

# Learning where to look: An empirical, computational and theoretical account of hidden target search performance

Leanne Chukoskie<sup>1</sup>, Joseph Snider<sup>1</sup>, Michael C Mozer<sup>1,2</sup>, Richard J Krauzlis<sup>3</sup> Terrence J Sejnowski<sup>1</sup>

1. Institute for Neural Computation, University of California at San Diego, La Jolla, CA 920932. Department of Computer Science and Institute of Cognitive Science, University of Colorado, Boulder, CO 80309 3. Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, MD 20892-44354. Howard Hughes Medical Institute, Computational Neurobiology Laboratory, The Salk Institute for Biological Studies, La Jolla, CA 92037 5. Division of Biological Sciences, University of California at San Diego, La Jolla, CA 92093

Submitted to Proceedings of the National Academy of Sciences of the United States of America

We choose where to look over 170,000 times per day, approximately three times per wakeful second. The frequency of these saccadic eye movements belies the complexity underlying each individual choice. Experience factors into the choice of where to look, and can be invoked to rapidly redirect gaze in a context and task-appropriate manner. Yet, remarkably little is known about how individuals learn to direct their gaze given the current context and task. To address this gap in our understanding of search behavior, we designed a new task in which participants search a novel scene for a target whose location was drawn stochastically on each trial from a fixed prior distribution. To focus on how participants learned this distribution, we made the target invisible on a blank screen, and the participants were rewarded when they fixated on the target location. In just a few trials, participants rapidly found the hidden targets by looking near previously rewarded locations and avoiding previously unrewarded locations. Learning trajectories were well characterized by a simple reinforcement-learning (RL) model that maintained and continually updated a reward map of locations. The RL model made further predictions concerning sensitivity to recent experience that were confirmed by the data. The asymptotic performance of both the participants and the RL model approached optimal performance characterized by an ideal-observer theory. These two complementary levels of explanation show how experience in a novel environment drives visual search in humans, and may extend to other forms of search such as animal foraging.

ideal observer | oculomotor | reinforcement learning | saccades | search

Our daily activities depend on successful search strategies for finding objects in our environment. Visual search is ubiquitous in routine tasks: finding one's car in a parking lot, house keys on a cluttered desk, or the button you wish to click on a computer interface. When searching common scene contexts for a target object, individuals rapidly glean information about where targets are typically located (1-9). This ability to use the "gist" of an image (3, 4) enables individuals to perform flexibly and efficiently in familiar environments. Add to that the predictable sequence of eye movements that occurs when someone is engaged in a manual task (10) and it becomes clear that despite the large body of research on how image salience guides gaze (2, 11), learned spatial associations are perhaps just as important for effectively engaging our visual environment (10, 12, 13). Ironically, however, little research has been directed to *how* individuals learn to direct gaze in a context and task-appropriate manner in novel environments.

Research relevant to learning where to look comes from the literature on eye movements, rewards and their expected value. Like all motor behavior, saccades are influenced by reward, occurring at shorter latency for more valued targets (14). In fact, finding something you seek may be intrinsically rewarding (15).

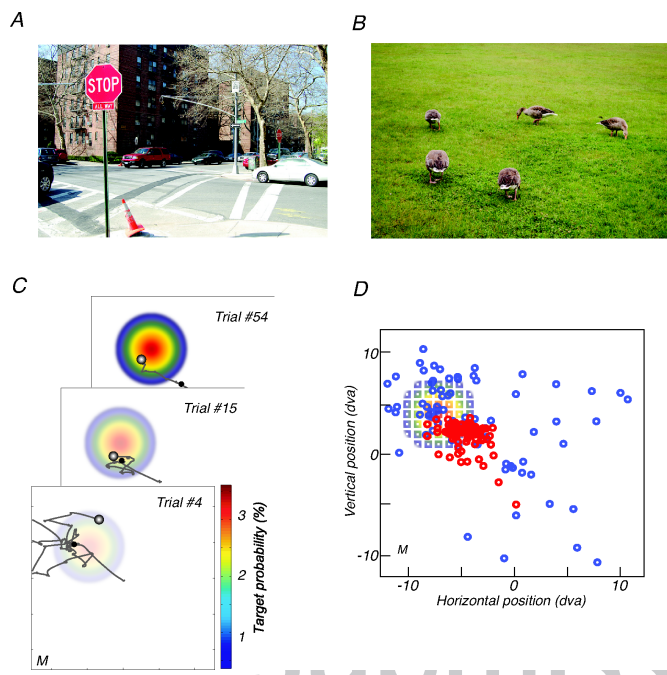
Refining the well-known canonical main sequence relationship between saccade amplitude and velocity, the value of a saccade target can alter details of the motor plan executed either speeding or slowing the saccade itself depending upon the value of that target for the subject (16, 17). This result is especially interesting in light of the research indicating that the low-level stimulus features, which have an expected distribution of attracting fixations (18), are different (19) and perhaps also differently valuable (20) depending on their distance from the current fixation location. Taken together these results underscore the complex interplay of external and internal information in guiding eye movement choice.

Two early foundational studies from Buswell (21) and Yarbus (22), foreshadowed modern concepts of a priority or salience map by showing that some portions of an image are fixated with greater likelihood than others. Both researchers also provided early evidence that this priority map effectively changes depending on the type of information sought. Yarbus observed that the patterns of gaze that followed different scene-based questions or tasks given to the observer were quite distinct, suggesting that the observer knew where to find information in the scene to answer the question and looked specifically to areas containing that information when it was needed. Henderson and colleagues (23) have replicated this result for the different tasks of visual search and image memorization. However Wolfe and colleagues (24), using a slightly different question and task paradigm, failed to find evidence that saccade patterns were predictive of specific mental states. Regardless of specific replications of Yarbus' demonstration, it is clear that scene gist—context specific information about where objects are typically found—emerges very quickly, and guides target search of a scene with a known context (4). For example, when shown a street scene, an observer would immediately know where to look for street signs, cars and pedestrians (Fig. 1A).

Castelhano and Heaven (9) have also shown that in addition to scene gist itself, learned spatial associations guide eye movements during search. Subjects use these learned associations as well as other context-based experience, such as stimulus probability, and past rewards and penalties (25-27) to hone the aim of a saccadic eye movement. A recent review and commentary from

Reserved for Publication Footnotes

137  
138  
139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155  
156  
157  
158  
159  
160  
161  
162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204

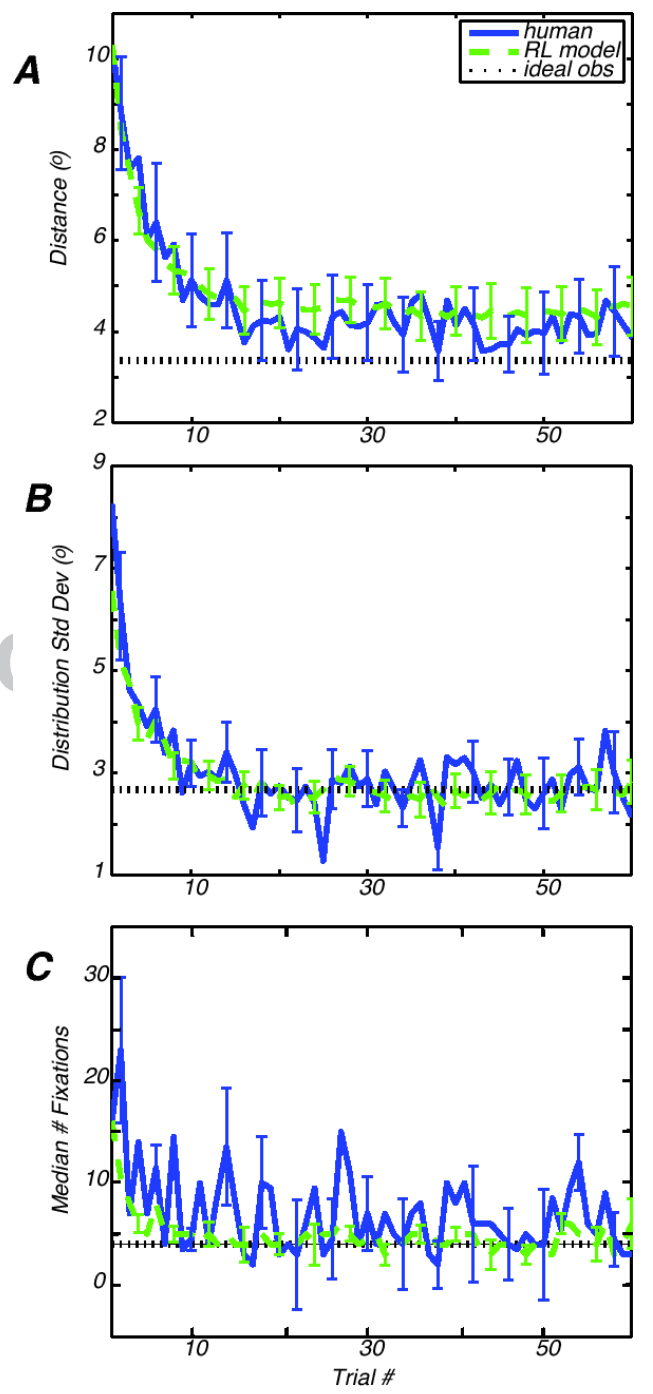


**Fig. 1.** Visible and hidden search tasks. (A) An experienced pedestrian has prior knowledge of where to look for signs, cars and sidewalks in this street scene. (B) Ducks foraging in a large expanse of grass. (C) A representation of the screen is superimposed with the hidden target distribution that is learned over the session as well as sample eye traces from three trials for participant M. The first fixation of each trial is marked with a black dot. The final and rewarded fixation is marked by a shaded gray-scale dot. (D) The region of the screen sampled with fixation shrinks from the entire screen on early trials (blue circles; 87 fixations over the first 5 trials) to a region that approximates the size and position of the Gaussian-integer distributed target locations (squares, color proportional to the probability as given in A) on later trials (red circles; 85 fixations from trials 32-39). Fixation position data is from participant M.

Wolfe and colleagues explores the notion of “semantic” guidance in complex, naturalistic scenes (28) as providing knowledge of the probability of finding a known object in a particular part of a scene. This perspective relates work on scene gist together with more classic visual search tasks, offering a framework for considering how individuals might use past experience to direct gaze in both real-world scenes as well as in the contrived scenarios of our laboratories.

Quite distinct from the literature on visual search is the literature on another sort of search that is commonly required of animals and people: foraging. Foraging agents seek food, which is often hidden in the environment in which they search (Fig. 1B). The search for hidden food rewards changes not only with the position of the reward, but also with the size of the distribution of rewards (29). Other work has cast foraging behavior in terms of optimal search (30). What distinguishes foraging from visual search tasks is that visual search tasks have visible cues that drive search, in addition to contextual information that specifies probable target location. In order to make visual search more like foraging, we can strip the visible cues from visual search. A visual search task devoid of visual cues would allow us to determine whether there are underlying commonalities between these two types of search and whether general principles of search might emerge from such an investigation.

The reinforcement-learning framework has become widely accepted for modeling performance in tasks involving a series of movements leading to reward (31, 32). In addition, for organisms across many levels of complexity, reinforcement learning has been shown to be an excellent framework to consider adaptive



**Fig. 2.** Learning curves for hidden target search task. (A) The distance between the mean of the fixation cluster for each trial to the target centroid, averaged across participants, is shown in blue and green indicates the result of 200 simulations of the reinforcement-learning model for each participant's parameters. The standard error of the mean is given for both. The ideal observer prediction is indicated by the black dotted line. (B) The standard deviation of the eye position distributions or “search spread” is shown for the average of all participants (blue) and the RL model (green) with standard error of the mean. The dashed line is the ideal observer theoretical optimum in each case, assuming perfect knowledge of the target distribution. (C) The median number of fixations made to find the target on each trial is shown (blue) along with the RL model prediction (green) of fixation number. The standard error of the mean is shown for both.

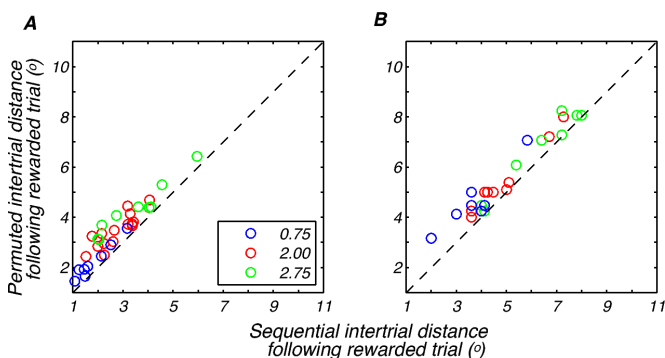
behavior in complex and changing environments (33, 34). Here we describe performance in our task in terms of a reinforcement-

205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215  
216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272

**Table 1. Performance at asymptote of learning for participants, the ideal-observer theory, and a reinforcement-learning model**


|                              | Target Spread Condition (°) | Mean distance from target centroid to fixations (°) on trials 31-60 | Search spread (°) on trials 31-60 |
|------------------------------|-----------------------------|---|-----------------------------------|
| Participant data             | 0.75                        | 1.97  | 1.14                              |
|                              | 2.00                        | 4.08  | 2.80                              |
|                              | 2.75                        | 4.39  | 3.70                              |
| Ideal Observer Theory        | 0.75                        | 0.70  | 0.56                              |
|                              | 2.00                        | 3.36  | 2.68                              |
|                              | 2.75                        | 4.74  | 3.78                              |
| Reinforcement-learning Model | 0.75                        | 3.21  | 1.56                              |
|                              | 2.00                        | 4.46  | 2.61                              |
|                              | 2.75                        | 6.07  | 4.29                              |

Data, theory, and model statistics for the mean fixation distance and search spread for 0.75°, 2.0°, and 2.75° target distribution conditions.



**Fig. 3.** Sequential effects in the human data and predictions of the RL model. (A) For each subject, we plot the mean *sequential intertrial distance* (the distance between the final fixation on trial  $n$  and the first fixation on trial  $n+1$  when trial  $n$  yields a reward) versus the *permuted intertrial distance* (the distance between the final fixation on a trial and the first fixation of another randomly drawn trial). Each circle denotes a subject, and the circle color indicates the target-spread condition (blue:  $\sigma = 0.75$ , red:  $\sigma = 2.00$ , green:  $\sigma = 2.75$ ). Consistent with the model prediction (B), the sequential intertrial distance is reliably shorter for than permuted intertrial distance, as indicated by the points lying above the diagonal. All intertrial distances are larger in the model, reflecting a greater degree of exploration than in the participants, but this mismatch is orthogonal to the sequential effects. (C) The effect of previous trials on search in the current trial is plotted as a function of the number of trials back. An exponential fit to the data is shown in green.

learning (RL) perspective. Participants' learning trajectories were well characterized by a simple RL model that maintained and continually updated a reward map of locations. The RL model made further predictions concerning sensitivity to recent experience that were confirmed by the data. The asymptotic performance of both the participants and the RL model approached optimal performance characterized by an ideal-observer theory assuming perfect knowledge of the static target distribution and independently chosen fixations. These two complementary levels of explanation show how experience in a novel environment drives visual search in humans.



**Fig. 4.** Length distributions of saccades in the hidden target task. A turning point algorithm applied to raw eye movement data yields a distribution of step sizes for all participants (see SI Materials and Methods for details). Very small "fixational" eye movements comprise the left side of the plot and large larger saccadic jumps on the right for three different sizes of target distribution. The points and lines (Loess fits with 95% confidence interval shading) for each search distribution size, all share a similar shape, particularly a bend at step sizes approaching 1 degree of visual angle.

## Results

**Humans Rapidly Learn to Find Hidden Targets.** In visual search, previous experiments failed to isolate completely the visual appearance of a target from the learned location of the reward; in all cases a visual indication of a target, or a memory of a moment-ago visible target (26) and its surroundings were available to guide the movement. To understand how participants learn where to look in a novel scene or context where no relationship exists between visual targets and associated rewards or penalties, we designed a new search task in which participants were rewarded for finding a hidden target, similar to the scenario encountered by a foraging animal (Fig. 1C).

Participants repeatedly searched a single unfamiliar scene (context) for a target. However, to study the role of task knowledge in guiding search apart from the visual cues ordinarily used to identify a target, the target was rendered *invisible*. The participants' task was to explore the screen with their gaze and find a hidden target location that would sound a reward tone when fixated. Unbeknownst to each participant, the hidden target position varied from trial to trial and was drawn from a Gaussian distribution with a centroid and spread (target mean and standard deviation, respectively) that was held constant throughout a session (see Fig. 1C).

At the start of a session, participants had no prior knowledge to inform their search; their initial search was effectively "blind". As the session proceeded participants accumulated information from gaining reward or not at fixation points and improved their success rate by developing an expectation for the distribution of hidden targets and using it to guide future search (Fig. 1D).

After remarkably few trials, participants gathered enough information about the target distribution to direct gaze efficiently near the actual target distribution, as illustrated by one participant's data in Fig. 1C,D. We observed a similar pattern of learning for all participants: early fixations were broadly scattered

throughout the search screen; after approximately a dozen trials, fixations narrowed to the region with high target probability.

A characterization of this effect for all participants is shown in Fig. 2A. The average distance from the centroid of the target distribution to individual fixations in a trial drops precipitously over roughly the first dozen trials. Fig. 2A shows this distance for all participants in the 2° target spread condition. The asymptotic distance from centroid increased monotonically with the target spread (Table 1).

A measure of *search spread* is the standard deviation of the set of fixations in a trial. The search spread was initially broad and narrowed as the session progressed, as shown in Fig. 2B for all participants in the 2° target-spread condition. The asymptotic search spread monotonically increased with the target-spread condition (Table 1). These data suggest that participants estimated the spread of the hidden target distribution and adjusted their search spread accordingly. Also, the median number of fixations that participants made to find the target (on target-found trials) decreased rapidly within a session to reach an asymptote (Fig. 2C).

**Humans Approach Ideal Observer Performance.** We now consider the behavior of participants once performance stabilized. Taking trials 31-60 to reflect asymptotic behavior, we examined the efficiency of human search in comparison to a theoretical optimum. An ideal observer was derived for the HTST assuming that fixations are independent of one another, and that the target distribution is known, and the expected number of trials is minimized. The dashed lines in Figs. 2A-C mark ideal observer performance. Ideal search performance requires a distribution of planned fixation “guesses” that is  $\sqrt{2}$  broader than the target distribution itself (35, 36). As seen in Figs. 2B,C, the performance of participants hovered around this ideal search distribution after about a dozen trials. In Fig. 2A, the mean for the human data from trials 31-60 trend higher than the theory suggests, but the theory presumes stationarity of the target distribution. However, individuals must be responsive to nonstationarities in natural environments and this responsivity yields an increase in uncertainty (37) consistent with observed human performance. Across different target distribution spreads, the ideal-observer statistics qualitatively matched those of the human participants (Table 1), and the quantitative match was excellent for 2.00° and 2.75°.

**Reinforcement Learning Model Matches Human Learning.** In addition to the ideal-observer theory, which characterizes the asymptotic efficiency of human search, we developed a complementary, mechanistic account that captured the learning, individual differences, and dynamics of human behavior. Reinforcement-learning theory, motivated by animal learning and behavioral experiments (38), suggests a simple and intuitive model that constructs a value function mapping locations in space to expected reward. The value function is updated after each fixation based on whether or not the target is found, and is used for selecting saccade destinations that are likely to be rewarded.

We augmented this intuitive model with two additional assumptions: First, each time a saccade is made to a location, the feedback obtained generalized to nearby spatial locations; second, we incorporated a proximity bias that favored shorter saccades. A preference for shorter saccades was present in the data (Fig. S4), and has been noted by other researchers (22, 39), some of whom have shown that it can override knowledge that participants have about the expected location of a target (40). Incorporating a proximity bias into the model changed the nature of the task because the choice of the next fixation became dependent on the current fixation. Consequently, participants must plan fixation sequences instead of choosing independent fixations.

We modeled the task using temporal difference methods (31), which are particularly appropriate for Markovian tasks in which

sequences of actions lead to reward (see *SI Materials and Methods* for details). The model’s free parameters were fit to each subject’s sequence of fixations for each of the first 20 trials. Given these parameters, the model was run in generative mode from a de novo state to simulate the subject performing the task.

Fig. 2 shows the mean performance of the model side-by-side with the mean human performance. The model also predicted an asymptotic search spread that increased with the target spread (Table 1), consistent with the participants’ aggregate performance. Similar to the human performance observed in Fig. 2A, the RL model approaches, but does not reach, the theoretical asymptote. Like the human participants, RL model is responsive to nonstationarity in the distribution, whereas the ideal observer theory assumes that the distribution is static. In addition, the model accounted for individual differences (see *SI Materials and Methods*). Because the model had emergent dynamics and the data used to fit the model were quite different in nature from the statistics derived from the model, the observed consistency between participants and simulation suggests that the model is more than an existence proof of a mechanism, but could provide insight into the biological mechanisms of learning (41).

Fig. 2 suggests that participants acquire the target distribution in roughly a dozen trials and then their performance is static. However, in the RL model the value function is adjusted after each fixation, unabated over time. A signature of this ongoing adjustment is a *sequential dependency* across trials – specifically, a dependency between one trial’s final fixation and the next trial’s initial fixation. Dependencies were indeed observed in the data throughout a session (Fig. 3A), as predicted by the model (Fig. 3B) and explained some of the trial-to-trial variability in performance (Fig. 2 and *SI Materials and Methods*). Participants were biased to start the next trial’s search near found target locations from recent trials. The influence of previous trials decreases exponentially, with the previous two or possibly three trials back possibly influencing the current trial’s saccade choice (Fig. 3C). This exponential damping of previous trials’ influence is approximated by the memoryless case (35), allowing both the RL model and ideal planner to coexist asymptotically.

**Bimodal Distribution of Saccade Lengths.** Our motivation in designing the hidden target search task was to link the visual search and foraging literatures. Performance in our task had features analogous to those found in the larger context of animal foraging (Fig. 4). While individual trials look like Lévy flights—a mixture of fixation and sporadic large excursions—that are known to be optimal in some cases of foraging behavior (42-44) the length distribution of all straight line segments is not Lévy like, but separates into two distinct length scales like the intermittent search popularized by Bénichou (30). The shorter length scale, fixations less than about 1°, corresponds to a local power law search with a very steep exponent, making it a classic random walk that densely samples the local space. That local search is combined with the larger, but rarer, saccades represented by the peaked hump at step sizes larger than 1°. These are the distinct choices from the planned distribution described already, i.e., the guess distribution, or value function. The distinctive knee shape in Fig. 4 is similar to that found in other demanding visual search tasks (35), as well as intermittent foraging by a wide range of animals (30, 43).

## Discussion

Human search performance can be put into the more general context of animal foraging, which has close connections with reinforcement-learning models (34) and optimal search theory (29). The hidden target search task introduced here has allowed us to separate the influence of external cues from internal prior information for seeking rewards in a novel environment (45). Our experimental results on how the distribution of hidden tar-

gets was learned through experience were well described by a reinforcement-learning model and the asymptotic performance approached the theoretical bound from a theory of optimal search behavior. The primate oculomotor system has been well studied, which will make it possible to uncover the neural mechanisms underlying the learning and performance of the hidden target task, which may be shared with other search behaviors.

In our hidden target search task, participants explored a novel environment and quickly learned to align their fixations with the region of space over which invisible targets were probabilistically distributed. After about a dozen trials, the fixation statistics came close to matching those obtained by an ideal-observer theory. This near-match allowed us to cast human performance as optimal memory-free search with perfect knowledge of the target distribution. As a complement to the ideal-observer theory that addresses asymptotic performance, we developed a mechanistic account of trial-to-trial learning from reinforcement. Our reinforcement-learning (RL) model characterized the time course of learning, attained an asymptote near ideal-observer performance, and tied the problem of visual search to a broader theory of motivated learning.

**Natural environments.** The ideal-observer and reinforcement-learning frameworks provide the foundation for a broader theoretical perspective on saccade choice during natural vision, in which people learn to search in varied contexts for visible targets, where visual features of the scene are clearly essential. In a Bayesian framework, the subjects in our task learned the prior distribution of the hidden targets. In a natural environment, the prior distribution would be combined with visual information to determine the posterior distribution, from which saccadic targets are generated.

Naturalistic environments are non-stationary. For example, an animal foraging for food may exhaust the supply in one neighborhood and have to move on to another. A searcher must be sensitive to such changes in the environment. Sequential dependencies (Fig. 3) are a signature of this sensitivity (37, 46, 47): recent targets influence subsequent behavior, even after the searcher has seemingly learned the target distribution, as reflected in asymptotic performance. Sequential dependencies were predicted by the RL model, which generated behavior remarkably close to the participants as a group, and also captured individual idiosyncrasies (see *SI Materials and Methods*). Sensitivity to non-stationary environments can explain why our participants and the RL model attained an asymptotic search distribution somewhat further from the target centroid than is predicted by an ideal-observer theory premised on stationarity.

**Neural Basis of Search.** The neurobiology of eye movement behavior offers an alternative perspective on the similarities of visual search behavior and foraging. The question of where to look next has been explored neurophysiologically, and cells in several regions of the macaque brain appear carry signatures of task components required for successful visual search. The lateral intraparietal area (LIP) and the superior colliculus are two brain regions that contain a priority map representing locations of relevant stimuli that could serve as the target of the next saccade. Recordings in macaque area LIP and the SC have shown that this priority map integrates information from both external (“bottom-up”) and internal (“top-down”) signals in visual search tasks (48, 49).

Recently, Bisley and colleagues have used a foraging-like visual search task to show that area LIP cells differentiated between targets and distracters, and kept a running estimate of likely saccade goal payoffs (50). Area LIP neurons integrate information from different foraging-relevant modalities to encode the value associated with a movement to a particular target (51, 52) The neural mechanisms serving patch stay-leave foraging decisions

have recently been characterized in a simplified visual choice task (53), providing a scheme for investigations of precisely how prior information and other task demands mix with visual information available in the scene. Sub-threshold microstimulation in area LIP (54) or the SC (55) also biases the selection saccades toward the target in the stimulated field. Taken together, these results suggest that area LIP and the SC might be neural substrates mediating the map of likely next saccade locations in our task, akin to the value map in our RL model.

We asked how subjects *learn* to choose valuable targets in a novel environment. Recent neurophysiological experiments in the basal ganglia provide some suggestions on how prior information is encoded for use in choosing the most valuable saccade target in a complex environment (56). Hikosaka and colleagues have identified signals related to recently learned, and still labile, value information for saccade targets in the head of the caudate nucleus and more stable value information in the tail of the caudate and substantia nigra, pars reticulata (SNr) (57,57). As the cells carrying this stable value information appear to project preferentially to the SC, these signals are well-placed to influence saccade choices through a fast and evolutionarily conserved circuit for controlling orienting behavior. These results provide a neurophysiological basis for understanding how experience is learned and consolidated in the service of the saccades we make to gather information about our environment approximately 3 times each second.

**Conclusions.** In our eye movement search task, subjects learn to choose saccade goals based on prior experience of reward that is divorced from specific visual features in a novel scene. The resulting search performance was well described by a reinforcement-learning model similar to that used previously to examine both foraging animal behavior and neuronal firing of dopaminergic cells. In addition, the search performance approached the theoretical optimum for performance on this task. By characterizing how prior experience guides eye movement choice in novel contexts and integrating it with both model and theory, we have created a framework for considering how prior experience guides saccade choice during natural vision.

## Methods

We defined a spatial region of an image as salient by associating it with reward to examine how participants used their prior experience of finding targets to direct future saccades. We took advantage of the fact that the goal of saccadic eye movements is to obtain information about the world and asked human participants to “conduct an eye movement search to find a rewarded target location as quickly as possible.” Participants (N=7) were also told that they would learn more about the rewarded targets as the session progressed and that they should try to find the rewarded target location as quickly as possible. The rewarded targets had no visual representation on the screen and were thus invisible to the subject. The display screen was the same on each trial (either a blank mean grey or 1/f noise surrounded by mean grey) within a session and provided no information about the target location. The location and the spread ( $\sigma = 0.75^\circ, 2.0^\circ, \text{ and } 2.75^\circ$ ) of the rewarded target distribution were varied with each session. Each trial began with a fixation cross on a blank screen. After the subject acquired the fixation cross with her gaze and maintained fixation for 300 ms, the fixation cross disappeared, leaving the search screen, which was either blank or 1/f “pink” noise. No quantitative difference in the pattern of fixations was observed between the conditions across subjects, so the two conditions were collapsed here. Gaze position was monitored in real time. When a subject’s gaze fell within  $2^\circ$  of a hidden target for at least 50 ms success was signaled with a reward tone and the promise of greater compensation, but no visual feedback. If after 20 seconds of search the target was not acquired, the trial ended with no tone. More details about observers, stimuli, equipment, and procedure are given in *SI Materials and Methods*.

**ACKNOWLEDGEMENTS:** Wish to thank Krista Kornylo and Natalie Dill for technical assistance. This work was supported in part by NSF grant #SBE 0542013 to the Temporal Dynamics of Learning Center, an NSF Science of Learning Center (LC, JS, MM, TS), a Blasker Rose-Miah grant from the San Diego Foundation (LC), ONR MURI Award No.: N00014-10-1-0072 (JS, TS), and Howard Hughes Medical Institute (TS)..

681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701  
702  
703  
704  
705  
706  
707  
708  
709  
710  
711  
712  
713  
714  
715  
716  
717  
718  
719  
720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748

1. Castelinho MS & Heaven C (2011) Scene context influences without scene gist: Eye movements guided by spatial associations in visual search. *Psychon Bull Rev*.
2. Castelinho MS, Mack ML, & Henderson JM (2009) Viewing task influences eye movement control during active scene perception. *J Vis* 9(3):6 1-15.
3. Yarbus AL (1967) *Eye Movements and Vision* (Plenum, New York).
4. Oliva A & Torralba A (2006) Building the gist of a scene: the role of global image features in recognition. *Prog Brain Res* 155:23-36.
5. Torralba A, Oliva A, Castelinho MS, & Henderson JM (2006) Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search. *Psychol Rev* 113(4):766-786.
6. Potter MC (1975) Meaning in visual search. *Science* 187(4180):965-966.
7. Itti L & Koch C (2000) A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Res* 40(10-12):1489-1506.
8. Neider MB & Zelinsky GJ (2006) Scene context guides eye movements during visual search. *Vision Res* 46(5):614-621.
9. Rayner K, Castelinho MS, & Yang J (2009) Eye movements when looking at unusual/weird scenes: are there cultural differences? *J Exp Psychol Learn Mem Cogn* 35(1):254-259.
10. Vo ML & Henderson JM (2010) The time course of initial scene processing for eye movement guidance in natural scene search. *J Vis* 10(3):14 11-13.
11. Castelinho MS & Heaven C (2010) The relative contribution of scene context and target features to visual search in scenes. *Atten Percept Psychophys* 72(5):1283-1297.
12. Hayhoe M & Ballard D (2005) Eye movements in natural behavior. *Trends Cogn Sci* 9(4):188-194.
13. Parkhurst DJ & Niebur E (2003) Scene content selected by active vision. *Spat Vis* 16(2):125-154.
14. Tatler BW & Vincent BT (2009) The prominence of behavioural biases in eye guidance. *Visual Cognition* 17(6-7):1029-1054.
15. Chun MM & Jiang Y (1998) Contextual cueing: implicit learning and memory of visual context guides spatial attention. *Cogn Psychol* 36(1):28-71.
16. Milstein DM & Dorris MC (2007) The influence of expected value on saccadic preparation. *J Neurosci* 27(18):4810-4818.
17. Xu-Wilson M, Zee DS, & Shadmehr R (2009) The intrinsic value of visual information affects saccade velocities. *Exp Brain Res* 196(4):475-481.
18. Shadmehr R (2010) Control of movements and temporal discounting of reward. *Curr Opin Neurobiol* 20(6):726-730.
19. Shadmehr R, Orban de Xivry JJ, Xu-Wilson M, & Shih TY (2010) Temporal discounting of reward and the cost of time in motor control. *J Neurosci* 30(31):10507-10516.
20. Reinagel P & Zador AM (1999) Natural scene statistics at the centre of gaze. *Network* 10(4):341-350.
21. Tatler BW, Baddeley RJ, & Vincent BT (2006) The long and the short of it: spatial statistics at fixation vary with saccade amplitude and task. *Vision Res* 46(12):1857-1862.
22. Acik A, Sarwary A, Schultze-Kraft R, Onat S, & Konig P (2010) Developmental Changes in Natural Viewing Behavior: Bottom-Up and Top-Down Differences between Children, Young Adults and Older Adults. (Translated from eng) *Front Psychol* 1:207.
23. Buswell GT (1935) *How people look at pictures: a study of the psychology of perception in art* (University of Chicago Press, Chicago, IL).
24. Greene MR, Liu T, & Wolfe JM (2012) Reconsidering Yarbus: a failure to predict observers' task from eye movement patterns. *Vision Res* 62:1-8.
25. Schutz AC, Trommershauser J, & Gegenfurtner KR (2012) Dynamic integration of information about salience and value for saccadic eye movements. *Proc Natl Acad Sci U S A* 109(19):7547-7552.
26. Stritzke M & Trommershauser J (2007) Eye movements during rapid pointing under risk. *Vision Res* 47(15):2000-2009.
27. Geng JJ & Behrmann M (2005) Spatial probability as an attentional cue in visual search. *Percept Psychophys* 67(7):1252-1268.
28. Wolfe JM, Vo ML, Evans KK, & Greene MR (2011) Visual search in scenes involves selective and nonselective pathways. *Trends Cogn Sci* 15(2):77-84.
29. Charnov EL (1976) Optimal Foraging, The Marginal Value Theorem. *Theoretical Population Biology* 9(2):129-136.
30. Bénichou O, Coppel M, Moreau M, Suet P-H, & Voituriez R (2005) Optimal Search Strategies for Hidden Targets. *Phys. Rev. Lett.* 94(198101):1-4.
31. Sutton RS (1988) Learning to predict by the method of temporal differences. *Machine Learning* 3:9-44.
32. Montague PR & Sejnowski TJ (1994) The predictive brain: temporal coincidence and temporal order in synaptic learning mechanisms. *Learn Mem* 1(1):1-33.
33. Lee D, Seo H, & Jung MW (2012) Neural Basis of Reinforcement Learning and Decision Making. *Annu Rev Neurosci* 35:287-308.
34. Niv Y, Joel D, Meilijson I, & Ruppin E (2002) Evolution of Reinforcement Learning in Uncertain Environments: A Simple Explanation for Complex Foraging Behaviors. *Adaptive Behavior* 10(1):5-24.
35. Snider J (2010) Optimal random search for a single hidden target. *Physical Review E* 83(1):13.
37. Yu AJ & Cohen JD (2009) Sequential effects: Superstition or rational behavior? . *Advances in Neural Information Processing Systems*, (MIT Press, Cambridge, MA), pp 1873-1880.
38. Sutton RS & Barto AG (1998) *Reinforcement learning : an introduction* (MIT Press, Cambridge, Mass.) pp xviii, 322 p.
39. Rayner K (1998) Eye movements in reading and information processing: 20 years of research. *Psychol Bull* 124(3):372-422.
40. Araujo C, Kowler E, & Pavel M (2001) Eye movements during visual search: the costs of choosing the optimal path. *Vision Res* 41(25-26):3613-3625.
41. Schultz W, Dayan P, & Montague PR (1997) A neural substrate of prediction and reward. *Science* 275(5306):1593-1599.
42. Humphries NE, et al. (2010) Environmental context explains Levy and Brownian movement patterns of marine predators. *Nature* 465(7301):1066-1069.
43. James A, Plank MJ, & Edwards AM (2011) Assessing Lévy walks as models of animal foraging. *Journal of the Royal Society Interface*:1-15.
44. Viswanathan GM, et al. (1999) Optimizing the success of random searches. *Nature* 401(6756):911-914.
45. Adams GK, Watson KK, Pearson J, & Platt ML (2012) Neuroethology of decision-making. *Curr Opin Neurobiol* 22(6):982-989.
46. Fecteau JH & Munoz DP (2003) Exploring the consequences of the previous trial. *Nat Rev Neurosci* 4(6):435-443.
47. Wilder MH, Mozer MC, & Wickens CD (2011) An integrative, experience-based theory of attentional control. *J Vis* 11(2).
48. Bisley JW & Goldberg ME (2010) Attention, intention, and priority in the parietal lobe. *Annu Rev Neurosci* 33:1-21.
49. Fecteau JH & Munoz DP (2006) Saliency, relevance, and firing: a priority map for target selection. *Trends Cogn Sci* 10(8):382-390.
50. Mirpour K, Arcizet F, Ong WS, & Bisley JW (2009) Been there, seen that: a neural mechanism for performing efficient visual search. *J Neurophysiol* 102(6):3481-3491.
51. Klein JT, Deaner RO, & Platt ML (2008) Neural correlates of social target value in macaque parietal cortex. *Curr Biol* 18(6):419-424.
52. Platt ML & Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400(6741):233-238.
53. Hayden BY, Pearson JM, & Platt ML (2011) Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci* 14(7):933-939.
54. Mirpour K, Ong WS, & Bisley JW (2010) Microstimulation of posterior parietal cortex biases the selection of eye movement goals during search. *J Neurophysiol* 104(6):3021-3028.
55. Carello CD & Krauzlis RJ (2004) Manipulating intent: evidence for a causal role of the superior colliculus in target selection. *Neuron* 43(4):575-583.
56. Nakahara H & Hikosaka O (2012) Learning to represent reward structure: a key to adapting to complex environments. *Neurosci Res* 74(3-4):177-183.
57. Yasuda M, Yamamoto S, & Hikosaka O (2012) Robust representation of stable object values in the oculomotor Basal Ganglia. *J Neurosci* 32(47):16917-16932.

749  
750  
751  
752  
753  
754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809  
810  
811  
812  
813  
814  
815  
816