A Visualisation Framework for Design and Evaluation

Benjamin J. Blundell[a], Gary Ng [b] and Steve Pettifer[a]

[a]Advanced Interfaces Group, The University of Manchester, UK;
[b]Network Inference Inc. 5963 Laplace Court, Carlsbad, California. USA

ABSTRACT

The creation of compelling visualisation paradigms is a craft often dominated by intuition and issues of aesthetics, with relatively few models to support good design. The majority of problem cases are approached by simply applying a previously evaluated visualisation technique. A large body of work exists covering the individual aspects of visualisation design such as the human cognition aspects visualisation methods for specific problem areas, psychology studies and so forth, yet most frameworks regarding visualisation are applied after-the-fact as an evaluation measure.

We present an extensible framework for visualisation aimed at structuring the design process, increasing decision traceability and delineating the notions of function, aesthetics and usability. The framework can be used to derive a set of requirements for good visualisation design and evaluating existing visualisations, presenting possible improvements. Our framework achieves this by being both broad and general, built on top of existing works, with hooks for extensions and customizations. This paper shows how existing theories of information visualisation fit into the scheme, presents our experience in the application of this framework on several designs, and offers our evaluation of the framework and the designs studied.

Keywords: framework, visualisation, design, evaluation, requirements

1. INTRODUCTION

To date, few frameworks exist that cover visualisation design in both breadth and depth. The methodology is still evolving as advances are made in areas such as human cognition, human computer interaction and software engineering. Our approach presents a unified vocabulary and a set of guidelines for stating and deriving visualisation requirements respectively.

The aim of our framework is to help engineers in identifying:

- Where graphics should be used within an analytical process
- What the graphics should represent (raw versus derived data)
- How to construct the appropriate graphics, given a set of tasks, goals and data
- How the graphics relate to interactive controls

The value of our framework may be summarised as follows:

- Structures the design process and increases decision traceability
- The same framework can be used for design and evaluation
- Overhauls and is complimentary to the existing body of work
- Extensible depending on individual project demands

An overview of the framework is described in Section 2. Sections 3 reports two case studies where the framework has been applied and the experience of using the framework. Section 4 relates existing work to this framework and finally Section 5 concludes with a discussion of future work.

Further author information: (Send correspondence to Benjamin Blundell)
Benjmain Blundell.: E-mail: blundell@cs.man.ac.uk, Telephone: +44 (0)161 2756266
Gary Ng.: E-mail: gary.ng@networkinference.com
Steve Pettifer.: E-mail: srp@cs.man.ac.uk, Telephone: +44 (0)161 2756259

Visualization and Data Analysis 2006, edited by Robert F. Erbacher, Jonathan C. Roberts, Matti T. Gröhn, Katy Börner,
Proc. of SPIE-IS&T Electronic Imaging, SPIE Vol. 6060, 606008, © 2006 SPIE-IS&T · 0277-786X/06/$15

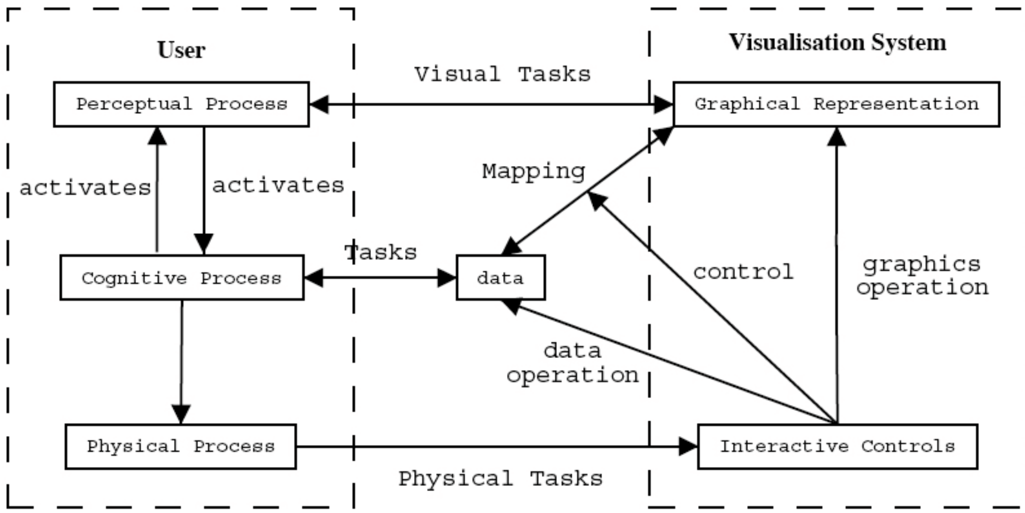

Figure 1. A Conceptual overview of Visualisation Components.

2. THE FRAMEWORK

Visualisation is part of an information analysis process. Visualisation aids this process by minimising cognitive load, maximizing the externalisation of information, and reducing information access cost. It is easy to apply existing graphical techniques onto the entire set of raw data in an aesthetic way, in the hope that something useful will result. The purpose of the graphics is implicit, and the usability of the eventual system is unknown.

A more focused approach targets only the information pertinent to the task at hand, presenting to users information at the right level of detail to complete the task. In addition, much of the data analysis tasks could be done programmatically without any graphics. To avoid dead ends and unnecessary work it is important to identify where graphics can add the most value for the user at the requirements analysis stage.

Visualisation can be seen as a form of human computer interaction, which is often a cyclic process involving user decisions then computer feedback prompting further user decisions. Our framework reflects this aspect insofar as the process of obtaining the requirements for designing a visualisation solution is also cyclic. At the highest level, effective visualisation is an outcome of proper interaction between users tasks, data characteristics, interactive controls and graphical techniques.[7,8,11,14,17] We identify ten relationships between these concepts, summarised in Figure 1.

To simplify the approach we consider these relationships to form two problems within visualisation design, given an analytical goal:

- Where graphics should be used and on what data
- How to use the graphics with interactive controls to meet the goal

This framework interprets visualisation design as a process of iterative refinement of answers to the two aspects, evolving towards a solution. The former can be realised by a task analysis style approach, where data, users and tasks are considered. The latter can be answered by understanding how different symbols convey information to the user through visual implications, and how interaction can play a significant role through graphical/data operations.

2.1. The "Where" of Visualisation Design

The characteristics of data, users, and the tasks to be performed form the initial core of the framework as captured in Figure 1. To identify where visualisations can help a user, the first step is to understand what the user is trying to achieve.

2.1.1. Data Characteristics

Understanding the structure and properties of data will help to determine the type of visual encodings required to give the desired visual implications for given tasks.

Taxonomies detailing the various aspects and characteristics of data have been documented; the relevant sections are mentioned here. Data may be atomic or composite,[18] the latter being divided into set and structure types. In a set, each component is considered only in relation to the set and not among the components. In a structure, relationships among components are important. Under these terms, a relational database is therefore a composite object consisting of a structure (schema) of data sets (tables). Bertin[2] identified that the type of properties a data value exhibits has implications on how that data may be represented in graphics. There are four such properties: absolute zero, equal interval, magnitude and category.

Displaying absolute values is different to displaying relative intervals. Displaying a set of data properties is different to displaying a data structure. A structure can be treated as a data set when the properties of objects in the structure are the focus of analysis. Derived data can be used instead of raw if the gap between the tasks and the raw data is too great. Data characteristics directly affect the type of visual implications necessary and the amount of programmatic processing required for a task.

2.1.2. Users

A user's approach often switches between goal directed and opportunistic due to the cyclic nature of perceptive and cognitive processes. Their nature affects how the high-level tasks are performed, thus affecting the kinds of visual tasks which are most suitable. The level of users experience also determines the nature of interaction.

Several cognitive models of Human-Computer interaction have been developed over the years in an effort to understand the various, complex user behaviours. Under visualisation circumstances, user behaviour can be divided into two sections:

1. Goal Directed. Interaction with the system is part of a decision making process with preconceived questions[2]
2. Opportunistic. Interaction is a discovery process. Users react to events in the surrounding environment.[13]

Some models see certain tasks as both goal directed and opportunistic. This view stems from the fact that another cyclic relationship exists between goal directed and opportunistic tasks. One usually leads to the other and vice versa. One may consider Shneidermans Visual Information Seeking mantra[12]: First gain an overview of the entire collection of data (a opportunistic perceptual process), leading to a decision (activates a cognitive process) to focus on items of interest. This in-turn produces another view of the data (another perceptual process, goal directed, looking for something specific). Finally, a decision is made (a cognitive process) to look at some items in detail (switch to a different representation). During the information drill-down process, the user may opportunistically find something else of interest. How often the users swap roles or how long they spend in each role is often a question of experience.[10]

2.1.3. High Level Tasks

Tasks have 3 aspects and fall under one of 6 high-level tasks. These high-level tasks would normally be performed cognitively but may be off-loaded with appropriate visualisation. Visualisations that satisfy such tasks may be considered useful.

To date, at least three aspects of task have been described in various works as relevant to visualisation design:

- Scope of tasks.
- Type of tasks.
- Decomposition and distribution of task load.

According to Bertin, Scope may be divided into 3 subtypes: elementary, intermediate and global. Each signifies a different amount of data that a task is required to operate over. Types of tasks may be viewed from several levels. Yu[17] defines the following as the set of all high-level tasks appropriate to visualisation under which all other tasks fall:

- Spotting outliers.
- Discriminating Clusters.
- Checking Assumptions.
- Examining relationships and interactions.
- Comparing group differences.
- Navigate and Traverse.

These tasks may be "off-loaded" from being purely cognitive tasks the user undertakes, to perceptual processes via the use of effective visualisation. These tasks may be accomplished through a concert of visual and physical tasks; an appropriate decomposition of high-level tasks into visual and physical tasks will reduce the cognitive load.

2.1.4. Visual Tasks

Visual tasks form the core of the visualisation and are derived from the high-level tasks. It describes how a high-level task can be accomplished visually. To accomplish a visual task, the visualisation must afford and offer the appropriate visual implications.

This section contains material adopted from Zhou and Feiner[19] only. They describe a set of visual tasks that enable the previously mentioned High-Level tasks to be accomplished. These Visual tasks are satisfied by the *visual implications* of different visual techniques. For example, the *spotting outliers* task may be accomplished by visually locating an object that is different relative to all others. However, this abstract task of locating can be satisfied through a number of concrete visual techniques, including highlighting the object in colour, or spatially segregating the object. Visual tasks form part of the bridge between visualisation techniques and the requirements drawn from our framework. There are a total of 15 such tasks described in.[10, 19]

Visual tasks may be derived from examining existing methods of visualisation, predicted by analysing the data, and by breaking down the high level visual tasks described previously. Certain sets of visual implications may satisfy the visual tasks but only a small number of these sets will prove useful. For a complete list of the visual tasks, refer to.[19]

2.1.5. Physical Tasks

Physical Tasks describe the physical act of manipulating some interactive controls, which may change the visualisation and thus the visual implications it affords.

As the name suggests, physical tasks refer to the users manual input using some form of peripheral device (such as mouse) in order to manipulate the visualisation. We define three such physical tasks:

- Select a region.
- Select a location.
- Specify a value.

Physical tasks may be categorised as either data operations or graphical operations depending on which domain the physical task affects. A data operation only changes the set of data to be visualised whereas a graphical operation changes how the visualisation is rendered. Again, like visual tasks, physical tasks can be goal directed or opportunistic.

2.2. The "How" of Visualisation Design

The previous sections have provided the tools for answering the first question; where interactive graphics may be used to reduce cognitive processing. The answers to this question in a given visualisation problem form its requirement set. For example, a requirement might be phased as: Given the data is a set of temperatures over time, to spot the outliers (high-level task) within the entire set of data (scope) opportunistically (behaviour), by empowering the user to clearly identify and locate (visual tasks) such pieces of data from the corpus. Interactive controls (physical tasks) may be used to adjust the threshold of what is considered an outlier or to reduce the scope of the data.

To complete the framework, one must consider how a graphical system supports the user user by meeting such requirements. This section outlines what constitutes a graphical representation and how visual implications are delivered.

2.2.1. Visual Implications

Visual implications refers to the implications arising from the organisation and presentation of the graphical primitives at varying levels. Since each graphical primitive in a visualisation represents some piece of data, it is directly suggesting to the user some relationships about the underlying data. From a different view point, a user can make inference about the data based on what he/she sees from and among the graphical primitives. This is where perception, cognition, gestalt psychology and individual differences comes into play.

Every graphical primitive affects the overall visual implications at three different levels:

Primitives level. This corresponds to the abstraction of one graphical item (or glyph) representing a set of the underlying data. At this level, the nature of the abstraction may be *noisy* or *smooth*. An example of this would be the notion of representing temperature over time. If the task required is simply to gauge the relative magnitude of multiple temperatures, a noisy plot would involve simply viewing the raw data represented as a two column table or a scatter-plot, whereas a smooth abstraction may involve using a coloured bar ranging from a blue colour representing a cold temperature to a red representing a warm temperature. The latter abstracts away the detail of the values, offloading the task of categorizing, associating and comparing magnitudes and presents a smoother overview. Whether an abstraction is noisy or smooth depends on the visual tasks expected to be performed on it. A smooth plot for one task may be noisy for another task, which demands higher level abstractions.

Within-Visualisation level. This corresponds to the grouping of individual glyphs within a single scene. At this level, there are two kinds of visual implications: *grouping* and *attention*. Visual grouping is derived from Gestalt psychology[9] which states that people tends to group visual elements together by: proximity, similarity, continuity and closure. Visual tasks such as cluster and distinguish are facilitated by the implications of grouping.

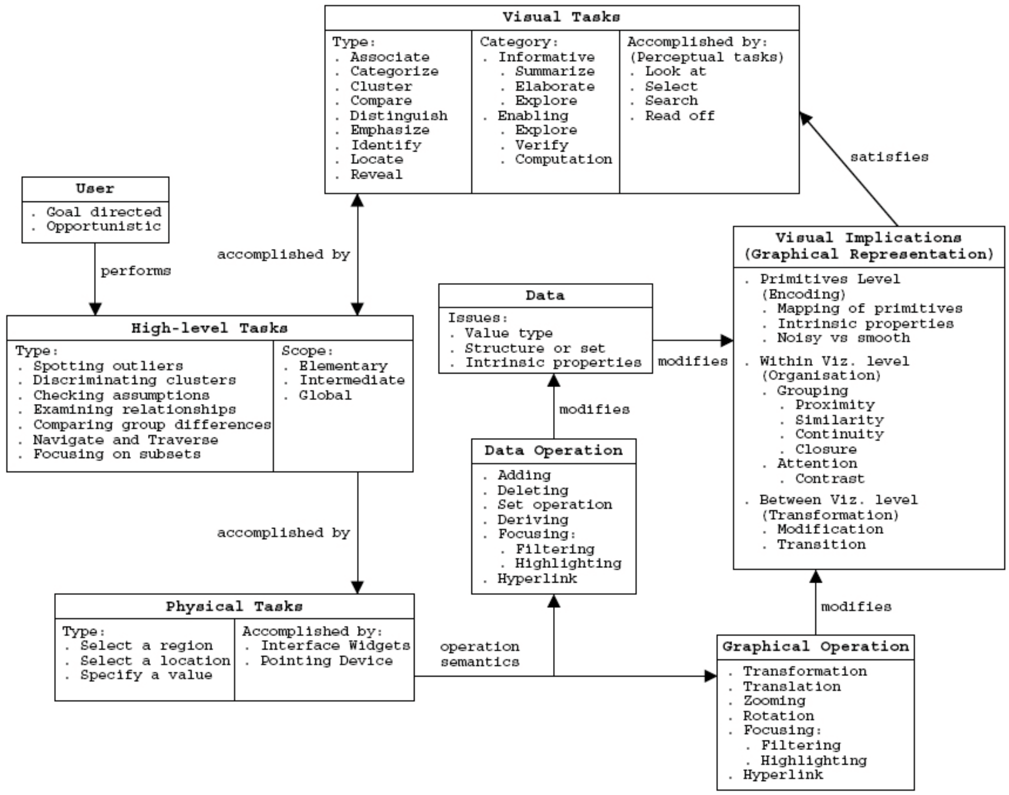

Figure 2. The Framework Summary

Visual attention refers to the variety of visual fixations one may make within a scene. Visual tasks such as select, identify and locate are facilitated by a scene which allows such level of visual attention.

Between-visualisation level. This refers to the relationship between successive images of data, usually as the results of some physical tasks, or as part of an animation. The new image can be a modification of the previous one (to incorporate new data), or a complete transition to another image.[14] Visual tasks such as associate and compare are usually performed between visualisations.

To address the requirements laid out in the preamble of this section, a scatter plot of temperatures over time will allow outliers to be opportunistically identified as a small number of graphical markers, located away from the visual grouping of the other markers. Whether such marker is located far away enough to be considered separate from the group depends on the scale chosen to plot the temperatures, which can be adjusted interactively.

2.3. Appropriateness of a design

Visual implications can either facilitate or hinder some visual tasks to be performed. Given a visualisation on which a variety of visual tasks can be performed, none may contribute to solving a users high-level tasks. Given that a high-level task may be solved with several visualisations, there are questions of effectiveness and efficiency. Lastly, a visualisation can be purpose built, or it may serve multiple purposes. There is a need for a finer definition of what makes an overall usable visualisation.

This framework defines the functionality of a visualisation as the set of high-level tasks that can be performed on it. Appropriateness is defined as whether such design matches the users task goals for a given problem. Given an appropriate visualisation, effectiveness is defined as the accuracy of task performance; efficiency is defined as the effort required to accomplish the task. We believe visualisations can be compared both qualitatively and quantitatively using these definitions.

While issues of aesthetics should not affect the appropriateness, effectiveness and efficiency of a given visualisation, an aesthetically pleasing visualisation is favourable over a less pleasing one, given all other measures being equal. The usability of a visualisation is in turn the summative score of all of the above.

2.4. Extensions to the framework

We have presented a meta-framework of visualisation design where individual aspects overlap with existing fields of study; user behavior and tasks in HCI, visual implications and perception in psychology, and existing visualisation designs offering possible reuse. As such advancements in these fields will no doubt contribute to better visualisation designs. The main contribution of this framework is not to impose a fixed set of design vocabulary and tools based on a snapshot of the current state of the art. Rather, it outlines how these individual aspects of visualisation relate to each other, such that any new knowledge in these areas is immediately applicable.

Finally, this framework has highlighted that visualisation design is the iterative process of matching visual implications against users high-level tasks, by breaking the tasks down into individual visual tasks and physical tasks. This framework has only presented an initial sample of visual and physical tasks. With this framework, the study of visualisation becomes the study of task decomposition and solution matching.

3. CASE STUDIES

To clarify our approach, we present two case studies. The first employs the framework to extrapolate a set of requirements from scratch. The second uses the framework as an evaluation; possibly leading to a set of requirements for a new version of the software being evaluated.

3.1. Case Study 1: DAITO

Difference Analysis Imagery in Taxonomies and Ontologies (DAITO) is a project currently undertaken by the author at the time of writing.[4] The aim is to ease the cognitive load placed on users of OWL ontologies where the merging and understanding of separate versions of the same ontology are key problems. Visual aids have often been used in this field but the adoption rate of such visualisations has been very low. To date, none support analysing two or more versions of the same ontology. Being able to recognise, understand and verify changes is extremely important to the ontology community. From both literature and user surveys, it is possible to gain a good understanding of the inherent problems and the general tasks ontologists wish to perform.

3.1.1. High-Level Tasks

The first consideration is the high-level tasks the users need to perform. What exactly do users want to do with their ontologies? Two high-level tasks were identified:

1. Exploration (What has changed / What are the differences in this ontology?)
2. Verification (What does this change mean / Why has “i” changed?)

These tasks have varying scope and fall into varying categories. For example, Exploration falls under *Navigate & Traverse* and *Discriminating Clusters* since users may recognise one area as interesting and the rest as irrelevant and traverse towards that cluster. Each of these tasks has similar aspects. Exploration and Verification have *navigation, examination* and *comparison* sub-tasks. In summary, Exploration represents a more opportunistic aspect of the three defined tasks, whereas Verification is much more Goal Directed.

3.1.2. Data and User Characteristics

Users fall into roughly two categories; these experienced with an ontology and these who are not. New users tend to spend more time in opportunistic exploration, whereas experienced users tend to ask more questions about particular items and return to some sort of start point more often than new users.

The data within an ontology has both implicit and explicit relationship types. There are several different atomic types with relationships constructed of at least two and often more atomic types. The data is largely nominal and does not lend itself well to smooth plots. This combination of atomic and composite at various levels within the ontology leads directly to *Detail vs. Overview* visual implication. When considering changes, more implicit relationships exist. The atomic data falls into the "category" subtype as defined in the framework. When these items are grouped together (as they always are in ontologies), *closure* and *proximity* become important visual implications.

3.1.3. Visual Tasks

There are several, potential visual tasks derived from the high-level tasks and these which might satisfy the visual implications already defined. *Reveal,Identify* and *locate* are the main visual tasks when performing exploration, as defined by Ng.[10] During examination, users potentially need to *Categorise, Generalise* and *Emphasise.* Finally comparing, as its name suggests, has a large set of potential visual tasks such as *Cluster, Distinguish, Associate, Compare* and *Correlate.*

3.1.4. Physical Tasks and Graphical Operations

As no actual implementation exists, the physical tasks are difficult to define. It is known, however, that verification requires the system to respond to users questions. Certain visual tasks imply some sort of animation or transition from one state to another. Cognitive science also tells us that interaction is important in creating a users mental model. It is here where existing taxonomies, visualisation patterns and qualitative aspects need considering. This is discussed later. To summarise, Navigation typically considers selecting an item, area or relationship. Examination considers a single data item and Comparison, an area.

3.1.5. Visual Implications

Following on directly from navigation, examination also suffers from Detail versus Overview. In addition, Visual Attention, Visual Grouping and Modification between visualisations are further considerations. Visual Attention arises from the generalise and emphasise tasks. Certain circumstances may mean items should be generalised whilst others are emphasised in order to guide the user to a particular goal, or suggest alternatives or to ease the user in deciding which data to ignore. Visual Grouping's roots are buried deep within Gestalt Psychology.[9]

Finally, the Comparison aspect suffers from similar implications although the Visual Grouping implication is more pronounced than that of Examination. Visual Grouping, specifically proximity, similarity, continuity (especially over successive visualisations) and closure are all important, linking directly from the Cluster, Associate and Compare tasks. Correlating various items over changes suggests particular attention should be paid to Visual Composition. Finally distinguish falls into two categories. Items and groups of items may be distinguished via proximity or similarity.

3.1.6. Summary of Daito

There are several visual implications, visual tasks which may support these implications, interesting data characteristics and scope for transition or animation.

- High-Level Tasks: Navigate, Examine , Comparison (aspects of Exploration and Verification)
- Physical Tasks: Select Area, Item or Relationship (Navigation). Select Item (Examination). Select Area (Comparison)
- Visual Tasks: Reveal, Identify, Locate (Navigate). Associate, Compare, Generalise, Emphasise, Categorise (Examine). Cluster, Distinguish, Associate, Compare, Correlate(Comparison)

- Visual Implications: Detail vs Overview (Navigate / Examine). Visual Attention (Examine). Visual Grouping, Visual Modification between Visualisations (Examine / Comparison). Visual Composition(Comparison)

These tasks and implications form the basis for selecting appropriate graphical operations and solutions. From the above, is is clear that animation or visual transition will play a strong role in DAITO, along with interactive, meaningful grouping of the atomic items that compose the larger, composite items. The data is highly relational with many intrinsic relationships, some of which tend strongly towards certain visualisation techniques. Detail vs. Overview is particularly important, leading to some form of visual abstraction and guidance. In comparing items, visual support needs to be given to associating said items, as well as providing easy comparison. Ideas such is highlighting, side-by-side views and similar would support such visual tasks.

The *where* visualisation can be used can be extrapolated from the visual tasks. Aligning, comparing and grouping the various categories of data are perhaps the most obvious. The *How* is more difficult to discern, but a selection of ideas can be compared against the visual implications identified.

3.2. Case Study 2: Infolens

InfoLens was designed to visualise a highly connected graph of semantic knowledge in the clinical domain.[6] The driver of the problem was the medical knowledge base becoming so big that no single knowledge engineer knows for sure what is actually in it, and whether what is in it is correct.[10] The primary problems to address were: to opportunistically identify patterns within the graph of knowledge, and to verify certain assumptions about the captured knowledge. These tasks are repeatedly performed by knowledge engineers, experienced or new, during the development of the model.

The most immediate intuition was to visualise the knowledge base as a node-link graph. There are certainly plenty of examples where such approach is aesthetically pleasing and usable, as many of graph based visualisations are now commercialised. However, the visualisation does not solve any of the user problems.

3.2.1. High Level Tasks

The high level tasks identified where visualisation can be used can be categorised into two kinds, given here as examples:

- Given that body structures (lungs, heart, liver) are three dimensional, all have certain attributes describing its relation to other parts and subparts, and other attributes describing its function. What is the family of attributes in the graph that are usually used to describe body structures?
- Given that all drugs must have an active ingredients, verify that all antibiotics for the liver must have an active ingredient of X

The former is opportunistic with an exploratory nature where as the latter is goal directed. The visual tasks to the above two problems can be generalised into, respectively:

- Find the common characteristics among a set of items (find clusters within a set)
- Confirm all individuals within a set must have a certain characteristics (spot outliers within a set)

3.2.2. Visual Implications

The kinds of visual implications to solve these visual tasks have been solved previously, as demonstrated by Bertin's Re-order-able Matrix.[2] Thus a matrix layout was chosen to be the visualisation form:

- A matrix layout of each item in the set down the rows, and unique attributes across the columns, common characteristics of the items are immediately obvious as columns of filled cells.
- The same matrix layout with a appropriate colour coding can lead the user to confirm that if all items has that colour, his question is verified.

3.2.3. Physical Tasks

- Data operation: The construction of a set from a graph, through a variety of semantic queries over the knowledge graph (get me all the descendents of body parts which are hollow 3 dimensional structures)
- Graphical operation: To lay out and highlighting in colour based on results of queries. (highlight all antibiotics in red if the active ingredient is NOT of type X)

3.2.4. Design and Implementation

The two physical tasks were eventually addressed by an implementation of Magic Lens,[3] where stacks of the lens filters were used in a variety of purposes.

- A stack of lens implements a data flow model constructing different sets of data based on queries, and present to the user in a matrix layout.
- Other types of lens were used to perform queries on the data and have the result highlighted in colours.

User evaluation suggests that while the visualisations did help them to answer the questions, the tool itself was very cumbersome to use, and the lens were hard to manipulate to achieve certain tasks.

InfoLens was a perfect case to demonstrate the strengths and weaknesses of this framework. It helped the designer of InfoLens to focus on what is important to the user, identifying where graphics should be used, what needs to be shown, how it should be shown. The emphasis on analysing visual tasks has helped identifying that a tabular format was functionally much more suitable to the tasks at hand, rather than trying to fit the tasks over a node-link graph. Based on the results from InfoLens, the DOPAMINE tool was constructed to visualise the drugs knowledge base.[16]

At the same time, the framework gave too much freedom in the actual implementation of the visualisation tool, leading to sub-optimal usability and aesthetics. For example, the data set construction aspect could be implemented using existing effective techniques, such as dynamic query sliders, or several well known data-flow paradigms. Indeed, some verification tasks may require simply a "yes" or "no" answer instead of a visualisation. Including such constraints into the framework of consideration could give rise to a better design.

4. RELATED WORK

Recently, there have been efforts to understand why certain visualisations fail or are not adopted. These, more qualitative aspects are addressed by Ernst[5] among others, and may form a second tier of analysis. Extra work in this area may prove fruitful since such aspects of visualisation are relevant to future usability and analysis.

Amar and Stasko's work on Representational Primacy[1] approaches the design problem from another direction. Although they agree that task analysis and taxonomies for visualisation are useful, visualisations still fail to bridge what they define as the "Rationale" and "World-View Gaps". Visualisations that do not address these gaps will most likely fail to address the key problems users wish to solve. The tasks and evaluation criteria presented by Amar and Stasko elegantly expands mainly on the Appropriateness aspect of this Framework, yet the effects of their ideas can be felt throughout the framework.

This framework provides an overall design methodology with hooks for attaching further accepted methods. One such important contribution are the Visualisation Patterns presented by Wilkins.[15] The core of Wilkins' PhD thesis describes a set of visualisation patterns which might be used to correctly apply proven visualisation techniques to satisfy the set of requirements derived by this framework.

5. SUMMARY AND FUTURE WORK

We have presented a unified, general framework which derives a set of requirements for a visualisation. Secondly, we presented two case studies to clarify the framework. Lastly, related work, specifically with the more qualitative aspects of design were considered. The aim of visualisation is to minimise cognitive load, to maximise externalisation of information, and to reduce information access cost. The aim of this framework is to provide some structure to the design process. The Framework begins with considering *Where visualisation may be used*, leading to *How Visualisation may be implemented*, leading back to where visualisation may be used in an iterative manner when needed. The relationships between the various tasks, implications and the nature of user and data are shown in Figure 2.

One first considers the in-variants; the data, the users and their high-level tasks, leading to an initial set of visual and physical tasks. Secondly, the visual implications are considered along with the graphical operations, satisfying the visual tasks, possibly resulting in a refinement of these tasks. Finally, when all iteration is complete, a set of task and implications become clear, forming the basis for a good visualisation. For a more in depth analysis and breakdown of the various sections of this work, refer to Gary Ng's thesis.[10]

Our framework is not exhaustive nor is it designed to replace all existing techniques. Rather, the aim is to structure existing approaches from differing fields allowing designers to benefit from a variety of different techniques without losing sight of the overall goal. Users of this framework can customise the level of detail required for each area, bringing in relevant theories and techniques.

There is certainly scope for future work in the area of qualitative analysis and visualisation patterns. Arriving at a set of formal requirements and considerations is only the first step. Qualitative guidelines need to be followed whilst designing a visualisation in keeping with the past experience that patterns may provide.

REFERENCES

1. Robert Amar and John Stasko. A knowledge task-based framework for design and evaluation of information visualizations. In *Infoviz*, 2004.
2. J Bertin. Semiology of graphics: Diagrams, networks, maps. Technical report, University of Wisconsin Press, 1983.
3. E. A. Bier, M .C. Stone, K. Pier, W. Buxton, and T.D. DeRose. Toolglass and magical lenses: The see-through interface. In *SIGGRAPH '93, International Conference on Computer Graphics and Interactive Techniques*, pages 73–80, Anaheim, CA, 1993. ACM Press.
4. Benjamin Blundell and Steve Pettifer. Requirements for the visualisation of ontological evolution. In *EGUK Theory and Practice of Computer Graphics*, pages 18–23, Bournemouth, 2004. IEEE Press.
5. Neil Ernst, Neil Ernst Margaret-Anne Storey, Robert Lintern, and David Perrin. A preliminary analysis of visualization requirements in knowledge engineering tools. Technical report, CHISEL Group University of Victoria, 2003.
6. Ian Horrocks, Alan Rector, and Carole Goble. A description logic based schema for the classification of medical data. In *The 3rd Workshop KRDB96*, pages 24–28, 1996.
7. Jessie B Kennedy, K.J. Mitchell, and P.J Barclay. A framework for information visualisation. *SIGMOD Record*, 25(4):30–34, 1996.
8. N Murray, Carole Goble, and Norman Paton. A framework for the graphical visualisation of database interfaces. *Journal of Visual Languages and Computing*, 9(4):429–456, 1998.
9. Keith V. Nesbitt and Carsten Friedrich. Applying gestalt principles to animated visualizations of network data. In *Sixth International Conference on Information Visualisation (IV'02)*, page 737, 2002.
10. Gary Ng. *Interactive Visualisation Techniques for Ontology Development.* PhD thesis, The University of Manchester, 2000.
11. D.J. Peebles, P.C.H. Cheng, and N. R. Shadbolt. Multiple processes in graph-based reasoning. In *Twenty First Annual Conference of the Cognitive Science Society*, Hillsdale, NJ, 1999. Lawrence Erlbaum.
12. B Shneiderman. The eyes have it: A task by data type taxonomy for information visualization. Technical report, University of Maryland. Department of Computer Science Institute for System Research, 1996 1996.

13. L.A. Suchman. *Plans and Situated Actions The Problem of Human-Machine Communication*. Cambridge University Press, 1987.
14. L Tweedie. Interactive visualisation artifacts: How can abstractions inform design. In M Kirby, A Dix, and J Finlay, editors, *HCI, People and Computers*, pages 247–265, Huddersfield, UK, 1995. Cambridge University Press.
15. Barry Wilkins. *MELD: A Pattern Supported Methodolgy for Visualisation Design*. Doctor of philosophy, University of Birmingham, 2003.
16. Chris Wroe, W. Solomon, Alan Rector, and Jim Rodgers. Dopamine: A tool for visualizing clinical properties of generic drugs. In *The Fifth Workshop on Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-2000)*, Berlin, Germany, 2000.
17. C.H. Yu and J.T. Behrens. Applications of scientific multivariate visualization to behavioural sciences. *Behavioural Research Methods, Instruments, and Computers*, 27:264–271, 1995.
18. Michelle X. Zhou and S. K. Feiner. Data characterization for automatically visualizing heterogeneous information. In *InfoViz*, pages 13–20, San Francico, CA, 1996. IEEE Computer Society Press.
19. Michelle X. Zhou and Steven K. Feiner. Visual task characterization for automated visual discourse synthesis. In *Conference on Human Factors in Computing Systems, CHI-98*. ACM Press, 1998.