PERGAMON

Vision Research 41 (2001) 3613–3625

Vision Research

www.elsevier.com/locate/visres

Eye movements during visual search: the costs of choosing the optimal path

Christian Araujo [a], Eileen Kowler [a,*], Misha Pavel [b]

[a] Department of Psychology, Rutgers University, 152 Frelinghuysen Road, Piscataway, NJ 08854, USA
[b] Department of Electrical and Computer Engineering, Oregon Graduate Institute, 20000 NW Walker Road, Beaverton, OR 97006, USA

Received 2 March 2001; received in revised form 17 July 2001

Abstract

Saccadic eye movements are usually assumed to be directed to locations containing important or useful information, but such assumptions fail to take into account that planning saccades to such locations might be too costly in terms of effort or attention required. To investigate costs of saccadic planning, subjects searched for a target letter that was contained in either one of two clusters located on either side of a central fixation target. A target was present on each trial and was more likely (probability = 0.8) to appear in one cluster than the other. Probabilities were disclosed by differences in cluster intensities. The distance between each cluster and central fixation varied (60′–300′). The presentation time was limited (500 ms) to ensure that a successful search would require a wisely chosen saccadic plan. The best chance of finding the target would be to direct the first saccade to the high-probability location, but only one of the six subjects tested followed this strategy consistently. The rest (to varying degrees) preferred to aim the first saccade to the closer location, often followed by an attempted search of the remaining location. Two-location searches were unsuccessful; performance at both locations was poor due to insufficient time. Preferences for such ineffective strategies were surprising. They suggest that saccadic plans were influenced by attempts to minimize the cognitive and attentional load attached to planning and to maximize the number of new foveal views that can be acquired in a limited period of time. These strategies, though disastrous in our task, may be crucial in natural scanning, when many cognitive operations are performed at once, and the risk attached to a few errant glances at unimportant places is small.

Keywords: Eye movements; Saccadic eye movement; Saccades; Visual search; Visual attention; Optimal search; Planning; Decision-making

1. Introduction

The decision human beings make most often during the course of a day is undoubtedly the decision about where to look next. Saccadic eye movements occur once every several seconds when prolonged viewing of an object is required, and as often as three times each second when new information has to be acquired quickly. In an ideal world, decisions about where to direct saccades would be made rationally, on the basis of the demands of the ongoing task, so that the line of sight will be brought without delay to those locations containing the most useful information. Often, saccades appear to be doing just this, judging by how well the locations examined are suited to the task (Epelboim et al., 1995; Ballard, Hayhoe, & Pelz, 1995; Epelboim & Suppes, 2001; Land & Hayhoe, 2001; Melcher & Kowler, 2001). But there are also some striking exceptions, examples where cues to the optimal saccadic path are ignored and saccades are instead directed to locations known to be useless and uninformative (Zelinsky, 1996; Hooge & Erkelens, 1998, 1999). The rules guiding such saccades are less clear.

Understanding the basis for saccadic decisions, whether the sequences of movements appear to be sensible or irrational, requires taking into account that saccades are constrained by inherent and unavoidable limitations—visual, cognitive and motor in origin—on the ability to perform particular sequences of movements. For example, planning saccades, in and of itself,

* Corresponding author. Tel.: + 1-732-445-2086; fax: + 1-732-445-2263.
E-mail address: kowler@rci.rutgers.edu (E. Kowler).

0042-6989/01/$ - see front matter
PII: S0042-6989(01)00196-1

is an activity that takes time and consumes processing resources. Thus, the sequence of eye movements that is best for obtaining important visual information might not be performed because planning or executing the sequence is too difficult or too costly (in terms of the demands on limited processing resources). A complete understanding of saccadic decisions requires identifying the relevant constraints and costs, and finding out how they are taken into account.

We set out to study saccadic decisions, not in the complex and rapidly changing environment that characterizes most natural tasks, where decision rules might be hard to infer, but rather in a simpler, single brief 'snapshot' of time, under conditions selected to incorporate some of the demands and constraints that determine natural saccadic decisions.

We used a visual search task to study the relative importance that subjects would assign to two different aspects of the task: (1) the probability of finding the target in one or another location and (2) the distances of the locations from the line of sight. One of us (Pavel) had preliminary results indicating that in a manual search task involving arm movements, probability was the principle determinant of performance. But when constraints were imposed in the experiment that limited the speed of the movements, information about probability was often ignored, despite its value for finding the target quickly. Instead, preferences for moving the arm to the closer location began to emerge, and probability played a relatively smaller role. A preference to move to closer locations makes sense when slow arm movements are involved because considerable time is expended covering large distances. Saccadic eye movements, by contrast, are so fast that spatial distance, provided it is not excessively large, can be safely neglected. We expected that in a search task mediated by eye movements, where the spatial distances involved were modest, only probability would be important.

These intuitions were wrong. Subjects unexpectedly took into account spatial distance when choosing where to look, and made much less use of probabilistic information than they should have. As a result, search performance suffered. Analysis of both the eye movement patterns and the success at finding the target suggests that subjects brought to our task some of the strategic conventions and decision rules that are usually helpful in more naturalistic situations: they attempted to plan saccadic sequences and to minimize effortful planning preceding each saccade.

2. Method

2.1. Overview of the task

The task required search for a target (a tilted letter T) located at the center of either one of two clusters of items (tilted Ls). The clusters were located on either side of the central fixation point (distances 1–5°) and the displays were presented very briefly (0.5 s), long enough to permit one saccade, but seldom two. Cues present in the display indicated which cluster was most likely (probability = 0.8) to contain the target. Success at finding the target was demonstrated by a report of target orientation at the end of the trial.

2.2. Display apparatus

The display was generated by digital-to-analog converters and shown on a display monitor (Tektronix 608, P4 phosphor) located directly in front of the subject's right eye. The display was refreshed every 20 ms, a rate that was high enough to prevent visible flicker. The luminance of the display, measured by a UDT photometer (model 61) from a 2.2×2.2 cm region containing 1600 points refreshed every 20 ms, was 36 cd/m^2. Luminance was reduced to 9 cd/m^2 in some cases (see below). Displays were viewed against a dim, homogeneous background (luminance = 1.8 cd/m^2) generated on a second display monitor located at right angles to the first. The two displays were combined by means of a pellicle beam splitter. The room was dark except for the displays. The stimulus was viewed through a collimating lens that placed it at optical infinity.

2.3. Stimulus and procedure

Before each trial, the subject fixated a central point and started the trial by a button press, when ready. After 500 ms, during which fixation had to be maintained on the central point, the critical frame appeared for 500 ms. The critical frame was followed by a mask (500 ms). Fig. 1 shows the sequence of frames.

The critical frame contained two 3×3 clusters of letters, as shown in Fig. 1. One cluster contained the target, a letter T, whose orientation was chosen randomly from eight possibilities (upright or tilted 45°, 90°, 135°, 180°, 225°, 270°, or 315° from upright). The remaining letters were Ls with an orientation chosen randomly from the same eight values. The mask contained two clusters of Xs occupying the same locations as the letters in the critical frame.

One letter cluster was located to the right and the other to the left of fixation. The center of one cluster was either 60′, 120′, 180′, 240′ or 300′, from fixation. The center of the other cluster was 360′ from the center of the first cluster. This created 5 different location-pairs: left cluster 60′ away from fixation and right cluster 300′ away; left cluster 120′ away and right cluster 240′ away; both clusters 180′ away; left cluster

240′ away and right cluster 120′ away, left cluster 300′ away and right cluster 60′ away. The location pair was selected randomly from these S possibilities on each trial.

The size of the letters in each cluster varied slightly among the subjects (range = 8′–10′ stroke). Sizes were chosen in preliminary testing so that subjects were able to easily discern the target letter in the cluster when they looked directly at it. Letters in the cluster were separated by 30′ (center to center). At these small letter sizes, and with so many nearby neighbors, the target letter orientation could not be discerned from the central fixation position.

One target letter T was presented in each trial. The probability of one of the clusters containing the letter T was 0.8, and the probability of the other cluster containing the T was, correspondingly, 0.2. The location of the high-probability cluster (left or right of fixation) was chosen randomly on each trial. The probability assigned to each cluster was disclosed to the subject by setting the luminance of one cluster in the critical frame to 36 cd/m^2 and the luminance of the other to 9 cd/m^2. In half the sessions, the higher luminance signaled the higher probability. In the remaining sessions, the significance of the luminance cue was reversed.

Subjects were told that one of the clusters contained the target letter T and that they had to report its orientation by means of a button press after each trial. They were also told the probability value associated with each intensity cue (e.g. high intensity = 0.8 and low intensity = 0.2 or high intensity = 0.2 and low intensity = 0.8) before each session began. Responses were given after the mask and were immediately followed by a display (500 ms) of the target letter in the same location it had occupied in the critical frame to disclose the correct response. A 'beep' was sounded after each trial that had an incorrect report.

2.4. Subjects

Six subjects were tested, all naïve with respect to the purpose of the experiments. Three of the subjects had normal vision and were tested without spectacle correction, one was tested while wearing soft contact lenses, and two wore spectacle correction and were tested with a negative lens in the optical path to improve display visibility.

2.5. Eye-movement recording

Two-dimensional movements of the right eye were recorded by a Generation IV SRI Double Purkinje Image Tracker (Crane & Steele, 1978). The subject's left eye was covered, and the head was stabilized on a dental bite board.

The voltage output of the Tracker was fed on-line through a low-pass 50 Hz filter to a 12 bit analog-to-digital converter (ADC). The ADC, controlled by an IBM-compatible PC, sampled eye position every 5 ms. The digitized voltages were stored for later analysis.

Tracker noise level was measured with an artificial eye after the tracker had been adjusted so as to have the same first and fourth image reflections as the average subject's eye. The filtering and sampling rate were the same as those used in the experiment. Noise level,

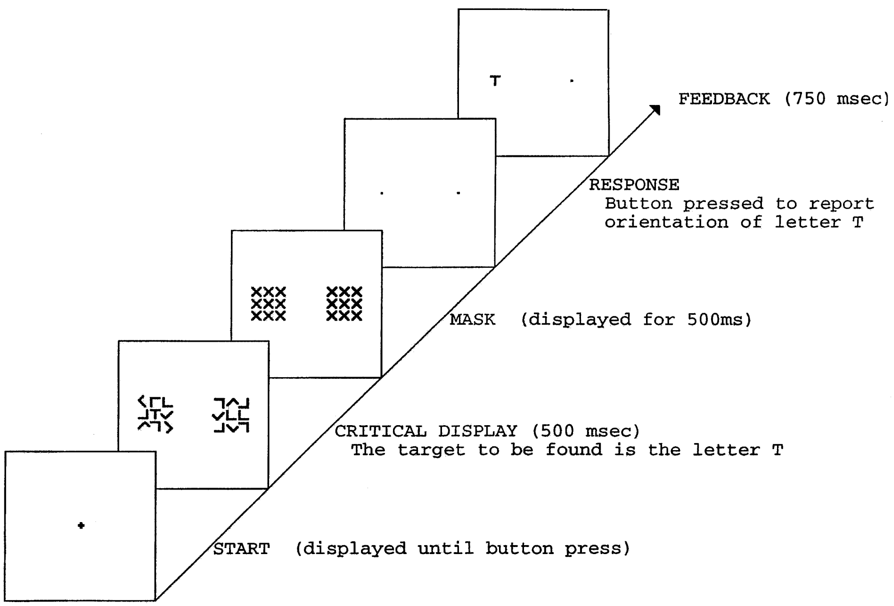

Fig. 1. Sequence of frames in a trial. The target letter T is shown in the left-hand cluster in the second frame. In the experiment, the orientation of the T was selected randomly from eight possible values.

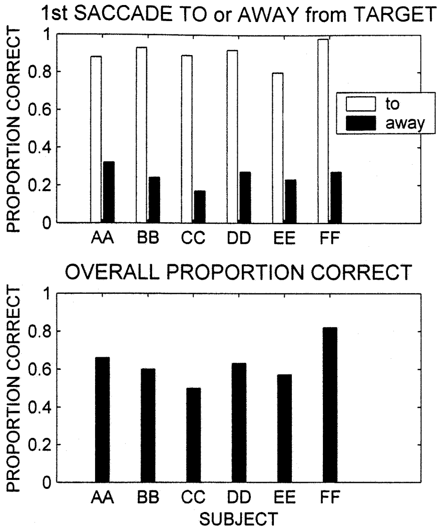

Fig. 2. (Top) Proportion correct reports for subjects AA–FF given that the first saccade was directed to the cluster that did (open bars) and did not (filled bars) contain the target. (Bottom) Proportion of correct reports over all trials.

expressed as a standard deviation of position samples, was 0.4′ for horizontal and 0.7′ for vertical position.

Recordings were made with the tracker's automatically movable optical stage (auto-stage) and focus-servo disabled. These procedures are necessary with Generation IV Trackers because motion of either the auto-stage or the focus-servo introduces large, artifactual deviations of Tracker output. The focus-servo was used, as needed, only during intertrial intervals to maintain subject alignment. This can be done without introducing artifacts into the recordings or changing the eye position/voltage analog calibration. The auto-stage was permanently disabled because its operation, even during intertrial intervals, changed the eye position/voltage analog calibration.

2.6. Analysis of saccades

The beginning and end positions of saccades were detected by means of a computer algorithm employing an acceleration criterion. Specifically, eye velocity was calculated for two overlapping 20 ms intervals. The onset time of the second interval was 10 ms later than the onset time of the first. The criterion for detecting the beginning of a saccade was a velocity difference between the samples of 300′ or more. The criterion for saccade termination was more stringent in that two consecutive velocity differences had to be less than 300′. This more stringent criterion was used to ensure that the overshoot at the end of the saccade would be bypassed. The value of the criterion (300′/s) was determined empirically by examining a large sample of analog records of eye position. Saccades as small as the microsaccades that may be observed during maintained fixation (Steinman, Haddad, Skavenski, & Wyman, 1973) could be reliably detected by the algorithm.

The size of the first saccade was defined as the distance between the mean position of the eye at the start of the trial and the position of the eye at the end of the saccade. By using eye position at the start of the trial, rather than eye position at the onset of the detected saccade, the estimate of saccade size also incorporated any drift (Kowler & Steinman, 1979) that might occur during the latency interval. The mean latency and mean size of the first saccade will be reported below. Saccades after the first were also examined to look for evidence that a search of both clusters was attempted. An attempt to search the second cluster was deemed to have occurred if a saccade made after the first was opposite in direction to the first and large enough to cross the vertical midline of the display.

2.7. Number of sessions tested

Experimental sessions contained 50–100 trials. Usually, two sessions were tested in any given day.

2.8. Eliminated trials

Some trials could not be analyzed for any one of the following reasons: loss of tracker lock (2.6%), saccade made before the critical frame (3.6%), and enormous error (> 60% of eccentricity) of the first saccade (1.6%), suggesting that the first saccade was not actually aimed toward a cluster. Analyses were based on the following number of trials/subject: 1226 (88%) for AA, 519(87%) for BB, 313(89%) for CC, 953(95%) for DD, 757(95%) for EE, 1468 (98%) for FF.

3. Results

3.1. Importance of looking at the target with the first saccade

The duration of the critical frame was set to a value (0.5 s) that would be short enough to allow search of one location, but not both, thus making successful search contingent on appropriate saccadic planning. Fig. 2 (top graphs) shows that, as expected, a successful search depended on the direction of the first saccade. All subjects scored better than 80% correct (average across subjects = 90%) when the first saccade was made to the cluster that contained the target and less than about 30% correct when the first saccade was made in the opposite direction. This outcome confirms that the best strategy to ensure the highest proportion of correct reports is to use the probability cue to direct the first saccade to the cluster most likely to contain the target.

3.2. Direction of the first saccade

The best strategy is obvious, but only one subject used it consistently. Fig. 3 shows the probability that the first saccade was directed to the right as a function of the eccentricity of the right-hand cluster. Performance is shown separately for trials in which the target was more likely to appear on the right and on the left. Subject FF (lower right-hand graph) made the best use of probability information. FF's first saccades went to the right on almost every trial when the high-probability cluster was on the right, and to the left on the vast majority of trials when the high-probability location was on the left. FF was not infallible, as Fig. 2 shows, but overall, excellent use was made of the probability cues.

None of the others did as well. Subjects AA, DD and EE made some use of probability information (as shown by the separation between the functions for the two different probability conditions), but were clearly influenced by distance, showing a preference to look at nearby locations even when they were not likely to contain the target. CC's performance was determined almost exclusively by distance, while BB, also influenced by distance, showed a strong bias to make the first saccade to the right, regardless of distance or probability.

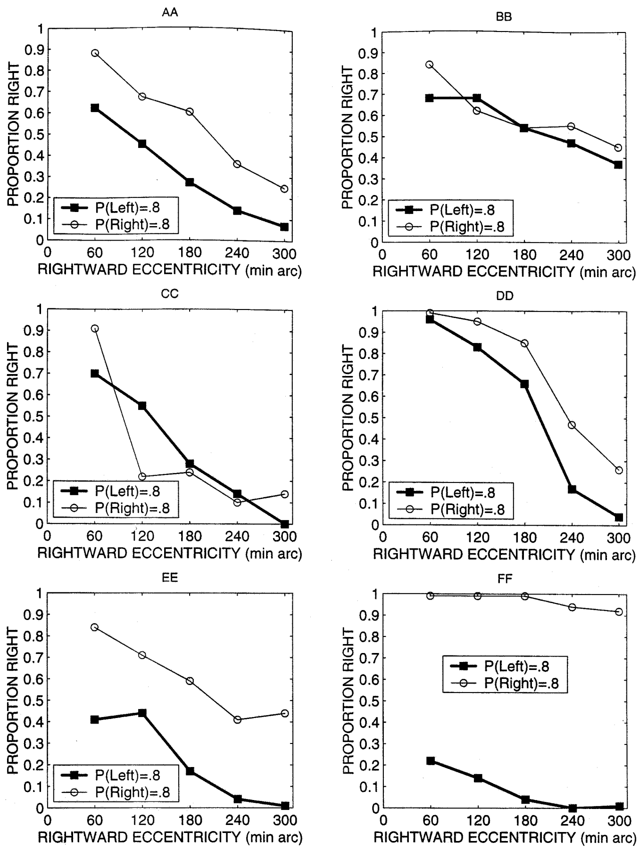

Fig. 3. Proportion of rightward saccades as a function of the eccentricity of the righthand cluster for all six subjects (AA–FF) for trials in which the target was likely (probability = 0.8) to be in the left-hand or right-hand clusters.

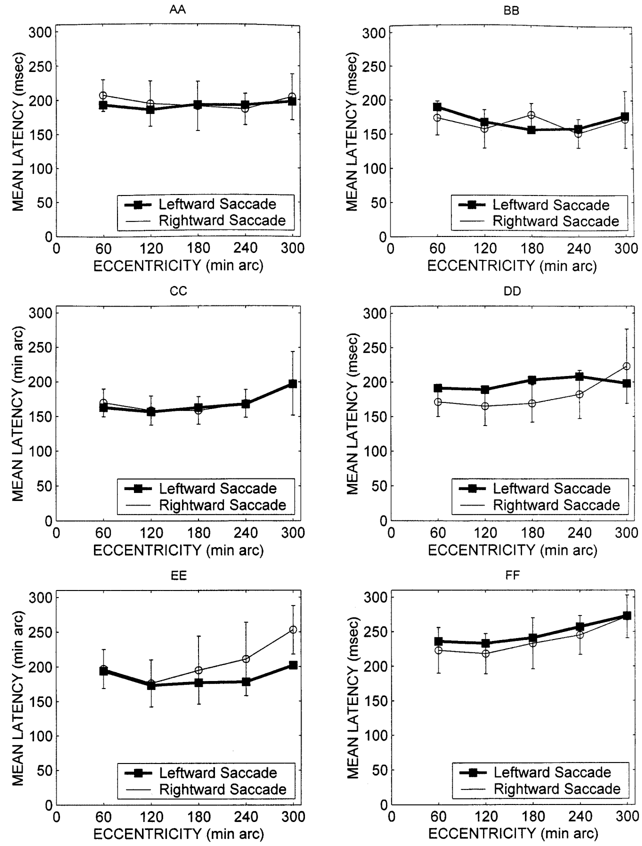

Fig. 4. Mean latency of the first saccade (leftward and rightward) as a function of eccentricity for all six subjects (AA–FF). Vertical bars represent ± 1 S.D., shown only for rightward saccades. Leftward S.D.s were similar.

Not surprisingly, the decision about where to direct the first saccade had a large effect on search performance, with FF doing much better than the other subjects at reporting target orientation (see overall proportion correct in Fig. 2, lower graph).

3.3. Latency and sizes of saccades

Preferences to look at nearby locations were not due to effects of distance on saccadic latency. Latencies did not vary appreciably as a function of eccentricity of the cluster (Fig. 4). Subject FF, the one who took probability into account almost all the time, had latencies somewhat longer (50 ms) than the others. Even with the longer latencies, enough time was left over for FF to identify the target.

Saccade sizes were affected by distance, with saccades to the larger eccentricities tending to undershoot, on average, and also to be more variable (Fig. 5). The effects of eccentricity on saccadic accuracy may have had some effect on the reports. The proportion of correct reports, given the subject looked at the cluster, fell somewhat with increasing target eccentricity for three of the subjects (CC, DD and EE), while for the other three, it remained about the same, regardless of the eccentricity of the cluster (Fig. 6). While the difficulties experienced by three of the subjects at the larger eccentricities may have discouraged an initial saccade to the distant target, it is clear that this was not a good strategy because the drop off in performance with eccentricity was small in comparison to the huge decrement observed when first saccades were aimed away from the cluster containing target (Fig. 2).

3.4. Effect of probability on latency

In contrast to the modest effects of eccentricity on the latency of saccades (Fig. 4), probability was quite influential. Fig. 7 shows that the four subjects who took probability into account when planning saccades (AA, DD, EE, FF; see Fig. 3) did so with clear cost: latencies of saccades to the high-probability locations were increased by as much as 50 ms, particularly at the farthest eccentricities. BB and CC showed no increase in latency, consistent with their pattern of ignoring the probability cue (Fig. 3). The increase in the latency of saccades to the high-probability location can be ascribed to the time required to decode the intensity cue and decide where to look. These extra operations may have discouraged taking probability into account, even though the magnitude of the latency increase was small relative to the duration of the critical display.

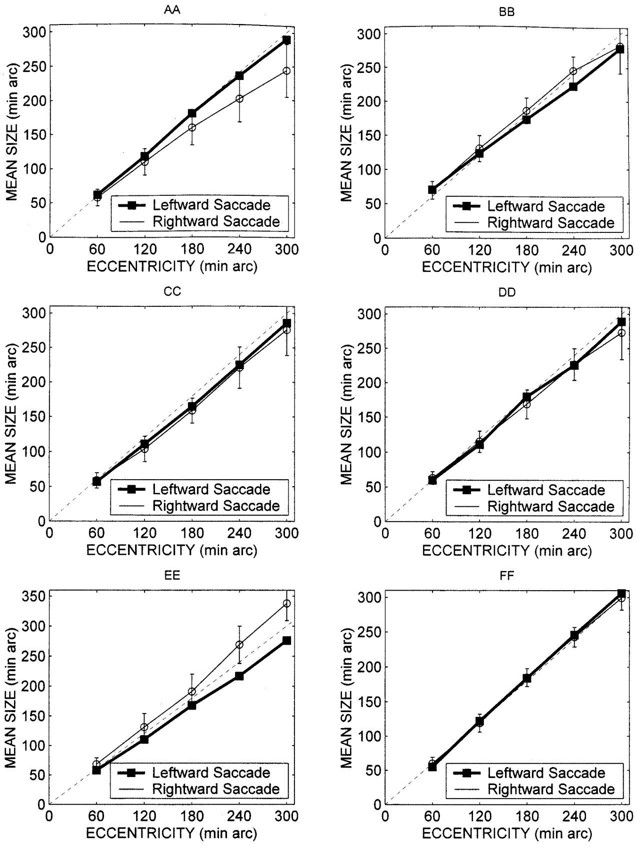

Fig. 5. Mean size of the first saccade (leftward and rightward) as a function of eccentricity for all six subjects (AA–FF). Vertical bars represent ±1 S.D., shown only for rightward saccades. Leftward S.D.s were similar.

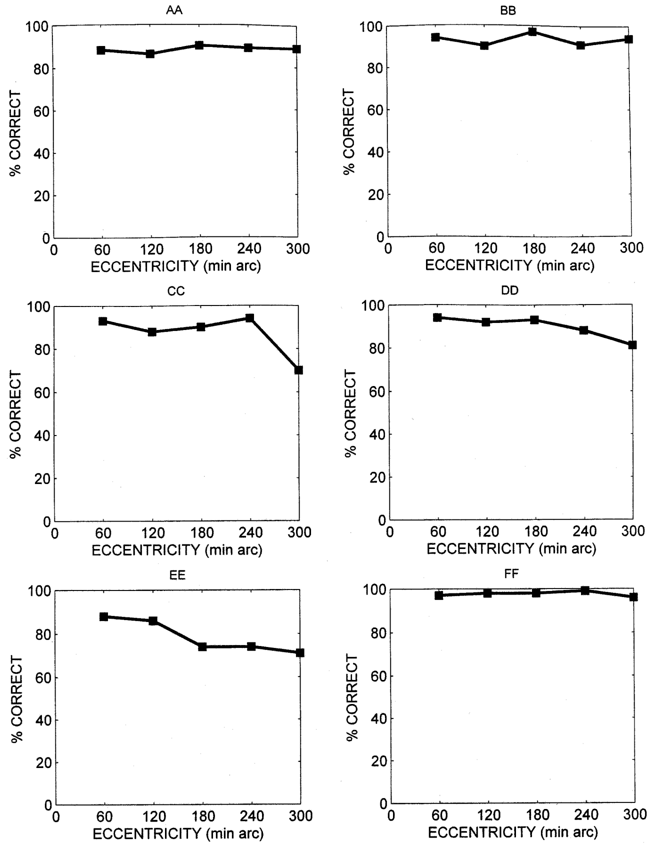

Fig. 6. Proportion of correct reports as a function of the eccentricity of the cluster in which the target was located, given that the first saccade was made to the cluster containing the target, for all six subjects (AA–FF).

3.5. One- vs. two-location search

In contrast to the reluctance to delay the first saccade, there was considerable enthusiasm for attempting to search both locations. We defined an attempted two-location search as one in which the onset of the saccade to the second location occurred within 200 ms after offset of the critical frame. This liberal criterion allowed us to capture attempts that failed because they were too late, but nevertheless represented genuine efforts to look at both locations.

Except for FF, who rarely searched both locations, the rest of the subjects searched, or attempted to search, both locations on about one-third to one-half of the trials (see Proportion of two-location searches in Table 1). On the vast majority of the trials (76–95%) in

Table 1
Comparison of one- and two-location searches

Subject	One-location search: Proportion correct	Proportion of one-location searches	Proportion correct given target in first location searched	Two-location search: Proportion correct	Proportion of two-location searches	Proportion correct given target in first location searched	Proportion correct given target in second location searched
AA	0.79	0.68	0.88	0.38	0.32	0.30	0.40
BB	0.76	0.70	0.94	0.24	0.30	0.67	0.22
CC	0.76	0.54	0.89	0.19	0.46	0.80	0.17
DD	0.81	0.58	0.92	0.38	0.42	0.76	0.31
EE	0.66	0.71	0.80	0.35	0.29	0.43	0.34
FF	0.83	0.94	0.98	0.72	0.06	0.63	0.74

which a search of both locations was attempted, the first saccade had been directed to a nearby [eccentricity 180′ or less], low-probability location. And, of those trials in which the first saccade was directed to a low-probability location, the majority contained a subsequent saccade to search the second location (Fig. 8).

The strategy of looking at two locations seldom led to a correct report of target orientation. Table 1 shows the proportion of correct reports in trials with one- and two-location searches. It also provides a further breakdown of performance by showing the proportion of correct reports when the target was in either the first or second location searched. Performance was clearly poorer for two-location searches, regardless of which location contained the target. The poor performance occurred because there was too much to do—including planning and executing both saccades, and completing the required visual analyses of the both clusters—in the limited time available. This low level of performance for two-location searches does not justify the preference for following such a strategy so often.

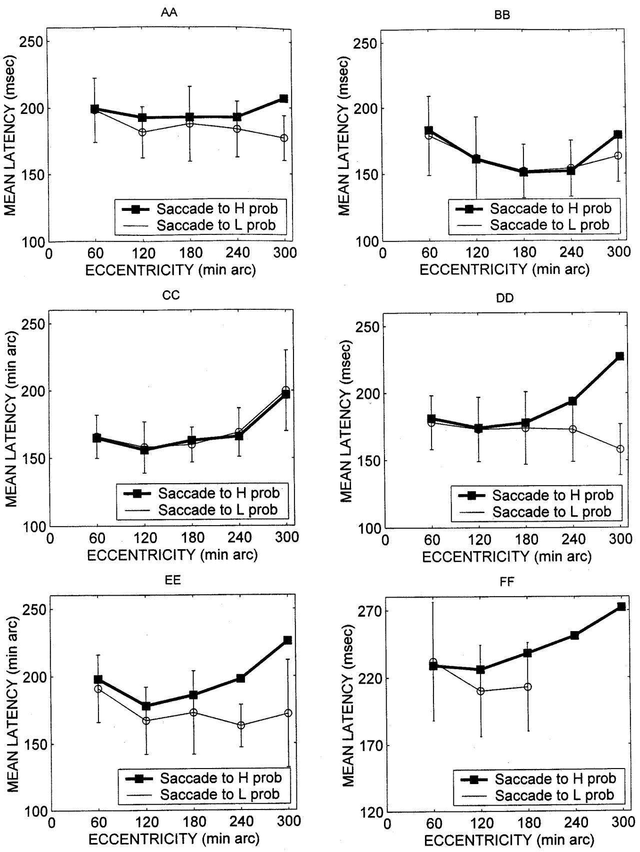

Fig. 7. Mean latency of the first saccade to high and low probability locations as a function of eccentricity for all six subjects (AA–FF). Vertical bars represent ±1 S.D., shown only for saccades to the low probability location. S.D.s to high-probability locations were similar.

3.6. *What prompted the two-location search?*

A search of the second location could have been initiated by a failure to find the target at the first location searched. Alternatively, the sequence of two searching saccades could have been planned at the outset of the trial, in an optimistic attempt to capture both of the locations before the critical frame was taken away. Such a strategy would surely lead to an accurate report, if only there were enough time.

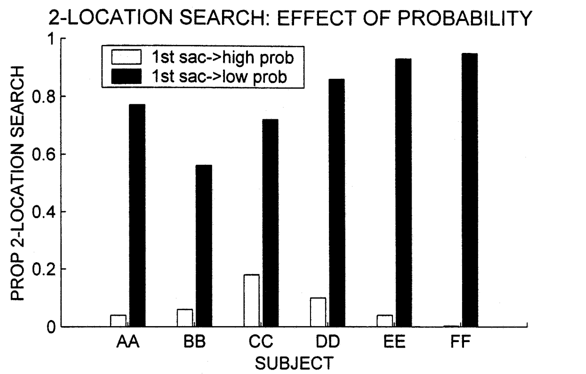

Fig. 8. Proportion of trials with an attempted search of the second location when the first saccade was directed to the cluster that was likely (probability = 0.8, open bars) and not likely (probability = 0.2, filled bars) to contain the target for all six subjects (AA–FF).

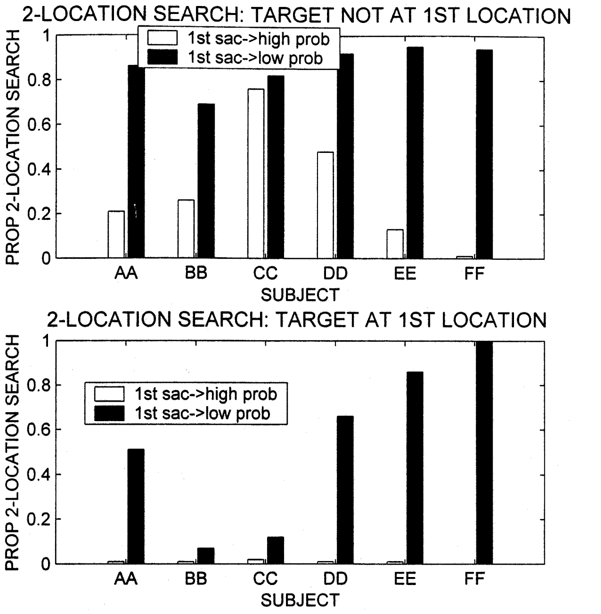

Fig. 9. (Top) Proportion of trials with an attempted search of the second location when the first saccade was directed to the cluster that was likely (probability = 0.8, open bars) and not likely (probability = 0.2, filled bars) to contain the target, given that the target was not at the first location searched, for all six subjects (AA–FF). (Bottom) Same, except the target was present at the first location searched.

To find out whether the search of the second location was planned in advance, or, alternatively, was provoked by failure to detect a target at the first location inspected, we examined the proportion of two-location searches in trials where the first saccade was directed to locations that either did or did not contain the target. If the absence of the target triggered the search of the second location, then the probability level associated with the first location (0.8 or 0.2) should not matter.

Fig. 9 (top graph) shows that the probability level was quite important. Attempted searches of the second location occurred much more often when the target was missing from a low-probability location than when it was missing from a high-probability location. The effects of probability were large for all subjects except CC, who was likely to try to search the second location when the target was missing from the first, regardless of the probability levels. The performance of the other five subjects shows that they took probability into account in planning the two-location search. In fact, the effect of probability was so strong that even when the low-probability location contained the target, four of the six subjects frequently (> 50%) went on to search the second location (Fig. 9, bottom). Comparing the proportion of two-location searches when the target was either present or absent at the first location (i.e. comparing the top and bottom graphs in Fig. 9) shows that four subjects (AA, DD, EE, FF) frequently carried out two-location searches when the first saccade was directed to a low-probability location, regardless of whether a target was discovered at the first location.

These results are consistent with the conclusion that two-location searches could be, and often were, initiated as part of a global plan (Zingale & Kowler, 1987), rather than occurring at the spur of the moment in response to the failure to find a target at the first location searched.

3.6.1. Distance minimization

It is possible that the two-location searches, which most often began with the close, low-probability location and ended at the more distant high-probability location, reflected a plan to search both locations in a way that minimized the total distance covered (i.e. minimize the sum of the sizes of the two saccades). Monkeys use such a strategy when collecting food hidden at various locations (Cramer & Gallistel, 1997). Minimizing distance is important with time-consuming activities, such as moving the limbs, so perhaps such preferences (in the interest of ensuring coordination among motor systems) carry over to the faster saccades.

We evaluated distance minimization using the results of the four subjects whose two-location searches were not solely responses to a failure to find a target at the

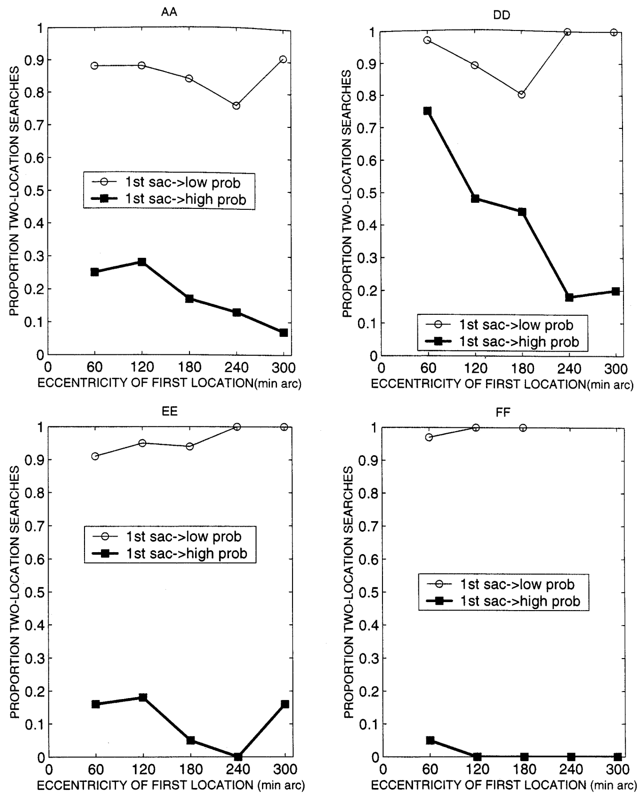

Fig. 10. Proportion of two-location searches as a function of the eccentricity of the first location searched when the first saccade was directed to a cluster that was likely (probability = 0.8) and not likely (probability = 0.2) to contain the target for four subjects who made two-location searches frequently enough to permit meaningful analyses.

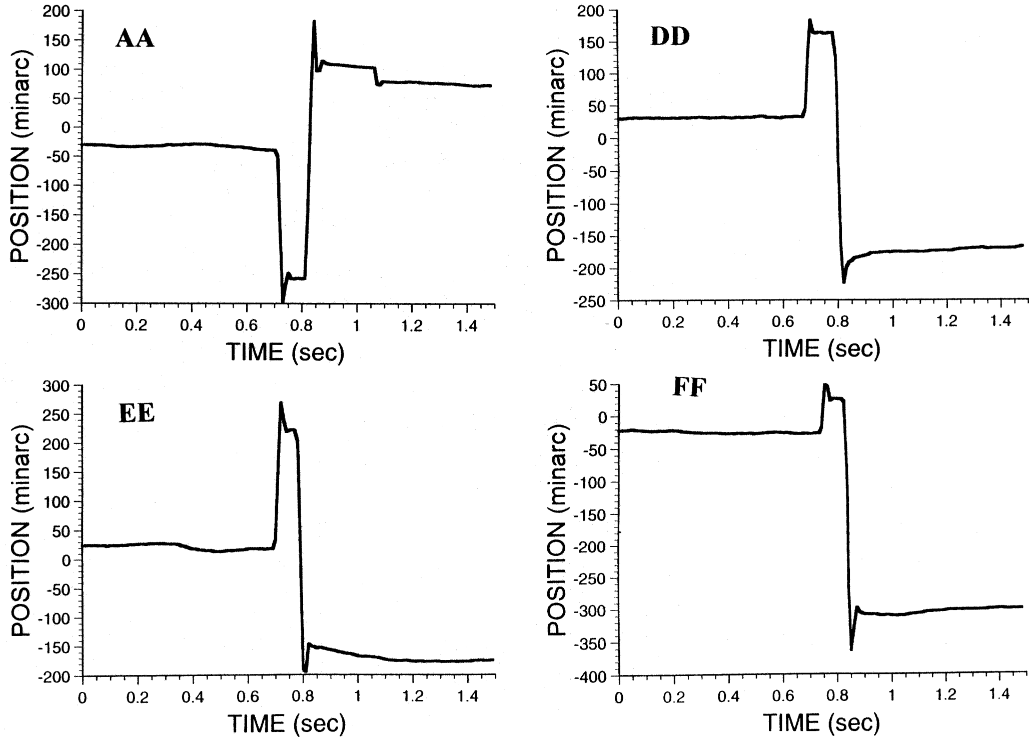

Fig. 11. Examples of very brief inter-saccadic intervals during two-location searches. The trial starts at time = 0. The critical frame containing the target appeared 0.5 s after trial onset and was displayed until 1 s after trial onset.

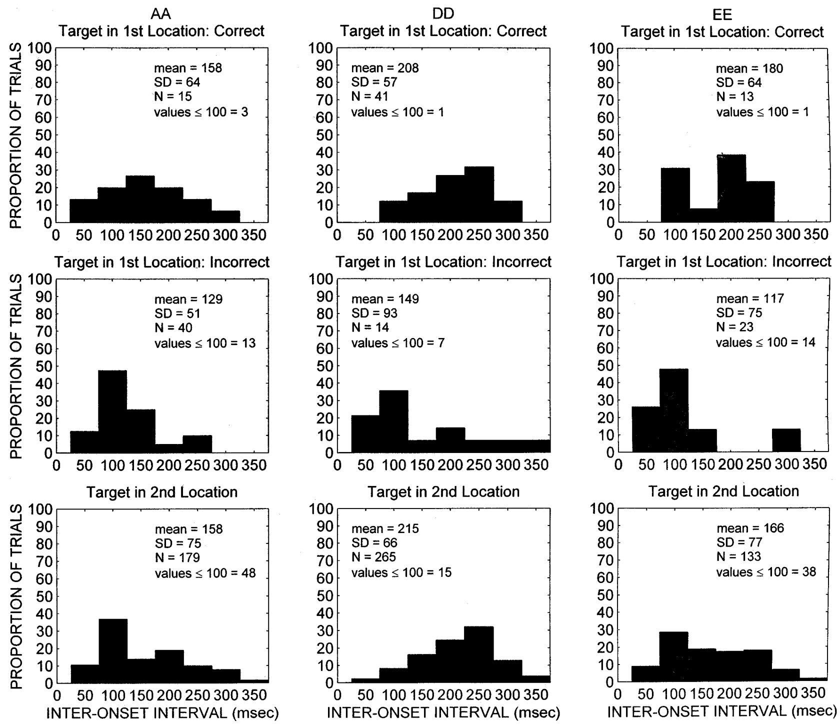

Fig. 12. Distributions of intervals between the onsets of successive saccades for two-location searches for three subjects. Trials with the target in the first location and a correct report are shown in the top graphs, trials with the target in the first location and an incorrect report in the middle row and trials with the target in the second location searched in the bottom row. Data are shown for three subjects (AA, DD and EE) who attempted two-location searches frequently enough to permit meaningful analyses.

first location examined. Fig. 10 shows the probability of a two-location search, given that the first saccade was directed to either a high- or low-probability cluster at different eccentricities. In order to obtain a fair assessment of the effect of probability, analyses were restricted to those trials where no target was present at the first location searched. If two-location searches were the result of a strategy of distance minimization, we would expect those searches to occur most frequently when the eccentricity of the first location was small. Instead, Fig. 10 shows that eccentricity played some role, but was not very important when the first saccade was directed to the low-probabflility location. Something other than distance minimization was responsible for the two-location search.

3.6.2. Very brief inter-saccadic intervals

The attempted two-location searches produced a substantial number of trials with very brief intervals (100 ms or less) between saccades. Fig. 11 shows some examples of these brief pauses. Distributions of intervals between the onsets of successive saccades are shown in Fig. 12. The trials shown are limited to those in which the second saccade occurred while the critical frame was still visible, so the intervals represent the time available to process the first location searched. Only three subjects made such saccades frequently enough to permit meaningful analyses, and only their distributions are shown. The trials are broken down into three categories: those in which the target in the first location was reported correctly, those in which it was reported incorrectly, and those in which the target was not present in the first location.

The inter-onset intervals were shortest for trials in which the target in the first location was reported incorrectly. This implies that the failure to identify the target in these trials was due to the brevity of the intersaccadic pause, as the subject rushed ahead to the second location, lending further support to the notion that many of the two-location searches were planned in advance. The instances of very brief pauses between the two saccades is characteristic of pre-planned sequences (Zingale & Kowler, 1987), which sometimes appear during search (Viviani & Swensson, 1982; Hooge & Erkelens, 1996; McPeek, Nakayama, & Skavenski, 2000; Findlay, Brown, & Gilchrist, 2001). It clearly was not an effective strategy in our task: not only was it difficult to complete a search of the second location in the allotted time, but we now see that the time available to search the first location was compromised as well.

4. Discussion

4.1. Summary of major findings

The best strategy to adopt in this two-location search task was to use the probability cue to aim the first saccade to the location most likely to contain the target. Only one of the six subjects tested followed this strategy consistently. The others often aimed the first saccade to the nearer location, regardless of whether it was likely to contain the target. Since there was ample time to process the probability cue, make an accurate saccade, and recognize the target's orientation, it would have been surprising to find that any of the six subjects frequently ignored information about probability when planning the first saccade, let alone finding that nearly all of them did.

Most saccades that were aimed to a low-probability location were followed by a second saccade to the high probability location. In four of the six subjects, two-location searches were planned and executed with little regard for the visual information acquired during the intersaccadic pause, showing that the search of the second location was not simply a response to failure to find a target at the first location. The two-location searches were seldom successful. This is because there was not enough time to look at both locations and to complete the necessary visual processing.

Why would any of the participants in this experiment make an initial saccade away from the more probable location, and what does an analysis of the strategies reveal about mechanisms of saccadic planning and execution? We identify and critically examine four factors, any or all of which could have influenced the saccades.

4.1.1. Failure to appreciate the significance of the probability cue

Saccades could have been influenced by decisions or strategies that underestimated the value of using the probability cue to direct saccades. The literature on decision-making is replete with examples of failure to make optimal decisions, often because biases provoked by the particular circumstances surrounding a choice hindered the appreciation of the true probabilistic structure of the decision problem (e.g. Kahneman, Slovic, & Tversky, 1982; Collier, Johnson, & Berman, 1998). However, optimal performance has been observed in relatively simple situations. Using a multi-location visual search task in which eye movements were not permitted, Shaw and Shaw (1977) found that attention was distributed over the display locations in proportion to the probability of target appearance, a strategy that they showed would produce optimal performance. In that situation, however, unlike ours, probabilities were not explicitly cued, but were instead disclosed by the past history of target appearance.

In our experiment, the presence of the target at the low-probability location in a noticeable portion of the trials (20%) could have encouraged a strategy of neglecting the cue and relying instead on the immediate past history of target locations to predict where the target might appear next. Such a strategy would produce occasional saccades to the low probability location. Prior work has shown that past history can affect both anticipatory smooth eye movements and saccades during target step-tracking tasks (Kowler, Martins, & Pavel, 1984), so an effect of past history in the present experiment would not be surprising. Planning saccades based on immediate past history can explain neglect of the probability cue, but cannot account for the preference to look at the closer location.

4.1.2. Preference to avoid effortful saccadic planning

Choosing to use the probability cue, as four of the six subjects did to varying degrees, took some effort because the cue had to be detected and interpreted on each trial. The extra effort was reflected in the longer latencies of saccades to the high-probability location (Fig. 7). Looking at the closer location was easier, reducing the computational load associated with choosing the goal position of the saccade. Preferences to reduce the computational load associated with saccadic planning have been noted before (Zelinsky, 1996; Hooge & Erkelens, 1998, 1999; Melcher & Kowler, 2001) in tasks that required careful inspection of the contents of a display. There was clear incentive to minimize effortful saccadic planning in these tasks because attention paid to planning reduced the attentional resources available to analyze the foveal view. Minimizing planning effort would not seem to be important in our study because we did not require processing of foveal information before the first saccade. Nevertheless, subjects may have acted as if they needed to limit the resources allocated to planning saccades, perhaps because of built-in preferences or learned habits. Minimization of effortful saccadic planning could be a default option that requires an explicit decision to override.

4.1.3. Greater signal strength of nearby targets

Analysis of eye movements made in different visual tasks shows that in the absence of incentives to choose one or another saccadic path, saccades often follow a path that systematically takes the line of sight to nearby locations (Epelboim et al., 1995; Hooge & Erkelens, 1996; Sommer, 1997; Motter & Belky, 1998; Melcher & Kowler, 2001). Preferences for nearby locations in those prior studies, as well as in the present experiment, may stem from visual factors, such as the higher visibility or higher resolution of elements nearer to the fovea. These nearby elements could provide stronger input signals to the attentional or saccadic systems and thus increase the likelihood of attracting the saccade in any competitive 'race' between attended locations (Kowler, Anderson, Dosher, & Blaser, 1995; Theeuwes, Kramer, Hahn, Irwin, & Zelinsky, 1999). We think that visual factors contributed to redistributing attentional weights, rather than evoking an automatic or reflexive attention shift or saccade, because there already was considerable incentive to attend to each location due to the structure of the task (i.e. the target could have been located in either cluster and each cluster contained information about probability). Salient stimuli tend not to attract attention (Yantis & Egeth, 1999) or saccades (Kowler & Steinman, 1979; He & Kowler, 1989) unless they are relevant to the task, as our stimuli were.

By assuming differences in relative attentional strength due to stimulus eccentricity, the notion of the effort required to program saccades presented above can be made more concrete. Effort represents the cost in processing resources and time (Fig. 7; see also Sperling & Dosher, 1986) required to interpret the probability cue, redistribute the attentional weights according to probability rather than eccentricity, and direct the saccade to the more probable location.

4.1.4. Two-location searches: divided attention

So far, we have discussed the factors that can affect the planning of the first saccade. Our results also showed that four of the six subjects attempted a rapid search of the both locations quite frequently, without regard for whether the target was discovered at the first location. The saccadic pairs making up the two-location searches resemble pre-planned sequences (Zingale & Kowler, 1987), in which the preparation of the second saccade begins before the first is completed (Sommer, 1997; Theeuwes et al., 1999; McPeek et al., 2000).

An explanation for the two-location searches needs to take into account that such searches were frequent only when the first saccade was directed to the low-probability location, regardless of whether the target was present at the first location searched (Fig. 9) and regardless of the eccentricity of the low probability location (Fig. 10). The effect of probability on the number and order of locations searched could have been a consequence of a division of attention between the clusters according to the following scheme. Immediately after display onset, it is likely that some attention is paid to each cluster so that its location and probability level can be encoded. The attentional strength assigned to each cluster would depend on factors we have already discussed, namely, interpretation of the probability cue, cluster eccentricity, and a subjective assessment of the likelihood that either location would contain the target. Attentional strength could change over time as analysis of the probability cue proceeds. The saccade would be drawn to whichever location had the higher attentional weight during the critical portion of the latency period (Kowler et al., 1995; Theeuwes et al., 1999). If we assume some random variation in assigned attentional strength, then either location could dominate and become the target of the first saccade in any particular trial. In cases where the first saccade is made to the low-probability location, attention would continue to be allocated to the other, high-probability location even after the first saccade is completed because of its continued relevance to the task, allowing the high-probability location to quickly attract a second saccade. When the first saccade is made to the high-probability location, however, the remaining low-probability location preserves little of its original attentional strength (whatever strength it had due to proximity to the line of sight would be diminished after the first saccade took the eye to a new location). A second saccade to the low probability location would thus be rare.

The scheme we have outlined is similar to the 'race model' proposed by Theeuwes et al. (1999) and the 'pipeline programming model' by McPeek et al. (2000), both based on the idea that programming of the second saccade of a pair can begin before the programming and execution of the first saccade is completed. In their experiments, both attention and the initial saccade were drawn to an irrelevant item either because of its abrupt appearance in the display (Theeuwes et al., 1999) or because it shared features with relevant stimuli seen on prior trials (McPeek et al., 2000). In our experiment, both locations were relevant, but were assigned different attentional strengths depending on eccentricity and the probability of containing a target.

A noteworthy aspect of the divided attention proposal we have outlined is that by allowing attention to be distributed in parallel among the display elements, and by allowing the attention allocated to at least some elements to persist across saccades, it is possible for an orderly temporal pattern to be imposed on a saccadic sequence without explicitly representing order in the original saccadic plan. This is a useful feature that can lead to more efficient, less effortful scanning, and may be broadly applicable to natural oculomotor tasks encountered in daily life.

4.2. Summary and conclusions

Saccadic eye movements have long been of interest because they hold the promise of revealing underlying cognitive operations as they unfold over time, phenomena not readily observable by other means. The easiest way to use eye movements to infer underlying processing events is to assume that the eye-movement patterns reflect the momentary demands of the task. This assumption has been challenged in the past (O'Regan, 1990; Suppes, 1990; Viviani, 1990), and our results add to the questions raised about how best to link saccades and cognitive task demands.

Using a simple, two-location visual search task we found that saccadic patterns were influenced by a stimulus characteristic–spatial distance—that was not related to the probability of finding the target. The best strategy—looking directly to the location most likely to contain the target—was rejected in a large proportion of trials.

We proposed several factors to account for the saccadic patterns observed: (1) decision strategies that failed to recognize the significance of probability cues, (2) built-in preferences to minimize effortful saccadic planning, (3) attraction of attention and saccades to nearby locations, and (4) initiation of saccades while attention remained divided between cluster locations. Any or all of these could have contributed to the saccadic performance we observed.

There is a common characteristic that links the various factors we have proposed as possible contributors to the formulation of saccadic plans. All act to facilitate rapid scanning of a series of locations, in contrast to a deliberate, one-by-one selection of saccadic goals. The saccadic system, through its links to attention and its capacity to launch eye movements at a rapid rate, is well suited to execute high-speed, sequential scans of potentially important locations. This capacity could well be important in natural tasks. The risk of some errant glances at unimportant places may be a small price to pay for having a system that is able to deliver a large number of new, foveal views in a limited span of time with minimal effort devoted to saccadic computation and plans.

Acknowledgements

Supported by the Air Force Office of Scientific Research, Life Sciences Directorate, F4960-99-1-0080 to E. Kowler.

References

Ballard, D. H., Hayhoe, M. M., & Pelz, J. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, *7*, 66–80.

Collier, G., Johnson, D. F., & Berman, J. (1998). Patch choice as a function of procurement cost and encounter rate. *Journal of the Experimental Analysis of Behavior*, *69*, 5–16.

Cramer, A. E., & Gallistel, C. R. (1997). Vervet monkeys as travelling salesman. *Nature*, *387*, 464.

Crane, H. D., & Steele, C. S. (1978). Accurate three-dimensional eyetracker. *Applied Optics*, *17*, 691–705.

Epelboim, J., Steinman, R. M., Kowler, E., Edwards, M., Pizlo, Z., Erkelens, C. J., & Collewijn, H. (1995). The function of visual search and memory in sequential looking tasks. *Vision Research*, *35*, 3401–3422.

Epelboim, J., & Suppes, P (2001). A model of cognitive processes during eye movements in geometry. *Vision Research*, *41*, 1561–1574.

Findlay, J. M., Brown, V., & Gilchrist, I. D. (2001). Saccade target selection in visual search: the effect of information from the previous fixation. *Vision Research*, *41*, 87–95.

He, P., & Kowler, E. (1989). The role of location probability in the programming of saccades: implications for 'center-of-gravity' tendencies. *Vision Research*, *9*, 1165–1181.

Hooge, I. T. C., & Erkelens, C. J. (1996). Control of fixation duration in a simple search task. *Perception & Psychophysics*, *58*, 969–976.

Hooge, I. T. C, & Erkelens, C. J. (1998). Adjustment of fixation duration in visual search. *Vision Research*, *38*, 1295–1302.

Hooge, I. T. C., & Erkelens, C. J. (1999). Peripheral vision and oculomotor control during visual search. *Vision Research*, *39*, 1567–1575.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty*: *heuristics and biases*. New York: Cambridge University Press.

Kowler, E., Anderson, E., Dosher, B., & Blaser, E. (1995). The role of attention in the programming of saccades. *Vision Research*, *35*, 1897–1916.

Kowler, E., Martins, A. J., & Pavel, M. (1984). The effect of expectations on slow oculomotor control—IV: anticipatory smooth eye movements depend on prior target motions. *Vision Research*, *24*, 197–210.

Kowler, E., & Steinman, R. M. (1979). The effect of expectations on slow oculomotor control—I. Periodic target steps. *Vision Research*, *19*, 619–632.

Land, M. F. & Hayhoe, M. M. (2001). In what ways do eye movements contribute to everyday activities? *Vision Research*, *41*, 3559–3565.

McPeek, R. M., Nakayama, K., & Skavenski, A. A. (2000). Concurrent processing of saccades in visual search. *Vision Research*, *40*, 2499–2516.

Melcher, D. & Kowler, E. (2001). Visual scene memory and the guidance of saccadic eye movements. *Vision Research 41*, 3597–3611.

Motter, B. C., & Belky, E. J. (1998). The guidance of eye movements during active visual search. *Vision Research*, *38*, 1805–1815.

O'Regan, J. K. (1990). Eye movements and reading. In E. Kowler, *Eye movements and their role in visual and cognitive processes* (pp. 395–453). Amsterdam: Elsevier.

Shaw, M. L., & Shaw, P. (1977). Optimal allocation of cognitive resources to spatial locations. *Journal of Experimental Psychology*, *3*, 201–211.

Sommer, M. A. (1997). The spatial relationship between scanning saccades and express saccades. *Vision Research*, *37*, 2745–2756.

Sperling, G., & Dosher, B. A. (1986). Strategy and organization in human information processing. In K. R. Boff, L. Kaufman, & J. P. Thomas, *Sensory processes and perception*. In: *Handbook of perception and human performance*, vol. 1. New York: Wiley Chapter 2.

Steinman, R. M., Haddad, G. M., Skavenski, A. A., & Wyman, D. (1973). Miniature eye movements. *Science*, *81*, 810–819.

Suppes, P. (1990). Eye movements models for arithmetic and reading performance. In E. Kowler, *Eye movements and their role in visual and cognitive processes* (pp. 455–477). Amsterdam: Elsevier.

Theeuwes, J., Kramer, A. F., Hahn, S., Irwin, D. E., & Zelinsky, G. J. (1999). Influence of attentional capture on oculomotor control. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 1595–1608.

Viviani, P. (1990). Eye movements in visual search: cognitive, perceptual and motor aspects. In E. Kowler, *Eye movements and their role in visual and cognitive processes* (pp. 353–393). Amsterdam: Elsevier.

Viviani, P. & Swensson, R. G. (1982). Saccadic eye movements to peripherally descriminated visual targets. *Journal of Experimental Psychology*: *Human Perception and Performances*, *8*, 113–126.

Yantis, S, & Egeth, H. E. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 661–676.

Zelinsky, G. J. (1996). Using eye saccades to assess the selectivity of search movements. *Vision Research*, *36*, 2177–2187.

Zingale, C. M., & Kowler, E. (1987). Programming sequences of saccades. *Vision Research*, *27*, 1327–1341.