

Business Intelligence Analytics

Danyel Fisher, Steven Drucker, and Mary Czerwinski ■ Microsoft Research

Businesses are increasingly monitoring and tracking data about what it takes to keep themselves running. They collect and maintain increasingly available data, such as

- transaction and sales data stored in data warehouses,
- server log files tracking visitors,
- data from sensors tracking delays on factory floors,
- IT data logs, and
- data on their competitors and industrial sectors.

Data-driven decision making—orienting business decisions around data—drives major IT initiatives across all business sectors.

Although buzzwords such as “data-driven decisions” and its “big data” and “data science” cousins might be overused, underneath them lie real opportunities for organizations to understand and reflect on their processes and customers.

Collecting the information is only the first step. A great deal of research has investigated the process by which people work with data: how they ask questions about it, make sense of it, and communicate with it. A critical part of this chain is information visualization tools. Visualization provides a powerful way to make sense of data. By mapping data attributes to visual properties such as position, size, shape, and color, visualization designers leverage perceptual skills to help users discern and interpret patterns in data. For a first foray into the information visualization field, see *Readings in Information Visualization: Using Vision to Think*¹ and “Interactive Dynamics for Visual Analysis.”²

Business Intelligence

Business intelligence (BI) is the practice of in-

terpreting and visualizing data to make useful business-oriented decisions. BI tasks occur in offices, universities, and data centers and provide data-oriented lifeblood to research and business organizations worldwide. BI systems must often appeal to broad audiences, from knowledge workers to CEOs to stockholders. They must allow for rapid analysis for decision making, developing insights, and communicating those insights' results.

BI tasks fall into three categories:

- exploring and analyzing data,
- monitoring ongoing dataflows through dashboards, and
- communicating insights to others, both in and outside a company.

Much BI work happens around data that's stored specifically for BI analysis, often called a *data warehouse*. Data in those warehouses might be stored in a precomputed structure called an OLAP (online analytical processing) cube. OLAP cubes make it easy to look at various data aggregates, filtering on a number of selected fields.

The most common data structures are frequently addressed through commercial off-the-shelf tools. Both companies dedicated to visualization, such as Tableau, SAS, and Qlik, and established data companies, such as IBM, Oracle, SAP, and Microsoft, produce tools targeted at BI needs related to structured data stored in data warehouses.

As more data becomes available, it comes in many forms that differ widely from the typical OLAP structure. For example, geographic data, log data, and network data are increasingly common. Although information visualization researchers have suggested many ways to visualize these types of data, they have less often created tools adapted for BI needs.

In This Issue

Here, we turn the spotlight on BI as an area of inquiry and explore beyond the current standard practices. We look to learn about the processes that practitioners currently follow in this area and how new BI techniques and capabilities will help users understand and act on widely disparate types of data. We also look to see how organizations make decisions around datasets, what purposes visualizations are used for, and what different representations people use to show and explore data.

This special issue differs in two critical ways from many other issues of *IEEE CG&A*.

First, we selected the articles for their ability to represent interesting business challenges. We haven't edited any of the visualizations. If a practitioner has been able to get his or her audience to understand his or her representation, we wished to see and share the visualization. We're sure that many readers will have other ideas on how data could have been presented for many of the visualizations.

Second, many business decisions require confidentiality. Although most of these articles are based on real scenarios encountered in daily work, some authors weren't able to release exact data or the precise decisions that were made with it. Nonetheless, we trust you'll find their accounts illustrative and illuminating.

The first article, "Visual Business Ecosystem Intelligence: Lessons from the Field," functions as an overview. The author, Rahul Basole, is an internal consultant on visualizations at the Georgia Institute of Technology and works with a variety of teams with different decision needs. Here, he presents his approach to help them narrow down their tasks and data and to choose appropriate representations that let them make decisions related to those tasks. He concludes with challenges visualization designers face in BI scenarios that might differ from the challenges faced in other fields.

In "Applying a Sunburst Visualization to Summarize User Navigation Sequences," Kerry Rodden examines a specific task at YouTube. YouTube users work their way through the website, clicking between videos, search pages, and user pages. This structure is unlike many conventional e-commerce sites, which funnel users toward a purchase. Her team wanted to understand how users were using the site. Rodden describes her design and development process and the complexities involved in balancing user needs, task understanding, and tool capabilities.

Although visualization is important, it's only a part of how work gets done in business set-

tings. In "From Data to Insight: Work Practices of Analysts in the Enterprise," Eser Kandogan and his colleagues report on business analytics and BI in large companies. Previous researchers have also looked at interactions with data in the workplace.^{3,4} Kandogan and his colleagues' perspective complements some of these past projects and casts light on limitations related to how visualization affects the use of BI in organizations. Readers might find the article a spur to overcome some of the obstacles this article reports—or might be challenged to think about how visualization and BI will continue to work in the future.

We then turn to two applications articles. First, in "A Visual-Analytics System for Railway Safety Management," Wallace Lira and his colleagues describe a visualization of run-over risk factors for train routes—a problem with life-and-death consequences. This article is a prime example of how geographic visualization can provide a context for understanding situations and events for communicating problems or successes to management.

Finally, in "Business Intelligence from Social Media: A Study from the VAST Box Office Challenge," Yafeng Lu and his colleagues report on a visual-analytics application of social media data. As part of the Visual Analytics Science and Technology Box Office Challenge, their academic team worked to predict movie revenue on the basis of tweets about the movies. They present both the visual and data analysis techniques in elegant detail. Although their research isn't based on a specific organization, its applications to commerce and marketing are obvious.

In these five articles, many sorts of common BI visualizations aren't represented. The area is simply too broad to do full justice in just a few articles. We encourage you to learn about the many other business uses of visualization. Perhaps two of the most common include the use of dashboards as monitoring tools and the adaptation of visualizations for presentations to large audiences.

This issue's articles show just a sliver of how visualization can help inform the ways that organizations use data. We hope these articles provoke new ideas and conversations. We also hope that researchers can see business as a fertile domain in which new visualization techniques can apply to a broad selection of use cases, audiences, and datasets. Finally, we hope that people looking at data in a business context can use these articles as inspiration and as a boost into new, exciting areas for working with data. 

References

1. S.K. Card, J.D. Mackinlay, and B. Shneiderman, *Readings in Information Visualization: Using Vision to Think*, Morgan Kaufman, 1999.
2. J. Heer and B. Shneiderman, "Interactive Dynamics for Visual Analysis," *Comm. ACM*, vol. 55, no. 4, 2012, pp. 45-54.
3. S. Kandel et al., "Enterprise Data Analysis and Visualization: An Interview Study," *IEEE Trans. Visualization and Computer Graphics*, vol. 18, no. 12, 2012, pp. 2917-2926.
4. D. Fisher et al., "Interactions with Big Data Analytics," *Interactions*, vol. 19, no. 3, 2012, pp. 50-59.

Danyel Fisher is a researcher in Microsoft Research's Visualization and Interaction Research Group. His main research interest is information and data visualization. Fisher received a PhD in information and computer science from the University of California, Irvine. Contact him at danyelf@microsoft.com.

Steven Drucker is a principal researcher in Microsoft Research's Visualization and Interaction Research Group. His research focuses on human-computer interaction for dealing with large amounts of information. Drucker received a PhD in arts and media technology and an MS in brain and cognitive sciences, both from MIT. Contact him at sdrucker@microsoft.com.

Mary Czerwinski is a research manager in Microsoft Research's Visualization and Interaction Research Group. Her research interests are emotion tracking, information worker task management, multitasking, and awareness systems for individuals and groups. Czerwinski received a PhD in cognitive psychology from Indiana University Bloomington. Contact her at marycz@microsoft.com.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.

