The Broker

Bob Spence

Department of Electrical and Electronic Engineering,
Imperial College London,
SW7 2AZ, United Kingdom
`r.spence@imperial.ac.uk`

Abstract. The task of a Broker is to interpret relevant knowledge acquired by cognitive and perceptual psychologists and bring it suitably to the notice of interaction designers, thereby avoiding the need for that designer to have knowledge of cognition and perception. The task is first illustrated by an example based on the concept of Design Actions and demonstrates the implication, for two different design challenges, of certain properties of the human visual processing system. It is then argued that the task of the Broker can be eased by the definition and classification of relevant concepts, in the illustrative example those of browsing, interaction and visualization. Finally, a current need for a Broker's expertise is illustrated in the context of the interactive and dynamic exploration of the relationships associated with a multivariable system.

To explain the title of this contribution I must point out that I'm an engineer who first became an interaction designer, with a concern for visualization, in 1968 when I started to design (Figure 1) an interactive-graphic CAD system for electronic circuit design [1] that eventually (Figure 2) became a commercial product in 1985. During that time I was very much aware that the perceptual and cognitive abilities of the human being – the circuit designer in my case – would be ignored at one's peril. Few others in the industry appeared to share my view but, as an engineer and not, definitely not, a psychologist, I had to get things done – in my case the creation, with colleagues, of the first commercial interactive-graphic CAD system for electronic circuit designers (MINNIE). I therefore had some sympathy with Isambard Kingdom Brunel because I did not have, at my disposal, all – or even a little – of the knowledge about human factors that would better inform my design. So a great deal of intuition had to be applied, and we know how dangerous a substance that can be. I did involve psychologists in my research to a limited extent, but their activities could only provide a very small fraction of the knowledge I was seeking.

So for a very long time I've been aware of the need for some means of interpreting relevant knowledge acquired by cognitive and perceptual psychologists and bringing it *suitably* to the notice of interaction designers. I shall call the person responsible for doing that The Broker.

To illustrate the task of a broker and how it might be supported I first present an example; to the same end I then advocate more attention to definitions and

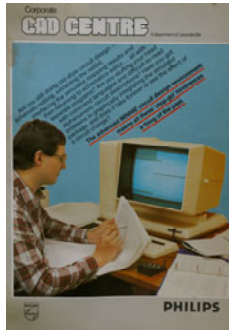**Fig. 1.** The MINNIE interactive-graphic CAD system, 1968

Fig. 2. The commercially available MINNIE system for CAD, 1985

classification; and, finally, I illustrate a current need for action by a broker. Because a keynote is often based on personal experience I make no apology for the fact that many of the references are to the work of myself and my colleagues.

1   An Example

My psychologist colleague Dr. de Bruijn and I had an idea that required the services of a broker. The idea was based on the fact that an interaction designer has to design an interface to support some type of human behaviour. The designer knows that such behaviour is controlled by cognitive and perceptual processes associated with the human user, *but may have little or no understanding of that subject.* Let's take an example of behaviour – that of search browsing. That behaviour will certainly be affected by the human visual processing system. However, as I have emphasised, the interaction designer almost certainly has no knowledge of such processes. So our view was that he or she would benefit from suitably presented guidance that is based on knowledge about the human behaviour that has to be supported by the designed interface. We suggested that it is possible – and this is where the Broker comes in – to identify the human processes relevant to the behaviour to be supported – one of which is human visual processing – and generate what we called Design Actions [2] (Figure 3).
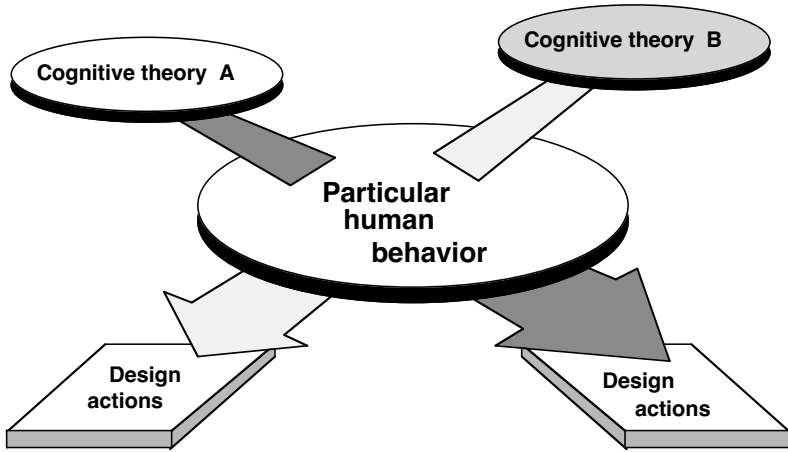
Fig. 3. The conceptual basis of Design Actions

They tell the interaction designer how to create a good design and they relieve that designer of the need to understand human cognitive and perceptual processes.

Very briefly, Design Actions are tables that provide design guidance, and at a useful level of detail. Figure 4a shows the structure of a Design Action. First there is a brief title, followed by a clarification of that title. The next line says that if you, the interaction designer, take a particular design decision, a particular outcome can be achieved. Next, because any design decision has its pros and cons, it is helpful for the designer to know what they are. Then there are related issues that the designer might beneficially consider. Finally there is a reference to the HCI principles involved in case the designer is curious. Figure 4b shows one example of a Design Action relevant to a situation in which you want a user to respond very rapidly to the appearance of a target image. Design Actions do not, of course, constitute the only methodology available to the interaction designer: a comparison with Alexander Patterns [3] is provided in [2].

The concept of Design Actions is extremely general – that's one of its potential attractions. To illustrate that potential we focused [2] on a particular behaviour, that of browsing, with two applications in mind. One concerned the rapid search, on a mobile, for an interesting news item among a collection of such items, using the technique of Rapid Serial Visual Presentation (RSVP) [4]. The other concerned involuntary browsing (defined later) where, for example, small icons moving very slowly around the periphery of a coffee table (Figure 5a) [5] might catch the *latent* interest of a person sitting at that table. This table was originally designed to be located within an area such as a pub or library: the little icons refer to activities going on in the local community: sight of one of them might trigger some latent interest. Pushing the little icon (Figure 5b) into the centre provides more information, in this case about a chess tournament.

Title	Design objective
Description	Clarification of title
Effect	"If you do this, that will happen"
Upside	The advantages of doing it
Downside	The disadvantages of doing it
Issues	Considerations
Theory	Reference to theoretical basis of Design Action

(a) The format of a Design Action

ID	DA2
Title	Presentation for immediate action
Description	If selectable information sources are each represented by an image (with or without brief text), ensure that each item can be fixated for at least 100ms, and that an immediate response can be initiated as soon as a relevant item is encountered.
Effect	This action gives reasonable certainty that the meaning of each fixated item can be established and that items that match a relevance threshold can trigger an appropriate action
Upside	Many information items can be presented in a relatively short period of time either sequentially or concurrently
Downside	The lower limit of 100ms will, for a given application, establish an upper limit to the number of images that can be presented concurrently with the expectation that, during the normal - and often random and involuntary - eye gaze activity, a relevant one will be identified.
Issues	(1) undirected/random eye-gaze activity is such that only a very short (e.g., 50 to 300ms) fixation on an information item may take place, further supporting the need for pictures of familiar objects and scenes. Conscious attention should not be necessary to recognise words and images within the information items. (2) The actions that are triggered in response to the interpretation of items need to be simple so that they can be executed without delay and consideration.
Theory	The ability to automatically and rapidly match an information need with a viewed image.

(b) An example of a Design Action

Fig. 4. The format of a Design Action and an example

The question we had to address was "what knowledge about human visual processing exists that will help us with these two quite different design challenges?" Answering that question is the task of a broker.

Human visual processing was obviously fundamental to both these applications. For that reason we found the book *Fleeting Memories* [6] to be most helpful, and especially the chapter [7] by Mollie Potter of MIT who was responsible for proposing the concept of Conceptual Short Term Memory (CSTM). Basically Potter proposed the CSTM model – our own illustration of which is shown in Figure 6 – to describe the processing of a visual stimulus in the first few hundred milliseconds after its initiation. Almost immediately following the arrival of that stimulus in sensory storage it is categorized by reference to previously stored knowledge in long-term memory (LTM): it may be categorized as 'a cat in front of a house' or 'a briefcase' or 'coloured wiggly lines'. That categorization
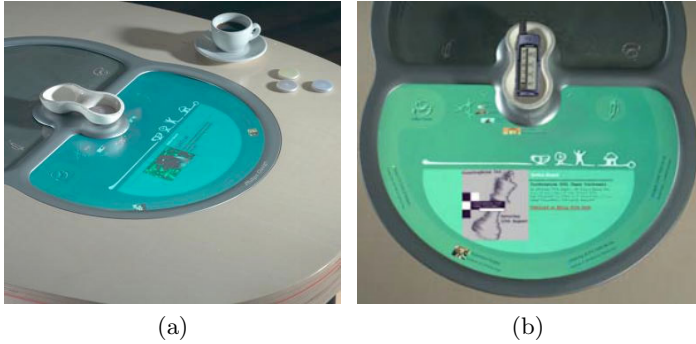
(a)   (b)

Fig. 5. A coffee table designed to trigger the latent interests of the coffee drinker
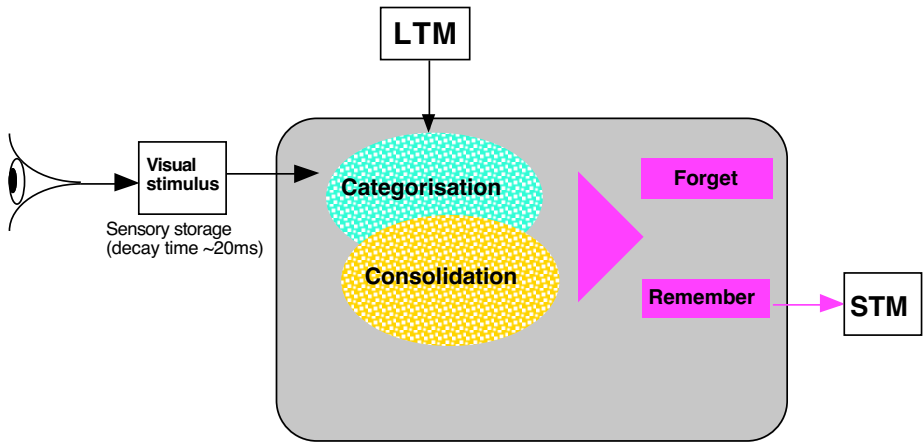

Fig. 6. A representation of Conceptual Short-term Memory

may well be 'primed' [8] by the task being performed: it is thought to occur in about 100ms and is performed essentially unconsciously.

The meaning or relevance of the stimulus is then interpreted, again unconsciously, in a process called consolidation. If there is relevance it is passed to short-term memory (STM); if there is not, then forgetting occurs. Potter points out that the whole cycle – identification of stimulus, memory recruitment, structuring, consolidation and the forgetting of irrelevant data – may occur in less than a second and, essentially, is achieved unconsciously.

We then examined various experiments that had been carried out by cognitive psychologists over a number of decades to see how they could be associated with CSTM. I'll identify three (Figure 7), although we made use of seven. One is the discovery [9] that if images are presented at a rate as fast as 10 per second, the search for a known target image will be essentially successful. Another is that if two targets are being sought during such a presentation, one will not be
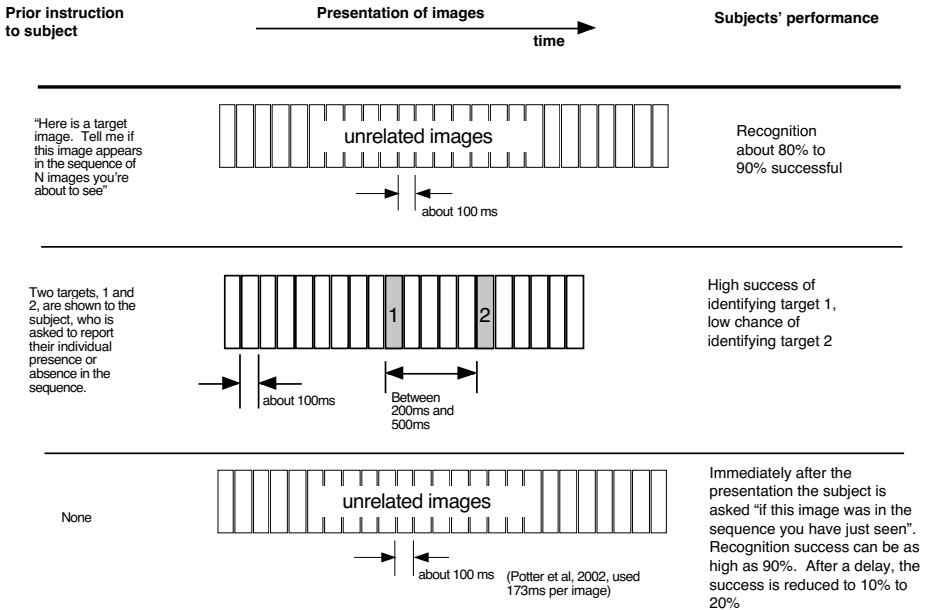
Fig. 7. Three experiments concerned with human visual processing

detected if it occurs shortly after the first to be detected, a phenomenon known as Attentional Blink [10]. A third concerns one's memory of what has been seen that had not been seen before – in general it is pretty poor [11].

To cut a very long story short, reference to these experiments enabled us to associate, with the original CSTM model, a scale indicating the earliest permissible onset of a following image if a certain task had to be successfully performed with the current image. I'll give three examples. If a 'satisficing' task is involved, a following image should not appear until 100ms after the appearance of the preceding image. If the task is to select the most appropriate image from a collection – we called that an 'optimising task' – then the separation must not be less than 500ms. If the task is of the 'preview-consider-select' type, the separation must be at least 1500 ms.

Using this model, we were able to derive, for the activity of browsing, various Design Actions [2] that should be brought to the attention of an interaction designer when designing, in one case, the RSVP presentation of news items and, in the other, a coffee table presentation of icons representing activities within a local community. The outcome was encouraging in the sense that what the extended CSTM model suggested turned out experimentally to be reasonable, though the framework we proposed requires much more use and evaluation.

Obviously there are hidden dangers in what we did as brokers, because psychologists are still trying to unravel the mysteries of human visual processing. Also, there are many other factors relevant to interaction design that we have not yet taken into account. Indeed, for any particular behaviour such as finding

a news channel of interest on a mobile, there will be more than one relevant cognitive theory and therefore more than one set of Design Actions. However, I'm talking here about engineering design: and as a designer I need answers *right away*. So I feel that HCI professionals have a duty to provide models – however approximate – that will inform design. Having something reasonably well understood to go by is better than having nothing. Brunel did not, to my knowledge, have access to a finite element analysis package, but he produced some remarkable results.

2   Definitions

In this second of three sections I want to talk about definitions. I have a view – a philosophy, if you like – as to how they can facilitate the transfer of knowledge from cognitive psychology to the interaction designer.

My opinion is that although we are all of a scientific persuasion we seem to shy away from precise definitions. Take the term *browsing*, for example. Ask an HCI specialist what is meant by this frequently used term and you often get a very fuzzy reply (usually containing the term 'casual') and a lot of arm waving. I think that is not good enough: and goodness knows what students trying to adopt a scholarly approach to HCI will think. It turns out that in the work I've just described, reasonably precise definitions and taxonomies were of considerable help to the brokers. For example, based on Norman's Action Cycle [12] we defined three types of browsing using, as their defining parameters, Goal, Intention and Action Plan. The three types of browsing we defined are Search Browsing, characterised by a goal, a conscious intention and an action plan; Opportunistic Browsing with no immediate goal, but with an intention and an action plan; and Involuntary Browsing which is characterised by a latent goal [13] (of which we each have many) and therefore no conscious intention or action plan.

With these reasonably precise definitions and classification – though acknowledging that the topic of browsing is still under discussion – we were able to be quite specific about the Design Actions we derived.

I have two more examples where I feel that definition is important. The first concerns the definition of 'visualization'. I'm frankly surprised that many experts seize upon the 'visual' in that term and, for example, totally ignore the encoding of data in sound, a tendency that is not helped either by the new term 'Visual Analytics' [14] or by the legacy of our past computer-graphics-based approach to visualization. As I've said on many previous occasions, visualization has nothing to do with computers and is simply defined in some very respected dictionaries as

Visualise: to form a mental image of something.

As a consequence we have (at least) four classes of visualization (vision, sound, touch and smell) depending upon the human sense by which encoded data is perceived.

Interaction is another term we frequently use, so we had better find some classification scheme. We will certainly have to if we're to achieve the goal, described by Thomas & Cook [14] of creating a Science of Interaction. My own suggestion for a classification of interaction [15], which I have found to be helpful for interaction design, is again based on dictionary definitions and leads to three types:

Stepped interaction (such as a mouse click) which causes a discrete movement in information space,

Continuous interaction, as occurs during scrolling, for example, and

Sensory interaction: here something happens between user and computer that is experienced by one of the senses (e.g., reading a book, inspecting an image).

We can use classifications of visualization and interaction to form a two dimensional space (Figure 8) which, for emphasis, has been simplified by the omission of incidental interaction [16] and the mention of the status-event concept [17]. We can then perhaps introduce our browsing classification to create a 3-dimensional space. I think such a space can be very helpful when seeking frameworks to support interaction design. What I think is unfortunate is that the scope of this space is not fully explored by researchers seeking a framework to inform interaction design, first because the very term Visual Analytics restricts the sense involved in the perception of encoded data; second because sensory interaction does not seem to be accepted as a class of interaction; and third because involuntary browsing receives little attention. By ignoring such a large part of this space there may be a danger that any Science of Interaction that might be discovered will be rather impoverished.
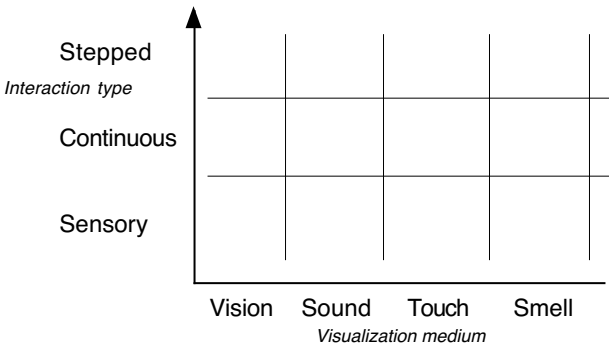

Fig. 8. The interaction-visualization space

So my plea, to all HCI practitioners, initially concerns science rather than engineering: "please can we have more precise definitions of familiar terms, because we can then more readily derive some frameworks that can inform interaction and visualization design". In any case there is no excuse for saying "everyone knows what we mean by. . . " – if everyone *does* know, then let's write it down.

3   Need

In this third section I want to identify – by an illustrative example – a specific need for the services of a broker. Again I base my example on work I have done with colleagues.

In 1970 we had devised an efficient algorithm [18] that allowed us to implement a scheme [19] whereby the designer of an electronic circuit could first point to a component whose effect on the performance of a circuit must be explored: then, following execution of the algorithm, the component's value could be varied manually and its effect *immediately* seen. The aim of that interface was to support exploration and the improvement of the designer's mental model.

Later, in 1996, my group went on to invent the Influence Explorer [20] – another interface designed to support the interactive exploration of data – though again without the benefit of knowledge of any relevant psychological literature. Let me first, and again necessarily briefly, illustrate the idea.

An electric lamp is to be designed [21]. It contains, internally (Figure 9), a structure supporting a filament. The structure is described by four dimensions (X_1 to X_4) that the lamp designer can choose. Interest lay in the effect of these dimensions on four stresses within the lamp (S_1 to S_4). The mathematical model (Figure 10) expressing all Ss in terms of all Xs provided no insight whatsoever for the lamp designer. So we (easily) simulated 400 random *but possible* designs and displayed the results (Figure 11) in the form of histograms. Alone they don't
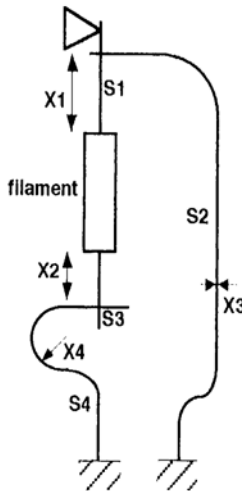

Fig. 9. A structure supporting the filament of an electric lamp

$$S_4 = 60.4 + 23X_1 - 3.8X_2 + 631.2X_3 - 26.4X_4 - 79.7X_1X_3 + 4.8X_2X_3 + 2.6X_1X_4 - 2.6X_1^2 - 278.2X_3^2 + 5.7X_1^2X_3$$

Fig. 10. An expression relating a structural stress to four designable dimensions
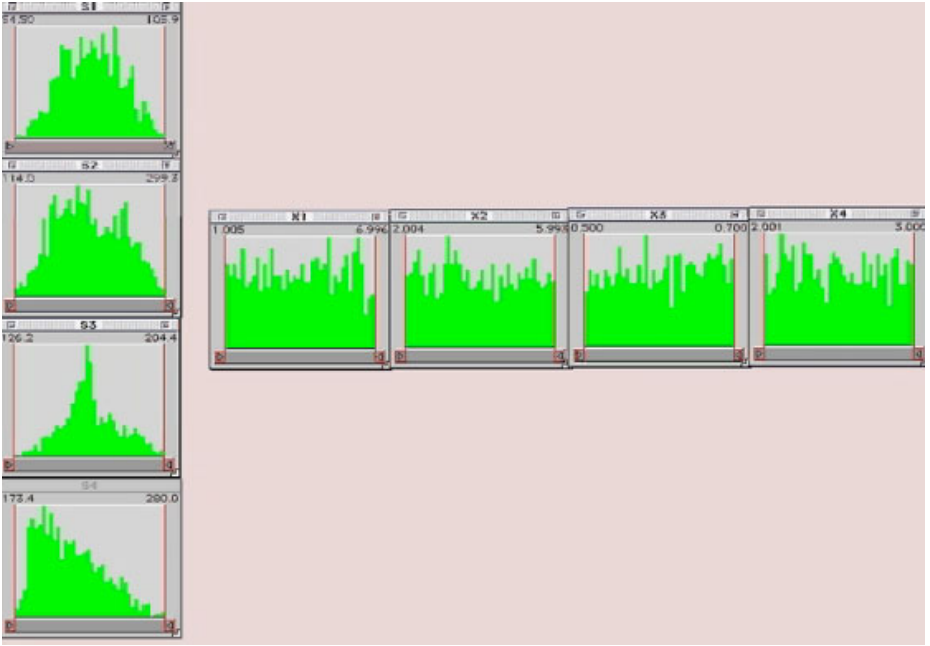
Fig. 11. Interactive histograms showing the result, for four stresses in the structure of Figure 9, of a random selection of designs (i.e., choices of X_1 to X_4)

tell you very much, but they certainly can if interaction is made possible. For example, if the designer interactively selects those designs that have low values of S_4, and especially move the range of selected values of S_4, they immediately notice a trade-of with S_3 (Figure 12) as well as a correlation with S_2 and with the designable parameter X_1. The mental model thereby gained by the designer of that lamp was immensely valuable and rapid. The same principle can be applied to many other disciplines including, for example, financial design.

What I have described is a very powerful technique allowing an investigator to gain insight into the complex relationships between variables and functions of those variables in a wide variety of multi-variable systems. Not surprisingly, some people are developing this idea [22] because the potential it offers is considerable. But I don't know – except from intuition based on experience – how best to design such an exploration tool. So I want the answers to many questions having to do with navigation, representation, 'near miss' information, the influence of task, how exploration can be automated [23] (in which case what is a good representation?), the awareness of desirable and undesirable regions of multidimensional space [24], design to provide useful overviews, useful computations – the list is long. These questions are not restricted to the Influence Explorer: recent exciting developments in interactive exploration [25] pose similar questions.

It may well be that some answers are already available in the psychological literature, in which case I want to know about them. If not, our design of
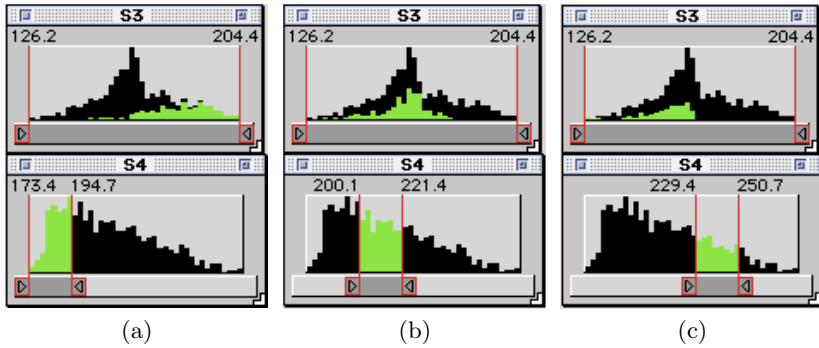
Fig. 12. The effect of moving the range of selected values of S_4 on the distribution of corresponding values of S_3, for the structure of Figure 9

these exploratory tools might not be anywhere near as successful as they could otherwise be.

4   Conclusions

My conclusion is simply stated: that there's a need to package knowledge gained by psychologists in a form that will inform an interaction designer, and that to do so effectively we must pay more attention to definition and classification.

References

1. Spence, R., Apperley, M.: The interactive-graphic man-computer dialogue in computer-aided circuit design. IEEE Transactions on Circuits and Systems 24, 49–61 (1977)
2. De Bruijn, O., Spence, R.: A new framework for theory-based interaction design applied to serendipitous information retrieval. ACM Transactions on Computer-Human Interaction 15, 1–38 (2008)
3. Borchers, J.O.: A pattern approach to interaction design. In: Gill, S. (ed.) Cognition, Communication and Interaction, pp. 114–131. Springer, London (2008)
4. Spence, R.: Rapid, serial and visual: a presentation technique with potential. Information Visualization 1, 13–19 (2002)
5. Stathis, K., de Bruijn, O., Macedo, S.: Living memory: agent-based information management for connected local communities. Interacting with Computers 14, 663–688 (2002)

6. Coltheart, V. (ed.): Fleeting memories: cognition of brief visual stimuli. MIT Press, Cambridge (1999)
7. Potter, M.C.: Understanding sentences and scenes: The role of conceptual short-term memory. In: [6], ch. 2, pp. 13–46
8. Neely, J.H.: Semantic priming effects in visual word recognition: A selective review of current findings and theories. In: Besner, D., Humphreys, G.W. (eds.) Basic Processes in Reading: Visual Word Recognition, pp. 264–336. Lawrence Erlbaum Associates Inc., Hillsdale (1991)
9. Potter, M.C., Levy, E.I.: Recognition memory for a rapid sequence of pictures. Journal of Experimental Psychology 81, 10–15 (1969)
10. Raymond, J.E., Shapiro, K.L., Arnell, K.M.: Temporary suppression of visual processing in an rsvp task: An attentional blink? Journal of Experimental Psychology: Human Perception and Performance 18, 849–860 (1992)
11. Potter, M.C., Staub, A., Rado, J., O'Connor, D.H.: Recognition memory for briefly presented pictures: the time course of rapid forgetting. Journal of Experimental Psychology: Human Perception and Performance 28, 1163–1175 (2002)
12. Norman, D.A.: The Design of Everyday Things. Basic Books, New York (2002)
13. Zeigarnik, B.: On finished and unfinished tasks. In: Ellis, W.D. (ed.) A source book of Gestalt psychology, pp. 300–314. Trubner & Co., Ltd, Kegan Paul (1938)
14. Thomas, J.J., Cook, K.A.: Illuminating the Path: The Research and Development Agenda for Visual Analytics, National Visualization and Analytics Center (2005)
15. Spence, R.: Information Visualization: Design for Interaction, 2nd edn. Pearson Education, London (2006)
16. Dix, A.: Incidental interaction (2002), http://www.comp.lancs.ac.uk/~dixa/topics/incidental/
17. Dix, A.: Status-event analysis (2010), http://www.comp.lancs.ac.uk/~dixa/topics/status/
18. Leung, K., Spence, R.: Multiparameter large-change sensitivity analysis and systematic exploration. IEEE Transactions on Circuits and Systems 22, 796–804 (1975)
19. Spence, R., Apperley, M.: On the use of interactive graphics in circuit design. In: Proceedings of the International Symposium on Circuits and Systems, pp. 558–563. IEEE, Los Alamitos (1974)
20. Tweedie, L., Spence, R., Dawkes, H., Su, H.: Externalising abstract mathematical models. In: CHI 1996: Proceedings of the SIGCHI Conference on Human factors in Computing Systems, pp. 406–412. ACM, New York (1996)
21. Su, H., Nelder, J.A., Wolbert, P., Spence, R.: Application of generalized linear models to the design improvement of an engineering artefact. Quality and Reliability Engineering International 12, 101–112 (1996)
22. Neufeld, E.M., Kristtorn, S.K., Guan, Q., Sanscartier, M., Ware, C.: Exploring causal influences. In: Erbacher, R.F., Roberts, J.C., Grohn, M.T., Borner, K. (eds.) Visualization and Data Analysis. SPIE, vol. 5669, pp. 52–62 (2005)
23. Colgan, L., Spence, R., Rankin, P.: The cockpit metaphor. Behaviour & Information Technology 14, 251–263 (1995)
24. Tweedie, L., Spence, R.: The prosection matrix: A tool to support the interactive exploration of statistical models and data. Computational Statistics 13, 65–76 (1998)
25. Elmqvist, N., Dragicevic, P., Fekete, J.D.: Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. IEEE Transactions on Visualization and Computer Graphics 14, 1148–1539 (2008)

Videos

It is difficult to convey the considerable potential of interactive exploration via the written word and static image. Some of the effects mentioned in this paper can be better appreciated by viewing short videos. The references to which they relate, and their source, are listed below

[19]: Video V2 on the DVD associated with reference [15]
[20,21]: Video V31 on the DVD associated with reference [15]
[22]: Video V33 on the DVD associated with reference [15]