

Interactive Visualization for Information Analysis in Medical Diagnosis

B.L. William Wong¹, Kai Xu¹, and Andreas Holzinger²

¹Interaction Design Center, School of Engineering and Information Sciences,
Middlesex University, London, UK, {w.wong, k.xu}@mdx.ac.uk

²Research Unit Human-Computer Interaction, Institute for Medical Informatics,
Statistics and Documentation (IMI), Medical University Graz, Graz, Austria,
andreas.holzinger@medunigraz.at

Abstract. This paper investigates to what extent the findings and solutions of information analysis in intelligence analysis can be applied and transferred into the medical diagnosis domains. Interactive visualization is proposed to address some of the problems faced by both domain. Its design issues related to selected common problems are then discussed in details. Finally, a visual sense making system INVISQUE is used as an example to illustrate how the interactive visualization can be used to support information analysis and medical diagnosis.

Keywords: Visualization, Visual Analytics, Medical Diagnosis

1 Introduction

In this paper we will briefly compare the similarities between information analysis in intelligence analysis and the medical diagnosis domains, and hence to draw from our work in intelligence and then examine how it might be applied to the medical domain. Our understanding of the cognitive processes in information analysis suggest that it is more than just the process of information search and retrieval, but incorporates a number of other features that is now characterized as *sense-making*. The more frequently cited Pirolli and Card model [15] of intelligence analysis while useful in helping us see the different cycles (i.e., foraging, hypothesis formulation and testing), can be complemented by Klein et al.'s Data-Frame model [11] which describes the process of creating plausible explanations for observed data. This similarity allows us to consider our work in the intelligence, legal investigation and e-discovery domains, in the context of medical diagnostic analysis, particularly in the area of reviewing a patient's medical history for the purpose of developing treatment plans. From our work in designing interactive visualizations for information analysis, we documented a list of 20 design problems [19], and will discuss five problems that we believe have a bearing in designing medical diagnostic displays that can assist in improving the review of a patient's medical history. We will then discuss the design issues and a number of possible designs currently under consideration in the context of a visual sense making system INVISQUE [20].

2 Information Analysis in Intelligence and Medical Domains

In trying to assess the level of maturity of intelligence analysis as a discipline, Fisher and Johnston [1] provide a brief account of the similarities between the process of medical decision-making and clinical judgement, with intelligence analysis. In their review they also reported on the beginnings of evidence-based medicine (EBM) and how EBM, with the use of statistical information, help make clinical judgement and treatment more systematic and able to draw on evidence of past cases, rather than on limited experience or anecdotal information.

But are they really similar? Let's take a look at what makes intelligence analysis difficult. There are many problems, but we'll briefly discuss those that are more relevant to the medical community (e.g. we will not focus on data that has been created for deception and to mislead an analyst).




Fig. 1. Illustration of the problems faced by an information analyst

Fig. 1 is a basic illustration of the information challenges facing intelligence analysts. Analysts often work alone, and are required to make sense out of large data sets that come from different sources and in different formats, and are often of varying quality and reliability. The information may be incomplete, out of sequence, changing as the situation changes, and misleading. For the analyst peering through the tiny viewport of his or her computer display to this very large data space, is like a person peering through the keyhole of a door to an enormous hall. To gain a sense of what the hall is like, the analyst must piece together his or her different views from memory, which can lead to many problems associated with memory and cognitive limitations, attention, and bias.

This can lead to errors of re-construction or simply forgetting what was seen previously.

In addition, the data that analysts are often presented with, is akin to a jig-saw puzzle, where the pieces have been emptied onto the table top, and the analyst has to find, organize and join the relevant pieces together. To complicate the problem, the box tops of the jig-saw puzzle, which provides the big picture of the puzzle, are almost always not available, i.e. there is no context in which to view the pieces of information. The analyst instead has to build the picture as he or she carries out the analysis. To further aggravate this already difficult situation, an analyst is often presented with not one jig-saw puzzle, but several at the same time, and each without its box top, where the puzzles may be related or may have absolutely nothing to do with one another.

Then, guided by their training, expectations, beliefs, goals, socio-cultural factors and other background factors, they will create frames that help piece together the information to create explanations or narratives that is able to account for what they have observed [11]. In the process of creating these data-frame relationships that help them make sense of the data, the analysts are looking for patterns, underlying relationships, and triggers in the data, that help them collate evidence to support possible explanations (or ‘hypotheses’, in the social science, and sometimes in the scientific sense of the word), in order to come to a conclusion.

In a public health context, there is probably a large similarity in trying to identify the source of an outbreak of an infectious disease in a populated area, and depending on various conditions, the situation could evolve very rapidly. Evidence will be collated from different sources such as news items, hospital and doctors reports, laboratory results, and so forth. In order to diagnose a patient’s medical situation, it may be necessary for the doctor, or in more complex cases, the medical team to collate and review the patient’s medical history, which can be long and complicated, e.g. a geriatric patient with a history of acute glaucoma, hypertension, arthritis and joint pains, liver sclerosis, may require treatment for breathing difficulties. Each of these areas may have been treated separately with records held in different specialist clinics.

3 Sense-Making and the Data-Frame Model

Pirolli and Card [15], based on a cognitive task analysis of intelligence analysts, explains the intelligence process as one primarily of search: searching for information, relationships, evidence to formulate or support a hypothesis; and in the process of searching, the analyst carries out a number of other processing actions as well, such as reading and extracting relevant or meaningful information, indexing and filing away data, creating schema by organizing and re-organizing the data in order to create understanding and insight that can lead further actions for building the case for the formulated hypothesis, and to then subsequently create a narrative that tells a story.




Fig. 2. Sense-making loop for intelligence analysis [15]

This model has been used and referred to considerably by analysts and researchers. It can however be usefully complemented by Kleins et al. [9] Data-Frame Model of Sense-making. The Data-Frame model basically explains that an analyst upon observing a set of data, will attempt to create an explanation for this data that they refer to a frame. The creation of this frame is contingent upon the analysts background such as, training and experience, goals and expectancies. Together with the data, this frame will determine how the data is combined and used to explain what has been observed. The frame also determines what the analyst will attend to, what data should be related, or filtered. Klein et al. [9] then describes additional strategies used to help the analyst to gain a deeper understanding by elaborating and filling in gaps or seeking more information, to query one's own assumptions and beliefs and in particular how earlier data has been used to generate the explanations, judging the plausibility of the arguments or narratives, and the quality of the data at the same time. There may be times when the data and the frame are so badly mis-matched that it requires the analyst to re-evaluate his frames, and to possibly revamp these frames or even seek a new frame. There is no real sequence in the process. Instead, it shows the variety of strategies that may be invoked to make sense of the data and one's frame and to assess the plausibility of the offered accounts.

In addition, Klein et al. [10] also explains that people are also engaged in another form of reasoning which they have also observed occurring in naturalistic environments - causal reasoning. While causal reasoning is characterized by the determination of causes for observed effects, this causality can sometimes be confused with correlation as events can co-occur, while not being the cause of the observed effects. Also, another characteristic of causality is mutability, or the




Fig. 3. The Data-Frame Model of Sense-making [11]

ability to engage in a reasoning strategy that allows one to investigate or imagine what might have happened if one or all of the causal factors did not occur or could be reversed, and this can also be used as a test for causality. Josephson and Tanner [8] explain another form of reasoning that is useful: abductive reasoning. Such methods of reasoning is akin to the strategies used by archaeologists, where based on fragments of evidence dug up from the ground, together with other known facts and history, enables them to piece together convincing accounts about life and civilization. This abductive inferencing method is also used in information analysis and medical diagnosis, e.g. given signs and symptoms of a disease, doctors are expected to infer the type of illness and therefore treat it.

4 Interactive Visual Sense-Making Design

In the course of our work in developing visual representations with which to represent datasets, to carry out and to report on analysis carried out on them, and to reason with the data, we have identified a list of 20 user problems that require new or better techniques for their visual representation.

These areas, though described as problems, represent areas where Information Analysts can benefit from more advanced science and technology [19]. The problems have been gathered and condensed from across several studies, and a series of interviews and focus group discussions with researchers and students interacting with library electronic resource systems, legal investigators, and information analysts.

1. The problem of seeing a large data set and reasoning space through a small keyhole.

2. The problem of handling missing data.
3. The problem of handling deceptive / misleading data.
4. The problem of handling contradictory data.
5. The problem of aggregating and reconciling multiple points of view or predictions.
6. The problem of evidence collation and evidential reasoning.
7. The problem of provenance and tracing analytic reasoning.
8. The problem of integrating data space, analytic space and hypothesis spaces.
9. The problem of handling strength of evidence (including subjective and objective measures of strength) + contribution of different pieces of evidence to a conclusion.
10. The problem of handling uncertainty in data and / or information.
11. The problem of representing and handling evidence over time and space.
12. The problem of annotating, remembering, re-visiting, and setting aside.
13. The problem of developing a sense of what is in the data exploring what is there.
14. The problem of predicting and representing emergent behaviour.
15. The problem of Identifying and representing trends.
16. The problem of recognising and representing anomalous changes.
17. The problem of finding the needle in the haystack (or knowing what is chaff i.e. info of no or low value)
18. The problem of predicting the path of cascading failures or effects.
19. The problem or representing the static and dynamic relationship between the data / information.
20. The problem of scalability and reusability.

In the following sections, we hope to explain some of these problems in the context of medical domain.

4.1 Aggregating and Reconciling Multiple Views or Predictions

This problem occurs particularly when analysts have to work together, and where their efforts need to be coordinated, while valuing independent inputs from the respective analysts. Some points of view may be very divergent. What is crucial in representing these differences in opinions or predictions, is not the differences in themselves, but rather the trace of the analytic reasoning process, i.e. how did one get to this conclusion? It should show or reveal how the different analysts have used the data and how the way they used the data contributed to the conclusions. In this way, it then becomes possible for a reviewer to seek out area of potential errors or errors of judgement with the given data.

4.2 Handling Evidence Strength and Contribution to a Conclusion

Unlike intelligence analysis, doctors have access to statistical indicators showing adverse reactions to particular medications or the susceptibility of, say, different types of people to certain diseases. In Evidence Based Medicine, such information

is sought to provide a base-line from which to evaluate the likelihood of observed signs and symptoms relating to particular diseases. What is needed in intelligence analysis are schemata (ways in which data may be structured and represented for further analysis) which help make obvious the reliability, quality or likelihood of occurrence in a given context, and their ability to show how their usage can lead them to various logic traps and other flaws such as false positives.

4.3 Annotating, Remembering, Re-visiting and Setting Aside

We are often not able to remember the myriad of small decisions we made along the process of a complicated analysis. There are also times where we use storytelling techniques to fill in missing data in the collection. Analysts as well as doctors (who see many patients often in short period of time), need to annotate for their own remembering purposes, as well as for a trace for other doctors or medical personnel to follow-up on the treatment. Sometimes the data is non-conclusive, or sometimes, data may have some use later, but the analyst or doctor may not want to re-create the search for that piece or collection of data and would like to set it aside, possibly with an annotation, for later use.

Complexity is the main problem in the medical domain, because most of the medical data is weakly structured or even unstructured and there is always the danger of modelling artefacts, which can then lead to wrong decisions. Let us look at standard medical documents for example: The broad application of enterprise hospital information systems amasses large amounts of medical documents, which must be reviewed, observed and analysed by human experts [3] (Kreuzthaler et al., 2011). All essential documents of the patient records contain at least a certain portion of data which has been entered in non-standardized format (wrongly called ‘free-text’) and has long been in the focus of research. Although such text can be created simply by the end-users, the support of automatic analysis is extremely difficult [2, 5, 13].

So, it is very likely that some interesting and relevant relationships remain completely undiscovered, due to the fact that the relevant data are scattered and no investigator is able to link them together manually [16, 4]. Consequently, there are a lot of relevant open research issues at the intersection of HCI and IR/KDD to help (medical) professionals to identify and extract useful information from data.

4.4 Developing a Sense of What is in the Data

One problem at the start of any investigation or review occurs when the analyst or doctor is presented with a large set of data, and he or she has to make sense of it. How does one know where to start if one does not know what is in the data set? Or at least, what are the main categories or methods of organization of the data? Tools are needed to summarise the data set in various ways that lend themselves to rapid exploration. Various forms of semantic maps of information clusters have

been used to show groups, group densities, group peaks, and relationships between and within groups (e.g. IN-SPIRE, <http://in-spire.pnnl.gov/>), with software tools that facilitate drill-downs as well as other methods of analysis.

A good example of a data intensive and highly complex microscopic structure is a yeast protein network. Yeasts are eukaryotic micro-organisms (fungi) with 1,500 currently known species, estimated to be only 1% of all yeast species. Yeasts are unicellular, typically measuring 4 μm in diameter. The first protein interaction network was published by [6]. The problem with such structures is that they are very big and that there are so many. A great challenge is to find unknown structures (structural homologies, see e.g. [7]) amongst the enormous set of uncharacterised data. Let us illustrate this process with a typical example from the life sciences: X-ray crystallography is a standard method to analyse the arrangement of objects (atoms, molecules) within a crystal structure. This data contains the mean positions of the entities within the substance, their chemical relationship, and various others and the data is stored in a Protein Data Base (PDB, <http://www.rcsb.org/pdb/>). This database contains vast amounts of data. If a medical professional looks at the data, he or she sees only lengthy tables of numbers.

However, by application of a special visualization method, such structures can be made graphically visible and the medical professionals can understand these data more easily and most of all they can gain knowledge—for instance, it may lead to the discovery of new, unknown structures in order to modify drugs, and consequently to contribute to enhancing human health. The transformation of such information into knowledge is vital for the prevention and treatment of diseases [17, 18].

To demonstrate that not only natural processes have such structures there is a nice example (http://datamining.typepad.com/data_mining/2007/01/the_blogosphere.html) which shows a visualization of the blogosphere (cf. also with [12]): The larger, denser part of the blogosphere is characterized by socio-political discussion the periphery contains some topical groupings. By showing only the links in the graph, we can get a far better look at the structure than if we include all the nodes.

4.5 The problem of Identifying and representing trends

The final problem is that of identifying and representing patterns in data as well as key trends over time, and whether there are correlating effects of those trends. While it is possible to show trends and patterns in quantitative data relatively easily, how do we reveal patterns in visual forms about qualitative data that the human perceptual system can readily discern?

Maimon & Rokach state in their book [14]: “Knowledge Discovery demonstrates intelligent computing at its best, and is the most desirable and interesting end-product of Information Technology”. Whereas this is true, using intelligent computing is necessary but not sufficient: Computers are (still) Von-Neumann machines and not endowed with any insight, and possess little knowledge of the real-world on which to check whether and to what extent the concepts they are

examining are worthwhile or useful. Consequently, the challenge is to enable effective human control over powerful intelligent machine services and to integrate statistical methods and information visualization, so as to support human insight, breakthrough discoveries, and bold decisions primary research objectives in the field of Human-Computer Interaction.

A further challenge is based on the fact that only a small percentage of data is structured most of the data is semi-structured, weakly structured or even unstructured. A common misconception is to confuse structure with standardization. While the closely related fields of IR/KDD have developed wonderful intelligent (semi)automatic processes and algorithms to extract useful knowledge from rapidly growing amounts of data, these methods fail when data are weakly structured. The problem is that we are faced with the danger of modelling artifices without being aware of it and this may lead to wrong decisions. One solution is to raise the quality of information, while at the same time make the medical professionals aware of the value of information quality; a possible solution is in a systematic documentation. That means that all treatment relevant data are collected in a quality process oriented manner. Most of all it must be possible to condense the data into information as a function of time to visualize it as longitudinal data; the visible patterns and trends can be used to make decisions and to meet predefined treatment goals e.g. in order to provide individualized treatment.

5 INVISQUE

INVISQUE [20] is an interactive visualization system designed to visual sense-making. It aims to address the challenges discussed earlier. Presented here are some preliminary results and future features. It is domain independent and can be transferred easily to biomedical domain.

INVISQUE is designed around a metaphor of physical index cards on a two-dimensional infinite canvas workspace. This is a departure from the traditional 1-dimensional list-style interfaces (such as Google), and the cards present basic information about each result. Rather than relying on static text boxes for input, INVISQUE allows the user to start a new search anywhere on the canvas. This is done simply by clicking on the white space and typing in the search term. Each set of search results are grouped into a cluster. Fig 4 shows an example of INVISQUE working with a library database and displaying the result of two searches: “energy” and “heating”.

By default, the index cards are ordered in both the x and y axes. The ordering attributes are domain dependent and can be set by the user. In the example of searching for journal articles, the y axis can represent the number of citations and the x axis can represent the date of publication. This provides the capability to interactively identify the trend within the data on the selected dimensions. By clicking on an index card, users are able drill down to find more information. In the case of a document search, users can view the content of the document.




Fig. 4. INVISQUE showing library database search results.

The infinite canvas allows users to visualize multiple searches (or clusters) simultaneously (see Fig. 4). This is a step away from traditional tabbed-browsing, and allows users to make visual comparisons between multiple search results. This capability will be further developed to facilitate visually aggregating and reconciling multiple points of view or predictions. Although we are only in a 2D space, the use of transparency creates a series of layers (see Fig. 4). The active search (i.e., ‘heating’) is opaque, giving it the impression of the closest layer and, therefore, the main focus. Remaining clusters are semi-transparent, giving the impression they are in the background, providing context to the active search.

INVISQUE has a few features to support ‘annotating, remembering, re-visiting, and setting aside’. Users can save an index card for later use or mark one as important. These are achieved by dragging the index card to the specific circles in the corners (dashed circles in Fig. 5). Its colour will then change to indicate it is saved (yellow) or marked (green). Users can easily invoke Boolean operations by dragging and dropping. For example, two clusters can be merged by dragging one cluster title on top of the other.

Currently text analysis functions are being integrated into INVISQUES. Once completed, INVISQUE will be able to extract significant phrases (i.e., popular topics) from a collection of documents. This will address the problem of ‘developing a sense of what is in the data’ by visually presenting the significant topics and their relationships. Another new feature being added to INVISQUE is provenance, which is the conclusion pathway that records the information about the reasoning process from the raw data to final conclusion. Part of the provenance is the information of the strength of evidences and how they contribute to a conclusion, which is another problem discussed earlier.




Fig. 5. Marking index cards.

6 Conclusions

In this paper we demonstrated the similarity between the information analysis during intelligence analysis and medical diagnosis. Based on the Sense-Making Loop and Data-Frame model, we discussed the key issues that need to be addressed when designing an interactive visualization system to support such information analysis. A visual analytics system INVISQUE is used to demonstrate the application of some of these design principles.

References

1. Fisher, R., Johnston, R.: Analyzing Intelligence, chap. Is intelligence analysis a discipline?, pp. 55–68. Georgetown University Press (2008)
2. Gregory, J., Mattison, J.E., Linde, C.: Naming notes: transitions from free text to structured entry. *Methods Inf Med* 34(1-2), 57–67 (Mar 1995)
3. Holzinger, A., Geierhofer, R., Errath, M.: Semantic Information in Medical Information Systems - from Data and Information to Knowledge: Facing Information Overload, pp. 323–330 (2007)
4. Holzinger, A., Geierhofer, R., Mdrtscher, F.: Semantic information in medical information systems: Utilization of text mining techniques to analyze medical diagnoses. *Journal of Universal Computer Science* 14(22), 3781–3795 (2008)
5. Holzinger, A., Kainz, A., Gell, G., Brunold, M., Maurer, H.: Interactive Computer Assisted Formulation of Retrieval Requests for a Medical Information System using an Intelligent Tutoring System, pp. 431–436. Charlottesville (VA): AACE (2000)
6. Jeong, H., Mason, S.P., Barabasi, A.L., Oltvai, Z.N.: Lethality and centrality in protein networks. *Nature* 411(6833), 41–42 (May 2001)

7. Jornvall, H., Carlstrom, A., Pettersson, T., Jacobsson, B., Persson, M., Mutt, V.: Structural homologies between prealbumin, gastrointestinal prohormones and other proteins. *Nature* 291(5812), 261–263 (May 1981), <http://dx.doi.org/10.1038/291261a0>
8. Josephson, J.R., Tanner, M.C.: *Abductive Inference: Computation, Philosophy, Technology*, chap. Conceptual analysis of abduction. Cambridge University Press (1996)
9. Klein, G., Moon, B., Hoffman, R.: Making sense of sensemaking 2: A macrocognitive model. *Intelligent Systems, IEEE* 21(5), 88–92 (2006)
10. Klein, G., Mueller, S., Rasmussen, L., Hoffman, R.: Naturalistic model of causal reasoning: Developing an experiential user guide (eug) to understand fusion algorithms and simulation models. Tech. Rep. AFRL-RH-WP-TR-2011-0018, Air Force Research Laboratory 711th Human Performance Wing, Human Performance Directorate, Wright-Patterson Air Force Base, OH 45433 (2010)
11. Klein, G., Phillips, J.K., Rall, E.L., Peluso, D.A.: A Data/Frame Theory of Sense Making. In: *Expertise out of context: proceedings of the sixth International Conference on Naturalistic Decision Making*. pp. 113–155 (2003)
12. Leskovec, J., McGlohon, M., Faloutsos, C., Glance, N.S., Hurst, M.: Patterns of Cascading Behavior in Large Blog Graphs. In: *Proceedings of the Seventh SIAM International Conference on Data Mining* (2007)
13. Lovis, C., Baud, R.H., Planche, P.: Power of expression in the electronic patient record: structured data or narrative text? *International Journal of Medical Informatics* 58-59, 101–110 (2000)
14. Maimon, O., Rokach, L.: *Data Mining and Knowledge Discovery Handbook*. Springer, 2nd ed. edn. (Oct 2010)
15. Pirolli, P., Card, S.: The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In: *Proceedings of International Conference on Intelligence Analysis* (2005)
16. Smalheiser, N.R., Swanson, D.R.: Using arrowsmith: a computer-assisted approach to formulating and assessing scientific hypotheses. *Comput Methods Programs Biomed* 57(3), 149–153 (Nov 1998)
17. Wiltgen, M., Holzinger, A.: Visualization in bioinformatics: Protein structures with physicochemical and biological annotations. In: Zara, J., Sloup, J. (eds.) *Proceedings of Central European Multimedia and Virtual Reality Conference*. pp. 69–74 (2005)
18. Wiltgen, M., Holzinger, A., Tilz, G.P.: Interactive analysis and visualization of macromolecular interfaces between proteins. In: *Proceedings of the 3rd Human-computer interaction and usability engineering of the Austrian computer society conference on HCI and usability for medicine and health care*. pp. 199–212. Springer-Verlag, Berlin, Heidelberg (2007)
19. Wong, B.W., Varga, M.: Blackholes, keyholes and brownworms: Challenges in sense-making. In: *11th NATO Networks of Expert Workshop, Visual Analytics and Network Operations and Health* (2011)
20. Wong, W., Chen, R., Kodagoda, N., Rooney, C., Xu, K.: INVISQUE: intuitive information exploration through interactive visualization. In: *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems*. pp. 311–316. CHI EA '11, ACM, Vancouver, BC, Canada (2011)