

# Exploring Global Terrorism Data: A Web-based Visualization of Temporal Data

Joonghoon Lee  
Computer Science Department & HCIL  
University of Maryland, College Park  
jhl@cs.umd.edu

## Abstract

*Recently, there has been a dramatic increase in both the amount of terrorism related literature and public interest towards the matter. Even so, it is difficult to gain access to relevant data. This is because much of the available data is not well organized nor of high quality, and available data is not very well presented.*

*The GTD (Global Terrorism Database) [4] is an open-source database that has information on 75,000 terrorist events around the world since 1970 through 2004. It is managed by the National Consortium for the Study of Terrorism and Responses to Terrorism (START) at the University of Maryland.*

*This report will introduce the GTD Explorer, a web-based interactive visual exploratory tool that deals with this data. It counts the number of incidents grouped over a certain criteria, and stack charts on top of each other to see both individual and accumulated patterns of incidents over time. This tool provides insight to experts while making the data approachable and informative. Making the visualization light-weight and universally accessible was one of the main concerns.*

*Since this is a web-based tool, there are certain challenges and limitations to the approach. Details of the implementation will be covered. Interesting facts that were found using the tool will be discussed.*

## 1 Introduction

Dealing with terrorism has become the quintessential social problem for civilized societies. The world has changed in many ways during the last decade in response to this increased awareness. Yet, terrorism has always been existent.

Recently, there has been a dramatic increase in both the amount of terrorism related literature and public interest towards the matter. Even so, it is difficult to gain access to relevant data. This is because much of the available data is not well organized nor of high quality, and available data is not very well presented.

The Global Terrorism Database [4] is an open-source database that has information of terrorist events around the world since 1970 through 2004. It is managed by the National Consortium for the Study of Terrorism and Responses to Terrorism (START) at the University of Maryland, which claims to be “the most comprehensive open source data set on terrorism”. The START intends to make this data more accessible to the public and this tool is one way of doing so.

The GTD includes information on more than 75,000 terrorist attacks that occurred around the world between 1970 and 2004. The GTD data were collected by trained researchers who recorded terrorism attacks from wire services (including Reuters and the Foreign Broadcast Information Service), U.S. and foreign government reports, U.S. and foreign newspapers (e.g., the New York Times, the British Financial Times), and in recent years, the

internet. The same coding scheme was used during the first 28 years of data collection. Researchers at the National Consortium for the Study of Terrorism and Responses to Terrorism (START) have been updating the data base during the past five years.

When this data was first released as a web-based version, the START GTD site had 36,000 visitors per month. START is expecting this number to increase substantially when a new version is released next year. In addition to supplying the data on line, this data is made available to policy analysts in the intelligence and defense communities and to academic researchers.

For the policy makers in particular, it is important to have a web-based version of the data. Given the longitudinal nature of the data, the database lends itself especially well to time series graphs.

This report introduces the GTD Explorer, a web-based interactive visual exploratory tool that deals with this data. the GTD Explorer basically takes key variables in the database and allows them to be visualized in data streams – the larger the number of attacks, or the number of weapons, the larger the stream. It counts the number of incidents grouped over a certain criteria, and stack charts on top of each other to see both individual and accumulated patterns of incidents over time. In this way, it is possible to see very rapidly how different units of analysis contribute to total outcomes.

This tool provides insight to experts while making the data approachable and informative. Making the visualization light-weight and universally accessible was one of the main concerns.

Since this is a web-based tool, there are certain challenges and limitations to the approach. Details of the implementation will be covered. Interesting facts that were found using the tool will be discussed.

## 2 Related Work

Martin Wattenberg's NameVoyager [1] is a web-based visualization that provided much inspiration for this project. It is an interactive visualization of the popularity of America's new born babies' names over time. Each name is represented by a stripe which corresponds to the timeline of the name, from 1900 to 2006. A stripe is thicker on years when which the name was popular.

It uses a bi-color scheme, with gradient colors for usage frequency ranges. "In keeping with contemporary American custom" [2], light blue colors are used for boys and pinks for girls.

The search functionality searches for prefixes in the entire dataset. The NameVoyager is a very compelling example of a web-based visualizations that provides much information to a domain expert in the subject matter, baby names, while also providing much pleasure to the general audience. It served as a guideline in building the GTD Explorer in many ways.

In [2], Wattenberg discusses how the NameVoyager can create an online social environment. It provides a thoughtful analysis of several factors that led to the popularity of the NameVoyager, based on comments collected from various web sources. Wattenberg analogizes NameVoyager commentators to MUD (multi user dungeon) game users [15], categorizing them based on their motive for interacting with other users.




Figure 1. NameVoyager [1]

HistoryFlow from IBMs Viegas, Wattenberg and Dave visualizes the editing history of Wikipedia pages. It is an effective exploratory data analysis tool that reveals interesting facts regarding the edits done to a file over time. It is a strong example of visualizations assisting in providing insights. Their visualizations resemble ThemeRiver™ and Inselberg’s parallel coordinates.




Figure 2. HistoryFlow [5]

“ThemeRiver™: In Search of Trends, Patterns, and Relationships” [6] visualizes thematic variations over time across a collection of documents. A three dimensional version was also presented in “Interactive Poster: 3D ThemeRiver’ [7]’.




Figure 3. ThemeRiver showing Castro data from November 1959 through 1961. [6]




Figure 4. 3D ThemeRiver [7]

The New York Times featured a well crafted visualization of the box-office record of movies from 1968 to 2007 called “The Ebb and Flow of Movies: Box Office Receipts 1986 - 2007” [3]. It’s appearance is also similar to that of the ThemeRiver visualization. It plots curves that build upwards and downwards along a horizontal time line. The height of the curves depict how movies have fared at the box office. They adjust the figures for inflation to make a more meaningful compromise.

There are two versions, on and off-line. The difference between the two media, screen and newspaper, raises interesting questions. In the offline version, they have a much larger presentation area, making it possible for better labeling, and also permitting more information to be immediately displayed. On the other hand, the online version

has interactions, allowing for zooming in, and more detailed information.




Figure 5. The Ebb and Flow of Movies: Box Office Receipts 1986 - 2007 [3]

Revisionist[9], visualizes the evolution of software projects. It treats projects as a collection of files. The addition, removal and reorganization of these files throughout the course of development is visualized. It is an interesting approach, although addresses the problem of displaying massive amounts of text in limited screen space.

Ringel et al, "Milestones in Time: The Value of Landmarks in Retrieving Information from Personal Stores"[8], discusses a means to explore events along a timeline. The strategy is to narrow the scope of the search so that users can use relevant, nearby events to assist in their search. This strategy suggests a possible expansion to our tool, if we were to make our search more fine-grained.

"Building a Global Terrorism Database"[4] explains the dataset this tool was built upon. This database originally comes from the PGIS Corporations Global Intelligence Service (PGIS). It is a unique database with extensive detailed data. There is a possibility of potential media bias. Other weaknesses include misinformation, lack of information beyond incident specific details alone, and missing data from lost cards (data for the year 1993, due to an office move).

### 3 The GTD Explorer

The GTD Explorer is a tool that lets users explore the trend of terrorism incidents over time through an interactive visualization. The visualization consists of the main visualization, the search control, grouping control, filter control and details table.




Figure 6. The GTD Explorer

The main visualization is a stacked chart image of the frequency of terrorist incidents each year. When the tool first loads there will be 203 stripes shown each corresponding to a country. These stripes each show how many terrorist attacks occurred in the corresponding country each year. The number of attacks is represented by the thickness of the stripe at the year on the x-axis.

Although the default grouping for stripes is the country the incident occurred in, it can be changed to other values. The combo box on the top right of the tool allows users to group the stripes using a different criteria such as Terrorist Organization, Target Nationality, Target Type, Type of Attack, Weapon Type, Region, etc.. In this chapter, we will assume we are using Country.

### 3.1 Interactions

The GTD Explorer follows Shneiderman’s mantra of “overview first, zoom and filter, details on demand”. [14] Once the program is loaded and displays the full set of data as described above, users may filter the data using two types of filters, or may demand further details of the data. Several user interactions are supported to make the data visualization more explorable and informative. The goal was to make these intuitive, so that first time users can easily use most of the functionality without a tutorial.

### 3.1.1 Search Box

Beneath the main visualization is the search control. It is a text box which users can edit. Whenever text is updated in here, the main visualization filters out the stripes to only show those that have a prefix in them that matches to the content of the text box.




Figure 7. (left) Countries view filtered by 's' in the search box  
(right) Countries starting with 'south'

### 3.1.2 Filter Control

The filter control is a vertical slider that has as many ticks as items in the visualization. This allows to filter out stripes that have a total count of incidents that falls outside the range of the filter control. The numbers on the slider represent the *rank* of the total count of incidents.



Figure 8. Rank Filter

This filter was originally designed to use the raw numbers. But in many cases the full range becomes very large, making it difficult to make minuscule changes. Another option would be to use a logarithmic scale. This could be more suitable in some cases, but inconvenient when there are only a few (< 20) items to filter, which is the case for several criteria, such as weapon type or attack type.

### 3.1.3 Mouse Interactions

When a user's mouse hovers over a stripe, the name of the country that matches to the stripe, along with the color of the stripe appears in the upper left hand side. The year at which the mouse is at, along with the number of incidents during this year also appear beneath the country name. Also the targeted stripe's border is highlighted in bright yellow, to let users see the selected stripe.




Figure 9. A highlighted stripe

If a stripe is clicked on, the name of the selected item is filled in the search box, as if users typed it in. The visualization shows only this stripe.

The ‘Region’ view behaves in a slightly different way. Instead of showing the corresponding stripe, it shows all countries within this region. This is because we believe the details of a region are the countries in it.

### 3.1.4 Animation

Whenever the visualization is updated, the transition is shown as a short animation. This allows users to follow how the trend and patterns change as a result of the update. It also masks the short delay that occurs as a result of the update. When many visualization updates are made the transition is not as smooth as intended. This is an area that needs further improvement.

### 3.1.5 Details Grid

At the bottom right of the tool, there is a data grid that shows the data that is currently displayed. It gives further details of the data currently displayed. Each row represents a single stripe. First comes the name, and then the color of the corresponding stripe in the visualization. After that comes the cumulative count of data for this item. This count is used in filters.



Name		Count
Spain		2669
Sri Lanka		2010
South Africa		1841
Somalia		167
Syria		147
Switzerland		90
Sudan		73
Soviet Union		71
Senegal		67

Figure 10. Details Grid

Every time the visualization is updated, so is the data in this grid. It always shows all items that should be displayed. For implementation reasons, some of this data is not in fact visible. This will be discussed further in ‘Coloring’ (4.2.3).

### 3.1.6 Reset

At the top right of the application is a reset button. This cancels out all filters so that the visualization displays the full range of data. Users who want to go view the initial state of the visualization can use this button to initialize the visualization to display the full range of data without any filters applied.

## 4 Implementation

### 4.1 Development Tools

The GTD Explorer was implemented as a Flash application using the Adobe Flex platform and the Flex 3 SDK, so that it can be run on a wide variety of systems and browsers. The layout was done in mxml, and the underlying data structures and actions were implemented in ActionScript. The underlying data was split into small files to be loaded into ActionScript objects.

The decision to use Flash was because it is nearly ubiquitous on internet-enabled desktop computers. Nearly every computer has flash installed on it, making it the first choice for an application which aims to attract a large audience. Adobe claims that over 98.8% of internet-enabled desktops worldwide have Adobe Flash Player installed as of March 2008. [20]

The flare visualization toolkit [17] was used for the main visualization and transition effects. It is a port of the prefuse visualization toolkit [18], to an Actionscript library, and is an open source project that is still in alpha testing stage. Prefuse is a widely used visualization tool that was used in SocialAction, NameVoyager, Vizster etc.. We also used flexlib, the Open Source Flex Component Library to make use of their advanced components.

Java is another viable choice. NameVoyager [1] has used Java, to much success. Java can run on most web browsers on most platforms and has good library support. Being a general purpose language, it can also deal with a much wider set of tasks, and allows better performance.

But since reaching the general public was one of our main goals, we chose Flash over Java for it's ubiquitousness.

### 4.2 Design Decisions

#### 4.2.1 Layout

Since the visualization should comfortably fit into most computer screens, the size of the application was set at  $800 \times 600$  pixels.

In the GTD Explorer, the pattern of change over time is the most important information. To most effectively convey this information, the stacked chart was plotted over a time range, and the chart was set to span the entire width of the application, and roughly half of the height.

Other components or labels are laid around this main visualization. At the top, there is a label saying "Incidents per Year Grouped By", followed by a pull-down combo box. This combo box lets users select the grouping criteria. The label and combo box are positioned so that they appear as "Incidents per Year Grouped By *Grouping Criteria*"

Beneath the visualization, are the filter controls and the details data grid. A vertical, not horizontal, slider was used for the incident count filter, since it seems to match the stacked charts, which are also stacked vertically.

It is a vertical slider with two control thumbs. The region between the two thumbs is colored to highlight the valid range. We used flexlib [19] version of the flex vertical slider, to allow for dragging the region in between the two sliders, not only the individual thumbs.

The details grid is positioned at the bottom right. We gave the the grid as much vertical space as we could without compromising the size of the main visualization, to display as many items as possible, and to minimize scrolling.

### 4.2.2 Missing Data

In the original GTD, all data for the year 1993 is missing due to an office move. [4] If the missing data was treated as zero, there would be a sudden gap in 1993, possibly distracting users.

Another possibility was to interpolate numbers based on the previous and following year(s). When using such an approach, it was best to draw the chart according to the interpolated value, but also account for the fact that the data for this year is missing.

In order to let users know that 1993 does not have accurate data, an asterisk was placed along the horizontal axis, where 1993 was supposed to be. Hovering the mouse pointer over this asterisk reveals a text box mentioning that data for 1993 is missing, and the current shown graph is drawn based on interpolated values. This indicates that data in 1993 is dealt specially, without disturbing the overflow of the visualization.

There are also several incidents where a few fields have empty values. Such cases were just grouped as having the value “N/A”.

### 4.2.3 Coloring

Color plays an important role in the visualization. Not only does it let users distinguish adjacent stripes, but can also convey additional information regarding the underlying data. Also, if properly used it can make the visualization aesthetically appealing, more enjoyable. Colors that would convey extra information, but not overwhelm users, were chosen.

When the visualization groups by the country attacked, a similar color was used for countries in the same region. These colors are based on the colors assigned to the stripes in the ‘Region’ view. Within a region, countries with high counts are colored in darker shades. This also makes the interaction between the ‘Region’ and ‘Countries’ mentioned in chapter 3.1.1. smooth.

When meaningful extra information could not be found, a random RGB value was generated as the color of each stripe. The random value was biased to have stronger blue and green values, a relatively weaker red value, and softer hues.

When users hover the mouse pointer over a stripe, the targeted stripe’s border is colored a bright yellow, and it’s width is thickened.

The random color generator is biased to avoid creating a strong yellow. This keeps the highlighting effective.

If adjacent stripes have similar colors, it is difficult to tell them apart. By drawing a thin white border around the stripe, this problem can be lessened. But when the data set is large and several hundreds of stripes have to be displayed, some parts of the chart become crammed, making individual stripes appear indistinguishable. This causes an unpleasant white region to be rendered, due to the dense collection of borderlines.

By assigning low alpha values to borderlines, both of the issues above are addressed. There still is separation between adjacent, similarly colored stripes, but crammed areas still display a hint of their original color. It is also apparent that it is a heavily filled area.

### 4.2.4 Performance

The original GTD raw data is over 60 megabytes in size and comprises more than 75,000 incidents. The loading and processing time for data of this size is significant on a web based program running on a personal computer. Also, as the current working set of data grows, many of the interactions become less and less responsive to user input.

Therefore the data was preprocessed to reduce the loading time. A python program was written that took the original data and counted the number of incidents per year grouped by certain criteria. A separate data file was created for each criteria, and is loaded each time the user asks for the corresponding criteria. This simplifies the

data, trimming it down to a manageable size. This also makes it fairly easy to extend the tool to support different types grouping criteria. But by preprocessing data, much of the detailed information is lost.

Runtime responsiveness is also an issue when dealing with data sets with more than 1,000 items. Updates to the visualization often take several seconds.

Transition animations serve a twofold purpose: changes become more perceivable, and redrawing delays can be masked.

Thin stripes are not displayed. These stripes do not affect the appearance of the chart, yet can severely reduce the speed. The threshold for determining such stripes is based on the relative thickness of the stripe in question with regard to the thickest stripe that is currently displayed.

## 5 Results

Graduate students enrolled in the University of Maryland Information Visualization Spring 2008 course provided feedback on how the tool may be improved. Several graduate students in the Computer Science department of the University of Maryland were asked to report interesting findings in the GTD Explorer, and were observed to see how users use the tool. Suggestions from such users will be discussed in the next chapter.

Researchers from START have been using the GTD Explorer for over six weeks in public presentations, and for their own understanding. They have been relying on the tool to get a better sense of major trends in the GTD data and to explicate to others.

The START group wishes to push this visualization tool out to the intelligence and defense communities, as well as the general public, to gain further feedback. The GTD Explorer is expected to be posted on the START web site. User comments to the page should provide more valuable feedback.

### 5.1 Findings

The tool is especially useful for observing very rapidly the major time-linked structures buried in the data.

For example, many “hotbeds” of terrorism from the 1970s and early 1980s – such as Italy – became far less important after the 1990s, despite the overall increase in terrorism activity during this time.




Figure 11. Italy. There is a sharp decline in the number of incidents in the 80's

Similarly, the visualization tool shows very clearly what a big impact the demise of the Soviet Union had on the nature of global terrorism. Following the collapse of the Soviet Union (1991), there are big declines in many politically oriented Marxist-Leninist terrorist attacks—which are increasingly replaced in more recent years with more religiously oriented jihadi-style groups. (Joonghoon: this came from Gary, I'm not sure how you can see this... I've sent an email asking about it)

The country-level visualization results show that the United States is somewhat unique in terms of terrorist attacks in that it continues to get attacked over the entire period spanned by the data – but by an ever-changing array of groups.




Figure 12. United States.

It is apparent that in many cases, less than 20% of the top items make up for the vast majority of incidents; whether it be country attacked, terrorist organization or target nationality. Figure 13 shows three images with the rank filter set at different ranges. The first image shows the full range of items. The second image shows only the items that are in the top 20% range, and yet the overall pattern is nearly identical. The third image shows the bottom 80% range. It is strikingly clear that the pattern is quite different, and the summed count is much smaller (around 800), at about one-fifth the total count (around 4000). Even a split at 10/90 would show similar results.




Figure 13. Take note at how the filtered range affects the visualization.

## 5.2 Usage

Five users were given the tool and observed. They were given a brief introduction to the underlying data, but were not introduced to the features of the tool. They were given 5 minutes to play with the tool, without asking questions, and were asked a few simple questions regarding the data. Each question was given one at a time.

1) Which country had the most incidents in 1980?

All users easily answered the first question.

2) Which country had the largest number of attacks in 1993?

Two of the users were baffled with the second question, and did not notice that data was missing in 1993. The other users properly used the asterisk to reveal information.

3) Find how many incidents occurred in 1992 in Germany.

For the third question, three users tried to first search for Germany by hovering the mouse pointer over stripes to find the target, be unsuccessful and use the search box. Only one person out of the five people used the search box

first.

*4) What is the total number of incidents, over the full time span, that occurred in Europe?*

For the fourth question, four people answered easily, by switching to the regions view, but one person did not switch to the region view, and tried to add the count of all countries that the person knew to be in Europe. When later asked why, the person responded that he forgot such a feature existed.

*5) Which is the more popular weapon type, Chemical Agents or Explosives?*

Everyone answered the final question easily, even the person who had trouble with the fourth question. Since the question asked for a completely different type of data, it was apparent that some different action needed to be taken.

After the questions, each user was asked what facts seem interesting, and what features seem lacking. Sudden increases in terrorism activity in certain regions or countries particularly caught their interest, such as 1987 and 1988 in Sri Lanka. Users noted that the 1970's data seemed exceptionally small, and questioned what the spikes accounted for. All users were curious of why data in 1993 was missing.

Overall, users found the tool easy to learn and use, and were eager to learn more about the underlying data. One user who had previously been exposed to NameVoyager commented that also she enjoyed using the GTD Explorer, it was not as enticing as the NameVoyager data which she was more familiar with and could associate with people around her. Most users also commented on the response time, suggesting that an increased responsiveness would make the tool more usable.

## **6 Software Improvements**

There are several software improvements that may be made.

In the visualization, stripes that have a wide region may be labeled. When dealing with larger data sets, the GTD Explorer can slow down quite a bit. Performance enhancements must be made to keep up with the growing data set.

One feature that was requested by multiple users was a comparison feature. Comparing several arbitrary items is not supported. One possible solution would be to have checkboxes in each row of the details grid. Users can select or deselect items to make corresponding stripes visible or not.

The details grid is convenient when users are searching for details of a particular item. Once they find this item, the corresponding stripes should be highlighted to lead the user.

When using interactive controls, users often want to undo/redo actions. History storing was a feature many users requested for.

## **7 Conclusion and Future Work**

The GTD Explorer is a web-based visualization of the Global Terrorism Database, a database that contains data over time. It supports mouse and keyboard interactions to explore the data, and is easily extendable to support different types of data. It has a text based search filter, data count based filter, and detailed information grid which gives an overview of the data in text.

The work in this paper is largely a work in progress, suggesting several issues that remain for future investigation.

First of all, these are issues that are related to the use of temporal data. A time-axis filter could be provided to zoom into specific time slots. A Time Searcher [16] style pattern searching feature could be useful in finding similar patterns.

Entity identification and resolution is also a problem. For example, a search for Germany suggests that the country was extremely peaceful until 1990, but a search for West Germany tells a different story. Czechoslovakia, Czech and Slovak or the Soviet Union and Russia would be other examples. Relating these entities is a problem to be solved.




Figure 14. Germany (top) and West Germany (bottom)

Users tend to view temporal patterns in light of real world events, such as the Munich Olympics or September 11 World Trade Center attacks. An advanced version of the GTD Explorer could incorporate meaningful event data to offer a more comprehensive view of the data. Allowing public users to annotate the chart could generate interesting outcome.

Also, when the data being displayed is geographical, it would be useful to have a world map view that shows where the incidents occurred using scattered dots. Heat maps may be used if there is much occlusion.

The GTD Explorer is a tool developed specifically for the Global Terrorism Database. A general framework that is applicable in visualizing time series or temporal data could be developed to aid future users or data providers.

In the future we would like to do an analysis of the social aspects of this tool. Feedback received from web users should be a valuable source in seeing what drives the users to use this tool and how they use it to communicate information.

## 8 Acknowledgements

I appreciate the invaluable guidance of Prof. Ben Shneiderman at the University of Maryland Human Computer Interaction Lab. I am also extremely grateful to Adam Perer for his insight and thoughtful feedback and continued assistance throughout the project.

Many thanks to researchers from START, Gary LaFree, Alex Jonas, Erin Miller, Laura Dugan, Gary Ackerman, Kathleen Smarick. They have provided the Global Terrorism Database and have given opinions as experts on the subject matter. Without either, this project could not have been pursued.

I also thank my fellow graduate students who participated in the user study and provided many useful comments.

## References

- [1] Martin and Laura Wattenberg, *NameVoyager*, “<http://www.babynamewizard.com/namevoyager/>” Accessed May 5, 2008.
- [2] Wattenberg, M., IBM Research, *Baby Names, Visualization, and Social Data Analysis* Proceedings of the Proceedings of the 2005 IEEE Symposium on Information Visualization
- [3] Bloch, M., Byron, L., Carter, S., Cox, A., “The Ebb and Flow of Movies: Box Office Receipts 1986 - 2007”, *New York Times*, [http://www.nytimes.com/interactive/2008/02/23/movies/20080223\\_REVENUE\\_GRAPHIC.html](http://www.nytimes.com/interactive/2008/02/23/movies/20080223_REVENUE_GRAPHIC.html), February 23, 2008
- [4] LaFree, G., Dugan, L., Fogg., Heather, Scott, J., *Building a Global Terrorism Database*, NTIS Selected Reports July 2006 - May 2007
- [5] Viégas, F., Wattenberg, M., Dave, K., *Studying Cooperation and Conflict between Authors with history flow Visualizations CHI 2004*, Volume 6 Number 1.
- [6] Havre, S., Hertzler, B., Nowell, L., *ThemeRiver<sup>TM</sup>: In Search of Trends, Patterns, and Relationships*, <http://infoviz.pnl.gov/pdf/themeriver99.pdf>
- [7] Imrich, P., Mueller, K., Imre, D., Zelenyuk, A., Zhu, W. *Interactive Poster: 3D ThemeRiver*, [www.cs.sunysb.edu/vislab/papers/3DThemeriver.pdf](http://www.cs.sunysb.edu/vislab/papers/3DThemeriver.pdf)
- [8] Ringel, M., Cutrell, E., Dumais, S., Horvitz, E. *Milestones in Time: The Value of Landmarks in Retrieving Information From Personal Stores*
- [9] Fry, B. *revisionist*  
<http://acg.media.mit.edu/people/fry/revisionist>
- [10] Inselberg, A., *The plane with parallel coordinates*. *The Visual Computer*, 1(2):69–92, 1985.
- [11] MacDonald, L. *Using Color Effectively in Computer Graphics*. *IEEE Computer Graphics & Applications*, July/Aug 1999
- [12] Brewer C., Harrower M. *ColorBrewer*. <http://ColorBrewer.org>
- [13] Heer, J., Viégas, F., Wattenberg, M. *Voyagers and Voyeurs: Supporting Asynchronous Collaborative Information Visualization*

- [14] Shneiderman, B., *The eyes have it: A task by data type taxonomy for information visualizations*. Proc. 1996 IEEE, Visual Languages.
- [15] Bartle, R. *Players Who Suit MUDs*, *Journal of MUD Research*, 1:1. 1996
- [16] P., Plaisant, C., Simeone, A., Aris, A., Shneiderman, B., Shmueli, G., Jank, W. *Similarity-Based Forecasting with Simultaneous Previews: A River Plot Interface for Time Series Forecasting*, Proc. of the 11th International Conference Information Visualization (IV '07), 2007, 191-196
- [17] *The flare visualization toolkit*.  
<http://flare.prefuse.org/>
- [18] *The prefuse visualization toolkit*.  
<http://www.prefuse.org/>
- [19] *Open Source Flex Component Library*.  
<http://code.google.com/p/flexlib/>
- [20] *Adobe Flash Player Version Penetration*.  
[http://www.adobe.com/products/player\\_census/flashplayer/](http://www.adobe.com/products/player_census/flashplayer/)  
[http://www.adobe.com/products/player\\_census/flashplayer/version\\_penetration.html](http://www.adobe.com/products/player_census/flashplayer/version_penetration.html)