Paper Landscapes: A Visualization Design Methodology

Richard Brath
Visual Insights Inc., 555 Richmond St. W. #1108, Toronto, Canada
richard.brath@visualinsights.com

Abstract

Paper landscape refers to both a iterative design process and a document as an aid to the design and development process for creating new information visualizations. A paper landscape engages all stakeholders early in the process of creating new visualizations and is used to solicit input; clarify ideas, features, requirements, tasks; and obtain support for the proposal, whether group consensus, market validation or project funding.

1. Introduction

Information visualizations have a wide range of potential representations and interactions. Hundreds of visualization have been proposed and presented at InfoViz over the last ten years, and hundreds more have been built at universities and visualization companies.

Often, visualizations are the result of multiple professional disciplines, such as technologists, data specialists, visualization specialists and intended users. With multiple stakeholders in the project, there can be significant differences in design intent that are not apparent while either discussing or referring to written visualization specifications. Uncovering differences in understanding early during the project lifecycle reduces risks and costs associated with the project and increases the confidence in the ability to deliver the target solution.

William Wright [Wri95,97] coined the term paper landscape in the early 1990s. A paper landscape refers to both a methodology and a document as an aid to the design and development process for creating new information visualizations. The paper landscape records in a visual document the target visualization thereby making visible what may be otherwise vague design intent. For example, a stakeholders may agree on a “map” component within a visualization, but one stakeholder may anticipate a chloropleth map while another user may anticipate a map with geographically located markers. If this is not resolved prior to coding the application, the level of effort to adjust the visualization and associated cost will be greater later in the project cycle. In addition, the paper landscape document acts as a design aid, enabling designers and stakeholders to record, explore and evaluate various design alternatives; and facilitates discussion for collecting feedback into an iterative design process.

In addition to the visual representation, a paper landscape document records the key goals of the visualization and the primary interactions; as shown in figure one. As a result, the paper landscape is a concise, easy to understand, self contained description of the overall visualization project.

2. Related work

The design process for creating information visualizations has not been extensively documented, perhaps because the process already occurs in some intuitive form in most projects. With the perspective of having created hundreds of paper landscapes at Visual Insights, the process has been identified and at times become somewhat formalized as well as taught to new staff members and key clients.

Architecture, Animation and Design

Many early employees at Visual Insights have educational and professional backgrounds in design.

Drawings and sketches are used throughout traditional design processes both to record ideas and as an aid to thinking about the design. These drawings and sketches are highly iterative evolving through several stages of development. Early stage “design drawings” record the key attributes of the intended object including form and function thus aiding communications with stakeholders such as clients, engineers, and review boards.

Human Computer Interaction

The development of interactive software typically requires iterative design practices to incorporate feedback early in the software lifecycle to inform design decisions that affect usability. [Dix93]. The field of human computer interaction (HCI) recognizes a wide variety of techniques for capturing user feedback very early in the software design process, including rapid prototyping techniques, such as paper mockups, story boards [Gou90], wizard of oz techniques, simulations and limited functionality applications.

Iterative design and feedback techniques such as paper mockups and storyboards are also valuable to the overall design process by capturing design rationale; as a record of various design decisions, design trade-offs between alternatives, and capturing context [Dix93].

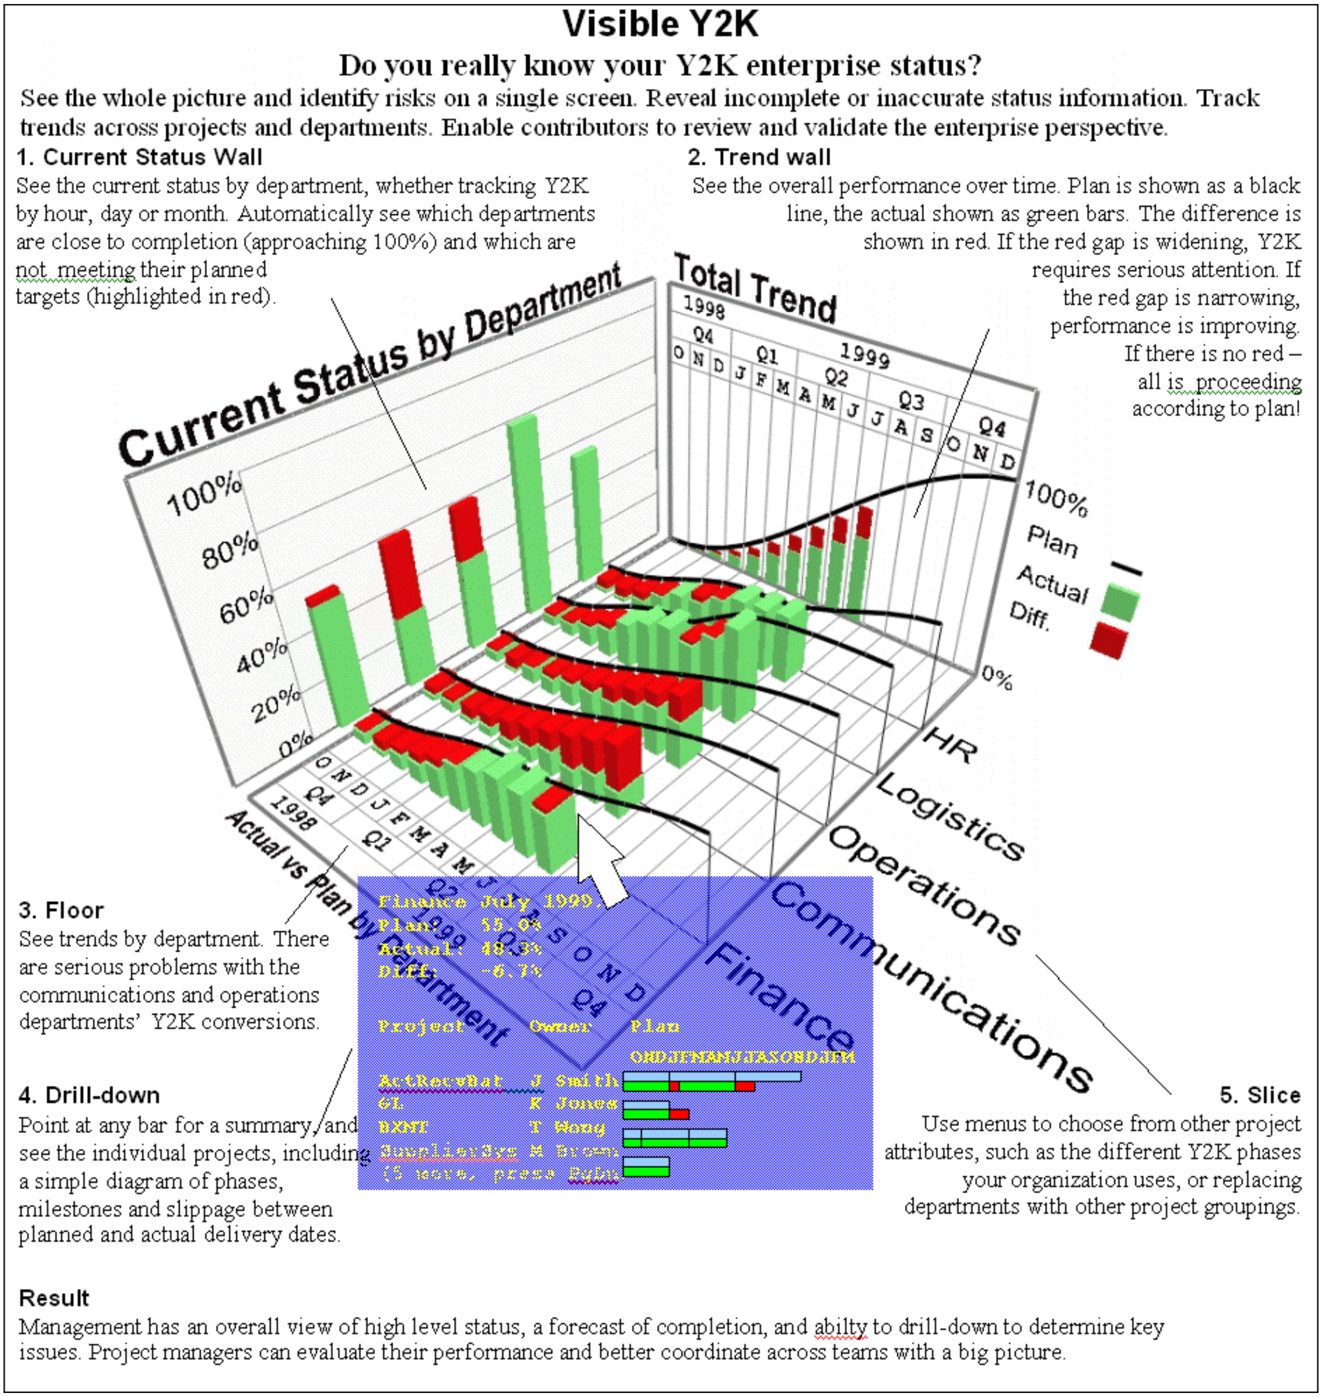

Figure 1. Visible Y2K paper landscape.

Visualization Guidelines

HCI recognizes rules, guidelines and standards can significantly enhance usability. In the field of information visualization, Bertin [Ber83,00], Tufte [Tuf83,90,96], Cleveland[Cle85] and others are recognized as for their contribution to visualization guidelines.

Bertin organizes the visual and data elements that are central to the mapping of data to visual representations and presents a set of rules within this organized framework. Tufte presents numerous case studies and consolidates findings into a series of guidelines. Similarly Cleveland, MacKinlay [Mac86], Brath [Bra97] and others [Rot94, Chu96, Zho96] have presented various other case studies and guidelines as well.

3. Creating a paper landscape

Designing a visualization requires a diverse set of skills:

- *Software application design* - what can and cannot be programmed. Knowledge of object-oriented programming, databases and the graphics programming are ideal.
- *User interface design* – how to map workflows into logical screens and interface components. A user interface designer with expertise in direct manipulation interfaces is desirable.
- *Visual design* - what works visually. Typically a graphic designer or architect with some visualization experience is desired.
- *Business requirements* - what questions need to be answered by this visualization. A business analyst or an end-user can usually articulate the deficiencies in their current processes.

Since it is unlikely that any one person has all the above abilities, the design of a visualization is typically a multi-disciplinary task involving the participation of a team of people.

Designing a paper landscape is an iterative process. A typical iteration will:

- *Collect requirements*, e.g. explicitly ask for requirements; or review existing analysis techniques with attention to areas of difficulty.
- *Propose possible designs*, e.g. sketch a concept on a whiteboard; or view an example visual; or walkthrough a demonstration of an existing related visualization
- *Test the design*, e.g. walkthrough a workflow; or consider how the design handles particular data scenarios, such as negative data values, null data, etc.
- *Collect feedback*, e.g. comments in a meeting within the design team; or formal review meetings with management, and customers; or informal, on-going email threads.

In the earliest stages, these iterations will be fast, typically done in a group setting using tools such as pen and paper, whiteboards, and interactive demos. In successive iterations the cycle will become slower, using more care and detail in the mockup of the visual design.

Software creating detailed visual mockups include categories such as:

- *Basic drawing*, such as Microsoft Word and Powerpoint
- *Illustration*, such as Corel Draw or Adobe Illustrator
- *Multimedia*, such as Powerpoint or Macromedia Flash
- *Rapid prototyping*, such as Visual Basic or Visual Insights Studio. AVS Express and VTK are likely useful as well.

We have not used general purpose 3D nor CAD tools. The effort has been too great to learn these tools in past attempts, and the speed of assembling a scene for the designer too slow.

The number of iterations varies with each project. In some projects, this phase has been completed in days. In one recent project, this phase required 6 months, due to changing data availability, which impacted the requirements for the visuals. Many man months worth of software development were saved by staying in a paper landscape phase.

A paper landscape is essentially finished when the team reaches a consensus on the design. In addition, the team should be aware that the design may still iterate (successively refined) through the development; for example, to accommodate unexpected data found during the implementation; or visual enhancements, etc.

4. Elements of a paper landscape

A simple paper landscape titled “Visible Y2K” is shown in figure 1. Paper landscapes generally contain the following information, all shown in this example.

4.1. Title

An easy way for everyone involved in a project to simply refer to it by a name. If a project fractures into a number of different designs, design alternatives or phases, titles are a means to easily refer to different groups of design ideas with a single label. In this case “Visible Y2K” immediately conveys what the visualization is about.

4.2. Goal

One paragraph outline of the key goals and/or the key business questions that the visualization is intended to answer. In this case, the key issue is “do you know your Y2K enterprise status?” Managers were worried in 1998 that they did not understand their Y2K project status. Subsequent points identify the issues in more detail, e.g. “See the whole picture and identify risks on a single screen” relates to the issue of dealing with massive spreadsheets and project plans that many Y2K project managers attempted to use but did not provide high level overviews.

4.3. Visual Diagram

A diagram in the center depicts the key visual elements of the visualization, including important interaction elements such as key dialog boxes and other components elements that may be located in separate windows. Items in the diagram should be labeled as they would in the intended application to facilitate understanding of the visual scene. Legends should be included if these are important. Visual accuracy, such as perspective and relative sizes of objects is not critical, particularly at early stages.

4.4. Description Sequence

For any visualization with more than a simple chart, it is important to provide the viewer with a description of the visualization. A numbered sequence can be used to provide an incremental framework for understanding the application, as well as one potential workflow for accomplishing a simple task.

4.4.1. What is initially visible

Prior to explaining any interactions, it is important to first explain what is visible. Explaining physically and logically separate areas of the visualization provides a means for the viewer to begin to construct a mental model of the application.

In each numbered element, the description first provides a name to the scene element (used to facilitate discussion). Then the initial sentence typically describes the element (e.g. Back wall shows current status by department). Further sentences describe the semantics of the visual (e.g. green bars indicate status OK and red bars indicate project lateness). Successive sentences may introduce interactive elements.

In this Y2K example, the walls and floors are explained in the first three numbered items.

4.4.2. What the interactions are

Subsequent steps outline interactions. In the Y2K example, the first interaction shows a “brushing” event. Brushing is an interactive technique for displaying quantitative data [Cle85].

4.4.3. Description of sample workflow(s)

Workflow can be implicit, through the description sequence or made explicit. In applications supporting multiple workflows it is valuable to point out different workflows, as different viewers may have different workflows in mind.

4.5. Technical considerations/limitations/tradeoffs

Technical considerations may limit the visualization application, for example, certain data may not be available, or response times for some events may jeopardize the utility of those interactions. These issues should be explicitly stated, thereby encouraging alternate solutions to be discussed.

4.6. Results

As a result of the visualization, the end-user will gain some benefits, implied in the initial goal. A results section should show how the goal is achieved through the use of the proposed visualization.

5. Example of paper landscape development at Statistics Canada

Statistics Canada is a national agency with the responsibility to collect, compile, analyze, abstract and publish information spanning the entire social and economic landscape of Canada. The agency sells over 1.2 million publications annually, and has a staff of over 2,200 including analysts and other professionals who specialize by industry and social disciplines.

The Monthly Survey of Manufacturers department uses Visual Insights to analyze data from thousands of companies. Errors and omissions must be eliminated from the sample as early in the process as possible. Since this data is used as the basis of modeling a national perspective, the effect of erroneous data is magnified.

Survey analysts use extensive statistical reports to seek out anomalies in the data. These are then used in conjunction with spreadsheets and charts to review and present survey results to management. Hundreds of published cells within which the data resides need to be reviewed in order to substantiate the data.

Over the course of one year, Visual Insights worked together with Statistics Canada to collect requirements, design and implement a visualization for the Monthly Survey of Manufacturers to facilitate this review of data.

5.1. Initial Paper Landscape Design

The initial contact with Statistics Canada was a meeting to provide an understanding of Visual Insights capabilities. Monthly Survey of Manufacturers outlined the nature of data – multiple timeseries of 11 different measures for 11,000 companies organized by geographic region and Standard Industry Classification (SIC codes); and high level requirements, such as the need to compare multiple timeseries simultaneously, at both the level of individual companies and at intermediate aggregations.

Based on this initial meeting a paper landscape was produced to respond to a “Request for Proposal” (RFP); and to be used in subsequent meetings to draw out discussion from the various participants. This paper landscape is shown in Figure 2.

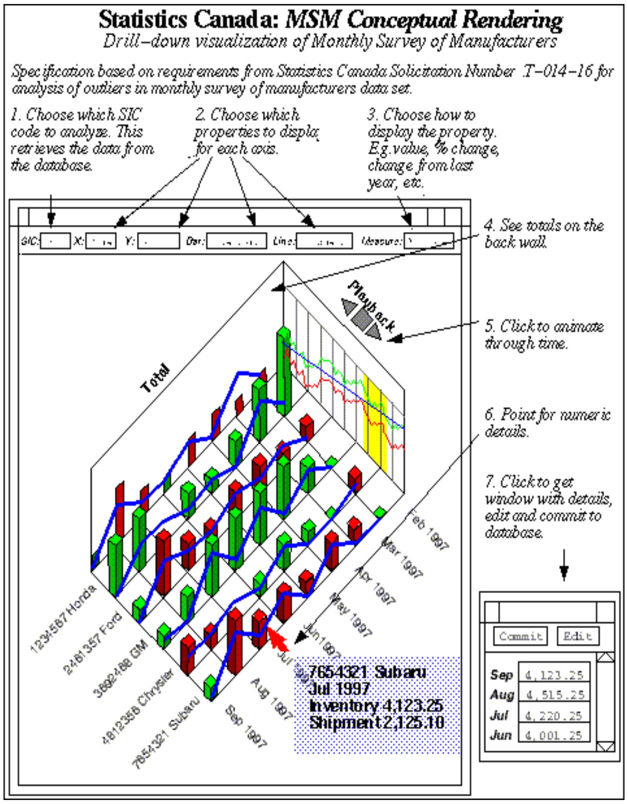

Figure 2: Statistics Canada, first design

This paper landscape is terse and does it completely follow the above guidelines because it is part of a larger proposal document where these other items are addressed. For example, the goals section simply refers back to the RFP; and the sequential description does not explain the visual diagram - it assumes that the diagram is self explanatory.

5.2. Detailed Paper Landscape Design

Following the success of the RFP, Visual Insights subsequently held a full-day meeting with key project stakeholders, including a half-dozen users, information technology staff, and management.

Significant findings were made that had a strong influence on requirements and design:

- The users are sophisticated statisticians with deep expertise in the domain of the data (Subject Matter Experts – SMEs).
- The stakeholders use terminology, such as "micro-data" to refer to the raw data at the company level, "estimate data" to refer to high-level aggregate data, etc.
- The primary task in the data analysis is to identify major changes in any reported index; determine the cause of the change; determine whether the change is legitimate; correct errors and explain the legitimate changes. E.g. "Did auto industry inventories decrease 3% due to an error in the transcription of the data or due to strike by workers?"
- Different groups of users have evolved different workflows and artifacts (scripts, reports, etc) to achieve the task.
- One predominant workflow is "top-down analysis". The userstarts with high-level metrics, focuses on anomalies, then drill-downs to successively smaller subsets of data until the underlying micro-data containing the anomaly is identified.
- Another predominant workflow is "bottom-up analysis". The user starts with a subset of micro-data (timeseries of individual companies) sorted to bring the largest contributors to change to the top. By successively reviewing the this micro-data, the estimate data will adjust.
- A third workflow is reviews timeseries of indirectly related data. These other timeseries are leading indicators that help explain why a change may have occurred in the data. E.g. a decrease in employment levels may precede a decrease in shipments.
- One long-established analysis technique is "alerts". Alerts are triggered at both micro and estimate level data when the data surpasses a threshold indicating further investigation is warranted.

Based on these findings and a review of the existing reports and analysis techniques; many design ideas and fragments were sketched on a whiteboard, including:

- *color-coded line segments* within a time-series to differentiate portions that are increasing vs. decreasing.
- *"flags"* to indicate pre-computed triggers.
- *color-coded grids* to indicate which data subsets contribute most significantly to the larger estimate.
- *data adjustment techniques* for seasonal adjustments, for approximating new estimates, etc.

At the end of this meeting, it was determined that at least two paper landscapes would be created, one to reflect top-down workflow, another to reflect bottom up workflows.

A third paper landscape was also created, combining elements from the first two, as a "hybrid" supporting both top-down and bottom-up analysis. This third 'hybrid' paper landscape is shown in figure 3. This third paper landscape was selected for the implementation, with some modifications.

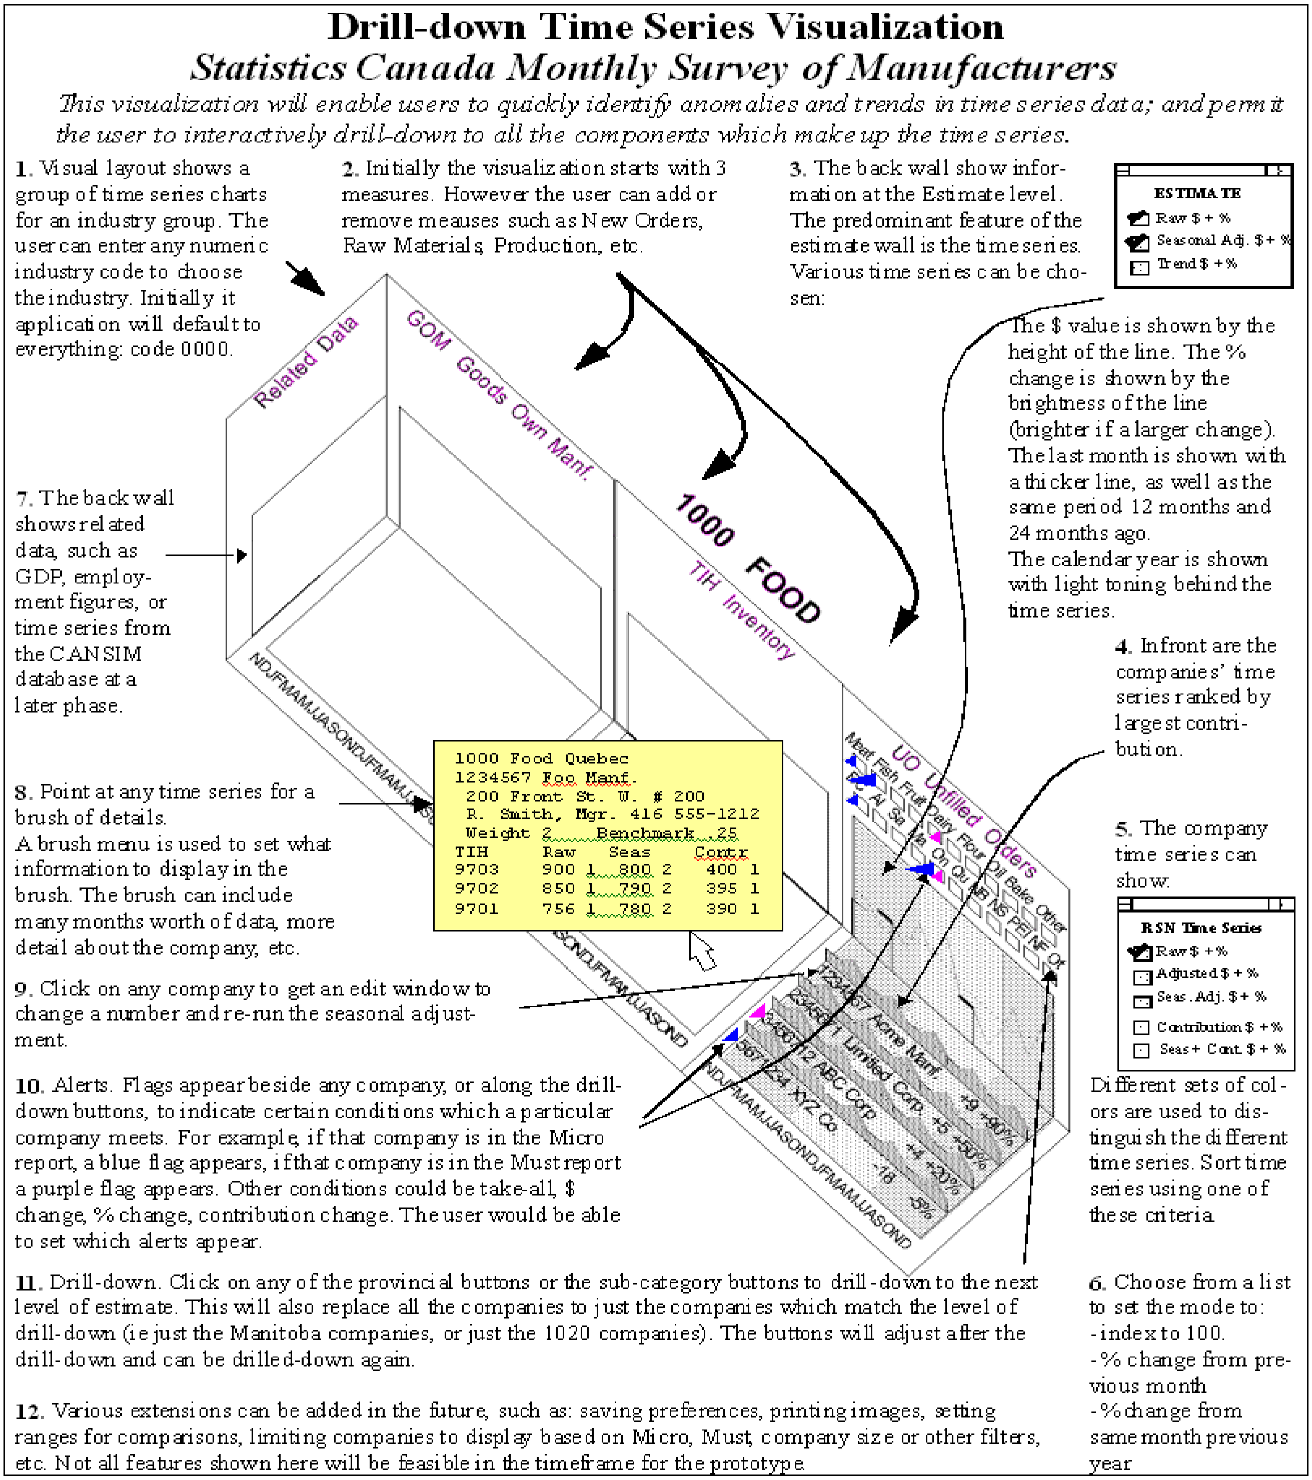

Figure 3. Statistics Canada detailed design.

This is a very detailed paper landscape – intended for detailed feedback and for construction. Familiar paper landscape elements are here, as well as highly evolved visuals and interactions:

- Graphical representations (glyphs) use a number of visual attributes to convey information, such as the timeseries (points 3 & 5), using line charts with brightness and hue and a shaded context.
- Note the many interactions for manipulating timeseries (point 3, 5 & 6) with the ability to view raw data, seasonally adjusted data, weighed data (contribution), indexed to a base date, previous month

of same month previous year. These many techniques are important to workflow (e.g. estimates are only recalculated on particularly days of the month, weights are re-balanced on a yearly basis) and accommodating the needs of different users (e.g. agricultural data has strong seasonality, manufacturing data does not).

Workflow is not explicitly outlined in the description; however, the ability of this third design to permit multiple workflows to work together is critical its acceptance:

- The right wall (estimates wall) facilitates top down workflow. The estimate timeseries and the flagged buttons provide information for deciding where to drill down next.
- The floor (micro data) facilitates bottom up workflow. Sorting, labels, side-by-side comparison and flagged series provide information for deciding which timeseries to focus.
- The left wall (related data) supports the addition of external data to the analysis.

The success of this solution is that it not only accommodates different workflows but enables users to move back and forth between different workflows - visually and interactively moving back and forth between micro data and estimate data.

5.3. Feedback and Implementation

The paper landscape described above was accepted, with the following modification: the button wall was replaced with grid wall for identifying changes and contribution by both industry and geography.

The paper landscape also facilitated discussion about data transformations. E.g. seasonal adjustment calculations are done by highly specific algorithms on a centralized computer at scheduled times. The paper landscapes showed interactive updates of estimates and seasonal adjustments. This expectation became obvious during discussions between users and technical staff viewing visuals on a paper landscape, as opposed to footnotes in a specifications document. A workaround was resolved before any code had been written. Attempting to resolve this during development would likely have added at least 2 more man months of effort to the project.

The initial pilot implementation is shown in figure 4. All the key features as described by the paper landscape can be seen.

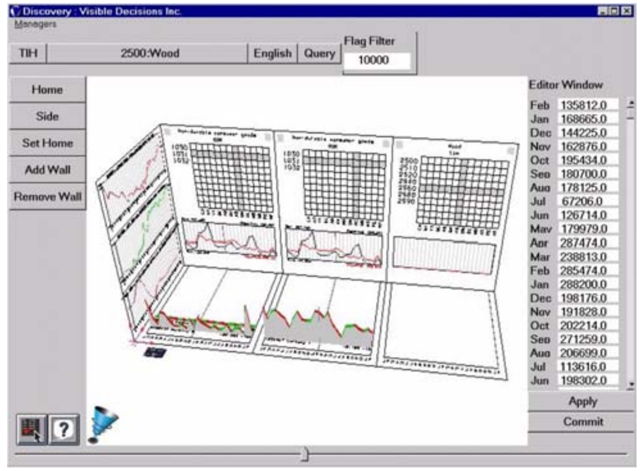

Figure 4: Statistics Canada Pilot Implementation.

The production implementation is shown in figure 5. All the same key features can be seen, with many additional enhancements.

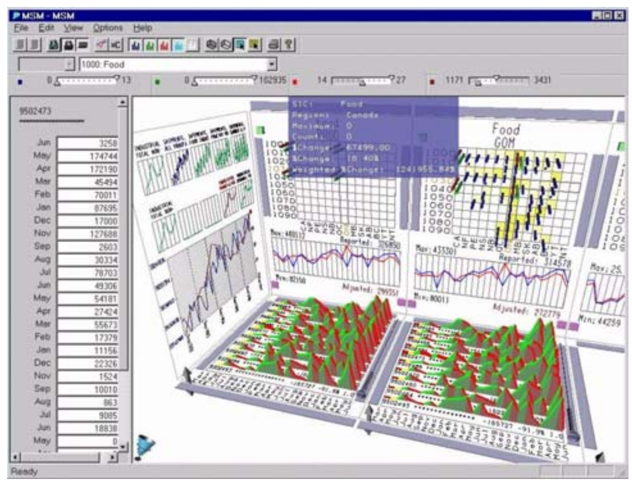

Figure 5: Statistics Canada Production Implementation.

5.4. Results

A very demanding user community, a data rich environment with analytic complexity, multiple workflows, and the pressure of short production deadlines made this particular visualization very challenging. The paper landscapes greatly facilitated early discussions about workflows, visual representations, and interactions. It is highly doubtful that such a successful solution could have been achieved in a “first” implementation without the use of paper landscapes. The visualization is in use today at Statistics Canada, six years after the initial meetings.

6. Conclusions and future work

Paper landscapes are valuable design time tool for stakeholders to validate their understanding of the project and assess downstream impact. Stakeholders can use the document to walk through workflows, effectively testing use case scenarios. Proposed solutions can be evaluated against each other. Project owners can evaluate whether sufficient resources (computing power, data, project staff, training resources) are available.

Paper landscapes can reduce the overall development effort. In one scenario the team was able to identify issues between data actually available and the intended display, saving more than two man months. In a number of cases, paper landscapes have identified different workflows and requirements among user groups. Without identifying and accommodating these needs early in the process would risk a visualization re-write or non-acceptance of the solution.

A difficulty with paper landscapes is the difference between the design artifact and the actual implementation. If the design document is too abstract (a whiteboard sketch or a drawing such as figure 3) some people may have difficulty understanding the representation. This is similar to the difficulty people may have relating a drawing of a floor plan to a building. When we use a rapid prototyping tool (Visual Insights Studio), the visuals can appear as they would in the target application. Unfortunately, in this scenario, many people have difficultly understanding that the prototype is not the finished application – it appears to be 3D, it appears to have data, and it may even appear to have interaction.

Paper landscapes, whether on paper or constructed with a rapid prototyping solution, are a stepping stone towards a solution. In an ideal world, visualization designers would compose the visuals and interactions for their target visualization in an easy to use, point and click environment against real data. Because of the level of complexity in understanding coding, databases, user domains, data modeling, client/server architectures, visual design, etc makes such an integrated solution a “holy grail”. Even traditional user interfaces (e.g. windows application, website) require a multi-disciplinary approach with designers far removed from the larger context of the application because of the complexity of various systems involved.

Visual Insights has done some R&D work in this area. It should be noted that software tools made for designers (e.g. Alias, AutoCAD, Flash) are very complex. The target tool may require significant effort create – more than the ability of the visualization market to support today.

In the future, better design and visualization tools should replace paper landscapes.

7. Acknowledgements

Visual Insights thanks Statistics Canada and the Monthly Survey of Manufacturers team for permission to refer this project.

References

[Ber83] J. Bertin. *Semiology of Graphics*, University of Wisconsin Press, 1983.

[Ber00] J. Bertin. Matrix Theory of Graphics. In *Information Design Journal,* Vol. 10, No. 1, 2000.

[Bra97] R. Brath. Interactive Information Visualization Guidelines. In *Proceedings of HCI International ’97*.

[Chu96] M Chuah and S Roth. On the Semantics of Interactive Visualizations. In *Proceedings of Information Visualization ‘96*: 29-36. IEEE, 1996.

[Cle85] W. Cleveland, *The Elements of Graphing Data*. Hobart Press, 1985.

[Dix93] A. Dix, J. Finlay, G. Abowd, R. Beale, *Human-Computer Interaction,* Prentice Hall, 1993.

[Gou90] J. Gould, S. Boies, S. Levy, J. Richards and R. Schoonard. The 1984 Olypmpic message system: A test of behavioural principles of system design. In *Human-Computer Interaction,* Prentice Hall, 1990.

[Mac86] J. Mackinlay. “Automating the Design of Graphical Presentations”, in *ACM Transactions on Graphics*, Vol. 5, No.2, April 1986

[Rot94] S F Roth, J. Kolojejchick, J. Mattis, J. Goldstein. Interactive Graphic Design Using Automatic Presentation Knowledge. In *Proceedings of CHI’94*: 112-117, 1994.

[Tuf83] E. R. Tufte. *The Visual Display of Quantitative Information*, Graphics Press, 1983.

[Tuf90] E. R. Tufte. *Envisioning Information*, Graphics Press, 1990.

[Tuf96] E. R. Tufte. *Visual Explanations*, Graphics Press, 1996.

[Wri95] W. Wright. Information Animation: Applications in the Capital Markets. In *Proceedings of Information Visualization’95*. IEEE, 1995.

[Wri97] W. Wright, Information Visualization: The Fourth Dimension. In *Database Programming & Design*, Vol. 10, No. 4. Miller Freeman, April 1997.

[Zho96] M Zhou and S Feiner. Data Characterization For Automatically Visualizing Heterogeneous Information. In *Proceedings of Information Visualization ‘96*: 13-20. IEEE, 1996.