7 Perception and Cognitive Aspects

7.1 Motivation

The human is at the heart of visual analytics human interaction, analysis, intuition, problem solving and visual perception. This chapter is entitled “perception and cognition”, and it is possible to have a narrow focus of this looking purely at the perceptual and cognitive aspects during the time when a user interacts directly with a visualisation or adjusts parameters in a model. However, there are many human-related aspects of visual analytics beyond those involved in the direct interactions between a user and a visual representation of data. Figure 7.1 presents a simplified view of the broad visual analytics process that emphasises some of the wider context and the human issues involved.

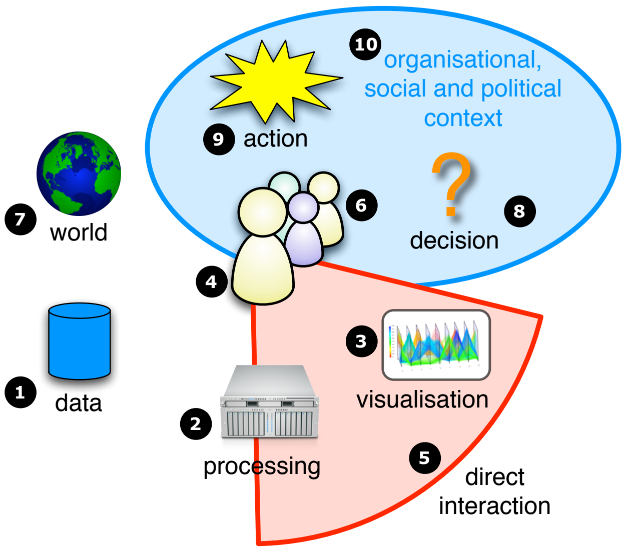

Figure 7.1: The human context of visual analytics

Working through the numbered parts of Figure 7.1, visual analytics involves some data (1), typically being processed (2) computationally (e.g., machine learning, statistics), then visualised (3) and interpreted by the user (4) in order to perform problem solving, analysis etc.. The pie-shaped region (5) represents the obvious direct interactions between the primary user, processing and visualisation. When multiple people are involved in this process, it can also be collaborative (6).

However, the role of people goes beyond direct interaction with visual analytics systems. The data being visualised comes from the world (7) (or some simulation of it) and is typically used by people, who may not be those involved in interacting with the visual analytics system, to make decisions (8) that influence actions (9) that ultimately affect the world.

This gives rise to a far broader organisational, social and political context (10): the stakeholders who use the outputs of visual analytics and those impacted by the decisions cannot be ignored by those using the systems and indeed, those involved in the technical process may be subject to social and political pressures as well as considering how the results of the visual analytics process can best be presented to others.

7.2 State of the Art

There is a substantial literature on specific techniques and systems for interactive visualisation in general, although fewer looking at human interaction when there is more complex non-visualisation processing as in visual analytics (with exceptions such as clustering or dimensional reduction). Looking beyond experience reports or simple user studies to detailed perceptual and cognitive knowledge the picture becomes more patchy. There is work on static visualisation (e.g., abilities to compare sizes), yet there is little on even simple interactive or dynamic visualisation let alone where this is combined with more complex processing. Again, whilst there is a longstanding literature of technical aspects of collaborative visualisation, social and organisational aspects are less well studied. For example, recent work on sales forecasting found that, perhaps unsurprisingly, issues of organisational context and politics were as important as statistical accuracy. Methodology is also important, even in more traditional visualisation areas issues, such as evaluation, are known to be problematic.

7.2.1 Psychology of Perception and Cognition

Psychological research on perception of visual information is based on the seminal work of Allan Paivio who asserted that the human perceptual system consists of two subsystems, one being responsible for verbal material and the other for all other events and objects (especially visual information). He emphasised the importance of mental images for human cognition. Even if some of his assumptions have been criticised, his considerations still provide an important reference point for psychologists investigating visual perception.

Distinction between high and low-level vision

Researchers in perceptual psychology usually distinguish between high and low-level vision. Activities related to low-level vision are usually associated with the extraction of certain physical properties of the visible environment, such as depth, three-dimensional shape, object boundaries or surface material properties. High-level vision comprises activities like object recognition and classification. Results from low-level vision research are finding their way now in visualisation and visual analytics[122], but results from high-level vision research are not yet adopted.

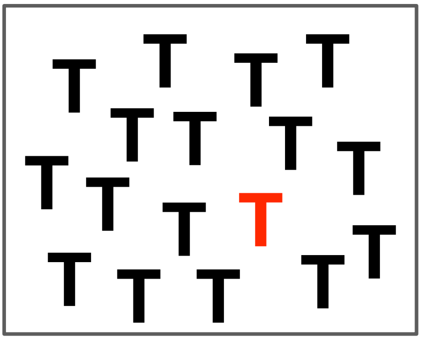

Figure 7.2: Preattentive processing – pop-out effect

Ware[122] discusses preattentive processing quite extensively. This theory tries to explain the fact that some elements of visual displays pop out immediately and can be processed almost automatically (see Figure 7.2). These elements can be distinguished easily from others, for example by their form, colour or orientation. Such processes are very important for visual perception because they support visual search considerably. Despite some criticism, this theory has been very influential for information visualisation and visual analytics because the quality of systems representing information on a computer screen depends to a considerable extent on whether they support search processes efficiently.

Preattentive processing makes items pop out the display automatically

The human visual system has by far the highest bandwidth (the amount of data in a given time interval) of any of our senses and there is considerable research into how we make use of this data about our immediate environment. Visual representations are generally very short lived (about 100msec) and consequently much of what we 'see' is discarded before it reaches consciousness. Evolution has given humans the ability to rapidly comprehend visual scenes, as well as text and symbols and much of this rapid, unconscious processing involves representations in our conceptual short-term memory[90] where small snippets of information (such as individual words) are consolidated into more meaningful structures. However, addition processing stages are required before we become aware of a particular stimulus and it survives in longer-term memory. Demands on this higher-level processing from rapidly presented sequences of visual stimuli can give rise to failures in retaining visual information, such as attentional blink and repetition blindness[33], and as such are important to designers of visual analytic systems.

Another theory of visual perception, which has some relevance for visual analytics, is Gestalt psychology. This assumes that visual perception is a holistic process and that human beings have a tendency to perceive simple geometric forms as illustrated by the examples in Figure 7.3. This implies that the structure underlying a visual display is more important than the elements

Humans tend to perceive simple geometric forms

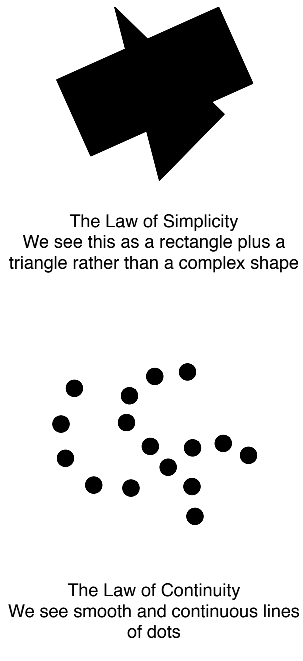

Figure 7.3: The Gestalt Laws imply that we tend to see simple, often connected structures within a scene. (Only a subset of the Laws is shown)

and is often summarised as ‘The whole is more than the sum of its parts’. These principles can be used for guiding attention efficiently in visual displays in order to help reason through the data, although we need to be aware that strong visual characteristics, such as bright colours or joining lines, can dominate or influence one’s reasoning processes.

Visual perception is an exploratory process

Recent research in the psychology of perception indicates that perception is an exploratory and active process. Gibson[48] pointed out that human perception is tied to the movement of the human body in a natural environment. We do not see a sequence of more or less static images but a continuous flow of changing scenes in this natural environment whilst we move around.

Neisser[81] developed a model of perception based on a cycle consisting of schemata, available information about objects and perceptual exploration (see Figure 7.4). The process described in this model is always influenced by past experience (schemata, expectations). Based on this experience, hypotheses are formulated which guide perceptual exploration. Our cognitive resources, especially our short term memory, are limited; therefore, we direct our attention only to objects we consider in advance to be interesting. If the information from the environment does not match these hypotheses, schemata in human memory are modified. This is an ongoing and iterative process.

In this context, the movement of the eyes, especially the so-called saccadic

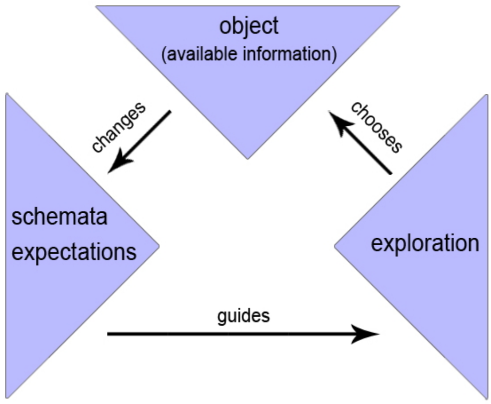

Figure 7.4: Model of Perception[81]

movements, plays an important role. The eyes do not move continuously, but in series of jumps (about four per second). Between these jumps, fixation occurs when people gaze at objects in the environment. Eye movements are especially important as peripheral human vision is rather inefficient. To resolve detail, an image has to be projected onto the fovea - a fairly small region on the retina, which is responsible for sharp central vision. Everything else in the visual field is quite blurred (see Figure 7.5). It is, therefore, not possible to get a comprehensive impression of the environment at one glance. In this context, eye movements play an important role. They enable human beings to see the necessary details in a series of several fixations. We have to look for information actively to get a fairly comprehensive image of the environment, in a process quite similar to the one described by Neisser (see above).

Eyes move in a series of jumps

Human peripheral vision is poor

Figure 7.5: Acuity is only high in the centre of the visual field.

These and similar approaches in the psychology of perception are especially suited for modelling the interaction of users with visualisations. The usage of such tools is often described as an active and exploratory process yielding complex insights[7]. In the course of this process, hypotheses are generated and tested on the basis of the data visualised by the tool. There is research that applies results from cognitive psychology to the design of visualisations, which adopts such an approach[102].

Change blindness means we often fail to see seemingly obvious changes

Recently, the phenomenon of change blindness has attracted much attention[102]. Change blindness describes the phenomenon that observers often fail to notice important changes in their environment, especially if they do not pay attention to these changes. Rensink also argues that humans do not possess a detailed, picture-like representation of the scenes they see. Nevertheless, observers gain the subjective impression that they have a stable representation of their environment. This is due to the fact that observers can get any information they need whenever they want it just by focusing their attention on the relevant object. It might be argued that observers use the environment as some kind of external memory to relieve their own limited capacities (especially short term memory and processing capacities). This approach also assumes that perception is active, not passive. Ware describes this as visual queries - the search for patterns in the outside world. This capacity of human information processing is very flexible and adaptive. Both Rensink and Ware argue that visual representations, especially visualisations on a computer screen, should be designed in a way to support these processes, and they both suggest a set of design guidelines for this.

Flexible and adaptive vision system searches for patterns

7.2.2 Distributed Cognition

Distributed cognition is a theoretical framework describing the interaction between (groups of) persons and artefacts[58, 61]. It builds on the information-processing concept of a problem space, but extends the boundaries of the problem space to incorporate knowledge in the mind of the user and knowledge in the world. It proposes that our everyday problem solving involves the coordinated use of knowledge structures in the mind, in our environment and from other individuals. The object of investigation is, therefore, not the single individual, but a system of cooperating individuals and artefacts. The model has been adopted in HCI to clarify problems of the interaction of users and computers. Distributed cognition argues that cognitive accomplishments are usually achieved in conjunction with artefacts. In these artefacts, representations of knowledge are embodied as, for example, in a thermometer or other measuring devices, which contain the accumulated information about this scientific area. Results achieved by using such cognitive tools emerge from the interaction between human and artefact and cannot only be attributed to human activity.

Representation of knowledge is embodied in everyday objects

In many cases, human users of information technology do not possess coherent and comprehensive mental models of the problem at hand. Such mental models only emerge in the process of using the technology because the information relevant for the solution of these problems is distributed among humans and computers. O'Malley and Draper[83] argue that computer users do not possess and also do not need such coherent models because they can, in many situations,

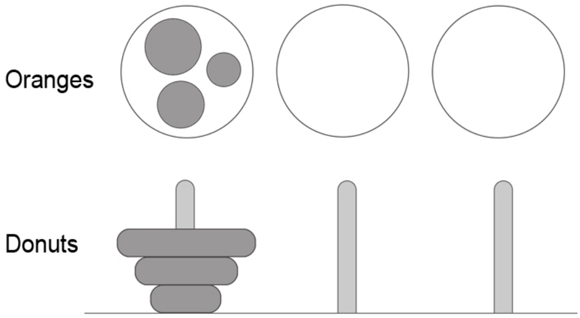

Figure 7.6: Oranges vs Donuts representation of Towers of Hanoi (adapted from Liu et al.,[72])

extract the relevant information from the environment. In this way, users can alleviate the burden on short term and long term memory. In free recall experiments, users of word processors are usually unable to recall many menu items. This low achievement is, according to O'Malley and Draper, due to the fact that users employ stimuli from the context of the interface to guide their search processes. They either use semantic relationships between header and menu item or information about spatial location to find relevant commands. The successful usage of a word processor (and similar programs) is, therefore, probably due to the interaction between user and system and depends on distributed representation of the relevant information on the computer as well as in the mind of the user. In this context, the strengths of each information processing system (humans and computers) are utilised and both systems complement each other. Hollan et al[58] argue that computer interfaces should be designed in a way to support this process efficiently.

Humans do not have to remember everything but extract visual clues from the environment

In relation to visual artefacts, probably the most compelling work comes from Zhang & Normans theory of distributed representations. Central to the theory is the concept of the Representational Effect: "The phenomenon that different (visual) representations of a common abstract structure can generate dramatically different representational efficiencies, task complexities and behavioural outcomes"[128]. It has been argued that the design of visualisations should carefully consider this effect[72] as every representation offers various possibilities and has specific constraints. The Towers of Hanoi problem can, for example, be represented as different sized donuts on pegs or oranges on a plate (see Figure 7.6). The donuts-on-pegs representation is inherently easier because constraints of the problem are part of the analogy, as only the topmost and largest donut can be removed (only one can be moved at a time). Users adopt situated solution strategies using previous practical experience rather than abstract mental plans. This phenomenon is more consistent with distributed cognition than with traditional problem-solving theories. It seems to be a plausible assumption that similar strategies are used in interacting with information visualisation tools as every object on the screen offers specific perceived affordances (e.g., a button looks like an object that can be pressed).

Problem solving depends on context rather than abstract plans

Interaction is, therefore, very important although Liu et al.[72] point out that interaction is still a concept which is not very well understood, and research is required into how people develop interaction strategies during sense-making and analytical reasoning.

7.2.3 Problem Solving

It has been argued that the exploration of data represented by visualisations is to a certain extent a problem solving activity. In problem solving, researchers usually distinguish between well-defined and ill-defined (or ill-structured) problems. The latter is where virtually no information about the problem and possible solutions are available, so the early stages of problem solving (recognition, definition, representation of a problem) are a challenging task. If the problem is well-defined, the emphasis of the problem solver's activity is on the later stages (development of a solution strategy, progress monitoring, evaluation of the solution). In addition, the solution path can often be described by an algorithm, which is not possible with ill-defined problems because they usually necessitate radical changes in problem representation.

Ill-defined problems are a challenge

An example for an ill-defined problem, which might necessitate radical change of representation, could be described as follows. Imagine a person going to work by car. One day, the car breaks down, and expensive repair is necessary. The person has to decide, whether they wants to repair the car or buy a new car. The problem to solve in this case, is the consideration of whether it is more expensive to repair the old car or buy a new (or used) car. But they might also consider not to buy a new car at all, but take the bus to go to work instead. Often, such radical reformulations of problem representations are not self-evident. In the case of ceasing to use a car, this has serious consequences for the life style of a person. This is, therefore, not an easy choice.

So far, research into problem solving (e.g., Simon's theory of problem solving) has concentrated on well-defined problems, although most problems in everyday life are ill-defined. Likewise, the problems for which interactive information visualisations are developed are often ill-defined. The Andrienkos[7] point out that a common goal in explorative data analysis is to 'get acquainted with the data'. This is a very general goal, and often more specific questions are only formulated after a general overview of the data. This usually is an iterative process of exploration. At the beginning, the problem is not defined in great detail, and radical changes of representation (e.g., another type of visualisation) in the course of the exploration of the data are possible.

Gaining insight is about discovery and is often unexpected

In this context, the concept of insight plays an important role. Increasingly, the term 'insight' is being used[82, 127] to denote that the exploration of information presented by visualisations is a complex process of sense-making. Saraiya et al.[95] define insight "as an individual observation about the data by the participant, a unit of discovery". They observe that the discovery of an insight is often a sudden process, motivated by the fact that the user observes something they have overlooked before. It is the purpose of visualisations to support this process and make the detection of insights easier. North[82] points out that the definition of insight used in information visualisation is fairly informal and that researchers tend to use implicit conceptualisations. He posits that important characteristics of insights are that they are complex, deep (building up over time), qualitative (not exact), unexpected and relevant. Yi et al[127] also argue that there is no common definition of the term ‘insight’ in the information visualisation community. They point out that insights are not only end results, but might also be the source of further exploration processes. At the beginning of such exploration processes, there is often no clearly defined goal, and insights might be gained by serendipity. They assume that a vital question is how people gain insights, and they identify four distinctive processes how this might be done: provide overview (understand the big picture), adjust (explore the data by changing level of detail/selection, e.g., by grouping, aggregation, filtering), detect patterns and match the user’s mental model (linking the presented information with real-world knowledge). The authors note that barriers to gaining insight include inappropriate visual encoding, poor usability and clutter.

How we gain insight is a vital question when designing visualisations

There is some similarity of the ideas about insight in information visualisation/visual analytics and the concept of insight proposed by psychology, especially in the area of human reasoning and problem solving[106]. The term insight was first used in psychology by Gestalt psychologists. Gestalt psychology conceptualises insight as a result of productive thinking, which goes beyond existing information. It often comes suddenly as a consequence of a complete restructuring of existing information. Gestalt psychology is based on holistic cognitive processes, which means that we do not solve problems by trial and error in a stepwise process (as behaviourism had assumed), but by detecting the meaningful overall structure of a situation.

Gestalt psychology suggests gaining insight is about restructuring existing information

Mayer[77] points out that research concerning insight concentrates on the first phases of the problem solving process (especially the representation of the problem) and on non-routine problems, that is problems, which problems solvers have not solved previously. He describes five interrelated views of insight based on the assumptions of Gestalt psychology:

- Insight as completing a schema
- Insight as suddenly reorganising visual information
- Insight as reformulation of a problem
- Insight as removing mental blocks
- Insight as finding a problem analogue

In principle, all of the above mentioned aspects are relevant for the clarification of the processes related to interaction with visualisations, but some of them seem to be especially important. ‘Insight as suddenly reorganising visual information’ is per se concerned about visual cues. It occurs when a person *looks* at a problem situation in a new way. Insight as the reformulation of a problem is related to that. In this case, a problem situation is represented in a completely new way. The suddenness of a solution is often seen as a characteristic of this theory of insights. It should be pointed out, however, that suddenness in this context does not mean that the solutions occur very quickly

Insight may occur suddenly but often requires much unconscious effort as restructuring may take some time, and even if a viable solution turns up, it usually requires some effort to realise it.

Whilst the importance of insight is for non-routine and ill-defined problems, in practice, laboratory experiments focus on well understood puzzles in order to make the empirical research more tractable. These puzzles are new to the subjects being studied, but typically have a single 'right' solution and all the information needed available (see for example, the puzzle in Figure 7.7).

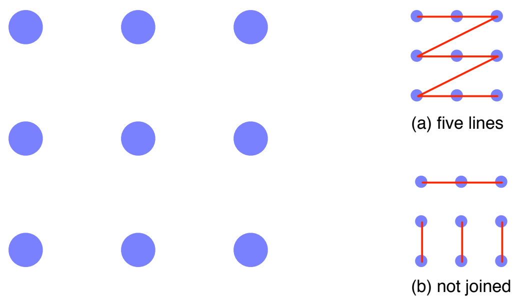

Figure 7.7: Nine dots puzzle: draw four straight lines that go through all nine dots, but without lifting pen from the paper. Note (a) and (b) show two incorrect solutions (a) has five lines not four and in (b) the lines cannot be drawn without lifting the pen. (see Figure 7.10 for solution)

Expert decision making often uses a highly developed mental model

Research into expert decision making in critical systems may provide an alternative path from understanding to insight. Klein has investigated how workers such as fire fighters, pilots and military personnel can resolve problems in high pressure environments[68]. He proposes that naturalistic decision making is often recognition primed, based on an individuals projected model of causal relationships. He provides a compelling example of how a naval officer was able to distinguish between an oncoming missile and friendly aircraft in a very primitive visual display. This difference would be impossible for a non-expert to identify as it involved the integration of both visual feedback and a highly developed mental model of the battlefield. This style of investigation is highly relevant for understanding the 'A-ha' moment that allows expert decision making to occur.

The usage of analogies also plays an important role for getting insights and is often mentioned as a source of creative thought[59]. In information visualisation, space is usually used as an analogy for other, more abstract phenomena (consider a scatterplot of engine size vs. miles per gallon). As human beings are highly capable of processing spatial information coming from their environment, space is a powerful analogy. In recent years, experimentation has taken place to clarify the concept of insight. The results of this research might form a valuable input for visual analytics, especially because it emphasises the reasoning processes associated with using information visualisations.

7.2.4 Interaction

The previous sections have concentrated on how humans perceive visual artefacts within abstract representation of data and try to make sense of these in order to gain information. We have also looked at work on modelling interaction and developing theoretical frameworks. The importance of interaction has been emphasised, as it is this that provides the opportunity for the user to explore the dataset. Whilst we can make good use of the large amount of research effort under the umbrella of HCI, there is not so much work focussed on visual analytics. Indeed, one of the recommendations from Illuminating the Path[111] was the creation of a new science of interaction to support visual analytics.

Interaction is vital in visual information discovery

A comprehensive review of the literature on interacting with visualisations is given by Fikkert et al.[45], although the authors do focus on virtual environments and associated display and interaction devices rather than information visualisations.

Attempts have been made to classify interaction for information visualisation[25]. More recently Yi et al.[126] identified the following categories of interaction:

We should think about the users' intentions when designing interactive systems

- select : mark data items of interest, possible followed by another operation,
- explore : show some other data e.g., panning, zoom, resampling,
- reconfigure : rearrange the data spatially e.g., sort, change attribute assigned to axis, rotate (3D), slide,
- encode : change visual appearance e.g., change type of representation (view), adjust colour/size/shape,
- abstract/elaborate : show more or less detail e.g., details on demand, tooltips, geometric zoom,
- filter : select or show data matching certain conditions,
- connect : highlight related data items e.g., brushing (selection shown in multiple views).

It useful to group together different interactive operations in this way, but possibly a more important outcome is a vocabulary to think about users' intentions when exploring datasets.

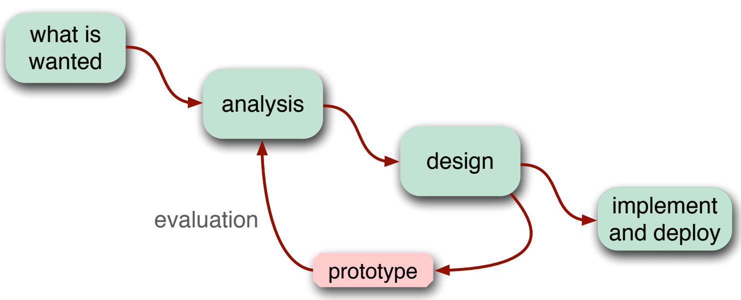

Figure 7.8: Typical user interface design process[36]

7.2.5 User Evaluation

User evaluation is at the heart of both research into human computer interaction and usability professional practice. This is because both researchers and practitioners recognise the limits of their ability to predict users responses to complex interactive systems and so there is always a need to test with real users, if possible, in real situations. For this reason, user interface design processes usually involve a tight cycle of prototyping and evaluation (see Figure 7.8).

Within HCI, evaluation techniques fall roughly into two styles:

- *quantitative evaluation* emphasises measurable outcomes, typically task completion time and error rate,
- *qualitative evaluation* emphasises more interpretative analysis of video and audio logs, or direct observations

Quantitative evaluation is often performed within well-controlled situations or laboratory settings, whereas qualitative evaluation is often performed ‘in the wild’ in real world situations, or artificial ones made closer to reality. Sometimes the two are seen as alternatives with strong proponents for each, but they can more productively be seen as complementary offering alternative insights.

While not diminishing the importance of effective evaluation, there is also a growing realisation that user evaluation, at least interpreted in a simplistic sense, is not always the most appropriate tool for all stages in the design and research processes. The tight cycle of prototyping and evaluation works well for refining and fixing the details of a design, especially in well understood domains. However, it is not so effective at arriving at novel designs or establishing the insights needed to drive new design ideas.

Within the information visualisation community, there is an ongoing discussion about methodological approaches for evaluation[13]. In this context, researchers argue that the measurement of time and error are insufficient to evaluate information visualisation techniques and tools because visualisation is typically exploratory in nature: interaction with information visualisation yields insights rather than information. This discussion is highly relevant for visual analytics because it emphasises the importance of the human reasoning processes as a whole, discovering new patterns within data rather than performing a known task in an ‘optimal’ way.

Quantitative evaluation is not appropriate for exploratory visualisation systems

One approach to this is to adopt more qualitative approaches. One example is grounded evaluation, an iterative design process that uses qualitative studies as a form of evaluation that can be carried out before initial design has been recommended. This is based on the recognition that to ensure the utility of visual analytics solutions it is necessary to ensure that the context of use is focused upon throughout the development life-cycle.

While this and other qualitative methods are better able to deal with the exploratory nature of the use of visual analytics, they still face the problem that users may not be able to appreciate the potential of radically new technology.

In such cases, it may be better to regard early prototypes as *technology probes*; that is being there to expose users to new ideas and then use this as the means to obtain rich, usually qualitative, feedback.

If evaluation is set within a wider context of ‘validating’ designs or research concepts, then these different approaches can be seen as building parts of a larger argument that may also include previous literature, published empirical data, theoretical models, and expert insights.

Chapter 8 discusses evaluation as an issue within visual analytics.

7.2.6 Early Application Examples

While basic theories and knowledge of human capabilities and behaviour can be applied from first principles, more applied knowledge is also needed. This is especially important when multiple factors come together. For example, complex interactions cannot be thought of as a combination of simple interactions often studied in pure science. In addition, it is only after protracted use that many issues become apparent, making it particularly hard to assess novel technology. This is of course the case for visual analytics as it is a new field. Happily it is possible to find much older systems and areas, which share essential characteristics with visual analytics and indeed would probably be termed as such if the phrase had been used when they were first established. Such systems are an opportunity to explore more applied issues with the benefit of hindsight and in some cases long-term use. They offer a wealth of existing knowledge to help us design for new visual analytics systems, and also a comparison point to assess the impact of changing factors such as massive data volumes or new interaction technology.

Can learn from early applications that share characteristics of visual analytics

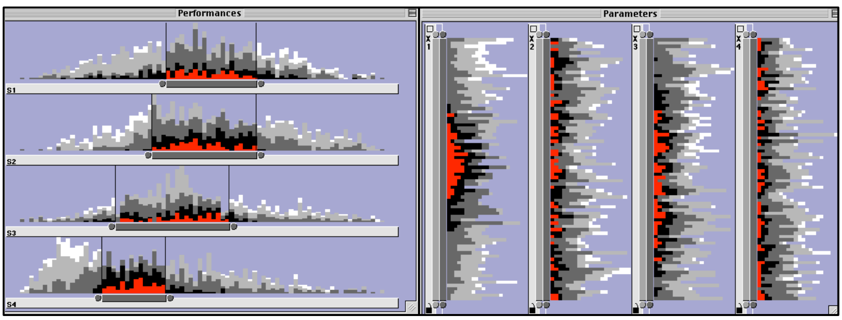

Figure 7.9: Parameter selection sliders of the Influence Explorer[117]

One example is the Influence Explorer shown in Figure 7.9[117], which used extensive computation to simulate the space of design parameters of light bulbs and then presented the results using interactive visualisations. Unlike much of the early work in visualisation, which was often framed around particular ideas for techniques, in the case of Influence Explorer the problem domain came first and the innovative interactive visualisation was developed in order to solve the problem. Influence Explorer embodies many critical features of visual analytics, notably the fact that there is a joint activity of analysis involving human perception and insight as well as computation and visualisation. It made use of sampling, as the complete data space of design parameters was too large to simulate exhaustively. It also allowed the user to 'peek over the horizon'. Most visualisations show you the effect of the current viewing parameters and rely on the user to actively interact to see alternatives temporally; in contrast, by exploiting a technique first introduced in the Attribute Explorer[118], the Influence Explorer's parameter selection sliders include miniature histograms, which let you know what the impact would be of alternative or future selections. The Influence Explorer is highly unusual, it is the first system of which we are aware that actively used sampling in visualisation and HiBROWSE[42] (a largely text-based faceted browser) is the only similar system allowing this 'peeking'.

Influence Explorer embodies many critical features of visual analytics

Business problems have long required complex analysis. One example is sales forecasting. In this case past sales data is typically modelled using various forms of time series analysis and the predictions from this visualised as simple graphs. However, the computer predictions cannot be used on their own as there are many additional internal and external factors such as sales promotions, advertising, competitors, and even the weather, all of which can influence future sales. Using the forecasting system involves the selection of data (e.g., do you include historic data on a product that had a major change?), the choice of forecasting algorithm (e.g., seasonal adjustments, kind of time series analysis), and then the inputting of 'adjustments', that is manual changes to the predicted outputs – effectively, an early example of visual analytics combining computation, visualisation and user interaction. The analyst may face pressure from members of the organisation, for example, a product division may wish to see higher forecasts in their area, and the results are not the end point of the process as they feed into decision making meetings where the forecasts are used to form plans for stocking, pricing etc.. Even the notion of accuracy that is central to the study of this process is problematic as the predictions feed into the process that is being predicted; indeed, forecasters are often more concerned that their forecast are reasonable and make sense to the recipient, and are less concerned with an ideal 'best' prediction.

Sales forecasting is an early example of visual analytics

There is a tight integration between users, computation process, organisational influence and the reflexive nature of visual analytics

As we can see, sales forecasting emphasises the need to take into account the whole picture in Figure 7.1: the tight interaction between users and computational processes, the organisational and political pressures that influence the analysis, and the reflexive nature of visual analytics, where its outcomes may affect the data on which it depends. These are lessons for visual analytics more generally. For example, in the homeland security applications that are the focus of Illuminating the Path[111], there may be a predicted attack by a terrorist group and the suspects consequentially arrested; the attack will therefore not take place, but this does not mean the prediction was 'wrong'.

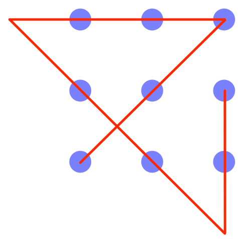

Figure 7.10: Solution to nine dots puzzle (see Figure 7.7)

7.3 Challenges and Opportunities

7.3.1 Stakeholders and Communities

The user facing a visualisation is not alone. As already noted, the results of visual analytics are often for other people, perhaps a manager making a decision or maybe for the general public. How can visual analytics systems support not just the problem solving and analytic task of the direct user, but also the production of static or interactive visualisation for these secondary users? In some situations, the end user may not be a visualisation expert but an ordinary day-to-day user of a computer. When issues of eDemocracy and open data are high on political agendas; we need to consider ways to give the general public tools to understand and to analyse vast volumes of data, in their own ways.

General public need tools to understand data in their own ways

Appropriate design methodologies have to be developed to support the reasoning processes of non-expert users. Simplicity, natural metaphors and intuitiveness will be important aspects in this context. Guidance is often required, especially for novice users, on what visualisations (scatterplot, parallel coordinates, treemaps, etc.) are appropriate for a given task; the focus should be on the problem to be solved rather than the type of raw data. Some visualisation tools (e.g., Tableau Software) attempt to do this by suggesting alternatives, albeit only simple visualisations, and some systems automatically generate the ‘appropriate’ visualisation.

Accepting that the immediate user is part of a community can help. As mash-up technology becomes commonplace, local experts could be one way in which more customised visualisations could be made available for those with less technical skills. For this we need to consider both the social means to share such components and also the technical means by which visualisations and processing can be slotted together in a plug-and-play manner. This is also an important requirement for effective research to prevent researchers wasting time implementing basic mechanisms in order to experiment with a single novel feature. While there are existing open-source visualisation systems a combination of poor documentation, reluctance to use external solutions, and lack of knowledge of what is available, makes reuse rare. Google’s visualisation framework is a notable example that has allowed much

Community effort is required to share visualisation components more mashing of visualisations – can we envisage richer variations of this approach?

7.3.2 Applying Psychological Theory to Real Applications

As discussed in Section 7.2, there are substantial amounts of existing and emerging perceptual and cognitive knowledge that is highly relevant to visual analytics. For example, in medical analysis, slides or X-ray plates are examined for signs of tumours. Eye gaze data of expert radiographers has shown that important artefacts were looked at but did not get reported by the person doing the analysis; that is the expert's level of visual attention was a better estimate of the presence of a tumour than the conscious choice. In other kinds of visual analytics is this likely to also be an issue? This example raises issues concerning foveal vs. peripheral vision, the importance of the latter in many forms of visual processing is only recently being understood and could be a powerful asset in visual presentation. Furthermore, implicit learning and more conscious reasoning are separate human processing systems. Often people have a gut reaction or make an instinctive decision but the process is unclear; modelling this type of decision making is a difficult problem, and this kind of issue is only recently being addressed in 'dual space' theories of cognition.

Make existing knowledge on perceptual and cognition more accessible as well as promoting new research in this area

At a more detailed level, this relates to design decisions such as ratios between blank space and 'used' space, text lists vs. graphic display, and display aesthetics. There is a real challenge to mine the literature to bring out these more general issues that are often buried in papers describing specific techniques and systems. However, some new fundamental knowledge is also needed. One such case arose in connection with dot densities as found in dense scatterplots. One of the key measures used in assessing perceptual stimuli is 'just noticeable difference', for example, when two slightly different sounds are played, when do subjects cease to notice that they are not the same. However, whilst the data was present for visual stimuli relating to solid blocks of colours or shades, no similar data was found in the base psychological literature for dot density. In order to assess the acceptability of differing sub-sampling regimes, Bertini and Santucci[15] had to perform fresh experiments to determine just noticeable difference for dot density.

7.3.3 Understanding the Analytical Process

We still do not have a complete understanding of information visualisation let alone visual analytics. There are examples: Illuminating the Path[111] cites several frameworks for understanding the sense-making and analysis process, Yi et al.[127] characterise the process of gaining insight, and work by de Bruijn and Spence[34] considers different classes of browsing (search, opportunistic, involuntary and perusal) and suggest interaction modes to support such behaviour.

A key challenge is to extend such frameworks to consider the entire analytic process. Most of the frameworks are focused on the stages of visual exploration, although the data/frame theory of Klein et al.[68] also considers the mental representations of the analyst and Pirolli and Card[85] have a dual loop model of the sense-making process, which takes into account the way mental schema give rise to hypothesis and interact with the exploration of external data. However, there seems to be an absence of a) consideration of the visualisation and understanding of the parameter setting of computational processes, and b) the externalisation of the analysts mental representations. Both are connected to distributed cognition and ecological perception (see Section 7.2) as a) they should take into account that the perception of a visualisation is related to the (interactive) setting of the parameters that gave rise to it, and b) the externalisation of mental constructs makes them available for perceptual and more explicit critique. Understanding this more clearly can help suggest appropriate interaction mechanisms, for example, making use of design rationale, annotations, history and provenance.

Further understanding of the analytic process is needed in order to design appropriate interaction mechanisms

7.3.4 The Need for Design Guidelines

There is a need to create clear design guidelines for designers of visual analytics systems and also the means to share practical design knowledge. Many writers on both visualisation and machine learning are wedded to their own particular techniques, so it is often hard to obtain unbiased views of the adequacy of techniques beyond the advocacy of their proponents. This is always a problem, but as the computational side of visual analytics is more complex, having clear advice is correspondingly more important.

There are some steps in this direction including work on design patterns for information visualisation and for visual analytics. However, given that visualisations are applied/developed to support specific contexts, how can design knowledge be re-used in alternative domains? As an example, whilst the concept of Fisheye interfaces goes back more than 20 years, there are still calls for clear design guidelines. This is emphasised in a recent review of challenges in information visualisation, which suggests that the entire field of information visualisation is in the pursuit of finding the most effective visual metaphors. The authors point out that one single metaphor is unlikely to overcome the problems of cognition (intuitive from a users point of view), very large datasets and/or a high number of dimensions. A closely related challenge is the choice of an optimal level of visual abstraction (e.g., from the low level 1-to-1 correspondence of a scatterplot to high levels that involve clustering); however, as with visual metaphors, the choice is very dependent on the user and their experience, knowledge and goals.

The effectiveness of different visual metaphors and levels of abstraction is very dependent on the user

To date, efforts to communicate design knowledge have tended to focus on the re-use of pre-designed solutions. However, the need to design visualisations that reflect contextual system constraints restricts the utility of this approach. Rather than prescribing design solutions the development of related taxonomies of cognitive work systems and appropriate design methodologies have been proposed. These can be used by a designer to classify contextual problems and to identify relevant design artefacts that will support the overall visualisation design process. While these taxonomies are only briefly outlined, their extension could enable and encourage the re-use of design knowledge across different work domains.

Interaction designers need guidelines based on underlying perception and cognition research

Cognitive empirical and theoretical knowledge is continually developing, but we cannot wait for this to be fully developed, but instead must create good engineering advice based on existing knowledge and update this as knowledge develops. Spence[105] highlights the need for ‘brokers’, people who are able to identify important factors in the perception and cognition literature and interpret these for the benefit of interaction designers. One example of this are ’design actions’ that provide guidelines for some specific cases[34]; and the design patterns mentioned above also can be seen in this light.

Another example is the Ecological Interface Design framework. This provides visual design guidelines that support specific levels of cognitive control, including diagnostic activities[21]. While the framework has been validated across a range of complex process control systems, its applicability to intentional decision making and analytical model building requires further investigation.

7.3.5 Defining the Language of Visual Analytics

Clear definitions are essential for the advancement of science, but many of the concepts used in visual analytics do not have precise definitions. While anything involving human capabilities inevitably has fuzzy edges, there is a clear need to attempt to develop clear definitions of core concepts, subject to understanding the limits of such definitions once formulated.

Definitions are important

One example is the concept of insight. Insight is an important concept for the perceptual and cognitive analysis of interaction with visual analytics methodologies reflecting the importance of reasoning processes in the task of interpretating large amounts of data, however, there is as yet no precise and systematic definition of this concept. Valuable input can be gained from research in cognitive psychology[106]. This research is influenced by Gestalt psychology, which conceptualises reasoning as a holistic and structured process.

Even the term ’visual analytics’ is itself potentially problematic. The adjective ’visual’ suggests the use of sight only whereas visualisation is the action of creating a mental model (in the user). Modalities other than visual are important as perception is a holistic process, encompassing sound, touch, smell and taste; these modalities should also be considered in visual analytics. It is evident that these other senses, most significantly aural representations, have a part to play, but this needs to be emphasised lest the term visual analytics accidentally marginalises them. There is also confusion between information visualisation and visual analytics suggesting that greater clarity is required to explain the new issues that arise.

One step in the direction of greater clarity is Thomas's[110] discussion paper on a proposed taxonomy for visual analytics. This effectively creates a lexicon of key terms and some structure to them, but each term really needs a complete definition.

7.3.6 Observability and Trust

There is a need for the user to be made more aware of the visual analytics process in order to gain confidence in the results; for example, business intelligence is commonly used in the context of a decision support system and suffers from poor user acceptance, with the user often ignoring evidence in favour of (potentially biased) past experience. It is suggested that this mistrust in the outputs of a decision support system may be overcome by making the users more aware of the automated decision making process. This is exacerbated by the fact that people are not necessarily good statistical thinkers and so a challenge of visual analytics is to make the statistical methods understandable so the user has enough confidence in the results to counter biased opinions. As a further example, when comparing dynamic queries[2] with Attribute Explorer[118], filtering throws away data and potentially makes it more difficult to obtain a mental model of the data. On the positive side, selecting a facet and filtering the results at least gives the user an idea of the amount of data related to that facet. Note that while faceted browsing allows the user to rapidly reduce the amount of data, the Attribute Explorer is unusual in that its miniature histograms allow you to see an overview of the complete dataset and also assess potentially what may happen as you make further parameter filtering selections. As noted earlier in Section 7.2.6, in general it is rare for visualisation to give a 'glimpse over the horizon' or some idea of the potential results of applying a filter before actually doing the filtering.

Understanding the analytic process can give confidence in the results

One particular issue is the visualisation of uncertainty. Uncertainty takes many forms and some can be estimated quantitatively such as the statistical variance of estimates, but others are more qualitative such as the level of confidence you have in a particular data source (e.g., BBC news vs. sales literature of a competitor). The dual space understanding of cognition is critical here as humans have a primitive-stimulus response learning system that learns through repeated exposure. This effectively learns probabilities, but very slowly. The other mechanism is our more explicit memories and reasoning over them through abduction. This higher-order memory gives us one step learning of new situations, but is relatively poor at probabilities, without explicit mathematical analysis. A challenge for visual analytics, is to use machine processing and visualisation, to complement the human analyst's abilities in understanding uncertainties.

Users should be made aware of sources of uncertainty in the data

7.3.7 Evaluation of Novel Designs

As discussed in Section 7.2 and Chapter 8, issues of evaluation are a hot topic within the information visualisation community with regular workshops on the topic[13]. Many different methods are used to study information visualisation methodologies[26], but more work is required to determine which of these methods are especially appropriate for visual analytics. It is an open question whether traditional methods of cognitive psychology or HCI are appropriate for the investigation of perceptual and cognitive aspects of visualisation. It maybe that we should develop entirely new methodologies that take into account both complexity of detail and context.

Evaluation of visual analytics is particularly difficult as problems are often ill-defined and open ended

Part of the problem is that visual analytics is about solving open ended problems and so it is hard to create meaningful tests as almost by definition these will be for known solutions: puzzles not problems. This is a problem in both research and real world application. For example, the management science literature on sales forecasting systems focuses almost entirely on 'accuracy' as the key evaluation parameter. However, as discussed in Section 7.2, forecasts affect the stocking, placement and advertising decisions of a company and hence sales (the full outer circle in Figure 7.1). That is, the visual analytics within this is itself part of the process it is predicting.

Problem solving involves gaining insight, and this occurs at different levels during the problem solving process. So we need to think about assessing the effectiveness of a design (in terms of interactivity and visualisation) on the generation of insight in: a) assessing the data and finding relationships, b) the capability to support hypothesis formulation, and c) how well the conclusions reached by the user at each stage of analysis can be traced so they can be verified by others.

New technologies may help evaluation

New technologies such as eye-tracking and even brain scanning, offer the potential for radically different ways of approaching evaluation. However, these are themselves areas of substantial complexity, for example, one problem is the appropriate interpretation of the data gained from eye-tracking studies and the definition of clear variables that can be measured by eye-tracking. It maybe that in the short term we need visual analytics to actually address some of the research challenges in these areas as they offer visual analytics new tools of investigation. It is an open question whether these new techniques actually offer any additional information than more qualitative methods such as cooperative evaluation. In general, evaluation of all kinds is also expensive and so in the world of practical visual analytics system design we also need low cost/resources methods.

7.3.8 Designing for the Analyst

Flexible designs allow analysts to customise the way they work

Visualisation designers have their own ideas about what constitute good designs and visualisations, and build these assumptions into the tools created for the analysts. However, analysts often do not think the same way as designers. While there may be a need for some standard tools for standard tasks, the challenge for the community will be the development of advanced tool sets for the analysts. These would enable the analysts to bring different functional capabilities together, enabling them to create visualisations and interact with them in ways that are flexible yet robust. This presumes we have a good understanding of the information handling strategies that users invoke when working with the different sorts of data and documents. Liu et al.[72] also points out that many of the current visualisation systems are not flexible enough to allow user customisation and hence may inhibit the analysts managing data sources and hypothesis creation in a way their feel as appropriate.

Dealing with biased opinions has already been noted as a problem. The work of an analyst is influenced by a host of cognitive biases (e.g., confirmatory bias and anchoring) and many of these biases are often set in motion by the way information is 'fed' into the perceptual-cognitive processes. How do interactive visualisations (designs) influence biases? We need to know and understand the effect of information designs that combine interactivity and visualisation on interpretation and analysis, and the inter-play of that with known cognitive and perceptual biases. The way in which a system presents patterns and cues, and how their significance and salience are rendered, can activate biases. Therefore, it is important to be aware of when they may occur and then develop appropriate controls to minimise such effects.

Need to be aware of and minimise cognitive and perceptual biases

7.3.9 Changing Interfaces: Users, Data and Devices

In current practice, the mathematical models used in decision support are processed offline and only the results are visualised by the user. There is a need to make this process more dynamic both in terms of parameter setting and also the choice of models; however, this will create demands on the underlying visual analytics architectures. Looking at the choice of visualisations, some are highly information intensive, but also very complex, whereas others give less information, but maybe more informative for a novice. There is a real challenge in adapting these visualisations to suit the user and the data, whether under direct user control or semi-automatically; and furthermore to transition smoothly between different levels of visualisation complexity. Similar issues arise when dealing with different devices and hardware from mobile phones to wall-sized multi-screen displays.

Systems need to adapt to a wide range of users, data types and sources, and input/output devices

In the business intelligence world, visual analytics is often presented as a set of visualisations (e.g., treemap, heatmap) from which people with 'data overload' can select an appropriate solution, with little consideration of either the problem to be solved or the process required. We clearly need to be able to offer more guidance as to which methods are better suited to particular classes of problems. The issue here is not the kinds of raw data (time series, categorical, network, etc.), but what we want to do with the data. Furthermore, there are different levels and timescales of problem solving in business (e.g., financial, sales) from everyday decision making to longer term corporate policymaking. Visual analytics is typically applied to 'bigger' decisions, but many systems do not take into account the long-term use and re-use, such as means to annotate past use to inform future interactive sessions. The use of visual analytics for much more moment to moment decision making is perhaps even more problematic and would likely require some automatic aid.

Users need guidance in choosing an appropriate visual analytic solutions for a given task

The Web was designed to ship fairly traditional data from CERN to physicists across the world. However, the Web has more recently given rise to very large-scale data such as folksonomies and tag data, co-occurrence data used in recommender systems and RDF ontologies for the semantic web. Web data presents new problems being both large scale, but also typically less-tabular, and more relational; in the case of semantic web there is the potential for inference and data to become intertwined. As with visual analytics itself, we can easily find 'Web-like' data before the Web, so there are places to look for inspiration, but certainly this is likely to pose fresh challenges for large scale visual analytics in the years to come.

7.4 Next Steps

From the previous sections, we can identify several necessary actions in order to progress understanding of human aspects of visual analytics:

- appropriate design methodologies need to be developed taking into account all the human issues impacting visual analytics as discussed in Section 7.2, the heterogeneity of devices and data as discussed in Section 7.3.9, and range of stakeholders (Section 7.3.1)
- these need to be backed up by design guidelines and clear definitions, especially for non-expert users of visual analytics systems (Sections 7.3.1, 7.3.4 and 7.3.5)
- of particular importance are the development of interaction and visualisation mechanisms that will enable analysis to assess more confidently the reliability of results of visual analytics systems, including issues of uncertainty and provenance of data (Section 7.3.6)
- these need to be backed up by appropriate evaluation mechanisms, potentially including emerging techniques such as eye tracking (Section 7.3.7)
- all of the above require an ongoing development of the basic human science of visual analytics including brokering existing fundamental psychological and social knowledge, generating new such knowledge and most importantly creating robust and applicable holistic models of the visual analytics process (Sections 7.3.2 and 7.3.3)

In general, the topic of perceptual and cognitive aspects of visual analytics is highly interdisciplinary and these very heterogeneous disciplines provide interesting input for visual analytics. Whilst we have gone some way in establishing contacts between these communities, there is much still to accomplish.