But when you approach the same information through mæve,



**Figure 4.** The GraphDice<sup>5</sup> multivariate graph visualization tool showing the current plot (right), the plot matrix (top centre) and the query layer box (bottom center). The integrated spreadsheet-style table (left) shows details on demand. A FaST-slider<sup>48</sup> (active in the current plot) allows for quick navigation between data dimensions.



**Figure 5.** Mæve<sup>8,9</sup> installation at the Venice Biennale 2008, with a detail (inset) of the interactive tabletop surface.

something happens.’ This ‘something,’ according to Bardzell et al., is the ultimate combination of form and content, and content and interface. The result is that the visualization becomes an experience rather than a (productive) tool. Even more, sharing some characteristics with cinema, mæve is set up inside a fully controlled, darkened performance space, displaying user interactions projected on the wall in the space.

In general, mæve shows how the concepts of fluidity can transcend beyond the traditional screen media. Whereas the typical interface constraints are more apparent and restricting when working on a multi-touch tabletop medium, there is little that withholds the actual translation of these principles to typical information visualization applications on classic screen media.

Mæve exhibits the following fluid characteristics:

- *Responsive and immediate* visual feedback to interaction;
- *Focus on experience* rather than productivity; and
- *Powerful and effective use of novel inputs and outputs* that encourages experimentation.

### *We Feel Fine*

We Feel Fine,<sup>10</sup> developed by Jonathan Harris and Sep Kamvar (<http://www.wefeelfine.org/>), is an interactive exploration of contemporary human feelings. Its data are based on over 12 millions blog posts starting with the terms ‘I feel...’ that have been published online since 2005. Using a series of playful interfaces, these feelings can be searched and sorted across a number of demographic slices, offering responses to specific questions like: ‘Do Europeans feel sad more often than Americans?’ or ‘Does rainy weather affect how we feel?’

The interface uses a kinetic metaphor of a continuously self-organizing particle system, where each particle represents a single feeling posted by a single individual. The particles’ properties, such as their colour, size, shape, and opacity, indicate the nature of the feeling inside. The particles move about randomly around the screen until requested to self-organize in pictograms or along a number of data axes, thereby plotting and expressing various patterns of human emotion. The online application, which also is disseminated in a documentary book,<sup>49</sup> provides an engaging view of predominantly qualitative data that are inherently interesting and relevant to the user. The interaction is immediate and playful, where clicking (and even just moving) the computer mouse becomes an overarching goal on its own, with a visceral experience as a joyful reward. While the content pushes the user’s curiosity, the real user experience is the interaction. The interaction flow never really ends: there is always more to select, discover and explore, immersing the user in a true information experience. The aesthetics are compelling, with a consistent visual language and well-executed details throughout (e.g. shivering particles). Fluidity finds its summon in terms of how interactivity, visual design and content complement and augment each other.

We extract the following characteristic properties of fluid interaction from We Feel Fine:

- *Immediate and playful* interaction;
- *Never-ending interaction flow* that encourages exploration; and
- *Holistic interaction, visual design, and content* that complement and reinforce each other.

## Interactive holographics in *Iron Man 2*

Since the rising popularity of computer interfaces, the movie and television industry has embraced the persuasiveness of complex data representations. For the movie *Iron Man 2*,<sup>11</sup> special effects design firm Prologue (<http://www.prologue.com/>) designed a collection of incredibly dense information dashboards and highly responsive real-time 3D interfaces that demonstrate an impressive futuristic view of intuitive depictions of complex data (Figure 6)

The core of their vision builds upon the current knowledge about tabletop interaction, and imagines a future wherein information can be presented in a truly pervasive and directly manipulable way. Hand gestures, finger snapping, and pinches make information immediately queryable, results can be pushed aside and even thrown, and the highly contemporary design of the interfaces makes mundane task-oriented actions fun to accomplish. In terms of fluidity, *Iron Man 2* showcases opportunities when screen and interaction media physically disappear, and information immerses the user in life-size, high-definition resolution.

While the interfaces shown in the movie might neglect some obvious usability and usefulness constraints, they effectively demonstrate how far the typical characteristics of fluidity can be pushed in terms of immediateness, smoothness, and, in particular, the expressiveness of data. As with multi-touch technology, the persuasiveness of this future vision will further inspire many research agendas in the immediate future, and influence the way we will interact with information. In fact, with the emergence of ever-increasing LCD display sizes and the recent commercial releases of affordable controller-less interaction devices such as the Microsoft Kinect, this future might not be that far off.

The interactive holographics in *Iron Man 2* exhibits the following characteristics:

- *Reality-based interaction*<sup>37</sup> that transfers natural interactions such as gestures and body

language to the interaction with data and visual displays; and

- *Immediate visual feedback* that responds instantly to user interaction.

## Design guidelines for fluidity

Drawing on our observations and insights from the above review of InfoVis design exemplars, we here present a set of design guidelines for how to design and build effective information visualizations that support fluid interaction. While this list is far from exhaustive, our hope is it will serve as a common ground for other researchers and designers to improve and extend upon in the future.

- **DG1:** *Use smooth animated transitions between states.* Animated transitions help the user maintain an accurate mental model of the system's current state. Avoid abrupt mode switches because they are potentially disorienting and may break the user's flow. However, animation is a very strong visual variable<sup>43</sup> and designers should be careful not to take this design guideline too far. There exists a rich literature on this topic; e.g.<sup>44,50-50</sup>
- **DG2:** *Provide immediate visual feedback on interaction.* Do this for *every* key press or mouse motion, not just 'major' events ones like mouse clicks and the Enter key. In particular, this means that the visualization must be able to respond in real-time; if this is not possible, consider precomputing or simplifying some computation so that it becomes possible.
- **DG3:** *Minimize indirection in the interface.* If at all possible, use direct manipulation<sup>31</sup> so that interaction operations (filtering, selection, and details-on-demand) are integrated in the visual representation. In particular, avoid control panels that are separated from the visualization, or, if this is not possible, only put seldom used controls there.
- **DG4:** *Integrate user interface components in the visual representation.* If you cannot use direct manipulation



Figure 6. Tony Stark (Robert Downey, Jr.) interacting with the information-rich holographic displays in *Iron Man 2*.



and must use traditional interface components, like text fields, sliders, or buttons, try to make them a seamless, nearly embodied, part of the visualization. The interface should be nearly ‘invisible.’

- **DG5:** *Reward interaction.* Users should be encouraged to interact with a visualization so that the dialogue between user and system is initiated and maintained. Visually indicate how the user can interact with the visual representation, and when they do, give them a visceral reward. These rewards are effects that trigger a positive user experience without actually having a function in terms of visual exploration or visual communication; instead, their purpose is to keep the user stimulated while exploring. Examples of rewards include animations, sounds, and pretty graphics.
- **DG6:** *Ensure that interaction never ‘ends.’* The user should never reach a dead end where they can no longer proceed; it should always be possible to continue exploring the data. The system should be robust so that it allows all interactions without fear of crashing, long response times, or irreversible operations.
- **DG7:** *Reinforce a clear conceptual model.* The user should always have a clear idea of the state of the visualization and all interactions should be designed to reinforce this model. Operations should be reversible, allowing the user to return to a previous state. For visualizations using coordinated multiple views, connections between views should be clear and visible to the user.
- **DG8:** *Avoid explicit mode changes.* Instead of introducing different modes, integrate all operations in the same mode. This includes avoiding both drastic visual changes and drastic interaction modality changes. Mode changes may break the user’s flow.

## Vision and research directions

Interaction has a central role in visualization despite typically receiving much less emphasis than visual aspects:<sup>14</sup> it is the catalyst for the user’s dialogue with the data and, ultimately, for attaining insight and understanding. This dialogue between the user and the system is what sets our field apart from statistical data graphics and infographics, and it is clear this is where the true potential of visualization comes to the fore. A research agenda focusing on fluid interaction for information visualization will help fulfil this potential. However, we have barely scratched the surface of this exciting new line of research, and much work remains to be done. In the following subsections, we outline our vision for future research directions in this area.

## InfoVis exemplar repository

In this article, we review six different tools that serve as exemplars of fluid interaction for information visualization, three from the research community and three from the design and web community. Naturally, there are many more tools that could have been included in this collection, and our selection was mainly based on space constraints and the familiarity of the authors with different tools. Furthermore, we also do not mean to imply that the six tools reviewed here are paragons of perfection, but rather that they possess the kind of fluid interaction features that we wanted to highlight. Therefore, we hope that the reader will accept this review in the spirit it was collected in: a set of excellent visualizations that together help define the concept of fluidity in information visualization.

Having said that, a worthwhile future research direction may be to continue collecting these InfoVis exemplars into repositories for students, researchers, and designers alike to use as sources of inspiration and reference for their own projects. Such a repository should probably be curated (like the examples in this article) and include both commentaries and reviews by people other than the creators themselves to highlight the strengths (and weaknesses) of each tool.

Exemplar repositories of this kind would be similar to Edward Tufte’s books, which essentially are catalogues of good graphic design. However, while the examples in Tufte’s books are curated by the author himself, a larger question would be who should curate these InfoVis exemplars? Should it be single expert, or groups of experts? Should we have a single exemplar repository, or several? There may be a need to split these examples into different categories for academic, commercial, design and art—as we have seen in this article, the criteria for excellence vary widely between these different communities.

Fortunately, there already exist a set of blogs and online collections that already serve as exemplar repositories in this regard, and with the blog authors as the informal curators. For example, the *information aesthetics* (<http://infosthetics.com/>) weblog, created and maintained by one of the authors of this article, ‘explores the symbiotic relationship between creative design and the field of information visualization’ Many of the design exemplars discussed in this article are derived from entries posted on the infosthetics website. Other similar resources include Manuel Lima’s *VisualComplexity.com* (<http://www.visualcomplexity.com/>), Michael Friendly’s *Gallery of Data Visualization* (<http://www.math.yorku.ca/SCS/Gallery/>, which also includes counter examples of bad visual design), Robert Kosara’s *EagerEyes* (<http://eagereyes.org/>), and the Potsdam Information Design Patterns website (<http://infodesignpatterns.com/>)



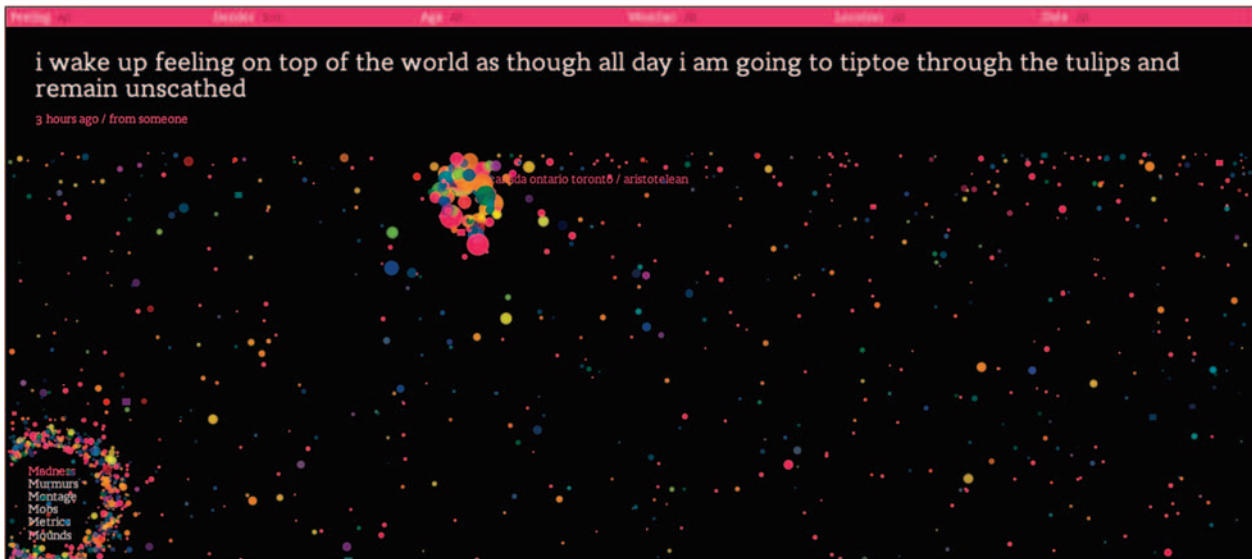


Figure 7. Sample interaction in the We Feel Fine application.

### Visualization design patterns

While our above design guidelines are short and practical guidelines on building effective and fluid information visualizations, an ultimate goal of this research may be to formalize the concept of *visualization design patterns* that build on the idea of ‘design patterns’ that was originally devised for architecture by Christopher Alexander in 1977,<sup>53</sup> but which has since been embraced by the software engineering community.<sup>12</sup> According to Alexander, a design pattern can be described as follows:

‘Each pattern describes a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice.’ — Christopher Alexander, *A Pattern Language* (1977).<sup>53</sup>

It is important to note the distinction between *visualization design patterns* and the *software design patterns for visualization* introduced by Heer and Agrawala<sup>54</sup>; the latter deals with pure software engineering strategies for building visualization software, whereas our proposed visualization design patterns focus on high-level design aspects of the visual representation, interaction, and user experience of the visualization itself. We think that there is a space and a need for defining this kind of pattern language (or at least a pattern catalogue) for information visualization in the same way as for the architecture and software engineering fields, thus standardizing terminology, perpetuating best practices, and improving visualization quality in general.

Adopting a pattern mindset is also advantageous because it allows us to define the notion of an *antipattern*<sup>45</sup>: standard solutions to common problems (in whatever field this is applied to) that simply do not work (or work poorly). Examples of such antipatterns may be the infamous rainbow colour scale (which both is not aesthetically pleasing according to most observers, and more importantly has perceptual limitations<sup>55</sup>), serious and distracting overuse of animation,<sup>43</sup> and overly complex visualizations that hide, rather than expose, important connections, or which show outright false relations. In general, counter examples are instructive in teaching people what *not* to do—examples include Huff’s classic *How to Lie with Statistics*,<sup>56</sup> as well as Michael Friendly’s gallery (discussed above), which includes not only the best but also the worst data visualizations on the Internet. Generalizing these to actual antipatterns would be beneficial towards increasing visualization quality overall.

### Towards visualization criticism

Both education and professional practice in mature design disciplines—such as architecture, industrial design, and graphic design—incorporate the notion of *criticism* as an integral component of the design process. This practice is slowly starting to be adopted for HCI and interaction design as *interaction criticism*.<sup>7,57</sup> The concept of ‘expert review’ has long been an established evaluation methodology in HCI, but these new advances consider aesthetics in addition to function and incorporate also the opinion of practitioners and users as part of the process.

Perhaps a similar notion of *visualization criticism* should be applied to our own field. In fact, a recent

survey<sup>58</sup> shows that one of the most common practical exercises in existing information visualization courses is to critique existing InfoVis tools. Furthermore, expert reviews have begun to gain traction in the visualization domain as well.<sup>59</sup> For the aesthetics component, the model proposed by Lau and Vande Moere<sup>60</sup> goes beyond subjective judgments to analyse the artistic influence and aesthetic engagement of a visualization. However, all these approaches tend to focus on the graphical aspects of a visualization rather than the interactive aspects.

At the core of this article lies the message that representation and interaction deserve equal treatment in visualization research. Therefore, we think that a worthwhile avenue for future work would be to combine all of the above approaches together with interaction criticism in an effort to identify visualization criticism as a skilled practice in its own right, one that focuses on the interplay between visuals and interaction along the lines of the discussions in this article. Not only would this provide a method for curating visualization exemplars (see above) as well as for educating the next generation of visualization designers and researchers, but it would also bring interactive—and not just visual—aspects of visualization tools to the fore.

## References

- Russell DM, Stefik MJ, Pirolli P and Card SK. The cost structure of sensemaking. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, Amsterdam, The Netherlands, 1993; pp. 269–276.
- Csikszentmihalyi M. *Flow: The Psychology of Optimal Experience*. New York: Harper Collins, 1991.
- Wattenberg M. Baby visualization names and social data analysis. In *Proceedings of the IEEE Symposium on Information Visualization, Minneapolis, MN, USA*, 2005; 1–6.
- Jetter H-C, Gerken J, Zöllner M, Reiterer H and Milic-Frayling N. Materializing the query with facet-streams – a hybrid surface for collaborative search on tabletops. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, Vancouver, BC, Canada, 2011.
- Bezerianos A, Chevalier F, Dragicevic P, Elmqvist N and Fekete J-D. GraphDice: A system for exploring multivariate social networks. *Comput Graphics Forum* 2010; 29(3): 863–872.
- Elmqvist N, Dragicevic P and Fekete J-D. Rolling the dice: multidimensional visual exploration using scatterplot matrix navigation. *IEEE Transactions on Visualization and Computer Graphics* 2008; 14(6): 1141–1148.
- Bardzell J, Bolter J and Löwgren J. Interaction criticism: Three readings of an interaction design, and what they get us. *Interactions* 2010; 17(2): 32–37.
- mæve. <http://portal.mace-project.eu/mæve/>, 2008. Accessed November 2010.
- Nagel T, Pschetz L, Stefaner M, Halkia M and Müller B. mæve – an interactive tabletop installation for exploring background information in exhibitions. In *Proceedings of Human-Computer Interaction International*, volume 5612 of *Lecture Notes in Computer Science*, San Diego, CA, USA, 2009, pp. 483–491.
- Harris J and Kamvar S. *We feel fine*, <http://www.wefeelfine.org/>, 2006. Accessed November 2010.
- Cone J. *Prologue: Iron Man 2*. <http://motionographer.com/features/prologue-iron-man-2/>, 2010. Accessed November 2010.
- Gamma E, Helm R, Johnson R and Vlissides J. *Design Patterns: Elements of Reusable Object-Oriented Software*. Reading, MA: Addison-Wesley, 1995.
- Amar RA, Eagan J and Stasko JT. Low-level components of analytic activity in information visualization. *Proceedings of the IEEE Symposium on Information Visualization* 111–117, 2005.
- Yi JS, Kang Ya, Stasko JT and Jacko JA. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions of Visualization and Computer Graphics* 2007; 13(6): 1224–1231.
- Bertin J. *Sémiologie graphique: Les diagrammes - Les réseaux - Les cartes*. Paris, France: Editions de l'École Hautes Etudes en Sciences, 1967, les réimpressions edition.
- Mackinlay JD. Automating the design of graphical presentations of relational information. *ACM Trans Graphics* 1986; 5(2): 110–141.
- Card SK and Mackinlay J. The structure of the information visualization design space. In *Proceedings of the IEEE Symposium on Information Visualization*, Phoenix, AZ, USA, 1997, pp. 93–99.
- Tufte ER. *The Visual Display of Quantitative Information*. Cheshire, Connecticut: Graphics Press, 1983.
- Tufte ER. *Envisioning Information*. Cheshire, Connecticut: Graphics Press, 1990.
- Tufte ER. *Visual Explanations*. Cheshire, Connecticut: Graphics Press, 1997.
- Tufte ER. *Beautiful Evidence*. Cheshire, Connecticut: Graphics Press, 2006.
- Few S. *Show Me the Numbers: Designing Tables and Graphs to Enlighten*. Oakland, USA: Analytics Press, 2004.
- Few S. *Information Dashboard Design: The Effective Visual Communication of Data*. Sebastopol: O'Reilly, 2006.
- Ware C. *Information Visualization: Perception for Design*, 2nd edition. San Francisco: Morgan Kaufmann Publishers, 2004.
- Ware C. *Visual Thinking for Design*. San Francisco: Morgan Kaufmann Publishers, 2008.
- Spence R. *Information Visualization: Design for Interaction*. Essex, UK: Pearson Education, 2007.
- Segaran T and Hammerbacher J. editors. *Beautiful Data: The Stories Behind Elegant Data Solutions*. Sebastopol: O'Reilly, 2009.
- Ward M, Grinstein G and Keim D. *Interactive Data Visualization*. Natick, MA: A K Peters, 2010.
- Few S. *Now You See It: Simple Visualization Techniques for Quantitative Analysis*. Oakland, USA: Analytics Press, 2009.
- Pike WA, Stasko JT, Chang R and O'Connell TA. The science of interaction. *Inform Visual* 2009; 8(4): 263–274.
- Shneiderman B. Direct manipulation: a step beyond programming languages. *Computer* 1983; 16(8): 57–69.
- Norman DA and Draper SW. editors. *User Centered System Design: New Perspectives on Human-Computer Interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1986.
- Beaudouin-Lafon M. Instrumental interaction: an interaction model for designing post-WIMP user interfaces. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, The Hague, Netherlands, 2000, pp. 446–453.
- Beaudouin-Lafon M. Designing interaction, not interfaces. *Proceedings of the ACM Conference on Advanced Visual Interfaces*, Gallipoli, Italy, 2004, pp. 15–22.
- Ishii H and Ullmer B. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, Atlanta, GA, USA, 1997, pp. 234–241.
- Dourish P. *Where the Action Is: The Foundations of Embodied Interaction*. Cambridge, USA: MIT Press, 2001.

37. Jacob RJK, Girouard A, Hirshfield LM, Horn MS, Shaer O, Solovey ET and Zigelbaum J. Reality-based interaction: a framework for post-WIMP interfaces. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, Florence, Italy, 2008, pp. 201–210.
38. Bederson BB. Interfaces for staying in the flow. *Ubiquity* 2004; 5(27).
39. Hutchins EL, Hollan JD and Norman DA. Direct manipulation interfaces. *Human-Comput Interac* 1985; 1(4): 311–338.
40. Dix AJ and Ellis G. Starting simple: adding value to static visualisation through simple interaction, In *Proceedings of the ACM Conference on Advanced Visual Interfaces*, L'Aquila, Italy, 1998, pp. 124–134.
41. Norman DA. *Things that Make Us Smart: Defending Human Attributes in the Age of the Machine*, chapter Visual Representations. Reading, USA: Addison-Wesley, 1994.
42. Pirolli P and Card SK. Sensemaking processes of intelligence analysts and possible leverage points as identified through cognitive task analysis. In *Proceedings of the International Conference on Intelligence Analysis*, McLean, VA, 2005.
43. Tversky B, Morrison JB and Bétrancourt M. Animation: can it facilitate? *Int J Human-Computer Stud* 2002; 57(4): 247–262.
44. Dragicevic P, Bezerianos A, Javed W, Elmqvist N and Fekete J-D. Temporal distortion for animated transitions. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, Vancouver, BC, Canada, 2011, pp. 2009–2018.
45. Koenig A. Patterns and antipatterns. *J Object-Oriented Program* 1995; 8(1): 46–48.
46. Young D and Shneiderman B. A graphical filter/flow representation of boolean queries: a prototype implementation and evaluation. *J Am Soc Inform Sci* 1993; 44(6): 327–339.
47. Byron L and Wattenberg M. Stacked graphs-geometry & aesthetics. *IEEE Transactions on Visualization and Computer Graphics* 2008; 14(6): 1245–1252.
48. McGuffin M, Burtnyk N and Kurtenbach G. FaST Sliders: Integrating Marking Menus and the Adjustment of Continuous Values. In *Proceedings of Graphics Interface*, Calgary, AB, Canada, 2002, pp. 35–41.
49. Kamvar S and Kamvar J. *We Feel Fine: An Almanac of Human Emotion*. New York: Scribner, 2009.
50. Hudson SE and Stasko JT. Animation support in a user interface toolkit: Flexible, robust, and reusable abstractions. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, Atlanta, GA, USA, 1993, pp. 57–67.
51. Heer J, Viégas FB and Wattenberg M. Voyagers and voyeurs: supporting asynchronous collaborative information visualization. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, San Diego, CA, USA, 2007, pp. 1029–1038.
52. Robertson GG, Fernandez R, Fisher D, Lee B and Stasko JT. Effectiveness of animation in trend visualization. *IEEE Transactions on Visualization and Computer Graphics* 2008; 14(6): 1325–1332.
53. Alexander C, Ishakawa S and Silverstein M. *A Pattern Language: Towns, Buildings, Construction*. New York: Oxford University Press, 1977.
54. Heer J and Agrawala M. Software design patterns for information visualization. *IEEE Transactions on Visualization and Computer Graphics* 2006; 12(5): 853–860.
55. Munzner T. Process and pitfalls in writing information visualization research papers. *Lecture Notes in Computer Science* 2008; 4950: 134–153.
56. Huff D. *How to Lie with Statistics*. New York: Norton, 1954.
57. Bardzell J. Interaction criticism and aesthetics. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, Boston, MA, USA, 2009, pp. 2357–2366.
58. Kerren A, Stasko JT and Dykes J. Teaching information visualization. In *Information Visualization – Human-Centered Issues and Perspectives*, volume 4950 of *Lecture Notes Comput Sci*, pp. 65–91, 2008.
59. Tory M and Möller T. Evaluating visualizations: Do expert reviews work? *IEEE Comput Graphics Appl* 2005; 25(5): 8–11.
60. Lau A and Moere AV. Towards a model of information aesthetics in information visualization. In *Proceedings of the International Conference on Information Visualization*, Zürich, Switzerland, 2007, pp. 87–92.