

A NEW FRAMEWORK FOR THEORY-BASED INTERACTION DESIGN, APPLIED TO SERENDIPITOUS INFORMATION RETRIEVAL

OSCAR DE BRUIJN

The University of Manchester

and

ROBERT SPENCE

Imperial College London

Abstract

The activities of opportunistic and involuntary browsing offer the potential for many of a user's latent problems to be resolved serendipitously, with negligible cognitive effort. In this paper we demonstrate how the design of two novel artefacts to support such behaviour was based on a set of Design Actions which were derived from a model of browsing behaviour in combination with a cognitive model of human visual information processing. We propose the concept of Design Actions as a way of avoiding the need for an interaction designer associated with these and similar artefacts to understand the cognitive theories underlying them.

Categories and Subject Descriptors: D2.2 [**Design Tools and Techniques**]: *User Interfaces*; H5.2 [**Information Interfaces and Presentation**]: *User Interfaces - User-centred design; Interaction style; Theory and methods*; H1.2 [**User/Machine Systems**]: *Human Information Processing*

General Terms: Design, Experimentation, Human Factors, Theory

Additional Key Words and Phrases: Human-computer interaction, Cognition, Opportunistic Browsing, Involuntary Browsing, Serendipity, Design Actions, Rapid Serial Visual Presentation, Table-top interaction, Mobile Web browsing

1. INTRODUCTION

A number of cognitive theories have been devised that are immediately relevant to the design of interactive artefacts for widespread personal and commercial use. As an example, mature theories of human visual processing can explain many aspects of a user's behaviour while searching for a specific image during a rapid sequential display of photographs on a PDA.

Unfortunately the potentially beneficial application of these cognitive theories to the design process is hindered by the apparent need for an interaction designer to understand them. To overcome this drawback, this paper develops an approach to interaction design

This research was supported by the UK Engineering and Physical Sciences Research Council (GR/N20881/01).

Authors' addresses: Oscar de Bruijn, Manchester Business School, The University of Manchester, Manchester, M15 6PB, UK, ; Robert Spence, Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2BT, UK.

Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2008 ACM 1073-0516/01/0300-0034 \$5.00

that allows interaction designers to make direct use of cognitive theories without having to develop an in-depth knowledge of human cognition.

To provide an illustrative context for the development of the new design approach we have selected, and examined in some depth, the common activity of serendipitous browsing. In the remainder of this section we define and discuss serendipitous and other forms of browsing to provide an illustration of the relevance of cognitive models, leading to a new framework having the potential to remove the need for interaction designers to become familiar with cognitive theories. As explained later in this section, a significant product of this framework is the concept of Design Actions.

1.1 Serendipity

Though the acquisition of information to satisfy a goal is a widespread and common activity, it is often far from effortless. Typically, and as Norman's Action Cycle (Norman, [1986]; Figure 1) reminds us, an intention has to be articulated, an action plan has to be formulated and, finally, that plan has to be executed. Available interfaces support these activities with a greater or lesser degree of success but, even so, the user always has to be consciously aware both of the need for information as well as the means

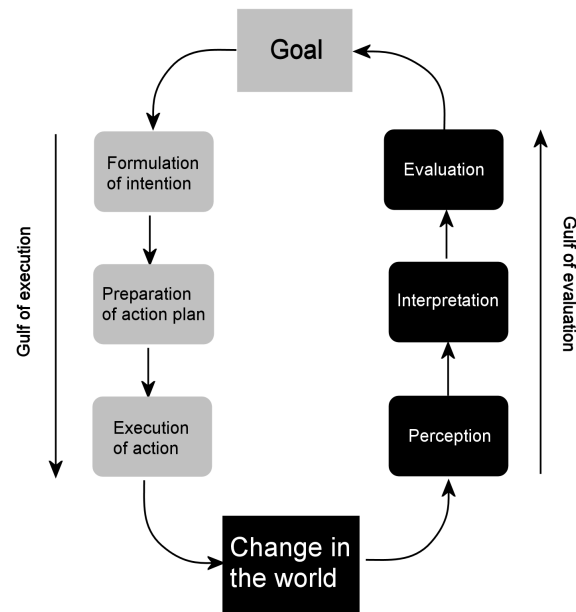


Fig. 1. This graphic shows the actions that span the gulf of execution (left) and evaluation (right) according to Norman's action cycle as it relates to goal oriented information seeking behaviour.

for its acquisition. The accompanying expenditure of cognitive effort restricts both the rate at which information can be acquired and the inclination to do so.

There is an alternative approach to the acquisition of information, and one which requires little or no cognitive effort, even to the extent that other activities can take place simultaneously. We refer to the *serendipitous* acquisition of information which typically occurs when a user's gaze happens to fall upon a representation – perhaps an image and two or three words - of some information of interest. Up until the point at which the relevance of the representation to some long-standing interest is recognised, current theories suggest that no conscious cognitive effort is involved in the discovery of such information (e.g., Lavie, [2000]; see also footnote 6). The cognitively demanding Gulf of Execution (Figure 1) need not be traversed in such a situation. The three stages of the Gulf of Evaluation are traversed (as highlighted in Figure 1), and only if the stage of evaluation identifies what has been seen as relevant to one of a user's many latent interests or goals, does some conscious action ensue.

Many of users' information seeking goals may be long-term in nature, and many may not be easily decomposed into sub-goals that can be translated into immediate action. Furthermore, at the time these goals were formed, suitable courses of action may not have been available or the user may have decided that the effort required to pursue the goal outweighed the benefit of its achievement (e.g., Payne et al. [1993]). In this respect interesting light is shed on long-term unresolved goals by Zeigarnik [1938] who showed that uncompleted tasks are remembered better than completed ones. The fact that a mere glance – perhaps only 100ms in duration – can trigger the awareness of a possible solution to a problem, coupled with the typical abundance of latent interests stored in a user's long term memory, suggests that the serendipitous acquisition of information is a phenomenon that deserves to be studied with a view to its exploitation. That is one purpose of this paper.

1.2 Browsing

The serendipitous activity referred to above can be either opportunistic or involuntary, and it is important to differentiate between these two activities. To do so we define two classes of browsing, where browsing is taken to mean the perceptual registration of visual content (Spence [1999]):

Opportunistic Browsing (OB) is intentional, but the user is unaware of any goal being pursued. The attitude underlying OB is 'let's see what's there'.

By contrast,

Involuntary Browsing (IB) is *unintentional*, though again the user is unaware of any latent goal that might be pursued. An example is provided by a scenario in which a user's eye gaze, naturally moving rapidly between a series of fixations, serendipitously fixates on an information item that may lead to the answer to a longstanding query.

There is a third type of browsing which is not serendipitous in intent but will be discussed later, and it is useful to define it here for comparison:

Search Browsing (SB) is intentional, and the user is aware of the goal being pursued, however precisely or otherwise it is defined. With SB the Gulf of Execution is fully engaged, with the inevitable expenditure of cognitive effort. SB is typified by a web search for specific information.

1.3 Criteria for definitions of browsing

It is helpful if we can set our definitions of OB and IB, as well as justify them, in the context of some of the many studies (e.g., Foster & Ford [2003], Catledge & Pitkow [1995], Carmel et al. [1992]) that have sought an understanding of browsing. About one class of browsing – that which we and others have called Search Browsing (SB) – there seems little disagreement. Such browsing is *intentional*, and the user is *aware of a goal*, however crisply or fuzzily it may be defined. Other forms of browsing have variously been classified as 'general purpose', 'leisurely' and 'casual', to name a few. The definition of general purpose browsing (Catledge & Pitkow [1995]) suffers from the qualification that it involves "consulting sources that have a high likelihood of items of interest", requiring that that likelihood be measured in some way if serious use is to be made of the definition. Equally, there is difficulty in seeking precise definitions of 'leisurely' and 'casual'.

However, useful light was shed on the definition of browsing by Foster and Ford [2003] in their extensive investigation of serendipity. Their use of this term is consistent with ours and includes – but clearly separates – the activities of 'stimulated browsing' (our OB) and 'encountered browsing' (our IB). They clarify their view by remarking that "this paper . . . reinterprets the notion of serendipity as a phenomenon arising from both conditions and strategies – as both a purposive and a non-purposive component of information seeking". Respectively, we have instead used the terms *intentional* and *unintentional*.

Our own definitions of browsing given above are broadly in agreement with those of Foster and Ford, and are adopted because they satisfy criteria which we feel to be essential to underpin any experimental study, namely (1) they either do or do not involve the Gulf of Execution; (2) they hold the promise of being linked to a theory of human visual processing; and (3) they do not assume the outcome of browsing. It is for the latter reason that we would say that OB and IB are not serendipitous *per se*: they only *may* lead to a serendipitous outcome.

1.4 Cognitive models

As we shall review later, current cognitive theory suggests that only 100 ms need elapse before an information item is recognised as relevant, and points to the crucial part that the human visual processing system plays in the serendipitous acquisition of information. Fortunately we know something about this process: over the last few decades cognitive psychologists have investigated the human visual system experimentally and have proposed models. Potentially, these models can be of value in the design of interfaces to support all forms of browsing. We shall in fact review some relevant experiments and a model to which they lead later. The question remains, however, as to how such a model can be harnessed to inform interaction design and thereby lead to effective interfaces.

1.5 Interaction design

However well a cognitive model can explain experimentally observed behaviour, and hence provide a basis for prediction, its value is limited unless it can beneficially influence the design of interfaces to support OB and IB. Notwithstanding the maturity of any model, a major problem is posed by the fact that an interaction designer cannot be expected to be familiar with, and judge the consequences of, any cognitive model. There is therefore a pressing need to develop approaches that allow designers to make direct use of cognitive theories, but without them having to develop an in-depth knowledge of human cognition. The development of such an approach is the principal aim of this paper, and will be illustrated by its application to the design of two novel artefacts intended to support OB and IB.

1.6 A new bridging framework

In an attempt to achieve this aim we present a framework which allows theoretical knowledge of cognition to be expressed in a way that can be applied directly to the design of new artefacts. We do this (Figure 2) by examining the implications of a cognitive theory in the context of a particular human behaviour, and in that way identify potentially

useful design actions that can be applied to the design of artefacts intended to support such behaviour. In that sense, a model or description of this human behaviour is acting as a *bridge* between cognitive theories on the one hand and, on the other, the benefits, drawbacks and other issues associated with a specific design action (cf., Walenstein, [2004]). Using the bridging method to provide an appropriate context, as illustrated in Figure 2, it is possible to translate the experimental results on which the cognitive theory is based into concrete design actions, thereby making the implications of theoretical knowledge directly available to the interaction designer. By using an appropriate format for describing the design actions, the interaction designer would not need to have in-depth knowledge of the cognitive theory in question.

1.7 This paper

In the remainder of the paper we explore this possibility. To do so we select for consideration those situations in which the serendipitous acquisition of information is achieved by arranging for potentially relevant information items to be encoded visually in such a way that their relevance is effortlessly assessed and without the need for an information requirement to be consciously formulated. To this end we identify, in section 3, a cognitive model concerned with the interpretation of rapidly glimpsed images, a model called Conceptual Short-term Memory (CSTM; Potter [1999]).¹ We then review, in section 4, the behaviour of Opportunistic and Involuntary Browsing in which a user may not be consciously looking for a particular image but serendipitously notices one that is relevant to their (long-established) interests. We then, in section 5, examine the relevance of the cognitive theory (CSTM) to this behaviour (OB and IB) in

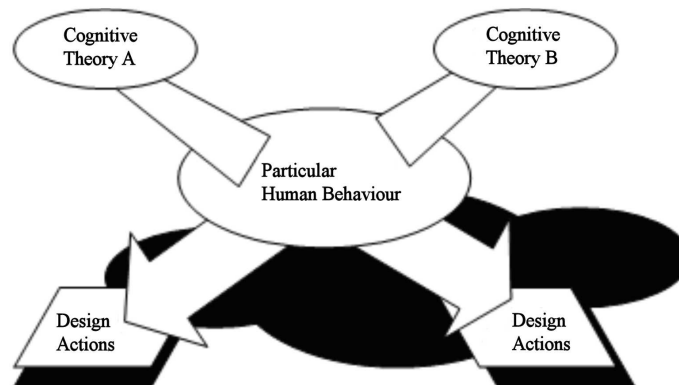


Fig. 2. Design Actions intended to inform the interaction designer are generated by applying cognitive theories in the context of the human behaviour being supported and described in a behavioural model.

order to identify some design actions that may usefully be taken by an interaction designer looking to provide electronic support for these two types of browsing. Although the resulting design actions are shown in section 6 to be useful by reference to recently designed artefacts, the principal aim of this paper is to establish the *framework* (Figure 2) by which design actions are related to a cognitive theory via the bridging mechanism of a model of specific human behaviour. Complete validation of the framework will require many more years of research effort, not just from us. We feel it is important, therefore, to present the framework now, and illustrate how it can be applied to the development of useful design guidance. However, before embarking upon the developments reported in sections 3 to 6, it is useful first to set them in the broader context of the problems posed by interaction design and the precise role that cognitive theory can play in such design.

2. BACKGROUND

As already implied, interaction designers continually search for concepts related to human information processing that could beneficially inform their design activity. While it is generally agreed that our understanding of many aspects of human cognition will remain incomplete for many years to come, considerable progress has been made in the development of relevant psychological concepts and models. In fact, the needs of the interaction designer are sufficiently urgent that even currently available models and concepts which are not fully developed and validated as yet, but which nevertheless are strongly supported by experimental evidence, may be regarded by the interaction designer as sufficiently robust to warrant experimental use. For example, even though research into the detailed operation of memory and attention is still ongoing, it appears that even approximate models of these theoretical mechanisms can already be applied (e.g., Miyata & Norman [1986]).

Therefore, to make use of predictive cognitive theories in HCI the following two problems need to be addressed:

1. We need to develop theories that can predict user interaction with complex multimedia systems. We describe one such theory (CSTM) in the next section, following a brief survey of this need immediately below.
2. We also need to develop methods that allow designers to directly apply these theories to their designs without requiring them to have any in-depth theoretical knowledge of human cognition. Again, a brief survey of available methods provides the context for the novel method of satisfying this criterion that is the principal focus of this paper.

2.1 The need for theories of interaction

Barnard et al. [2000] stressed the importance of developing cognitive theories to explain and predict at the right level of detail for the kind of application under development. Only then will it be possible to link cognitive characteristics of the user to system features *via* so called bridging theories (see also Barnard [1991], and more recently Walenstein [2004]). These bridging theories address the behaviour (on a cognitive level) of users as they interact with systems and they could therefore be used to guide the design of interactive systems. Although we are probably some way off developing full blown bridging theories, a number of models have recently been developed that explicitly address users' interactive behaviour. These include, for example, a model for successful interaction with complex graphical user interfaces (Sutcliffe and Ryan [2000]), virtual reality (Sutcliffe and Kaur [2000]), the navigation of information spaces (Spence, [2004]), information foraging (Pirolli [2005]; Pirolli and Card [1995]), and conversation grounding (Clark [1996]).

2.2 The need for design support

The development of bridging models does not address the second problem outlined above. Indeed, the kind of design guideline arising from bridging theories is still too generic to be translated directly into concrete design decisions in the absence of in-depth knowledge about relevant cognitive theories. The application of these guidelines is therefore a potentially hazardous activity for designers who are not necessarily interested in acquiring such detailed theoretical knowledge. To overcome this problem, Sutcliffe [2000] extended the notion of *claims*. Claims were first introduced (Carroll and Rosson [1992]) to provide generic examples of design decisions appropriate to specific interaction scenarios, and to illustrate these by reference to existing designed artefacts. Sutcliffe argued that claims become vehicles for capturing theoretical knowledge used in the design and evaluation of the specific artefacts on which claims are based. By creating generic versions of these claims and artefacts and subsequently using them in the design of new artefacts this knowledge is implicitly applied. In the same *genre* and with the same aim of supporting interaction design, the concept of patterns (e.g., van Welie, [2002]) has recently emerged. Patterns attempt to encapsulate proven good design practice by providing general guidelines classified under generic headings (e.g. Interaction, Dynamic Queries) as well as outline advice and examples. However, their relation to underlying cognitive theory is rarely explicit and the advice proffered tends to be qualitative rather than quantitative.

Other mechanisms have also been proposed for translating theoretical knowledge into design (e.g., GOMS, and cognitive models such as ACT-R and ICS). However, their application remains the domain of the expert, is generally very time-consuming or addresses only basic low-level task components (Sutcliffe, [2000]). The first of these drawbacks is particularly problematic because experts in GOMS and cognitive models are usually not the people with the design ideas and experience.

3. A COGNITIVE MODEL

The rapid interpretation of images offers considerable potential for many applications (Spence [2002]) including on-line shopping (Wittenburg et al. [1998]), TV channel selection (Wittenburg et al. [2003]), the search for a photo within an album (de Bruijn et al. [2001]) and the quick assessment of the *genre* of a video (Tse et al. [1998]). However, there is little guidance available to the interaction designer who is unaware of the underlying perceptual and cognitive mechanisms involved, and how they can be exploited to advantage. In this paper we suggest that the interaction designer can substantially be informed by guidance based on Conceptual Short-Term Memory (CSTM).

3.1 Memory

Our ability to remember or identify images is astonishing. We are able, for example, to recognize a target image within a number of images presented to us in a very rapid (e.g., 10 per second) sequence (Potter and Levy [1969]; Potter [1975], [1976]; Intraub [1981]), a mode of presentation called Rapid Serial Visual Presentation (RSVP). And, 30 years ago, Standing [1973] discovered that, after being exposed to as many as 10,000 images

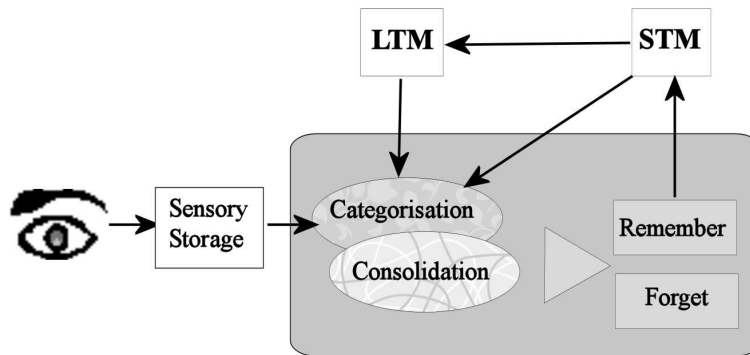


Fig. 3. A model of Conceptual Short-term Memory (CSTM) and its relation to visual perception, long term memory (LTM) and short-term or working memory (STM).

for as little as 5 seconds each, participants could later identify, with surprising accuracy, which one of a subsequently displayed pair of images was among the 10,000 already seen. Both of these examples are intimately associated with human memory, but with very different ‘lifetimes’. In the second example any model of the memory process is primarily concerned with Long Term Memory (LTM), as recognition occurred as much as 2 days after exposure. By contrast, in the first example, any cognitive model must account for phenomena occurring within about 10 to 1000 milliseconds following the onset of a visual stimulus. It is with such rapid processes that Conceptual Short-Term Memory (CSTM) is specifically concerned. Not surprisingly, these processes are often known as ‘fleeting memories’ (Coltheart [1999]). Equally unsurprising in view of the huge volume of visual and aural stimuli experienced by a conscious human being is the fact that the process of forgetting is an integral feature of CSTM.

In the remainder of this section we describe a model of CSTM and briefly identify the experimental results which support it; for a detailed exposition of matters relating to CSTM the reader is referred especially to the publications of Potter and her collaborators.

3.2 Conceptual Short-Term Memory

While conscious, we are continually in receipt of visual and aural stimuli (Potter [1993], Potter [1999]). We look at the table, the waiter across the restaurant, the lighting fixture and the carpet, and hear a variety of sounds, perhaps including conversation and the noise of passing traffic, all in quite rapid succession. Some images and sounds are remembered; most are forgotten.

Those images and sounds are first held in sensory storage in which a literal copy of the stimulus is held very briefly (Neisser [1967]). Because the lifetime of this storage is of the order of a few tens of milliseconds, any subsequent processing of a stimulus must occur with little delay.

Almost immediately following the arrival of a visual stimulus in sensory storage, the stimulus is *categorized* (Figure 3) by reference to previously stored knowledge held mainly in the user’s Long Term Memory (LTM).² Thus, the stimulus may be categorized as a cat in front of a house, or as coloured wiggly lines, or a briefcase. The categorisation is the result of a lifetime of exposure to visual stimuli³, and is little influenced by any task

² The process we have described may be equally applicable to sounds. However, in the remainder of this paper we confine our attention to visual stimuli.

³ In this section we use the term ‘visual stimulus’ quite generally. It can mean an image as in a painting by Turner, an icon such as that of IBM or NASA, the layout of a letterhead, a line drawing or even the layout of some text as in a motorway sign.

being executed. The categorization is carried out *effortlessly* and *unconsciously*, and is thought to occur in about 100 milliseconds. Supportive experimental evidence for this and other aspects of CSTM is discussed below.

Following the categorization of a stimulus its meaning/relevance is interpreted, again unconsciously.⁴ This process, also referred to as *structuring* or *consolidation*, is essentially one in which the relevance of a stimulus, either alone or together with other images, is established within the context of a task of which a user may or may not be aware. The processes of categorization and consolidation are incorporated in the simple model of Conceptual Short-term Memory (CSTM) shown in Figure 3, a model we shall refine and quantify.

The process of consolidation may take different forms and may occur with reference to information held in either Long-term Memory (LTM) or Short-term Memory (STM). Where a user is searching, for example, for a familiar photographic image within a collection of such images, the ‘relevance’ will be the identification of a match between a target image held in memory and a currently viewed image. If the target image is one seen some time (e.g., several hours or days) before the search, then that information usually resides within LTM. If, on the other hand, a user is searching for an image viewed just before the search (within seconds or minutes), or one of which they have a clear image in their head, then the relevant information will reside in STM. A positive outcome of the structuring process results in the meaning/relevance being consolidated in STM. In the event that no relevance is established the information will be lost from current awareness (forgetting).

Potter [1993] points out that “this whole cycle – identification of stimuli, memory recruitment, structuring, consolidation and the forgetting of non-structured material – may occur in less than one second when viewing a pictured scene or reading a sentence”. In case the word *memory* in CSTM suggests something static it is useful to point out that

⁴ We are not claiming that the interpretation of an image can start only after it has been fully identified in the categorisation process. Categorisation and interpretation may well occur in parallel to some extent. In some cases this could lead to the selection of an image as the target image before it has been fully identified on the basis of certain salient visual features that set it apart from the non-target items in the sequence (de Bruijn & Spence, 2000).

Potter identified an essential feature of CSTM with the remark that "... CSTM is part of a dynamic, structure building process, rather than a passive store".⁵

3.3 Experimental Evidence

Features of the CSTM model can be refined and quantified by reference to seven experiments:

1. A first experiment, illustrated in Figure 4a, is of the form "tell me if *this image* occurs in the rapid (e.g., 10 per second) sequence of N images I'm about to show you" (Potter and Levy [1969]; Potter [1976]; Intraub [1981]). In this experiment, detection is about 80% to 90% successful depending upon the precise duration of each image presentation. Here, following the categorization of each image, the interpretation is simple: no identification of structure is involved, merely the matching of a perceived image to a memory of the target image.
2. In a second experiment, illustrated in Figure 4b, the target image was specified again in advance of the presentation, but this time by giving only a description of its content (e.g., Intraub, 1981). Two levels of description were used; its name (e.g., "a dog") or its category ("an animal"). In this experiment a much longer presentation duration, i.e., 258 ms, was needed in order to reach approximately 90% detection accuracy. Again, the interpretation is simple: the process of matching a perceived image to a description of the target image held in memory can be disrupted by the following image for much longer than in the case of the first experiment.
3. In a third experiment (Tse et al. [1998]), the images in a rapidly presented sequence were key frames taken from a videotaped play (Figure 4c). It was found that even at a presentation rate of about 10 frames per second, a participant could understand and remember the gist of that play in view of the structure inherent in the video and reflected in the sequence of key frames.
4. In a fourth experiment (Figure 4d) a participant is shown a rapid sequence (e.g., 10 per second) of *unrelated* images, and then asked afterwards if they recognize, as having been within the sequence, images from a set containing both old and new images. In this case it was found that 'recognition memory' was good only when the

⁵ CSTM appears to operate as a buffer, linking together the quite separate entities of sensory memory, STM and LTM. To build the complete picture, we would need to include theories of sensory memory, LTM and STM. This, however, is outside the scope of this paper.

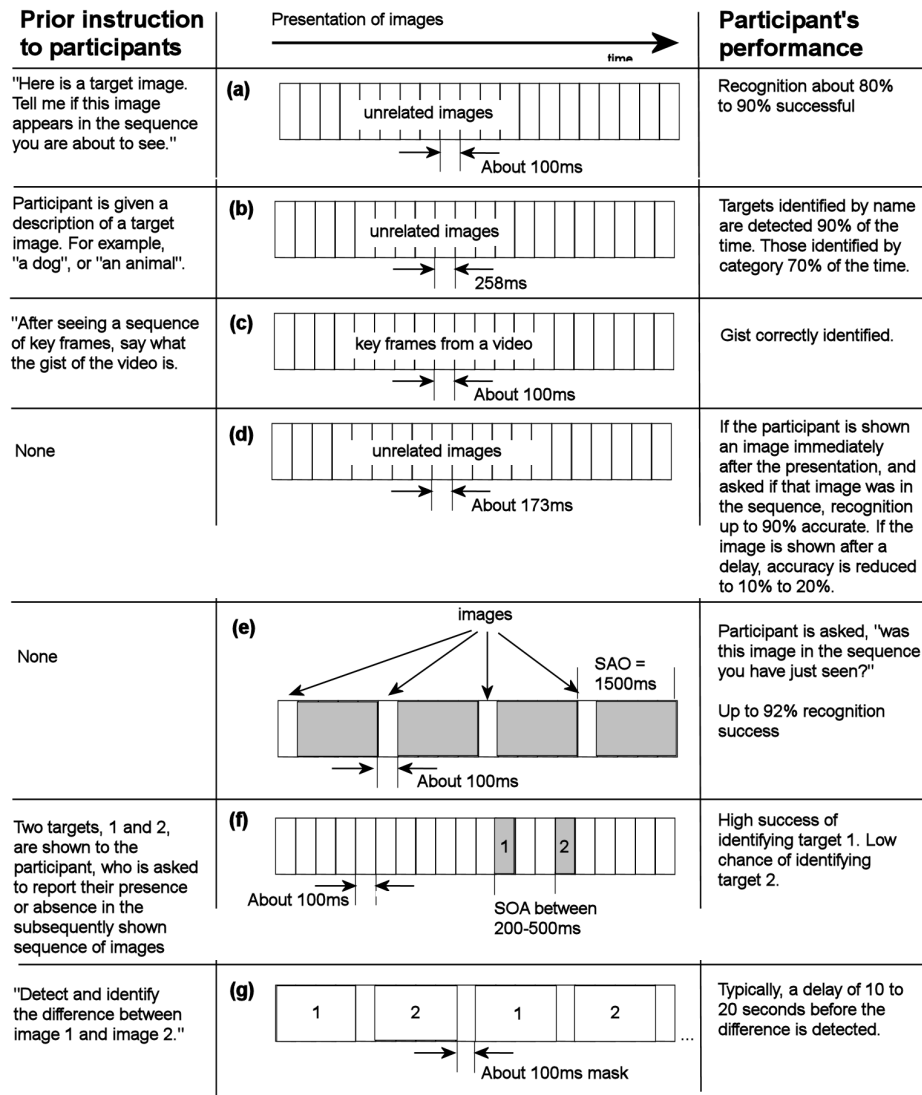


Fig. 4. A schematic representation of the experiments illustrating the concept of Conceptual Short-Tem Memory (SOA = Stimulus Onset Asynchrony).

decision was made immediately after the presentation of the sequence (Potter et al. [2002]). Any delay in making the decision led to a rapid decline in performance, with success rates of only around 10% to 20% (Potter and Levy [1969]; Potter [1976]; Potter et al. [2002]) after a delay as short as five seconds.

5. A fifth experiment (Figure 4e) demonstrates that the poor ability to recognize an image in the experiment of Figure 4d is *not* a result of its short presentation time. In this experiment the brief (e.g. 100 ms) presentation of each image was followed by a meaningless 'visual mask', with the result that images were easy to remember,

suggesting that whereas 100 ms would be sufficient for the user to comprehend an image and identify it as the image being sought, as in the experiment of Figure 4a, it would be insufficient for its consolidation into Short-Term Memory (Intraub [1980], [1984], [1999]; Potter [1976]).

6. Other experiments are relevant to briefly glimpsed images, and two are of particular interest. One such experiment, in which the process of consolidation can have interesting side-effects of which an interaction designer should be aware, led to the term Attentional Blink. For example, if a participant is presented with two target stimuli (Figure 4f), the presence of those stimuli in an RSVP presentation is likely to be reported correctly except when the second target appears within 200 to 500 milliseconds of the first target (Raymond et al. [1992]; Coltheart [1999]). The other experiment of interest (Figure 4g) demonstrates the phenomenon of Change Blindness (see, Nowell et al. [2001] for examples of how this may affect the effectiveness of visual interfaces). Two images that differ are shown in alternating sequence, separated by a gap of about 100ms. It is often the case that a participant in such an experiment will take as long as 20 seconds to detect the difference which, when known, is strikingly obvious.

The experimental results illustrated in Figure 4 allow some degree of quantitative characterization to be added to the CSTM model, particularly with regard to the arrival times before which an image can interfere with the processing of a previously perceived

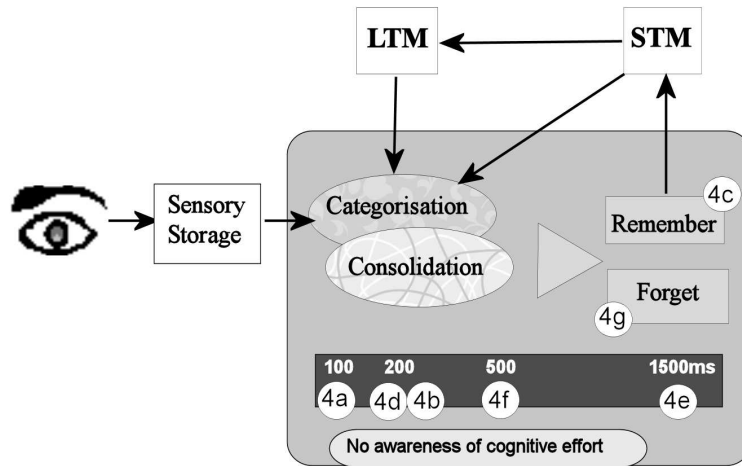


Fig. 5. A model of Conceptual Short-Term Memory (CSTM). Times in the dark gray box are the earliest arrival times of a following image to avoid interference with categorization and consolidation. White circles identify experiments leading to various features of the model.

image. The enhanced model is shown in Figure 5, where times in the scale are the earliest arrival times of following images to avoid interference with categorization and consolidation. Blue circles with an inset identify the relevant experiments of Figure 4 leading to various features of the model. The refined model also indicates that a user is not consciously aware of any cognitive effort during the CSTM process⁶, as well as the fact that categorization and consolidation occur in parallel to some extent.

Its success in explaining, with reasonable accuracy, both quantitative and qualitative aspects of human visual performance in the context of briefly perceived images is more than sufficient to establish the value of the CSTM model and related experimental evidence to the interaction designer. In section 5, we examine the process by which such a cognitive model can provide valuable guidance to the interaction designer. First though, in section 4, we examine a behavioural phenomenon to act as a bridge in the derivation of such guidance.

4. BEHAVIOURAL PHENOMENON

We now describe the way in which a behavioural model provides the context (Figure 2) in which the consequences of CSTM for interaction design can be established. For this purpose, we chose the phenomena of Opportunistic and Involuntary Browsing described in the introduction.

4.1 The behavioural model

Perhaps the first bridging model that was applied to interaction design is Norman's Action Cycle (Norman [1986]; Figure 6a). In this model, human interaction with computers is described as a cycle in which a user passes, in sequence, through a number of activities aimed at achieving a particular goal. The formation of an intention to act, the preparation of an action plan, and the execution of that plan span the so-called Gulf of Execution. The resulting change in the world (i.e., system response) is then perceived, interpreted and finally evaluated in relation to the goal, thereby spanning the Gulf of Evaluation.

As discussed in the Introduction, the behaviour associated with OB and IB is, like a lot of human activity, not consciously goal driven. A user merely reacts to events by

⁶ In a discussion headed "Is CSTM Conscious?" Potter (1999, page 41) reviews some relevant experiences but concludes by saying that " . . . such [conscious] experiences seem to be the exception rather than the rule. Thus, I adopt the working hypothesis that much of CSTM activation and selection and structuring happens before one becomes aware: it is the structured result, typically, of which one is aware, which is why perception and cognition seem so effortless and accurate."

unconsciously evaluating them in terms of the many goals, often ill-defined, that they may be pursuing on a daily basis. Therefore, to emphasize the nature of OB and IB, we show by the highlighting in Figure 6a how OB and IB correspond with the Gulf of Evaluation in Norman's Action Cycle and, most importantly, that *they do not invoke the cognitively demanding processes involved in traversing the Gulf of Execution*. When, as in OB or IB, Norman's World Change is the change in a visual display, that change is perceived when the display is fixated, the content of the display is then interpreted and its relevance is evaluated. Indeed, Figure 6b positions OB and IB in the context of the CSTM model in the sense that perception, involving fixation on a display, corresponds to the arrival of a visual stimulus, interpretation is equivalent to the categorization of its content, and evaluation becomes the process of structuring to assess the relevance of that content. In case the display is not already fixated at the moment of visual change, fixation on the display will either happen during the course of normal scanning behaviour (either planned or opportunistic), or can be achieved by a deliberate orienting response towards the change.

4.2 Essential features of the visual change

For the interaction designer whose task it is to *design* the visual change, the description of Opportunistic and Involuntary Browsing in Figure 6b may be helpful but is far from sufficient. More detail is required: we must establish the necessary or desirable requirements for the visual change designed by the interaction designer to support the various stages of OB or IB identified in Figure 6b. These requirements are discussed in outline below.

4.2.1. Perception. Perception requires that any visual change be registered without undue effort. If the user's attention is focused on the change, perception can occur in under 30ms. If, on the other hand, the user's eye-gaze trajectory is quite random (i.e., involuntary), then their eye-gaze will only fixate on a visual change by chance: it may therefore be some time before a change is perceived. In order to increase the chance that a visual change will be fixated, it must be sufficiently clear and abrupt so as to draw attention automatically and invoke an orienting response (Yantis and Jonides [1990]). The latter course must be used only in situations in which it is absolutely essential that a change be perceived, as it inevitably causes major disruption to ongoing activities and will lead to frustration and irritation if the change is deemed not important enough.

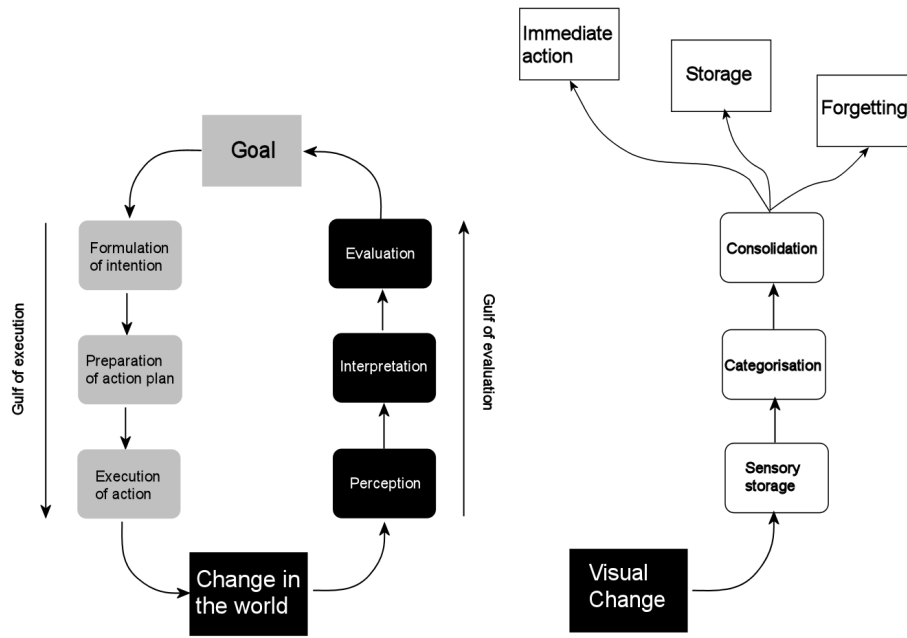


Fig. 6. (a) Norman's Action Cycle. The processes involved in OB or IB are highlighted, and (b) The processes within CSTM associated with OB and IB and their possible outcomes.

4.2.2. Categorisation. Categorisation of a visual change requires that the change be meaningful, and will be facilitated by the provision of appropriate visual cues, by good visual design and by adequate presentation in terms of size, contrast, orientation and duration.

4.2.3. Consolidation. Successful consolidation is the subject of many design issues. For example, the use of well-designed visual images can usually ensure that a visual change is sufficiently rich for it to be suggestive of what the change represents. This is a necessary requirement for the correct evaluation of an item's importance. Incorrect evaluation of an item's importance may lead, on the one hand, to relevant items being ignored and forgotten and, on the other, to irrelevant items being consciously considered for action. Both situations are clearly undesirable.

As part of the consolidation process, the content can either be selected for immediate action, stored temporarily for later selection or forgetting, or integrated with earlier changes of the display. The exact outcome of the consolidation process depends both on the decision strategy employed by the user and the characteristics of the presentation. For example, the user may employ a threshold matching (satisficing) strategy, by acting immediately upon perceiving an information item of interest. Alternatively, the user may

decide to postpone any action until a sufficient number of items have been assessed, whereupon the most relevant item is selected (optimisation). However, whether the latter strategy will work depends to a large extent upon the ability of the user to store all the items until the moment when a selection is made. In the next section, we provide *detailed* (i.e., quantitative) design advice for accommodating each of the possible outcomes of the consolidation process.

5. DESIGN ACTIONS

Notwithstanding any value to be gained from the general design guidelines discussed above, the support they provide to an interaction designer is still very limited, partly because of their qualitative nature and partly because the guidelines contain insufficient detail. Design trade-offs, for example, are not identified. What is needed is some way of describing design decisions that identifies not only their consequences but also establishes their advantages and disadvantages as well as the context in which they are relevant (cf. Sutcliffe [2000]).

For many of the requirements arising from the specification of an interface there may be no sufficiently well established cognitive theories. We cannot, for example, precisely quantify the requirement that “the visual change be sufficiently clear and abrupt to draw attention”. For other requirements, however, the experimental results that gave rise to the model of CSTM and shown in Figure 4 come into their own, because some degree of quantitative precision can be offered to the interaction designer. It is important to note however that it would be impossible to apply these results correctly without the context provided by the OB and IB model. For example, although we have established that a presentation duration of 100ms should be enough under certain circumstances, we also need to know what the user’s intentions are likely to be in order to decide that presenting an image for this length of time is indeed sufficient. In the next section we illustrate the approach to the development of a *Design Action* by addressing one requirement for OB and IB; that of exposing users to as many information items as possible without requiring the user to instigate changes in the world themselves. For that requirement we derive a *design action* by reference to the experimental results that support the CSTM model, and the model of OB and IB.

5.1. An example: Maximisation of information exposure

We shall explain the characteristics and derivation of a design action by presenting one in Table I and explaining it in detail.

Table I. An example of a Design Action.

ID	DA1
Title	Maximisation of information exposure
Description	Expose users to as many representations of information items as possible commensurate with maximising the likelihood of those representations being correctly interpreted in terms of the information being represented.
Effect	The more of these items being presented to the user the greater the likelihood will be that one is perceived that is a relevant to a user's interests.
Upside	(1) Queries need not be explicitly formulated, as the relevant items will be recognized as such when encountered. (2) Irrelevant items will be forgotten with no cost in cognitive effort
Downside	Any display area is finite, giving rise to a trade-off between the number and size of items being presented simultaneously, and their presentation duration.
Issues	Comment is often made concerning the effects of information overload. This is not so relevant here, because no additional cognitive effort is involved in perceiving, interpreting and then forgetting information that is irrelevant. However, any action triggered does require focused attention. It is therefore important that this DA be applied to situations in which the likelihood of encountering a relevant item is relatively small. If multiple items of interest are likely to be encountered, then it is important that either a) the items can be easily prioritised, or b) the items can be dealt with sequentially.
Theory	Current theory suggests that irrelevant information is rapidly forgotten at no additional cost.

The design action is headed by an **identifier** (here, DA1), and a brief **title** (here, “Maximisation of information exposure”) indicating as clearly as possible the expected result of applying the design action. In this instance the designer wishes to expose a user to as many relevant information items as possible.

Next, a **Description** clarifies the necessarily brief title.

Under **Effect**, the apparently obvious point is made that the more items that are presented to the user, the greater is the likelihood that one will be perceived that is relevant to the user's interests.

In the **Upside** section the advantages accruing from the design action are delineated. Here there are two. One - and the most important one - has already been emphasised in this paper, namely that there is no need for the user to explicitly formulate queries. The second is that there is no conscious cognitive effort associated with items that are forgotten.

Trade-offs are identified in the **Downside** section. Here the trade-off is imposed by the combined finite resources of display area and presentation duration.

Issues that are neither Upsides nor Downsides will often arise in the application of a design action. Here, one concerns the notion of information overload.

Intentionally, design actions are described in such a way that the underlying cognitive **Theory** need not be understood by the interaction designer. Nevertheless, to provide an opportunity for that designer to ‘dig deeper’, a reference to the relevant cognitive theory is provided. Here, the forgetting process of CSTM is of particular importance, because it suggests that most items will be forgotten without the user becoming aware of them.

In this section we have illustrated the derivation of a design action from cognitive theory (here, CSTM) via the context of a bridging model of human behaviour (here, OB and IB). Other design actions whose derivation is outlined in Figure 7 are collected in the Appendix for the same cognitive theory and behavioural model. These DAs address a variety of design issues. DA1, as discussed, suggests how a user’s exposure to information can be maximized. DA2 is concerned with presentation to support immediate responsive action within a ‘satisficing’ strategy. If an exhaustive examination of information items, followed by immediate appropriate selection (an ‘optimising’ strategy) is the aim, then DA3 is relevant. If, however, a user is expected to assess each information item in the context of the others and then make a considered decision (a ‘preview-consider-select’ strategy), then DA4 offers appropriate guidance. DA5 addresses a different type of user task, one in which the gist of a narrative must be deduced from a rapidly presented sequence of key frames. Finally, the need for a user to be aware of visual change is addressed by DA6. In the following section, we illustrate the application of these design actions to the design of two novel interfaces.

6. USING THE DESIGN ACTIONS

We illustrate the value of *design actions* by their application to the design rationale associated with two artefacts. These artefacts are very different: one involves focused Web browsing by a single user on a small display that is typically part of a pocket

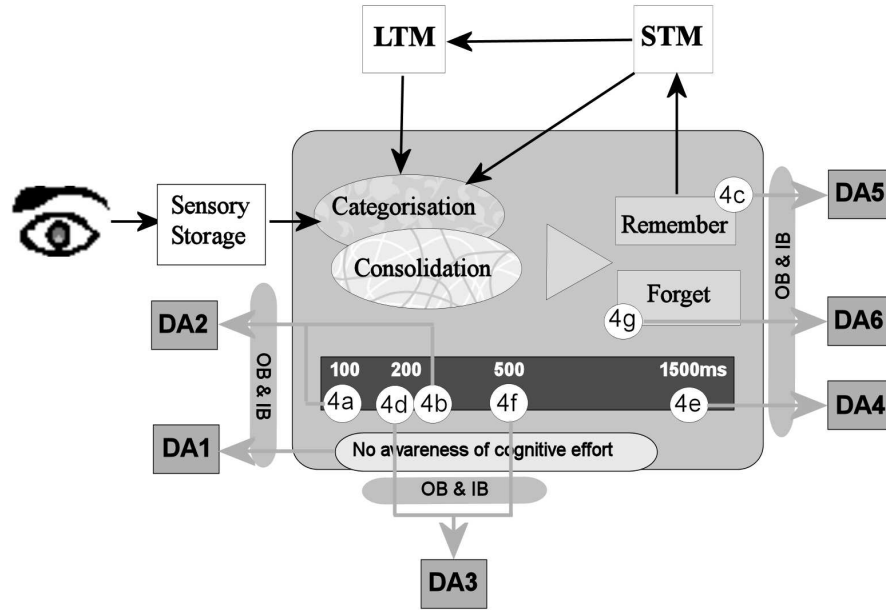


Fig. 7. This schematic representation shows how the Design Actions are derived from the model of CSTM through the behavioural model of OB and IB. The white circles represent the various experiments on which features of the CSTM model, such as estimates of the minimal delay between onset of consecutive stimuli to prevent interference, are based.

computer or mobile phone, while the other involves unfocussed multi-user interaction with a large display embedded in the surface of a coffee table.

6.1 The mobile Web browser

The DAs described in the Appendix were used to support the design of a Web browser for small screen devices such as mobile phones and personal digital assistants (PDA) (de Bruijn & Chong, [2002]). The browser was intended to support intentional browsing activities aimed at retrieving information from the Web, activities similar to those supported by conventional browsers running on desktop computers such as MS Internet Explorer and Mozilla Firefox. As we shall show, some of the Design Actions outlined in the appendix are specifically relevant to the design of this browser. A prototype of the mobile browser was implemented allowing users to browse daily news stories (ITV News; Figure 8a). In this prototype two levels of RSVP were employed. At the first level, RSVP was used to present links to five news categories (e.g., “Britain”, “World”) in the form of a one-word label and an image for each category (see Figure 8b). One deemed by the user to be appropriate could be selected after which a further description (Figure 8c) of the news category became available for viewing. A second level of RSVP presented a sequence of previews of news stories within the selected

category (Figure 8d presents one example of these). Each preview contained a readily identifiable image that was representative of the content of the story and a very brief description in the form of a headline. When a story preview at this level was selected, the first page of the news item was presented for view (Figure 8e). The news item could then be read at leisure, using a paging action if the extent of the content so required (cf. Öquist et al., [2004]). All content was taken from the ITV News Web site (www.itv.com/news).

6.1.1 Application of Design Actions

Five of the Design Actions (DAs 1, 2, 3, 4 and 6) presented in the Appendix are relevant to the design of the browser. **DA1**, for example, is relevant since the user will often wish to assess as many information items as possible without wasting time and effort. Such assessment of items can be beneficial in both search browsing and opportunistic browsing situations. In search browsing, the adoption of this DA allows the user to assess all possible actions and determine the best possible one in the acquisition of a target item. In opportunistic browsing, this DA allows quick and effortless assessment of options for their relevance to perhaps many unspecified goals.

Subject to the demands of other DAs, as discussed below, the need to present as many

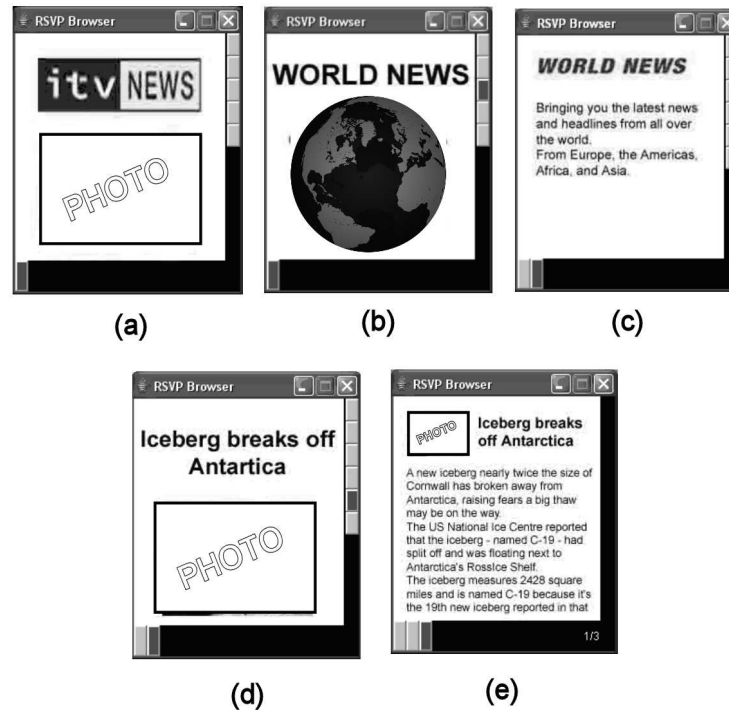


Fig. 8. The RSVP Browser presented information items on five levels of a content hierarchy; (a) the home page, (b) a news category option, (c) a news category, (d) a news story option and (e) a news story.

items within as short a time as possible can be satisfied in many ways. Given a large display, for example, information items can be presented concurrently, and rapidly assessed in the course of many sequential and rapid eye-gaze movements. But such an approach is not possible within the constraints of the limited display area associated with a mobile phone or PDA. However, one method of presentation that is ideally suited to a limited display area is that known as Rapid Serial Visual Presentation (RSVP; Spence et al., [2004]), in which each information item is displayed in full for, say, 100ms, and then replaced by the next item for a similar period of time. Thus, with RSVP, each item can be represented by both text and an image occupying almost the same display area as within a standard Web page. Indeed, as far as the user's visual processing system is concerned, there is much in common between an RSVP presentation and the succession of images presented to the foveal region of the eye as the user scans a much larger Web page. Alternative approaches to the presentation of information on small displays have been suggested, but suffer from significant disadvantages. For example, scrolling both horizontally and vertically to assess large amounts of information on a small screen puts unacceptable demands on the user (e.g., Jones et al. [1999]), while reducing the amount of information displayed to a list of brief textual items significantly increases the risk of misinterpretation of their relevance (e.g., Kaasten, [2002]).

Once we decided to use RSVP of information items, the next issues to be resolved were the appropriate presentation rate for each of the previews, and the appropriate actions for selecting items. A number of DAs were pertinent to these issues.

DA2 suggested that a fixation of 100ms on an information item may be enough to ensure that its relevance can be established with reasonable certainty. Thus, if immediate selection of a 'satisficing' item – possibly in the absence of further search – were intended, then a presentation rate of 10 per second would have been adequate. As soon as a relevant item were encountered, the user should then be able to stop the presentation and select the item. Alternatively, the user might adopt an 'optimising' strategy and wish to see all available items before selecting the 'best' according to some criterion. In this case, **DA3** pointed to a presentation rate of 2 per second. Like DA2, this DA also required that the interaction design be such that the user would have easily been able to select an item after the different options were assessed (optimising). For yet another strategy **DA4** was relevant. This design action was relevant to a strategy in which the user wanted to preview all the options and only think of what the best option was after all of them had been considered (preview-consider-select). In this case it would have been

important that users were able to remember all of the options for some time after presentation during which a decision about the preferred option could have been made. DA4 suggested that in such a situation a presentation rate of less than 1 item per second would have been appropriate, with a maximum limit of 16 on the total number of items presented before the user is given time to consider the alternatives and make a decision.

To summarise, RSVP rates of approximately 10, 2 and 1 item per second have been associated with different strategies that users might adopt. To cater for the fact that a user might well adopt *any or all* of these strategies, the slowest rate at first seemed appropriate for design purposes. However, the choice of presentation rate should be guided by the strategy that the users are likely to adopt. In the case of the RSVP Browser, and for the purpose of the planned evaluation study in particular (see section 6.2), it appeared unlikely that users would adopt a preview-consider-select strategy and we therefore settled on a presentation rate of two per second.

Finally, the detailed design of information items was also informed by the design actions. For example, DA1 required that both recognition and interpretation of information items presented using RSVP would have to be rapid and effortless. This suggested that each information item contain a clear image and no more than about five words of text, since it is the gist of each item that has to be interpreted rapidly (longer texts would take too much time to read). **DA6** suggested that any change in the display should involve the presentation of a new gist.

6.1.2 Operation of the RSVP Browser

The items available at both levels (news categories and stories) were permanently represented by buttons presented in a column along the right side of the screen (see Figure 8). These buttons allowed users to select items directly in case they were looking for a particular story and knew where to find it (search browsing), or, indeed, to perform ‘manual RSVP’, which may be preferred to support a satisficing strategy in opportunistic browsing.⁷ The RSVP of a sequence of information items was started by pressing the centre of the PDA’s directional pad. After a short delay the first information item was presented filling the majority of the screen. Then a new item appeared every 500 ms until the RSVP was stopped either by pressing the centre of the pad again, by tapping the screen in the area where the item was presented, or by tapping one of the buttons. Selection of an item was achieved by double tapping the main display area on the screen.

If the RSVP overshoot the target or if it was stopped prematurely, repeatedly pressing the top or bottom of the pad allowed the user to step back or forward, respectively, through the information items. Alternatively, the appropriate item could be identified during the first presentation and selected during a second presentation. For this purpose, the RSVP would loop to the first information item after the last one was presented for 500ms and the whole sequence would be presented again.

6.2 Evaluation of the RSVP Browser.

An evaluation of the RSVP Browser compared it with a Pocket version of Microsoft Internet Explorer ('Pocket IE'), both implemented on a Compaq IPAQ PDA (Figures 8 and 9). Thirty participants were each asked to retrieve information from a set of 32 news items divided into five categories. Half the participants were randomly assigned to use the RSVP Browser and the other half used the Pocket IE. Each participant answered a series of eight questions, of which a typical example is:

"The cross-border train link between Belfast and Dublin has been closed after several explosions were heard near the line. Has the cause of these explosions been identified?"

Participants using the RSVP Browser were asked the same questions as those using IE Explorer. Three sets of questions were used such that the answers to each of the questions in each set could be found in the same news category, but in different stories, as in the

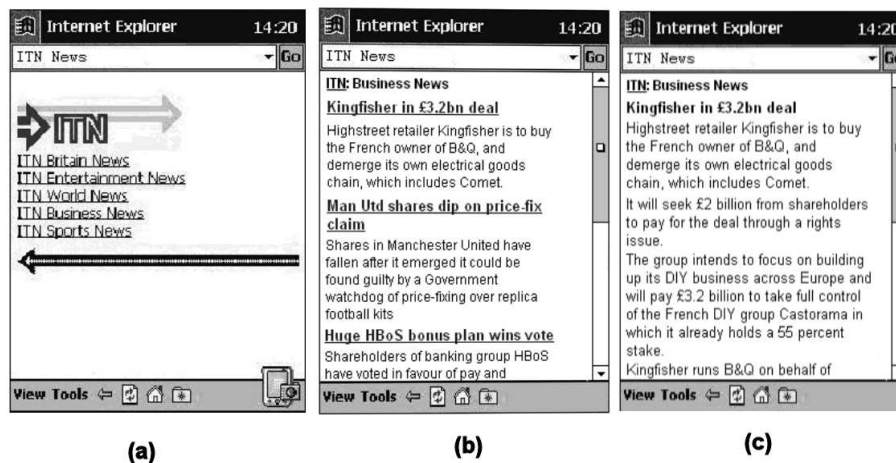


Fig. 9. The IE browser presented three levels of a content hierarchy; (a) the home page, (b) news headlines and extracts, and (c) news stories.

⁷ Automatic RSVP is likely to overshoot a relevant item before it can be stopped. If the number of choices is limited, the user may judge that the extra burden of manual RSVP is sufficiently compensated by not having to back track before the item can be selected.

corresponding questions of the other sets (making a total of 24 questions). This was done to minimise the chance that some questions would favour one type of browser over the other.

Participants were given no instructions about the functioning of the two browsers either prior to or during attempts to find the answer to the first question. Then, while looking for the answers to questions two to six, participants were given instruction on the use of the browsers and any questions they had were answered. When looking for the answers to questions seven and eight, participants were again not given any assistance. This method allowed us to assess both browsers' guessability during question one and their ability to support skilled user performance during the last two questions (Moyes and Jordan, 1993).

6.2.1 Performance

We scored the performance on each question by measuring the time to find the correct answer (Time), and the number of unnecessary steps taken to find the answer normalised with respect to the minimum number of steps required (Steps). Although participants using the RSVP Browser needed, on average, more than four times as long to answer question 1 (100 seconds) compared to those using IE (22 seconds), this difference was no longer apparent when answering questions seven and eight (24 and 25 seconds respectively). A similar comparative performance was obtained for the number of unnecessary steps (see de Bruijn and Tong, 2003, for detailed statistical analyses). These results show that for users who had very little practice, the RSVP Browser was as effective as the *de facto* standard for mobile Web browsing.

6.2.2 Users' preferences

Particularly relevant to design actions DA2, DA3 and DA4 was the result of eliciting users' preferences concerning the speed of RSVP. Of the 15 participants who used the RSVP Browser, only two indicated that they experienced the RSVP preview rate as being '*too fast*' on a five-point scale ranging from '*too slow*' to '*too fast*'. Furthermore, only four participants indicated, on a five-point scale ranging from '*very easy*' to '*very difficult*', that they found it either difficult (3) or very difficult (1) to select their desired choice using the RSVP function, whereas nine found this easy to achieve. These results suggest that the choice of 500ms display per information item was well advised by the DAs, but that it may represent a minimum for the RSVP Browser.

When asked whether they found it easy or difficult to go to a particular news story most IE users found it easy (13) whereas the RSVP Browser users' opinions were more

divided with some finding it very easy (5) or easy (6) while others were indifferent (4). However, when asked whether they found it easy or difficult to assess the available options for their relevance to the questions, slightly more RSVP Browser users found it either very easy (1) or easy (11) than IE users (2 and 8 respectively).

6.2.3 Support for browsing

The RSVP Browser supported both automatic and manual presentation of information items, and most participants who used this browser opted for the latter. They did this by clicking, using the device's stylus, the buttons displayed along the right side of the screen (usually) in a top-to-bottom rapid sequence. After approximately question 3, the automatic method was only rarely used in the selection of news categories, which were mostly selected by a directly clicking the one button representing the desired category for each question.⁸ Similarly, most news stories were assessed through manual RSVP. Only one third of the participants chose to use automatic RSVP to find the answers to questions seven and eight. The other participants used the preview buttons to assess news stories, and this was almost universally done by using the stylus to click the buttons in a top-to-bottom order until the relevant preview was found and selected. On the one hand, participants who manually explored news stories assessed on average only 0.15 previews more than a satisficing strategy would predict when answering questions seven and eight. Participants who used the automatic presentation of news stories, on the other hand, assessed on average 3.14 items more, but they did not consistently assess all stories as would be expected if they adopted an optimising strategy. Therefore, it appears that both methods were mostly used as part of a satisficing strategy, which may reflect an emphasis on search browsing induced by the nature of the task participants were given to perform.

6.2.4 Conclusions

From the evidence discussed it could be argued that there is, at least, a *prima facie* case for considering the potential of RSVP to enable web content to be assessed via a mobile display. Although the evaluation probably assessed the browser's support of search browsing rather than opportunistic browsing, insights into its support of OB can be gained. It is apparent, for example, that the Browser supported both the immediate selection of items of interest as soon as they were encountered (satisficing) as well as, on

⁸ This method was potentially error prone in the absence of textual labels on the buttons, and five errors were made in total while answering questions seven and eight.

some occasions, the exhaustive exploration of all alternatives and the selection of one of interest (optimising).

Support for the relevance of Design Actions is implicit rather than explicit, and is especially difficult to elicit when different types of browsing may be present: the existence of three user strategies that may be adopted further complicates the picture. Nevertheless, with regard to the choice of RSVP rate for which DA2, DA3 and DA4 offered differing guidance, the fact that only 2 out of fifteen subjects regarded the rate as too fast is encouraging, as is the absence of any negative indications regarding users' behaviour. Nevertheless we do not underestimate the challenge posed by the eventual need to disambiguate between the three types of browsing as well as the strategies adopted by users.

6.3 The Coffee Table

As stated in the introduction, people have a constant need for information, a need that is conventionally serviced by an explicit statement of interest according to some conventional scheme (e.g., Google). Nevertheless, though people have many interests, they cannot continuously, simultaneously and consciously articulate them. An alternative and very attractive way of servicing their information needs, in that it requires no articulation or conscious effort, is to employ opportunistic or involuntary browsing. In order to allow as much involuntary browsing as possible to take place we need to expose users to a large number of information items in order to significantly increase the ability of people to access information serendipitously.

In order to expose people to as many items of information as possible whilst they go about their daily business or pleasure, we designed an embedded interactive information display in the form of a coffee table. The Table has a dual role. It supports social intercourse by allowing people to sit around the Table - perhaps drinking coffee, eating donuts and engaging in conversation - while exposing those people to a range of information items. A video projector above the Table (Figure 10a) projects, on to the Table's surface, information items in the format of rectangular frames (Figure 10b) typically comprising a picture and about two or three words (Stathis et al., 2002).

Typically, these information items will be drawn from a large collection relevant to the locality in which the Table is established (e.g., local news about a village, with the table installed in the village's pub). Because the information is not selected with the interests of a particular user in mind, but instead for a community of which the users are



Fig. 10. (a) The Coffee Table, showing the video projector responsible for projecting information items onto the table surface. (b) Information items projected onto the surface of the Table move slowly around the perimeter. The interactive DiamondTouch surface (Ryall et al. [2004]) senses touch on a particular item which the coffee table software translates into manipulations of the information items.

likely to be members, the collection will almost certainly contain a mix of items that are relevant and irrelevant to any one person.

If a person sitting at the Table notices an item of potential interest a decision between three simple actions will have to be made based on the disruption to social intercourse that is deemed to be acceptable. In the order of low to high disruption these actions include (1) the item is ignored (no action is executed); (2) the item is moved by finger drag to the edge of the Table where that person is sitting; or (3) the item is moved to the centre of the Table by finger drag, whereupon it expands to be more easily studied. If action (2) is chosen the user can later copy such saved items onto some storage device (each touch is identified by seat identification). If action (3) is chosen, two options are now possible: (a) the magnified item can be ignored (it will fade away after about five seconds), or (b) the magnified item can be consolidated in the centre of the Table – perhaps for all to discuss – by tapping a finger on the surface where it appears. All the actions described are supported by a Table surface (Ryall et al., 2004) which is sensitive to the touch of a finger, and which can identify its owner.

Not surprisingly, there are aspects of browsing behaviour for which the Design Actions do not offer guidance. One such is the need to avoid disruption of people's natural eye gaze behaviour if it is the intention of the designer to support *involuntary* browsing rather than opportunistic browsing. For this reason the speed of image movement was chosen to be one revolution in about 2 minutes as this ensured that each item was presented for long enough to allow its assessment through intermittent glances

at the table. Some of the many design decisions can, however, be directly related to design actions, as we now discuss.

DA1 indicates how it is possible to expose users to many information items without imposing unacceptable demands on the user's attention. This DA provided the rationale for the interactive coffee table concept. The coffee table is just one example of a class of interactive artefacts intended to push information to potential users. The electronic coffee table exposes the user to much more information compared to the usual stack of magazines and newspapers present on ordinary coffee tables in this environment. Each of the information items is clearly visible without any overlap (unlike the way magazines are often stacked up). In order to optimize the available screen real estate, items that have been presented for some time are continuously, but slowly, replaced. **DA1** also indicates that there should be a suitable mix of items likely to be relevant to a particular user; this requirement is satisfied by the location of the Table in a community setting (e.g., a pub) and by the choice of content relevant to this setting rather than any particular individual within it (see Stathis et al., 2002).

Another Design Action that is pertinent to the design of artefacts intended to support involuntary browsing in particular is **DA2**, because this form of browsing is by definition a threshold matching activity. There are three consequences arising from DA2.

Firstly, DA2 requires that the amount of information in each rectangle be kept to a minimum; this requirement was satisfied by the nature of each item, involving a simple image and only two or three words. Either the image or the words can be recognized in a single glance. Secondly, DA2 requires that it be possible for an item to be recognized within the short time a user's gaze rests upon it, which means that a rectangle should not move significantly in about 100ms. For this reason it was important to create, as much as possible, the impression of smooth movement of the items presented on the table so that items could be easily tracked during fixation. Finally, DA2 requires that an information item be rapidly interpreted. Since an upside-down item may not easily be interpreted unconsciously, a circular trajectory was chosen for the items in order to give equal opportunity to every person sitting around the Table to view each item in an upright position.

DA6 requires that each information item should be as distinctive as possible when compared with other items in terms of both its visual appearance and (more importantly) its gist. The graphic design of the items attempted to satisfy the first requirement, while

the selection of items from a wide range of different topics ensures that the latter requirement is also satisfied

To test the result of applying DAs in this way, as well as to evaluate the performance of the Table itself, a study was undertaken to test some of the assumptions underlying the design of the table.

6.4 Evaluation of the coffee table.

The effectiveness of the design actions in the design of the Table was evaluated in a laboratory situation. Consequently, a major challenge was to create, in this artificial environment, the natural context for which the Table was designed - the support of information retrieval by involuntary browsing simultaneously with the activity of conversation. For this reason we organised our participants into pairs and instructed each pair to have a conversation while seated at the Table. However, in order to create specific unresolved queries in their memories we gave each participant, *prior to the conversation task*, some general knowledge questions. The reason for these questions was not revealed to the participants until after the conversation task.

6.4.1 Method

Eight groups of two participants (dyads) took part in this experiment. They were given a number of tasks to perform during one session lasting approximately half an hour.

PRIMING TASK: In order to plant some unresolved queries into their memory, each participant was given a “flag test” in which they had to identify the countries represented by five different flags. They were told that if they could not do so straight away they could think about the flags during the rest of the session and provide any additional answers at the end. Each correct identification carried a reward of £1 regardless of when the answer was supplied. The reward was designed to act as an incentive to maintain a high level of importance for the unresolved queries created during the priming task. Each of the participants in a dyad was given a different set of flags to identify.

FAMILIARISATION WITH THE TABLE: Quite separately from the priming task, participants were shown the functionality of the Table and were allowed to acquire experience of its operation. During this time a special set of information items was used which contained no items that were shown on the table during the subsequent conversation task.

CONVERSATION TASK: Participants were invited to sit on opposite sides of the Table and have a conversation about “global warming”, a topic that was anticipated to be interesting and unlikely to cause offence. They were also told that they were allowed to

interact with any of the items displayed on the Table while maintaining a reasonably coherent conversation.

THE EXPERIMENT: During the conversation, one flag from each participant's priming task was included in the continuous stream of 25 items presented on the Table, a selection from which is shown in Figure 11: for participant 1 of each dyad this was item 9 and for participant 2 it was item 11. These items contained only a picture of the flag, *without any text*. In order to find the name of the country represented by the flag, and thereby find the answer to their possibly unresolved query, participants had to push the item to the centre of the Table, at which point the country name was prominently displayed. These were the intended IB targets.

Other items were included among those displayed on the Table, and for a specific purpose. Items 7, 8 and 10 were also flags, but not the ones that occurred during the priming task. These were included to see if items similar to the IB targets would also attract attention from the participants. Other items were somewhat related to the topic of conversation, for example item 12 about the weather in Manchester (where the experiment was carried out) and item 25 about global disasters. It was expected that participants would interact more with items related to the query or the conversation than with non-related items. Items 1, 2, 4, and 5 had the same headline to see if participants would notice connections between items (gist).

MEMORY TESTS: It was suggested earlier that IB is particularly associated with a threshold matching strategy, as supported by DA2. If that is the case, one would expect that irrelevant items will not be remembered afterwards. Therefore, after the conversation task, participants were instructed to write down a brief description of any of the items seen on the Table that they could remember ('free recall test'). They were also given a two-alternative forced choice recognition test on five of the pictures displayed and five of the headlines to see what items had been remembered unconsciously. Additionally, participants were asked if they had seen the three flags similar to the IB targets and if they could remember the (displayed) names of those countries. It was expected that non-target flag items might be better remembered than entirely non-relevant non-flag items.

QUESTIONNAIRE: To end the session all participants were given a questionnaire in which they were asked about their experiences of having a conversation around the Table.

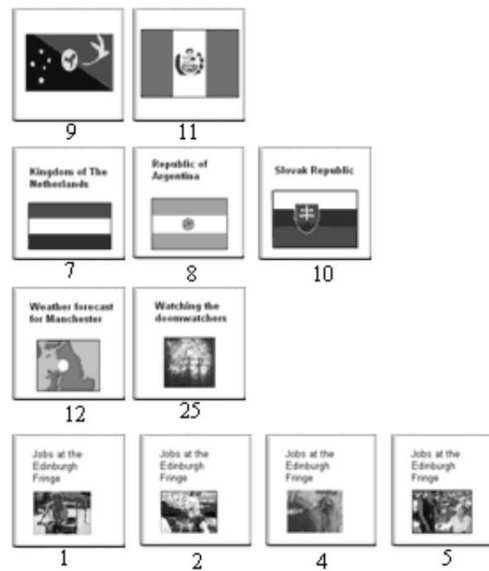


Fig 11. Items were chosen for presentation on the table based on various considerations. Five items were directly relevant to the primed query (i.e., the IB targets [9, 11], and the Non-target flags [7, 8, 10]), two were relevant to the topic of conversation [12, 25] and four had a common gist as indicated by identical headlines [1, 2, 4, 5]. All other items presented had no *a priori* relevance.

6.4.2 Results and discussion

EFFECT OF THE PRIMED QUERY ON TABLE INTERACTION: Nine of the sixteen participants knew the name of their target flag after the conversation around the table but not before. All of these participants actually interacted with the flag during the conversation. One participant knew the flag even before the conversation at the coffee table, but interacted with the flag anyway, presumably to confirm their previous answer. One additional participant indicated seeing the flag but did not think of pushing it into the middle, because he thought it was socially unacceptable. Therefore, for 10 out of 15 participants the table provided the information they were either consciously or unconsciously looking for. Whether this information was gathered by means of IB, OB or SB is hard to determine.

We can make a distinction between OB and IB on the one hand and SB on the other if we look at how often participants interacted with items related either to the query (non-target flag items) or ones related to the conversation, and compare this to interactions with entirely non-relevant items. Participants interacted 17 times in total with items related to either the query (8) or the conversation (9) out of the 40 times such items appeared on the table. They also interacted 17 times with entirely non-relevant items, but

out of the 112 times such items appeared on the table. Proportionally, therefore, relevant items had a much higher chance of being interacted with than non-relevant items (Wilcoxon signed rank test $N=8$, $Z=6$, $p(\text{one-tailed}) < .05$). It appears therefore that the coffee table did support information acquisition through either OB or IB.⁹

MEMORY PERFORMANCE: Free recall memory of any non-target items was generally poor; with only 8% of (non-flag) items being recalled on average. For non-target flag items the average recall was significantly higher (30%; Wilcoxon signed rank test $N=10$, $Z=5$, $p=.01$). This is consistent with a threshold matching strategy being employed by the participants, and with DA2. However, participants were reasonably good at correctly recognizing previously presented (non-flag) pictures in the recognition test. They were on average 68% correct (35% after adjusting for guessing; Wilcoxon signed rank test $N=16$, $Z=19$, $p<.01$).¹⁰ Participants were also able to distinguish between headlines that had been shown and ones that had not. They were on average correct 65% of the time (30% after adjusting for guessing; Wilcoxon signed rank test $N=16$, $Z=24$, $p<.02$). On the whole, there was no significant difference between the accuracy with which pictures and headlines were recognized (Wilcoxon signed-ranks test $N=16$ $Z=-48$, not significant). Items that participants interacted with were always reported correctly by those participants after the conversation. If interaction did not occur, then only 66% of (non-flag) items were correctly reported (30% after adjusting for guessing). It appears, therefore, that many more items attracted attention from the participants than just the primed target and items associated with it. These results indicate that categorization and consolidation of non-flag items had taken place, but their lack of relevance lead to forgetting as predicted by DA1 in the sense that they cannot be recalled easily.

WHERE ON THE TABLE INTERACTION STARTED: *Both* participants were observed to pick up items (both flags and other items) mostly when they had just appeared on the table. For participants sitting at the beginning of the stream, this is almost exclusively the case (80%), except for a few items picked up when items were almost half way through their trajectory. Remarkably, this was also the case for participants sitting near the end of

⁹ No difference was found between the number of interactions with gist items (3 out of 32) and non-relevant items (17 out of 112).

¹⁰ To conservatively adjust for the possibility of guessing we applied the following formula to the results of the two-alternative forced-choice recognition task: $S = R - W/(k-1)$, where S is the score after correction, R is the number of correct responses, W the number of incorrect responses, and k the number of alternative (i.e., two).

the stream. This suggests that glances at the table were directed mostly towards the point where new items appeared, despite the unfavourable viewing angle for participants sitting at the end of the stream. Only two participants sitting near the end, and who interacted with any items, picked up items exclusively from their own side of the table. It is not impossible that this pattern reflects the natural (random) pattern of eye gaze associated with IB, but it is more likely that it reflects OB behaviour. It seems likely that this bias was, at least in part, influenced by the novelty of the artefact.

6.4.3 Conclusions

The extent to which the DAs beneficially influenced the design of the Table can only be implicit. Supportive evidence is provided by (1) the proportion of participants for whom interaction with the Table yielded answers to their questions, and (2) the extent to which non-flag items attracted attention without the participant being aware of it, as revealed by the free-recall and recognition tests. We must point out, however, the inherent difficulty in measuring successful task outcome in implicit tasks like these.

7. SUMMARY AND CONCLUSIONS

As a method of informing design decisions we have proposed, and illustrated by examples, the concept of Design Actions. They are derived from cognitive theory (in our case CSTM) in the context of a behavioural model (in our case OB and IB). Intended to inform interaction designers, the design actions are intentionally presented in a format that does not require an understanding of the underlying cognitive theory. We suggest that the value of design actions is enhanced by the identification of both the beneficial and detrimental consequences of applying them. Undoubtedly, there may be debate about the identification of Issues. For example, the issue of information overload *clarifies* the relevance of DA1, and has therefore been identified as an Issue. However, information overload could clearly not be identified as a Downside, as that would imply using the DA in situations to which it is not relevant.

To place Design Actions in context we first emphasise that their use is not intended to replace the familiar iterative design cycle. Indeed, in the spectrum of approaches to interaction design, Design Actions may eventually come to be seen as complementary to other approaches. For example, whereas Design Patterns (e.g., Borchers, [2001]; Tidwell, [2006]; Van Welie, [2002]) offer broader granularity, reflect the subjective assessment of existing designs, are usually qualitative in their advice and offer scope for multiple solutions, Design Actions, by contrast, have narrow granularity, are based rigorously upon experiment and theory and are usually quantitative in their advice (see Table II).

Table II. A comparison between Design Actions and Design Patterns.

Characteristic	Patterns	Design Actions
Basis	Based on a necessarily subjective assessment of existing designs ("Experience based")	Based on specific experimental data or models derived from such data ("Theory based")
Granularity	Broad	Narrow
Identification	by the PROBLEM to be experienced by the USER	by the GOAL of the INTERACTION DESIGNER
Output	Usually QUALITATIVE in the form of guidance supported by examples	Usually QUANTITATIVE and specific
Interpretation	Scope for creativity and interpretation by the interaction designer	Intentionally, little scope if any for creativity

DAs are also complementary to claims which serve a similar purpose, but are derived differently (Sutcliffe, [2000]). Continuing work in claims (see e.g., Wahid [2006]) may well lead to developments in the formulation and use of DAs.

Proof of the validity of specific design actions poses significant challenges for experimentation, as has been demonstrated by the reported evaluations of the RSVP Browser and the Coffee Table. Leaving aside the value *per se* of these artefacts, our experimental results do not of course offer conclusive proof of the validity of the various DAs we have identified. Rather, in addition to the absence of any contradictory evidence, the experiments and related models have, in many respects, pointed the way in which future developments might proceed. The integration of experimental data to refine and quantify the CSTM model will no doubt continue in the light of recent knowledge of human visual processing (Ögmen & Breitmeyer, [2006]) and the definitions of three types of browsing may well help to clarify the nature of browsing behaviour. Similarly, the use of a priming task to induce IB illustrates an experimental method in which the challenging task of disambiguating between SB, OB and IB may be addressed. Despite these considerable challenges we feel that they are well worth addressing if there is the possibility that they may lead to useful tools for interaction designers, and we therefore advocate use of the framework we have proposed to generate new Design Actions from other cognitive theories within the context of a range of behavioural models.

By application of the framework advocated in this paper, the ideal situation might arise in which there exist a range of interconnected theories at both cognitive and

behavioural level (cf. Barnard et al. [2000]) giving rise to hundreds or even thousands of Design Actions. In such a situation it will be important that a mechanism be available by which DAs can be easily retrieved depending on their relevance to a particular design objective. However, this is a problem that applies to any form of design guidance, whether in the form of guidelines, patterns, claims or Design Actions. Two types of solutions have been proposed to solve this problem. The first is to organise design guidance according to how similar they are. For example, design patterns are often organised into hierarchical categories making up a pattern library. The main problem with this approach is the danger of organising guidance into categories that do not capture the context in which the guidance will be applied. The second type of solution is to link design guidance to behavioural models. Designers can walk through these models to find paths that match the purpose of the intended designs and thereby find associated design guidance (Sutcliffe & Carroll, [1999]) and resolve potential conflicts between applicable DAs. The latter solution is particularly applicable to the organisation and retrieval of DAs as their development is explicitly linked to models of interaction.

Future research efforts by us and others will have to extend and validate the framework presented here further. For example, the design implications of a range of cognitive theories in the light of several behavioural models need to be explored leading to additional DAs. However, some may feel that the use of cognitive theory as a starting point for deriving design actions risks the observation that most cognitive theories are still being revised and that very few have reached maturity. However, we do not agree with the conclusion that “. . . we have little prospect of developing final answers to questions about the nature of human activity . . . at the level of detail that would provide specific guidance to designers” (Carroll, [1995]). We are of the opinion that having something reasonably well understood to go by is better than having nothing.

ACKNOWLEDGMENTS

Part of the work reported here is an outcome of research funded by the UK Engineering and Physical Sciences Research Council (GR/N20881/01): this support is gratefully acknowledged. The research has benefited from discussions with many people including Alistair Sutcliffe and Andrew Howes. We are grateful to Kent Wittenburg, Kathy Ryall, Chia Shen and Frederic Vernier at MERL for the provision of, and assistance with the Mitsubishi DiamondTouch interactive surface. Kenny Tong implemented the RSVP Browser and carried out its evaluation. We further thank the anonymous reviewers of

earlier drafts of this paper whose comments contributed considerably to the quality of the paper.

REFERENCES

- BARNARD, P. 1991. Bridging between basic theories and the artifacts of human-computer interaction. In *Designing interaction: psychology at the human-computer interface*, Cambridge University Press, 103-127.
- BARNARD, P., MAY, J., DUKE, D. AND DUCE, D. 2000. Systems, interactions, and macrotheory. *ACM Transactions on Compute.-Human Interaction*, 7, 222-262.
- BORCHERS, J. 2001. *A Pattern Approach to Interaction Design*. Chichester, UK: Wiley.
- CARMEL, E., CRAWFORD, S. AND CHEN, H. 1992. Browsing in Hypertext: A cognitive study, *IEEE Transactions on Systems, Man and Cybernetics*, 22, 865-884.
- CARROLL, J.M. 1995. Introduction: The scenario perspective on system development. In *Scenario-based Design: Envisioning Work and Technology in System Development*, J.M. CARROLL, Ed. New York, John Wiley & Sons, 1-17.
- CARROLL, J. M. and ROSSON, M. B. 1992. Getting around the task-artifact cycle: how to make claims and design by scenario. *ACM Transactions on Information Systems*, 10, 181-212.
- CATLEDGE, L.D. AND PITKOW, J.E. 1995. Characterizing browsing strategies in the World-Wide Web, *Journal of Computer Networks and ISDN Systems*, 27, 1065-1073.
- CLARK, H. H. 1996. *Using Language*. Cambridge University Press, Cambridge, UK.
- COLTHEART, V. 1999. *Fleeting Memories*. MIT Press, Cambridge, MA.
- DE BRUIJN, O. AND SPENCE, R. 2000. Rapid serial visual presentation: A space-time trade-off in information presentation, *Proceedings of the ACM Working Conference on Advanced Visual Interfaces, AVI'2000* (Palermo, Italy, May 2000), ACM Press, pp. 189-192.
- DE BRUIJN, O., SPENCE, R. AND CHONG, M.Y. 2002. RSVP Browser: Web Browsing on Small Screen Devices, *Personal and Ubiquitous Computing*, 6, 245-252.
- DE BRUIJN, O. , SPENCE, R., AND TONG, C.H. 2001. Movement in the Web, *Extended Abstract Proceedings CHI2001* (Seattle, WA, 31 March – 5 April, 2001), ACM Press, 209-210.
- DE BRUIJN, O.. AND TONG, C. H. 2003. M-RSVP: Mobile Web browsing on a PDA. In *People and Computers XVII - Designing for Society*, E. O'NEILL, P. PALANQUE AND P. JOHNSON, Eds. Springer, Bath, UK, 297-311.
- FOSTER, A. AND FORD, N. 2003. Serendipity and information seeking: an empirical study, *Journal of Documentation*, 59, 321-340.
- INTRAUB, H. 1980. Presentation rate and the representation of briefly glimpsed pictures in memory. *Journal of Experimental Psychology-Human Learning and Memory* 6, 1-12.
- INTRAUB, H. 1981. Rapid Conceptual Identification of Sequentially Presented Pictures. *Journal of Experimental Psychology-Human Perception and Performance* 7, 604-610.
- INTRAUB, H. 1984. Conceptual masking - the effects of subsequent visual events on memory for pictures. *Journal of Experimental Psychology-Learning Memory and Cognition* 10, 115-125.
- INTRAUB, H. 1999. Understanding and remembering briefly glimpsed pictures: Implications for visual scanning and memory. In *Fleeting Memories*, V. COLTHEART, Ed. MIT Press, Cambridge, MA,
- JONES, M., MARSDEN, G., MOHD-NASIR, N., BOONE, K. AND BUCHANAN, G. 1999. Improving Web interaction on small displays. *Proceedings of the Eight International Conference on World Wide Web (WWW8)*, Toronto, Canada, May 1999, pp. 1129-1137.
- KAASTEN, S., GREENBERG, S. AND EDWARDS, C. 2002. How people recognize previously seen Web pages from titles, URLs and thumbnails. In *People and Computers XVI (Proceedings of HCI'02)*, X. FAULKNER, J. FINLAY AND F. DETIENNE, Eds. Springer-Verlag, 247-266.
- LAVIE, N. 2000. Selective attention and cognitive control: Dissociating attentional functions through different types of load, In *Attention and performance: Vol. XVIII. Control of cognitive processes*, S. MONSELL & J. DRIVER, Eds. Cambridge, MA: MIT Press, 175-197.
- MIYATA, Y. AND NORMAN, D.A. 1986. Psychological issues in support of multiple activities. In *User Centred System Design: New Perspectives on Human-Computer Interaction*, D.A. NORMAN AND S.W. DRAPER, Eds. Hillsdale, NJ, LEA, 265-284.
- MOYES, J. AND JORDAN, P.W. 1993. Icon design and its effect on guessability, learnability, and experienced user performance, In *Proceedings of the HCI'93 Conference on People and Computers VIII* , J.L. ALTY, D. DIAPER AND S. GUEST, Eds. Cambridge, UK: Cambridge University Press, 49-59.
- NEISSER, U. 1967. *Cognitive Psychology*. Prentice-Hall, Englewood Cliffs, NJ.
- NORMAN, D. A. 1986. Cognitive engineering, In *User Centered System Design: New Perspectives on Human-Computer Interaction*, D. A. NORMAN AND S. DRAPER, Eds. Hillsdale, NJ, LEA.
- NOWELL, L., HETZLER, E. AND TANASSE, T. 2001 Change blindness in information visualization: a case study. In *Proceedings of Information Visualization, 2001. INFOVIS 2001. IEEE Symposium on*, 15-22.

- OGMEN, H., AND BREITMEYER, B. 2006. *The First Half Second: The Microgenesis and Temporal Dynamics of Unconscious and Conscious Visual Processes*, Cambridge, MA, MIT Press.
- ÖQVIST, G., GOLDSTEIN, M. AND CHINCHOLLE, D. 2004, Assessing usability across multiple user interfaces. In *Multiple User Interfaces: Cross Platform Applications and Context-Aware Interfaces*, A. SEFFAH AND H. JAVAHERY, Eds. New York, NY, John Wiley & Sons, 327-349.
- PAYNE, J.W., JOHNSON, E.J. AND BETMANN, J.R. 1993. *The Adaptive Decision Maker*. New York, NY, Cambridge University Press.
- PIROLI, P. 2005. Rational Analyses of Information Foraging on the Web, *Cognitive Science* 29, 343-373.
- PIROLI, P. AND CARD, S. 1995. Information foraging in information access environments In *Proceedings of the SIGCHI conference on Human factors in computing systems* ACM Press/Addison-Wesley Publishing Co., Denver, CO, 51-58.
- POTTER, M. C. 1975. Meaning in visual search. *Science* 187, 965-966.
- POTTER, M. C. 1976. Short-Term Conceptual Memory for Pictures. *Journal of Experimental Psychology-Human Learning and Memory*, 2, 509-522.
- POTTER, M. C. 1993. Very Short-Term Conceptual Memory. *Memory & Cognition*, 21, 156-161.
- POTTER, M. C. 1999. Understanding sentences and scenes: The role of Conceptual Short-Term Memory. In *Fleeting Memories*, V. COLTHEART, Ed. MIT Press, Cambridge, MA.
- POTTER, M. C. AND LEVY, E. I. 1969. Recognition memory for a rapid sequence of pictures. *Journal of Experimental Psychology*, 81, 10-&.
- POTTER, M. C., STAUB, A., RADO, J. AND O'CONNOR, D. H. 2002. Recognition memory for briefly presented pictures: The time course of rapid forgetting. *Journal of Experimental Psychology-Human Perception and Performance*, 28, 1163-1175.
- RAYMOND, J. E., SHAPIRO, K. L. AND ARNELL, K. M. 1992. Temporary suppression of visual processing in an RSVP task - an attentional blink. *Journal of Experimental Psychology-Human Perception and Performance*, 18, 849-860.
- RYALL, K., FORLINES, C., SHEN, C. AND RINGEL-MORRIS, M. 2004. Exploring the effects of group size and table size interactions with tabletop shared-display groupware. In *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW)*, November 2004, ACM Press, 284-293.
- SPENCE, R. 1999. A framework for navigation, *International Journal of Human-Computer Studies*, 51, 919-945.
- SPENCE, R. 2002. Rapid, serial and visual: A presentation technique with potential, *Information Visualization*, 1, 13-19.
- SPENCE, R. 2004. Sensitivity, residue and scent. Concepts to inform interaction design for the support of information space navigation. *Information Design Journal*, 12, 163-180.
- SPENCE, R., WITKOWSKI, M., FAWCETT, C., CRAFT B. AND DE BRUIJN, O. 2004. Image presentation in space and time: Errors, preferences and eye-gaze activity. In *Proceedings of the ACM Working Conference Advanced Visual Interfaces, AVI2004* (Gallipoli, Italy, May 25-28, 2004), ACM Press, 141-149.
- STANDING, L. 1973. Learning 10,000 Pictures. *Quarterly Journal of Experimental Psychology*, 25, 207-222.
- STATHIS, K., DE BRUIJN, O. AND MACEDO, S. 2002. Living Memory: agent-based information management for connected local communities, *Interacting with Computers*, 14, 663-688.
- SUTCLIFFE, A. 2000. On the effective use and reuse of HCI knowledge. *ACM Transactions on Computer-Human Interaction*, 7, 197-221.
- SUTCLIFFE, A. G. AND CARROLL, J. M. 1999. Designing claims for reuse in interactive systems design. *International Journal of Human-Computer Studies* 50, 213-241.
- SUTCLIFFE, A. G. AND KAUR, K. 2000. Evaluating the usability of virtual reality user interfaces. *Behaviour & Information Technology*, 19, 415-426.
- SUTCLIFFE, A. G. AND RYAN, M. 2000. Model mismatch analysis: Towards a deeper explanation of users' usability problems. *Behaviour & Information Technology*, 19, 43-55.
- TIDWELL, J. 2006. *Designing Interfaces*. O'Reilly Media.
- TSE, T., MARCHIONINI, G., DING, W., SLAUGHTER, L. AND KOMLODI, A. 1998. Dynamic key frame presentation techniques for augmenting video browsing In *Proceedings of the working conference on Advanced Visual Interfaces* ACM Press, L'Aquila, Italy 185-194.
- VAN WELIE, M. 2002, Interaction Design Patterns, <http://www.welie.com/patterns/index.html> (accessed 01/08/2007).
- WAHID, S. 2006. Investigating design knowledge reuse for interface development. In *Proceedings of the 6th ACM Conference on Designing Interactive Systems* (University Park, PA, USA, June 26 - 28, 2006). DIS '06. ACM Press, New York, NY, 354-356.
- WALENSTEIN, A. 2004. Theory adapters as discipline coordinators. In *Proceedings of the 2004 ACM workshop on Interdisciplinary software engineering research*, Newport Beach, CA, USA, ACM Press, 8-17.
- WITTENBURG, K., ALI-AHMED, W., LALIBERTE, D., AND LANNING, T. 1998. Rapid-fire image previews for information navigation. *Proceedings of Advanced Visual Interfaces AVI'98* (L'Aquila, Italy, May 1998), ACM, 76-82.
- WITTENBURG, K., FORLINES, C., LANNING, T., ESENTER, A.W., HARADA, S. AND MIYACHI, T. 2003. Rapid serial visual presentation techniques for consumer digital video devices. *ACM Symposium on User Interface Software and Technology (UIST)*, November 2003, ACM Press, 115-124.

YANTIS, S. AND JONIDES, J. 1990, Abrupt visual onsets and selective attention - voluntary versus automatic allocation. *Journal of Experimental Psychology-Human Perception and Performance* 16, 121-134.
 ZEIGARNIK, B. (1938). On finished and unfinished tasks. In *A Source Book of Gestalt Psychology*, W. D. ELLIS, Ed. NY: Harcourt-Brace.

Received January 2005; revised March 2007 and October 2007

APPENDIX

ID	DA1
Title	Maximisation of information exposure
Description	Expose users to as many representations of information items as possible commensurate with maximising the likelihood of those representations being correctly interpreted in terms of the information being represented.
Effect	The more of these items being presented to the user the greater the likelihood will be that one is perceived that is a relevant to a user's interests.
Upside	(1) Queries need not be explicitly formulated, as the relevant items will be recognized as such when encountered. (2) Irrelevant items will be forgotten with no cost in cognitive effort
Downside	Any display area is finite, giving rise to a trade-off between the number and size of items being presented simultaneously, and their presentation duration.
Issues	Comment is often made concerning the effects of information overload. This is not so relevant here, because no additional cognitive effort is involved in perceiving, interpreting and then forgetting information that is irrelevant. However, any action triggered does require focused attention. It is therefore important that this DA be applied to situations in which the likelihood of encountering a relevant item is relatively small. If multiple items of interest are likely to be encountered, then it is important that either a) the items can be easily prioritised, or b) the items can be dealt with sequentially.
Theory	Current theory suggests that irrelevant information is rapidly forgotten at no additional cost.

ID	DA2
Title	Presentation for immediate action
Description	Presenting selectable information as images that can be inspected for at least 100ms allows a large number of items to be presented within an acceptably short time, provided 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 immediate action supportive of a satisficing approach to browsing.
Upside	Many information items can be presented in a relatively short period of time either sequentially or concurrently (see DA1).

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 recognize 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. (3) Given the limits to people's reaction time, a mechanism needs to be provided that allows retreat after overshooting the target.
Theory	The ability to automatically and rapidly match an information need with a viewed image

ID	DA3
Title	Presentation for exhaustive attention and selection from a group
Description	Where a <i>group</i> of potentially selectable information sources, each represented by an image (with or without brief text) is presented, ensure that each item can be fixated for at least 500ms if an optimising approach to selection is intended. At least 20 choices at a time can be presented in this way either sequentially or concurrently.
Effect	This action ensures that each in a sequence of information items can be assessed individually and an appropriate and reasonably confident selection made of one of them once all the options have been presented.
Upside	The probability of selecting the most relevant information item; i.e., the likelihood of efficiently evaluating each option in relation to a (unconscious) goal and selecting the best match.
Downside	Items may be forgotten if the decision process takes too long.
Issues	(1) Because there is a risk that information items are forgotten before a decision about their relevance can be made, it is therefore important that user be able to repeat the presentation without undue effort. (2) Some minimal form of representation of all the items have to remain visible on the display in order to allow users to select relevant items after presentation.
Theory	The buffering of information items in memory allowing a decision to be made after a considerable number of items has been presented, together with the need to avoid interference between two consecutively presented relevant items.

ID	DA4
Title	Presentation for consideration and selection from multiple items
Description	Where a <i>group</i> of potentially selectable information sources, each represented by an image (with or without brief text) is presented in a

	sequence in a 'preview-consider-select' strategy, ensure that each item can be fixated for between 1.5 and 2 seconds . At least 16 choices at a time can be presented either sequentially or concurrently in this way.
Effect	This action ensures that each in a sequence of information items can be assessed individually and in relation to other items, and an appropriate and a confident selection of one or more items can be made.
Upside	Selection need not be restricted to a single item.
Downside	Assessment of a large number of items requires a considerable amount of time for the presentation of all information.
Issues	(1) When a sequential presentation of items is used, some minimal form of representation of all the items has to remain visible on the display in order to allow the user to select relevant items after presentation. (2) Users are able to select relevant items far better and more easily when pictures are being presented compared to when text is used.
Theory	The ability to recall previously presented images (and to a lesser extent text).

ID	DA5
Title	Presentation to enhance the interpretation of the gist of narratives and themes.
Description	The selection and sequential presentation, at up to 10 per second , of an ordered group of key images representative of a narrative or theme.
Effect	The awareness of the gist of the narrative and extraction of the theme.
Upside	Provided key frames are appropriately chosen, acquisition of the gist of a narrative can be obtained in a very small fraction of the time it would take to view the entire narrative. Similarly, the theme of a collection of images can be discovered in a fraction of the time it would take to view all the images in the collection (provided the collection contains a certain minimum number of items).
Downside	The need to make a selection of key frames to effectively represent the narrative.
Issues	Algorithms for the automatic extraction of representative images exist, but these are not always appropriate. In many situations appropriate selection depends on the possession of an adequate knowledge of the domain of the theme. Alternatively, selections can be made manually by someone with the right amount of domain knowledge.
Theory	The integration of information over time into meaningful narratives.

ID	DA6
Title	Presentation to enhance the detection of changes
Description	Ensure that unrelated information items each have a visual impact, particularly in the sense that it involves a new gist. This usually translates into a requirement to treat each item in a sequence or collection as independent so that no comparison with a previous item

	is necessary in the interpretation of the significance of the current item.
Effect	This design action is particularly important when items are displayed or fixated intermittently. Intermittent displays may include the presentation of an irrelevant image in between images, which prevents any change from being perceived directly. When users are not constantly looking at a display, any change occurring during the time the eye gaze is diverted elsewhere may not be perceived. As long as the gist of each newly fixated item is significantly different from the previously fixated item, it is ensured that the new display is interpreted and evaluated even if the change itself is not detected.
Upside	The chance that a relevant image is correctly interpreted is increased.
Downside	This design action limits, to some extent, the kind of changes that can be designed. For example, subtle changes to the appearance of an icon reflecting a change of system state will often go unnoticed with the consequence that the user might be unaware of the correct system state.
Issues	Whether a visual change is sufficiently obvious for it to be detected is essentially unpredictable and depends very much on the nature of the visual media in the collection. The designer will have to experiment with the media to establish the visual impact of a change.
Theory	The forgetting of detailed visual representations of scenes after very short delays, combined with the much longer retention of conceptual representations of scenes.