
Interactive Parallel Coordinates for Collaborative Intelligence Analysis

Jeff Wilson

School of Computer Science
Carleton University
Ottawa, Canada
rjp.wilson@gmail.com

Judith M. Brown

School of Computer Science
Carleton University
Ottawa, Canada
mmjbrown@gmail.com

Robert Biddle

School of Computer Science
Carleton University
Ottawa, Canada
robert.biddle@carleton.ca

Abstract

Large multitouch surfaces are useful platforms for co-located collaborative analysis work, and we have been exploring their use for intelligence analysis. We produced a large multitouch surface prototype to support an intelligence analysis technique called Analysis of Competing Hypotheses (ACH). One challenge was exploration of multi-dimensional data, and we used a parallel-coordinates visualization. In this paper we present our exploration of the issues that arose in applying this visualization in the context of our multitouch collaborative tool.

Author Keywords

Information Visualization, Intelligence Analysis.

ACM Classification Keywords

H.5.2 [Information interfaces and presentation]: User Interfaces - Interaction styles, Input devices and strategies.

Introduction

Large touch surfaces are useful platforms for collaborative analysis work [3] but there is still research to be done before such platforms are widely adopted. We have been collaborating with a government partner who is interested in exploring ways to adopt this technology in service of intelligence analysis.




Figure 1: ACH Walkthrough in use on connected systems: 1) a 44-inch Sony TV with a PQ Labs touch overlay, and 2) a laptop computer.

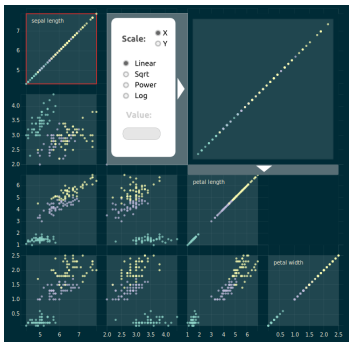


Figure 2: Generalized draftsman's plot (ruled out)

One particular challenge that has arisen in this research is the interactive exploration of multi-dimensional data. In this paper we present one example of an interactive visualization especially customized for large surface displays in a collaborative co-located intelligence task.

The design for the visualization flowed from requirements gathered after a 4-day field study involving real intelligence analysts. A subset of these analysts used a technique called Analysis of Competing Hypotheses (ACH). In answer to an apparent need for tools, we produced a large-surface, web-based prototype [7], (shown in figure 1) to help with this analysis task. Within this prototype we designed a parallel-coordinates visualization (shown in figure 3) that is especially suited to the particular tasks that emerged.

In this paper we present our exploration of the issues that arose in applying parallel-coordinates visualization in the context of our multitouch collaborative tool for intelligence analysis. We present some solutions to these issues that are based on theories of perception and communication.

ACH Walkthrough

Our prototype tool is based on a technique developed in the 1970s by Richards Heuer of the CIA called Analysis of Competing Hypotheses (ACH) [4]. ACH is a structured analysis tool that seeks to reduce cognitive bias in validating analytic theories or conclusions on the basis of available evidence. Heuer's technique follows eight steps:

1. Identify all hypotheses, ensuring mutual exclusivity
2. List significant evidence for and against each hypothesis
3. Produce a *diagnosticity matrix*, evaluating each item of evidence for its value in disproving each hypothesis
4. Refine the matrix, reconsidering hypotheses and deleting evidence with no diagnostic value
5. Draw tentative conclusions by comparing the likelihood of each hypothesis by attempting to disprove each one
6. Test your conclusions for sensitivity to critical evidence items, considering consequences of error
7. Report conclusions, including not just the most likely hypothesis but the relative strengths of each
8. Identify milestones for future observation

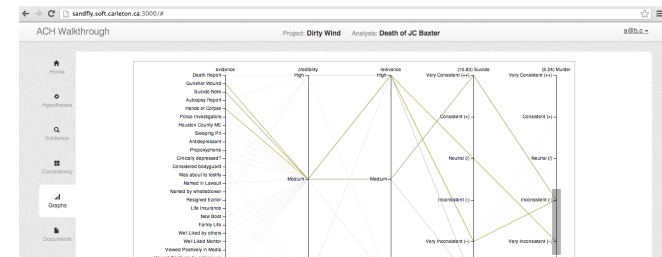




Figure 3: Parallel coordinates visualization in our prototype

Heuer's focus on *disproof* is the technique's greatest strength. The goal is to make an analyst more skeptical of their own favourite theories. Additional confidence in the conclusions comes from the structured evaluation of evidence as it relates to all possible theories.



(a) 3x3 Evidence Rating Widget



(b) Consistency Widget



(c) Widget in context

Figure 4: ACH-W Input Widgets

In our new software called *ACH Walkthrough* (so named for its focus on guiding users new to collaborative ACH) we chose to use interactive visualization as a core feature. Custom visualizations were developed to address the needs of data manipulation and team-based deixis, where team members gesture toward on-screen features and discuss/change their settings with respect to the overall analysis. These visualizations were used for the following tasks:

1. When entering evidence, users rank credibility and relevance on a 3x3 matrix (figure 4a)
2. When evaluating diagnosticity, users rank consistency using a *fretboard* (figure 4b and 4c)
3. When refining the matrix, users employed an interactive parallel-coordinates graph (figure 5)

The parallel coordinates axes come from the following dimensions of the application database:

1. Evidence name from the evidence table
2. Credibility from the Evidence rating widget
3. Relevance from the Evidence rating widget
4. One axis for each hypothesis from the Consistency widget


Without dwelling on the details of ACH, a few factors of the resulting data should be mentioned. The evidence ratings of relevance and credibility are saved on an ordinal scale of Low, Medium, and High. The consistency ratings – answers to the question: “*Is evidence item E consistent or inconsistent with hypothesis H?*” – are similarly rated on an ordinal scale

of: highly consistent (++) , consistent (+) , neutral (/) , inconsistent (-) , highly inconsistent (- -) , and irrelevant or unrated (n/r).


The combination of the two evidence scales and the multiple consistency scales (one for each hypothesis) comprise the multi-dimensional dataset upon which the analysis is based. We sought guidance regarding perceptual aspects from Ware’s book on visualization [6]. There were two principal candidate visualizations that had most potential: generalized draftsman’s plots (figure 2) and Parallel Coordinates [5] graph. We ruled out the first option on a couple of critical problems: it was difficult to do more than pairwise comparisons, and the data points did not trace back easily to the evidence text essential to the analysis.

We chose to visualize this dataset with a parallel coordinates graph (hereafter referred to as par-coords). We leveraged examples provided by the creator of the D3 [2] JavaScript visualization library, and the entire application was developed using the Meteor application development framework [1]. This toolset offered a solid foundation for the live rendering required for interactive collaboration.


Figure 5a shows the first version of a par-coords plot for ACH-W, and an important issue should be immediately apparent. The problem can be seen when examining the number of lines between the first and second axes (left to right) and the apparent loss of detail as lines in subsequent gaps are rendered one on top of another. This issue results from the fact that the data points are not floating point values but instead are categorical (the first axis) and ordinal (the remaining axes). This loss of information can be corrected by



(a) Original version



(b) Correcting for non-floating data



(c) Enlarged view showing improvement with curves

Figure 5: ACH-W Parallel Coordinates

using curved lines, as shown in figures 5b and 5c.

With this problem corrected, we can begin to see the value of colour in the overall process in ACH-W (we assume you are reading this in colour). The colours from the *Evidence Rating Widget* carry over into the par-coords graph. This can be seen most clearly by concentrating on the middle axis in figure 5b. Values at the top of the axis are in shades of orange and brown, values in the middle are shades of green, and values at the bottom are shades blue. Users familiar with this coding scheme will recognize this as relating to relevance, for example bluer lines are less relevant to the overall question than lines rendered in greens and oranges.

This use of colour offers an overall interpretive cue for analysis, but the real power of the tool comes from interacting with the graph's axes. Selecting sub-ranges of the axes using *brushing* allows the user to focus on meaningful subsets of a dimension's overall range. Lines that fall within these highlighted regions retain their brighter colours, while lines that fall outside the regions are rendered in muted grey. This operates as a filter that permits easier inspection of features of interest.

The brushed ranges are retained until deliberately removed. This allows the user to add additional brushes on other dimensions, effectively narrowing their query with a great deal of flexibility. For example, figure 6 is a spotlighted section from a larger graph. It shows the effect when seeking only highly credible and highly relevant evidence.

Brushing supports a surprising range of interaction tasks, especially as users become familiar with the

meaning of the graph's dimensions. Users new to par-coords graphs might at first be drawn to visual clusters and reinforcing trends across the display, and indeed in many domains this is a strength of analysis work like ACH, the real power comes from drawing one's attention to *individual* evidence items that fall within meaningful regions of the graph and then taking the time to consider one's evaluations from fresh perspectives.

Heuer's technique of ACH was designed to draw attention to highly *diagnostic* evidence when the natural human tendency is to explore evidence that confirms our suspicions. He explains diagnosticity with an example taken from medicine. A junior intern recently swamped by flu season might incorrectly diagnose flu again based on yet another patient complaining of fever, whereas a more experienced resident would know that fever is consistent with too many conditions for it to offer much diagnostic utility and they would ask more probing questions. Their focus would be on identifying symptoms that have the power to *rule out* competing hypotheses, rather than those that confirm the diagnosis they most easily suspect.

Interaction with par-coords supports this kind of diagnostic reasoning by making it easy to select items that help rule out a given hypothesis. The user can create a brush that selects for ratings of *inconsistent* or *very inconsistent* along a particular hypothesis axis and then draw their attention to the evidence associated with only the highlighted lines.

Evaluation and Design Implications

Informal usability sessions with a handful of users revealed opportunities for refinements of interactive features.

Our experience with using this technique has shown it to have important analytic value, but it does demand a certain level of patience and attention to detail. For users new to parallel coordinates in ACH, the availability of brushing is not immediately obvious. It may also go against their expectations if they assume that the visualization is merely a static aggregation of data.




Figure 6: Selecting for high credibility and relevance




Figure 7: Mockup of possible affordances.

In our original prototype there weren't any obvious affordances for a newcomer. In fact there was only one type of affordance offered to mouse users, namely a change of the cursor's appearance when hovering over an axis. Of course this feature assumes that hovering even takes place. Users of our touch interface lacked the ability to hover, and so they missed out on interaction cues altogether.

The mockup in figure 7 shows an enlarged view of one axis where we highlight the concept of touch tabs to assist in brushing. These affordances would hopefully also have the effect for new users of simply cueing them to the notion of brushing in general.

Another possible enhancement might be to offer a list of pre-set brush selections based on ACH-specific tasks (e.g. filter irrelevant items, confirm diagnostic items for hypothesis n , then $n + 1$, etc.) and then instruct the user to walk through each of these presets as an introduction to brushing. Also, the initial rendering of the par-coords graph could briefly show animated selections on each axis that quickly unfold

until they encompass their full range, leaving behind affordances for the user to adjust.

A final source of confusion revealed in our usability sessions was that multiple brush filters operate as an implied logical AND condition. In other words, when more than one axis had a sub-range selected, items that did not satisfy *all* filters were suppressed. In an informal trial with intelligence analysts, some users had assumed a logical OR condition would apply, and they expressed surprised at some of the results when multiple filters were active. A simple solution to this issue would be to offer an AND/OR toggle with an immediate re-rendering of the graph.

A somewhat advanced solution would be to give the user complete control over the query logic implied by brushing. For example we could allow users to save and name individual brush selections (filters) and then combine them using a separate expression editor. This editor could be a touch-based graphical formalism that could employ a tree structure for nested logical expressions, perhaps accompanied by a syntactical form of the entire expression in a text editor. A mockup of a graphical logic editor is shown in figure 8.




Figure 8: Mockup of logic editor.

Discussion

The expressive power of purely static multi-dimensional visualizations is already quite powerful. When we add interactivity we find opportunities for even greater discovery of knowledge. But not all data sources are equivalent, and not all multi-dimensional visualizations are a good fit for users' tasks.

Even within reasonable spatial limits, the question arises as to whether interaction should be unconstrained or guided. In our research we see the value of offering such actions as filtering, sorting, categorizing, and mathematically transforming raw data, but access to live computation offers the potential to support a much richer set of operations.

But there is a risk in attempting to build a comprehensive, general purpose tool containing too many data manipulation and rendering features. There are only so many potential operations and rendering features that would be useful. Beyond a reasonable subset of features, the best we might do would be to provide access to JavaScript itself. This would offer the most intrepid data miners the tools they might need, and it would also keep the interface within reach of moderately trained users. We note that such richness might be valuable where one might reasonably assume that user-driven design features could pay dividends (such as when exploring unfamiliar data), but in our example with ACH we see good reason for a more domain-specific approach to exploration.

The ACH data model imposes a fairly rigid set of constraints, and even within this structure users are faced with a significant cognitive challenge. We are, however, considering two potentially useful future additions. We believe that the capture and presentation

of interaction history may help foster more exploration and free users from the risk of false belief based on mere effort, and we believe that the addition of a touch-based logic editor might invite users to consider new ways of thinking in visual form. In our future work we plan to modify ACH-W to support features such as these and to validate their utility with lab experiments.

References

- [1] Meteor. <http://meteor.com/>.
- [2] Bostock, M., Ogievetsky, V., and Heer, J. D³ data-driven documents. *Visualization and Computer Graphics, IEEE Transactions on* 17, 12 (2011), 2301–2309.
- [3] Brown, J., Wilson, J., Stevenson, G., Hack, C., and Biddle, R. *Surface Computing and Collaborative Analysis Work*. Claypool and Morgan, 2013.
- [4] Heuer, R. *Psychology of Intelligence Analysis*. Psychology of Intelligence Analysis. Center for the Study of Intelligence, 1999.
- [5] Inselberg, A. Multidimensional detective. In *Information Visualization, 1997. Proceedings., IEEE Symposium on*, IEEE (1997), 100–107.
- [6] Ware, C. *Information Visualization, Third Edition: Perception for Design*, 3rd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2012.
- [7] Wilson, J., Brown, J. M., and Biddle, R. ACH walkthrough: A distributed multi-device tool for collaborative security analysis. In *First Workshop on Security Information Workers*, ACM SIGSAC (2014).