Visibility

Object visibility in dynamic scenes

route
outdoor media analytics

Biographies

Dr Paul Barber was employed at Birkbeck in the University of London where for over thirty years he taught courses on perception, psychological statistics and experimental design, computing and ergonomics. He was head of Birkbeck's Department of Psychology from 1988 to 1992. He is now Emeritus Reader in Psychology at Birkbeck. Dr Barber was research supervisor to sixteen PhD students, his own doctoral research being on visual input processes. He is an Associate Fellow of the British Psychology Society and is a Chartered Psychologist. He was the author/co-author of a number of psychology textbooks and many research papers in refereed scientific journals; he was Psychology Editor for the journal Ergonomics for over ten years. He has been a consultant for Route since 1995, closely identified with its programme of research on poster panel visibility.

Dr Paul Wilson teaches at the University of Lodz, Poland. He has also held the position of Honorary Research Fellow at Birkbeck (University of London) since gaining his PhD in Psychology there in 2000. He has also worked as a researcher on a number of projects at the University of East London and Birkbeck. He is the author/co-author of a number of research papers and chapters in refereed scientific journals and volumes. During his employment at Birkbeck between 2008 and 2011 he carried out research on poster panel visibility for Postar (now Route).

Contents

Foreword

Our industry has a proud tradition of independent eye-tracking research that dates back two decades. In 2013 and 2014, we re-published the original studies that sought to estimate the visibility of posters for car drivers and passengers[1] as well as for pedestrians[2].

Here we bring things up to date by collating all that we so far know about the effect of movement. By this we mean both the poster sites that incorporate a form of scrolling or digital display, as well as objects that themselves move such as buses and taxis.

What we have learnt from the research has been incorporated by Route and it helps to inform the way that we calculate the probability that people will see various forms of moving, or dynamic, image.

It is an under-studied area and much of the work that was carried out by the Department of Psychology at Birkbeck College, in the University of London, was experimental in nature. As it remains a work in progress, we have also gathered here an extensive review of academic literature on the subject that might help to inform the next steps for inquiry.

A full list of our various studies can be found in the appendix.

James Whitmore
Managing Director
Route

August 2015

[1] "Poster panel visibility for drivers and passengers: a first look" February 2013

[2] "Estimating the visibility of poster panels for pedestrians" March 2014

Abstract

Object visibility in dynamic scenes

Paul Barber and Paul Wilson
Department of Psychological Sciences
Birkbeck College
University of London

The question of the extent to which factors associated with dynamic properties of poster panels influence their visibility is addressed in this report. Three exploratory empirical studies are reported. The account is accompanied by a literature review of factors that may affect object visibility and visual attention, using basic and applied sources of evidence.

Results for comparisons between stationary and dynamic formats are reported for all three studies. A variety of dynamic formats and properties featured in video clips captured from a variety of locations in London for the studies. They included the ubiquitous scrolling operation commonly seen in roadside, travel and commercial environments. Digital screens were also included, as were panels on moving vehicles (buses and taxis). Because the research was very much exploratory in nature, any comparisons were hedged about with a greater deal of uncertainty than usual. The tendency in all cases was for dynamic material to out-perform the stationary forms by a modest amount. A more rigorous and extensive body of research is required to take this further.

The review of research is broadly scoped but is intended to bear on the use of billboards with an electro-mechanical scrolling system and electronic systems potentially supporting dynamic displays. The possibility that a moving stimulus can capture attention without the intention of the observer was uppermost in our considerations. The evidence on this question is reviewed.

A different perspective on the matter of dynamic stimuli has been posited. It is argued that there is generally a confounding between a moving stimulus and its onset/offset properties and this offers an alternative explanation for any ostensible motion effects. Experimental methods for dissociating motion and onset/offset effects, as well as other confounding operations, are noted, and the relevant evidence is provided.

Another section of the review deals with the findings from the area of Web-based advertising, which though quite separate from an interest in outdoor advertising does offer some lessons for the current area of research.

A more obviously connected area of interest is to do with driving safety issues arising from the introduction of electronic billboards. Most of this research is fully focussed on safety matters and only tangentially impacts the question of visibility. The conclusions of investigations are that safety concerns are not compromised by the use of electronic billboards but the research is roundly criticised by a number of experts. None of the safety studies reviewed has used billboards with truly dynamic properties, except for simple transitions between images. This question has been addressed by one study using a virtual reality platform.

Our three empirical studies were not closely steered by the literature but were informed by key findings such as the distinction between motion and change features. These were pilot studies and clearly established the need for more extensive and technologically well-founded research. This should take into account the substantial research literature that is now available to drive the topics addressed by the present studies.

Preface

This report provides an integrated account of three exploratory studies of the efficacy of dynamic imagery for poster frames/panels. There is also a review of the related literature. Unusually we have presented the review *after* the account of the empirical research to which it is attached. The reason for this is that the three investigations were to answer practical questions with some urgency. While they were preceded and to some extent steered by a contemporaneous literature review, the actual fabric of the research was mainly constrained by considerations about what was there in the real world to be evaluated. In addition funding priorities dictated that these studies should be exploratory, limited in scope, and deliverable on a tight schedule. In the quieter space of recent times, and the launch of Route, it has been possible to expand the scope and currency of the review, which is why it now occupies its unorthodox but honest position in the report. The three exploratory or pilot studies focussed on a number of aspects of moving imagery associated with the stock of UK poster advertising panels in 2008-9. The presumption had been expressed within the industry that dynamic imagery would confer a distinct advantage in terms of panel visibility, but empirical support for this was lacking.
The principal aim of the present studies was to provide data to fill this gap in the evidence.

Chapter 1
Phase One study summary

The Phase One study was the first of three investigations examining the contribution to poster panel visibility of their dynamic properties. This study considered two core aspects of movement (not including that of the observer): movement internal to the panel (as in scrolling panels) and movement of the panel itself (exemplified by bus poster panels on buses themselves on the move). Published research was reviewed and empirically supported trends were noted. The somewhat surprising tenor of this research was that movement per se was not a powerful attractor of attention; according to the evidence what seemed more likely to capture attention was some change property (like the shift from being stationary to being in motion). A more comprehensive and up-to-date literature review now forms part of the report. The method used for the empirical part of the research was a traditional laboratory speeded reaction time task, intended to emulate elements of the activity of driving a car; the participant responded as quickly as possible to signal the presence/absence of a visual target (presented peripherally) while monitoring the state of a sequence of centrally displayed items. The type of peripheral signal was varied according a four-factor randomised experimental design. These factors were Type (vertical scroll vs. lateral shift), Motion (moving vs. stationary), Size (two levels) and Offset (also on two levels). An even more complex design could be expanded (e.g., to include Side – left vs. right field of view), but this was not practicable. Results were expressed as reaction time measures but they were also translated into hit rate scores for the purposes of analysis. The findings from both reaction time and hit rate data contained some useful confirmatory indications of the sensitivity of the research method to key variables, but generally with qualifications. The complex battery of outcomes included a mixture of results that were expected as well as some surprises (e.g., moving stimuli were more effective, but only for laterally shifted signals; and larger targets were easier to spot, but only for vertically scrolling targets). The method could be applied to other configurations of the possible design factors on a larger scale than here. In the event the step was taken in the Phase Two study to devise a task incorporating actual video recordings, albeit with practical limitations on which design factors could be manipulated.

Introduction

The visibility research conducted for Postar (now Route) relied extensively on eye-tracking techniques. All but one of the investigations used stationary photographic images of scenes containing poster panels. The exception was the so-called driver attention study which in 1999/2000 used an in-car eye-tracking system, resulting in videotaped recordings of the visual behaviour of drivers and passengers on roads in Nottingham in the UK; this enabled the participant's attention to be ascertained by measuring gaze position as the vehicle proceeded. The videotapes were later re-analysed (reported as Barber and Sanderson, 2005/2006) to identify how frequently the gaze of the participants was directed at poster panels (fixed panels and panels on buses). This exploratory re-analysis took advantage of the existence of the eye movement recordings and it needs to be stressed that the study was not intended or designed as a basis for obtaining data on visibility hit rates. Although this did not diminish its validity – indeed arguably the opposite is the case – the study was undoubtedly limited because the routes were not chosen to maximise the participants' exposure to poster panels, which would have supplied a more extensive sampling of panel types and locations. The present three-phase project is intended to extend Route's visibility research by obtaining eye movement data while the participants are viewing images in which there is movement within the scene.

This aspect of dynamic imagery will be the focus of the research for the time being, not movement on the part both of the scene <u>and</u> the observer which we address later. In this respect the technical achievement is less than that of the driver attention study. Among the most important reasons for imposing a technical limit in this way is the fact that satisfactory in-car eye-tracking equipment is not readily available. This is also the case regarding software solutions for the complex problem of tracking dynamic targets.

A plan in three phases for the dynamic imagery research project was presented for comment in June 2008 and a revised version of the proposal for the first phase was approved in November 2008. The document consisted of a literature survey that is here integrated into the literature review forming a substantial later part of this report. The latter was a review of a set of options for taking dynamic panels into account and is presented next. It reflected the concerns of the moment: *"This paper outlines some possible directions for Postar's visibility research programme, taking account of the fact that interest has recently centred on how to incorporate the dynamic aspects of outdoor advertising artefacts and viewers' exposure to them."*

This study considered two core aspects of movement (not including that of the observer): movement internal to the panel (as in scrolling panels) and movement of the panel itself (exemplified by bus poster panels on buses themselves on the move).

Dynamic imagery research: Technical solutions

The 2008 review of possible technical solutions that formed the basis of the discussions leading to the studies referred to as Phases One, Two and Three is presented in this section.

The primary interest of the empirical work expected to be required concerned static structures that display moving images – that is, scrolling or rotating poster panels and digital screens of various sizes. The secondary interest was in moving structures that display advertising – for example, buses and taxis. Five approaches (technical options) towards these objectives were considered and evaluated.

In the event Options 1, 2 and 5 were for various reasons judged at the time to be inappropriate or untimely. This left Options 3 and 4, and the final recommendation was to select Option 3 followed by Option 4 (though they effectively overlapped). Option 3 would give an early indication of any visibility differential between scrolling and non-scrolling panels, and hence an interim visibility adjustment for dynamic imagery. It would also be very useful in providing a test-bed for methodological elements of Option 4. The latter would provide the framework for final visibility adjustments for dynamic imagery.

Technical options

1. Driving simulator/virtual reality approach

The virtual reality approach has been mooted from time to time as a possible vehicle for visibility research, and the driving simulator method can be seen as a variant on this. Both would support dynamic imagery and could be combined with eye-tracking and both would enable poster panels to be inserted in the virtual/simulator experience. There had previously been little enthusiasm for these approaches but the option was tabled as a practical option to explore. The driving simulator option was brought back into contention by a report on advertising and driver distraction (Young and Mahfoud, 2007). The illustrations and task description from that report indicated the quality of the driving experience likely to be delivered by then affordable current technical platforms. This report was not thought to be the place for a full critique of this study, but it is worth noting that the scenarios depicted seemed simplistic, even cartoon-like. A virtual reality approach would raise similar concerns and others about the tolerance of respondents to the experience. Both methods seemed likely to entail a lengthy set-up time, and neither was therefore recommended for the immediate purposes of the research, and they did not continue in contention.

2. Real-world exposure with eye-tracking

There was one other approach that was excluded though it is arguably the final and optimal goal for visibility research. This was to record the eye movements of people in real-world situations, exposed to the range of panel types of interest. This appears to have been the methodology used in Australia for their industry-sponsored visibility research. It was tabled as the “ultimate” option, but it was judged to be unaffordable (a dedicated and suitably equipped car would need to be available) and practically daunting in the short-term since numerous caveats about implementation would apply. In due course it will be worth reviewing this and the two previous options to identify the scale of the obstacles that would be now encountered when the cost and time frame for the development of this option might have become acceptable.

The three remaining methods (Options 3, 4 and 5) were ordered as steps towards the “ultimate” version (Option 2). All three involved the presentation of dynamic images (beginning with the ubiquitous scrolling type of panel, but progressing to include digital panels too). The last pair both entailed use of eye-tracking. All three entailed the presentation of moving images. The task for use in Option 3 is illustrative of the options then available; it was chosen to provide an enabling platform to Option 4.

3. Laboratory-based proxy

This approach requires a means of recording responses to visual stimuli with designed properties; this could be done by custom-built software but was conveniently provided in the form of a software package for creating and running experiments: E-Prime 2.0 (http://www.pstnet.com/products/e-prime/), a new release of a recommended and widely used system.

The preliminary literature review (now integrated into the updated account that features as the fourth main part of this report) identified stimulus properties most likely to influence attention, and which suggest formats that would be likely to perform well through their dynamic properties. The choice of task and its design should reflect the range of properties to be investigated (change onset, motion, etc.). These properties are represented in panels with internal motion (e.g., scrolling) and those that are in external motion (e.g., moving vehicular panels).

Option 3 employs a laboratory speeded-decision task as used in our direct search studies. The participant would search for specified peripheral targets (possibly looking for roadside posters), while monitoring a central location for other key signals. This simulates crucial aspects of a driving task – attention must be focussed directly ahead but stimuli of interest to the left or right may also be inspected when the opportunity occurs. Performance is to be measured by the speed of response (reaction time). In our previous studies in which an active search task was used, reaction times were converted to visibility hit rates.

Differentials between conditions may be established by this means. For instance, the difference in hit rates for scrolling and non-scrolling 6-sheets could be measured. The study would need to include anchor contrasts, such as stationary 6-sheets vs. 48-sheets, to provide a basis for calibrating the findings back to results from our static image eye-tracking studies.

The stimulus material could include presentations of panels on the move to represent panels on vehicles. These stimuli could either be video clips of actual panels or some other graphic representation thereof. A wide range of formats and other relevant site features could be assessed using designs that assessed them individually and in combination.

This method could be implemented much earlier than the next because it should entail less software development and simpler data processing. In fact Approach 3 is a stepping-stone to Approach 4 because, depending on its implementation, it could achieve the preparation of stimulus material and research design needed for the latter. It certainly would deliver indicative findings much sooner.

4. Laboratory-based proxy with eye-tracking

This approach requires the means of recording eye movements to dynamic scenarios. Film or video sequences would be created that contained poster panels with prescribed properties (size, eccentricity, distance) according to a balanced experimental design. Suitable eye-tracking software (and hardware) with the capability of handling video images and substantial amounts of raw data would be required. Crucially, the locations of the panels would have to be specified, and preferably these location files would need to be automatically tracked by the software. This would otherwise be a rather lengthy and painstaking task. Depending on the implementation of Option 3 its stimulus material could be used, with minor adaptations of the procedure. A shift to a less abstract approach might lose some experimental control but ensure a more realistic experience for the participants.

5. Laboratory-based quasi-reality with eye-tracking

The final approach requires a means of recording eye movements to dynamic scenes with natural uncontrolled properties. Because the trajectories of the panels would be much more complex than in the previous case, the task of specifying their locations would be much more difficult. Work is in progress on developing algorithms for tracking objects that are moving in time and space, but this was out of reach at the time of this study. The method is essentially the same as used by TAB in the US for its visibility research, which seems partly for technical reasons to have been a very lengthy affair. The scale of the study as described here would be smaller, even so the time-consuming elements of this approach would be daunting given the pressure to deliver a visibility scheme to accommodate the burgeoning sector utilizing dynamic imagery. The video sequences to be used would need to be very carefully recorded and edited, and sub-sequences would no doubt have to be spliced together. Eye-tracking – preferably in a wide-screen format – would be required. In addition the results to be analysed would probably present the most demanding and protracted part of the process. The overall cost of a study would be high.

Chapter 2

Phase One study report

Health warning

This section begins with a health warning instead of ending with one or burying it as a footnote. The Phase One study was preparatory; it did some important groundwork for the dynamic imagery project. It put in place a very rough approximation of the research to follow, tackling technical problems that needed to be confronted, for example, by exploring experiment control options for managing the complex procedures involved in presenting and coordinating video inputs with respondents' actions. It also provided the opportunity to specify the potential scope of the later phases of the research, what variables should be manipulated, what needed to be controlled, what should be measured. The variety of poster panel types that was eventually included was sampled and represented schematically, as illustrated below. The task to be used was the direct search task as developed for Route's research on static imagery. This should alert the reader to the fact that eye-tracking was not involved at this stage, and the empirical yield of the study was at best a set of estimated hit rates. It is useful to bear in mind who the reader might be; more than likely someone with a professional interest in outdoor advertising, even someone who is routinely on the watch for poster panels in all their diversity. This will warn all readers against over-emphasizing their own experience, intuitions and expectations, now and later as this vein of research develops. The findings should not violate reasonable intuitions, but they will not be definitive; they may even be indicative, they should not be wildly out of kilter with established results; however, it should not be a requirement that they must make comfortable reading!

The study provided pointers as to possible performance levels and differences, but most importantly it enabled informed research design decisions to be made about the scale and likely demands of subsequent research.

Phase One study: Objectives

The objective of the dynamic imagery research was to provide acceptably precise estimates of the contributions to visibility hit rates of poster panels with dynamic/motion properties. The research progressed to this end via a set of steps; it began with the Phase One study, which was relatively light on resources, and developed some structural features of the final study, as well as providing data to decide the scale of data collection needed for later studies. This study was therefore a platform for subsequent research; as a preliminary investigation it bore an appropriate "resemblance" to the eventual endpoint of the dynamic imagery research, the Phase Three study using video clips.

The (partial) resemblance is seen first in the structure of the task used: this represents a driving-related activity; it required the execution of a central task (reacting to a brake light) in a cluttered environment; in addition there was potential attention capture by sundry non-driving-critical objects. There was also a secondary task, consisting of a response to the appearance on screen of certain objects that served as "targets" for attention. These targets embodied properties of dynamic images to be included in Phase Three. Finally, the task instructions induced the respondents to adopt an appropriate driving-related "mind set".

The study provided pointers as to possible performance levels and differences, but most importantly it enabled informed research design decisions to be made about the scale and likely demands of subsequent research.

Phase One study: Method

Equipment

The experiment was written with E-Prime 2.0 Professional. A Dell PC was used for the stimulus presentation and responses were logged with the use of a Psychology Software Tools serial response box and foot pedal.

Stimuli

This preliminary study using dynamic imagery was carried out using specially constructed stimuli (see Figure 1.1) containing "targets" embedded in a cluttered display of non-targets (Figure 1.2).

On a given trial, a pre-target stimulus was presented for either 8, 6 or 4 seconds prior to the onset of the target stimulus; this consisted of a small green circle presented in the centre of the display. Subjects were instructed to focus their attention on this. This was ensured by asking subjects to press a foot pedal whenever the centre circle changed to red (this change in colour was random and for a duration of one second). The centre circle never changed to red during the presentation of the target stimuli. Throughout both the pre-target stimulus phase and the target stimulus phase grey circles (in three shades of grey and two sizes) were presented in parts of the screen where the centre circle and then the target stimulus did not appear. Interspersed among these grey circles were blank spaces. The various types of grey circles and spaces were presented randomly. This background changed every 200 milliseconds (ms).

The target stimuli (Figure 1.1) were presented immediately after the pre-target stimuli. The target stimuli differed in terms of orientation (vertical vs. horizontal), colour (blue-and-black striped vs. yellow-and-black striped), movement (moving vs. static), direction of movement: up vs. down (for vertical panels); left vs. right (for horizontal panels), size (large vs. small), side of presentation (left vs. right), offset (nearer to the centre of the screen vs. further away from the centre of the screen). The horizontal panels had a horizontal striped pattern, and the lateral movement of these panels (both right and left) was designed to emulate the movement of a bus. The vertical stimuli had both a vertical and a horizontal striped pattern.

Figure 1.1: Targets used in the Phase One study

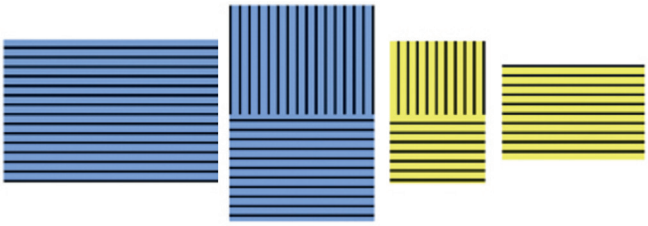

The movement of the vertical panels was designed to emulate the internal movement of scrolling panels, which was achieved by a vertical striped pattern moving (up or down) "over" the background of a horizontal striped pattern. In contrast, the static vertical panels were constructed to emulate scrolling panels in mid-scroll and comprised a vertical striped top half and a horizontal striped bottom half (see Figure 1.1). The height x width dimensions (in pixels) were 90 x 60 (for the large vertical panels), 60 x 40 (for the small vertical panels), 60 x 90 (for the large horizontal panels), and 40 x 60 (for the small horizontal panels). This ensured an equal area for the two types of large stimuli, and for the two types of small stimuli. The display size was 1020 x 768 pixels, and it was viewed at a distance of 48 cm from the participant's eyes.

The maximum period for the presentation of the target stimuli was three seconds, during which the moving vertical panels were designed to "scroll" twice and the moving horizontal panels to make one lateral movement (from left-to-right or from right-to-left). However, in most cases the presentation of the target stimuli was terminated prematurely by the target detection response of the subjects. The speed of movement (in pixels per second) was as follows: 60 (for the large vertical panels), 40 (for the small vertical panels), 30 (for both the large and the small horizontal panels).

Each pre-target and target phase together comprised one trial for the purpose of analysis. However, the presentation of a target stimulus immediately followed its pre-target stimulus, which was immediately followed by the presentation of the pre-target stimulus of the next trial and so on. This gave the impression of a continuous display. The trials were presented in a random order.

The appearance of the screen with one target item along with a background of non-targets is shown as Figure 1.2. Exceptionally this shows the central target as it would appear in the middle of the screen; this is for the purposes of illustration only since it would never normally accompany one of the peripheral target stimuli.

Figure 1.2: Example of display showing a target in the midst of randomly located non-target (background) items

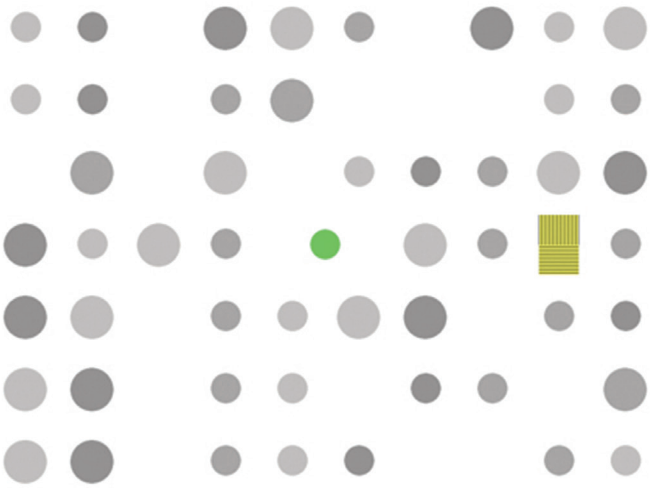

Task

The task was modelled on the direct search procedure used for an earlier study in this series; the participants were instructed to look for targets and report their presence as quickly as possible. Their reaction times were recorded along with a log of the actual response made as an accuracy check.

There were two trials per stimulus type, however, there were 16 trials for each of four factors included in the results analysis (type of panel, motion status, panel size, and offset position) – see below. There were 256 trials in total, divided into four blocks of 64 trials. Although all types were represented it was not feasible to run the study with a sufficient number of trials for every combination of conditions to be examined in the results analysis.

The targets were presented against a background of non-targets. This background changed every 200 ms to provide a suitable overall appearance of the search array for the task.

The construction of short video sequences for each of the stimuli was a painstaking and lengthy task, and the resulting code took up a considerable amount of space, and it imposed a substantial load on the PC's processing capacity.

Participants

Recruitment of individuals to serve as research participants was done using posters positioned in the vicinity of Birkbeck supplemented by the Birkbeck Psychology subject panel. The research laboratory at Birkbeck is in the central Bloomsbury area of London that contains much of the University of London campus. As this was a preliminary investigation the aim was not to cast the widest possible net as to demographic sampling or representativeness, however, it was apparent that the result was very similar to the usual sampling outcome. Participants were paid £10 for taking part in a single session of about 45 minutes.

There were 23 participants (16 females, 7 males). The mean age was 31 years. The mean age for the females (31.9 years) was slightly higher than that of the males (28.3 years). The subjects' occupations were varied but there were few manual workers (accountant; administrator, 2; book seller; civil servant; computer specialist; digital photographer; events organiser; executive officer; lab manager; management consultant; nanny; PhD student, Archaeology; PhD student, Biochemistry; PhD student, Chemistry; PhD student, Genetics; post-doctoral researcher; research assistant, Psychology; student, Business; temporary catering assistant; trainee embryologist; university administrator; waiter). Their nationalities were also varied with a majority from the UK or Europe (British, 10; East Timor; El Salvador; German; Indian, 2; Israeli; Italian, 3; Japanese; Polish; Spanish; USA).

Phase One study: Research design

Actual design

To explain the design of the study some further details of the targets will be useful. The type of panel was intended to allow the most salient dynamic properties of vertically scrolling panels and laterally moving bus panels to be portrayed, so the contrast was between vertical and horizontal rectangular shapes. The actual stimuli were not otherwise representations of these panel types. The movement of the scroller (when it moved) was internal so there was no change of position of the panel itself. In the case of the bus panel, the motion (when it moved) was horizontal and there was a noticeable position change. Moving and stationary versions of each panel type were presented.

In the following account, the label Type refers to the contrast between the two types of panel, those capable of vertical scrolling vs. those capable of lateral movement, irrespective of their motion status. It will be evident from Figure 1.1 that the difference between types is one of shape (portrait vs. landscape; e.g., 6-sheet vs. 48-sheet); it does not mean a difference due to motion status.

Both panel types were seen in moving and stationary states. A difference in motion status means just that (moving vs. stationary). If the effect of motion is greater for one panel type than the other, this corresponds to an interaction between Type and Motion.

The two other factors that were represented in the design and were assessed statistically were Size (of panel) and Offset (from screen centre). There were two sizes, the larger panels occupied an area 90 pixels x 60 pixels (=5400) and the smaller occupied an area of 60 pixels x 40 pixels (=2400); the former therefore being 2¼ times bigger than the latter.

The list of factors that needed to be controlled or manipulated was as follows:

• Type of panel/ stimulus motion	Vertical format (for internal scrolling) vs. horizontal format (for lateral position change)
• Panel Size	Large vs. small
• Motion status	Moving vs. stationary
• Offset position	Inner vs. outer
• Target colour	Yellow vs. blue
• Screen side	Left vs. right

These properties were fully balanced but when combined with every other variable the representation of any one factor was too small for the purposes of statistical analysis and so only the four underlined variables featured in the reported data analysis (Type, Size, Motion and Offset).

Potential design size

A design with complete factorial representation would entail 2 x 2 x 2 x 2 x 2 x 2 = 128 design "cells". In addition there are varieties of movement of special interest (which could themselves be elaborated by other directional variations), namely:

• Motion type - Lateral	Away vs. Towards Screen Centre
• Motion type - Scrolling	Upwards vs. Downwards

Together this constitutes 256 possible condition combinations. For an adequate assessment of performance on the task, something like 15-20 trials per task would be required. In the present instance 16 trials were used. For the complete design on this scale this would have resulted in 4096 trials which would take about 8 times longer than a typical visual experiment – perhaps 5 hours of testing. Generally this kind of demand would be handled by testing over a number of sessions or by testing different participants on subsets of conditions. It will be appreciated that for the preliminary nature of the present study, the complete design was not feasible. In the event the focus of the study, and the results analysis was on the first four factors listed above (and as underlined).

Phase One study: Results

Performance on a direct search task is generally evaluated by recording the reaction time (RT) to each target, typically using for analysis the mean or some other measure of central tendency for each target condition. The raw data may be used but in instances such as the present task where the RT distribution was likely to be substantially skewed, a transformation of RT may be employed. In a previous instance of using a direct search task (i.e., what was referred to as the Wave 3 Study), the RT distribution for each image was used to obtain a hit-rate like measure. This entailed the simple device of finding the proportion of RTs less than an arbitrary cut-off value. This method recommended itself here because it obviously produces results that are in the tradition of Route's eye-tracking studies, as was the case for the Wave 3 study. This is not to say they may be directly compared with conventional visibility hit rates, but the outcome of the study is more readily understood if this form of measurement is used. Accordingly the data were analysed in this form in addition to being analysed as mean RTs which would normally be the measure of choice. A brief summary of the results of the analyses is now provided, focussing on the principal statistical outcomes.

The first analysis was of the basic RT data. An analysis of variance was computed with 23 participants contributing mean RTs for 16 trials of the Type x Motion x Offset x Size design. Only two factors were significant at the 5% level: Type (F[1,22] = 6.342; p=0.020); and Offset (F[1,22] = 4.915; p=0.037). The first of these, Type, corresponded to the contrast between scrolling target and lateral moving targets, performance in response to the latter being slightly faster than to the former. The effect of Offset reflected quicker responses to targets when closer to the screen centre. Three other interactions were approaching significance and flagged themselves for further examination; namely Type x Motion, Motion x Size, and Type x Motion x Size; however, they were too far from the critical 5% significance level to merit closer scrutiny.

These results were suggestive rather than diagnostic and it seemed wiser to focus on the results that survived other treatments, notably the conversion from reaction times to hit rates since this is the measure that is practically most important. Deciding on a cut-off on the reaction time data is a somewhat arbitrary affair. If the cut-off is too low, all reaction times will be translated into misses, there are no hits and no discrimination between conditions is possible - whereas if the cut-off is too high, and all responses are counted as hits, there is again no scope for discrimination between conditions. For intermediate settings, there is the possibility that hit rates will differ between conditions, and have values somewhere between zero and one. In the case of Wave 3, settings were used that yielded hit rates in the region of those established by the stationary eye-tracking studies. There was no parallel in the present case, and the intermediate range had to be explored to find a value that optimised the discrimination between conditions. There was no directional bias in this other than seeking to find a position that produced the strongest indications of differences between conditions – there was no freedom in this to influence the direction of the differences. In the event, it was found that the profile of the results of the statistical analyses used was stable for a wide range of intermediate values, obviously avoiding the upper and lower limits.

Hit rates derived from the reaction times for each participant using a cut-off of 550 ms were subjected to the analysis of variance. The main effects that were significant on this occasion were again Type (F[1,22] = 4.717; p=0.041); and Offset (F[1,22] = 7.731; p=0.011). The other two main effects that were not significant were Motion (F[1,22] = 0.631; p=0.436); and Size (F[1,22] = 2.232; p=0.142). In addition there were two significant interactions: Motion x Offset (F[1,22] = 4.914; p=0.037); and Type x Size (F[1,22] = 7.879; p=0.010). It is of interest that Motion is not itself significant and Size is only bordering on significance. However, as noted both variables were implicated in interaction effects: there was a significant effect of Motion taken in conjunction with offset; similarly Size was significant when taken together with Type. The form of these interactions will be examined presently.

Evidence of the robustness of these findings was found in the results obtained by using other cut-offs. It would be expected that as the cut-off increases, the size of any significant effects will decline in the upper (or lower) reaches of the reaction time distribution, and that the probability values associated with the various terms in the analysis will wax and wane, and there will not be abrupt changes in the pattern of the findings. This was confirmed for the 600 and 650 ms cut-offs, the most durable effects being those of Offset and Motion x Offset. The mean profiles of all of the effects listed above were effectively the same, albeit at different overall hit rate levels. The overall average hit rates were 20.9%, 36.8% and 51.3% at the 550, 600 and 650 cut-offs respectively.

The mean hit rates for the variables and combinations of variables that achieved or approached statistical significance are of interest but their applicability in real world settings is quite limited; the findings need to be considered against the objectives of the study (see Introduction). For this purpose the means for the 500 ms cut-off were used, so modest hit rate levels were evident, and any hit rate differentials were small.

The effect of Offset reflected quicker responses to targets when closer to the screen centre.

Type indexes the difference between the panel formats used for Vertical Scrolling (19.9%) and Lateral Position Change (21.8%) and was significant; this reflected a very small advantage to the latter, but because this effect was obtained in the aggregation across motion status, and there was no sign of a Type x Motion interaction, it had to be attributed to a difference between vertical and horizontal formats. That, or the finding is spurious – a chance occurrence. In any event it should be emphasized that the statistical significance of Type does not mean that there was a difference between formats with scrolling and lateral motion properties.

Offset was also significant, a small but reliable advantage to the inner locations (22.2%) over the outer ones (19.5%). Size was not significant but the direction of this very marginal effect was in the expected direction (21.7% for the larger panels, against 19.8% for the smaller ones). The overall effect of Motion was not significant; the scores were 20.5% for moving panels and 21.3% for stationary ones. However, both Motion and Size were involved in significant interactions.

The mean hit rates for the two significant interactions are shown in Tables 1.1 and 1.2 below. The Type x Size interaction effect is probably spurious since it was only evident when a 500 ms cut-off was imposed. It does not simply mean that small scrolling panels were at a disadvantage relative to the other conditions tabulated, because the results were pooled across both moving and stationary conditions. It is necessary to look elsewhere for an explanation for the deficit, if indeed it was reliable. A factor that could come into play is aspect ratio; conceivably the disadvantage was associated with the small portrait format compared with the other conditions.

Table 1.1: Hit rates (%) as a function of target type and size

	Large	Small
Scroll	21.9	17.9
Lateral	21.5	22.1

The combination of motion and offset was involved in the other interaction that reached significance. In this case the data were pooled across panel types - vertically scrolling and laterally moving. It transpired that the overall effect of offset – favouring the inner position over the more peripheral one – was limited to stationary panels; by contrast there was no effect of offset for moving panels, as if movement of either type compensated for the disadvantage of being located further into the periphery.

Table 1.2: Hit rates (%) as a function of target motion and offset

	Outer	Inner
Moving	20.2	20.7
Stationary	18.8	23.8

The form of this interaction (for the 550 ms cut-off) is shown graphically in Figure 1.3, along with the same interaction as obtained at the 600 and 650 ms cut-offs. The profile was evidently resistant to the change in cut-off; while the height of the columns increased as the cut-off lengthened, the hit rate pattern between conditions was essentially undisturbed. Although this is not entirely surprising since the segmentation of the reaction time distributions is progressive and inevitably there is a dependency between successive cuts, it is reassuring that this degree of consistency occurred for this and other effects and interactions.

The mean hit rates for the variables and combinations of variables that achieved or approached statistical significance are of interest but their applicability in real world settings is quite limited.

Figure 1.3: Motion x Offset interaction at three cut-offs in the RT distribution

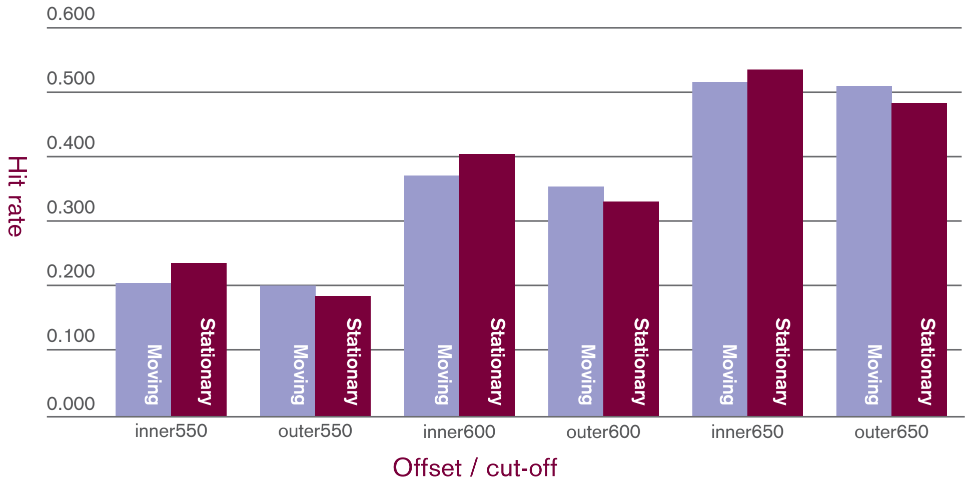

The remarkable thing about the transformation from reaction times to hit rates is that it resulted in analyses that appear more sensitive to the variables of interest than are found when mean reaction times are used. The difficulty is to know what cut-off to employ and why, if this methodological approach is to be applied.

Phase One study: Conclusions and learnings

The key lesson from these analyses is that the paradigm itself is sufficiently sensitive to the influence of variables that would, broadly speaking, be expected to have an effect on performance. The observed effects were sometimes masked by the fact that variables operated in concert or in conflict with one another. There were no instances where the outcome was bizarrely opposed to expectations. The chances are that the more subtle interaction effects will be moderated in the complex mix of real-world influences, so the Phase One Study is not likely to foreshadow later outcomes in any detail. The measure of its success at this stage is that the methodology did not militate against effects on hit rate; indeed it enabled certain effects to be detected and thus cleared the way for further development of the methods. The task as implemented was not beyond the participants' capabilities, they could see and identify moving (and stationary) targets in dynamic displays; it remains to explore whether this holds up with the use of video material that is more representative of real-world settings, as in Phases Two and Three.

The Phase One study was followed by two further investigations in which the methodology was still situated in the laboratory but was more closely geared to the visual experience of observers in real-world environments. In Phase Two, eye movements were monitored while the research participants viewed a series of video clips created from a fixed camera setting in a roadside location; the camera was fixed for the purpose while the surrounding environment was moving normally. The scenes depicted scrolling or stationary panels, and two panel sizes were represented. In Phase Three video clips were created using a Steadicam system to capture the experience of a pedestrian on the move. Another two sets of video clips were constructed using a dashboard-mounted camera so that the experience of car drivers and passengers could also be assessed. These studies were informed by the experience of the Phase One study but adopted a methodological approach that was indirect relative to the task to be performed (i.e., participants were not instructed to look for targets) and a measurement approach in the tradition of our stationary image eye-tracking studies (i.e., recording eye movements and deriving visibility hit rates).

Chapter 3

Phase Two study report

Phase Two study: Summary

The Phase Two study reverted to the use of eye movement recording for the purpose of data acquisition. Eye-tracking data were acquired from 34 participants, viewing video clips containing 6 and 48-sheets, half of which had a scrolling mechanism. Much was learned about the technology, the nature of the data, and the analyses thereof that map optimally on to past results and visibility models. It was stressed that the findings were extremely tentative, but it seemed that scrolling panels had a small advantage over non-scrollers. The usual finding that 48-sheet panels have higher hit rates than 6-sheet panels was replicated, suggesting as a bottom line that the methodology demonstrated sensitivity to a variable that it should. The treatment of the results is more important for what it illustrated by way of analytical options. Scores may be aggregated for a panel over the duration of the viewing interval but such comparisons need to be properly calibrated and conducted on an equivalent basis. Greater understanding of any effect a panel with movement may have is provided by an examination of the various scrolling phases (before/during/after). The most important issue for resolution was how to derive panel visibility from raw hit rate scores, given that these may, depending on how Phase 3 was to be implemented, be based on viewing intervals of variable lengths. Correcting for this variation is not straightforward. Even for a fixed interval (as in the previous static image eye-tracking studies) there remains a question of how to produce something akin to a standardized measure of visibility, and hence of "visibility adjusted contacts". The Phase Two study produced an agenda for Phase 3 full of interest and practical significance.

Phase Two study: Operational and design considerations

The objective of Phase Two was to complete the preparations for the Phase Three study. This process included work on the technical platform to be used, the video material to be used, and how this would be utilised in an eye-tracking study.

The panel types assessed in Phase Two were as follows:

Scrolling vs. static 6-sheets – Pedestrian

Scrolling vs. static 48-sheets – Pedestrian

Buses – Pedestrian

For Phase Two we compiled a sizeable amount of video material of scrolling panels from a pedestrian/roadside perspective. The combination of roadside settings and the viewpoint of the pedestrian were most convenient for the study because this was relatively accessible, and did not require permission for photography, etc. Buses served as an incidental element in the videos, which were recorded from a stationary camera; the technical requirements of the task for Phase Two were capable of being met without professional assistance. Some video recording was undertaken of digital panels to ensure that this material could be presented satisfactorily on the eye-tracker's screen, but this did not feature in the data collection part of Phase Two.

This package was appropriate for Phase Two since the study was principally geared to the development of the methodology. Nevertheless the study did provide preliminary indications of differences between scrolling and static panels. It also showed what needed to be done to obtain suitable footage of panels generally and of buses in particular, and enabled an agreed specification of what formats and circumstances should be represented in the video material for Phase Three. The final scale of the Phase Three study was decided on the basis of results from Phase Two with some foreshadowing of this by Phase One.

Phase Two study: Research design

The planned design for the study is shown in Table 2.1 below. The target presently was for 12 clips per "cell" in the matrix, slightly more than half the size planned for Phase Three, and this objective was met. A set of 12 decoy clips were recorded that contained no commercial content; the aim of this was to help avoid the participants becoming aware of the study's focus on poster panels. In the Phase One Study the equivalent sample size was 16 images per cell. Clip duration was limited to 6 seconds; much of the interesting action within a scene is inevitably protracted and some care was needed to ensure a natural look about the final composite video. It was provisionally estimated that 10-20 seconds would be the target duration for Phase Three.

Table 2.1: Design of video clip selection for Phase Two study

		Scrolling	Stationary
Pedestrian	6-sheet	12	12
	48-sheet	12	12
	Buses (opportunity sample)	-	-
	Decoy	-	12

A set of 12 decoy clips were recorded that contained no commercial content; the aim of this was to help avoid the participants becoming aware of the study's focus on poster panels.

Phase Two study: Methodology

Apparatus

Stimulus presentation and response recording were carried out using a DELL PC interfaced to an infrared EyeGaze eye-tracker (LC Technologies, Inc., 2003). Video stimuli were displayed with a screen resolution of 1024 x 768 pixels on a 19-inch Sony monitor with image onset synchronised with screen refresh at a refresh rate of 60Hz. The EyeGaze system measured direction of gaze using the "pupil centre corneal reflection" method without an attachment to the subject's head. The subject's eye, illuminated by an infrared LED at the centre of the camera lens, was monitored by a video camera mounted below the computer screen on which images were viewed; the screen was viewed at a distance of 49 cm from the participant's eyes. The centres of the corneal reflection from the LED and pupil were located by software and these data enabled the subject's gaze point to be established, by measuring the intersection between the optic axis and the screen. Fixations were defined by a gaze deviation within 25 pixels for a minimum of six samples (100 ms).

(This was basically very similar to our old system used to record eye data when viewing static images but could support video presentation and data capture, with a faster central processor and eye-tracker, as well as better control and analysis software.)

Material

A portfolio of video recordings was created and a selection of these recordings was made for editing into a series of six-second "video clips" featuring one scrolling or one stationary panel. The source material consisted of 121 recordings (containing 26 x 6-sheet non-scrollers, 25 x 6-sheet scrollers, 23 x 48-sheet non-scrollers, 20 x 48-sheet scrollers and 27 decoys with no panel in view). From this substantial array of options a final selection of 60 was made, resulting in 12 from each category. Selection was naturally rather more burdensome than for static images studies, with every step being more time-consuming. The process was speeded by using still images to represent the several clips. We attempted to collect images with digital displays but they were not available in sufficient numbers to be included at this stage of the research.

The usual stricture applied to the sourcing of video material, that it should represent typical rather than idealized views of the panels and their settings.

Figure 2.1: Snapshot from a sample video clip showing a roadside 6-sheet panel

The usual stricture applied to the sourcing of video material, that it should represent typical rather than idealized views of the panels and their settings. The intended focus of interest for the viewer was a roadside scene depicting a possible pedestrian route ahead or a setting in which to wait for a bus, for example.

It was evident from the preliminary video recording experience that it would prolong the acquisition of suitable material considerably if a balance of panels on the left and right were to be achieved. In the event, it was possible to have a partial balance with four on the right and eight on the left in all categories.

Participants

As for the Phase One study, recruitment of individuals to serve as research participants was done using posters positioned in the vicinity of Birkbeck supplemented by the Birkbeck Psychology subject panel. There were 33 successfully participants (22 females, 11 males). The mean age was 31 years. The mean age for the females (29.5 years) was slightly lower than that of the males (33.5 years). They were paid £10 for taking part in a session that typically lasted 45-60 minutes. They were from a wide range of occupations: administrator (7), PhD student, student (9), accountant (2), unemployed, teacher (5), swimming instructor, waiter, research scientist, lecturer, international trading manager, secretary, waitress, research assistant.

Procedure

The material was divided into two blocks of 20 clips, randomly composed with the constraint that there was equal representation of all categories in each block. Two such sequences were constructed to increase the random nature of the overall presentation of images.

Calibration of the equipment for each participant was done before eye-tracking commenced and the equipment setting was checked between blocks. Each block began with the presentation of a black central fixation cross on a uniform white background for 1.5 seconds. The same fixation stimulus appeared between each of the clips for the same amount of time. Participants took breaks between blocks and this generally lasted about a minute.

Instructions

The participants were given an informed consent form indicating the nature of the task and stated that they would be free to terminate the session and leave at any time. They were also orally instructed as follows:

All of the videos you will see were taken in everyday urban environments in the UK, and are representative of what someone could encounter on their day-to-day travels. Many are of road scenes, and some are in commercial settings. You should think of yourself as about to walk into the scene ahead, making up your mind whether or not to proceed in that direction, looking where you might be going, and doing so as naturally as possible in that setting.

The important thing to remember is that when each picture comes up it is important that you look for the first few moments at where you would be focussing if you were in such a scene.

You have just a few seconds to look at each scene so you will just have a few moments to first quickly decide where you're going, watching out where you walk. Then carry on looking at anything else that you might normally view when walking in such a scene. We'd like you to look at each scene as you normally would when making your way about town.

Phase Two study: Technical learnings

The scroller clips (6 seconds duration) were constructed so that the scroll occurred at three different points after the start of the clip: 0.5, 1.5 or 2.5 seconds. The eye-tracking data confirmed that these planned values were achieved. These data also indicated some variation in the scroll times as shown in Table 2.2. We assumed that scrolling times would be about 3 seconds, leaving a bit of time to assess any effect in the immediate post-scrolling period (as would occur if the scroll caught the participant's attention). Evidently we need to have rather longer clips and to expect quite a degree of scrolling time variability. Buses did come into view in a number of the clips, but this was very much an uncontrolled opportunity sample. The viewpoint for the bus shots was restricted to what might be expected for a pedestrian/ pavement perspective, namely of Superside and rear panels. The clips containing bus panels were not analysed for the purpose of getting hit rates for this panel subset.

Table 2.2: Durations per scroll phase per video clip

Panel	Timing	Clip number	Prescroll	Scroll	Postscroll
Scroller48	Half	MOV01A	0.52	5.48	0.00
Scroller48	Half	MOV01E	0.52	3.32	2.16
Scroller48	Half	MOV03F	0.68	5.32	0.00
Scroller48	Half	MOV041	0.56	3.44	2.00
Scroller48	OneandHalf	MOV02E	1.52	4.48	0.00
Scroller48	OneandHalf	MOV04F	1.52	3.28	1.20
Scroller48	OneandHalf	MOV061	1.52	3.28	1.20
Scroller48	OneandHalf	MOV062	1.56	3.36	1.08
Scroller48	TwoandHalf	MOV009	2.56	3.44	0.00
Scroller48	TwoandHalf	MOV00D	2.52	3.28	0.20
Scroller48	TwoandHalf	MOV016	2.52	3.48	0.00
Scroller48	TwoandHalf	MOV03D	2.56	3.44	0.00
Scroller6	Half	MOV00A	0.44	2.88	2.68
Scroller6	Half	MOV022	0.52	2.32	3.16
Scroller6	Half	MOV02C	0.52	1.88	3.60
Scroller6	Half	MOV039	0.48	2.28	3.24
Scroller6	OneandHalf	MOV005	1.48	2.24	2.28
Scroller6	OneandHalf	MOV006	1.44	2.32	2.24
Scroller6	OneandHalf	MOV014	1.52	2.56	1.92
Scroller6	OneandHalf	MOV052	1.52	2.52	1.96
Scroller6	TwoandHalf	MOV010	2.48	2.32	1.20
Scroller6	TwoandHalf	MOV012	2.52	3.12	0.36
Scroller6	TwoandHalf	MOV026	2.52	3.48	0.00
Scroller6	TwoandHalf	MOV031	2.56	2.24	1.20

Means for the above table of individual clip measures underline the picture that emerged from the raw scores (see Table 2.3). The totals are 6 seconds as planned. The pre-scroll durations in the last column (0.53, 1.51 and 2.53 seconds) are also close to the planned values (0.5, 1.5 and 2.5 seconds). The mean scrolling time for the twelve 48-sheet panels was 3.72 seconds and for the twelve 6-sheet panels 2.45 seconds. Post-scrolling times made up the balance of the 6 seconds per clip. In a number of cases the scroll was so slow that there was nothing left in the clip for scoring the post-scrolling effect.

Table 2.3: Durations per scroll phase as a function of panel size and scroll onset time

Timing of scroll onset	Mean duration	Panel Scroller 48	Scroller 6
0.5 sec	Pre-scroll	0.57	0.49
	Scroll	4.39	2.34
	Post-scroll	1.04	3.17
	Total duration	6.00	6.00
1.5 sec	Pre-scroll	1.53	1.49
	Scroll	3.60	2.41
	Post-scroll	0.87	2.10
	Total duration	6.00	6.00
2.5 sec	Pre-scroll	2.54	2.52
	Scroll	3.41	2.79
	Post-scroll	0.05	0.69
	Total duration	6.00	6.00
Overall averages	Pre-scroll	1.58	1.52
	Scroll	3.72	2.45
	Post-scroll	0.70	2.03
	Total duration	6.00	6.00

There were occasional difficulties with the video playback through the eye-tracking computer; notably some images were unstable and the scene jumped at unpredictable points. It is possible that this was due to the recording, and it was clear that a higher quality video camera was needed for the next phase of research.

Despite precautions in the composition of the instructions (and the inclusion of decoy clips), a number of participants commented that the study appeared to be about advertising. The clips in this study were from a stationary camera with movement of traffic and pedestrians ahead, in addition to any internal movement of the panel, no doubt emphasising the presence of a panel. This problem was mitigated by the use of a moving camera in Phase 3, but the instructions needed to be more carefully composed to avoid allowing the study objectives to be as transparent as this one's evidently were. No comments were made to suggest any participants identifying the contrast between stationary and moving panels as important.

A new software facility that was not needed for this study was to be able to specify an "area of interest" or AOI, defined for the first frame of the clip, and then recomputed automatically as the scene changed. This was not used here since the panel was at a fixed distance.

Phase Two study: Analyses of results

The data from the Phase 2 study are analysed in the following sections[3]. Extra caution is appropriate about the data first because they have been obtained using a new eye-tracking system, and second because the radical step was taken into an environment in which the imagery was dynamic. This "dynamism" relates to the manner in which the images were presented (video-based instead of still camera photos) and to the properties of some of the panel formats assessed.

This was a preliminary study and so a less rigorous approach to results analysis than normal was adopted to see if anything at all might be there of potential importance and also that might emerge from the "woodwork". The following should not be treated as a source of headline figures one way or another. It was essentially exploratory but may give a flavour of things to come, subject to various caveats that will be noted as the account proceeds. This applies equally to the findings from the Phase Three study.

There were 12 video clips with a scrolling panel and 12 with a non-scrolling panel; six in each case featured a 6-sheet panel, and six featured a 48-sheet panel. There were 34 participants in the research, each of whom viewed all clips, and whose scores were pooled for the purposes of analysis; the data for one of them was discarded as noted immediately below so the eventual participant sample size was 33. The duration of each clip was six seconds. The scrolling clips were divided into three scroll phases (pre-scroll, scroll, and post-scroll). They were composed so that the scroll was initiated after 0.5 seconds, 1.5 seconds or 2.5 seconds, and there were 8 clips for the 6-sheet condition for each panel onset time, and similarly for the 48-sheet condition. This provided a fully balanced design.

Fixation durations

Table 2.4 shows the frequencies of fixation durations pooled over 33 participants. These data are for all fixations above 80 milliseconds (ms), and are aggregated over all locations in each scene regardless of content. The data for one of the original 34 participants contained a handful of exceptionally long fixations, possibly for some unidentified technical reason, and his data were excluded.

It can be seen that the distribution was highly skewed; the modal interval was 100-199 ms, and the median duration was just above 200 ms. This is rather shorter than in many

[3] Hit rate scores are expressed in this Results section alternatively as a proportion (between 0 and 1) or as a percentage as it best seems to suit the account.

scene viewing studies but at the time of this study very few of them had used dynamic images. A closer analysis may have revealed whether a shorter minimum value than 80 ms might be warranted, but this is a nuance too far at this juncture. It should be noted that the majority of fixations (83.7%) lasted between one-tenth and one-half of a second (100 and 499 ms).

Table 2.4: Frequency distribution of fixation durations

Duration		Frequency (all Ss – 1)	
Lower limit	Upper limit	Frequency	Percentage
0.080	0.099	1809	3.98%
0.100	0.199	15058	33.11%
0.200	0.299	13102	28.81%
0.300	0.399	6503	14.30%
0.400	0.499	3397	7.47%
0.500	0.599	1765	3.88%
0.600	0.699	1035	2.28%
0.700	0.799	638	1.40%
0.800	0.899	403	0.89%
0.900	0.999	322	0.71%
1.000	1.099	251	0.55%
1.100	1.199	206	0.45%
1.200	1.299	175	0.38%
1.300	1.399	149	0.33%
1.400	1.499	124	0.27%
1.500	1.599	121	0.27%
1.600	1.699	119	0.26%
1.700	1.799	92	0.20%
1.800	1.899	49	0.11%
1.900	1.999	29	0.06%
2.000	2.999	90	0.20%
3.000	3.999	21	0.05%
4.000	4.999	8	0.02%
5.000	5.999	8	0.02%

Hit rates: Preliminary results

The raw data were software-packaged as large text files (929 MB in total), each too big to be accommodated by the then available version of Excel because its maximum number of rows was exceeded! The detail was even more extensive than hitherto with the static eye-tracking technique, and would be even greater if longer clips were used.

Analysis of aggregate hit rates

The first analysis of the results, summarized in Table 2.5, was on the aggregate hit rate scores for scrolling and non-scrolling panels (regardless of the phase of the scroll, which was taken into consideration in the second analysis). An analysis of variance was performed on these scores with panel size (6-sheet vs. 48-sheet) and panel type (scrolling vs. non-scrolling) as factors. Only the first of these factors was statistically significant and the interaction between the two was not significant.

Table 2.5: Analysis of hit rates as a function of panel size and panel type (scrolling vs. stationary)

Source	Sum of squares	df	Mean square	F	Sig.
Panel size (6-sheet vs. 48-sheet)	2686.5	1	2686.517	12.79	0.001
Panel type (scroll vs. non-scroll)	250.7	1	250.7102	1.19	0.280
Panel size x Panel type	58.3	1	58.30021	0.28	0.601
Error	9239.1	44	209.9803	-	-
Total	168471.8	48			

The mean hit rates for 6 and 48-sheets were 49.6% and 64.5% respectively. Scrollers scored an average of 59.3% and non-scrollers 54.8%; the difference of 4.6% was too small to be considered reliable.

The “interaction” term indicates whether scrolling is more (or less) effective with one panel size rather than the other. The means are in Table 2.6: there was a small numerical advantage to scrolling in both cases; 3.37% for 48-sheets, and 6.78% for 6-sheets. The fact that this was not significant should be emphasised – the result was, in effect, the way the cookie crumbled, and the entire pattern could change if we repeated the study. Of course the same applies to the other findings but it would be reasonable to anticipate a significant panel size effect in a repeat study. Although it is not the stuff of headlines, this particular finding should not be underestimated – it was the first sign of what may apply when dynamic images are used.

Table 2.6: Mean hit rate (HR%) as a function of panel size and panel type (scrolling vs. stationary)

	Scrollers	Non-scrollers
6-sheet	53.0	46.2
48-sheet	65.7	63.4

Analysis of hit rates by scrolling phase timing

The next analysis of the results, summarized in Table 2.7, was of the hit rate scores for scrolling and non-scrolling panels, also taking account of the timing of the scroll for the scrolling panels. An analysis of variance was performed on these scores with panel size (6-sheet vs. 48-sheet) and panel condition (scrolling 0.5 second onset vs. scrolling 1.5 second onset vs. scrolling 2.5 second onset vs. non-scrolling) as factors. Only the first of these factors was statistically significant, and the interaction between the two was not significant.

Table 2.7: Analysis of hit rates as a function of panel size and pre-scroll phase

Tests of Between-Subjects Effects

Dependent Variable: 6-second hit rate

Source	Sum of squares	df	Mean square	F	Sig.
Panel size	1844.27	1	1844.27	8.324	0.006
Pre-scroll duration	434.43	3	144.81	0.654	0.585
Panel size x pre-scroll duration	251.30	3	83.77	0.378	0.769
Error	8862.41	40	221.56	-	-
Total	168471.79	48			

The first three levels on the panel condition factor reflected the different balance of any scrolling effects. The 0.5-second condition had most of its contribution (5.5 seconds worth) to hit rate in the scroll and post-scroll phases. By contrast, the 1.5-second and 2.5-second conditions had less (4.5 and 3.5 seconds respectively) opportunity for hit rate to be enhanced by the scroll. Accordingly, if hit rate increases because of the scroll and its immediate aftermath, the scores will decline across the three conditions. There was no evidence of this as seen in Table 2.8.

Table 2.8: Mean hit rate (%) as a function of pre-scroll phase duration

Pre-scroll duration (s)	Mean HR%
Half	59.7
One and half	55.8
Two and half	62.5
Non scroll	54.8

This possible indicator of the effect of scrolling could differ as a function of panel size, but the absence of an interaction effect suggests that this was not the case. The numerical pattern for the scrollers was the same for 6- and 48-sheets (see Table 2.9).

Table 2.9: Mean hit rate (%) as a function of panel size and pre-scroll interval

Pre-scroll interval (s)	48-sheet	6-sheet
Half	64.0	55.4
One and half	66.2	45.4
Two and half	67.0	58.1
Non scroll	63.4	46.2

An intuitively appealing next step would be to calculate an average hit rate per second, by dividing each mean hit rate by the presentation duration, and scale up or down to get the score for some other time period. However, this would be a mistake since one could not then simply scale up or down arithmetically. For example, consider what happens when the score for a ten-second interval is obtained by simply extrapolating from the six-second averages. For the 6-sheet pre-scroll phase (see Tables 2.13 and 2.14, coming later), the per-second result was 0.160/1.50 = 0.106, and so for 10 seconds the hit rate would be 1.06. Indeed in every case in those two tables, the result exceeds 1, which is impossible – just as much beyond the possible as it is for a sportsman or woman to deliver on a promise to deliver a performance at 110%. Evidently this simple scaling operation is wrong.

Our solution was to develop the accumulated hit rate method applied to obtain visibility scores in the stationary eye-tracking studies, as described in the next lengthy section. After that we return to the analysis of Phase Two data.

Accumulating hit rates for dynamic imagery studies: Exploratory analysis

The main objective of the visibility research is to provide inputs needed in the modelling process. A key element in this exercise is the method of accumulating visibility hit rates over the opportunity-to-see (OTS) interval for a panel. The OTS episode may be decomposed into various elements (phase/duration, viewpoint, OTS distances, speed of travel, OTS behaviour, etc.) which vary between panels but the eventual common feature for all is HR accumulation. The following account establishes a framework for hit rate accumulation to apply to dynamic panels, and that accommodates the varying OTS values that characterize them and their component phases.

It was important for the groundwork of this task to be prepared because the dynamic imagery project introduced new complications as to how accumulation should be effected. This bears on what data are collected and provided as inputs to the modelling task. It was important also to align this part of the work with the inputs and outputs from the previous scheme operated by Postar Ltd., the predecessor of Route.

This contribution to the topic was to focus attention on the issues and in particular to demonstrate how accumulation could be done using the data from the Phase Two study. The fundamental problem was to develop a means for interpolating and extrapolating from recorded visibility hit rates to exposure intervals other than that used for the eye-tracking research, a problem that increased in complexity in the context of the dynamic imagery studies.

Table 2.10 illustrates various aspects of the HR accumulation process, and associated dilemmas, a basic process familiar to many Route users. The table uses values from the accumulation process for the non-scrolling 6-sheet and 48-sheet data from Phase Two; the "raw" hit rates were 0.462 and 0.634 respectively. Some details of the Table, including the labelling, are explained later.

Starting values for the accumulation process are required; to kick-start the discussion and illustrate the process, values estimated from the Phase Two data for one second's worth of viewing were used; for 6 and 48-sheets the respective starting-values are 0.077 and 0.106. The smart/suspicious reader may already have baulked at this as leading to numerical shenanigans – which it is, as was explained above – however, it is an instructive starting-point for seeing how to proceed.

Table 2.10: Accumulated hit rates: An example

	6-sheet		48-sheet	
Interval	Accumulated miss rate	Accumulated hit rate	Accumulated miss rate	Accumulated hit rate
N seconds	$(1-p)^N$	$1-(1-p)^N$	$(1-p)^N$	$1-(1-p)^N$
1	0.923	0.077	0.894	0.106
2	0.852	0.148	0.800	0.200
3	0.786	0.214	0.715	0.285
4	0.726	0.274	0.640	0.360
5	0.670	0.330	0.572	0.428
6	**0.618**	**0.382**	**0.512**	**0.488**
7	0.571	0.429	0.458	0.542
8	0.527	0.473	0.409	0.591
9	0.486	0.514	0.366	0.634
10	0.449	0.551	0.327	0.673

Consider first the columns of data for the sample of 6-sheet panels in Table 2.10. The probability of a hit in a 1-second interval for this 6-sheet sample is taken as 0.077. Therefore the probability of NO hit after one second is 1 – 0.077 = 0.923. The value of 0.148 for 6-sheets after 2 seconds is derived as follows: the probability of no hit in a 1-second interval is 0.923 and this applies to both 1-second intervals in turn, so the probability of no hits in the first two seconds is 0.923 x 0.923 (= 0.852). It follows that the probability of a hit in the first two seconds = 1 – 0.852 = 0.148. The next value of 0.306 for 6-sheets after 3 seconds is 1 – 0.923 x 0.923 x 0.923 = 1 – 0.786 = 0.214. And so forth. The second column for a 48-sheet sample proceeds similarly but from a starting value of 0.106.

Note that the accumulated hit rates after 6 seconds are 0.382 and 0.488, which are lower by some margin than the recorded values of 0.462 and 0.634. A closer examination of the problem is needed. This was done by developing a more general version of the process.

It was important for the groundwork of this task to be prepared because the dynamic imagery project introduced new complications as to how accumulation should be effected.

Accumulating hit rates over an OTS period: Generalizing the account

The aim is to obtain the hit rate for a target exposed for T seconds, accumulated (i.e., aggregated) over an entire OTS period.

We begin by dividing the OTS period T seconds into intervals of equal duration t seconds. We assume that the probability of fixating the target in any sub-interval is a constant, p, and we estimate p from the eye-tracking hit rate data as follows.

The problem – of deriving an expression for the *accumulated hit rate* - is approached from the complementary perspective, by focussing on the *accumulated* *miss* *rate*.

If we know the probability that the target has been missed at the end of a given interval (i.e., there have been no hits) then we can simply subtract that number from 1 to get the probability of a hit having occurred. This is done as follows:

1. The probability of a hit by the end of interval 1 is p. So the probability of zero or no hits by the end of interval 1 is 1 - p.
2. The probability of no hits by the end of interval 2 is the probability of no hits in interval 1 *and* no hits in interval 2; that is, $(1 – p) * (1 – p)$ or $(1 – p)^2$.
3. By similar reasoning the probability of no hits by the end of interval 3 is $(1 – p) * (1 – p) * (1 – p)$ or $(1 – p)^3$.

   In general, the probability of no hits by the end of the Nth interval is $(1 – p)^N$.
4. It follows that the probability of at least one hit by the end of the Nth interval (the *complement* of this value) is $1 - (1 – p)^N$.

This is a transparent computationally convenient formula that is sufficiently general for present purposes. Figure 2.2 illustrates how hit rate grows under this simple accumulation process for two hit rate values estimated from the data. Provisionally the estimate was done by taking a raw hit rate value from the eye-tracking data for a 6 second exposure (the raw hit rate, HR6), and dividing it by 600 to give a probability for a 1/100th second exposure. The two examples from the Table are for raw HR6 = 0.462 and 0.634, respectively corresponding to base HR values of 0.0467 and 0.106.

Figure 2.2: Example of accumulated hit rates for two base rates

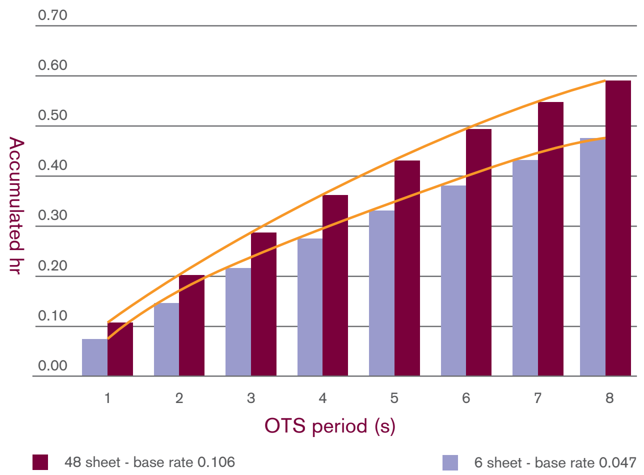

The growth of hit rate is clearly shown in Figure 2.2. This is consistent with the findings of the Duration Study that showed empirically how hit rates "grow" with duration. What is not so obvious is the undershoot in both cases relative to the raw HR6 values. The accumulated hit rate at 6 seconds is less than the raw hit rate values on which the starting values for the calculations were based: resulting in 0.382 instead of 0.462 for the 6-sheet example, and 0.488 instead of 0.634 for the 48-sheet example. This results from a flawed assumption about the base rates used as starting values, which was addressed and corrected in the next section.

What should be the starting point for accumulation? Estimating glimpse probabilities

What we need to establish is the correct base hit rate, the value we assume to be constant over any interval into which the OTS period is divided, and which will eventuate in the empirically estimated HR6 value after the accumulation process is applied. It will be termed the "glimpse probability" to signify the likelihood that the panel will be looked at in a brief interval.

The undershoot problem was caused by adopting an inappropriate starting value, which was obtained by dividing the HR6 score first by 6 to get a per-second score, then by a further 100 to get a per-hundredth-second score. (The size of the resulting interval, say 1 second or 1/10th or 1/100th of a second, is for us to choose. It matters partly because intermediate values may well be required.)

Recall, however, that we have earlier stressed that per-second hit rates cannot simply be multiplied by N to produce hit rates for intervals beyond N seconds because hit rates may exceed 100%. This might caution us to be circumspect about dividing hit rates too.

Table 2.10 showed cumulative values for successive one-second intervals including 6 seconds – which is the exposure duration used for the static imagery eye-tracking studies, and also used for the clip lengths in the Phase Two study. This exposure duration was a way of placing a limit on the OTS period. But the figures are based on the use of an inappropriate starting value.

How better to proceed? What starting-points would more appropriate? An alternative would be to search by trial and error for values that actually culminate in the observed hit rates from the eye-tracking studies (HR6). For example, if HR6 = 0.462, the starting value could be found just by trial and error, and it would turn out to be 0.0982 (somewhat higher than the 0.077 which we get by dividing 0.462 by 600). So how else than by trial and error could this value come from? We need an analytical solution to this part of the problem.

Estimating glimpse probability from HR6

To estimate glimpse probability p, we start with the recorded hit rate HR6 for the panel, and then find p by reversing the accumulation process. If we have the correct glimpse probability value p, the accumulation process in its normal direction produces exactly this value as the accumulated hit rate AHR6. The process of reversing the accumulation process to derive p is as follows:

1. We showed that the hit rate after 6 seconds is $HR6 = 1 - (1 - p)^6$.
2. This can be rearranged to give $(1 - p)^6 = 1 - HR6$.
3. Hence $1 - p = (1 - HR6)^{1/6}$.
4. Finally $p = 1 - (1 - HR6)^{1/6}$ where p is the starting value from which the value of HR6 can be derived by accumulation[4].

We can now apply this simple process to find the appropriate values of p for the Phase Two scores.

The respective glimpse probability values for 6-sheets and 48-sheets were 0.098 and 0.154, and Figure 2.2 was redrawn with these values.

Figure 2.3: Accumulated hit rates for the base rates for the 6- and 48-sheet non-scrolling panels in the Phase Two study

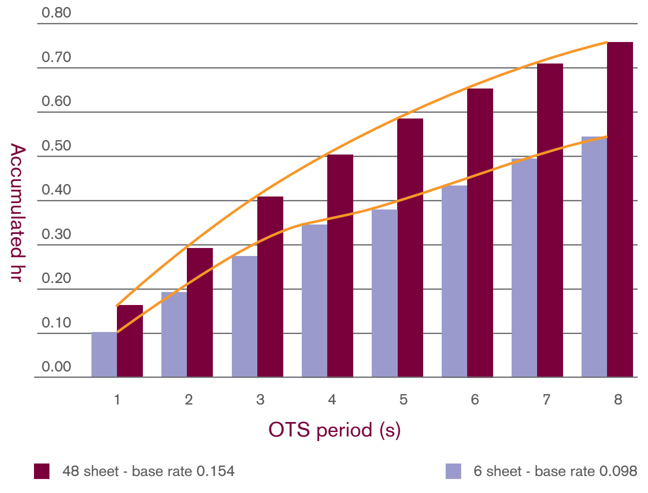

Further analysis of Phase Two data

Having dealt with and learned from the example of the non-scrolling panels, we are now in a better position to apply the accumulation process to the rest of the Phase Two data, including its novel aspects.

The scrolling panel presents a more complicated case, first because there are three phases with potentially different hit rates; and second, because (unexpectedly) there were found to be different durations per phase. Although one post-scroll phase is in practice the pre-scroll phase for the next scroll, they are distinguished here because they were visually separate in the video clips. Notwithstanding the various hit rates for sub-categories were available and a summary of hit rate scores per phase and panel size in Table 2.11, with the corresponding durations in Table 2.12.

Table 2.11: Observed hit rates per condition: Phase Two study

	Pre-scroll HR	Scroll HR	Post-scroll HR
6-sheet	0.160	0.321	0.221
48-sheet	0.258	0.525	0.200

Table 2.12: Durations per condition: Phase Two study

	Pre-scroll	Scroll	Post-scroll
6-sheet	1.50	2.51	1.99
48-sheet	1.55	3.80	0.65

What may seem to be an intuitively sensible approach is to divide each observed hit rate by the corresponding OTS duration. Unfortunately (as was exemplified by the undershoot problem) difficulties arise for hit rate calculations if simple averaging, multiplying or dividing operations are applied. An extra though minor hurdle is posed by estimating glimpse probabilities from observed hit rates for intervals of varying duration. In the previous section we developed a more suitable basis for obtaining accumulated hit rates and this can be applied to the problem of scrollers and their components. To this end we used the information from Tables 2.11 and 2.12, taking each pair of cells in turn, beginning with the observed pre-scroll hit rate for 6-sheets of 0.160 for an OTS duration of 1.50 seconds.

[4] Note: the power 1/6 means the 6th root of the quantity in brackets. The Excel function 1-POWER(1-HR6,1/6) will calculate this quantity, where HR6 is the nominal hit rate for a 6 second exposure. The function can of course be generalized.

Table 2.13: Glimpse probabilities (per one-second exposure) for the Phase Two study

	Pre-scroll	Scroll	Post-scroll	Non-scrollers
6-sheet	0.110	0.143	0.118	0.098
48-sheet	0.175	0.178	0.291	0.154

Table 2.13 arguably gives a “fair” basis for comparing conditions because the contribution of duration is removed. A more familiar way to make the comparison would be for the scores to be rated as if for a 6 second exposure, the presentation duration used in the stationary eye-tracking studies. It should be clear that an exposure duration this long is not feasible in many instances (e.g., the scrolling phase for scrollers). Notwithstanding, this way of “normalizing” the results has the advantage of putting matters on an even footing. This was therefore done in Table 2.14 which shows the results of accumulating hit rates based on the glimpse probabilities for Phase Two scroll phases (and non-scrollers) for a 6 second interval. The leading cell of 0.502 is the accumulated hit rate after 6 seconds for a glimpse probability of 0.110.

For the next step in the analysis, we refer again to the formula $p = 1 - (1 – HR6)^{1/6}$ to estimate p, the starting value from which the value of HR6 was derived by accumulation.
A general version of this expression may be used for other intervals of N seconds duration (N may be a non-integer):
$p = 1 - (1 – HRN)^{1/N}$

Table 2.14: Standardized hit rates (per 6 second exposure) for the Phase Two study

	Pre-scroll	Scroll	Post-scroll	Non-scrollers
6-sheet	0.502	0.604	0.529	0.462
48-sheet	0.685	0.691	0.873	0.634

This is a convenient stopover point, at which the reader may pause to consider the relative contributions of the various phases. No statistical analysis was attempted for these results because this is very much a preliminary study, and it is not wise to treat the figures as more than extremely tentative findings. Nevertheless the signs are that hit rates for scrollers are higher than for non-scrolling panels. More importantly for future application, the analysis signals a way in which the various contributions and conditions may be assessed.

Future concerns will include the increasing presence of digital panels which do not have to have the equivalent of a scrolling phase although it seems likely that there may be some merit in inserting a transitional element between advertisements in a sequence, and this no doubt will be programmable as to content and duration.

Aggregating over components

Finally for completeness we derived an aggregate score for scrollers from the phase component scores. This did not entail further recourse to the accumulation calculations since the probabilities were statistically combined over the actual component scores and durations. The components were presented in Table 2.11 and the key results were presented in Table 2.15. The composite probability score was obtained by combining component probabilities in the usual way, treating them as independent. For a 6-sheet, the component probabilities were 0.160, 0.321 and 0.221 (see Table 2.11). The probability of a hit by the end of the second phase was $0.160 + (1 - 0.160) \times 0.321 = 0.430$. The probability of a hit by the end of the third phase was therefore $0.430 + (1 – 0.430) \times 0.221 = 0.556$. In the same way the probability of a hit by the end of the third phase for a 48-sheet was 0.718.

Table 2.15: Hit rates (normalized to a 6 second exposure) as a function of panel size and scroller type for the Phase Two study

	Composite HR for scrollers	Non-scrollers
6-sheet	0.556	0.462
48-sheet	0.718	0.634

The composite hit rates for 6-sheet and 48-sheet scrollers were respectively about 20% and 13% higher than the scores for the corresponding non-scrollers. The other way of viewing the table was also potentially of interest: while 48-sheet scrollers scored 29% higher than 6-sheets, this advantage was 37% for traditional non-scrollers.

It is important not to see any of these figures as headline results. Of course they have to be interpreted in the context of how the study was done, and the actual hit rates were for intervals that varied substantially from condition to condition. The advantage to scrollers appears to be small but systematic. It remains to see whether it was also demonstrated in Phase Three.

When does accumulated hit rate reach a specified level (say, 99%)?

Finally, the opportunity arose as the accumulated hit rate analysis was being developed to enjoy a bonus result, the answer to another enduring question.

Defining the problem

The examples in the above account point to the possibility of extrapolation beyond the 6-second exposure interval used in the eye-tracking research. This is a question that was regularly raised in the context of poster panels that are in view while the observer waits in their vicinity, as on tube platforms or in tube trains. It is in any case of interest to determine when the accumulation hit rate process delivers a given hit rate value. For example, when the reported HR6 value is 0.60, when does the accumulated hit rate reach 99%? Similarly when HR6 is 0.20? Or when the target level is 75%? These questions can be tackled by reversing the general expression for accumulated hit rate.

The general solution

The formula for accumulated hit rate is used for this purpose, rearranged to express N in terms of p and AHRN.

1. The general version of the formula is AHRN = 1 - $(1 – p)^N$
2. This can be rearranged to give 1 – AHRN = $(1 – p)^N$.
3. Taking logs of both sides of this equation we get log (1 – AHRN) = log $(1 – p)^N$

or log (1 – AHRN) = N log (1 – p), which rearranged gives[5]:

N = log (1 – AHRN)/ log (1 – p)

Some specifics

Starting with a value for HR6 of 0.60, when does AHRN reach 0.99?

We first have to get the value of p corresponding to the given value of HR6. This by the accumulated hit rate formula is p = 1 - $(1 – \text{HR6})^{1/6}$ which is 0.1416.

Next the value of p and the value of AHRN are substituted in the formula

N = log (1 – AHRN) / log (1 – p) = log 0.01 / log 0.8584 = 30.16.

In other words, for a panel with a raw hit rate of 0.60, the accumulated hit rate will reach 0.99 just after 30 seconds.

Table 2.16 reports further examples for a range of percentage hit rates (HR6) and two target levels (99% and 75%).

[5] The Excel function log (1-AHRN)/log(1-p) will calculate N for given values of AHRN and p. The result is the same regardless of the base of the logarithms, so long as they are the same!

Table 2.16: Times (in seconds) at which accumulated hit rate reaches a target level: Examples (in seconds) for various raw hit rates and for 99% and 75% target levels

HR6	0.90	0.80	0.70	0.60	0.50	0.40	0.30	0.20	0.10
Target HR=0.99	12	17	23	30	40	54	77	124	262
Target HR=0.75	4	5	7	9	12	16	23	37	79

The table illustrates the fact that as hit rate (HR6) declines, the time taken to reach a given target level increases. The two rates of increase can be seen to be non-linear. It seems apparent too that such curves will be only asymptotic to 100%. To take one example, if the observed hit rate is 10% for a 6-second presentation, the accumulated hit rate will reach 99% after just over four minutes of viewing (262 seconds).

Phase Two study: Discussion and recommendations

The step from an abstract representation of the dynamic imagery problem was taken in the Phase Two study by using video clips of scrolling and non-scrolling 6-sheet and 48-sheet panels. Each clip was edited to a 6-second duration, but in the instance of the scrolling panels, the editing enabled the movement to be seen after different starting points in the viewing interval. Performance was scored for the entire viewing interval, and in the case of scrollers for three movement phases (pre-scroll, scroll and post-scroll). Performance measurement for the scrollers presented a new analytical challenge because their varying durations had to be taken into account, notably in the accumulated hit rate process.

Statistical analyses were used tentatively and sparingly to explore the contribution of the various design factors. In the initial overall hit rate analysis, panel size was statistically significant as it routinely has been found to be, but movement (scrolling) was not. The advantage to scrolling panels of the same size was found to be pitched at roughly 33%.

Our suggested solution to the variable duration problem was to use the probability accumulation model as used in the visibility modelling schemes developed for Route. This enabled generalization to other intervals and provided a more reasonable basis for comparisons between panel formats. The method enables the scaling back to an estimate of glimpse probability and scaling up again to intervals of interest. This was facilitated in the first place by the use of fixed durations for the video clips. Inevitably the overall exposure times of the video clips in Phase 3 will not be capable of being edited to a fixed duration. However, it seems likely that the accumulation process would be needed to achieve comparability, by scaling the results for different conditions to a common duration. There is an important qualification to this because target panels in Phase Three will by changing distance and hence retinal position and eccentricity will not be fixed as in Phase Two. This will present a formidable analytical challenge for a future visibility modelling exercise. Nonetheless the core of the method and the underlying philosophy should not be modified or discarded except on clear rational grounds.

It was not clear that the scrolling phase per se contributed to hit rate. Of course this is the phase when one image is replaced by another and the screen is in motion. It may be the scrolling action – the actual motion - that captures attention. Alternatively, it may be the change from a stationary to a moving state, or vice versa, that attract attention – always assuming that something about the system is to some degree attentionally compelling. It was not clear from Phase Two that there were any differences between the scrolling phases. The issue could be revisited in future with the use of digital panels because they enable the duration of stationary and transitional phases to be manipulated, and for the latter to be eliminated if desired.

It may be the scrolling action – the actual motion - that captures attention. Alternatively, it may be the change from a stationary to a moving state, or vice versa, that attract attention – always assuming that something about the system is to some degree attentionally compelling.

Chapter 4

Phase Three study report

Phase Three study: Summary

The Phase Three investigation represented another substantial advance by enabling our visibility research to accommodate more facets of dynamic imagery related to poster panels. Video recordings were made of roadside scenes depicting the viewpoint of car occupants. Other recordings were made from a pedestrian perspective of roadside and station concourse scenes. These were edited to form a collection of video clips each of about 20 seconds duration and containing an advertising poster panel. Some roadside shots included views of mobile advertising surfaces (buses or taxis), and some additional clips were created with no advertising in view to serve as decoy material. The panels were of 6-sheet or 48-sheet size, and were fitted with a scrolling mechanism or were traditional stationary panels. A set of digital panels of both sizes were also presented via video clips. While the research participants viewed the clips their eye movements were recorded. Hit rates per panel (measured as the proportion/percentage of participants who fixated it) were compiled from the eye-tracking data. Accumulated hit rate measures were derived taking into account the opportunity-to-see durations of the panels. Interest lay chiefly in any performance advantage attributable to the different forms of movement. To assess the relative performances of the various panels a measure called the dynamic frame factor (DFF) was defined. Two versions of this factor were devised. Modest increments due to movement were noted, of the same size as obtained in the Phases One and Two studies. It was stressed nevertheless that these three investigations are on a modest scale, panel technology was changing (particularly digital), and the video material was not optimal. Nonetheless the research as a whole delivered important lessons, interesting insights and it points to possible directions for further investigation.

Phase Three study: Introduction

The next step to be added in Phase Three of this series of studies was to require the observer's viewpoint to move in the scene as an observer would, moving forward as if on foot or in a vehicle. Specifically therefore the perspectives (and mind-sets) of a car driver and car passenger were investigated in addition to that of the pedestrian. This presented a major challenge in terms of image acquisition as one set of video sequences would have to be acquired from inside a moving vehicle, another by using a camera to capture the movement of a pedestrian.

Another objective of the Phase Three study was to expand the range of panel types, specifically to include the digital panels that were appearing in greater numbers in the UK, and to attempt a wider coverage of bus and taxi advertising. This new study involved major changes in the nature of the material to be investigated and the associated image capture methods. In addition it was conceived, like Phases One and Two as an exploratory study, from which lessons could be learned as to the conduct of a normative study on the scale and scope of our static image studies.

Phase Three study: Methodology

This was broadly the same as for Phase Two. The main differences lay in the composition and presentation of the video sequences. The research design was expanded to include digital material.

Operational issues

Equipment was hired for the purpose of making videos: a Steadicam and camera for the walking shots, and mounting brackets for the car dashboard for the in-car shots. A bump-free video from a car was considered too costly as this was a preliminary investigation, it would have involved use of a special vehicle. In the event a satisfactory solution was to use a camera mounted inside the car. A shortcoming of this solution was that the video would be from a passenger perspective, but this was obviated by centring the camera position on the dashboard as much as possible. A camera/lens system was hired to ensure a good fit to the Steadicam and car mounting brackets. The preparation for capturing the videos included a review of routes to identify obstacles, verifying that views of panels were not occluded, and retakes when panels were inadvertently obscured by other vehicles/pedestrians. This was a lengthy process. The video capture process was piloted and the resulting material reviewed. The Steadicam shots were marked by a noticeable degree of shakiness before the main video recording stage, but it persisted to a still marked extent into the video clips used for the pedestrian condition. This would need to be improved in a future study, probably necessitating more expensive kit and possibly a professional operator. Apart from the shakiness, the pedestrian shots were generally satisfactory.

Apparatus

The same apparatus and recording equipment were used for presenting stimulus material and for acquiring the eye-tracking data as for the Phase Two study.

Material

A preliminary audit of a set of possible locations was conducted and potential routes were identified for video recording. These included various non-roadside venues in London: Liverpool Street Station, Victoria Station, and Westfield Shopping Centre. The yield of panels fitting the research requirements was sometimes problematic, as shown by the results of one typical foray for pedestrian material which located 37 sites with 70 6-sheet panels but a mere four sites with 48-sheets. The station concourse setting was eventually chosen as the most suitable and offering the wider range of required formats.

A portfolio of video recordings was created and a selection of these recordings was made for editing into a series of video clips of variable durations but averaging around 20 seconds. They featured a variety of scrolling and non-scrolling 6-sheet and 48-sheet panels, along with some digital versions of each. The collection of panels presented is summarized in Table 3.1, but some additional details follow:

Driver/Passenger: 8 x 48-sheet stationary, 8 x 48-sheet scrollers, 5 x 48-sheet digital (non-animated), 8 x 6-sheet stationary, 8 x 6-sheet scrollers, 8 x buses (1 Superside, 7 T-sides), 8 x taxis (4 supersides, 4 wraps).

Pedestrian: 8 x 48-sheet stationary, 8 x 48-sheet scrollers, 8 x 48-sheet digital (2 non-animated, 6 with varying degrees of animation/flashing), 8 x 6-sheet stationary, 8 x 6-sheet scrollers, 8 x buses (3 Supersides, 5 T-sides), 8 x taxis (all supersides).

Table 3.1: Number of panels contained in Phase Three video clips

		Mobile / decoy	Digital	Stationary	Scrolling
Driver / passenger	6-sheet	-	Not available	8	8
	48-sheet	-	5	8	8
	Buses (opportunity sample)	8	-	-	-
	Taxis (opportunity sample)	8	-	-	-
	Decoy	8	-	-	-
Pedestrian	6-sheet	-	8	8	8
	48-sheet	-	8	8	8
	Buses (opportunity sample)	8	-	-	-
	Taxis (opportunity sample)	8	-	-	-
	Decoy	8	-	-	-

As usual considerable care was taken to represent typical rather than idealized views of the panels and their settings. The intended focus of interest for the viewer was either a scene depicting a possible pedestrian route ahead or one showing a street view from a car. The tasks for the participants were conveyed by oral instructions as set out below.

Participants

As for the previous studies, recruitment of individuals to serve as research participants was done using posters positioned in the vicinity of Birkbeck supplemented by the Birkbeck Psychology subject panel. They were paid £10 for taking part in a session that with instructions, calibration, testing and debriefing, lasted about 45 minutes. For the Driver condition 24 subjects (16 females, 8 males) were tested; average age 27.5; 13 of the 24 were of European origin. Following the session it was found that 18 were still naïve as to the purpose of the experiment, but six thought it was about advertising.
For the Passenger condition 15 subjects (13 females, 2 males) were tested; average age 29.4: 10/15 of European origin. Debriefing revealed that 9 remained naïve, six were no longer naïve. For the Pedestrian condition 37 subjects (27 females, 10 males) were tested; average age 28.9; 21/37 of European origin. Debriefing revealed that 25 remained naïve, 12 were no longer naïve. The participants were from a wide range of occupations: academic adviser, administrator (5), advertising campaign designer, editor, editorial assistant, executive officer, human resources, lawyer, librarian, PhD student (3), research assistant/technician (4), researcher (3), secretary, student (6), supervisor, tailor, teacher (4), translator, waiter (2), writer.

Procedure

The panels appeared alone or with one or more other panels in the video clips. The duration of the clips depended on the number of panels present but varied around 20 seconds. The clips for the two conditions (Driver/Passenger and Pedestrian) were presented in two blocks. There were two practice blocks at the beginning of a session to ensure the instructions had been understood. Each block had eight decoy clips. Block 1 of the Driver/Passenger condition had 20 clips with panels; Block 2 had 22 clips with panels. Block 1 of the Pedestrian condition had 28 clips with panels; Block 2 had 26 clips with panels. The order of presentation of the conditions was counterbalanced and the clips were presented in a random order within the blocks. Each block lasted between 10 and 13 minutes.

Calibration of the equipment for each participant was done before eye-tracking commenced and the equipment setting was checked between blocks. Each block began with the presentation of a black central fixation cross on a uniform white background for 1.5 seconds. The same fixation stimulus appeared between each of the clips for the same amount of time. Participants took breaks between blocks and this generally lasted about a minute.

Instructions

The participants were given an informed consent form indicating the nature of the task and stated that they would be free to terminate the session and leave at any time.

The participant was then told whether he/she should view the videos as “driver” or “passenger” before being orally instructed as follows:

We are interested in the various factors that have an effect on people’s vision in everyday scenes, such as brightness/darkness, the relative speed of an observer, and the shakiness of the video.

All of the videos you will see were taken in everyday environments in the UK, and are representative of what someone could encounter on their day-to-day travels. They are generally urban scenes and you will need to think of yourself travelling in a car, driving or being driven, and looking at the scene ahead as naturally as you would in a car.

Your eye movements will be recorded using a small video camera that tracks the reflections from your eye. It is a safe and standard procedure.

Before any of the scenes are presented, we have to make sure that the camera is in focus on your eye. We then go through a short calibration sequence, so that the computer can interpret the signals that the tracker picks up from your eyes. The nature of these movements is unique to each individual.

After the calibration is complete, you will be shown a few ‘practice’ videos so you understand the nature of what we are doing.

The important thing to remember is that when each video comes up it is important that you look for the first few moments at where you would be focussing if you were in such a scene. You have just a few seconds to look at each scene so you will just have a few moments to first quickly think about where you’re going, watching out for any hazards. Then carry on looking at anything else that you might normally view when going along in a car in such a scene. We’d like you to look at each scene as you normally would when going in a car on your way about town.

In the experiment itself, while your eye movements are recorded, you will be shown a series of videos. Each video is shown for about 20 seconds. There will be two groups in all, and the experiment will take about 20 minutes.

When you are ready, the researcher will give you detailed instructions on what to do.

Participants serving in the “pedestrian” condition received the oral instructions that were the same as those for the driver/passenger condition with the exception of the second and sixth paragraphs:

All of the videos you will see were taken in everyday urban environments in the UK, and are representative of what someone could encounter on their day-to-day travels. You should think of yourself as walking into the scene ahead, looking where you might be going, and doing so as naturally as possible in that setting.

The important thing to remember is that when each video comes up it is important that you look for the first few moments at where you would be focussing if you were in such a scene. You have just a few seconds to look at each scene so you will just have a few moments to first quickly look where you're going, watching out where you walk and for any possible hazards. Then carry on looking at anything else that you might normally view when walking in such a scene. We'd like you to look at each scene as you normally would when making your way about town.

Phase Three study: Data and analyses

Caution continued to be exercised about the data, first because they were obtained using a new eye-tracking system, and second because the radical step was taken into an environment in which the imagery was dynamic. This "dynamism" related to the manner in which the images were presented, and to the properties of some of the panel formats assessed. In addition the viewpoint in the video image was itself on the move, completing the range of dynamic aspects of the method and material.

Expectations of the data based on the use and presentation of moving images with static and dynamic panel formats could be based: (a) most simply on extrapolation from the stationary eye-tracking precursors; supplemented by (b) an assumption as to the beneficial contribution of dynamic formats; possibly tempered by (c) an assumption that viewing a moving image may impose a more restrictive viewing style, and by the vaguer notion that "real" dynamic images may differ from those obtained for the stationary paradigm.

Results for driver, passenger and pedestrian perspectives were reported. It is important to stress that some of the impressions described in the following paragraph were not based on statistical evaluation of the results; statistical analysis of some key aspects of the data was carried out and reported. The descriptive account was driven by various considerations: a) established findings from the stationary eye-tracking studies; (b) what were judged industry-based interests; and (c) what individuals expressed in discussions of the potential of dynamic imagery.

The attempt was made to provide a comprehensive account of the findings. Eye-tracking data were obtained for three "perspectives": driver, passenger and pedestrian. The same video clips were used for the first two conditions, so the same panels were viewable and for the same amount of time (opportunity to see or OTS durations).

Phase Three study: Results

The hit rate scores are reported in full in Tables 3.2-3.4. The leftmost column in each Table is for stationary 6-sheet and 48-sheet panels, which served as "benchmark" conditions for all comparisons with the dynamic panels. The second column shows the hit rates for "mobile" panels, namely buses and taxis. The next two columns are used to show the pre-change and post-change hit rates. There follows a column for the proportion of respondents who fixated the panel in both pre- and post-change phases. Finally there is a column showing the hit rates for the dynamic panels irrespective of the change, scroll or digital transition.

Table 3.2 is for the driver data, followed by Table 3.3 for the passenger data and Table 3.4 for the pedestrian data; there were 24 driver respondents, 15 passenger respondents and 37 pedestrian respondents. There were 8 video clips depicting each of the conditions reported in the table, with the exception of the somewhat elusive digital driver/48-sheet combination which was represented by just 5 clips. The clips also contained eight instances each of buses and taxis, to allow hit rates to be estimated for mobile panels.

It is important to note the range of OTS durations in the various conditions. Understandably the 6-sheet and 48-sheet panels tended to be longer in potential view for pedestrians than for drivers/passengers. But some panels were available for nearly the whole viewing interval. This reflected the reality on the ground, and to obtain a more balanced and even selection would have required an immensely increased effort in capturing the source video material and editing it. In the event the analytical approach provided by the accumulated hit rate process goes some way to removing the unevenness in the OTS values.

The possibility of calibrating back to that longstanding and substantial database was the main reason for including stationary fixed poster panels in the video scenarios used for the new eye-tracking research.

Table 3.2: Hit rates and OTS durations for Phase Three: Driver condition

Phase Three Driver	Measures	No change panel	Mobile panel	Hit rate for pre-change phase	Hit rate for post-change phase	Hit rate for hits in both pre- & post-change phases	Hit rate regardless of phase
Stationary 6-sheet	Mean hit rate	0.159	-	-	-	-	-
(N=8)	Mean OTS duration	6.5	-	-	-	-	-
Stationary 48-sheet	Mean hit rate	0.252	-	-	-	-	-
(N=8)	Mean OTS duration	12.3	-	-	-	-	-
Scrolling 6-sheet	Mean hit rate	-	-	0.105	0.078	0.031	0.151
(N=8)	Mean OTS duration	-	-	4.9	3.2	8.1	8.1
Scrolling 48-sheet	Mean hit rate	-	-	0.219	0.328	0.094	0.453
(N=8)	Mean OTS duration	-	-	4.0	5.9	10.0	10.0
Digital 48-sheet	Mean hit rate	-	-	0.277	0.252	0.117	0.275
(N=5)	Mean OTS duration	-	-	8.4	8.9	17.3	17.3
Bus	Mean hit rate	-	0.316	-	-	-	-
(N=8)	Mean OTS duration	-	10.3	-	-	-	-
Taxi	Mean hit rate	-	0.349	-	-	-	-
(N=8)	Mean OTS duration	-	8.6	-	-	-	-
Overall mean hit rate		0.205	0.332	0.189	0.215	0.075	0.296
Overall mean OTS duration		9.4	9.5	5.4	5.6	11.0	11.0

Table 3.3: Hit rates and OTS durations for Phase Three: Passenger condition

Phase Three Passenger	Measures	No-change panel	Mobile panel	Hit rate for pre-change phase	Hit rate for post-change phase	Hit in both pre- & post-change phases	Hit rate regardless of phase
Stationary 6-sheet	Mean hit rate	0.233	-	-	-	-	-
(N=8)	Mean OTS duration	6.5	-	-	-	-	-
Stationary 48-sheet	Mean hit rate	0.482	-	-	-	-	-
(N=8)	Mean OTS duration	12.3	-	-	-	-	-
Scrolling 6-sheet	Mean hit rate	-	-	0.199	0.237	0.075	0.342
(N=8)	Mean OTS duration	-	-	4.9	3.2	8.1	8.1
Scrolling 48-sheet	Mean hit rate	-	-	0.272	0.531	0.158	0.633
(N=8)	Mean OTS duration	-	-	4.0	5.9	10.0	10.0
Digital 48-sheet	Mean hit rate	-	-	0.453	0.573	0.080	0.627
(N=5)	Mean OTS duration	-	-	8.4	8.9	17.3	17.3
Bus	Mean hit rate	-	0.408	-	-	-	-
(N=8)	Mean OTS duration	-	10.3	-	-	-	-
Taxi	Mean hit rate	-	0.361	-	-	-	-
(N=8)	Mean OTS duration	-	8.6	-	-	-	-
Overall mean hit rate		0.357	0.408	0.287	0.429	0.108	0.521
Overall mean OTS duration		9.4	10.3	5.4	5.6	11.0	11.0

Table 3.4: Hit rates and OTS durations for Phase Three: Pedestrian condition

Phase Three Pedestrian	Measures	No-change panel	Mobile panel	Hit rate for pre-change phase	Hit rate for post-change phase	Hit in both pre- & post-change phases	Hit rate regardless of phase
Stationary 6-sheet	Mean hit rate	0.803	-	-	-	-	-
(N=8)	Mean OTS duration	13.5	-	-	-	-	-
Stationary 48-sheet	Mean hit rate	0.805	-	-	-	-	-
(N=8)	Mean OTS duration	14.6	-	-	-	-	-
Scrolling 6-sheet	Mean hit rate	-	-	0.609	0.677	0.480	0.777
(N=8)	Mean OTS duration	-	-	4.0	6.8	10.8	10.8
Scrolling 48-sheet	Mean hit rate	-	-	0.631	0.791	0.507	0.872
(N=8)	Mean OTS duration	-	-	4.9	17.0	21.9	21.9
Digital 6-sheet	Mean hit rate	-	-	0.492	0.622	0.348	0.753
(N=8)	Mean OTS duration	-	-	4.5	6.8	11.3	11.3
Digital 48-sheet	Mean hit rate	-	-	0.457	0.630	0.291	0.767
(N=8)	Mean OTS duration	-	-	6.4	13.3	18.2	18.2
Bus	Mean hit rate	-	0.372	-	-	-	-
(N=8)	Mean OTS duration	-	5.5	-	-	-	-
Taxi	Mean hit rate	-	0.167	-	-	-	-
(N=8)	Mean OTS duration	-	7.2	-	-	-	-
Overall mean hit rate		0.804	0.270	0.547	0.680	0.406	0.792
Overall mean OTS duration		14.1	6.3	4.9	11.0	15.5	15.5

Before considering what the Phase Three study said about any multiplier that might be applied to proceed from a set of stationary panel scores to their dynamic panel equivalents, it is important to see whether the study produced results that were commensurate with those from the stationary eye-tracking studies. The possibility of calibrating back to that longstanding and substantial database was the main reason for including stationary fixed poster panels in the video scenarios used for the new eye-tracking research. For this purpose the simplest benchmark conditions were the 6-sheet and 48-sheet stationary panels viewed in driver and pedestrian scenarios, but it would be surprising if the comparison between the dynamic versions of these two formats did not yield the usual advantage to the larger of them. In fact in every case the score for 48-sheet panels was higher than for 6-sheet panels. The smallest difference (0.803 vs. 0.805), and it was indeed very marginal, was for the stationary panels viewed by pedestrian respondents. These comparisons are made via the first data columns in Tables 3.2 to 3.4.

A gross measure of hit rate was obtained by noting whether there was a hit on the panel regardless of movement phase. The data are in the final column in each case. The picture was far from uniform. Scrolling 6-sheets did not score appreciably higher than their stationary counterparts. Scrolling 48-sheets appeared to be at an advantage over stationary 48-sheets; the advantage for pedestrians was small but this may have reflected the high score for the stationary format for the viewer on foot. Digital 6-sheets were only seen by pedestrian respondents and were at the same level as stationary 6-sheets. Digital 48-sheets were only appreciably higher scorers than stationary 48-sheets for passenger respondents.

Pre-change scores were mostly lower than post-change scores; however, it should be remembered that the observer was nearer the panel after it had scrolled/transitioned. The driver condition seems to have been an exception, at least partially, because there was no post-change improvement in two cases.

A score was obtained for all respondents as to whether they viewed both faces of a scrolling panel. For drivers this averaged about 8%, for passengers about 11% and for pedestrians about 41%.

There was something of a progression in hit rates (for non-mobile panels) between the three perspective conditions, with driver hit rates the lowest, pedestrian hit rates the highest, and passenger hit rates in between. It should be remembered that the drivers and passengers saw exactly the same material so it would not be surprising if their hit rates were in close resemblance, but this was not apparent in the results. Pedestrians scored higher than passengers who in turn scored appreciably higher than drivers on dynamic panels (79% vs. 52% vs. 30%); this pattern was roughly sustained for the stationary benchmark panels (80% vs. 36% vs. 21%); however, it was quite different to that for mobile panels (27% vs. 41% vs. 33%).

Estimating the DFF

The objective of the Phase Three study was to provide a first estimate of what was to be termed the "dynamic frame factor" (DFF); that is, a multiplier to be applied to the visibility score for a static (fixed) panel to provide an estimate of the visibility score for a dynamic version of the panel. Since there were at least two variants on how dynamism was implemented (scrolling and digital), at least two values of DFF may eventually be required.

Two ways of calculating the overall hit rates for dynamic panels were used. First, the raw data were used to compile a record of all respondents who looked at the panel in both its pre-change and post-change phases. Second, a probability formula was employed to combine the hit rates for the two phases (shown to the far right of the tables, along with DFF values). As to DFF itself, this was defined as the hit rate for the dynamic panel divided by the hit rate for its static equivalent. There were other options, but this was an intuitively supportable choice, and was the one used as DFF Method 1.

DFF Method 1

Preliminary indications of what values DFF might take can be derived from Tables 3.2, 3.3 and 3.4. The simplest thing was to compute the ratio for the mean hit rate recorded for a particular dynamic format relative to that for its static fixed panel equivalent. If dynamic formats were (in terms of their hit rates) better than, equal to, or worse than their static fixed equivalents, their DFF values would respectively be greater than one, equal to one, or less than one.

For example, from the driver data, the 6-sheet scroller had a DFF of 0.151/0.159 = 0.95, whereas the 48-sheet scroller earned a DFF of 1.80, and the digital 48-sheet got 1.09. These results are shown in Table 3.5 together with those for the passenger and pedestrian conditions. Common sense prevailed in the cases of those which were less than one and the DFF was set as 1.00 in each case. From these results only the dynamic 48-sheet panels would merit a substantially positive DFF.

Table 3.5: Dynamic frame factors by DFF Method 1

	Driver	Passenger	Pedestrian
Scrolling 6-sheet	0.96	1.45	0.90
Scrolling 48-sheet	1.81	1.32	1.03
Digital 6-sheet	-	-	0.94
Digital 48-sheet	1.63	1.47	0.95

DFF Method 2

The basis for hit rate scoring for this method was "within" each respondent and across the images comprising each format. Method 1 entailed aggregating within images and across respondents. The two may well not agree with great precision because of their computational basis. Table 3.6 shows the results in the same layout as Table 3.5. The basis for calculating these measures also enabled confidence limits on each DFF to be established; these values (for a 95% confidence interval) were reported in Table 3.7.

Table 3.6: Dynamic frame factors by DFF Method 2

	Driver	Passenger	Pedestrian
Scrolling 6-sheet	1.06	1.35	0.93
Scrolling 48-sheet	1.63	1.70	1.22
Digital 6-sheet	-	-	0.94
Digital 48-sheet	1.33	2.01	1.00

Table 3.7: 95% confidence limits for dynamic frame factors using DFF Method 2

	Driver	Passenger	Pedestrian
Scrolling 6-sheet	0.71 - 1.41	0.90 – 1.80	0.92 – 1.06
Scrolling 48-sheet	1.32 - 1.93	1.11 – 2.30	0.97 – 1.57
Digital 6-sheet	-	-	1.47 - 1.75
Digital 48-sheet	0.81 -1.17	0.73 – 2.86	0.85 - 1.15

Finally, Table 3.8 shows the percentage of respondents whose hit rates in a given dynamic condition were greater than those in the corresponding baseline condition (stationary panel). This was reflected by a DFF greater than one. There was no case in the Table where the percentage falls below 50%, suggesting that there was a general advantage to dynamic panels. Finally, Table 3.8 shows the percentage of respondents whose hit rates in a given dynamic condition were greater than those in the corresponding baseline condition (stationary panel). This was reflected by a DFF greater than one. There was no case in the Table where the percentage falls below 50%, suggesting that there was a general advantage to dynamic panels.

Table 3.8: Percentage of respondents for whom the DFF was greater than 1

	Driver (N=24)	Passenger (N=15)	Pedestrian (N=37)
Scrolling 6-sheet	96%	87%	78%
Scrolling 48-sheet	63%	53%	57%
Digital 6-sheet	-	-	57%
Digital 48-sheet	63%	80%	57%

Statistical analysis (repeated measures analysis of variance) of the DFF data, comparing the various conditions for Driver, Passenger and Pedestrian respondents, revealed significant differences in DFFs between formats for Drivers ($F(2,23) = 6.61$; $p < 0.01$) and Pedestrians ($F(3,35) = 14.1$; $p < 0.001$) but not Passengers ($F(2,14) = 0.63$; not sig.). The locus of the effect in the first two cases was evidently the superiority of the scrolling 48-sheet format. The Passenger data suggested no such difference, but a general advantage to the dynamic panels, with an overall average DFF of 1.47. The mean DFF value for Drivers and Pedestrians was 1.22 for both groups.

Phase Three study: Discussion and recommendations

Application of DFF

If the present approach is accepted, the question eventually is whether the results are applicable, and if so, how? Should the goal be a generic multiplier, or one that is format-specific? Is the outcome sufficiently well specified for Route's needs and purposes?

These questions are not for us to answer in this paper but they do underpin some of the final remarks.

Final caveats

- The opportunity to see each panel was obviously not subject to control, but its duration was noted and under each hit rate in the tables, the average duration of the panel being potentially in view is recorded.
- It should be noted that Route has available a pertinent source of eye-movement data in dynamic conditions from its own study of hit rates and OTS durations, as reported in a Postar Technical Report in 2006 (Barber and Sanderson, Video analysis of driver eye-behaviour). The observers in this study had their eye movements monitored as they were driving or were being driven in a specially equipped car. The results from this study do not enable comparisons with dynamic panels, but they are relevant to the performance levels achieved when recording is over the entire OTS interval, and while the observer is actually on the move carrying out a real-world task.
- Hit rate obviously varies with duration, and its effect has not been statistically "partialled" out. Further analysis is required to specify what this would imply for generic or format-dependent conclusions.
- The hit rates relate to the whole OTS or potential viewing interval, and so the scores in the tables are accumulated hit rates.

So, is the conclusion that "further research is needed"?

- Further analysis of the dynamic eye-tracking data would seem worthwhile, expanding the range of panels analysed. Some clips contained more than what was counted as the target panel – as when a bus was in view at the same time as another stationary panel, so there is more to extract from the un-analysed panel eye-tracking data.
- The tabled results are particularly interesting as a source of data on OTS durations which on the face of these data range from 4 to 22 seconds. Obviously the range is skewed to reflect the requirements of editing the material into video clips, but the unedited videos contain a substantial number of panels that were not included in the eye-tracking portfolio of clips.
- A detailed specification of format differences in DFF would obviously require a suitably expanded study.
- Specification of a generic DFF is possible, but with due caution regarding the sample of panels used in the present study. The error of estimate could be reduced by increasing sample sizes (of respondents and panels); the present data are sufficiently comprehensive for the increase in size needed to be estimated with confidence.

The first two of these options require desk-based effort. The second two could entail a wider use of what the video archive made available. Alternatively it could require an expansion of the archive (more video recording) and more eye-tracking research. Further investigation should desirably also be guided by the growing body of relevant research which is the subject of the final section in this report.

Chapter 5

Review of literature: Visibility, visual attention and dynamic images

Introduction

This report ends with a review of research that is an extension of the preamble to the dynamic imagery research that was tabled for discussion prior to the laboratory studies. The literature search was conducted in 2015 adding to the one in 2008 ahead of the empirical research. We have addressed in the report thus far the contribution of various forms of movement to the efficacy of outdoor advertising, not forgetting the changes brought about by the digital revolution to outdoor vehicles for communication. We do not consider in this report how this may be mediated or moderated by execution factors; this is a matter for poster design though much of what we discuss is relevant to such concerns. The following review considers the more general role of movement in its several forms in the control and mediation of visual attention.

Movement is intrinsic to our visual environment. People move in relation to poster panels, and panels (for example, on buses and taxis) move in relation to the people with an opportunity to view them. An increasing proportion of fixed location panels have dynamic properties, the most familiar of which are scrolling panels, which are being joined in numbers by their digital counterparts. Their rationale seems to be twofold – they offer multiple opportunities for advertising on a single structure (a fundamental commercial feature), which they achieve by the mechanism of scrolling (a basic visibility feature). Scrolling – along with a multitude of related variants - is the aspect that is claimed to have visibility benefits by capturing the viewer's attention. Any such advantage could be due to the movement intrinsic to the change mechanism, the onset of the change process itself, or the change of content (colour, graphical form, etc.). This rather academic-sounding breakdown of the scrolling process is justified by the research to be reviewed below. For the moment, it is sufficient to underline that one can properly distinguish between change as a discrete event, and movement as a continuous property of a display, either, both or neither of which could capture attention.

Aside from the 2001 Sutton Study of JCDecaux, which assessed a single scrolling 48-sheet panel, there appears to be no industry research to evaluate scrolling panels in terms of visibility or "impact"; whether there is anything in the pipeline is not clear. The situation regarding viewers on the move and vehicular panels is considerably more fluid and active. First, historically, there is Route's own study (Barber and Sanderson, 2005) which analysed the eye movements of drivers and passengers in an instrumented car driving on the roads of a UK city as they passed buses and roadside panels (the data collection was not intended for this purpose). This was a seminal achievement for the industry. Subsequently there have been studies using video presentations of scenes as viewed from a vehicle driving past fixed roadside panels (e.g., Crundall et al., 2006), and others of in-car recordings of drivers passing digital billboards (e.g., Lee, McElheny and Gibbons, 2007).

One key assumption underpinning the introduction of dynamic imagery in advertising structures such as poster frames is that movement within or between frames attracts attention. It is intuitively compelling that this should be so, yet the question of whether movement does attract attention is not a closed matter. Although most observers would very likely contest that it is obvious that motion works as a powerful focus for attention, this is an untested assumption. Research would first be required to demonstrate that visual attention is captured by dynamic imagery. This all seems so fundamental that it is perhaps not surprising that there has been little research on the matter until relatively recently.

On the other hand, as shown in the opening paragraph, it quickly becomes clear that the core question masks many less straightforward issues. For instance, it is possible that motion captures attention but does not sustain it, or sustains it for a limited amount of time. Moreover, until it is demonstrated empirically, we are not in a position to measure the force of this attentional factor, and this would be crucial if the requirement is to derive a measurement system for object visibilities. Assuming positive results from this research, a further question for investigation is to establish *how* motion achieves its effects in capturing/sustaining visual attention. Careful analysis and clever experimentation would be needed to underpin the investigation of this latter question. It is apparent from the distinction between the potential capture and maintenance of attention that the former may be mediated by the onset of movement and the latter by the actual movement itself.

For most people it may seem beyond doubt that attention *is* captured "automatically" by movement; this is surely consistent with everyday experience. It also fits well with preconceived notions of movement detection being essential to survival. From an evolutionary standpoint one might further hazard the idea that movement would be especially effective outside the central region of the eye – wouldn't the detection of predators be mediated here? This line of reasoning based on the idea that motion detection has survival value, raises the interesting possibility that different forms of movement may have different degrees of attentional power. Distinctions that have been addressed experimentally include those between animate and inanimate forms, rotational and linear, and objects that are looming vs. retreating (via representation of changing distance or eccentricity relative to the viewer).

Any such advantage could be due to the movement intrinsic to the change mechanism, the onset of the change process itself, or the change of content (colour, graphical form, etc.)

Another set of issues are represented by mobile advertising forms, such as panels on buses and other types of transport. There is also the thorny question of the observer's own movement, in a car, on a train, on a bus or on foot. It is not difficult to see that the bodily and visual dynamics in the last case are likely to introduce highly complex dynamical complications; notwithstanding, an investigator may take some reassurance that the visual world is not subjectively as disjointed to the observer as seems likely to be the case on considering the jerky motion of a pedestrian.

But the discussion is already pointing to distinctions that are made by academic and applied research, yet to be described in this report, and we turn therefore to this research. Scientific psychological research on attention has been centre stage since the 1960s, and a vast literature has accumulated since that time. Carrasco (2011) estimated in a review of this literature that close to 2500 articles on visual attention were published after 1980, more than half of them between 2005 and the time of his review. One shortfall of this flourishing area of research was that little consideration was given to the contribution of dynamic changes to the guidance and attraction of attention until the 1990s. However, in one of the earliest studies on the possibility that stimulus movement could guide visual attention reflexively, Hillstrom and Yantis (1994) stated that "the experiments reported in this article seriously challenge the common belief that motion captures attention in a stimulus-driven fashion". Most of the evidence subsequently suggests it is considerably more complicated than this, however, motion *per se* may not attract attention automatically, although it can be used voluntarily to guide attention to its location. This distinction and the associated research literature are more comprehensively considered below.

Finally, it will be noted that there recently has been a resurgence of interest in driver distraction – with much attendant publicity; an important aspect of this concern with safety is the development of new methods for presenting advertising material, including motion/dynamic imagery. The direction of future research will no doubt need to take account of the issues raised by the various interest groups, including any relevant research findings that are instigated.

Aside from providing some empirical answers to basic questions about movement and attention, there are lessons regarding the various factors that may be confounded or that covary with movement, and the conditions under which they compete or combine.

Review of literature

Our literature review is based primarily on open source reports on topics related to the plan to extend visibility research in new directions, notably to deal with dynamic image advertising techniques (such as mechanical and electronic scrolling, but not forgetting advertising panels on vehicles). The bulk of the relevant publications (Section A), located in academic journals, were concerned with basic issues in psychological and visual science, and the research methods used reflect this orientation. A few studies (Section B) were also examined that are more directly related to advertising on the web though with findings that seem to have some significance for outdoor advertising. Other pertinent reports from industry sources on digital advertising are also noted (Section C). What is learned from these studies does not of course enable the direct development of a metric for the attentional value of dynamic imagery. But aside from providing some empirical answers to basic questions about movement and attention, there are lessons regarding the various factors that may be confounded or that covary with movement, and the conditions under which they compete or combine. There are also exemplars of research methodology – and technological developments – that may drive or at least inform future applied research into visibility and attention (Section D).

A: Basic research

Readers who acknowledge the importance of a poster site's capacity to compel attention will see the potential relevance of academic research on visual attention and visual search to outdoor advertising, even though applications have yet to be delivered on any scale. The topics covered by this wave of research are what factors may capture attention or guide search, and many such factors have been investigated. It will be easier to appreciate whether or not the findings and research techniques of this research area - which is owned for the most part by academic researchers - are relevant to present concerns, if examples of the methodology are described.

Theoretical treatments of attention and search are underpinned by data from a variety of experimental paradigms, including one in which the participant searches for a target in a display containing a number of non-target items; some displays will typically contain no target, so the task simply amounts to reporting (as fast and accurately as possible) whether or not a target is present. The task might for example involve searching for a red colour patch in an array of patches of other colours; or searching for a letter in an array of single digit numbers. The former is an easier task, particularly if the non-targets are in just one other colour, blue say. Performance is measured by the time taken to locate the target (commonly the task will be to say whether or not the display contains a target). In practice, the task also involves measuring performance for a number of display sizes (number of items in the display); search time often increases as display size increases; but not always, since sometimes when the task is very easy and the target "pops out" from the background, search time is more or less flat as display size increases. Indeed the performance measure of principal interest is the slope of the line relating search time to display size (flat for

Posner and Cohen concluded that they had shown: first, that a peripheral signal can summon attention; and second, that this comes to serve to inhibit further processing at that location.

the easiest of tasks - i.e., when there is “pop out” – and steep when the task is difficult). Those visual attributes that are attentionally most effective may in principle be identified using methods like this.

Another approach has been taken by research in which attention is likened to a spotlight that enables the observer to select a part of the visual environment, though whether this is a satisfactory metaphor is a moot point. Indeed the general topic of attention is replete with debate and controversy (see Pashler, 1999; Styles, 1997); the range of experimental approaches, the variety of conceptual distinctions, and the sheer amount of published work, are daunting to anyone seeking to extract pithy summary points or lessons. Sticking with the idea of a spotlight for the time being, we could ask whether the spotlight can be attracted to an object or region in the visual field; whether this can be achieved at will or involuntarily; and whether the span of the spotlight is adjustable. Posner (1980) introduced the spotlight concept to explain his finding that directing visual attention to a location facilitates the processing of a target that appears at that location.

A connection could be assumed to exist between the attentional spotlight and eye fixations/movements, since the positioning of the eye seems to be intrinsic to the visual orienting process. On the last point there is some evidence that we may attend elsewhere than to where our eyes are pointing (e.g., Kaufman and Richards, 1969) but this introduces a nuance too far for this discussion, and for this purpose we shall accept the quite strong link between eye position and attention, but we must continue with our preamble on the orienting of attention.

The work on the mediating role of motion on attention is a subset of the research field taken as a whole. To introduce this work it is useful to provide a brief overview of some seminal research on attention reported by Posner and Cohen (1984). For this study, a procedure was adopted that contains a basic format that has been widely used and adapted, including to studies on attention and motion, abrupt onsets, and so forth. The participant has to respond when a specified target appears (in this case a dot inside one of three square frames, arranged about a central point on a screen). The dot does not always appear, and the response must be withheld. To signal the start of a test trial, the square empty frames are shown. This “fixation signal” is shortly followed by a “cue signal”. The peripheral cue is a brightening of one of the cue signal squares. In the following example, the cue is the left-hand box and is invalid because the target appears in the right-hand box. A complete description of the experiment would have to show all the other combinations of cue and target position and the one chosen in Figure 4.1 is simply an illustration of the possibilities:

Figure 4.1: Canonical research design: Inhibition of return

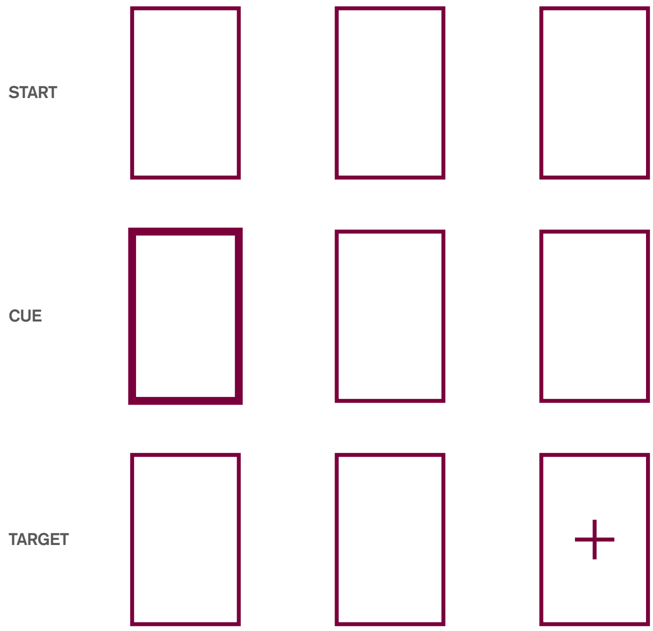

The other variations include where the target is validly cued (right-hand box). The intervals are all short, but the cue-target interval is varied (and is crucial for the results). The reaction time is generally less than half a second. Another important feature is that the probabilities of the cue being on the left, at the centre or on the right differed; they were 0.1, 0.6 and 0.1 respectively, with the remaining 0.2 of trials with no target. So a valid cue and an invalid cue were equally unlikely (p=0.1). Incidentally this form of cuing was described as exogenous, to be contrasted in other experiments with endogenous cuing, where the cue was an arrowhead in the central box. This is a means of contrasting processes depending on peripheral (sensory) processing with those relying on central (cognitive) processes. This is an important distinction that is thought by some to reflect the difference between bottom-up and top-down processing.

The results were considered to be highly significant from a theoretical point of view. The principal findings were twofold: responding at the shortest intervals between cue and target was quicker when the cue was valid than responding when it was invalid; but there was a crossover effect at longer intervals such that responding when the cue was valid was significantly slower than when it was invalid. Posner and Cohen concluded that they had shown: first, that a peripheral signal can summon attention; and second, that this comes to serve to inhibit further processing at that location.

This effect is known as "inhibition of return" (IOR), and has been much studied subsequently with variants of the basic procedure. This package of two findings (along with much additional detail) is of interest mainly for its theoretical consequences, but also because it is a combination of what seems intuitively obvious (the effect at short intervals) and something that surely is quite counter-intuitive (the effect at long intervals). The full detailed findings will not be given here but it is worth adding that when the peripheral cuing signals were replaced by a directional signal (an arrowhead pointing left or right) in the central box, the inhibition of return effect disappeared.

The basic IOR experiment is an experimental paradigm that has spawned many others, including studies of the contribution of motion as a property mediating the capture of attention. A swift convergence on applied topics does not seem to have taken place but the literature does increasingly feature references to real-world concerns (e.g., Kawahara, Yanase and Kitazaki; Mital et al., 2010) while busily expanding its subject matter. The remainder of this section is devoted to a summary account of these investigations where they make contact with practical aspects of visual attention, with particular reference to motion and related stimulus attributes.

Motion and attention

In this section the focus is on the potential involvement of stimulus motion in the control of attention. As in the previous discussion of basic research, we skirt round most of the methodological and empirical minutiae of the studies mentioned, however, some detail is provided to illustrate the experimental approaches that have been adopted.

Research on the role of motion in the capture and control of visual attention has flourished in the past 20 years or so as new theoretical and methodological avenues have been explored. A fundamental point made early on in this period was that the onset of motion and the motion itself are ordinarily confounded. This is true for other properties such as luminance but for present purposes the focus will be on motion onset (and offset). Abrams and Christ (2003), among others, devised a method whereby de-confounding of motion and motion onset could be achieved. The findings from their study suggested that a continuously moving target was no easier to find than a static one, but one defined by motion onset was. This was a task in which the observer was seeking a target whose motion-related properties were manipulated. In a later study Abrams and Christ (2006) concluded that motion onset is more efficacious than continuous motion in influencing the direction of attention.

Other experimental evidence bears more directly on the question of whether attention can be steered by irrelevant movement-related stimulation, and much of it applies experimental procedures derived from the IOR paradigm. The emphasis on "irrelevant" is important in the context of the use of dynamic imagery in advertising because the movement of or within a poster frame is typically no more relevant to the observer's primary task of navigating his or her way in the environment than is the poster itself. It is nevertheless of practical interest to discover what influence image motion, relevant or otherwise, may exert on attention; if any, as the next paragraph makes clear.

We reported above that Hillstrom and Yantis (1994) posited a challenge to "the common belief that motion captures attention in a stimulus-driven fashion". Their experiments are of interest not just because of the overall conclusions reached but also because of the "implementations of motion" their experiments entailed, and the fact that the focus of interest was on the effects of these manipulations on attention. Again we shall see a clear link with outdoor advertising forms in which motion is implemented in various ways. In one experiment three types of motion were tested "in which the contours of the letter remained stationary, but elements or textures within or near the letter moved: (1) diagonally striped texture inside the moving element moved smoothly from right to left; (2) dots revolved around the moving element; or (3) random dots inside the moving element rapidly and randomly repositioned (producing scintillation)." In addition there were two types of motion of the whole letter: "(4) horizontal oscillation; and (5) looming vs. receding (in which the moving element respectively gets larger or smaller)." (p. 400, with some minor editing of the original text).

On the basis of their experiments, Hillstrom and Yantis reached the following conclusion: "Motion joins a growing list of features that can guide attention when they constitute a target's defining attribute... but do not capture attention when they are irrelevant to the observer's task... These features are sometimes unseen when attention is focused elsewhere in the visual scene …". This summary statement is prophetic as to the focus of much later research but not to the current state of play as will be seen.

Following Hillstrom and Yantis's seminal research on motion and attention, different types of motion (e.g., looming, translating, receding) have been subjected to further experimental scrutiny. Franconeri and Simons (2003) claimed that motion in the form of looming and translating captured attention but receding did not. This could be interpreted in terms of the lesser biological urgency of animate objects that are retreating or departing, although this "explanation" does not seem to apply to the hunter in pursuit of prey. Indeed although early evidence suggested that receding stimuli did not exert an influence on attention, this has been qualified by further research and some recent studies have demonstrated positive evidence of a receding effect. Some results suggest that it is weaker than the effect of looming (e.g., von Mühlenen and Lleras, 2007), however, equivalence of the two has been claimed in a study using motion in depth by Skarrat, Cole and Gellatly (2009). In a further nuance in this developing story, while confiming the earlier finding of equivalence in size of effect, Skarrat, Gellatly, Cole, Pilling and Hulleman (2014) have argued that the effect of looming implicates motor control while that of receding is perceptually meditated. The practical implications will need to be carefully considered if these findings and analysis are supported by further research. At the outset of this section it was stressed that motion and motion onset are ordinarily confounded and clever experimentation is required to separate their effects. Dynamic imagery as implemented in poster panels exhibits just this form of confounding, so we next review some of the research that has focussed on abrupt onset/offset effects whether or not movement is involved. This research illuminates important conceptual issues that bear on the practical concerns of this report.

Abrupt onset/offset effects on attention

People appear preferentially to select objects that have just arrived in the visual field at the expense of objects that are already present in the visual field; this is sometimes characterized as the "preview effect". The idea is that people are able to prioritize the selection of new items over old items in visual search and attention. Several competing explanations have been put forward – visual marking, temporal segregation and attentional capture hypotheses – the last of which has survived with considerable experimental support. Much of the evidence does not relate to moving stimuli but it raises issues and concepts that are important to an understanding of motion and attention.

The story has developed slowly around a key theoretical contrast between bottom-up and top-down control of visual attention. It has proved difficult to devise experimental procedures to separate these (and other) competing accounts unequivocally. While Donk and Theeuwes (2001) showed that abrupt onsets attracted attention their findings were explicable in terms of both bottom-up (involuntary or stimulus-driven) and top-down (intentional or goal-directed) processing; the former depends on activation of selected elements in the visual field whereas the latter involves active goal-directed inhibition of such elements. Further research by the same researchers (Donk and Theeuwes, 2003) led them to conclude that the prioritized selection of new elements "is mediated by a bottom-up process. It seems as if the abrupt onsets accompanying the appearance of the new elements generate a large bottom-up activation, biasing observers to prioritize the processing of new elements over old ones".

Generally there seems to be an asymmetry in the ability of onsets and offsets to affect attention, and in a report making the point more specifically Pratt, Theeuwes and Donk (2007) noted that there are some studies that show that offsets produce weaker attentional capture than do onset stimuli. Significantly for practical issues, Boot, Kramer, and Peterson (2005), using an eye-movement paradigm, found that saccades were much more likely to be made to onset distractors (i.e., irrelevant stimuli) than offset distractors. Similarly, incorrect pro-saccades in an antisaccade task (in which an eye movement had to be made away from the direction signalled) were much more likely to be made in response to onset distractors than to offset distractors (Pratt and Trottier, 2005).

Brockmole and Henderson (2005) were the first to use eye tracking of real-world scenes to investigate the impact of new objects on the direction of attention. The basic components of gaze behaviour (fixations and saccades) provide a test-bed for exploring the role of the low-level transient signals that often accompany the appearance of an object; a new object arriving during a fixations elicits the transient signal whereas if it arrives in the course of a saccade, the signal is much reduced due to a process known as saccadic suppression (Matin, 1974). Brockmole and Henderson found that onsets routinely captured attention even in the absence of transients (i.e., when the onset occurred during a saccade), while offsets only captured attention when salient transients occurred or if the offset object was occluding another object (in which case the "offset" may be likened to the onset of the revealed object).

Pratt et al. (2007) noted that while onsets and offsets are both typified by concomitant luminance changes, a likely reason why offsets are less effective in capturing attention is that they lack the extra attention-capturing effect of onsets that is associated with the appearance of new objects. This throws some useful light on the position reached by some researchers who concluded that offsets are as effective as onsets so long as they are accompanied by a change in luminance. More important from a practical point of view is the caveat that the earlier research (as cited in this paragraph) was based on different methodologies. The Brockmole and Henderson, using real-life scenes, represents an important step towards external generalization of the findings.

A visual search task like the one described at the start of this section was used by Kawahara, Yanase and Kitazaki (2012) to assess the roles of bottom-up (stimulus-driven) and top-down (attentional set) factors in mediating the effect of motion on attention. A single green letter was the target embedded in a sequence of variously coloured letters while a background of dots was presented forming an expanding visual flow. The motion was arranged so a contrast was possible between commencement and cessation of movement (onset/offset), and a condition was included that varied the speed of movement. The experiments showed both onset and offset effects, but no effect of speed of the irrelevant background motion. The evidence of offset as well as onset effects is somewhat unusual but may indicate that the type of motion cue is important. The findings were interpreted as favouring a bottom-up explanation, in which attention can be biased by movement that is not relevant to the observer's primary activity.

Other experimenters have added to the evidence and it seems that there is an effect of motion when it acts as a cue to a target itself defined by motion, that is the cue property is relevant but not necessarily when the target is defined in another dimension (e.g., colour) and the cue property is irrelevant to the target. This was demonstrated by Folk, Remington, and Wright (1994) for the much-researched dimension of colour. Their conclusion was that a stimulus will capture attention only if it shares an attribute with the target. For example, if the task is to find a red target amid grey stimuli, a cue or distractor will only affect attention if it too is coloured red. However, Al-Aidroos, Guo, and Pratt (2010) have reported experiments purporting to show that new motion may capture attention even when the target is defined by colour and the distractor motion is irrelevant.

The key thing about much of this research is that the tenor of at least some accounts is that change effected by the onset and the offset of items in a display can be drivers of visual attention; the motive force of offset probably being less compelling than that of onset. As reviewed above the research otherwise points to nothing so compelling on the part of ongoing motion as a stimulus attribute. Excellent comprehensive reviews of the research considered here have been published but they are understandably uncompromisingly detailed, and they cannot yet paint a picture that shows closure across the board. Hence, as in so many instances of this research, it is incumbent on an interested reader to continue to "watch this space" while waiting for further research, and the eventual completion of the big picture.

Motion and onset/offset effects revisited

It seems likely that further studies of movement properties will add more distinctions to the evidence base and the possible inferences to be drawn, as indicated by the final investigation considering the role of motion in the mediation of visual attention.

This is a study by Mital, Smith, Hill and Henderson (2010) in which they used an innovative approach to the question of dynamic scene viewing as monitored by eye movement recording. Their subjects were asked simply to view a series of unconnected video clips (including advertisements, music videos, news items, etc.) while their eye movements were recorded, each clip being followed by a four-point rating of the degree to which they had liked the content. An important feature of the study was that the participants were not required to search for anything, to follow the main actors or actions; such instructions tend to bias the participants' viewing towards "exogenous control" – and presumably thereby enhancing attention to motion. A novel method for classifying the eye movement data was used (see Section D), and complex methods were used for computing the visual features present at foveated locations in the videos. Luminance, colour, edges, corners and orientation properties were the static features investigated, while flicker and object motion were the dynamic features assessed. Baseline foveations were formed for comparison with the measures at the actual foveations by taking a random but equal-sized sample of the features at all foveated locations pooled over participants for each movie. The result that stands out from a complex set of findings is that motion (optical flow) was the strongest predictor of gaze location. From the perspective of this review imposing a relaxed instructional set was a useful feature of the study, and would no doubt be in line with the "bottom-up" manner in which drivers and pedestrians view outdoor advertisements. The authors do speculate that a contrast with a "top-down" approach would be of interest in order to assess the relative contributions of exogenous and endogenous control of attention. Although we have assumed that our vehicle occupants and pedestrians are operating on a bottom-up basis, there is that sub-population of ad-avoiders who perhaps should be assumed to operate "top-down" if they are to be successful!

Other interesting developments include the work of Brockmole and Henderson (2005) discussed above in which real-world scenes have been used to explore attention phenomena. This research has spawned a series of further studies on what is termed "oculomotor capture". For example Matsukura, Brockmole and Henderson (2009), with a real-world scene viewing paradigm, showed that colour changes in a scene that were devised so that they did not result in physically new objects, nevertheless captured attention, but not so effectively as actual new objects. There appears to be a modest increase in the potential applicability of attention research, although interest does not yet seem to have extended to a study of motion as an attention-related property in such scenes.

Other factors

The potential for various stimulus attributes in capturing or otherwise mediating visual attention has also been thoroughly examined; aside from "motion" and stimulus onset and offset, they include "blinking" (display elements flash on and off successively), luminance, colour, and "looming" (part of the display expands and thus seems to grow in the direction of the viewer). It is worth noting that all of these different threads of evidence may have some potential for application.

Summary of attention research and its application

This very brief description should help the reader to assess the direct applicability of this kind of research to outdoor advertising. The examples from the literature on visual attention illustrate the type of abstract procedure that researchers adopt in studying attention and the nuanced theoretical issues that emerge from the research. It should be emphasised that this research has normally not been conducted with such applications (or indeed any other!) in mind; moreover it often addresses finer distinctions that are useful for practical purposes. What cannot be gainsaid, however, is its methodological sophistication and more importantly, the sensitivity of the methods used to variations in stimulus attributes that undoubtedly are practically important. This is best illustrated by the subtleties regarding how movement might work (or not) that are revealed by this research.

In summary - and in the specific context of scrolling displays - the research seems to justify a reasonable *a priori* position wherein scrolling may catch attention - mostly by virtue of the change to a new image, less by the change from an old image, and probably none at all because of the motion associated with the scrolling mechanism. It should be emphasized that this requires a leap of faith from a piecemeal assembly of research findings to a real-world application. The evidence does in aggregate support the assertion that scrolling is an effective attention-getting device. It does not clarify two other issues: first, whether or not scrolling retains attention to the point where road safety is compromised (that arguably depends on the potency of the execution); second, what is the magnitude of the effect on attention from a practical point of view (to put this more crudely, by how much would one expect visibility hit rates be uplifted?).

The research makes some important distinctions about the nature of attentional processing and establishes a wide range of incidental factors that influence attention. One important conceptual approach that is required for an understanding of the academic research reviewed above relates to the contrast between bottom-up (also referred to as data-driven, stimulus-driven or exogenous) processing and top-down (also known as endogenous or focally attentive) processing. This contrast may be roughly aligned with the expectations and knowledge required for the "top-down" control involved driving a car or navigating one's way as a pedestrian in an environment containing stimuli that may or may not capture one's attention in a "bottom-up" (involuntary) fashion.

This review has not addressed the full range of factors that may be implicated in the mediation and control of visual attention (e.g., it excludes colour, one of the factors that has been most thoroughly investigated). The key factor for the review is that of movement, which is quite nuanced according to the analytical and empirical treatment that researchers have devoted to it. For instance, it is important to recognize the contrast between abrupt onsets and offsets on the one hand and the motion itself on the other, and their possible separate contributions to attention.

It is also worth noting that luminance and other changes may accompany the occurrence of onset and offset events, further factors that should be unpacked from the general flux of variables.

B: Applied research: Attention and web-based advertising

No studies were found that focused on attention to adverts in public areas (roadside, underground, retail, etc.), however, a few studies were found that examined attention to banners on internet sites that offer some insights into factors associated with dynamic imagery.

In the first of these studies Hong, Thong, and Tam (2004) asked their subjects to use an online grocery shopping system developed for the purpose of the experiment. The display contained banners, some of which flashed. In one condition (task-relevant) the target item was flashed, which was contrasted with a condition in which the non-target item was flashed (task-irrelevant). Two types of web page were constructed – one with information in a list format and a second in a matrix; this corresponded to a contrast between a high local density environment and a low local density environment. Performance was measured by response time. This measure was not affected overall by flashing in the task-relevant condition, although response time to a flashing target item decreased in the high local density environment, balanced by an increase in the low local density environment. In the task-irrelevant condition, response time was significantly longer when a non-target item was flashed than when it was not flashed, but response time did not differ between the high local density and the low local density environments. The findings demonstrate some possible complications to the outcome that may arise when a target is signalled by flashing, and how this may be moderated by a variable linked to clutter.

Burke, Hornof, Nilsen, and Gorman (2005) reported two experiments in which they tested whether standard animated banner ads affect web users' visual search speed, perceived workload, memory for ad content and gaze patterns. Results for the first and fourth measures are most directly relevant here.

In the first experiment subjects searched for short words while two banners appeared within the search area. Banners included (a) an animated commercial, (b) static versions of the commercial banners, (c) novel cyan banners that flashed big text, and (d) dummy (blank) invisible banners. Banner animation failed to capture attention – subjects found the words just as quickly when an animated banner was presented as in all other conditions. However, the authors conceded that the lack of an effect might have been a consequence of the task as the subjects were given an incentive not to look at the banners.

In the second experiment eye tracking was used together with a more ecologically valid task – searching for news headlines. Subjects performed two types of searches: "exact" in which the target headline was known, and "semantic" in which the first few sentences of a full story appeared and the best matching headline had to be found. The main finding was that search times were no different between animated and static commercial banners; eye fixation results also showed no significant differences between the banner types. The authors concluded that "graphics and animation in the commercial banners did not attract participants' gaze" and "though the static and animated banners did increase search time, the increase cannot be attributed to participants looking directly at the banners and thus processing their detailed content. Instead, the delay might be caused by graphics and animation viewed peripherally".

The most important finding from this limited slice of the literature is that there was no significant difference between the animated and the static banners. This is at odds with the findings of Hong et al. (2004), so some further clarification is to be expected. Meantime these studies provide some useful pointers regarding technical and methodological aspects of a parallel study on outdoor advertising formats.

Two more recent reports regarding the use of animation in web advertising are of interest despite having no direct link to outdoor advertising.

In the first of these Lee and Ahn (2012) addressed the question of why the results of studies of animation in Web advertising were generally poor. The following is a quotation from this paper (with references removed):

> **One of the most popular attention-grabbing tools employed in Internet banner ads is animation, which is known to make objects salient and stimulate higher levels of user involvement. Paradoxically, animation may alert Internet users to the location of a banner ad, triggering ad avoidance behaviour. In addition, animated ads are known to require more of the reader's cognitive resources than static images, resulting in weaker memory performance. A number of studies have shown that animation in banner ads is not an effective tool. According to these studies, animation either does not affect memory or worsens it.**

Lee and Ahn recorded "attention data" (i.e., eye-tracking measures – fixation frequency and duration) from participants viewing the Internet at their own pace in a natural setting. They argued for the use of a self-controlled exposure setting procedure as providing more reliable results than forced-exposure settings, especially in the case of advertising in such very-low-involvement situations as when banner ads are used. Participants were asked to read news items via Internet pages on which banner ads appeared. These ads were static or with animation (the speed of which also varied – four vs.10 per four seconds). Animation affected eye fixation frequency and duration significantly – however, it should be noted that static ads attracted more fixations and for longer durations. Although no effect of animation on subsequent ad recognition was reported there was an overall carry-over from fixation frequency to memory. Animation speed did not have an effect on attention. The authors link the effects of animation to ad avoidance: "we found that animation, the Internet's most popular attention-attracting tool, drives user attention away".

The last study to be considered is arguably the weightiest from a theoretical and methodological point of view. Their review of the literature is extensive and useful, drawing attention to the mixed nature of the evidence on animation, for example. For this investigation Simola et al. (2011) designed Web pages with static and animated advertising images adjacent to a text section that the participants were asked to read. Contrary to some previous studies they found that online ads are not ignored when reading or browsing. The findings were a complicated mix of task (reading vs. browsing) and ad location (above vs. to the right) and the details are not of interest here. However, animation did distract reading and abrupt onset of advertising captured attention during reading. One key point about the study is that a complex set of options within an advertising brief were differentiated empirically using eye-tracking. Another is that the study was driven by thinking derived from mainstream academic research, a trend that we can only applaud and echo relative to the application to advertising topics such as visibility and attention, and in doing so quote the following:

> **From the theoretical perspective, the present study demonstrated that findings from basic attention research and the theories of visual attention derived from them, especially research pointing to the significance of visual saliency (Itti & Koch, 2000), the distinction between bottom-up versus top-down control of attention (Theeuwes, 1994; Theeuwes & Burger, 1998), the principles of the central capacity theory (Kahneman, 1973), and the research related to abrupt stimulus onsets (e.g., Theeuwes et al., 1999; Yantis & Jonides, 1984) can be successfully applied also to the Web environment when the influences of advertising are studied. This is relevant, as arguments to the contrary have been made (Diaper & Waelend, 2000; Zhang, 2000).** (Simola et al. 2011, p189)[6].

C: Applied research: Billboards and safety issues

Much of the research on digital billboards has been carried out in the US, and has focussed on safety concerns for drivers (in contrast and perhaps not surprisingly in light of the preoccupation with traffic safety, we discovered no research on digital billboards as viewed by pedestrians). The majority of the investigations made use of billboards with static images, many at actual roadside locations, often converted from conventional structures. Such evidence – because it is focussed on safety issues – has to be treated with considerable sensitivity if lessons are to be derived about the visual attention value of digital billboards. Frankly this is not within the remit of the research, however, it is possible to draw inferences about matters of interest for this review.

1. Tantala and Tantala (2007)

Among the most prominent contributors to the debate are Tantala and Tantala in a series of studies conducted under the auspices of the US organization FOARE – the Foundation of Outdoor Advertising Research and Education (part of the US Outdoor Advertising Association).

In the first of this series Tantala and Tantala (2007) reported a comparison of digital and conventional billboard panels, with the brief in particular to consider safety matters. The main conclusion of the study was that safety was not compromised by the presence of digital billboards. While the authors did not draw inferences about visibility *per se*, it is a simple step to make the connection between accident rate, driver distraction and increased visibility.

The data analysed were culled from traffic collision reports. Post hoc accident analyses are criticised as being prone to under-reporting of the true incidence of safety-related episodes and the full state of affairs regarding many reported accidents. There are various other concerns about the methodology used by Tantala and Tantala (see Wachtel 2009). For example, they excluded accidents near interchanges on the grounds that this is where drivers are coping with additional tasks such as lane changing. But this is just where an extra distraction such as a prominent billboard might increase risk by diverting attention from actual safety threats.

The core question addressed by the study was "Are accidents more, less, or equally likely to occur near digital billboards compared to conventional billboards?" For the purposes of what was termed a "temporal analysis", a comparison was made between traffic accident rates before and after the introduction of seven digital billboards (converted from conventional structures) on highways in Ohio (where there were 131 conventional billboards on similar roads). Changes in the digital displays occurred every eight seconds and there was no animation, no flashing lights, scrolling or video. For the purposes of a "spatial analysis" accident rates in the vicinity of the billboards were analysed together with other measures including the distance from a given accident site to the nearest billboard. The core question above raises a key statistical/logical point, namely that equality of outcome cannot be demonstrated statistically. The classical formulation of statistical hypothesis testing requires the specification of a null hypothesis that may be disproved, but cannot be proved: https://en.wikipedia.org/wiki/Null_hypothesis

Interestingly the incidence of traffic accidents did not increase after the conversion to digital, at least in numerical terms. Strictly speaking, basic statistical assumptions (regarding sampling, independence of observations, etc.) could not be met by the data for a before vs. after comparison, hence the comparison could not be formally evaluated, although this problem was not addressed. Accepting the results notwithstanding such objections, the data suggested a small fall in the traffic accident rate. On the other hand no data were given for conventional billboards for the before vs. after contrast, precluding a parallel assessment with digital billboards, which might well reflect any global chronological change in accident rate and would be shown in both data-sets.

The following is an excerpt from the report:

"The number of accidents within the (billboards') visible ranges for one year was 174 accidents for an estimated 85 million vehicles that drove by; this represents one accident for every 481,000 vehicles. If we exclude statistical bias (accidents from known causes), there are only 53 accidents in the year after the … signs were converted for 85 million vehicles; this represents one accident for every 1.5 million vehicles. The values per sign suggest an average of 7 accidents near a digital billboard per year for the same 85 million vehicles; this represents a rate of one accident per 12 million vehicles per

[6] The articles cited in the quotation are not included here.

year. Comparing a year before and after, the peak number of accidents on any given month decreased from 247 to 174, after the introduction of the digital billboard at the location; the peak number on any given month decreased from 14 to 8. Similar results were obtained for the longer 36-month windows. Based on the data and analysis, no significant change in accident occurrences can be attributed to the conversion of these billboards to digital format."

For the "spatial analysis" the Tantala study computed various correlations between accident-density rate (number of accidents per mile marker) on the one hand and billboard density (billboards per mile), viewer reaction distance (the distance in which drivers have time to notice/react to a billboard in the their field of vision), and proximity to the billboard (distance from mile marker to the nearest billboard). Tantala and Tantala used correlation coefficients to show how well the data "compared". As before the analyses were run excluding instances such as those where there was a known accident cause. Summary statistics did not include how many observations the correlation coefficients were based on. The correlations were small, generally below 0.2, and the conclusion was reached that there was no statistical relationship between vehicular accidents and billboards (including conventional and the seven, digital billboards). Unwisely the authors state (p3, p99) that the correlations "strongly suggest no causal relationship between the billboards and vehicular accidents", a remark that flies in the face of the conventional understanding that correlational evidence does not per se support an interpretation of causality. The investigators know this, witness their comment (p81) that "It is important to note that correlation is not necessarily causation, even though it may be an indicator". Their trenchant conclusions elsewhere are completely at odds with this comment.

Several procedural difficulties may also be identified: there is no indication of what statistical tests were applied (though a reasonable guess would be that some form of t-test was used); there is the curious reference to statistical bias attributable to accidents from known causes (hence the exclusion of related data); and there are no baseline data for other conventional billboards (before or after the conversion of the about-to-be digital ones). It is not surprising that the study attracted negative reviews (especially Wachtel, 2007).

2. Tantala and Tantala (2009a, 2009b, 2010a, 2010b)

These authors have added to their 2007 study on behalf of FOARE a series of four more studies of traffic and crash data (listed in a 2010 report on digital signage for OMA - the Outdoor Media Association of Australia), including an update of their 2007 investigation. Essentially the same method was used in each case, and this was an extended version of the 2007 method. The data sets in each case were traffic and crash data at billboard sites analysed with respect to the occurrence of crashes around the billboards before and after having been converted to digital (temporal analysis) and correlations between crash data and distance measures for digital billboards (spatial analysis). This entailed a broadening of the scope of the upstream and downstream distances around the billboard sites and the use of a wider set of accident indices and traffic measures. The findings of all four studies were very similar, notably that crash rates near digital billboards decreased at all distances over the several years for which data were analysed, and that crash rates did not increase when the billboards were converted. What was omitted from these studies, as from the 2007 study, is parallel evidence concerning traditional billboards. This is crucial since the results are hard to interpret without an indication of long-term trends in accident rates. It is of interest therefore that traffic fatalities (per 100,000) in the US showed a year-on-year decline in the eight years 2002 to 2009 with one small upturn. The following table is an extract from http://en.wikipedia.org/wiki/List_of_motor_vehicle_deaths_in_U.S._by_year:

Year	Fatalities per 100 million miles travelled	Fatalities per 100,000 population
2002	1.51	14.95
2003	1.48	14.78
2004	1.44	14.63
2005	1.46	14.72
2006	1.42	14.31
2007	1.36	13.70
2008	1.26	12.31
2009	1.15	11.05

The trends in accident statistics such as these show why appropriate baseline data are needed to interpret the reported findings of the Tantala and Tantala series of investigations.

If, setting aside the trenchant criticisms of this research approach by Wachtel (2009), one reaches the conclusion that safety is not affected by the introduction of digital billboards, then by the same token questions about their advertising effectiveness may be raised. If the conclusion were otherwise, that safety is compromised, then it would seem reasonable to infer that visual attention was affected. On the other hand it remains possible that safety is not compromised but attention is safely attracted. The problem is that conclusions, one way or another, are not safe for the reasons described.

A scathing critique of the 2007 research was made by Wachtel (2007), partly on the grounds of the bias intrinsic to this and the following study. If the conclusions of the studies are indeed on the conservative side, it would seem to follow that digital billboards are more visible than their conventional counterparts.

3. Lee, McElheny and Gibbons (2007)

This is another major study sponsored by FOARE for the US outdoor industry but using a very different methodology dependent critically on eye tracking. Eye movements were recorded for 36 drivers passing digital and conventional billboards. A range of sites were used on a 50 mile route, including a mere five with digital billboards and another 15 with conventional billboards; there were also empty sites for comparison and others described as comparison sites (a set of locations with what the authors considered similar properties to the target sites). Measures of eyes-on-road and number of glances were obtained for the eight seconds prior to each site. Driving performance measures were also recorded unobtrusively. While there were no differences in the frequency of glances at the two billboard types, these glances tended to be longer towards digital billboards, which is commonly and not surprisingly thought to be a safety threatening state of affairs.

The results for the four different types of event (or site) indicated that there were no differences between percent eyes-on-road or glance frequency measures. Digital billboards were glanced at no more frequently than conventional billboards or anything else, but the glances on them tended to be longer than at baseline sites or conventional billboards; and to complicate the matter this applied too to the "comparison" sites. This confusing nature of this picture was mirrored by the summary data for the driving measures. There was little or no consistency in the two sets of data patterns, so the report was unable to establish any clear differences or trends as a function of site type.

The study included a small-scale preliminary examination of night-time conditions. Descriptive data suggested that digital billboards and the "comparison" sites attract more frequent and longer eye-glances. To this extent there was some agreement with the daytime results, in that comparison sites are linked with digital billboards as having gained more attention. The authors concluded that the differentiating feature of both types of site was their intrinsic lighting characteristics. Their conclusions were of course framed in terms of safety concerns, and the report underlined the possibility that some panel sites may degrade driving performance.

4. Chattington et al. 2009

This was a substantial study, using a driving simulator in conjunction with an eye-tracking system, by the UK's Transport Research Laboratory. Fourteen video billboards were compared with an equal number of conventional billboards positioned on a simulator "drive" while eye movements and driving performance were monitored. The advertising (48 sheets) was presented to the left, centre or right of the road, and there was a fourth arrangement with panels in all three locations. Exposure duration (contrived by varying the occluding building structures) varied between two, four and six seconds. "Glances" at video billboards were more frequent and lasted longer than at standard billboards. The right side billboards received fewest glances and the "all three" received the most. Results were also reported on lateral lane control, driving speed and deceleration when approaching the video billboards; in all cases the outcomes were interpreted negatively relative to safe practice. The impairment of driving in the simulator was compared with the ingesting of cannabis and writing/sending a text message.

D: Miscellaneous studies, including technical innovations

Driving simulator/Virtual reality research

See Chattington et al. (2009) in the previous section. This study had an immediate forerunner in a study by Young and Mahfoud (2007).

In the preamble to our dynamic imagery research, a range of technical platforms/methods were briefly reviewed. This included the driving simulator and related approaches using virtual reality. In fact the virtual reality approach was mooted from time to time as a possible vehicle for visibility research, and the driving simulator method can be seen as a variant on this. Both could support dynamic imagery and could be combined with eye-tracking; both also enable poster panels to be inserted in the virtual/simulator experience. After reflection about the pros and cons, there was little enthusiasm for such approaches. The driving simulator approach was brought back into contention by a report on advertising and driver distraction (Young and Mahfoud, 2007). The illustrations and task description from that report indicate the quality of the driving experience able to be delivered by then affordable technical platforms, and the technology has advanced apace. This is not the place for a full critique of the study, but it is worth noting that the scenarios depicted were rather simplistic, even cartoon-like. The Chattington et al. (2009) in the previous section made use of a driving simulator with more convincing visuals and mechanics. The TRL Driving Simulator is "one of the most advance simulators in the UK. This entails an electrically driven motion system to provide "limited motion in three areas (heave, pitch, and roll)". Importantly this provides the driver with an impression of driving-like acceleration and vibration. The importance of replicating the kinaesthetic and related movement features of driving should not be underestimated. State-of-the-art simulators for training pilots demonstrate what is possible to represent the visual world and the dynamic experience of flying; such facilities exist on university campuses in the US (e.g., University of Iowa).

Such methods would seem likely to have a lengthy set-up time unless an existing facility (e.g., a high-quality commercially available simulator with good technical support) was available. A virtual reality approach would raise concerns about the tolerance of respondents to the experience; in the Chattington study eight out of 55 participants could not complete the testing because they experienced simulator sickness, a form of motion sickness.

Another concern is the quality and realism achieved in the visual presentation. Illustrations in reports are not exactly reassuring, as exemplified by the two reports cited here. There is a considerable potential advantage in enabling the implementation of new formats and properties (cycle rate, transition type, etc.). On the other hand the visual experience can be tailored to the ends of the researcher/client, with a spin on the results that is harder to engineer using a real world based approach.

A novel view of eye movement measurement arising in the context of a study on the effect of motion

The article by Mital et al. (2010) (see Section A) was discussed earlier because of its interest as an investigation of attention during the inspection of dynamic scenes. It is also notable by virtue of its novel approach to the "parsing" of eye movement data, considered to be required by the occurrence of "pursuit" eye movements when the eyes track an object in a moving scene. Traditionally the eye-movement data are parsed into a fixation-saccade sequence, a method the authors contend is inappropriate for stimuli that move relative to the observer. In this case they adopt a classification algorithm that identifies in the raw data a group of eye movements they call "foveations". The constituents of such a grouping are marked as eye blinks, saccades, or non-saccadic eye movements, that is, fixations, smooth pursuit movements and corrective optokinetic nystagmus. This is a highly technical business that we have not come across elsewhere[7] but seems to be a convincing method of preserving the movement contained in the gaze data, as well as providing a more accurate account of what is being inspected by the eyes.

Future research using eye tracking to investigate the viewing of dynamic scenes will need to take account of such advances in technique. The study by Mital et al. was about the viewing of videos by observers who were stationary (and saw the videos on a viewing screen with what appears to be a relatively stabilized chin and headrest (albeit described as "unrestrained"). It is not clear what else might be required to achieve an equivalent and effective monitoring of observers themselves on the move. Further developments - and accessible commentaries - on the technical options to become available will be of great interest. One issue that is not mentioned by Mital et al. is the extent to which their findings would correlate with those based on a traditional saccade-fixation classification.

It should be noted that substantial developments on the related problem of object tracking do not seem to have taken place. A solution to that particular problem does not seem to have entered the public arena. It is apparent that TV companies have technologies that allow them quite swiftly to respond to recorded events in which it is necessary to mask faces or items by which individuals may be identified (e.g., car registration plates). The ability to track moving objects therefore exists as a practical achievement but is yet to be brought into contention as a technical resource for eye tracking.

[7] Except that this may be linked to technical developments in imaging and eye-tracking which enable a display to be presented to an observer at high resolution only at the point of gaze and at lower quality elsewhere (see http://live.ece.utexas.edu/research/realtime_foveation/index.php)

Chapter 6

Conclusions and recommendations

The literature review has covered four main topics. It began with a survey of what is principally academic research on visual attention, and the factors related to movement that impinge on the deployment of attention. This research depends on a variety of pared-down experimental paradigms, often requiring the participants to search a display or scene, or otherwise to focus their attention on a task alongside the manipulation of extraneous (e.g., visually peripheral) sources of attention that may sway or distract the performance of the focal task. Important exemplars of the methodology employed were presented, together with sample findings and the conceptual issues that arise in their explanation. Evidence from a variety of experimental approaches was considered.

Research on attention in the context of web-based advertising was reviewed via a sample of the most pertinent studies. A much more substantial body of evidence is on safety concerns about electronic billboards. The literature is fraught with dispute and closure seems no nearer that the universal switch to driverless cars. One important aspect of the research that should not be overlooked is that methods and technical platforms for research are continually being refined and extended. In addition to pointers distributed through the review some relevant developments were addressed in a short final section.

The research paradigms used in Phases One, Two and Three were to a degree informed by research elsewhere but they did not address issues in the mainstream of basic research. Hence it was not possible, for example, to draw conclusions about the relative size of the effect on performance due to movement or change of state (onset/offset). This distinction, and other issues raised in the review, should be taken into account when planning further empirical studies.

References

Abrams, R. A., and Christ, S. E. (2003). Motion onset captures attention. *Psychological Science,* **14**, 427-432.

Abrams, R. A., and Christ, S. E. (2006). Motion onset captures attention: A rejoinder to Franconeri and Simons (2005). *Perception and Psychophysics,* **68**, 114-117.

Al-Aidroos, N., Guo, R., and Pratt, J. (2010). You can't stop new motion: Attentional capture despite a control set for colour. *Visual Cognition,* **18**, 859–880.

Barber, P.J. and Sanderson, M.W. (2005). Video analysis of driver eye-behaviour. Postar Technical Report: June 2005 (updated 2006).

Boot, W. R., Kramer, A. F., and Peterson, M. S. (2005). Oculomotor consequences of abrupt onsets and offsets: Onsets dominate oculomotor capture. *Perception and Psychophysics,* **67**, 910- 928.

Brockmole, J. R., and Henderson, J. M. (2005). Object appearance, disappearance, and attention prioritization in real-world scenes. *Psychonomic Bulletin and Review,* **12**, 1061-1067.

Burke, M., Hornof, A., Nilsen, E., and Gorman, N. (2005). High-cost banner blindness: Ads increase perceived workload, hinder visual search, and are forgotten. *Transactions on Computer-Human Interaction,* **12**, 423-445.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research,* **51**, 1484-1525.

Chattington, M., Reed, N., Basack, D., Flint, A., and Parkes, A. (2009). Investigating driver distraction: The effects of video and static advertising. TRL Published Project Report PPR409.

Crundall, D, van Loon, E., and Underwood, G. (2006). Attraction and distraction of attention with roadside advertisements, *Accident Analysis and Prevention* **38**, 671-677.

Donk, M., and Theeuwes, J. (2001). Visual marking beside the mark: Prioritizing selection by abrupt onsets. *Perception and Psychophysics,* **62**, 891-900.

Donk, M., and Theeuwes, J. (2003). Prioritizing selection of new elements: Bottom-up versus top-down control. *Perception and Psychophysics,* **65**, 1231-1242.

Folk, C. L., Remington, R. W., and Wright, J. H. (1994). The structure of attentional control: Contingent attentional capture by apparent motion, abrupt onset, and color. *Journal of Experimental Psychology: Human Perception and Performance,* **20**, 317-329.

Franconeri, S.L., and Simons, D.J. (2003). Moving and looming stimuli capture attention. *Perception and Psychophysics,* **65**, 999–1010.

Hillstrom, A. P., and Yantis, S. (1994). Visual motion and attentional capture. *Perception and Psychophysics,* **55**, 399-411.

Hong, W., Thong, J. Y. L., and Tam, K. Y. (2004). Does animation attract online users' attention? The effects of flash on information search performance and perceptions. *Information Systems Research,* **15**, 60-86.

Kaufman, L., and Richards, W. (1969). Spontaneous fixation tendencies for visual forms. *Perception and Psychophysics,* **5**, 85-88.

Kawahara, J., Yanase, K., and Kitazaki, M. (2012). Attentional capture by the onset and offset of motion signals outside the spatial focus of attention. *Journal of Vision,* **12**, 1-13.

Lee, S.E., McElheny, M.E., and Gibbons, R. (2007). Driving performance and digital billboards. Blacksburg, VA, USA: Foundation for Outdoor Advertising Research and Education.

Lee, J.W., and Ahn, J. (2012). Attention to banner ads and their effectiveness: An eye-tracking approach. *International Journal of Electronic Commerce,* **17**, 119-137.

Matin, E. (1974). Saccadic suppression: A review and an analysis. *Psychological Bulletin,* **81**, 899-917.

Matsukura, M., Brockmole, J.R., and Henderson, J.M. (2009), Overt attentional prioritization of new objects and feature changes during real-world scene viewing. *Visual Cognition,* **17**, 835-855.

Mital, P.K., Smith, T.J., Hill, R.L., and Henderson, J.M. (2010). Clustering of gaze during dynamic scene viewing I predicted by motion. *Cognitive Computation,* **3**, 5-24.

Posner, M.I., and Cohen, Y. (1984). Components of visual orienting. In H. Bouma, and D. G. Bouwhuis (Eds.), Attention and performance X (pp. 531–556). Hilldale, NJ: Erlbaum.

Pratt, J., and Trottier, L. (2005). Pro-saccades and anti-saccades to onset and offset targets. *Vision Research,* **45**, 765-774.

Pratt, J., Theeuwes, J., and Donk, M. (2007). Offsets and prioritizing the selection of new elements in search displays: More evidence for attentional capture in the preview effect. *Visual Cognition,* **15**, 133-148.

Simola, J., Kuisma, J., Öörni, A., Uusitalo, L., and Hyönä, J. (2011). The impact of salient advertisements on reading and attention on Web pages. *Journal of Experimental Psychology: Applied.* **17**, 174-190.

Skarratt, P. A., Cole, G. G., and Gellatly, A. R. (2009). Prioritization of looming and receding objects: Equal slopes, different intercepts. *Attention, Perception and Psychophysics,* **71**, 964-970.

Sunny, M.M., and von Muhlenen, A. (2011). Motion onset does not capture attention when subsequent motion is "smooth". *Psychonomic Bulletin and Review,* **18**, 1050-1056.

Tantala, A.M., and Tantala, M.W. A Study of the Relationship between Digital Billboards and Traffic Safety in Cuyahoga County, Ohio," Submitted to: The Foundation for Outdoor Advertising Research and Education, July 7, 2007.

Tantala and Tantala (2009a, 2009b, 2010a, 2010b). Reports cited in: Digital billboards and road safety: An analysis of current policy and research findings. Discussion paper prepared for OMA (Outdoor Media Association), Australia.

Von Mühlenen, A., and Lleras, A. (2007). No-onset looming motion guides spatial attention. *Journal of Experimental Psychology:* Human Perception and Performance, **33**, 1297-1310.

Wachtel, J. A critical, comprehensive review of two studies recently released by the Outdoor Advertising Association of America. Prepared for Maryland State Highway Administration, October 18, 2007.

Wachtel, J. (2009). Safety impacts of the emerging digital display techology for outdoor advertising signs. Report for AASHTO (American Association of State Highway and Transportation Officials).

Young, M.S., and Mahfoud, J.M. Driven to distraction: Determining the effects of roadside advertising on driver attention. Final Report to The Rees Jeffreys Road Fund. Brunel University, September 2007.

Acknowledgements

We are most grateful to Professor Anne Richards for providing access to the eye tracker in her laboratory, and for her help in providing technical advice regarding the Phase Two and Three investigations. We also thank Richard Hanson for advice and assistance in making the video clips.

Appendix: Visibility studies undertaken to date

Driver visibility study (1995-1996): OSCAR 2 measuring visibility hit rates of roadside panels, using infra-red eye-tracking methodology. Introduced the basic concept of visibility hit rates for poster panels. Modelled visibility in terms of panel size, eccentricity (offset from road) and distance. Respondents: drivers and passengers.

Maximum visibility study (1996-1997): assessing the furthest distance at which a panel can be seen with full concentration on the panel, using psychophysical methods.

Pedestrian visibility study (1998-1999): measuring visibility hit rates for poster panels in roadside and pedestrian environments, using infra-red eye-tracking methodology. Respondents: pedestrians.

Nottingham driver attention study (2000-2001): establishing how drivers' & passengers' attention is distributed down the road ahead – using real-world in-car eye camera technology. Respondents: drivers and passengers.

"Inclusivity" pilot (2002): comparing a set of active search methods as alternatives to passive eye-tracking methods (for speed, convenience and portability).

Wave 1 (aka Travel Wave) (2003-2004): using an active search method selected on the basis of the "Inclusivity" pilot to estimate hit rates for panels from transport media (buses, tube, rail, taxi). Respondents: pedestrians.

Wave 2 (aka Retail Wave) (2003-2004): using the active search method to estimate hit rates for panels in retail environments (supermarket car-parks, malls, pedestrian shopping precincts, petrol stations, telephone kiosk). Respondents: pedestrians.

Video analysis of driver eye behaviour (2004-2005): using video analysis of gaze data from Nottingham driver attention study to assess hit rates on roadside panels and buses. Respondents: drivers and passengers.

Pedestrian visual behaviour: walking speed and head-up study (2005): specifying key aspects of walking for use in pedestrian visibility modelling via literature searches and observational data.

Wave 3 (2006): using the active search method to provide supplementary data on panel hit rates in key transport environments (buses and tube). Respondents: pedestrians.

Wave 4 (2007-2008): using a passive eye-tracking method to estimate panel hit rates in key transport and retail environments, with contemporary roadside panels, providing an up-to-date database across environments with new eye camera technology. Respondents: drivers and pedestrians.

Wave 5 (2008): a passive eye-tracking method to update estimates of panel hit rates for telephone kiosks and taxis. Respondents: drivers and pedestrians.

Dynamic Imagery Research Phase 1 (2008-2009): Pilot study to explore technology for presenting moving images (scrolling displays) while recording eye movements. Respondents: unclassified.

Dynamic Imagery Research Phase 2 (2009): Investigation of effect of dynamic images (scrolling poster panels and bus panels) on hit rates, using a stationary view of the scene. Respondents: pedestrians.

Dynamic Imagery Research Phase 3 (2009-2010): Investigation of effect of dynamic imagery (scrolling and digital poster panels, and bus panels) on hit rates, using a dynamic view of the scene. Respondents: drivers and pedestrians.

Visibility of poster panels seen through bus and train windows (2010): using a passive eye-tracking method to estimate panel visibility when viewing through a bus or train window. Respondents: pedestrians and public transport passengers.

route

outdoor media analytics