

Factors affecting visibility distances for poster panels

Technical Report: POSTAR Ltd.

*Authors: P J Barber and A Dickenson, Department of Psychology, Birkbeck College,
London*

Abstract

The POSTAR visibility model was implemented with upper limits for visibility distances taken from its precursor (OSCAR) that had not been set on the basis of any systematic estimate of the distances at which posters can be seen. The present study aimed to review evidence on visibility and to devise and evaluate methods for measuring visibility applicable to posters to provide an empirical basis for setting maximum visibility distances (MVDs) for use in visibility modelling for POSTAR. In the light of the published research on visibility and legibility the existing MVDs need to be drastically revised. This was confirmed by the present experimental work which explored a range of methods for the purpose of investigating and setting maximum visibility distance parameters for the visibility model. Two methods were devised that produced the evidence on which recommendations are based for revised maximum visibility distances. In the first task the observers had to report the orientation of masked rectangular shapes and in the second they were asked to report using a numerical scale how easy it was to spot the target rectangle on each trial. Both tasks were implemented so that other potentially relevant variables could be examined, including target area and aspect ratio, to permit a mapping of the performance in the tasks on to realistic object shapes and distances of the size of posters. In both studies performance improved as a function of the logarithm of the target area. Estimates for MVDs were provided by both from the two methods and an estimate was also made on the basis of a standard sign legibility method. It is argued that poster viewing characteristically lacks intentionality on the part of the viewer, and that this key feature of the poster viewer's visual behaviour is best captured by the prominence judgement task. The most conservative estimates of MVDs supplied by the prominence judgement task draw on the upper reaches of the scale where the target is so obvious that it tends to jump out of the scene. Accordingly the recommended values for MVDs are 63, 183 and 259 metres for 6, 48 and 96 sheet panels respectively. It is evident that implementing these values in the POSTAR visibility model must increase the visibility value of all posters and some panel formats will be affected more than others. Some cautionary observations are made about the recommended distances. The distribution of attention at different distances ahead of a viewer (driver) is not known with any certainty, and it would seem sensible to incorporate weightings to reflect this in the visibility model pending further consideration of this factor. A provisional proposal for weightings is made in the report.

1. Introduction: The visibility problem and outdoor advertising

1.1. Distance as a limiting factor in visibility modelling

The visibility modelling component of POSTAR supplies a tool for estimating the visibility of roadside posters for passing drivers. The model was developed on the understanding that information about poster site parameters would be available in the site classification database. Factors that are plainly important in this connection are viewing distance, poster size, offset of the poster framework from the road, and the degree of competition and contrast of the poster relative to the rest of the visual field. These factors enter the model as parameters for the calculation of visibility estimates.

As initially implemented, the model operates computationally with distance (driver to site), poster size, horizontal displacement (road-centre to leading edge), and visual clutter (corresponding to the shopping vs. arterial vs. residential site classification of poster sites). Estimates are computed for all distances less than what is referred to as the maximum visibility distance (MVD), which is 30 metres for the standard 6 sheet size, and 100 metres for 48 and 96 sheet sizes. On the face of it the rationale for MVD as a limiting factor for the visibility model is simple enough: an object may be removed to such a distance from a viewer that it can no longer be seen; at and beyond that point the object may reasonably and for all practical purposes be counted as having zero visibility. The practical purposes include visibility computations in POSTAR.

The existing MVDs were taken from the OSCAR scheme and were not reassessed for the purpose of developing POSTAR. It is evident that changes in MVD settings would in general change visibility scores. It will be argued here that the only direction of change that makes sense from the standpoint of research or everyday experience is upwards. As MVDs for all panels would thus be increased, so visibility scores would increase across the board. The balance of any such changes would of course need to be considered.

1.2. OSCAR and maximum visibility distances: Practical and conceptual issues

In the exercise that preceded the specification of poster site visibility for OSCAR, the precursor of POSTAR, MVDs were decided on through the agreement of an industry committee. According to one account, possibly apocryphal in nature, there was a promenade around certain poster sites and the distances in current use were settled by a small group of experts considering when what was on the poster could be made out, and to some this seems to have been synonymous with *what was on the poster being legible*. It should be noted that this does not seem to have entailed consideration of *whether the poster was visible*, i.e., could be seen (given an unoccluded view of it of course). In the discussions about MVD, reference seems also to have been made to the distance used for assessing a driver's vision in the UK Highway Code, which is 25 yards. There appears to have been no justification in the light of existing research, and certainly no research tailored for the purpose, and the discussion proceeded with a rather ill-defined notion of visibility. In the event, it was decided by inspection that a 6 sheet poster could be made out at 30 metres, and that

this would be true for 48 sheet and 96 sheet posters, the two larger formats of main concern, at a distance of 100 metres.

No criticism is intended of the process whereby MVD values for OSCAR were arrived at, indeed there may be witnesses to challenge this rather loaded account. Irrespective of historical accuracy, this industry lore serves to highlight a number of key aspects of the problem of visibility and the issues it raises. The first question is whether or not the agreed values are adequate now, and whether there is a basis for defending them on the grounds that they are still adequate measures of the upper limits of poster visibility. Accepting that they should at least be checked because of the introduction of POSTAR, the second question is raised as to a method of doing this. Other issues follow, such as what might be the global effect on the output of the model of changing MVDs, and how might the balance of current site visibility values be affected. These issues relate to policy matters, and while these matters are not the province of this report, this account could not be written in total ignorance of them.

Finally, and very crucially, there remains the fundamental issue as to what is to be understood by “visibility”. Plainly if there is a view that it is the same as “legibility”, then the implications need to be made clear and the view be countered to the extent that it is mistaken. To this end a better articulated conception of visibility will need to be developed, and although this is not the place to do so in a comprehensive fashion, there needs to be a refining of the concept that is in use and this report will take steps in this direction too.

2. Research on visibility

A brief discussion of previous research is included, more as a scene-setting exercise than a comprehensive review. A more complete account should attend the formulation of an improved conceptual model of poster visual behaviour, that will include a more satisfactory conception of visibility mentioned above.

2.1 Sensitivity of the human visual system: Basic data and measurement techniques

A stock chapter in any textbook on visual perception will be found to contain statements about humans’ remarkable sensitivity to visual stimuli. Clinical and laboratory specifications of *visual acuity*, which is particular interest in the present context, are measured in different ways and expressed in different terms, but they both assess the resolving power of the eye. The clinical method will be familiar, depending as it does on the use of charts of letters of different sizes as used by opticians. It turns out that most people tested can discern two parts of a target (like the gap in a C shaped figure) separated by one minute of arc at 20 feet. This is the origin of the idea of 20/20 vision - someone possessing this can see at 20 feet what it is found most people can see at 20 feet. If another individual needed the two target elements to be separated by two minutes of arc when presented at 20 feet, which is equivalent to seeing the elements separated by one minute at 40 feet, then their vision is rated as 20/40 (this is the visual requirement for a driver’s licence in much of the US). Letters are used for clinical testing for various reasons, including the ease of

indicating whether one is seeing or not by calling out the name of the letter. Problems with the method are the bias that results simply because letters are familiar so the subject may be able to guess without really seeing properly, and its unsuitability for testing illiterate subjects.

Laboratory methods (and there are several) solve this problem, but introduce various complications, some of which have to be contended with for present purposes. One finding that astounds most readers is that the least values for visual resolution for a fine line may be as little as 0.5 seconds of arc, which is equivalent to a telephone wire viewed against a clear background at a distance of 2 km (Schiff, 1970). This is not really a proper basis for thinking about poster visibility - to begin with, posters are not long thin lines. Some more relevant comparisons are as follows: a letter on this page is about 12 minutes of arc if viewed at about arms-length (say 50 cm), the moon is about 30 minutes of arc, and a one pound coin subtends about 13 minutes of arc at 5 metres. But these figures are also misleading if we have in mind the one minute of arc standard. This "standard" is for an observer looking purposefully for a target, asked merely to discriminate some property of it (like whether its boundary contour contains a small gap), and doing the task under very favourable conditions. The poster viewer is generally doing none of these things. He or she is unlikely to be asked to look deliberately for a billboard in the distance at the limit of his or her vision, and under optimal conditions. If this were to be done then plainly MVDs would need to be increased dramatically. As will be seen in the next section there is a difficulty too in taking legibility, the ability to read letters on a display, as a working definition of visibility. On the other hand if some visual limit is to be set that corresponds to "maximum visibility distance", then something like the tasks used to assess acuity will need to be adopted.

2.2 Research on visibility of road signs

Forbes (1972) summarized a substantial body of research on what he referred to as visibility and legibility of road signs. The fact that he made this contrast is significant because it echoes a distinction that has been made in discussions of poster visibility, and that perhaps underlay the deliberations leading up to the decision about MVDs for OSCAR. The definition of legibility was the ability to read the letters on a sign. Visibility was considered to refer to attention-getting characteristics. Forbes emphasized that visibility did not mean greatest detection distance; "the practical problem for the traffic engineer is whether the sign attracts enough attention to be read when it is within legible distance". Indeed Forbes also used the term attention value as an alternative to visibility, and there is some merit in this usage (another term with the right kind of resonance is "conspicuity") given the tendency for discussions to treat visibility as synonymous with legibility. He analysed visibility/attention value into target value and priority value. Target value is determined by characteristics making a sign stand out against its background, and visual contrast is the main factor of interest here. Priority value is affected by location and mounting position along the highway, number of signs, and other psychological factors (e.g., search procedure, reading habits, and the mental set of looking for a target).

As noted above there have been echoes of much of this in discussions about poster visibility. Legibility is plainly to be considered as an upper limit of some kind

applying to circumstances where reading letters/text is a concern. Not every advertisement carries text, so some caution is needed in this connection, and this will be returned to below. A standards setting study was conducted under “full-scale outdoor observation(al)” conditions, for the US Bureau of Public Roads by Forbes and Holmes (1939); they used 6-letter words printed in black-on-white letters of various sizes and the observers viewed them at different distances. The legibility of these stimuli was determined, and Forbes reported that a legibility distance of 50 feet per inch of medium wide letters corresponded to the 20/20 vision acuity standard; 50 feet/in is a somewhat conservative value in the light of other studies reviewed by Forbes. What might this imply for MVDs for posters? Consider the example of a 6 sheet poster bearing a single six letter word, spanning more or less the full horizontal extent of the poster, so that each (square shape) letter is about 10 inches high. This word would be legible at 500 feet or 152 metres. It goes without saying that 48 and 96 sheet legibility distances would be much greater, and all three are well above current settings. However, just how to base any distance calculations for these two formats is not obvious since both are more elongated than a 6 sheet and at the same time they are normally encountered landscape style. A rough estimate of their legibility distances (assuming six-letter words in a 36 and 72 inch square font as in Forbes’s research respectively) would give legibility distances of 547 and 1094 metres. Of course advertisers do not typically use posters in this fashion (however, see recent campaigns by Open Direct). No doubt as a way of setting MVDs this method would be found wanting by the poster industry. In passing it is worth noting that “glance legibility” measurements, giving the observer about one second to view the sign, reduced the legibility distances by up to 15%; about three to four short familiar words can be recognized in this short time.

What of visibility as defined by Forbes and others in the traffic sign research community? First, it follows from the legibility distance studies that a sign positioned at a distance less than the legibility distance must be capable of being seen, and hence in an obvious sense of the term, it is visible. But it may not have attention value. In the case of traffic signs, this is principally determined by two contrast factors; the brightness ratio of sign-to-background C_{SB} , and the brightness ratio of letter-to-sign C_{LS} . For a sign with the legend occupying about half of the area, the visibility distance in a specially designed driving-type task was predicted by the average of these two ratios multiplied by the expected recognition distance of the sign (ER). The latter is not explained fully by Forbes but is itself specified as 1200 x the small dimension of the sign in feet. Whether or not this calculation should be applied to poster visibility is a moot point, however two points should be noted. First, the research draws attention to factors that affect a sign’s attention value, and this includes brightness contrast between the sign and its surrounds (to a degree a matter of site selection and design) and between the sign and its contents (essentially this is a matter of creative design). Second, the sign may be visible yet not be legible, it may catch the eye before its contents are discerned. It is a matter of observation that motorway signs are visible well in advance of the point at which one can read off information from them. The advertiser is arguably best served by a poster that catches attention as early as possible, that puts down a marker for further visual attention, as it were. Although it seems therefore that the notion of visibility from traffic sign research may not be directly adaptable for the purpose of the setting MVDs for posters, there is much to be learned from this research.

3. Poster visibility

3.1 Methodological options

The most obvious method for assessing the point at which an object ceases to be seen is to move it towards or away from the viewer, and to record this point over a series of runs. Equivalently one could move the viewer towards or away from the object, and this is a very satisfactory way of getting a feel of the distances which are being mentioned in this account. This type of method belongs to a family of techniques from the study of psychophysics. The critical distance at which the object is detected can be thought of as a threshold of perception.

A notional experiment that could serve present purposes would need to take account of many factors. One complication is of course what object to use, and where to conduct the experiment. In fact the method was under serious consideration, for example using an airfield as the site for an experiment to give an adequate range of distances in an uncluttered environment. Although the method was not adopted, it is a useful illustration of the practical problems that have to be solved. Many of these are not insoluble (like recruiting and transporting observers to the site, but some need careful thought. The aspect ratios of 6, 48 and 96 sheet panels are obviously different and the same goes for the orientation of their longer sides. So, what shape of panel should be used? It might be agreed that a square or circular panel should be used and one could be constructed for the purpose of an experiment. How big should it be? Clearly this ought not to matter if it is only the size of the object's image on the viewer's retina that matters. But this would certainly need to be checked (indeed target size was a factor in the road sign visibility research discussed above). Moreover there are many other factors to take into consideration, such as the contrast between the target and the background, the illumination and texture of the visual environment.

The task for the viewer would need to be devised with some care. The verbal reports of an observer need to be checked, and one might contemplate a procedure in which the target is sometimes omitted, in the manner of a "catch" trial, simply because it is easy for people to persuade themselves that the target can be seen even when it is not there. One could probably exclude the likelihood of observers guessing or dissembling in such an experiment but it does raise questions about individual differences and observer biases in research of this type. It was decided not to proceed with the study in this form for practical reasons, but all of the issues it raised are present for other approaches to the problem.

Before considering the solutions that we finally adopted, there is another method to consider that is a potentially very powerful alternative. The problem of interpreting observers' voluntary responses to the question of what they can see was solved in the POSTAR visibility experiment by using eye tracking to ascertain what they actually looked at. Eye movement recording ought to be adaptable for the purposes of setting MVDs. Indeed eye tracking has been used in research to measure the visual acuity of new-born infants. This research tapped into the natural tendency

of the eyes to follow a moving object. On the assumption that if a moving object cannot be seen then the eye will not follow it, the infants' eye movements were recorded for a range of stripe patterns. The width of the stripe capable of initiating visual following was taken as an index of acuity. Plainly this method itself is not applicable for present purposes, however, it would be possible to present pictures containing poster panels at progressively greater distances and to assess the size at which the panel was no longer fixated. This is not quite so straightforward as it seems because the threshold here would need to be the size at which the panel was not fixated more often than would be expected on the basis of chance. The method remains a viable option, indeed it has many compelling features, but it should be recognized that there are many procedural problems to overcome and a problem of statistical logic to contend with. Setting aside the procedural problems for the present - they do not appear to be insuperable - one can focus on the statistical issue. The difficulty is to know when the eye fixation data for an observer is random, and he/she can be considered as not having seen the image. In essence this is not far removed from the statistical (mal)practice of trying to prove the null hypothesis. Suppose we have a set of measurements that are characterized naturally by a degree of random variation, and we wish to make a decision as to whether a measure taken under some circumstances belongs to the same population as our base set. The common practice is to set sufficiently stringent membership limits, typically specified by the *improbability* of the target value being obtained from the base population by chance. In other words, it would be an unreasonable person who did not accept that the target measure was too extreme to be counted as belonging to the base population. A decision that is permitted is to reject the "null hypothesis" (of common statistical membership) given a sufficiently extreme target measure, the limit for the decision is one of extreme probability say of 99% or more. If the target measure is less deviant than this, the matter is to be left open, we are not entitled on this logic to conclude that the null hypothesis is true, that the target measure does come from the same population. Even if we were willing to do so, the problem is that these limits may be impractically large.

There are several other problems that militate against the ready application of this method. It is very likely that the target will need to be very small for it to be not visible in the usual sense of the term. This will obviously apply to an eye tracking approach to the matter too, an object would need to be very small on screen to fail to attract any fixations at all. Eye movement recording for very small objects necessitates very careful calibration and measurement and plainly there are limits on the accuracy of such recording methods anyway. The method was considered initially to be unsuitable for present purposes for a variety of other reasons, but was retained as having some future potential. A flavour of the problem can be gauged from the fact that there were never in a 6 second viewing of the POSTAR scenes many more than about 20 fixations. If the panel size were to be 5% of the total screen size (which is actually a considerable overestimate), then one would expect just one hit on the panel if scanning were completely random (a better account would be to reflect the actual distribution of fixations in estimating chance hit rates because people have quite regular scanning patterns, rarely visiting the "corners" of a scene). In any event the 1/20 hit rate is what fixation rate would have not to exceed for viewing to be considered to be random. A further methodological obstacle is - as usual in the context of POSTAR - the question of what to do about creative content. It is evident

from the work on road sign visibility that contrast is important, and no doubt so too would be colour contrast. Even at a great distance, a bright orange blob of colour might well attract more fixations than a grey or green one.

To adopt the kinds of acuity techniques that serve in clinical and laboratory environments without tailoring them for the concerns of poster visibility modelling would be quite mistaken. This would inevitably lead to the use of unsuitable and unrepresentative objects as target material, and to an investigation of the far limits of visual sensitivity of no serious relevance to poster advertising. However the general methodology of threshold measurement would seem broadly appropriate. For example, a method could be adapted for present use involving the gradual incrementation and decrementation of the size of a target object to ascertain a *relevant* threshold in *representative* conditions. This method was one of the approaches used for the present research. The task was an adaptation of the acuity method, seeking to assess when the target reached a size that enabled it to serve as an perceptual object. The notion was that there is no useful attention value in an undifferentiated object, but that when it was sufficiently perceptually articulated for its shape to be made out then it could be assumed to be capable of functioning as an object capable of guiding visual attention. Accordingly a task was devised for the present research that used rectangular shapes, varying inter alia in terms of aspect ratio, size and orientation (the longer dimension was vertical or horizontal), the observers being asked to judge whether the rectangle was upright or not. The target size was increased until the observer could reliably judge the rectangle's orientation.

A method that was considered but has not been adopted because it is not difficult to see what kind of conclusions would be reached, is the use of a legibility test. In Section 2.2, it was estimated that the MVD for a 6 sheet panel based on the legibility data from road sign research would be 152 metres, and it does not seem that this and the larger distances for larger panels would be viable for use in connection with the POSTAR visibility model. In any case it is not immediately clear what task would be suitable for the purpose of setting MVDs. Legibility is not considered further in this report, however, if agreement were reached on a standard font and size for lettering for a legibility study then this could be re-examined.

A very flexible and attractive method would be to use a visual search technique. This would entail constructing scenes containing objects of various sizes and asking observers to search for a specified target, on the assumption that objects less than a particular size would be missed or only found after an extensive search. It would be possible to set some time limit, and to take the average size of all objects not found in an allowed time interval as indicative of the limiting visual size for visibility. The MVD could then be set from knowledge of the geometry of the situation. The preparation of stimulus material would be time-consuming and would need a good deal of pilot work. The task would also call for focussed visual attention, and the deployment of scanning strategies that are unlikely to be fully representative of poster viewing. This is an obstacle for most methods that might lend themselves to use for the purpose of MVD setting.

In the course of the research a wide range of materials was used, and various relevant factors were investigated, all in the search for a technique that allowed the

perceptual performance of observers to be estimated and for the stimulus sizes and shapes of interest (6, 48 and 96 sheets) all to be differentiated in terms of measured performance. A recurring difficulty, that can readily be appreciated in the light of the quite substantial difference in the areas and aspect ratios of these panels, was that when performance for one was at an intermediate scaleable level, that on another was at a ceiling level (or a floor level). The contrast and hue of the targets were varied to this end, and clutter was introduced in the form of masking rectangular shapes, and in another form of randomly angled masking straight lines. The rationale for these manipulations was partly to solve the ceiling/floor problem, but also to attempt to emulate the clutter of the customary visual environment of poster panels.

There is another class of psychological methods known as psychophysical scaling techniques that are in wide use; instead of eschewing the use of verbal reports, these methods actually use them directly to help externalize the observer's perceptual experience. These methods involved the use of various scaling techniques, in which the observer makes an assessment of some aspect of the world on a verbal, numerical or some other scale. For example, the observer may be asked to assess the intensity of a noise on a numerical scale running from 1 to 100, systematically using the numbers to represent systematic differences in the noise level. The psychological "numbers" are found to vary systematically with actual physical measurements of the sound. This is found to apply for a wide range of perceptual continua; Stevens (1970) reported data for over 15 such continua including loudness, and visual area. Highly systematic relationships have also been evinced in cases where the underlying physical property is not obvious (Barber and Legge, 1976) and this recommended consideration of such a method for scaling visibility.

The observers were given the task of assessing the perceptual prominence of rectangles on a numerical scale and various properties of these targets were manipulated. These properties were the same as those used in the orientation task, namely size (area), aspect ratio, orientation and clutter. Essentially the task for the observers was to report on the attention value or conspicuity of the targets. The term prominence was used in instructing subjects taking part in the experiment since it was judged that this would convey the requirements of the task well. The instructions also made this a good deal more explicit (see Section 4.2) and subjects seemed to grasp easily enough what was required of them. It is to be expected that prominence so judged would increase as area increases, and decrease as clutter increases; should this not be the case, the method would be considered a failure. The nature of any relationship that may emerge is of central interest. The task provides an opportunity to discover whether aspect ratio and orientation affect prominence - it might be the case that the upright format of the 6 sheet has an attentional advantage, and that the long thin shape of the 96 sheet is advantageous. The scale was from 0 to 100 and from the data aggregated over all subjects it was hoped that it would be possible to specify the size needed to produce a given prominence value. Depending on the prominence vs. area relationship, it should then be possible to derive estimates of the distances for objects of different sizes to produce a prominence value that was seen to have some practical utility relevant to the visibility modelling exercise.

In fact, more than one method was adopted for the study to test the extent to which there might be a convergence on common or at least commensurate conclusions. Details of the experiments are supplied in Section 4.

3.2 The rationale for visibility modelling and demonstrations of visibility at a distance

Surprise is sometimes expressed when distances such as have been mentioned in the context of road sign visibility research are suggested as having possible application to poster advertising. It may be helpful to have some more down to earth illustrations of what such distances might involve instead of the dramatic but rather meaningless examples such as telephone wires being visible at a distance of more than a mile. An essential exercise for a poster industry expert is to pace out the distances mentioned in previous studies and the present one, to gain some relevant visual experience. It goes without saying that the view should not be obstructed at any point, not always an easy condition to satisfy in urban environments.

Although legibility is not immediately a serious contender as a method for present purposes, it is instructive to see what people can read when text is viewed at a distance. An A4 sheet of paper in portrait format has the same shape as a 6 sheet poster, and a 6 letter word (Appendix 1 is in 80 point font) should be legible at a distance of about 25 metres. Appendix 1 can be used to illustrate this, the first page is in 80 point font should be viewed at this distance, the second page contains an inset which is scaled for half the distance. Familiar objects such as a matchbox, cigarette packet, envelope or postcard can also be used in this way to demonstrate the remarkable prowess of the human visual system. A matchbox will be found to recognizable as such at a distance of about 15 paces. Postcards are quite good substitutes for posters in this regard because they carry pictorial information, and the viewer can test whether the contents can be made out at a far distance, at the legibility limit this will be about 12.5 metres. All of these examples have the 6 sheet shape, for a 48 sheet an A5 envelope can be used. They do not settle anything but are useful ways of illustrating what the debate concerning MVDs is about.

4. Research design and methods: Maximum viewing distance

4.1 Speeded orientation discrimination task

Basic description:

The task involved the discrimination of the orientation (horizontal or vertical) of the longest axis of a displayed rectangle with responses made with single key-presses. Stimulus rectangles were briefly displayed on a background of distracting, randomly-drawn, lines with the number of distracting lines being reduced each time the stimulus appeared.

Apparatus:

All data collection was carried out using an IBM compatible 486 PC situated in a dimly lit room illuminated by a single 60W lamp positioned above, and pointing away from, the computer monitor.

Stimuli were displayed on a Viglen 14 inch IBM PC VGA colour monitor operating with at a resolution of 640 by 480 (horizontal x vertical) pixels. The screen was covered with a mesh filter to reduce glare, distracting reflection, and to attempt to increase the sensitivity of the display to small contrast changes. Subjects were seated 4m from the screen at a table with two telegraph style response keys operated by their left and right index fingers. Reaction times were recorded with millisecond accuracy.

The rectangle stimuli were displayed in grey on a 'black' (blank screen) background with masking or distracting lines at three levels of contrast. The distracting lines, positioned randomly on the screen, were 2mm wide and of random length. The rectangles were of three aspect ratios - 0.25, 0.50, and 0.66 (corresponding to the 96, 48, and 6 sheet posters respectively), and four different areas (200, 400, 1000 and 2500 pixels). The design factors for this experiment are summarized in Table 1. [These dimensions were decided upon with reference to numerous pilot studies in which ceiling effects of size and contrast were determined.] The actual screen dimensions and respective (approximate) visual angles are given in Table 3.

Table 1: Factors and conditions for the speeded orientation discrimination task.

Factor	Levels	Level descriptions:
Aspect ratio	3	0.66 (6 sheet) 0.50 (48 sheet) 0.25 (96 sheet)
Target area	4	200, 400, 1000, 2500 square units
Orientation	2	Horizontal Vertical
Contrast (see Table 2)	3	Low Medium High

Table 2: Contrast and screen position of the rectangle stimuli employed in the orientation discrimination task

Block	Contrast (RGB number)	Rectangle position.
1	High (14)	Screen centre
2 & 3	High (14)	Random
4 & 5	Medium (10)	Random
6 & 7	Low (6)	Random

[Values for contrast (brightness) in Table 2 are those used by the palette control procedures in Turbo Pascal controlling RGB (red/green/blue) values within a range of 0 to 64. Using the same value for the red, green and blue variables allowed the production of a scale from blank screen to white (0 to 64).]

Procedure:

Trials took the form of a series of displays made up of the stimulus rectangle and distracting lines shown briefly, then removed and then re-displayed with fewer distracting lines. Stimulus rectangles and distracting lines were displayed concurrently (synchronised with screen refresh) and remained on the screen for 250ms before the rectangle was removed and the lines redrawn. The distracting lines remained visible for a further 1,000ms before they were removed with the screen then remaining blank for 1,000ms before the stimulus rectangle was displayed again (the inter-trial interval), this time with 10 less lines. The process was repeated until the same response was made to two successive displays. If two similar responses were not produced then the displays were repeated as above until all lines had been removed, the rectangle was displayed three more times if no further responses were made. Responses were accepted at all times during the rectangle display, lines only display or the inter-trial interval. At the end of each trial (after two identical responses or after the three rectangle only displays) a white circle was displayed on the screen for 2,000ms to alert the subjects to an upcoming full display.

Trials were arranged in blocks of 24 in such a way as to ensure that each of the 24 conditions (see Table 2) were presented once in randomized order. Each subject completed seven blocks of trials with contrast and rectangle position as shown in Table 2.

Table 3: Screen dimensions and visual angles for each rectangle condition

Aspect Ratio	Pixel Area	X dimension (mm)	Y dimension (mm)	X Visual angle (°)	Y Visual angle (°)
96 Sheet (0.25)	200	11.5	3	0.0029	0.0008
	400	16	4	0.0040	0.0010
	1000	24	6	0.0061	0.0015
	2500	40	10	0.0100	0.0025
48 Sheet (0.50)	200	8	4	0.0020	0.0010
	400	11	5.5	0.0028	0.0023
	1000	17	8	0.0042	0.0021
	2500	27	13.5	0.0068	0.0034
6 Sheet (0.66)	200	7	4.5	0.0018	0.0011
	400	9	6	0.0022	0.0015
	1000	15	10	0.0037	0.0025
	2500	24	16	0.0060	0.0040

Subjects were instructed to respond with a left or right key-press to the orientation of the longest axis of the rectangle (horizontal or vertical) if they could

make it out. They were told only to respond if they were sure as to its orientation but to press the correct button as quickly as possible once they were confident of the correct response.

Subjects:

23 subjects (10 male, 12 female, mean age 26.8 years) were recruited from the Birkbeck College Subject Panel and with the use of posters and flyers distributed around the University on London. All subjects reported having normal or corrected to normal vision, and this was checked for the purposes of the research using a standard letter chart read at 4m. All subjects were able to read 7mm letters correctly).

[1 additional subject completed the experiment but failed to follow the instructions correctly, their results were not further analysed. The data for three subjects of the 23 who completed the experiment were removed from the analysis as described later]

4.2 Prominence Judgement Task

Basic description:

Subjects were asked to make judgements as to the relative prominence of stimulus rectangles, orientated with the longer axis either vertical or horizontal, displayed on a background of different numbers of distracting lines.

Apparatus and stimuli:

Stimuli were displayed on an IBM PC in the same conditions as for the Orientation Discrimination Task above. Rectangle stimuli and distracting lines were also identical to those used in the previous experiment and were displayed on a black screen in the high contrast (14) grey. The factors employed in this experiment are shown in Table 4 below. Prominence judgements were recorded by the subjects on A4 answer sheets with numbered boxes corresponding to each trial in a block.

Table 4: Factors and instances of conditions employed in the prominence judgement task.

Factor	Levels	Description
Aspect Ratio	3	0.66 (6 sheet) 0.50 (48 sheet) 0.25 (96 sheet)
Area	5	200 (pixels) 400 1000 2500 None present
Orientation	2	Horizontal Vertical
Distracting lines	4	100

		60
		20
		None

Procedure:

Subjects were given written instructions as follows:

We are doing a series of experiments on the visibility of objects in various conditions. The task today is to judge the ease with which you can spot a rectangle displayed on the computer screen. You will see a variety of rectangle shapes and sizes that may appear anywhere on the screen. There will only ever be one rectangle on the screen at any time but, on some occasions there will be none at all. The display will also contain a collection of lines, spread around the screen, with the rectangle (if there is one) somewhere in their midst. You can think of your task then as signalling the prominence of the rectangle, put another way, judge how well it stands out.

After each display, please respond by giving a number between 0 and 100. If the rectangle is completely obvious, and it takes no effort at all to spot, then you should assign a very high number. A large rectangle in the middle of an otherwise blank screen would get 100. A rectangle that pops out at you with little or no effort on your part would get a high number, and if it is a bit less prominent, then a lower number would be needed. The more you have to do to find and make it out, the lower the number you should give. If on balance you judge that there was no rectangle but are not completely sure then use a low number, but if you are quite sure that there was no rectangle, then use an even lower number, while a blank screen with definitely no rectangle would get 0. Try to be consistent with your use of numbers but don't worry about it.

Please make sure that you make a judgement for each display and be careful to match the number of the display with the correct one on the response sheet.

Each trial was initiated by a key-press by the subject which cleared the computer screen for 1,000ms. The stimulus rectangle and distracting lines were then displayed together for 500ms before the screen was blanked again for a further 1,000ms. A message was then displayed instructing the subject to write their judgement in the box corresponding to the number of the trial.

Trials were arranged in five blocks of 108 such that each condition of lines, area and orientation appeared once in each block (the rectangle not present conditions naturally did not have any different orientation conditions). Subjects were instructed that they may take a short break between blocks if they wished.

Subjects:

18 subjects completed the experiment (7 male, 11 female, mean age 29.9 years) all of whom had normal or corrected to normal vision (all completed the letter reading task used in the previous experiment).

4.3 Orientation discrimination threshold: Staircase method.

A further experiment was conducted using a classical staircase method to estimate the point at which an observer could discriminate the orientation of a rectangle. The assumption was that this point would signal when the rectangle moved

from being something with a mere presence to having a tangible property. The design of the study is summarized in Table 5.

Apparatus and Stimuli:

The experiment was conducted using an IBM compatible 486 personal computer as with the orientation discrimination task above. The stimulus rectangles started with their longest axis being 120 pixels long.

Subjects respond with key-presses to the orientation of screen-displayed rectangles of the three aspect ratios. Each correct response reduced the size of the rectangle by 20%. After one in four of the displays the rectangle size was increased by 20%. The threshold for each was determined by 6 incorrect responses made in the last 10 steps. The rectangles remained visible until a response was made.

Table 5: Factors and conditions for the orientation discrimination threshold task.

Factor	Levels	Level descriptions:
Aspect ratio	3	0.66 (6 sheet) 0.50 (48 sheet) 0.25 (96 sheet)
Background brightness	3	Blank screen (0) Medium (10) High (20)
Rectangle contrast	3	Low (5) Medium (10) High (20)

This task failed to discriminate sufficiently between screen sizes, performance for many subjects reaching a ceiling level under what seemed the most adverse conditions. While the data may have been useful for setting visibility distances for 6 and 48 sheet panel sizes, it transpired that the task in effect was incapable of differentiating between 48 and 96 sheet sizes. Accordingly there is no point in reporting the detailed findings for the 11 subjects tested on the task. Performance on the task rivalled that obtained in laboratory studies of visual acuity, and pressed the equipment used to present and control the stimuli to its technical limits. No results for this task are therefore reported.

5. Results

5.1 Orientation discrimination task

The data for the orientation discrimination task were analysed using a repeated measures analysis of variance (ANOVA). There were four independent variables as described above. The dependent variable was the number of masking lines left on the display when the observer reached the task criterion of making identical successive responses as described above. A high score means that the observer could make out the orientation of the target when there were relatively large numbers of masking

lines on the display. Thus, an observer with a high mean score can be thought of as doing better than someone with a low score. The analysis of variance summary table is included as Appendix 2. On inspection of the data three subjects were found to have so many errors that their data were unusable.

The principal interest of the report hinges on the effects of target area, and so this section of the report will focus on the most directly relevant aspects of the data. Area itself was a highly significant source of variance ($F(3,57) = 28.89$; $p < 0.001$). There were two significant interactions involving Area; these were Aspect Ratio \times Area ($F(6,114) = 5.23$; $p < 0.001$) and Contrast \times Area ($F(6,114) = 14.47$; $p < 0.001$). These interactions are depicted in Figures 1 and 2 (they are also tabulated in Appendix 2; Tables 2.2 and 2.3 respectively). The Contrast \times Area interaction is attributable to a diminution of the effect of area under medium and high contrast conditions. Indeed when the low contrast data are excluded, the effect of area is not significant. A decrease in the effect of area is also manifest in the other interaction. It appears that the most elongated of the three formats made the task easier and the effect of area was correspondingly most marked for the least elongated format. This is hardly surprising since the task might reasonably be expected to be easier for more extreme rectangular shapes.

Figure 1: Orientation discrimination as a function of area moderated by aspect ratio

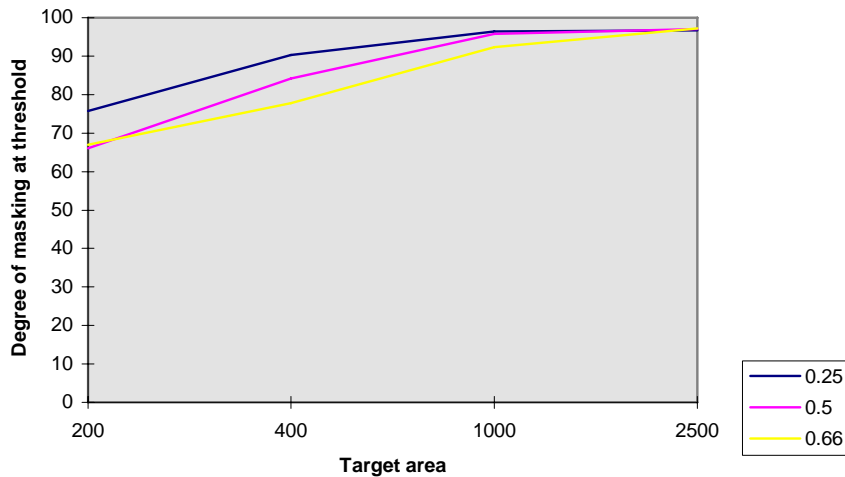



Figure 2: Orientation discrimination threshold as a function of area moderated by contrast



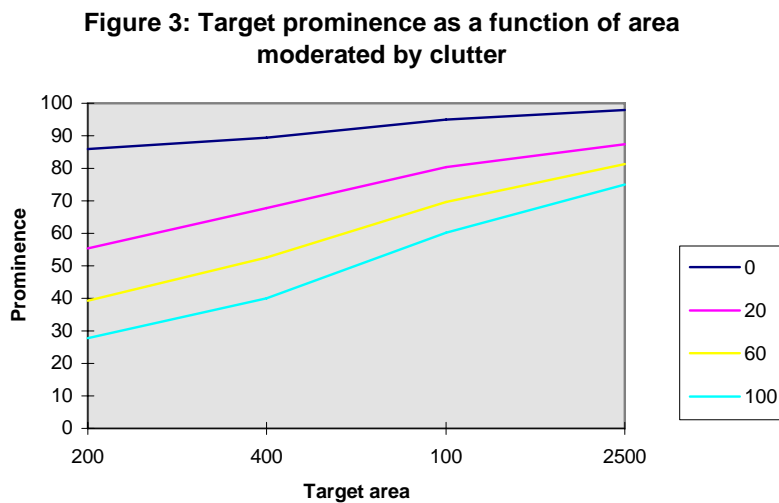
All three of the other main effects were also significant. For Aspect Ratio ($F(2,38) = 8.26$; $p < 0.01$), performance declined as the target shape became less elongated (89.8 for $AR = 0.25$, 85.8 for $AR = 0.50$, and 83.6 for $AR = 0.66$; a progression from 96 to 6 sheet style). It should be borne in mind that this is for targets equated in area. The effect of Orientation was significant ($F(1,19) = 5.78$; $p = 0.027$), performance for the horizontal format being rather better than for the vertical format (87.1 vs. 85.6 respectively). This finding appears to be at odds with a study by Bijl and Koenderink (1993) except that in their study it was absolute visibility that seems to have been in issue and not shape discrimination as here. Contrast was also significant ($F(2,38) = 20.84$, $P < 0.001$); this was attributable to a marked reduction in performance under low contrast (78.2), with medium and high contrast conditions producing very similar levels of performance (91.0 vs. 89.9 respectively). There was one significant interaction not involving Area, namely Aspect Ratio \times Orientation \times Contrast ($F(4,76) = 3.68$; $p = 0.009$). Additional analysis revealed that this interaction was not robust, it disappeared when the low contrast condition was excluded, and it will not be discussed further.

It is evident that performance was for the most part very high. Despite having contrived extreme viewing conditions, observers performed at close to ceiling levels in many conditions. This is likely to be at least partly the basis of the appearance of some of the interactions described in this section. The effect of area under the low contrast condition is arguably the best indicator of how performance depend on area for the purpose of estimating visibility distances, and this will be continued in Section 5.4. A pertinent comment regarding the use of any other data but these will be found at the end of the next section. Because the suspicion was aroused as the experiment was in progress that the task was not fully proofed against ceiling effects, a further task was devised, and this was the prominence judgement task.

5.2 Prominence judgement task


The data for the prominence judgement task were analysed using a repeated measures analysis of variance. There were four independent variables as described above. The dependent variable was mean of the numerical judgements made by each observer in each condition. The analysis of variance (ANOVA) summary table is included as Appendix 3.

As before the effects of target area will be the focus of the analysis. Area itself was a highly significant source of variance ($F(3,51) = 111.50$; $p < 0.001$). There was one significant interaction involving Area, namely for Clutter x Area ($F(9,153) = 22.00$; $p < 0.001$). This interaction is depicted in Figure 3 and is tabulated in Appendix 3 (Table 2.2). It can be seen that the effect of area is markedly increased as clutter increases.



The effect of Clutter was significant ($F(3,51) = 141.80$; $p < 0.001$), but this was not so for Orientation ($F(1,17) = 2.67$; $p = 0.12$) or Aspect Ratio ($F(2,34) = 0.31$; $p = 0.74$). There was only one significant interaction not involving Area, namely Clutter x Orientation ($F(3,51) = 4.44$; $p = 0.008$). The interaction is depicted in Figure 4 (see also Table 3.3 in Appendix 3). There is no systematic difference between orientations and it is difficult to see that any great practical importance attaches to this finding, which although statistically reliable is not particularly striking.

Figure 4: Target prominence as a function of clutter moderated by orientation



The conditions of the experiment were quite demanding for observers, as the more extreme conditions in the orientation discrimination task proved to be. It seems reasonable to take the most demanding condition again as the soundest basis for estimating visibility distances. It is clear that there is a clear and comprehensive separation between target areas at the highest level of clutter. Distances will be estimated from the data for this condition. It should be emphasised that data from any other conditions would produce smaller separations between panel sizes, but at the price of even greater visibility distances than these do.

5.3 Estimating visibility distances

Both experiments demonstrate a clear effect of image size on a perceptual judgement. The first task involved a fine judgement about an object's orientation, and the experiment sought to establish the limiting point at which this could be done. In this sense the experiment may be thought of as a means of estimating when an object achieves objectness for the observer, however, not merely as a matter of when it is visible but when it is sufficiently identifiable as an object *with orientation*. It seems reasonable to take this limiting point as a conservative estimate of visibility since more than the mere presence of an object is necessary for the response. The second experiment asked observers directly about their perceptual experiences, requiring them to indicate how easy it was to spot the target object on the screen. Estimating target prominence in principle supports the application of graded criteria regarding visibility. The scaling of the observers' perceptual judgements is done by themselves, and the scale may be taken to reflect a gradation from strict to lax criteria. The experiment clearly showed that image size affects these judgements.

A relationship with image size is not of course surprising - indeed it would have been singularly awkward had such a relationship not been obtained. What is of concern is the nature and robustness of the relationship. The evidence in both cases is that the relationship is subject to some influences that are pertinent to an analysis of visibility, but these influences do not seem to invalidate the use of the data for the estimation of visibility distances for present purposes. The nature of the relationship in both cases was examined by regression analysis of different measures of the performance in the two tasks against various alternative measures of area. The simplest approach would be to consider a linear relationship between performance

and area, and in fact such a relationship would be supported by the data for both experiments. A much stronger relationship, however, was obtained in both cases by regressing performance on log area.

The relationships obtained for the overall mean performance are summarized in Table 6. It will be recalled that in both cases there was a key factor that moderated the performance vs. area relationship, namely clutter in the case of orientation discrimination and contrast in the case of prominence. Low levels of clutter made the first task easier, as did high levels of contrast for the prominence task. Visibility distances would accordingly be greater than if high clutter and low contrast were applied respectively. Estimates of visibility distances for both tasks were therefore calculated for the two most extreme conditions in each case, and are shown in the lower half of Table 6.

Table 6: Summary of regression analyses

Task/condition	Regression coefficient	Intercept
Orientation: Mean	24.72	16.53
Prominence: Mean	30.74	-17.77
Orientation: Low contrast	42.01	-40.50
Prominence: High clutter	43.82	-73.04

To set visibility distances using these functions, consider the first of them:

$$Y = 24.72 \log \text{ area} + 16.53.$$

This expresses performance on the task Y in terms of the target area. For a score of 95%, that is, almost all the masking lines remain on the screen and the target's orientation is correctly judged, then the image size has to have $\log \text{ area} = 3.174$, so the screen area is 1494. Reference to Table 1 illustrates screen sizes, but it needs to be borne in mind that the screen was positioned 4 metres from the observer. Simple geometry produces a visibility distance for a 96 sheet in the region of 420 metres. For a performance criterion of 99%, the visibility distance would be about 350 metres. These are rough figures to give an impression of how visibility distances would vary with the visibility criterion. Visibility distances for 6, 48 and 96 sheet panels have been estimated from the mean performance scores in the two experiments and again for the more demanding conditions in each case, that might be thought to be the most representative of the visual difficulty of the real world of posters of those used in these studies. In fact it is arguable that the actual conditions of the experiment were more like looking for an unlit 6 sheet at dusk. In any event the estimates, summarized in Table 7, are considerably at variance with what obtains at present in the POSTAR visibility model. It will be seen that the estimates based on orientation discrimination are almost twice those for the prominence task. It will be seen that a

considerable adjustment to bring maximum visibility distances into register with what these two visibility studies indicate.

Table 7: Maximum visibility estimates for orientation and prominence tasks

	Panel size	6 sheet	48 sheet	96 sheet
Orientation	means/95	155	446	631
	means/ asymptote	123	354	500
	low contrast/95	146	421	595
	low contrast/asymptote	127	367	519
Prominence	means/95	87	251	355
	means asymptote	72	208	294
	high clutter/95	72	209	295
	high clutter/asymptote	63	183	259

The two studies agree in respect of the performance vs. log area relationship, however, they are at odds with respect to the actual distance estimates they produce. Since the orientation task entails more focussed effort to identify the target properties, it would seem less suitable as a basis for setting MVDs than the prominence task. The latter requires the observer to signal the ease with which the target can be made out. For it to “pop out” at the observer, means that it is blatantly obvious. The target has to be larger when viewing conditions are relatively unhelpful to the visual system, and this is what the most extreme condition (Prominence/high clutter) represents. Accordingly it is recommended that, all else being equal, new MVDs should be adopted in line with the asymptotic performance in this task and condition: namely, 63 metres for 6 sheet panels, 183 metres for 48 sheet panels, and 259 metres for 96 sheet panels.

Table 7 shows the maximum visibility estimates based on the two tasks, for the scores averaged over all conditions in each experiment, and for low contrast and high clutter conditions for orientation and prominence tasks respectively. Distances estimated when performance is 95% (orientation) and 95 (prominence) are given together with asymptotic scores in each case. It should be noted that most threshold studies adopt performance criteria of 50% or 75%, and that 95% is a very conservative criterion. Considerably larger MVDs could be supported on the basis of the results. The score of 95 for the prominence task (which is not to be equated with the 95% of the orientation task) means that the target was *extremely easy to see*. The observers would be saying in effect that the rectangle jumped out at them under those conditions.

Notwithstanding these recommendations, there are other considerations relating to how people progress through the world, and how they deploy their visual competencies, and these are reviewed briefly in the final section.

6. Discussion

Evidence from all sides supports an increase in the MVD settings in POSTAR. There are some considerations, however, that should temper any move in this direction. It is clear that considerable advantage may be derived by ensuring that observers are able to catch early sight of a poster.

Further research might be considered to check the present findings, and to assess the recommended MVDs from a different perspective. The obvious direction to consider at this juncture is an eye movement study using whatever is practically feasible. Knowing where the eye is pointing within say half a degree is quite satisfactory for most purposes. Applying the same laboratory conditions as used in for the POSTAR visibility research, this corresponds to a 1.2 metre real-world extent (the width of a 6 sheet by way of illustration) viewed at about 140 metres. It would seem likely that measurement accuracy and visibility considerations for larger panels would rule out the use of this approach. A feasibility study would seem worth considering using the smallest panels, on the understanding that one would need to generalize to the larger panel sizes.

A visual search task might be a particularly useful and convenient way of checking the present findings. It could require subjects to look specifically for posters in real world scenes, and although this is not the mental set of most people with respect to posters, by giving a group of subjects very limited exposure times to the scenes, it would be possible to set or check visibility distances. It would for instance be of interest to determine if posters can be detected at distances approximating the proposed MVDs, given limited opportunity to scan for them. If as suggested below, people tend to scan in a preferred viewing plane at some 30 or so metres from them (say), then they may well miss posters at greater (or nearer) distances, and the distance at which they spot most posters would be practically important to know in light of what is noted in the final paragraph. A fairly simple study could elucidate much of this.

The effect on visibility scores is evidently severe to judge from runs of the model using MVDs as recommended. This raises several issues. First, the proposed MVDs may be “wrong”. Second, it may be that the model is deficient in some sense. Third, maybe both are “right”, and the conversion of visibility scores to the “currency” of the medium is inappropriate. Some form of checking and adjustment to some component of this system is evidently needed, but this will not be pursued here.

The foregoing discussion illustrates the complexity of the issues that underlie poster visibility, and does so in the circumscribed context of deciding what should serve as the maximum visibility distance for a poster. It does not touch on such germane issues as whether what an observer gets from a poster varies as his or her viewing distance changes. This question has not been addressed by the visibility exercise so far, but it should be noted that there has been agreement that a simple unitary model would be easiest to justify and convey to others. The assumption has therefore hitherto been that the information conveyed by a visual object is equivalent at all distances from the observer even though there may be more or less of that

information as distance changes. The change has been assumed to be merely quantitative so far, and plainly the assumption may need to be reconsidered.

A second important point is that it is assumed in the visibility model that the visual world is sampled equally at all distances from the observer. Suppose our observer is shown two uniformly coloured circular panels placed at 30 metres and another at 150 metres, constructed so that the images they cast at the observers' eyes are equal in size (in other words, the panel at 150 metres is five times bigger than the nearer one). A question that has not been posed is whether the observer is equally likely to see the two panels. They have equal image sizes and randomly fixating the visual scene would surely not favour one over the other. But this assumes that the eye focusses back and forth in space between the foreground of the scene and the horizon, ranging over all intermediate distances with equal frequency. This may not be a good assumption. It is arguable that there is a preferred viewing distance at which the eyes range laterally, with adjustments of a forward/backward kind of a more limited kind. If this is so, then objects at a great distance will be visible if the eye reaches that far, but ordinarily it will only operate at a more humdrum distance that is appropriate for getting around the world. Some evidence consistent with this kind of account was reported by Wagner, Baird and Barberesi (1981) for pedestrians, who were found to report looking maximally at objects between 1½ and 3 metres ahead of them. While they also reported looking at very far objects, this was relatively seldom in the totality of things. Whether this applies to drivers is not known but it would seem inevitable that the distribution of attention as a function of distance would reflect the operational demands of driving, and any home distance around which looking behaviour may be organized would be much greater than that for people walking. In the light of this possibility, some kind of weighting factor might need to be incorporated into the drivers' visibility model. This is an overdue issue for consideration that will inevitably be brought into focus by the pedestrian visibility modelling study. A simple three step weighting might be considered, such as 1.0 for distances to 30 metres, 0.5 for 30 to 100 metres, and 0.1 for distances in excess of this. There is no immediate empirical basis for such a proposal but something of the sort may provide a way of reining back the apparently excessive increases in visibility scores when the recommended MVDs are used.

7. References

- Barber, P.J., and Legge, D. (1976) *Perception and information*. London: Methuen.
- Bijl, P., and Koenderink, J.J. (1993) Visibility of elliptical Gaussian blobs. *Vision-Research*, **33**, 243-255
- Forbes, T.W. (1972) Visibility and legibility of highway signs. In T.W Forbes (Ed.), *Human factors in highway traffic safety research*. New York: Wiley.
- Forbes, T.W. and Holmes, R.S. (1939) Legibility distances of highway destination signs in relation to letter height, letter width and reflectorization. *Proceedings of the Highway Research Board*, **19**, 321-335.

Schiff, W. (1980) *Perception: An applied approach*. Boston: Houghton Mifflin.

Stevens, S.S. (1970) Neural events and the psychophysical law. *Science*, **170**, 1043-1050.

Wagner, M., Baird, J.C., and Barbaresi, W. The locus of environmental attention. *Journal of Environmental Psychology*, **1**, 195-206.

Appendix 1: Examples of legibility task material

RANGOM

POBTAR

LKTTER

BUY
BETTA
BURGERS
THEY
WON'T
KILLYA

Appendix 2

Table 2.1: Analysis of variance of mean number of lines left for successful discrimination: orientation discrimination task

Source	Degrees of freedom	MS treatment	MS error	F-ratio	tail probability
Aspect	2,38	4730.75	572.61	8.26	0.001
Area	3,57	56694.35	1962.42	28.89	0.000
Orientation	1,19	810.00	140.15	5.78	0.027
Contrast	2,38	24214.18	1161.86	20.84	0.000
Aspect x Area	6,114	1300.72	248.71	5.23	0.000
Aspect x Orientation	2,38	669.43	347.99	1.92	0.160
Aspect x Contrast	4,76	216.22	106.44	2.03	0.098
Area x Orientation	3,57	287.69	118.72	2.42	0.075
Area x Contrast	6,114	6595.34	455.72	14.47	0.000
Orientation x Contrast	2,38	336.93	84.46	3.99	0.027
Aspect x Area x Orientation	6,114	72.46	214.25	0.34	0.915
Aspect x Area x Contrast	12,228	89.42	119.46	0.75	0.703
Aspect x Orientation x Contrast	4,76	405.10	110.09	3.68	0.009
Area x Orientation x Contrast	6,114	122.67	79.05	1.55	0.168
Aspect x Area x Orientation x Contrast	12,228	204.00	133.92	1.52	0.117

Table 2.2: Orientation discrimination threshold as a function of target area moderated by aspect ratio

Aspect ratio	Target area			
	200	400	1000	2500
0.25	75.67	90.25	96.38	96.71
0.50	66.04	84.25	95.75	97.08
0.66	66.92	77.79	92.33	97.17

Table 2.3: Orientation discrimination threshold as a function of target area moderated by contrast

Contrast	Target area			
	200	400	1000	2500
Low	50.96	72.50	92.50	96.79
Medium	79.50	91.25	96.58	96.79
High	78.17	88.54	95.38	97.38

Appendix 3**Table 3.1: Analysis of variance of mean judgements for prominence task**

Source	Degrees of freedom	MS treatment	MS error	F-ratio	tail probability
Aspect	2,34	42.88	138.82	0.31	0.736
Area	3,51	93837.23	841.61	111.50	0.000
Orientation	1,17	265.45	99.34	2.67	0.121
Clutter	3,52	136344.54	961.55	141.80	0.000
Aspect x Area	6,102	53.74	49.72	1.08	0.379
Aspect x Orientation	2,34	24.15	82.86	0.29	0.749
Aspect x Clutter	6,102	49.85	49.03	1.02	0.419
Area x Orientation	3,51	88.63	45.17	1.96	0.131
Area x Clutter	9,153	5154.99	234.35	22.00	0.000
Orientation x Clutter	3,51	186.38	42.02	4.44	0.008
Aspect x Area x Orientation	6,102	108.60	44.71	2.43	0.031
Aspect x Area x Clutter	18,306	40.97	49.90	0.82	0.675
Aspect x Orientation x Clutter	6,102	56.07	55.20	1.02	0.419
Area x Orientation x Clutter	9,153	85.86	48.80	1.76	0.080
Aspect x Area x Orientation x Clutter	18,306	51.33	55.65	0.92	0.552

Table 3.2: Prominence as a function of area moderated by clutter

Target area	Clutter (number of masking lines)
-------------	-----------------------------------

	0	20	60	100
0	2.55	8.01	15.39	15.72
200	85.86	55.31	39.35	27.84
400	89.49	67.83	52.58	40.06
1000	95.05	80.44	69.69	60.16
2500	97.89	87.41	81.36	75.02

Table 3.3: Prominence as a function of orientation moderated by clutter

	Clutter (number of masking lines)			
Orientation	0	20	60	100
Upright	92.22	73.75	61.59	50.35
Horizontal	91.93	71.75	59.90	51.19