

Variable search for orientation, uniformly optimal search for identity

Anna Nowakowska ^{1,4}, Alasdair D.F. Clarke ², Josephine Reuther ³, & Amelia R. Hunt ¹

¹ School of Psychology, University of Aberdeen

² School of Psychology, University of Essex

³ Department of Experimental Psychology, University of Göttingen

⁴ School of Psychology and Vision Sciences, University of Leicester

Supplementary materials are available at: <https://osf.io/9edx6/>

Correspondence should be addressed to: a.nowakowska@leicester.ac.uk

Anna Nowakowska

School of Psychology and Vision Sciences, University of Leicester

Abstract

We compare eye movement strategies across a range of different stimulus sets to test the prediction that eye movements are guided by expected information gain. When searching for a simple target that has been defined based on orientation, inter-individual variability is high, and a large proportion of eye movements are directed to locations where peripheral vision would have been sufficient to determine whether the target was present there or not. In contrast, when searching for a target defined based on identity, eye movements are similar across individuals and highly efficient, being directed almost exclusively to the locations where central vision is most needed. The results suggest that for most people, the way they search for a simple feature (orientation) is not directly representative of the way they search for objects based on their identity. More generally, the results highlight that because humans are adaptable, contradictory theories can be accurate descriptions of search in particular contexts and individuals. For a complete and accurate account of human search behaviour to be achieved, the conditions that shift us from one mode of behaviour to another need to be part of our models.

Keywords: Visual search, eye movements, strategy, individual differences

Word count:11950

Eye movements can be characterised as a series of rapid decisions that both determine, and are determined by, the flow of information into the visual system. This process of goal-driven information sampling is often studied in the context of visual search tasks, in which the target and the distractors or the background can be systematically manipulated to identify the factors that influence performance. Of particular interest is identifying general-purpose strategies or principles that drive the sequential selection of locations to fixate. One class of visual search theories is grounded in *guidance*, that is, how information about the scene and what the target looks like can be used to narrow down the search set to a smaller set of candidates that are then sampled sequentially, either by attention alone (if eye movements are not needed) or by attention and eye movements (if central vision is required). Our understanding of guidance is informed by Feature Integration Theory (Triesman & Gelade, 1980), which identifies the kinds of visual information that can be processed in the absence of focused attention and therefore be useful for guiding attention during search of scenes. These features, termed “pre-attentive”, include simple visual features like colour and orientation. Relatedly, scene context (Neider & Zelinsky, 2006) and predictable spatial relationships (Chun & Jiang, 1999) guide attention and the eyes to more plausible or likely target locations. The implementation of many of these guidance processes is formalized in the Guided Search model (Wolfe, 2021), which also elaborates on details specific to search, such as the comparison of visual details against stored templates to determine if a candidate is the target or not.

A distinct class of visual search theories addresses the extent to which attention and eye movements are guided by more context-independent strategies, such as a principle of maximizing information gain (Gottlieb, 2023). One specific example is a model devised by Najemnik and Geisler (2005; 2008) that minimizes the number of eye

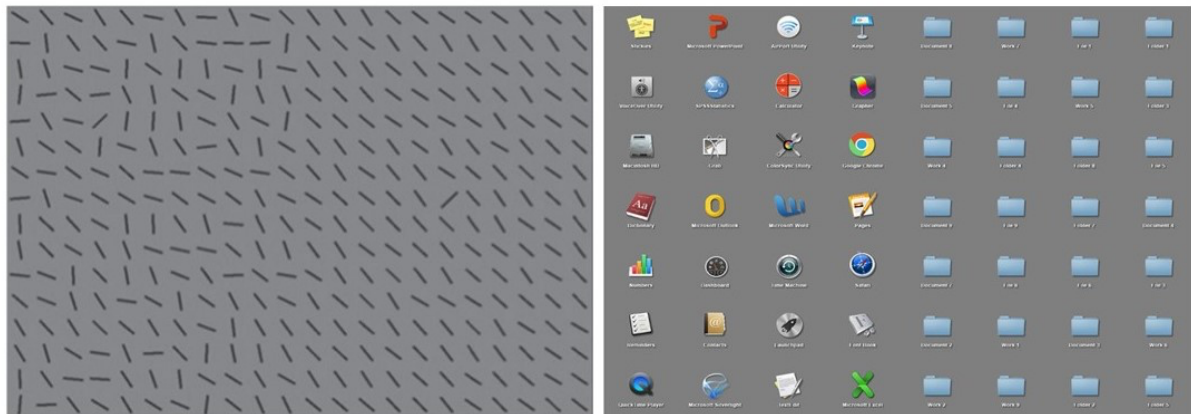
movements needed to find a Gabor patch hidden in visual noise, by selecting each fixation in a sequence based on an estimate of which location produces the largest information gain. Their model closely matched human searchers in how many fixations were executed before responding, providing some indirect evidence that expected information gain plays a role in determining fixation selection. However, human performance was also well-described by a model in which the computationally-expensive process of estimating information gain was replaced with a simple process of random sampling from a population of fixations (Clarke et al., 2016). The explanation for this seeming contradiction is that human searchers have effective general-purpose habits and biases, such as a tendency to fixate the center of images more than the edges (Clarke & Tatler, 2014). By sampling from the population of eye movements human searchers tend to make, the model contains at least some of these habits and biases, thereby producing search performance on par with a noisy optimal model, without having to compute expected target visibility across the visual field or keep track of which parts of the field had been sampled with which parts of the retina.

Given that the behaviour of these two different models provides equivalent good matches to human search behaviour, discriminating between theories seems to require a more experimental approach. However, the experimental evidence for one or the other account of eye movement selection is also mixed. While some research suggests we can choose eye movements that maximise information gain (e.g., Hoppe & Rothkopf, 2019), other studies reveal profound failures to do so (e.g., Morvan & Maloney, 2012; Verghese, 2012). In an attempt to resolve this debate, Nowakowska et al. (2017) devised the *split-half visual search* paradigm (Figure 1, left panel). Participants searched for a vertical line rotated 45° clockwise among an array of distractor lines. On one side of the array the variation in orientations of the distractors was narrow (*homogeneous*), so the target

popped out if it was present on that side. On the other side the variation was much wider (*heterogeneous*), requiring foveal inspection of the elements on that side in order to determine if the target was present. The optimal strategy for searching these stimuli is to only fixate the heterogeneous (hard) side. Consistent with the stochastic model (Clarke et al., 2016), the results showed nearly half of the first five saccades on each trial were made to the homogeneous (easy) side. Underlying this average, however, was large variation between participants, with some performing near optimally, some highly inefficiently, and the rest falling in between. The conclusion was that different individuals in the sample produced eye movement behaviour that was consistent with different visual search models.

Figure 1.

Example Trials of the Split-Half-Line Segment (SHLS) and Icon Experiments



Note. The left panel is an example of the split-half line segment (SHLS) stimuli used in the current set of experiments and in previous work described in the main text. The right panel is the icon version of the task, introduced in the current paper. For both, the heterogeneous side is on the left and the homogeneous side is on the right.

While these individual differences offer a potential resolution of contradictory evidence for competing theories of fixation selection during search, they also lead to

harder questions about why these individual differences arise and how to construct a model of search that accounts for them. Counter to the idea that differences in motivation or speed-accuracy trade-offs might explain such variation, neither using tight response deadlines, nor rewarding faster responses financially, reduced variability or improved the inefficient searchers' performance (Nowakowska et al., 2021). Given that large individual differences have been reported by other researchers across a range of visual search tasks (Araujo et al., 2001; Irons & Leber, 2016; Clarke et al. 2020; Lonnqvist et al., 2020), another tempting explanation is an "ability trait", like those reported in face recognition (Russell et al., 2009; Zhu et al., 2010) and visual comparison processes (Cooper, 1976). The intuitively appealing notion that some people are just better than others at searching in a general sense can also be ruled out: Search performance on the split-half task has high test-retest reliability, but is weakly correlated, at best, with performance on the other search tasks (Clarke et al., 2022).

A lack of correspondence in performance from one search task to another, despite the individual differences on each of these tasks separately being stable over time, suggests these particular search tasks tap into unique sources of variability. Aside from all involving search in some way, the tasks compared in Clarke et al. (2022) were otherwise quite different one from the next, and what was considered an effective strategy in one of these tasks bore no resemblance to an effective strategy on the other two tasks. Whether shared variance could be observed in more similar search tasks or with individual characteristics is a question several investigations have started to address (e.g., Li et al., 2022; Clark et al. 2022; Clarke et al., accepted pre-registration). On a similar line, the impetus for the current series of experiments, at least at the start, was the extent to which search strategies generalise from simple line segments to more complex and realistic search arrays. Nowakowska et al (2017) and the many related

investigations (Nowakowska et al. 2016; 2017; 2019; 2021) used line segments as search stimuli as they are considered to be a primitive feature, forming the basis for figure-ground assignment and object recognition (Rogers, 2017). The implicit assumption is that findings with line segments readily generalise to assemblies of line segments, such as objects in scenes. Moreover, the human visual system is highly familiar with line segments, and this high familiarity across the board should reduce sampling noise due to individual differences in experience. However, the vast individual differences observed in Nowakowska et al. (2017) and replicated in subsequent studies is a challenge for the idea that simplified stimuli evoke a basic, uniform response upon which a model for more complex scenes could be built. We therefore sought to measure the extent to which these individual differences would generalise across more closely related contexts, as a basis for building a causal explanation for them.

To that end, we devised new stimuli for the split-half search task (Figure 1, right panel), designed to mimic the task of searching a computer desktop for a specified target icon. On one half of the desktop, the icons are all folders, ensuring the target icon will pop out if it appears on this side. On the other half, an assortment of different icons is presented, making it necessary to move the eyes around on this side to find the target. The data were piloted to ensure that, in the absence of the ability to move the eyes, detection of the target was at ceiling on the homogeneous side, and close to chance on the heterogeneous side (see supplementary information for a full description of the pilot experiments). The same group of participants participated in both the line segment and icon search tasks to allow comparison of strategies across these two contexts. Focusing search on the heterogeneous side of the search array is always optimal across the two kinds of search stimuli in these experiments; the type of stimuli used as targets and distractors has no bearing on the optimal strategy or how easy it is to implement.

To foreshadow the results of this experiment, behaviour during search for a line segment oriented 45° to the right bore almost no resemblance to the same sample of participants searching for a desktop icon. For line segment search, we replicated the variable and largely sub-optimal strategies observed previously. For desktop icon search, participants were uniformly efficient. We went on to conduct two follow-up experiments to evaluate the potential role of colour and familiarity in explaining the large differences in search efficiency between line segments and desktop icons, and these follow-ups indicated that neither of these factors presented viable explanations. We first present the results of these three experiments as a single set, and then a second set addresses the question of why search strategy might differ dramatically between line segments and desktop icons, despite the fact that that the optimal strategy is similarly beneficial and straightforward to implement across both display types. The key insight from the series of experiments, which we return to in the general discussion, is that search behaviour is flexible and adaptive, and can shift dramatically with small changes to the stimulus array. These shifts provide an opportunity to isolate the aspects of the search context that can trigger this adaptation and thereby allow us to construct more complete models of visual search.

Experiment 1a: Lines and Icons

Methods

Participants

Thirty participants (*females=17, males=13*) completed Experiment 1a (*mean age = 25.27, SD = 5.02*) and a total of 29 participants took part in the two follow-up experiments (Experiment 1b [*n = 15 [females=12, males=17], mean age=22.33, SD=1.84*]; Experiment 1c [*n =14 [females=10, males=4], mean age=23.73, SD=3.62.*) Age and gender information

for this and all the experiments reported here was verbally requested from participants in an open-ended question and recorded by the experimenter next to the participant number. An additional 26 participants completed one of three pilot experiments. No participants contributed data to more than one of these experiments. Participants were recruited through word of mouth and the use of Aberdeen Psychology Research Participation Scheme (PRPS). Participants recruited through PRPS were awarded course credit, and all other participants received £5 reimbursement for their time. In this and all other experiments reported below, participants had normal or corrected-to-normal vision, provided informed consent, and were debriefed following their session. All protocols were reviewed and approved by the Aberdeen Psychology Ethics Committee.

Apparatus

A 19 inch CRT ViewSonic Graphics Series G90fB monitor with a resolution of 1024 x 768 and refresh rate of 100Hz was used to display stimuli. MATLAB running PsychToolbox (Pelli, 1997; Brainard, 1997) and EyeLinkToolbox was used to present stimuli and record data on a Macintosh PowerMac. A desktop-mounted EyeLink 1000 eye tracker sampling at 1000Hz was used to record the position of one eye.

Stimuli

A split screen array was used in which one half consisted of homogeneous distractors and the other half heterogeneous distractors. The homogeneous and heterogeneous halves were presented equally often on the left and right side of the screen and presented in a random order. The target was absent on half of the trials. In the remaining part of the paper, we will refer to the target among homogeneous distractors as a *homogeneous-side target*, target among heterogeneous distractors as *heterogeneous-side target*, and trials with no target present will be referred to as *target absent*.

Line segments were aligned in 22 columns and 16 rows and presented on a mid-grey background (see Figure 1, left panel). The search target was a line tilted 45° to the right. Homogeneous distractor line segments had an 18° range in orientation and heterogeneous distractor line segments had a 106° range in orientation. The mean orientation of the distractors in both heterogeneous and homogeneous conditions was orthogonal to the target (i.e., 45° to the left). The target could appear anywhere in the array other than the first and last row and column and the middle two columns.

Icons for common Apple OSX desktop applications were aligned in 8 columns and 7 rows on a mid-grey background (see Figure 1, right panel). Homogeneous distractors were all the same icon (folders, given generic titles like “Work 1”) and the heterogeneous distractors were taken from a set of 29 different icons. The target was a different icon on each trial (as described in the procedure below) and could appear anywhere other than the middle two columns. We exclude the middle two columns to ensure the target is surrounded by either heterogeneous or homogeneous distractors on all sides.

Piloting the stimuli

A key requirement for our paradigm is that when the target is among the homogeneous distractors, it can easily be detected using peripheral vision, while a fixation is needed to detect it against heterogeneous distractors. The line segment stimuli were already vetted to meet this requirement (see the pilot experiment in Nowakowska et al., 2017). To verify the suitability of the new stimuli, we conducted pilot experiments that measured the detectability of targets appearing among the homogeneous and heterogeneous distractors. In the pilot, each trial began with a target icon presented at screen center, informing participants what icon to look for on the upcoming trial. The target icon remained on the screen until the participant pressed any button on the keyboard.

Participants then fixated the center of the display (marked by a fixation cross) and pressed any key to begin the trial. The arrays were the same as described for the main experiment above, except that they were not split, but uniformly homogeneous or heterogeneous. A display duration of 200ms was used. This is not long enough to make a visually guided eye movement, therefore participants made their judgments based on the information they could pick up from central fixation. Eye movements were monitored to ensure participants started every trial in the centre of the screen. Participants had to report by key press whether the target icon was present or absent. There were 128 trials, divided into 4 blocks. Eight observers took part in the *icons* pilot (verifying the stimuli for Experiment 1a), ten observers in the *grey* pilot (1b), and eight in the *mosaic* pilot (1c). For the stimuli to be usable in our study, accuracy should be near-ceiling for the homogeneous stimuli, demonstrating they are detectable in the periphery without needing to move the eyes. Accuracy should be close to chance for the heterogeneous stimuli, demonstrating that eye movements would be required to find the target on the heterogeneous background. Pilot experiments confirmed that these conditions were met. For all the conditions the accuracy was under 57% for heterogeneous arrays and above 90% for homogenous arrays (detailed accuracy data can be found in supplementary materials). In theory, using exposure durations longer than 200ms may have further improved detection even without eye movements, but in practice we found it was extremely difficult to maintain fixation in the heterogenous condition for durations any longer than this, further reinforcing the need to move the eyes to find the target in this condition.

Procedure

Participants were seated in a dimly lit room using a chin rest set 50 cm in front of the monitor. A nine-point calibration was completed prior to beginning each block. The researcher was present in the room for each calibration. The participant was then left alone in the room to complete each block. Both tasks began with ten passive viewing trials presented for 5 seconds in which there was no task requirement, in order to assess where in the search array participants looked when not undertaking a search task. Following this, six practice trials were completed, followed by four blocks of experimental trials. Before each trial participants were required to fixate the center of a fixation cross then press any key to begin. Trials would not begin unless participants were fixating the cross. Participants reported if the target was present or absent using the *up* (present) and *down* (absent) arrow keys. Participants were told to respond as quickly and accurately as possible. Each array was presented until the participant made a response, and the trial timed out after 60 seconds and the fixation cross for the next trial appeared on the screen. Visual feedback was provided for an incorrect response in the form of a red screen.

Before each icon search trial, one of the icons was identified as the target on that trial. Participants were presented with this target icon for the upcoming trial in the center of the screen and were required to press any key in order to proceed to the fixation cross. Each icon was the target an equal number of times in the experiment. With eight repetitions of each icon, there were 232 icon search trials in total. The line search task consisted of a total of 160 experimental trials. The two stimulus types (line segments and icons) were blocked and counterbalanced across participants to control for order effects. Participants completed both tasks in one session. Each session lasted approximately 90 minutes.


Procedure for Follow-up Experiments

We carried out two follow-up experiments to rule out differences between the line segments and icons that could explain differences in search strategy. The methods and procedure were exactly the same as for the icon stimuli in the main experiment except for the details specified in the following.

Experiment 1b: Greyscale icons. In the line segment search task, we assume participants can use orientation to rapidly guide their attention to the target when it appears on the homogeneous side. In the icon search, the unique colour of each icon can provide a distinguishing feature that sets it apart from the distractors. Colour has long been known to be a guiding feature in visual search (e.g. Horowitz & Wolfe, 2014) that may be more powerful than orientation (Hulleman, 2020). This experiment therefore tested the possible contribution of colour to optimal search of icons. On a new set of participants, we repeated the icon experiment but with greyscale versions of the original icons (Figure 2, left panel).

Figure 2.

Examples of Stimuli in the Follow-up Experiments



Note. Left panel is an example of the greyscale display; Middle panel is an example of the mosaic display. Examples of pixelated dictionary and SPSS icons are presented in the top right and bottom right panels respectively.

Experiment 1c: Mosaic icons. A possible explanation for the improved performance with icons is that participants are highly practiced in searching for these specific objects. Practice with line segments leads to gradual improvements in efficiency (Nowakowska et al., 2019; 2021). When we test already-familiar stimuli like desktop icons, we may be observing participants who have already practiced searching for these particular stimuli many times before, and are therefore highly efficient. To test this hypothesis, we removed visual familiarity but retained the same visual properties as the icons by creating *mosaics* using the Adobe Photoshop mosaic filter (Figure 2, middle and right panel). The filter created patterns of coloured squares that are no longer recognizable desktop icons: we debriefed participants about their subjective impression of the stimuli after the experiment and only one participant recognised one desktop icon among the stimuli presented. If practice with the stimulus set is the primary reason that icon search is efficient, search of mosaics should produce similar idiosyncratic search behaviour as line segments.

As noted above, we piloted these stimuli with a different set of participants to ensure our assumption that participants should only fixate the heterogeneous side was valid. As a further check of this assumption, at the end of *mosaic* experiment we also presented trials that were uniformly homogeneous and uniformly heterogeneous (the method and procedure were exactly the same as the pilot experiment). Target detection accuracy was expected to be above 80% on the homogenous background, and below 65% on the heterogeneous background. This allowed us to remove from the analysis any participants for whom this is not the case. In the *mosaic* experiment one participant did not meet the accuracy criteria for the homogeneous search, and hence is removed from further

analysis. The mean percentage correct for homogeneous search for the remaining 13 participants was 86.95 ($SD=33.70$), and for heterogeneous search 50.61 ($SD=50.03$).

Analysis and Pre-processing

Our key measure of search efficiency is the proportion of fixations made to the heterogeneous side of the display on target absent trials only. Excluding target present trials ensures all the fixations are related to “searching” rather than verifying the target identity. We use fixations 2 - 6 because nearly all target absent trials contain at least six fixations, and because early fixations tend to vary most widely (as the trial progresses, search focuses on the heterogeneous side). Fixation 1 is not used, as the trial starts with a fixation in the centre of the screen. This measure is the same as that used in many previous published implementations of the SHLS task and therefore allows direct comparison (specifically, with Clarke et al., 2022a; Clarke et al., 2022b; Nowakowska et al., 2021; Nowakowska et al., 2017).

Transparency and Openness

Pre-registered methods, hypothesis, and analysis plan for the main experiment can be found here (<https://osf.io/gfxth>) and the two follow-up experiments can be found here (<https://osf.io/sqceu/>, <https://osf.io/x7rsq/>). The pre-registered analysis plan specified t-tests and correlations, separate for each experiment. Where applicable, these are presented in the supplementary materials for transparency. However, as the results show extremely large effects, the statistical analysis has been largely omitted from the main text in favour of full graphical representation of the data. The data from all the experiments presented here are publicly available at <https://osf.io/9edx6/>. We use R (v3.6.1) with the brms (v2.12) package (Bürkner, 2017) for data analysis.

Power Analysis

The main experiment is similar in design to Experiment 1 in Nowakowska et al. (2021), which compared two experimental conditions of the line segment task within subjects, in a blocked and counterbalanced order. As such, the power analysis from that study can be reused. In brief, we conducted 50 simulations of a sample of 15 participants with an underlying simulated difference in efficiency of 10% (specifically, a shift from around 50% to 60% efficiency within each participant). For each simulation we calculated the probability that the manipulation improved efficiency, given the data, and found this in around 95% of the simulations.

Results

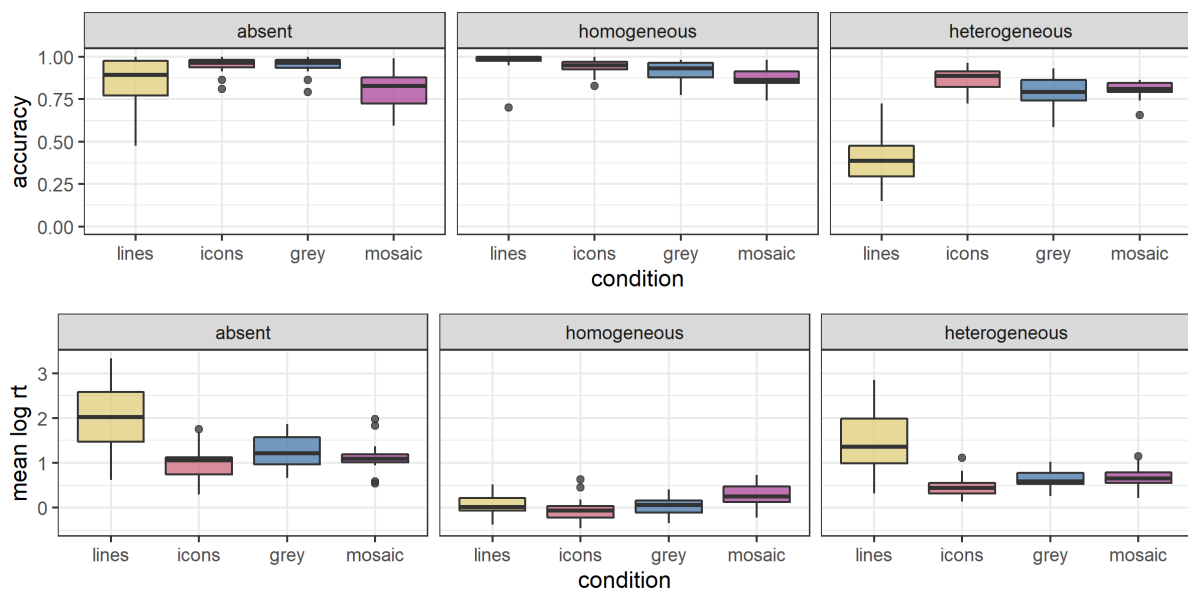
Accuracy and Reaction Time.

The distribution of accuracy and mean log reaction time are shown in *Figure 3*. Although results from all participants are presented in this figure, we exclude from further analysis participants with accuracy of 10% or less for detecting the target when it was present on the heterogeneous side. Two participants met this criterion. These participants do not try to find the target when it is present on the heterogeneous side, instead, use their peripheral vision to check for target presence on the homogenous side and if target is not detected they hit the target absent response key. They are therefore very fast and achieve an overall accuracy around 75%. This behaviour represents a dimension of individual differences in strategy (as we will return to in the discussion), but it is of limited usability in this analysis as we focus on analysing eye movements in search-related target absent trials; if participants are not searching at all on these trials (because their strategy is to guess the target is absent if they do not immediately detect it), their eye movements are not meaningful indicators of their strategy. We maintain this low accuracy threshold for

exclusion on the hard trials of 10%, however, because we are interested in individual differences in search efficiency and want to keep as representative and complete a sample as possible.

Figure 3.

Proportion Correct (top) and Log Reaction Time (s) (bottom) for Four Different Search Stimuli.




While the accuracies and reaction times for the line segment stimuli are in line with the results that we have previously reported (Nowakowska et al., 2017), there are marked differences between the *lines* and *icon/grey/mosaic* stimulus conditions in Figure 3. In particular, observers are faster and more accurate with the three variations of the *icon* stimuli for heterogeneous targets. They are also faster to correctly respond that the target is absent when there is no target present in the stimuli. An important question for our hypothesis is whether the faster reaction times are driven by differences in search efficiency.

Search efficiency.

The main variable of interest in our study is the search efficiency metric, defined as the proportion of fixations 2-6 directed to the heterogeneous side of the display during target absent trials. We start by looking at the results from Experiment 1a, in which we can clearly see that there are large differences in how observers approach searching these two classes of stimuli (see Figure 4). For the line segment stimuli (in yellow), we replicate the large individual differences documented by Nowakowska et al. (2017), Clarke et al. (2022), and others, finding a full range of behaviours from optimal (looking predominantly at the heterogeneous side), through chance (equal fixations on both sides), and even counter-optimal (looking at the homogeneous side more than the heterogeneous side). This wide range is seen across the first five fixations in the panels on the left side of Figure 4, and in the summary boxplot on the right side. For icon stimuli (red), however, there is very little variation between participants: every individual in our sample executed an almost-perfectly optimal search strategy. Highly efficient search was also observed for the greyscale icons (blue) and the mosaic icons (purple) in the follow-up experiments.

Figure 4.

Individual Strategy (Left Panel) for Lines and Icons and Average Strategy for all Conditions (Right Panel).



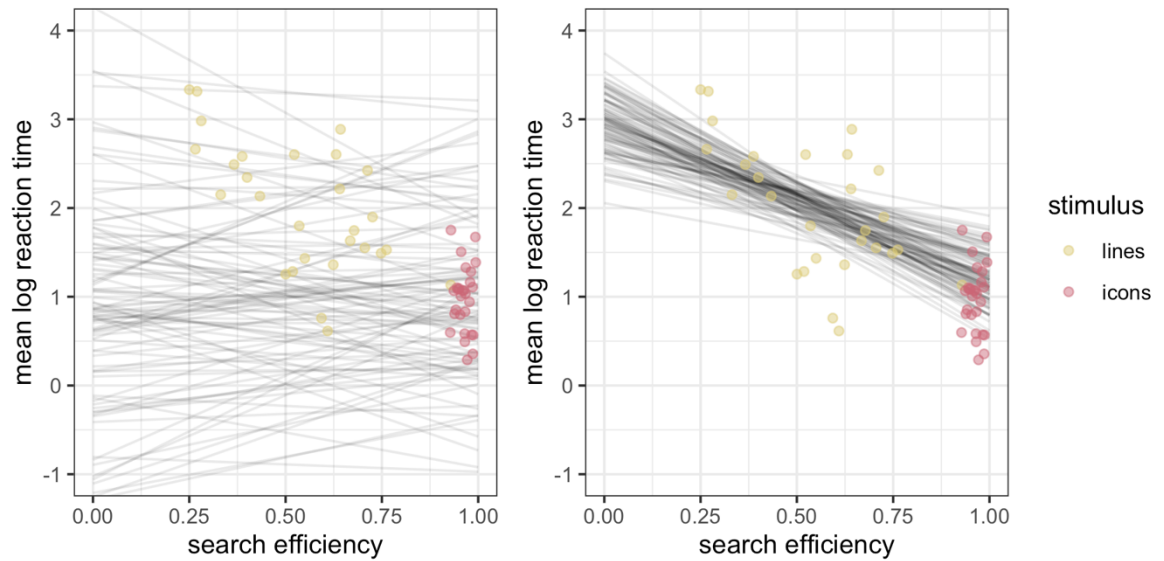
Note. The left panel shows facet plot for each participant's proportion of fixations on the heterogeneous side, for fixations 2-6 on target absent trials. We find a full range of individual differences for the line search (yellow lines), from optimal, through chance, to counter-optimal. All participants were near-optimal when searching for an icon (red lines). The right panel shows the boxplots for the proportion of fixations to the heterogeneous side for the four stimuli types. Points are the proportion of fixations 2-6 on all target absent trials for each participant.

Relationship between Search Strategy and Reaction Time.

Reaction times on icon search trials are faster than for lines, and participants also search the icon stimuli optimally, according to our efficiency measure. One key question is whether we can explain the faster reaction times for icons solely based on differences in strategy, or if, even accounting for poor strategies, search for lines is slower than for icons. This is important because a difficult search task might both contribute to poor strategies, and slow down search. To investigate this further, using the data from the line segment stimuli only, we fit a Bayesian linear model to predict an observer's reaction time based on their search efficiency (the details of the model can be found in supplementary materials, and the priors are shown in Figure 5, left panel). We can then use this model to predict reaction times for a given level of efficiency, and see where the icon data fall relative to this prediction. This is shown in Figure 5, right panel: the lines show the predicted relationship between efficiency and RT conditioned on the line segment data, and the red points show the empirical icon data. As can be seen from this figure, the faster reaction times for icon stimuli can indeed be predicted based on a model of the efficiency/RT relationship based solely on search among lines. While the model shows that the data is in line with the causal explanation we provide, there are other possible explanations that are also valid (for example, a third variable could drive both higher efficiency and faster reaction time).

Figure 5.

Predicting Reaction Time Based on Search Efficiency.



Note. Each dot represents an observer’s mean performance over the target absent trials. Each line represents a sample from a Bayesian linear model. The plot on the left illustrates our choice of the weakly informative priors for the intercept, $N(1,1)$, and slope, $N(0,1)$. The model is then trained on the target absent data from the line condition, and the posterior predictions are shown on the right. We can see that observers with more efficient search strategies typically have faster reaction times. Furthermore, the high search efficiencies seen for the icon stimuli result in reaction times that are in line with the predictions made from efficient search in the line condition.


Comparison of search to passive viewing.

As described in the methods, each participant started the experiment with ten passive viewing trials, before they had been given any instructions or targets to look for. These trials were included to measure whether there were baseline differences between the different stimulus arrays in how eye movements were distributed that might account for any differences in how participants distribute their eye gaze during active search for a specified target. The results of the passive viewing, showing the proportion of the first five fixations directed to the heterogeneous side on each trial, are in Figure 6, shown together with the same measure from each search condition (for ease of comparison, this duplicates the data shown in the Figure 4 boxplot). The results show a bias towards the heterogeneous side during passive viewing across all four stimulus types. However,

during search for line segments, this bias disappears, and eye movements are distributed roughly equally to both sides. For the icons, greyscale icons, and mosaics, the bias increases to approach 1. The passive viewing condition suggests small differences in passive viewing behaviour do exist between lines and icons, but the difference is much larger in the search task.

Figure 6.

Proportion of Fixations to Heterogeneous Side During Passive Viewing and Search.



Note. Each facet is a different stimulus type.

Discussion of Experiment 1

We replicated large individual differences in search efficiency for simple line segments, but when we replaced abstract line segment stimuli with icons, we observed a shift from highly variable strategies to uniformly efficient search behaviour across all participants. This stark change in efficiency arose despite the differences between stimulus sets not being relevant to our key manipulation: guiding search to the heterogeneous side of the search array is always optimal, and there are no obvious reasons why this strategy would not be equally simple to implement across both search contexts. The two follow-up experiments ruled out colour and explicit familiarity with the search objects as candidate explanations for the more optimal strategies observed with icons.

The simplest description of the overall pattern across the conditions is that the majority of participants switch between different search *modes* depending on the context of search. The mode observed when searching for the icons could be described as “ideal”, because participants are directing their eye movements to where they are most needed. The other mode, during search for line segments, is more challenging to describe and explain, because what makes it distinctive is its lack of similarity across participants. Out of the 30 participants who completed both the desktop icon and line segment versions of the search task, just one individual applied the ideal strategy, while the rest could be applying a collection of different strategies or policies, including near-ideal, random and heuristic-based eye movements that aim to cover the search area in a flexible and exploratory way. We could therefore tentatively label the line segment search an “exploratory mode”, a distinction alluding to the classic *exploitation* and *exploration* modes of behaviour in the management and economics literature (e.g., March, 1991), where exploitative choices are made with the intention of maximizing utility, and exploratory choices serve the function of gathering information about possible choices. A similar distinction is made in the animal learning literature (e.g., Atkinson, 1986), with *actions* defined as behaviours that are selected based on simulating expected outcomes and choosing the best of these. The alternative behavioural selection strategy is to rely on *habits*, executed on the basis of repeating previously rewarded behaviours. The distinction in both these cases is between a mode of behaviour that is purposeful, efficient, and computationally expensive versus one that is exploratory, noisy in terms of outcome, and computationally cheap (see the review presented in Clarke et al., 2019 for a more detailed description of this distinction). Similar to these distinctions, in visual search one mode is not necessarily better than the other in general terms, but there are circumstances in which one or the other is better suited. From the current set of results,

it seems that the icon search context induces a more purposeful mode of searching, while the mode of searching induced by the line segment context is more variable from one person to the next (but as demonstrated previously by Clarke et al., 2022, how a participant approaches the line segment task is relatively stable over time).

The passive viewing baseline condition reinforces this interpretation; participants on average have a slight preference for the heterogeneous side when they are not explicitly instructed to search, but when they start the search task for line segments this preference goes down. In the line segment context, participants' goal may be to distribute their fixations evenly across the search area, irrespective of how much new information those fixations will provide. For icons, participants have a similar moderate preference for the heterogeneous side when they are passively viewing, but when they are searching this preference increases dramatically, reflecting the optimal strategy for finding a target. Virtually no eye movements are allocated to the homogeneous side during search for icons.

Although it seems clear from the first set of experiments that the majority of participants switch between an efficient search strategy with icons and an inefficient search strategy with lines, what is not yet known is what feature of the stimulus it is exactly that induces this switch. We have ruled out colour and familiarity, but a number of other properties distinguish lines segments and icons. Two properties in particular stand out. First, when viewing the two arrays of stimuli in Figure 1, the array of lines looks more uniform and texture-like, while the set of icons (even when distorted into mosaic patterns) looks more like a scene full of objects. A scene-like setting may induce a search mode that relies more on a *nonselective pathway* for guidance. This pathway, labelled by Wolfe et al. (2011), allows for scene gist to guide search to likely target locations, and complements a

separate *selective pathway*, in which individual candidate items are inspected to determine if they are the target. The texture may encourage this selective pathway to dominate search because each item matches the target in all ways except orientation, and this more selective route will vary more because there is no clear hierarchy to follow. A second property that differs between the two arrays is the way the target is defined: participants are searching for an object of specific identity (*find this icon*) or for a feature of an object (*find a line with this tilt*). How the target template is defined has been shown to affect search performance, with targets defined by showing the visual item on the screen speeding responses and reducing the effect of set size (Vickery et al., 2005). Similarly, a classic paper by Bacon and Egeth (1994) demonstrated that participants can engage in what they refer to as a “singleton detection mode”, where attention is guided by uniqueness as opposed to a particular feature like colour or tilt, and is thereby also directed to unique distractors. This was presented in contrast to “feature detection mode” where the guiding feature is a specific characteristic (like a color or shape). In the same vein, the set of eye movement strategies available when searching for a pre-defined orientation may be a different set than those available to a participant who is holding in mind a particular visual template of an object.

Experiment 2: Lines and Pens

In this second set of experiments, we aimed to replicate and extend the switch in search mode observed in the first set, and to test whether it is the nature of the array or the target itself that drives the switch. In the first of these experiments (2a), we compared search among individual objects with the search among lines, while aiming to match the stimuli as closely as possible to eliminate differences in sizes and shapes between stimulus sets and present a set of sparser individual objects. Crucially, the task here was

the same for the two classes of stimuli: Participants were looking either for a pen tilted 45 degrees to the right among a set of other pens, or for a coloured line tilted 45 degrees to the right among a set of other coloured lines. Similarly to Experiment 1, we created search arrays split into halves. On one side the target was easy to spot (homogenous side) while on the other half the search required central vision in order for the target to be detected (or its absence reported). As before, we piloted the stimuli to ensure this assumption was valid. If icon search is optimal because it presents a set of clearly individuated objects, then we should see optimal search for the pen search task, and possibly the line search task as well. An intermediate option is that the pens will be closer to optimal than the lines, because it is more intuitive to see the set of individual photographic images of pens as unique objects than a set of coloured lines. Alternatively, it is not the set of stimuli that matters, but the way the target is defined that makes the icon search more efficient. If that is the case, we should see search performance that is variable and largely suboptimal for both pens and lines, because the target in both cases is defined based on orientation, rather than identity. To foreshadow the results, results were consistent with the latter hypothesis, with participants being highly variable and largely suboptimal in their search, both for 45°-tilted pens and for 45° tilted lines (matched in size and spacing and both displayed in colour).

In a follow-up experiment (2b), we therefore ask participants to again search through an array of pens, but now the target of search is a particular pen, and the orientation is irrelevant. This final experiment was designed to confirm that it is indeed the definition of the target that seems to drive the mode of search: if so, participants should return to being uniformly efficient when searching for a particular pen, as opposed to a pen of a particular orientation. Consistent with this prediction, participants are uniformly optimal in the final experiment.


Methods

Participants.

Thirty participants (20 females, 10 males) took part in Experiment 2a (mean age = 25.13, $SD = 5.08$). Two participants took over an hour to complete one condition (where most participants took an hour to complete both conditions) hence had no time to complete the other condition. The data sets of these participants have not been included in the analysis. An additional sample of 15 participants (9 females, 6 males) completed Experiment 2b (mean age = 23.37, $SD = 5.40$). In addition, there were 30 participants in total for the pilot experiments verifying that the three stimulus arrays met our assumptions (10 for each, subsequently labelled as *pens*, *lineL* (for larger lines), *penID*; the procedures for these pilots was the same as in the pilot experiments reported in Experiment 1). No participants contributed data to more than one of these experiments. Participants were recruited through word of mouth. Participants received £10 reimbursement for their time in the orientation experiment and £5 in the identity experiment.

Figure 7.

Example Line Stimuli (Left Panel) and Pens Stimuli (Right Panel).



Stimuli and Apparatus

The same apparatus as in the previous set of experiments was used. The logic of the displays was the same as well, with search arrays split into halves. On one side the target would be easy to spot (homogenous side) while on the other half the search would require central vision in order to for the target to be detected (or its absence reported).

In this experiment we aimed to match the stimuli as closely as possible. To introduce variation in colour to the line stimuli, lines were presented in three different (irrelevant) colours: yellow, violet and blue. Pen images were edited from a set of 81 different images of pens and pencils obtained through google image search. We also matched the displays in the number of elements that were presented between lines or pens. Compared to the line stimulus used in the first set of experiments, we reduced the density of stimuli to 10 columns and 8 rows and presented on a mid-grey background (see Figure 7). The search target was a line (left panel), or a pen (right panel) tilted 45° to the right. Similar to the original line stimulus, homogeneous distractors had an 18° range in orientation and heterogeneous distractors had a 106° range in orientation. The mean orientation of all distractors in both conditions was orthogonal to the target (i.e. 45° to the left). The target could be presented anywhere in the array other than the middle two columns.

Experiment 2a Procedure (pens and lines)

The procedure was the same as for the first set of experiments apart from for the following changes. The two stimulus types (coloured lines and pens) were blocked and counterbalanced across participants to control for order effects. Each pen was the target an equal number of times in the experiment. With two repetitions of each pen, there were 160 pen search trials in total. The line search task consisted of a total of 192 experimental trials. The target line was random colour from trial to trial and the colour was irrelevant


to the task. The identity of the pen was also irrelevant to the task, as both pen and line targets were defined based on their orientation. The task was exactly the same between the two conditions, i.e., participants were looking either for a pen tilted 45 degrees to the right or for a line tilted 45 degrees to the right.

Experiment 2b procedure (identity-based pen search)

We randomly selected 40 images of pens from the *pens* stimulus set used above to create heterogeneous side, and we created another 40 uniform pens stimuli to be used on homogenous side. The homogenous pen images varied in shade from teal to green, and otherwise were identical.

Figure 8.

Example Stimuli from the Identity-search Pen Task.



The pens were aligned in 10 columns and 8 rows and presented on a mid-grey background (see Figure 8). The target pen was always the same identity (Figure 8, 1st row 7th column). All distractor pens and the target pen had a 106° range in orientation with mean orientation being 45° to the left. The target could appear anywhere in the array other than the middle two columns. The orientation of the target and distractors was irrelevant to the task.

The experimental procedure was the same as in previous experiment except the following changes. Before each search trial, participants were presented with the target pen (it was always the same dark blue pen) in the centre of the screen and were required to press any key in order to proceed to the fixation cross. The task consisted of a total of 160 experimental trials and lasted approximately 30 minutes.

Piloting the stimuli

To validate our stimuli, we conducted pilot experiments that measured the detectability of targets appearing among the homogeneous and heterogeneous distractors, using the identical procedure as described for the first set of experiments. The pilot study for the *pens* and *penID* stimuli consisted of 128 trials and for the *lineL* condition consisted of 96 trials that were presented across 4 blocks. Ten participants completed each pilot experiment. The results showed accuracy above 90% for both conditions for homogeneous trials and below 55% for heterogeneous trials (full pilot accuracy data can be found in supplementary materials).

Post experiment accuracy checks.

Similarly to the *mosaic* experiment (Experiment 1c), following the experimental trials we presented trials that were uniformly homogeneous and uniformly heterogeneous (the method and procedure were exactly the same as for the pilot experiment). In these trials, target detection accuracy was expected to be above 80% on the homogenous background, and below 65% on the heterogeneous background. This allowed us to remove from the analysis any participants for whom this is not the case. We had pre-registered this criterion for inclusion of participants in the analysis but found that for 18 of the 30 participants in the main experiment this was not the case, either because of poor performance for the homogenous trials or performance on the heterogeneous side that

exceeded 65%. To be consistent with the prespecified criteria we performed the analysis on the remaining sample of 12 participants. However, because these exclusions were high, we report the full set of results to demonstrate that the results are the same whether they were included or not (and show this in full in the supplementary materials). For the line search condition, the mean percentage correct for homogeneous search for the remaining participants was 91.74 ($SD=27.56$), and for heterogeneous search 52.90 ($SD=49.97$). For the orientation pen search the mean percentage correct for homogeneous search was 92.03 ($SD=27.10$), and for heterogeneous search 59.38 ($SD=49.15$). For identity-based pen search two participants did not meet the accuracy criteria, and hence were removed from further analysis. The mean percentage correct for the remaining 13 participants was 95.80 ($SD=20.00$) for homogenous trials, and 52.50 ($SD=50.00$) for heterogeneous trials. Preregistered details of the study design and analysis can be found here (<https://osf.io/de8j2/>, <https://osf.io/4yzhp/>).

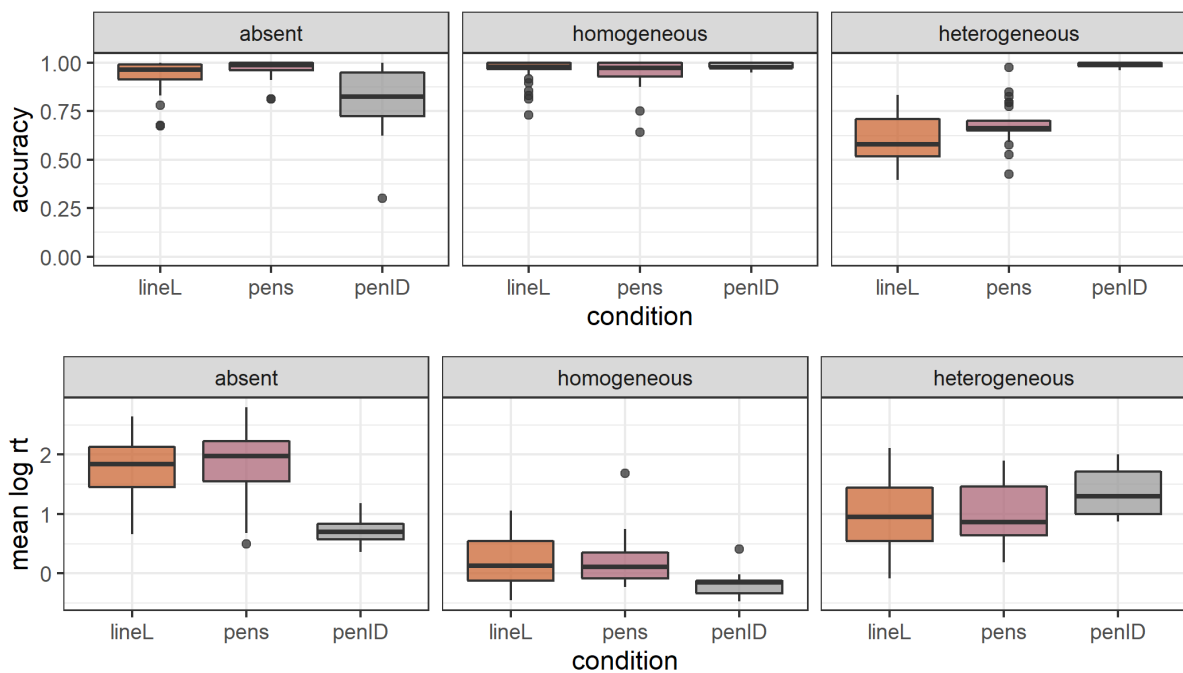
Results

Reaction Time and Accuracy.

As can be seen in Figure 9, performance is comparable between the line and pen conditions in which participants search for the target based on orientation. When searching for a pen of a specific identity, participants are more accurate when it is present on heterogeneous side (compared to search for pen/line orientation) but they are also less accurate and faster for target absent trials. This suggests only that participants were more likely to confuse targets and distractors when searching for a particular pen than when searching for a particular orientation.

Figure 9.

Proportion Correct (Top) and Log Reaction Time(s) (Bottom) for Three Search Conditions.

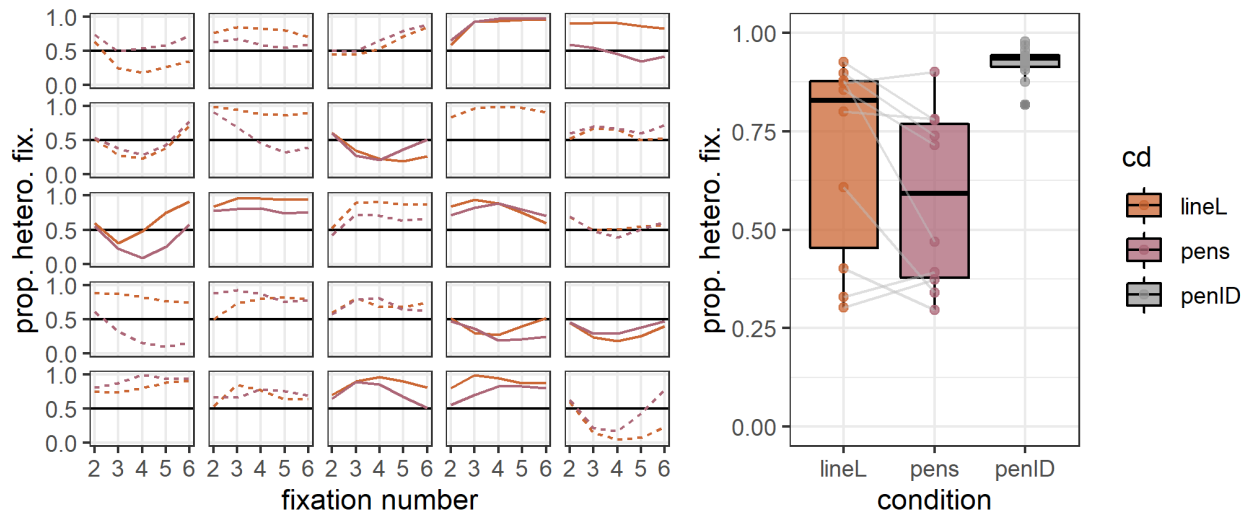


Search efficiency

Like in the line segment search in Experiment 1a, when searching based on orientation in Experiment 2a, participants showed a full range of search efficiency from optimal, through chance, to counter-optimal, and there was no systematic difference between line and pen stimuli (see Figure 10). Indeed, it is clear from the facets in Figure 10 showing individual participant data that efficiency is closely correlated across conditions. From these results we can conclude that searching through an array of individuated objects (*pens*) is not associated with a more efficient strategy than searching through line segments; when participants are searching for object feature (orientation) they are stochastic on average, variable on an individual level, and an individual is relatively consistent across similar stimulus sets. In contrast, searching for a pen based on its identity (the *penID* condition in Figure 10) invokes a uniformly optimal strategy, with all participants directing their eye movements almost exclusively to the heterogeneous side.

Figure 10.

Individual Search Efficiency (Left Panel) and Mean Efficiency (Right Panel).



Note. Left panel shows a facet plot for each participant's proportion of fixations on the heterogeneous side, for fixations 2-6 on target absent trials and the right panel shows the boxplots for the proportion of fixations to the heterogeneous side for the three stimuli types. In the left panel participants who either performed above chance on heterogeneous trials or below 80% on homogenous trials on post experiment detection accuracy are marked with dashed lines, all other participants are denoted by solid lines. We find a full range of individual differences for both the line (orange line) and pens search (purple line). Points are participants and only the participants who met the inclusion criteria are shown, although the plot looks similar with all 30 participants included (see supplementary materials).

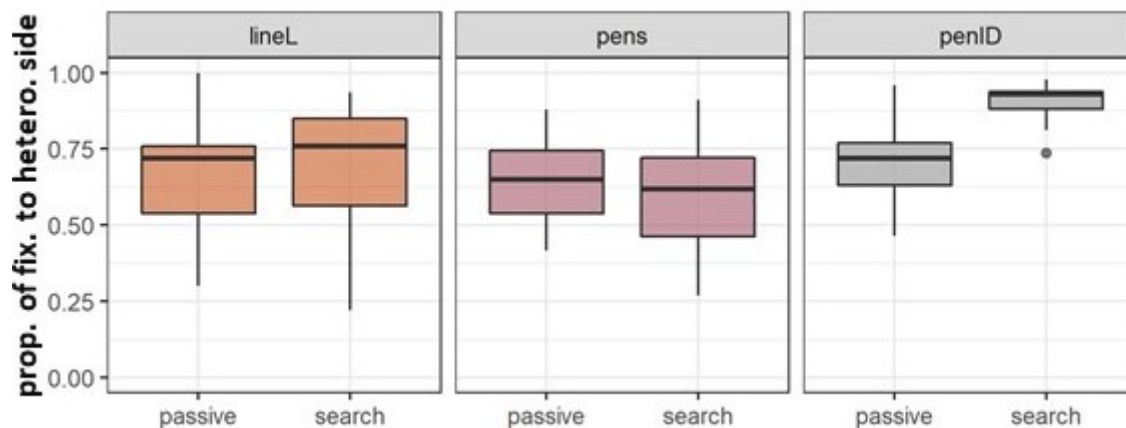
Comparison to a passive viewing baseline.

We included a series of ten passive viewing trials to measure any default bias to one side or the other. The results across the three conditions are presented in Figure 11.

Similar to Experiment 1, the stark difference between the orientation and identity-based search conditions emerges most clearly when participants are engaged in active search for the defined target.

Figure 11.

Proportion of Fixations to Heterogeneous Side During Passive Viewing and Search.



Note. Each facet is a different stimulus type.

Control experiment: Discriminability of the two sides of the display

Visually inspecting the different versions of the display used in the experiments above, a concern arises that the discriminability of the two sides of the search array (the heterogeneous relative to the homogeneous) might be lower for those conditions that are further from optimal (the small black lines, the larger colored lines, and the orientation-based pen search). To check that our participants could readily discriminate which side of the display was more variable (as a basis for deciding which side to fixate), we designed a simple follow-up experiment in which we presented split-half arrays for 200 ms and simply asked participants “which side is [more/less] variable: Left or right?” (asking the “more” version of the question for half of the blocks of trials and the “less” version for the other). We showed four different display types (icons, lines, colored lines, and pens) in eight separate counterbalanced blocks of 20 trials each (160 trials in total). They received no feedback, but they had 10 trials of practice to start each block with unlimited viewing.

A group of 10 participants (*8 females, 2 males, mean age 24.4, SD=3.6*) completed this follow-up. The accuracy in this experiment was at ceiling for all stimulus types: icons $M=96.5, SD=18.4$; lines $M=98.2, SD=13.1$; colored lines $M=97.8, SD=14.8$; pens $M=98.0, SD=14.0$, demonstrating that with 200ms exposure, participants could judge one side from the other across all versions of the search array.

General Discussion

In the experiments presented here, two dominant patterns of results emerged: when searching for a particular icon or pen, participants were uniformly efficient, directing their eye movements towards the locations where central vision was most needed to find the target. In contrast, when searching for a line or a pen oriented in a particular direction, participants were highly variable and mostly inefficient. The striking contrast between these conditions could not be explained through the colour, familiarity, spacing or size of objects, and differences between eye movements across conditions were much less pronounced during passive viewing than for active search. The two sets of experiments show independent and complementary results that demonstrate the “modes” identified in the first set of experiments comparing line segments and icons generalised to a new set comparing larger coloured lines to a closely matched array of pens. The final experiment further narrowed down a determining factor of the distinctive search behaviours to be the nature of the target of search, rather than differences in the search array.

In the series of experiments we have presented, one factor has emerged (i.e., how the target is defined) that appears to dramatically influence the search strategies that are implemented, and we have ruled out several others (colour, familiarity, textures vs

objects). For our initial aim, to identify the policies that drive selection of information in search, our results show clearly that the policies depend on characteristics of the target that seem, on their surface at least, unrelated to what policy is effective or how easy it is to implement. The conclusion we can draw is that a single policy is not sufficient to describe search behaviour, and point to an urgent need for further research to refine and broaden the set of search conditions that seem to trigger particular strategies.

Why might the way a target is defined trigger different search modes? An important point to keep in mind when considering this question is that in abstract terms, measures of search performance like reaction time and set-size effects cannot be straightforwardly related in an *a priori* way to efficient eye movement strategies. That is, a particular manipulation that is known to impede search performance might also make eye movements less efficient; equally, it could be associated with *better* eye movement strategies, because a good strategy is more beneficial under tougher conditions. Keeping this in mind, the key difference between conditions associated with the two modes is that in the variable and inefficient conditions, participants are required to distinguish one element from the others of the same kind; In the “ideal” mode conditions, participants are looking for an object of specific, memorised identity. Distinctive neural mechanisms have been argued to support visual object individuation and identification (Xu & Chun, 2009). *Looking through* the objects for one in a particular state (tilt) might therefore involve different mechanisms to *looking for* the specific target identity. Another point of distinction is that in our experiments, the four experiments with “ideal” search strategies all provided participants with a visual image of the search target at the start of each trial. Vickery et al. (2005) demonstrated that a visual template that precisely matches the target facilitates search relative to a visual image of the target in a different state, or a verbal cue defining the target. Taking these two together, it could be the nature of the

target (as one of a set of objects different only in its orientation, vs a particular object) or the way the template was presented (visually or verbally), or perhaps both, that influenced the way participants embarked on their search. Importantly though, neither of these factors has any bearing on how visible the target would be in the periphery, nor do they make the efficient strategy any more or less effective or easy to implement. An interesting but speculative possibility is that the mechanism that switches participants between efficient and inefficient modes is the expected (but not actual) difficulty of detecting the target. Having only a verbal template and knowing they are searching for a particular feature among otherwise similar objects may cause most participants to default to a more “thorough” mode of searching that entails covering the whole search area. When the target is expected to be easy to find, as when searching for a unique object they have just been shown an image of, participants rule out large areas of the search array more readily and do not waste time scrutinizing them.

An additional factor that could be determinative in how fixations are selected is the particular attentional requirements associated with different search templates. As Feature Integration Theory proposes (Treisman & Gelade, 1980), identifying objects requires binding, and searching for them typically involves serial inspection of distractors. In contrast, detecting a feature (like orientation) can be done in parallel across the search array. While it seems counter-intuitive that search strategies would be less efficient for features than for objects, this could be resolved under Load Theory (Lavie et al., 2004): when searching for a simple feature like orientation, participants might be more susceptible to distraction, leading to inefficient behaviour because they have spare attentional capacity. With the identity task, participants' attention is engaged with binding, and thus they are less susceptible to distraction and better able to focus the eyes where central vision is needed. More generally, understanding the interplay of the

attentional resources involved in object individuation and identification and those involved in strategic control of visual search behaviour will be fruitful ground for future work to explore.

Being able to easily see the distinction between the two halves of the array is a prerequisite for being able to direct the eyes to the side that provides more information, and the control experiment demonstrates that this requirement is met for all the different stimulus arrays, with accuracy to discriminate the two sides near ceiling with a 200ms exposure duration. Nonetheless, one speculative reason for individual differences in searching for line segments could be that some participants are more or less likely to notice the halves or see them as important. The control experiment only shows us that participants *can* discriminate them when asked, and not whether they spontaneously do so during search. This relates to a more general question of why some participants have reliably sub-optimal search strategies, continuing to fixate the side of the search array that provides no new information about the target, while other participants are capable of reliably producing optimal search by the same metric and under the same conditions. While we have not addressed this question here, the present results show that participants who produce sub-optimal search for line segments are in fact capable of producing optimal search given a different set of stimuli. This presents an even bigger puzzle as to why they fail to do so for lines. This puzzle is not addressed by the current set of experiments, but we recently completed a pre-registered study (DOI 10.17605/OSF.IO/5AQ4C, Clarke et al., 2023) in which over 300 participants completed the line segment version of the search task, among a battery of other tasks and questionnaires. A particularly intriguing aspect of these results is the free-text responses participants gave when asked to report the strategies they used, if any, in searching for the line segment. Out of the 267 participants who answered the question, 146

spontaneously mentioned the two halves of the search area (using words like “side” or “half” in describing their strategy). For most of these participants, they describe using these two halves in precisely the opposite of an optimal manner (e.g. *“First I would look at the side where all the lines were going the same way. If I didn't see it there, I would go row by row on the other side. It just seemed to be the most efficient.”*). These responses provide a preliminary indication that many of the participants who are inefficient with the line segments have all the information they need, but deliberately use the information in way that undermines their search. We plan to interrogate the data to better understand how well their reported strategies align with their fixation behaviour, among other interesting questions. This rich dataset will be an important resource for understanding what contributes to more or less efficient search strategies in search.

A range of other tasks have also revealed puzzlingly large variations between participants in terms of visual search behaviours (e.g., Araujo et al., 2001; Irons & Leber, 2016; Kristjansson et al., 2014). In these studies, as in most visual search experiments, the search stimuli are visually simple lines and colored shapes, designed to present an environment free of meaningful objects that might evoke different reactions in different people. It is therefore surprising that it is precisely these more “meaning-free” conditions that seem inclined to produce the largest individual differences in the current set of experiments. As Clarke et al (2022) showed, these individual differences are reliable when tested in two sessions with a week between them, so they seem to reflect some fairly stable trait. This can also be seen in the high correlation in search efficiency between the matched pen and line experiments, demonstrating that a person’s strategy does generalize across similar search arrays.

The difficulty of visually identifying the target in different conditions was not precisely the same across conditions in our experiments, but what was carefully controlled was the general property that the target would pop out in the periphery on the homogeneous side, not needing central vision to be detected. This ensured that the optimal strategy of only fixating the heterogeneous side was the same across conditions. In the first experiment, where there was a large difference in reaction time for the same group of participants when they search for lines compared to icons, we showed that this RT difference was closely correlated with poor strategies in the line segment task. We think the most likely interpretation is that participants were slower to find the line segments because they tended to waste eye movements on the homogeneous side when searching for lines (as we also argued in Nowakowska et al., 2017). We do not have data to determine whether it was objectively more or less easy to find different kinds of targets on the heterogeneous background alone (the pilot experiments only established participants were close to chance when they could not make any eye movements). However, the pilot experiments and the accuracy session at the end of the experiments clearly established that for all participants, the target was visible in the periphery on the homogeneous background, making further fixations on this side consistently low in informational value across all the different stimulus types.

An important way in which search for objects differs from search for orientation is familiarity; the target of search in daily life is typically some particular object, not an object in a particular state. Learning through reinforcement has been implicated as the mechanism driving the complex process of gathering information during natural tasks (Hayhoe & Ballard, 2015; Gottlieb et al., 2014, Paeye et al., 2016). In visual search, evidence accumulation is an active process whereby observers make the decision which location (e.g., in an array) to sample next by weighing the relative informational value of

the options. In the initial stages of learning, eye movements may be stochastic, but through a process of trial and error, they gradually become more structured as the system learns the value associated with each response in a given context. Because participants are unfamiliar with these specific arrays of line segment stimuli, strategies are initially highly variable. As participants get more practiced with the task, the strategies become more constrained and the behaviour more optimal (Nowakowska et al. 2019; 2021; see also Lonnqvist et al. 2020). In the case of icons, where the context and stimuli are more familiar (at least to our sample of undergraduate students) the strategies are already constrained and guided by information-value. The explanation of the difference in strategies between lines and icons in terms of reinforcement learning on the surface contradicts the results of *mosaic* experiment: participants are as efficient with the mosaic stimuli as they are with the original icon stimuli. According to their reported impression of the mosaics, collected at the end of the experiment, only one participant recognised one of the desktop icons from the set. Eight participants mentioned some of the pixelated icons resembled familiar objects (e.g., a tree, a flower, an object from a Mario Brothers game). Thus, the mosaics were not explicitly recognized, like the original icons, but may have been implicitly familiar to some of the participants. More importantly, the act of searching based on identity is, at an abstract level, more familiar than searching based on orientation, which may trigger a more “ideal” mode by tapping into a wealth of prior experience, even though the objects are novel.

To expand on this point, we have identified two distinct modes of visual search in this series of experiments, but it is likely that the “exploratory” mode we see in search for line segments is in fact a collection of different modes driven by individual differences, which more targeted research will be able to identify. For example, as noted in the results of Experiment 1, two participants had to be excluded for engaging in a low-effort version of

search, where they responded that the target was present only if it was immediately visible to them, and responded that it was absent otherwise. As a strategy, this yields an overall accuracy of around 75% correct and reaction times that are considerably faster than participants who choose to explore the scene before responding. In an individual differences study by Irons and Leber (2016) they use the label “effort minimizers” to describe participants with similar tendencies in their visual search task, and a similar label could be applied here. These clearly relate to error tolerance, bridging our strategy measures to the substantial literature on strategy and search termination (e.g., Chun and Wolfe, 1996). On the flipside, we also have anecdotally observed several participants with extremely slow reaction times tend to have scanpaths revealing an effortful “reading” mode of examining each item in a line-by-line fashion. It is an interesting, and open, question whether we can create conditions that might encourage all the participants in a sample to engage in either “low-effort” or “reading” modes of search, the way the icon task produces “ideal” search. Other modes could relate to the nonselective and selective search pathways (Wolfe et al., 2005), to singleton and feature search modes (Bacon & Egeth, 1994), or to serial dependencies (such as searching where the target was last seen, Walthew & Gilchrist, 2006). While we see the same group of participants switch their search behaviour dramatically from lines to icons, it is also an open question whether conditions could produce a graded shift from one to the other, if they can change over time and with experience, and whether participants can flexibly switch between modes on a trial-by-trial bases. Finally, understanding what other conditions produce uniform search modes and which encourage each participant to default to their own preferred approach is an interesting open question that could contribute usefully to our ability to predict search performance in untested conditions.

Conclusions

The experiments reported here suggest that at least two distinct patterns of results can be elicited from the same group of participants, and that these two patterns hinge on what has been designated as the search target. We use the term *mode* to describe each of these two search patterns. The results challenge the assumption that one can easily generalize visual search with simple, artificial stimuli to more complex ones, and that a single model based on one context can account for behaviour in another. This is not a limitation, however, but an opportunity to understand what distinguishes these search modes and what other properties or conditions might trigger the switch between them. This approach holds great promise for modelling search and related behaviours not only within a context, but to more precisely account for how behaviour can adapt across a range of different contexts.

Constraints on Generality

A key implication of our results is the constraints that limit generalizing from one visual setting to another, but of course constraints could also limit generalizing from one group of people to another. We have recruited participants from the University of Aberdeen community for the experiments reported in this series. This sample is, by definition, well-educated adults living in a European democracy. Although we did not directly record demographic characteristics of our sample aside from the gender and age information we reported in the methods of each experiment, about 46% of the population we sampled self-identify as white (based on 2021/2022 student survey data). The key feature of the results from which we draw our main conclusions is the diversity of strategy in some conditions, and uniformity of strategy in others. It is possible these are characteristics of the population we sampled from that may not be universally observed in all groups.

Public Significance Statement

The current research tests the predictions of theories of eye movement control that are based on a principle of expected information, that is, how much new information one can expect to gain from making each eye movement. We tested people's strategies across a range of different search displays, from simple line segments to computer desktop icons, and designed the displays to ensure a strategy based on expected information should be equivalently easy to implement across all these situations. Our results show two distinct patterns of behaviour: when participants are searching for a particular object (e.g., one specific pen among other pens) they clearly and uniformly match the predicted optimal strategy. But when they are searching for an object defined based on its orientation (e.g. a pen tilted a particular way) they exhibit a broad range of different behaviours. The results demonstrate the striking adaptability of human strategies and show us that, in order to devise a unifying theory that can explain and predict eye movements, we need to account for how they change across situations.

Authors' note.

Parts of the data and ideas appearing in this manuscript were presented at four conferences in 2022: the Scottish Vision Group meeting in Edinburgh (abstract 11), the Vision Sciences Society meeting in St Pete's Beach (abstract 4238); the European Conference on Eye Movements in Leicester, and the European Conference on Visual Perception in Nijmegen.

The data from all the experiments presented here are publicly available at <https://osf.io/9edx6/>.

This work was supported by the Economic and Social Research Council [ES/S016120/1 to ARH, PI, ADFC Co-I]. Additional support came from the James S. McDonnell Foundation (Scholar award to ARH). We wish to extend our gratitude for the contribution of Jessica Christie, who collected and prepared the icon stimuli and collected the data for the first experiment. We also wish to thank Jacqueline von Seth for helping with piloting the experiments and additional data collection.

For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission.

Author Contributions Statement

Conceptualization	AN, ADFC, ARH
Methodology	AN, JR, ADFC, ARH
Software	AN, JR, ADFC
Validation	AN, ADFC, ARH
Formal analysis	AN, ADFC
Investigation	AN, JR
Resources	AN, JR, ADFC, ARH
Data Curation	AN
Writing - Original Draft	AN, ARH
Writing - Review & Editing	JR, ADFC
Visualization	AN, ADFC
Supervision	ARH
Project administration	ARH
Funding acquisition	ARH (PI); ADFC (co-I)

References

- Araujo, C., Kowler, E., & Pavel, M. (2001). Eye movements during visual search: the costs of choosing the optimal path. *Vision Research*, 41, 3613-3625.
- Bacon, W.F. & Egeth, H.E. (1994). Overriding stimulus-driven attentional capture. *Perception & Psychophysics*, 55, 485-496.
- Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, 5, 617-629.
- Brainard, D.H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433-436.
- Bürkner P (2017). "brms: An R Package for Bayesian Multilevel Models Using Stan." *Journal of Statistical Software*, 80(1), 1-28.
- Chun, M.M. & Jiang, Y. (1999). Top-down attentional guidance based on implicit learning of visual covariation. *Psychological Science*, 10, 360-365.

- Clark, K., Birch-Hurst, K., Pennington, C.R., Petrie, A.C.P., Lee, J.T. & Hedge, C. (2022) Test-retest reliability for common tasks in vision science. *Journal of Vision*, 22, <https://doi.org/10.1167/jov.22.8.18>
- Clarke, A.D., Green, P., Chantler, M.J., & Hunt, A.R. (2016) Human search for a target on a textured background is consistent with a stochastic model. *Journal of Vision*, 16, 4.
- Clarke A.D.F, Nowakowska, A., & Hunt, A.R. (2019) Seeing beyond salience and guidance: the role of bias and decision in visual search. *Vision*, 3, 46.
- Clarke, A.D.F., Irons, J.L., James, W., Leber, A.B., & Hunt, A.R. (2022a). Stable individual differences in strategies within, but not between, visual search tasks. *Quarterly Journal of Experimental Psychology*. <https://doi.org/10.1177/1747021820929190>
- Clarke, A.D.F., Nowakowska, A., & Hunt, A.R. (2022b). Visual search habits and the spatial structure of scenes. *Attention, Perception & Psychophysics*, 84, 1874-1885.
- Clarke, A.D.F & Tatler, B.W. (2014). Deriving an appropriate baseline for describing fixation behaviour. *Vision Research*, 102, 41-51.
- Clarke, A.D.F., Sauerberger, K., Nowakowska, A., Rosenbaum, D.A., Zentall, T.R. & Hunt, A.R. (in-principal, stage one acceptance). Does precrastination explain why some observers are sub-optimal in a visual search task? A pre-registered report. *Royal Society Open Science*. <https://osf.io/5aq4c/>
- Cooper, L.A. (1976). Individual differences in visual comparison processes. *Perception and Psychophysics*, 19, 433-444.
- Gottlieb, J. (2023). Emerging Principles of Attention and Information Demand. *Current Directions in Psychological Science*, 32(2), 152-159.
- Gottlieb, J., Hayhoe, M., Hikosada O., & Rangel, A. (2014). Attention, reward, and information seeking. *Journal of Neuroscience*, 34, 15497-15504.
- Hayhoe, M., & Ballard, D. (2005). Eye movements in natural behavior. *Trends in Cognitive Sciences*, 9, 188-194.
- Hayhoe, M., & Ballard, D. (2014). Modelling task control of eye movements. *Current Biology*, 24, 622-628.
- Hayhoe, M., & Rothkopf, C.A. (2011). *Vision in the natural world*. Wiley Interdisciplinary Reviews: Cognitive Science, 2, 158-166.
- Hoppe, D. & Rothkopf, C.A. (2019). Multi-step planning of eye movements in visual search. *Scientific Reports*, 9, 144.
- Irons, J.L., & Leber, A.B. (2016). Choosing attentional control settings in a dynamically changing environment. *Attention, Perception and Psychophysics*, 78, 2031-2048.

- Irons, J.L., & Leber, A.B. (2018). Characterizing individual variation in the strategic use of attentional control. *Journal of Experimental Psychology: Human Perception and Performance*, 44, 1637-1654.
- Kristjansson, Å., Johannesson, O. I., & Thornton, I. M. (2014). Common Attentional Constraints in Visual Foraging. *PLoS ONE*, 9, e100752.
- Land, M.F., & McLeod, P. (2000). From eye movements to actions: how batsmen hit the ball. *Nature*, 3, 1340-1345.
- Lonnqvist, B., Elsner, M., Hunt, A.R. & Clarke, A.D.F. (2020). Modelling individual variation in visual search with reinforcement learning. [10.31234/osf.io/suj28](https://doi.org/10.31234/osf.io/suj28)
- Morvan, C., & Maloney, L.T. (2012). Human visual search does not maximize the post-saccadic probability of identifying targets. *PLoS Computational Biology*, 2012, 8, e1002342.
- Najemnik, J., & Geisler, W.S. (2005). Optimal eye movement strategies in visual search. *Nature*, 434, 387-391.
- Najemnik, J., & Geisler, W.S. (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *Journal of Vision*, 8, 1-14.
- Lavie, N., Hirst, A., de Fockert, J. W., & Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology: General*, 133(3), 339–354. <https://doi.org/10.1037/0096-3445.133.3.339>
- Neider, M.B. & Zelinsky, G.J. (2006). Scene context guides eye movements during visual search. *Vision Research*, 46, 614-621.
- Nowakowska, A., Clarke, A.D.F., Sahraie, A. & Hunt, A.R. (2016). Inefficient search strategies in simulated hemianopia. *Journal of Experimental Psychology: Human Perception and Performance*, 42, 1858-1872.
- Nowakowska, A., Clarke, A.D.F. & Hunt, A.R. (2017). Human visual search behaviour is far from ideal. *Proceedings of the Royal Society: Biology*, 284, 20162767.
- Nowakowska, A., Clarke, A.D.F. & Hunt, A.R. (2019). Practice-related changes in eye movement strategy in healthy adults with simulated hemianopia. *Neuropsychologia*, 128, 232-240.
- Nowakowska, A., Clarke, A. D. F., Hunt, A. R., & von Seth, J. (2021). Search strategies improve with practice, but not with time pressure or financial incentives. *Journal of Experimental Psychology: Human Perception and Performance*, 47, 1009–1021. <https://doi.org/10.1037/xhp0000912>
- Nowakowska, A., Clarke, A. D. F., Hunt, A. R. Visual search strategies in lab and life, 2018-2023 [Data set]. Open Science Framework. DOI 10.17605/OSF.IO/9EDX6

- Paeye, C., Schutz, A.C., & Gegenfurtner, K.R. (2016). Visual reinforcement shapes eye movements in visual search. *Journal of Vision*, 16, 1-15.
- Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10, 437-442.
- Rogers. B. (2017). *Perception*. A very short introduction.
- Russell, R., Duchaine, B., & Nakayama, K. (2009). Super-recognisers: people with extraordinary face recognition ability. *Psychonomic Bulletin & Review*, 16, 252-257.
- Tatler, B.W., Hayhoe, M.M., Land, M.F., & Ballard, D.H. (2011). Eye guidance in natural vision: Reinterpreting salience. *Journal of Vision*, 11, 1-23.
- Treisman, A.M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97-136.
- Vergheese, P. (2012). Active search for multiple targets is inefficient. *Vision Research*, 74, 61-71.
- Vickery, T.J., King, L-W, & Jiang, Y. (2005). Setting up the target template in visual search. *Journal of Vision*, 5, <https://doi.org/10.1167/5.1.8>
- Walthew, C., & Gilchrist, I. D. (2006). Target location probability effects in visual search: an effect of sequential dependencies. *Journal of Experimental Psychology: Human Perception and Performance*, 32(5), 1294.
- Wolfe, J.M. (2021). Guided Search 6.0: An updated model of visual search. *Psychonomic Bulletin & Review*, 28, 1060-1092.
- Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 15(3), 419-433.
- Wolfe, J. M., Butcher, S. J., Lee, C., & Hyle, M. (2003). Changing your mind: on the contributions of top-down and bottom-up guidance in visual search for feature singletons. *Journal of Experimental Psychology. Human Perception and Performance*, 29(2), 483-502.
- Xu, Y. & Chun, M.M. (2009). Selecting and perceiving multiple visual objects. *Trends in Cognitive Sciences*, 13, 167-174.
- Zhu, Q, Song, Y., Hu, S. et al.(2010). Heritability of the specific cognitive ability of face perception. *Current Biology*, 20, 137-142.