



## Lessons Learned from Designing Visualization Dashboards

**Maria-Elena Froese**

*University of Victoria*

**Melanie Tory**

*Tableau Research*

A vast array of knowledge is available in the visualization field about which visual designs are more effective for which tasks.<sup>1,2</sup> Yet, the actual design of visualization interfaces and tools is still an art, and the practice of implementing known design guidelines is nontrivial. In this article, we discuss our lessons learned from designing applications for small start-up companies and institutions.

In each of our short-term design projects, we started with a vague description of the requirements and proposed solutions in close collaboration with our industrial partners. We followed an Agile software development approach, meeting our industrial partners every week or two. Our process included gathering and documenting requirements, creating static wireframes that were iteratively improved with feedback from expert users, implementing functional prototypes using D3.js,<sup>3</sup> and conducting usability testing. In every company there were at least two expert users involved throughout the design process.

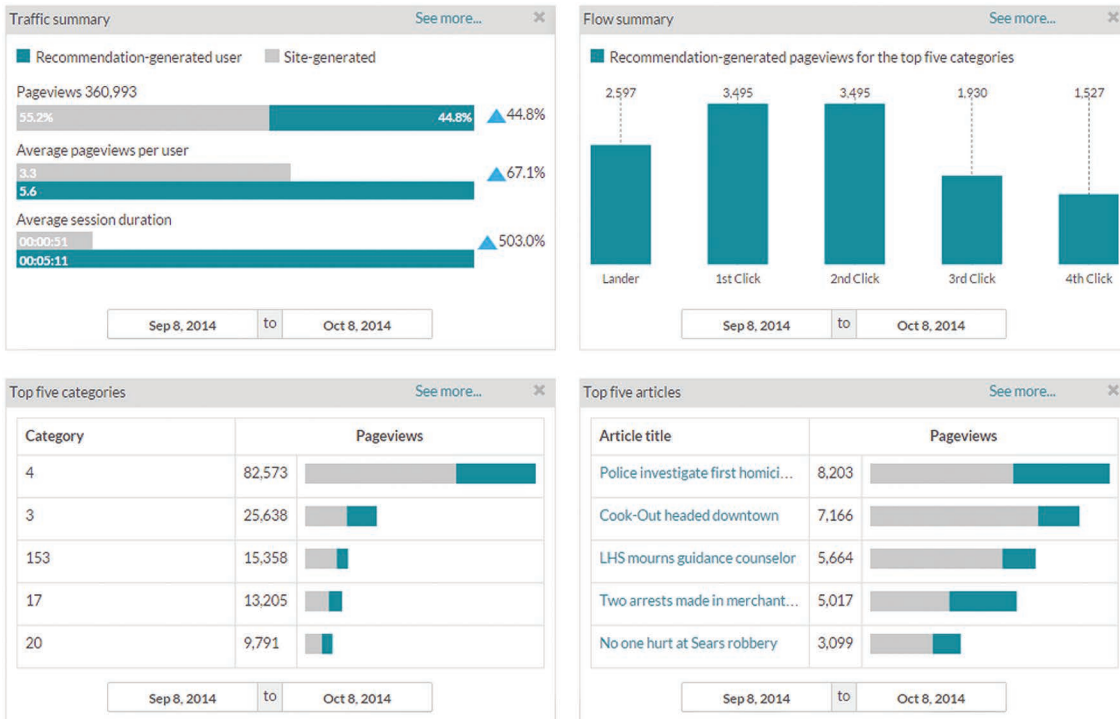
We noticed several commonalities in the design and development of the visualization dashboards. First, we always considered existing approaches, but in each case we concluded that standard, off-the-shelf visualization construction interfaces were insufficient: either the data needed too much preprocessing to be ready for those tools and/or the learning curve was too steep for the target audience. Second, the target audiences of the visual analytics applications were not experts in visualization or statistical reasoning. The users were mostly middle- to high-level executives who were typically knowledgeable about the measurements

for their field of interest, had a small number of specific questions about their data, had limited time to explore the data, and had previous experience with other dashboards. Third, the main challenge in all these projects was to design solutions with an easy point of entry that did not require specialized training and would encourage users to explore the data and customize the visualizations.

In this article, we briefly describe some of the representative requirements for each of our partners and describe our design solutions. As a part of this discussion, we also emphasize the need for a quicker way to create functional prototypes, point out frequent misconceptions on the scope of a functional prototype, discuss how performance can affect prototyping, and discuss our industrial partners' resistance to provide us access to their customers to establish requirements.

### Design Examples

We have worked with several small organizations. For the purposes of this discussion, we will review the challenges we encountered when designing dashboards for four companies: Youneeq, a company that provides a recommendation engine for online media publishers designed to get readers more engaged in their media content; Trusterra Technologies, a company that created a platform for documenting safety events in the workplace; Tutela Technologies, a company concerned with telecommunications network performance; and Ocean Networks Canada (ONC), a nonprofit that collects and analyzes ocean and seismic data. For simplicity, we will refer to all of them as our "industrial partners."



**Figure 1. Youneeq media content dashboard.** This overview shows four panels (traffic summary, flow summary, top five categories, and top five articles) summarizing the behavior of two personas: media consumers who followed links into recommended articles (darker teal color) and those who did not (grey color). The user can drill down into a visualization (by clicking on “See More...”), choose a new date range, adjust the location of the panels, and add or remove panels.

### Youneeq

Youneeq created a recommendation engine to be used by online media publishers. The recommendation engine provides content tailored to users’ profiles, making a website more competitive for digital advertisement and e-commerce. Youneeq monitor the traffic of its customers’ websites by collecting a set of key performance indicators (KPIs) daily, such as the total number of pages visited (page views), how many uniquely identified clients arrived (visitors), how many pages were viewed per visitor (page views per user), how often users left the site after reading an article (exit rate), and how long visitors spent per visit (session duration). Our partners wanted a dashboard to assess whether their recommendation engine was having a positive impact on e-commerce KPIs by comparing the behavior of visitors who were using the recommendations with those who did not. Target users were managers, executives, and salespeople.

Existing Web analytics dashboards offer the visualization of a comprehensive set of KPIs, but they do not provide a means to compare the traffic for different personas. Here are a few examples of user stories:

- “I want to compare the most important KPIs

for visitors who use the recommendations with those who do not.”

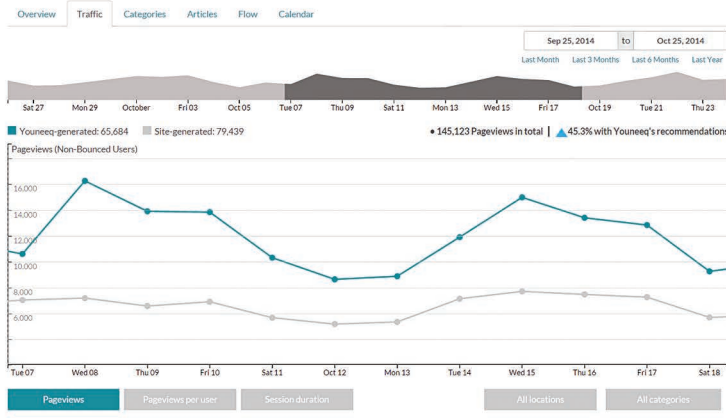
- “I want to know what content was most successful, and I want to know the impact of the recommendations on that success.”
- “I want to compare the traffic of the user during summer 2015 versus summer 2014.”

Figures 1 and 2 show the dashboard we designed for Youneeq. The overview in Figure 1 is aimed to provide a quick summary for high-level business executives who wanted to get a quick sense of how the recommendation engine was affecting their website’s traffic. The overview initially displayed five panels:

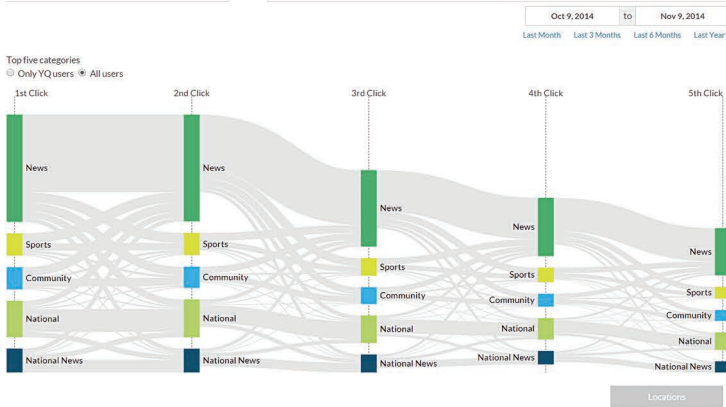
- a traffic summary showing three KPIs (page views, page views per user, and average session duration) for the two personas,
- a flow summary of traffic movement among the most popular categories,
- page views for the top five categories,
- page views for the top five articles, and
- a histogram of page views per day of the week (not shown).

The user can relocate, add, and remove panels. Each panel displays data from the last 30 days by

# Applications



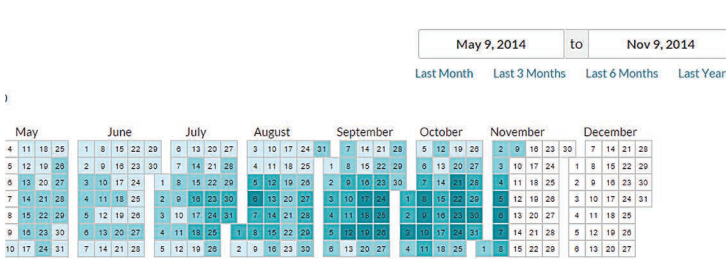
(a)



(b)

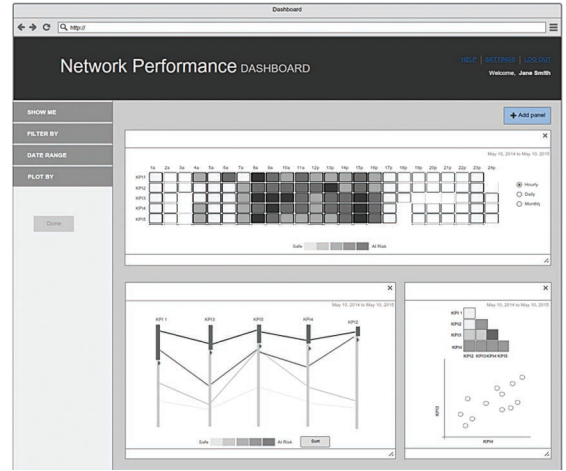
Rank	Category name	Pageviews (Non-Bounced Users)	Pageviews per user (Non-Bounced Users)	Visitors (Non-Bounced Users)
1	4			
2	3			
3	153			
4	20			
5	38			
6	17			
7	39			
8	18			
9	12			
10	37			

(c)



(d)

**Figure 2. Detailed panels for the Youneeq media content dashboard.** The user can see (a) KPIs over time, (b) the sequence of user clicks, (c) the list of the most visited categories, and (d) the site's traffic in a calendar view.



**Figure 3. Tutela Technologies network performance dashboard.** In its initial state, the dashboard displays three panels containing a heat map, a parallel coordinates plot, and a correlation matrix for the five most important KPIs. The dashboard can then be customized based on individual user needs via the settings panel on the left. For example, the user can add or remove KPIs on a given visualization, drill down into a visualization, filter data, choose a new date range, adjust the size and location of the panels, and add or remove panels.

default, but the date range is adjustable. Changing the time range on one panel intentionally does not affect the other panels, allowing users to duplicate a visualization to compare time ranges.

Figure 2 shows several of the detailed views. The traffic view compares traffic over time for recommendation-generated users and site-generated users, along with an interactive overview of all traffic over time (see Figure 2a). The flow view shows a Sankey diagram with the top five categories through the first five user interactions (see Figure 2b). The categories table shows KPIs for the most popular categories (see Figure 2c). Lastly, the calendar reveals page views over time as a calendar heat map (see Figure 2d).

## Tutela Technologies

Tutela develops software products to allow network service providers to monitor network performance. Its software enables field testers to collect simultaneous samples of about 40 KPIs such as bit rate, rebuffering time, and dropped packages. They also record information on the test settings, including location, connection type, device, and operating system. Tutela's customers need to explore the field test data and identify underlying causes of poor network performance. The end users are network performance analysts, including technicians, call center personnel, and midlevel executives. Here are a few examples of user stories:

- “I want to find the correlation between a group of KPIs.”
- “I want to see a report with a high-level view of the five most important KPIs.”
- “I want to identify the main networking problems.”

We designed a customizable dashboard with a two-column layout (see Figure 3). We implemented five types of visualizations:

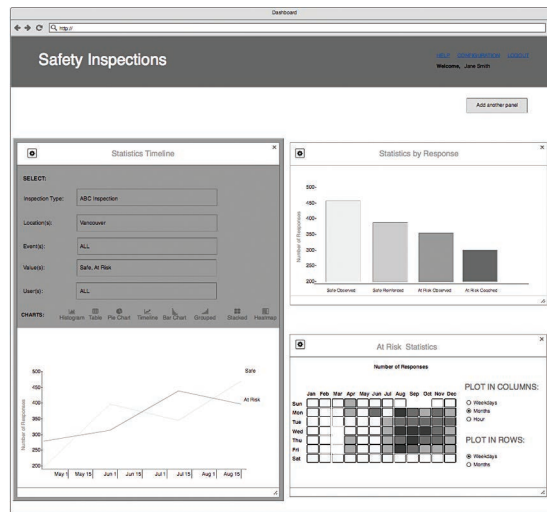
- The heat map provides an overview of the selected KPIs, where a darker color represents at risk values (outside a threshold). Rows in the heat map correspond to the selected KPIs, and columns correspond to units of time.
- The parallel coordinates visualization supports finding trends and outliers across all KPIs. We allowed brushing (selecting data subsets with a mouse) on the KPI axes.
- Multiple timelines allow the user to see changes in selected KPIs. We colored the lines to distinguish KPIs in relation to their threshold: darker colors represent at risk values. Below the stacked timelines we added a heat map overview based on the Line Graph Explorer.<sup>4</sup>
- The correlation matrix enables users to explore relationships between selected KPIs. Darker colors indicate higher correlation. Users can choose a cell to see the corresponding scatterplot.
- Histograms depict the distribution of values in the selected KPIs and support brushing and linking (updating one view based on selections in another).

In its initial state, the dashboard displays only three panels, as shown in Figure 3. Panels can be added, moved, and deleted. The user can choose a new segment of time for any given visualization (by default, it shows the last 30 days). To help analysts compare two segments of time using the same visualization, a change in one panel does not affect the others.

### Trusterra Technologies

Trusterra develops a tracking system that is used by its customers to report safety inspections (or incidents) in the workplace. The data is collected using customizable forms, which mainly contain multiple-choice fields but can also contain text, images, audio, and video. The collected data is used by business managers to identify problems for which they can devise corrective or preventive actions.

Trusterra needed a customizable tool that could allow its customers to create a range of interactive



**Figure 4. Trusterra Technologies safety inspections dashboard. In this view, the user is adding a third panel, which has its settings section expanded. The user can choose the name of the safety inspection, the locations, the description of the incident or event, the corresponding values, and the name of the employee who reported the incident. The user can then choose among suitable charts for the selected data. The user can then drill down into a visualization, choose a new date range, adjust the location of the panels, and add or remove panels.**

reports and that would be flexible enough to handle the multiple kinds of data that their forms could collect. Here are a few examples of user stories:

- “I want to see the overall statistics for the selected type of safety inspection.”
- “I want to see statistics for selected locations, types of observations, and who reported a selected type of safety inspection.”

Our visualization dashboard contains three panels showing basic statistics and allows the user to add, relocate, and remove panels as desired. Each panel has two sections: settings and view. Within the settings section, the user can choose data for five selection boxes (see Figure 4): the type of safety inspection, one or more locations, one or more events or observations, the associated values, and the employee who reported the incident. Each selection box can contain zero, one, or many data items. When the user does not make a selection of data items, the system includes all the data items.

Our design principle for this dashboard was to allow users to focus on choosing the data they wanted to see. We defined charting rules so that the system only allowed charts that were appropriate for the selected data. For example, a histogram is allowed only when the data to plot

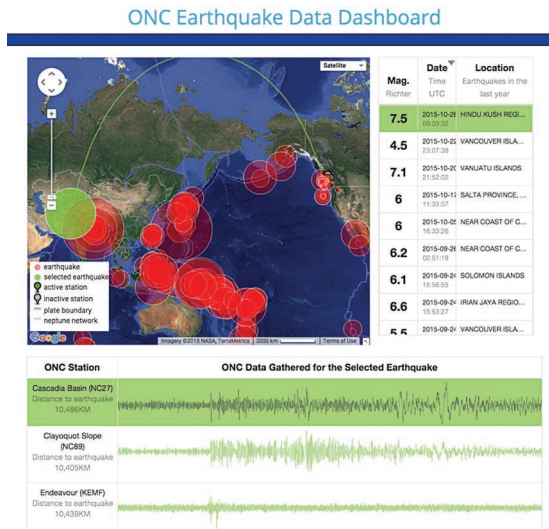


Figure 5. Ocean Networks Canada dashboard for seismic data. The ONC dashboard contains three coordinated panels: a map with the geographic location of the earthquakes; a table with the earthquake’s magnitude, data, and location; and a table with the raw sensor signals. A selected earthquake is highlighted in green in all panels. By default, the station (and its corresponding plot) on the top of the list is highlighted. The user can adjust the size and location of the panels. The dashboard shows data within the last year, and the seismic events table is initially sorted in descending order.

consists of one numerical item. A table is recommended when we have two selection boxes with multiple items (for example, multiple cities and multiple values, such as “safe” and “at risk”) or one large selection (such as a list of 15 locations). When the user selects up to five measures but no independent variable, a timeline is recommended. To cover all possible cases, we needed to create an extensive list of charting rules.

**Ocean Networks Canada**

ONC operates a network of sensors to detect seismic events, both inland and beneath the ocean’s surface. The time series data collected from their sensors is made available to the public through the Incorporated Research Institutions for Seismology (IRIS) website. Our partners needed a dashboard to display when and where their sensors had detected seismic events. The users included both members of the public and scientists interested in seismic events. Here are a few examples of user stories:

- “As a member of the public, I want to identify the geographic location of significant and recent seismic events.”
- “As a member of the public, I want to explore local, regional, and global seismic activity.”

- “As a scientist, I want to see the signal captured by the seismometers.”

We designed a dashboard with three panels: a map showing the location of the seismic events (as circles) and distance from the sensors (as lines); a sortable table of seismic events including magnitude, date, and location; and a table with the list of stations (including distance to the seismic event) and the signals gathered from the stations (see Figure 5). See the live dashboard at [www.oceannetworks.ca/data-tools/earthquake-data-dashboard](http://www.oceannetworks.ca/data-tools/earthquake-data-dashboard).

**Lessons Learned**

Working with our industry partners on dashboard applications revealed several practical issues, particularly in the areas of design, prototyping, and establishing expectations, requirements, and usability studies in a competitive commercial environment.

**Dashboard Designs**

Each our industrial partners had a set of KPIs that they needed to make available to their customers for exploration. Their customers were usually not interested in seeing all the KPIs at once, but preferred to see just the subset that supported the specific and concrete questions they were investigating. Our designs therefore included the ability to customize the dashboard by adding or removing KPIs and charts, rearranging the layout, drilling down into the data, brushing and linking, and applying filters. We took care to ensure an easy entry point for each user. Thus, we typically prepopulated the dashboard with default charts that could be customized easily so that users were not faced with the mystery of a blank canvas.

Another design goal was to release the user from making unnecessary design decisions. Although we found that it was easy for the users to choose which KPIs they want to explore, choosing the most effective type of chart is not trivial. Previous research has shown that users often choose charts they are already familiar with,<sup>5</sup> and we similarly noted a frequent request for pie charts. Moreover, mapping the selected KPIs to the visual elements of the chart can be challenging.<sup>5</sup>

Thus, our approach was to design the visualizations semiautomatically, building on pioneering early work in automated information presentation.<sup>6</sup> Once we identified the tasks and the relevant KPIs for those tasks, we could predetermine which visualizations were best. Then we either allowed the user to choose from premade charts or choose their data, based on which data the system rec-

ommended for that particular the chart. In both cases, we relieved users from making visualization design decisions and let them simply choose what they wanted to see.

When choosing the dashboard layout, we found it is necessary to consider visualization size and density. Dashboard layout is easy when using simple charts like bar charts, histograms, pie charts, maps, and small tables like those designed for the Youneeq overview and the Trusterra and ONC dashboards. However, more complex charts such as overview plus detail timelines, Sankey diagrams, TableLens, parallel coordinates, scatterplot matrices, and small multiples charts are harder to pack into a single dashboard. For Tutela, we tried a few alternatives. Dense charts can be downsized to fit into the dashboard like simple charts but at the risk of losing legibility and fidelity. More detail can be provided when full-sized charts are stacked vertically, but this limits the number of charts that can be easily compared at once. Different tabs allow room for even more detail, but they make it harder to compare charts. The best approach depends on the visualizations to be supported by the dashboard and the level of detail needed in them.

### **Prototypes**

We found that prototype performance can hinder the user experience. We frequently encountered performance problems when implementing our prototypes. For Youneeq and ONC we used full datasets. For Tutela we used a small sample of data. Regardless of the size, the responsiveness of a visualization prototype can be suboptimal for the user experience. For example, in the Youneeq dashboard, we let users create dynamic queries by choosing the time range, location, and category of the media publications. To support this, our prototype had to access the server to fetch new data. For each query, updating took about 10 seconds, which was too long for some users.

In general, the performance of our prototypes depended on several factors: the time needed to fetch data from the server, the complexity of the computations necessary to analyze the data, and the time taken to create the SVG elements that make up each visualization. We explored two approaches: sampling to reduce the data size and partitioning the data so that it could be loaded in pieces. Both approaches have tradeoffs: sampling may miss important details or show a nonrepresentative distribution, and partitioning may require multiple server retrievals. Server-side processing may be a complementary solution, but in general, research into prototyping tools that support rapid

iterative development with a higher level of system performance would be highly beneficial.

Quicker, lightweight prototyping mechanisms are also needed. Discovering that a chart is effective for the particulars of the data can take too long. Sometimes, design ideas portrayed in static wireframes work well for the specification of the tasks and data, but when implementing those ideas with real data, we may discover unforeseen issues. For example, in the Tutela project, we wanted to see the distribution of multiple KPIs at once. Based on the paper prototype, it looked like small multiple histograms would be appropriate, but after implementing this view with the real data, we discovered that the distribution of the data was far from normal because most of it fell into only one or two bins of the histograms.

To achieve an adequate visualization earlier in the design process, the visualization designer needs to have access to the data as soon as the requirements gathering stage starts and should create quicker, lightweight, throw-away prototypes with real data using tools such as Tableau, R, or Matlab. Other researchers have similarly noted this need.<sup>7</sup>

Lastly, it is important to establish expectations regarding prototypes. We used a variety of prototypes, ranging from static wireframes to functional applications. Using static wireframes let us effectively convey design ideas in early iterations. They helped us with what tasks to accomplish, what data to represent, and what visual elements to use. These static prototypes were also used to run early usability studies.

However, because the wireframes frequently looked similar to a running application, it was sometimes hard for users to remember that they were not dealing with a functional prototype. In contrast, creating functional prototypes was an excellent strategy for illustrating interaction design and delivering a proof of concept tool to our industrial partners. But this type of prototype was sometimes mistaken for a production-ready product that could be released “as is” to their clients.

### **Partners and End Users**

Working with our industry partners, we learned it is critical to establish expectations around research versus development. Doing research with small start-up companies is challenging because they have limited resources and an urgency to release their products to the market. It was difficult to explain to our partners that a research project requires exploration of multiple solutions and proper evaluation. Once we provided our partners with one visualization solution, they became excited

about it and consistently truncated the exploration and assessment of other visualization strategies or research problems. For example, while working on the implementation of a scatterplot to display time series data, we discovered that the sample data was missing values for long periods of time and was too dense in other periods. This made our scatterplot unattractive and potentially misleading. Visualizing unevenly spaced data has been scarcely investigated in the literature, making this an excellent research opportunity. However, our partner preferred that we finish the dashboard design with the current solution.

Our industry partners were also reluctant to include target users in the design process. We interacted closely with internal stakeholders who were experts in the problem domain and clearly understood their users' needs. However, contrary to established design practices and despite our best efforts to convince them otherwise, our start-up company partners were reluctant to let us interact directly with end users. Customer acquisition had only just occurred or was still ongoing, and the partners were concerned that showing customers a low fidelity mockup would lower their credibility and impose on the customers' time.

However, as one would expect, failing to interact directly with end users during requirements gathering led to problems such as missing important requirements and designing noncritical components. For example, in one case we designed several detailed views, but we later learned a well-designed overview was sufficient.

Our partners were much more willing to have us interact with their customers once we had established our credibility by producing an interactive prototype (that could also be a selling point for the company). At this point, we began involving users in usability studies and adjusted the designs as needed. Although effective, this process was inefficient. In retrospect, perhaps we might have done more to convince our partners of the importance of early and consistent end user involvement.

**T**he visualization community has developed a wealth of design knowledge, toolkits, and techniques for building interactive visualization tools. But for practitioners of visual analytics, it is important to know not only the latest visualization techniques but also the challenges that may be encountered during the process of designing and implementing a visual analytics tool.

Here, we have summarized the main chal-

lenges we found from our experience with four small partner organizations. We hope that our report of these experiences will help visual analytics practitioners save time and effort in the design and development of new tools. We also hope that, by publishing our experiences, we have helped make other practitioners and researchers aware of these important issues before they undertake their next project so they can consider possible solutions in advance. ❏

---

### Acknowledgments

*We thank our industrial partners, Maxwell Terpstra, and the VisID group from the University of Victoria for their thoughtful feedback and suggestions. This work was funded by MEOPAR, MITACS, and NSERC Engage grants.*

---

### References

1. T. Munzner, *Visualization Analysis and Design*, CRC Press, 2014.
2. S. Few, *Information Dashboard Design*, O'Reilly, 2006.
3. M. Bostock, V. Ogievetsky, and J. Heer, "D<sup>3</sup> Data-Driven Documents," *IEEE Trans. Visualization and Computer Graphics*, vol. 17, no. 12, 2011, pp. 2301-2309.
4. R. Kincaid and H. Lam. "Line Graph Explorer: Scalable Display of line Graphs Using Focus+Context," *Proc. ACM Working Conf. Advanced Visual Interfaces*, 2006, pp. 404-411.
5. L. Grammel, M. Tory, and M.A. Storey, "How Information Visualization Novices Construct Visualizations," *IEEE Trans. Visualization and Computer Graphics*, vol. 16, no. 6, 2010, pp. 943-952.
6. J. Mackinlay, "Automating the Design of Graphical Presentations of Relational Information," *ACM Tran. Graphics*, vol. 5, no. 2, 1986, pp. 110-141.
7. D. Lloyd, J. Dykes, and R. Radburn, "Human-Centered Approaches in Geovisualization Design: Investigating Multiple Methods through a Long-Term Case Study," *IEEE Trans. Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2498-2507.

**Maria-Elena Froese** is an information architect at the University of Victoria, Canada. Contact her at [maleh2@uvic.ca](mailto:maleh2@uvic.ca)

**Melanie Tory** is a research scientist at Tableau Research. Contact her at [mtory@tableau.com](mailto:mtory@tableau.com).

Contact department editor Mike Potel at [potel@wildcrest.com](mailto:potel@wildcrest.com).