

Exploration Views: understanding dashboard creation & customization for visualization novices

Micheline Elias^{1,2} and Anastasia Bezerianos¹

¹ MAS Laboratory, Ecole Centrale Paris, France

² SAP Research, Business Intelligence Practice, France

{micheline.elias | anastasia.bezerianos }@ecp.fr

Abstract. With the increase of visualization platforms targeting novices, researchers are now focusing on gathering insights regarding novice user practices. We describe the design and evaluation of Exploration Views (EV), a system that allows novice visualization users to easily build and customize Business Intelligence information dashboards. EV provides an intuitive environment for dynamically creating, rearranging, searching and exploring multiple visual data representations from diverse data-sources. These aspects aid users to better retrieve, experiment and familiarize themselves with their data. We evaluated EV with both novice and expert dashboard designers and report here (i) how novice users interact with the system, (ii) differences in how novice and expert users react to a dashboard systems that targets both, and (iii) provide new design guidelines for practitioners building dashboard applications, on the needs of novice visualization users.

Keywords: synchronized views, interface customization, novice users, visual queries, business intelligence dashboards.

1 Introduction

Business Intelligence (BI) deals with the collection of processes and software that supports organizations in understanding large datasets, retrieving and analyzing information and making decisions. The value of visual presentation of data was identified early on in this field and usually takes the form of dashboards, interfaces resembling an automobile's dashboard that organize and present information in an easy to read manner [27]. BI visualization dashboards provide collections of multiple visual components, such as charts, on a single view so that information can be monitored at a glance [6]. Appropriate visual representations in dashboards, using colors, size and shape, are combined with interactive exploration [9] to amplify human cognition and enhance information understanding [1].

The current life-cycle of a BI dashboard involves multiple actors [3], including end-users and business analysts. In user-centered BI [11] end-users intervene and provide feedback to the business analysts that create customized dashboards to meet user needs. This feedback comes at different stages of the dashboard design and setup,

and involves a large amount of communication between business analysts and end-users, in order to define functional specifications and a positive user experience. Thus there is an intrinsic delay introduced in any end-user requirement change or customization request. Moreover, current market trends (seen in the Gartner Survey [24]), recognize dashboards and data analysis are one of the driving forces adding value to companies. This indicates a tendency for an expanding end-user population with diverse needs, requiring quick access to customizable dashboard technology, bypassing dashboard designers. This is reflected in advances in easily customizable dashboard visualization systems (e.g. [28,29]), it coincides with advances in easily accessible information visualization environments for novices (such as ManyEyes [31]), and is also reflected in the increasing focus of the InfoVis research community on how to support novice users [12,15]. Nevertheless, there are still many questions about how novice users interact with actual information visualization systems.

To better understand how novice users interact and build visualizations on their own (without dashboard designers or other human mediators) we created Exploration Views (EV). EV is a dashboard prototype (Fig 1) build specifically to empower end-users, following guidelines from previous work on visualization novices. To evaluate its effectiveness, we compared how dashboard experts and novices reacted towards EV when creating visualizations for realistic tasks. Finally, through our observations we derived a set of additional design guidelines for dashboards targeting novice users.

Our contribution consists of: (i) putting together all guidelines for novice users (empirical or tested) in a fully functional system that can be tested in practice, (ii) creating a system that also includes advanced functionality required by expert users, and (iii) testing the system under realistic tasks to gain insights.



Fig 1. The Exploration Views consists of: (a) the main Dashboard and its components. (b) The Clipboard drawer widget that includes templates for charts (top) and data widget samples (bottom) that the user has created as a pallet. The drawer can be closed when not needed. (c) A visual representation of the global data filters that are active at any point. (d) And a set of other functions, such as search, save, share, etc.

2 Related Work

In this section we will present related research in the fields of dashboard creation and visualization construction, especially in the context of novice visualization users.

Dashboards. The majority of the commonly used commercial BI dashboards (e.g. [8,23,26]) assume that from the end-user perspective the dashboard design will remain consistent, and focus more on easy report generation. To create new or customize existing dashboards, end users need to either contact trained dashboard designers, or install separately provided dashboard design components and spend time to train themselves. There are some notable exceptions. Tableau [29] incorporates the design functionality in the dashboard. Spotfire [28] also accommodates novice users with a flexible drag-and-drop interface for customizing visualizations on the fly. Finally, the prototype Choosel environment [4] allows for easy visualization creation and further exploration of the visualized data, but does not yet provide advanced functionality (e.g. hierarchical data representation as in OLAP data models [22], with drill-up/down functions), required in real life business data analysis.

Visualization creation. A variety of visualization toolkits allow the creation of sophisticated visualization charts and combinations of them (e.g. Flare¹, Silverlight², JavaScript InfoVis tk³, ivtk⁴). These do not target novice visualization users, and usually require additional programming to create the visualizations, to synchronize them in dashboard views, and to connect them to data-sources. Several environments, such as Improvise [32] or Snap-Together [21], allow the construction of highly customized, powerful and fully linked visualizations (for a comprehensive list see [25]). Nevertheless their setup cost and learning curve is usually high and they target visualization experts rather than novice users.

Visualization for novices. On the other hand, there is a set of visualization platforms that specifically target novice users. For instance, the Google analytics platform [6] provides a highly customizable dashboard with good default visualization charts for web-related data. Other web-based visualization platforms easily accessible to novice visualization users, like ManyEyes [31], Sense.us [16], or Polstar [5], usually restrict users to a viewing and sharing a single visible chart at a time. Thus creating a dashboard is a laborious process performed outside the platform, and linking of multiple charts is no longer possible.

Novices in visualization. With the increase of visualization platforms targeting novices, researchers have started gathering insights into their practices and behaviors. Heer et al. [15] group different visualization systems by how skilled their users have to be, and note that tools supporting detailed data exploration target mostly experts, whereas it is mostly communication/reporting systems that accommodate novices. They also suggest that visualization tools for novices should: allow user-friendly data input for common data formats; automatically provide visualizations or reasonable defaults based on data types; and use contextual information to clarify the displayed data and encoding. Grammel et al. [12] investigate how novices create single charts,

¹ <http://flare.prefuse.org/>

² <http://www.silverlight.net/>

³ <http://thejit.org/>

⁴ <http://ivtk.sourceforge.net>

using a human mediator to interact with the creation interface. They found that the main activities performed by novices (apart from communicating their specifications to the mediator) are data attribute selections, visual template selections, and viewing and refinements of the above. They observed that subjects often had partial specifications for their desired visualization, and faced three major barriers: selecting correct data attributes, choosing appropriate visual mappings, and interpreting the visual results. They also verified suggestions from [15], and provided additional guidelines, such as facilitating searching of attributes, automatically creating visualizations when possible, providing explanations, and promoting learning.

Our work extends the above in the following ways: Using guidelines from [12,15], we build a fully functional dashboard visualization system, and observe how novice users use the interface to create dashboards (without a human mediator [12]). We compare our observations of novices to reactions from experts, and derive additional design guidelines for visualization dashboards that target both user groups.

3 Exploration Views (EV) Prototype

In this section we describe the design rationale behind Exploration Views (EV) and the functionality that supports novice users, based on guidelines from previous work. We will first explain the main goals of EV: easy creation, customization and interaction with multiple visual representations in a unified dashboard environment.

Although dashboard users may have clear questions regarding their data, they can be novices when it comes to visualization, and naive as to how dashboard software can support their needs. An *easy creation* process is essential for novices to experiment with dashboard designs fast. As seen in [12], visualization novices often have partial mental specifications for their visualization needs and tend to refine and change their designs. To ensure a user-friendly dashboard creation, the sequence of steps needs to be simple, with continuous visual feedback on the final outcome of user choices. Moreover, as novice users may have no previous knowledge of what visual templates and representations are possible for different data types, EV must provide chart suggestions and templates to choose from (also suggested by [12,15]). Finally, EV should support common data formats (recommendation in [15]), and users should not be restricted by data storage architecture, but be able to use data from multiple data-sources (e.g. personal data in excel and OLAP data from their enterprise).

As novice users create dashboards, they may need to explore and try out alternative visual templates and representations to learn what meets their needs. Requirements and tasks can also change over time, requiring different data representations and data filtering mechanisms, or new data sources. EV should support iterative visualization specifications (suggested in [12]) by being *easily customizable and adaptable*.

It is important for novices to see the effect of their choices immediately, promoting exploration and experimentation with visualizations. Thus in EV customizations *become active immediately*, ensuring that customization is tightly coupled with the *visual analytics process* (recommended by [12]). EV further supports the visual analytics process with functions such as saving and sharing, text search mechanisms, fully linked visualizations and visual queries, and other data exploration mechanisms.

Finally, if the design of EV supports both low-cost experimentation with visual templates and visual analytics processes, it can help novice users to become more accustomed to visualization creation and analysis, *promoting learning*.




Fig 2. (a) Creation dialogs with visual examples of templates or saved dashboards. (b) Chart creation dialog with recommendations of charts appropriate to data categories and analysis types, sorted from simpler to more complex. (c) Data Category and measures selection.

3.1 Easy creation

When creating a new dashboard the system provides miniature icons of saved dashboards and several template dashboards to choose from (**Fig 2.a**). The visual representation of possible dashboards helps users recall designs (as opposed to remembering dashboards solely based on names [20]), and their titles are descriptive so novices can start from a reasonable existing dashboard design, minimizing new content creation. These dashboard templates can be also shared with others.

If none of the existing dashboard templates meets the user needs, she can choose to create a dashboard from scratch. The user is guided by the system by first asking her to select data sources of various format types that are supported. When the data sources have been selected, the users can populate their dashboards with charts.

The dashboard work area (**Fig 1.a**), acts as both the design space for the dashboard, as well as the active dashboard when content is added. Here users can add charts, tables, and data filters. Possible components to add can be found in the Clipboard drawer (**Fig 1.b**), which can be closed when creation (or customization) ends. Components from the Clipboard can be dragged-and-dropped in the work area, and EV automatically realigns existing components to accommodate new additions.

When adding charts (i.e. visual templates), users need to select data attributes (e.g. sales, year of sale, etc), and match them to visual mapping (a visual property in the template). This is a challenging task for novice users [12] and EV aids them (as in [9,35]) by providing reasonable mapping and visual template recommendations (**Fig 2.b**). After selecting one or more data attributes in the form of measures (data attribute, e.g. sales prices **Fig 2.c**) and categories (possible groupings of the attribute, e.g. per state, year), the system presents appropriate analysis types for this data. For example for a dataset of sales per state over the period of 2 years, possible analysis types include comparisons (sales by state or year), contributions (the percentage of sales per state over all sales), or trends (the evolution of sales over time). If the user selects one of these analysis types, the system automatically recommends visual templates (charts) suitable to the specific analysis and to the nature of the attributes

(e.g. how many measures and categories are selected, if they are categorical, ordinal, etc), using a default visual mapping. These recommendations are ranked from simple and more commonly used charts to more complex ones, and aid novices to explore reasonable alternatives. Available visual templates include charts, gauges and tables.

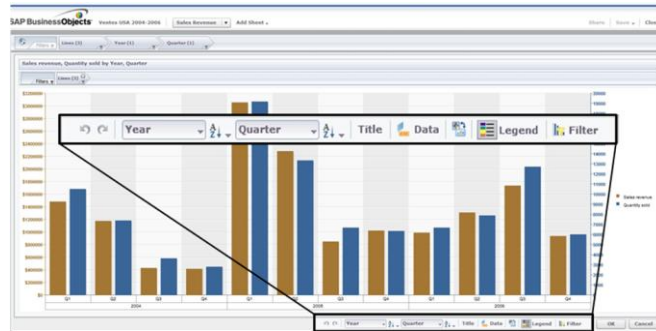


Fig 3. The “Explore” mode is a detailed view of a chart, where users can undo/redo actions, choose data categories, sort items, change data sources, switch chart templates, and create local data filters. To avoid switching to the “explore” mode, part of this functionality is also available as a drawer toolbar for each chart in the main dashboard, with added options to lock the chart.

3.2 Ease of customization

During the life of a dashboard, user needs may change; new hypothesis may need to be explored and new data sources to be added. Especially given novice users’ tendency for iterative visualization specifications [12], EV has several features for user friendly customization, such as actions to add/remove components, change their properties, duplicate their content, and explore sub-parts of the data using data filters.

Users can add new components on the dashboard or drag-and-drop existing ones from the Clipboard. If the components are visualization templates, the user has to go through data, attribute and mapping selection as described before. If the components are saved in the Clipboard, we use the properties of the component as it was stored.

Existing components on the dashboard can be refined and customized by clicking the “explore” icon of the component. This provides a detailed view of the component (**Fig 3**), together with choices for datasets, visual templates, attributes and mappings (as described in the **Sec 3.2**). Changes are immediately reflected on the component in an automatic preview, so that users can experiment with different choices [12].

Users can also customize the dashboard layout itself through functionality for hiding/showing dashboard components, dragging to rearrange them, resizing them, etc. EV intelligently resizes components to fit in the dashboard (e.g. **Fig 4**).

At any point users can save content for later use (in this or other dashboards), by dragging them in the Clipboard Drawer. This allows exploration of alternatives without risking losing existing components. To further ensure that novice users are not afraid to experiment with their dashboards, all actions are reversible through undo history functionality applied on the entire dashboard and locally on specific charts.




Fig 4. (a) Users select subset of their data, that get highlighted in all charts. (b) They drag their selections outside the chart to create a cut-out clone visualization (c), of the same type (here bar chart) with the selected data only. A blue highlighted bar (b) indicates the new chart will be placed on the right of the original, with existing charts shift to accommodate it (c).

3.3 Supporting Visual Analytics process

Apart from the creation and customization functionality, EV provides support a series of functions that aid the visual analytics process. EV has text search support that follows the guidelines of dynamic queries [33]: as the user searches text the corresponding text components on the dashboard that match it (axis labels, categories, tags and any other textual information) get dynamically highlighted.

All our visual representations are essentially linked coordinated views [24,32] of the underlying datasets (unless users give them local filters, discussed later). We thus extend the dynamic query functionality to any visual selection or query in the dashboard: if the user selects in one chart bars representing the 3 quarters of a year, these data instances are also highlighted on all other charts (**Fig 4.a**).

EV further encourages users to explore their data by providing global and local data filters, i.e. constraints on the data entries displayed. Filters can affect the entire dashboard (global), or specific components (local). Thus users can explore “what-if” questions without changing completely their dashboard. Our filters focus on data dimensions: enterprise data usually follow the multidimensional OLAP model [22], where data is organized in hierarchical categorical or ordinal dimensions (e.g. sales grouped per country and state, per year or month, etc). Each dimension has a set of values, which we call categories. Users can drag-and-drop Data-Filter components from the templates (**Fig 4.a,b**), that can be used to restrict data categories.

To enable users to dynamically explore filtered data and compare categories, EV provides faceted navigation components [13,15]. Data-Filter-Controller components are visual components, which get attached to a data dimension, similarly to filters. But instead of constraining the dimension categories, they provide a visual list of the categories (**Fig 5.d**) that users can toggle on/off. Users can easily create and dismiss controllers to ensure that only the desired ones are visible at any given time, reducing clutter in the faceted navigation space (a challenge of faceted navigation [14]).

During data analysis and exploration, users may want to clone and explore in parallel interesting data. EV allows visualization cloning by using the Clipboard (full cloning) or by dragging and cloning parts of the data. For example in the bar chart in **Fig 4**, the user wants to explore in detail sales for the 3 quarters of 2005. By selecting

and dragging only these bars, she creates a cut-out clone visualization of just the selected data, which keeps the initial component properties (dataset, visual template, categories, mappings). She can thus explore data subsets without creating components from scratch. This cloning process is essentially a fast local filter creation.

If the underlying data-sources follow the classic OLAP data organization, users have access to the traditional OLAP functionality [22], such as sorting, drill up/down, pivot, roll-up, etc. Users can also flip the chart axis to experiment with different chart layouts. This exploration functionality is necessary for expert users. It is present in the “explore” chart mode, but also in the main dashboard in the form of a widget drawer for each chart to avoid opening the “explore” view (e.g. in [29]). At any point users can undo/redo their actions on each chart. Finally EV provides functions to support long term and collaborative visual analysis, such as save and share.

To help users follow visual changes during data exploration, when using OLAP functions or filters, changes are smoothly animated to ensure visual momentum [34].




Fig 5. Filters and Faceted navigation creation. (a) The user can drag-and-drop a new global filter and choose the data categories that will be visible on the dashboard. The filter is either directly applied, or can become a filter control component for faceted navigation (b). All active global filters are visible at the top of the dashboard (c), and local filters on the top of charts (d).

4 Exploration Views Technical Platform

EV is not just a lab prototype, but a fully operational BI dashboard that can handle large amounts of data from different sources. Although this paper focuses on the interaction aspect of EV, we briefly describe its architecture. EV has 3 layers. The first is a web client based on the Adobe Flex that contains the dashboard layout, the exploration and faceted navigation components, and the chart manager. The chart manager CVOM (Common Visualization Object Model) provides a set of services for chart creation and recommendation based on the analysis users want to perform and data measures. The second layer is a server containing a user session manager (as in [5]), a dashboard state management, and a component that synchronizes different data sources to enable the multi-chart visualizations. The third layer consists of an in-memory database for quick data storage and retrieval [30], and a semantic layer that helps us to connect visual queries from multiple charts and have synchronized views.

5 Study Design

In this section, we define the purpose of our study, its design and methodology.

Our goal is to examine how novice visualization users create dashboards. We attempt to answer three questions: (i) is it easy for novices to discover visualization and dashboard functionality in an interface specifically designed using previous guidelines for novices; (ii) are the existing guidelines effective and what is missing from the literature, especially in the context of dashboard visualizations; and (iii) do novice and expert visualization users react similarly to such an interface.

Q1. Can novice users discover functionality and build dashboards? We created EV following guidelines by [12,15], and asked users to construct visualizations to answer questions, without the help of a human mediator (contrary to [12]). We wanted to see if novices can discover functionality, learn the interface and answer questions based on the created visualizations. Thus we provided no interface training and observed difficulties in interface discovery and visualization creation process.

Q2. Are the existing guidelines enough? Our users went through the entire process of dashboard design and customization, including data and visual template selection and attribute mapping, to answer multiple questions. We observed this process to verify existing guidelines and collect potentially new ones specifically for designing dashboard systems for novice users.

Q3. What are the differences between novice and expert users? Finally, we wanted to identify differences between the interaction needs and expectations of novices and BI visualization experts. Experts were knowledgeable in understanding which visualization templates and attributes mappings can answer specific questions, but not necessarily in how to interactively construct or customize a dashboard or visualization. In this respect all our users are novices in creation and customization of dashboards, but not in visualization choice and interpretation.

5.1 Participants and Pilot studies

Overall 15 participants (4 women), ages between 24 and 43 years old (median of 29) took part in our study: 8 were novice visualization and dashboard design users, and 7 were BI dashboard experts. Visualization and BI dashboard novices had no previous knowledge of visualization creation and were only familiar with charts (e.g. pie-charts) seen in everyday life in newspapers or presentations. They reported being exposed to charts on a monthly basis. When asked, only 2 novice users reported that they knew what a dashboard was, and one of them had used the Google analytics dashboard. All novice users were familiar with the concept of data filtering, but at least half did not understand concepts such as drill down-up. Novice users were recruited from a university community, with a wide professional expertise (IT, medicine, and pharmacology). Experts were product users of a BI dashboard provider with different areas of expertise (finance, human resources, analysts, etc). Experts reported using charts and dashboards on a daily basis and were familiar with multiple types of charts, data analysis types (such as correlation, trends, comparisons, etc), and data exploration concepts (filtering, drill-down/up, pivot, roll-up, sorting, etc).

We conducted a pilot study with two novice participants and found that both had trouble understanding the difference between a measure and a category, as well as some data exploration concepts such as data filtering and drill-down/up. We thus decided to include a graphical explanation of these aspects as training in the actual study. We also observed some interface problems: some icons were not clear enough, and were changed to more easily understandable icons and text descriptions.

5.2 Procedure and Apparatus

Participants performed the study in a usability lab using a 15.6" HD screen laptop. Sessions was audio/video recorded and screen captured. One observer was always present to give instructions, explain tasks and observe participants.

Each session lasted between 1.5 to 2 hours. It started with a survey on participant's background experience in computer usage, chart and dashboards, and specific data exploration functionalities such as filtering, drill down/up, etc. As one of our goals is to see how easy it is for novice users to discover functionality and create dashboards with EV, we provided graphical explanations and examples of concepts such as measures and categories, data filtering and drill-down/up, but no information about how to perform them in EV. Next, participants were introduced to a hypothetical situation and dataset (see [Sec 5.5](#)). We followed a think-aloud protocol, requesting users to vocalize their thoughts and actions while conducting their tasks. This phase lasted between 35-55 minutes. After finishing all the tasks in our scenario, users were asked to complete a Likert type questionnaire to elicit their opinion on specific aspects of EV, and were interviewed to clarify observations made by the experimenter, identify problems users faced, and suggest areas for improvement.

5.3 Tasks and dataset

Our dataset is an e-fashion dataset (as in [17]), containing 3860 records with 9 measures (e.g. sales revenue, quantity sold) and hierarchical categories (e.g. state/city/store, year/quarter/month, line/category). The dataset was new to both novice and experts so the impact of the dataset and domain on task performance is limited. The dataset is also complex enough to allow complicated analysis questions.

Our tasks were presented within a larger scenario. Participants were told they work in the financial department of a global chain of garment sales, monitoring financial aspects using a dashboard tool. We encouraged participants to think aloud, and if they failed to do so we prompted them to vocalize their insights and expectations.

Based on our pilot, we provided participants with a sheet explaining graphically of the nature of measures and categories, as well as filtering and drill down/up. We also give them a brief explanation of the dataset (what measures and categories exist), and told them to use the dashboard to create and rearrange visualizations.

Five tasks were performed, ordered from simple to complex ([Table 1](#)). Each task encouraged users to use specific operations and data exploration functionalities such as filtering. Tasks began by loading an existing dashboard and participants were given time to familiarize themselves with the dataset and dashboard.

At the end of the session participants took part in an interview and filled out a questionnaire. The goal of the interview was to clarify observations made during the study, prompt users to identify difficulties/barriers they faced, and suggest improvements. The goal of questionnaire was to rate EV on a 7 point Likert scale on aspects such as its utility, ease of use, overall satisfaction of use, etc.

#	Task script	Purpose of Task
T1	You are working on sales data for your company with a new dashboard tool (EV). Find and load the dashboard “Sales USA 2004-2006” and explore the new tool.	Create a new EV dashboard and explore its contents.
T2	Your manager calls to request an overview of <i>Sales Revenue</i> and <i>Quantity Sold</i> in the <i>Year: 2006</i> in the <i>States: California and Texas</i>	Select visual templates, measures and categories, and data filtering.
T3	He later requests more information on the <i>States</i> of <i>California</i> and <i>Texas</i> in <i>Year 2006</i> : <ol style="list-style-type: none"> Asks for a chart of the top 3 stores generating most <i>Sales Revenue</i>. Requests to store a copy of this chart. Asks for a comparison between <i>Sales Revenue</i> and <i>Quantities Sold</i> for <i>Items accessories and dresses</i>. Asks for an alternative visualization of the chart stored in T3.3 to choose from. 	Customize charts, using options such as visual template changing, sorting, chart saving and re-use. Conduct different types of common data analysis (e.g. comparisons, extreme values)
T4	Your manager contacts you again to request the <i>Year</i> that your company sold most (<i>Quantities sold</i>) items from the product <i>Line: "Leather"</i> .	Conduct different types of common data analysis (e.g. comparisons, extreme values)
T5	Your manager needs a view showing the U.S. <i>Sales</i> for <i>Years 2004-2006</i> based on a dashboard image he likes. He asks you to create a new dashboard looking like this one (Fig 6) with the data from your dataset.	Design of a complete EV dashboard that incorporates global/local data filters, and conduct layout management

Table 1. Tasks users performed during each session of the study, and their purpose.




Fig 6. E-fashion dashboard that participants were asked to replicate.

5.4 Data Analysis

We used four different data capturing and analysis methodologies: functionality checklists, observations, questionnaires and open-ended interviews. Observations made during the study were verified, and expanded upon, in a second pass through the recorded material (video, audio and screen captures).

To answer **Q1** on **discoverability** and **usability**, we kept a checklist of functions in the interface and observed how easy it was for participants to discover and use them (results in **Sec 6.1**). Questionnaire answers also contributed towards answering **Q1**.

To answer **Q2** on **guidelines for building dashboard systems** targeting novice users, we noted success events and barriers faced during dashboard creation and customization. These observations are augmented by user comments in the open-ended interview session. Findings are summarized in **Sec 6.2**, and the resulting new guidelines are presented in **Sec 7**. To answer **Q3** on **reactions from novices and experts**, we report explicitly only the differences between the two

6 Findings

We present our main findings on ease of use and discoverability of the EV interface (observed and user self-reporting, **Sec 6.1**), as well as broader findings (based on observations), applicable to dashboard design for novice users in general (**Sec 6.2**). We then provide design guidelines for creating dashboards for novices (**Sec 7**).

6.1 Discoverability and ease of use of interface

All 8 novice (and 7 expert) participants were able to successfully perform all tasks.

To examine if the functionality of EV was easy to use and discover, an experimenter observed the study session and noted difficulties encountered by novice users (8 users). Almost all users easily discovered and performed actions to create new charts (7/8), take snapshots (7/8), manipulate measures, categories and visual templates (7/8), and exploration actions like sorting and drill-down/up (7/8). All novice users applied local filters easily (8/8), although 2 had initially added global filters, as they expected both global and local filters to be added in the same way.

Functionality that novice users had trouble discovering, was chart cloning (4/8), locking (8/8, see section on **Locking** below), and activating chart explore mode (2/8 see section on **Terminology**). Users who did not discover this functionality on their own, were able to complete actions after a small hint from the experimenter.

We examined if users reacted to the interface and functionality differently, by analyzing a 7-point Likert scale questionnaire. Analysis was conducted using a Wilcoxon rank sum test with continuity correction as our samples (Experts, Novices) had different sizes. Overall (**Table 3**), novices ranked EV higher, which can be attributed to the novelty of using a dashboard. Responses were statistically different for: functionality, consistency, and overall satisfaction (all $p < .05$). Experts would have been more satisfied with extra functionality (mentioned next), while both groups had some trouble with representation consistency of global vs. local filters (**Sec 6.2**).

In the follow-up interview both groups of users reported EV was easy to use and customize. They appreciated that customization and data exploration actions became active immediately for quick feedback and verification of their choices. They stated that EV is intuitive and doesn't require training. Finally, mostly our expert users requested additional functionality, like extra customization (e.g. changing colors) and annotation features, and the ability to export dashboards in a printable version.

	Utility	Functionality	Ease of Use	Consistency	Satisfaction
Novice	5 (std. 1.12)	6* (std. 1.12)	5.5 (std. 1.58)	5.7* (std. 0.83)	6.5* (std. 0.74)
Expert	5 (std. 1.63)	4* (std. 0.75)	5 (std. 0.53)	6.1* (std. 0.95)	5* (std. 0.69)

Table 3. Average ratings of EV for different attributes, by Novice and Expert users (* sig).

6.2 General Observations on dashboards for novices

6.2.1 Chart creation and customization

Chart recommendations. As suggested in previous work on visualization novices, we provided a chart recommendation feature. All participants commented on the value of EV in helping them create appropriate charts fast. Although all users chose charts from the recommendation, only novices experimented with different alternatives, in order to “see what this new visual looks like”, “look at a chart I’ve never seen before”. We’ve had at least 2 users discovering and learning in this way a new visualization that they used often until the end of the study. Nevertheless, in most cases, novice and expert users ended up using familiar visualizations (bar and pie charts) for answering questions (“it looks complex, let’s stick with what we have”).

Base views and experimentation. We often found that users (novice and experts) created basic charts with measures they thought important (e.g. one for sales over time, one for number of items sold per state, etc). These charts (*base views*), usually 2-3 in number, were used as a starting point for customizing and refining to answer different questions. When customizing a chart, we observed that novices experimented frequently with filters, drill-up/down operations and measures, to “play with things and understand the different options”. On the other hand experts attempted to create the appropriate visualization without much exploration.

Chart per question. We observed that both user groups created a single visualization per task question. When asked if they see a benefit in using a dashboard, since one chart was the focus of the analysis at a time, they all mentioned that it “helps keep the context of all data” (benefit identified in BI dashboard literature [9]).

Use of categories. Novice participants often did not recognize the effect of adding hierarchical categories on a chart. In EV when a refined category is added in a chart, it affects aspects such as sorting, aggregation (averages), etc. For example a user had created a chart with sales sorted by year. She then added the category quarter, which is a drill-down category for year. At this point sorting was applied by quarter, but the user still expected sorting to be done by year. Thus automatically calculating metrics based on the finer categories of a hierarchy was not understood by novice users.

Data filtering. As we mentioned EV allows both global dashboard data filtering, as well as local filtering on a single chart. All users found this functionality useful, nevertheless it was often hard for novice users to make correct use of it.

There are 2 creation options for global filters: select part of the data on a chart and then apply the option “filter” which will filter data based on the selected categories and values; or select the option “filter” and then go over the process of selecting on which categories and values to filter. We found that novices generally preferred the second option because they want to make sure they “selected all correct values”. Contrary, experts created global filters with selections that they then refined further.

Another aspect where expert and novice users differed was the presentation of possible filter categories. Experts wanted categories to be presented grouped

thematically, e.g. grouping of time (year/quarter), or geographical categories (state/city), while novices preferred an alphabetic presentation.

Users found the existence of both global and local filters useful and relevant when analyzing data as it allows advanced analysis, for example comparison of different products sold by stores for several periods. All participants commented that presenting global and local filters consistently is important, but that local filters should be further highlighted in charts. The EV interface gave such detailed filter views, and users commented on the bulkiness of the filter panels. There is thus a tradeoff between presenting filter details and space. Finally, we found that sometimes users chose local filters that were in conflict with global ones. This caused problems for several novice users who could not understand why no data was displayed.

6.2.2 Interaction

Linking & coordinated views. Most participants (experts and novices) expect a link between charts. Thus when selecting a portion of one chart to highlight, they expect other charts with the selected data in EV to update accordingly. Nevertheless, the expectation of what “update” means differs across users. Most expert users expected selections on one chart to be reflected as selections in all the others, in the classic brushing and linking approach in coordinated views (default behavior in EV, Fig.4). Nevertheless, most novice users expected selections to act as explicit filtering: when partial data was selected in a chart, they expected all charts to update such that only the selected data is reflected throughout the dashboard. This interpretation would result in changes to measures such as averages, sorting, etc for different charts.

Undo as exploration strategy. All participants were familiar with undo/redo, but we found that it was novice participants that made extensive use of undo per chart. Undo was not used as a correction mechanism when mistakes were made, but rather as an exploration aid. Novice users reported that if they wanted to try something and see if their predictions were correct (e.g. drill-down), they could use undo to quickly dismiss their tests (instead of reversing their actions, e.g. drilling-up). They also used undo as a way of returning to *base views* (charts used as starting points for further exploration): from base charts they would drill-down/up or filter briefly to answer a question, and then use undo to return to the original chart.

Clipboard. This feature allows users to create a palette of active charts that they can later reuse. All participants found it useful for storing, but they reported it is most useful for sharing chart snapshots with others. This remark, reinforced by requests for ways to email or export charts to presentations, indicates users also saw the clipboard as a storytelling creation platform for sharing their findings with others. Participants from both groups also requested features such as annotation, putting titles, adding comments, and assigning the current date to a chart. Experts in particular requested stored charts to also clearly and visibly describe context (e.g. dataset, measures, filters, etc). These requests indicate that participants want to use the clipboard to bookmark their exploration process: keep track of interesting charts they explored, comment them, and store the date they conducted their exploration.

Some novice users mentioned they were unsure exactly what is captured when a selection is active on a chart (i.e. does the capture only include the selected data or not). This is an interpretation problem similar to that seen in coordinated views.

Visualization locking. In our prototype we provided functionality for visualization locking (not allowing editing or global filters). Expert participants expected this locking to apply to filters, but also visual properties like size and position in the dashboard. Some novices felt that locking a visualization “creates confusion because it takes out this particular chart from the context of the dashboard”. They requested very clear visual representation to indicate locked charts, such as depressed charts.

Terminology. Overall users were satisfied with the terminology and labels used in EV. Some novice users suggested more precise terminology even for some features that are familiar from other UIs. For example our original icon for editing a chart was confusing and novice users were not sure if they were going to edit the title, the chart or the entire dashboard. Thus they often requested tool-tips detailing the function of icons and their context. Finally, while experts were accustomed to English terms (in terms of functionality, dataset measures and categories), this was not true for novices with English as a second language who requested support in their mother tongue.

6.2.3 Design tradeoffs

The design of EV balances several tradeoffs to support both expert and novice users. Functionality required by experts, such as chart locking, axis flipping, and extended chart recommendations, often confuses. In our work, we attempted to always provide the simpler and most used functionality in a more prominent way. For example options like sorting, filtering, and visualization changing are first in the chart widget drawer, while axis flipping or locking come last. In the chart recommendation list, most common charts always appear first. Moreover, when possible we use reasonable defaults, for example in assigning color and other visual variables to categories. Nevertheless, we often found that functionality required by expert users still confused novices occasionally (e.g. the locking option).

7 Discussion and design guidelines

All novice participants in our study were able to successfully perform the required tasks in EV, a system built using guidelines for novice user interaction. Based on our observations we have additional guidelines when designing visualization systems for novice users, especially in the context of dashboards.

Basic training on concepts such as categories: Our novice users were successful in creating visualizations and solving their tasks. But, based on our pilots, they were given small tutorials on the differences between measures, categories, and the nature of basic OLAP and data exploration functionality (drill-down/up, filtering, etc). Thus dashboard interfaces for novices should provide small tutorials of these concepts. This training is particularly important for hierarchical categories. Novice users need to know if metrics (averages, sorted lists, etc) are recalculated in charts when adding finer hierarchical categories. Based on user comments we recommend either not recalculating metrics unless users request it, or recalculating per category on demand.

Considerations in coordinated views: When creating a dashboard, novice users assume that charts are linked by default, and expect brushing and linking

functionality. Its implementation for single data sources where charts share measures and categories is fairly easy. Nevertheless, if measures and categories need to be matched across data sources, then matching approaches need to be provided [21,32]. When providing coordinated views, designers should remember that novice users confuse linked selections of visual queries to applying global data filters based on the selections. To help disambiguation, they need clear visual differentiation between selection for linking and for data filtering. A linking selection should highlight selections without affecting the visual representation of unselected data; a filtering selection should also alter the visual representation of the unselected data (e.g. fade them out) to indicate that all data are being altered due to the selection.

Extensive exploration aids: Novice users created a small number of *base view* charts, most often of familiar visualization templates, that they later copied and refined to answer questions. When possible, such basic views should be provided. Novices also made extensive use of chart undo/redo as an exploration aid, thus this functionality should be provided on a per chart basis. Furthermore, clipboards were seen as good ways to share data for storytelling and communication, as well as bookmarks of the exploration process. For clipboard snapshots to be useful this way, they require functionality such as annotations and labeling, as well as ways to report their metadata, such as creation date and details of the context of the chart (i.e. dataset, active filters, etc). Finally, although locking charts can be useful during long analysis (as indicated by experts), this locking should be very clearly visible to not confuse novices (e.g. fading the entire locked chart, or making it appear as depressed).

Data filtering considerations: Data filters, although very useful for data exploration, can be challenging for novice users. We found that novices preferred to create global filters from scratch, select categories, values, etc. While experts preferred to create them graphically (through visual selections) and then refine them. Both options should be available. When selecting categories to filter, we recommend presenting them in thematic groups (expert choice) that promotes dataset learning. But this organization should be supplemented by text search functionality. Novice participants want filter details to always be visible (both local and global), while experts found they take up too much space. Designers should consider an auto-hide, or on/off visual representation of filters. Approaches such as interactive legends [17] could also reduce clutter of local filters, since most charts include labels of some sort. For both groups, being able to distinguish between global and local chart filters is important, and it should always be clear which categories and values contribute to existing filters. Finally, since novice users have trouble predicting conflicts in global and local filters, it is important to clearly mark such conflicts.

Language & terminology: Novice users are often not aware of international terms for different charts, metrics or functions. Thus local language support is important. Moreover, buttons, labels and icons (even universally understood ones) are not enough to provide novices with adequate information for their behavior in the context of dashboards. Detailed explanations need to be provided on demand (e.g. tool-tips).

8 Conclusions

We have presented Exploration Views, a prototype system that allows users to easily create, customize and interact with Business Intelligence visualization dashboards. EV was build following previous guidelines for novice visualization users. Through a user-study with novice and expert visualization users we verified that EV enables novices to create dashboards for answering complex questions. We found some differences in how experts and novice users interact with the interface, reported in our study. More importantly, based on our observations we provide new guidelines that augment previous work on designing for visualization novices, especially in the context of interactive visualization systems in the form of dashboards.

Of course further work is needed to gain more insights in how novices use dashboards. As study participants were not data experts, it could be interesting to see if our findings extend to data experts. Moreover, it would be interesting to observe novice user exploration behavior in a long-term study. Finally, our work focused more on the user interface design and use in a visualization environment for novices, and less on the insights gained from the data. Clearly this an avenue of future work.

9 Acknowledgments

We'd like to thank our users for their time and comments, Alexis Naibo and Didier Bolf for their support, and our reviewers for their useful suggestions. This research was funded by SAP BusinessObjects Innovation Center and SAP Research.

References

1. Card, S. K., Mackinlay, J. D., and Shneiderman, B. 1999. Readings in information visualization: Using vision to think. Morgan Kaufmann Publishers.
2. Capra, R. and Marchionini, G. 2008. The relation browser tool for faceted exploratory search. In Proc. JCDL08, ACM Press (2008), 420-420.
3. Chiang, A. 2009. Creating Dashboards: The Players and Collaboration You Need for a Successful Project, In Business Intelligence Journal, Vol. 14, No. 1.
4. Choosel 2009. <http://code.google.com/p/choosel/>
5. Cubramic, D. 2008. Polstar: Assisted navigation for exploring multi-dimensional information spaces. In Huma-Computer Information Retrieval (HCIR) 2008.
6. Cutroni, J. 2007. Google Analytics. Publisher O'Reilly Media, Inc., 2007.
7. CVOM patent: <http://www.freshpatents.com/Apparatus-and-method-for-data-charting-with-an-extensible-visualization-library-dt20080731ptan20080180458.php>
8. Dundas dashboard <http://www.dundas.com/Microsite/Dashboard>.
9. Few, S. 2004. Show Me the Numbers. Analytics Press.
10. Few, S. 2006. Information Dashboard design, O'Reilly.
11. Foody, P. 2009. User-Centered Business Intelligence, In Business Intelligence Journal, Vol. 14, No. 4.

12. Grammel, L., Tory, M., Storey, M-A. 2010. How Information Visualization Novices Construct Visualizations. In IEEE InfoVis - Transactions on Visualization and Computer Graphics 17(6): 943-952, 2010.
13. Hearst, M. 1999. User Interfaces and Visualization, Ch.10 Modern Inf. Retrieval, Baeza-Yates, R. & Ribeiro-Neto, B (edt).
14. Hearst, M. 2008. UIs for Faceted Navigation, Recent Advances and Remaining Open Problems. In HCIR 2008.
15. Heer, J., van Ham, F., Carpendale, S., Weaver, C., and Isenberg, P. 2008. Creation and collaboration: Engaging new audiences for information visualization. In Information Visualization—Human-Centered Issues and Perspectives, volume 4950 of LNCS State-of-the-Art Survey, pages 92–133. Springer Verlag, 2008.
16. Heer, J., Viégas, F.B., Wattenberg, M.: Voyagers and voyeurs: Supporting asynchronous collaborative information visualization. In Proc. of the Conference on Human Factors in Computing Systems (CHI), pp. 1029–1038.
17. Henry Riche, N., Lee, B. and Plaisant, C. 2010 Understanding Interactive Legends: a Comparative Study with Standard Widgets. In Proc. Eurovis 2010, 29(3), 1193-1202.
18. Mahyar, N., Sarvghad, A., and Tory, M. 2009. Roles of Notes in Co-located Collaborative Visualization. In Workshop on Collaborative Visualization on Interactive Surfaces (CoVis 2009), Oct. 2009.
19. Marchionini, G. 2006. Toward Human-Computer Information Retrieval, Bulletin of the American Society for Inf.Science.
20. Nielsen, J. 1994. Usability Engineering. San Diego: Academic Press. pp. 115–148.
21. North, C., Shneiderman, B., "Snap-Together Visualization: A User Interface for Coordinating Visualizations via Relational Schemata", In Proc. AVI 2000.
22. OLAP and OLAP Server Definitions. The OLAP Council. 1995.
23. Oracle BI 10g dashboard www.oracle.com/technology/products/bi/enterprise-edition-platform-components.html.
24. Raskino, M., Lopez, J. 2009. Early Findings From the 2010 Gartner CEO and Business Executive Survey. December 2009.
25. Roberts, J. C. 2007. State of the Art: Coordinated & Multiple Views in Exploratory Visualization. Proc. of the Fifth international Conference on Coordinated and Multiple Views in Exploratory Visualization. IEEE Comp. Society, pages 61-71.
26. SAP Business Objects Xcelsius www.sap.com/solutions/sapbusinessobjects/large/business-intelligence/dashboard-visualization/xcelsius-enterprise/index.epx.
27. Sorenson, M. 2002. Dictionary definition of dashboards, SearchCIO.com, 25/2/2002.
28. Spotfire <http://spotfire.tibco.com/>.
29. Tableau <http://www.tableausoftware.com/products/desktop>.
30. Willhalm, T., Popovici, N., Boshmaf, Y., Plattner, H., Zeier, A., Schaffner, J. 2009. SIMD-Scan: Ultra Fast in-Memory Table Scan using on-Chip Vector Processing Units. VLDB 2009.
31. Viegas, F. B., Wattenberg, M., van Ham, F., Kriss, J., and McKeon, M. 2007. ManyEyes: a Site for Visualization at Internet Scale. IEEE Transactions on Visualization and Computer Graphics 13, 6 (Nov. 2007).
32. Weaver, C. 2004. Building highly-coordinated visualizations in improvise. In IEEE Symposium on Information Visualization (InfoVis 2004), pages 159–166.
33. Williamson, C. and Shneiderman, B. 1992. The Dynamic HomeFinder: Evaluating Dynamic Queries in a Real-Estate Information Exploration System. In Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pages 339 – 346, ACM.
34. Woods, D. 1984. Visual momentum: a concept to improve the cognitive coupling of person and computer. Int. J. Man-Mach. Stud., 21(3):229–244, 1984.
35. Zelazny, G. 2006. The Say It With Charts Complete Toolkit. McGraw-Hill, 2006.