

The role of uncertainty and reward on eye movements in a virtual driving task

Brian T. Sullivan

Smith-Kettlewell Eye Research Institute,
San Francisco, CA, USA
Department of Psychology,
University of Texas at Austin, Austin, Texas, USA



Leif Johnson

Department of Computer Science,
University of Texas at Austin, Austin, Texas, USA



Constantin A. Rothkopf

Frankfurt Institute for Advanced Studies,
Goethe University Frankfurt, Germany
Institute of Cognitive Science, University Osnabrück,
Osnabrück, Germany



Dana Ballard

Department of Computer Science,
University of Texas at Austin, Austin, Texas, USA



Mary Hayhoe

Department of Psychology,
University of Texas at Austin, Austin, Texas, USA



Eye movements during natural tasks are well coordinated with ongoing task demands and many variables could influence gaze strategies. Sprague and Ballard (2003) proposed a gaze-scheduling model that uses a utility-weighted uncertainty metric to prioritize fixations on task-relevant objects and predicted that human gaze should be influenced by both reward structure and task-relevant uncertainties. To test this conjecture, we tracked the eye movements of participants in a simulated driving task where uncertainty and implicit reward (via task priority) were varied. Participants were instructed to simultaneously perform a Follow Task where they followed a lead car at a specific distance and a Speed Task where they drove at an exact speed. We varied implicit reward by instructing the participants to emphasize one task over the other and varied uncertainty in the Speed Task with the presence or absence of uniform noise added to the car's velocity. Subjects' gaze data were classified for the image content near fixation and segmented into looks. Gaze measures, including look proportion, duration and interlook interval, showed that drivers more closely monitor the speedometer if it had a high level of uncertainty, but only if it was also associated with high task priority or implicit reward. The interaction observed appears to be an example of a simple mechanism whereby the reduction of visual uncertainty is gated by behavioral relevance. This lends qualitative support for the primary variables controlling gaze allocation proposed in the Sprague and Ballard model.

Keywords: eye movements, cognition, driving, natural tasks, reinforcement learning

Citation: Sullivan, B. T., Johnson, L., Rothkopf, C. A., Ballard, D., & Hayhoe, M. (2012). The role of uncertainty and reward on eye movements in a virtual driving task. *Journal of Vision*, 12(13):19, 1–17, <http://www.journalofvision.org/content/12/13/19>, doi:10.1167/12.13.19.

Introduction

The deployment of visual attention, in particular eye movements, involves an interplay between top-down and bottom-up determinants of gaze (Fecteau, 2007; Folk, Remington, & Wright, 1994; Hayhoe & Ballard, 2005; Itti & Baldi, 2006; Knudsen, 2007; Wolfe, Butcher, Lee, & Hyle, 2003). It is known that image salience plays a role in the guidance of visual attention

and eye movements (Folk & Remington, 1998; Forster & LaVie, 2008; Leber & Egeth, 2006; Theeuwes, 2004; Triesman & Gelade, 1980; Wolfe & Horowitz, 2004; Yantis & Egeth, 1999) and that top-down goals can modulate this role (Folk & Remington, 1998; Lu & Han, 2009). Stimulus-driven control is well studied and has been used as a framework for computational modeling of human vision (Bruce & Tsotsos, 2009; Itti & Baldi, 2006; Itti & Koch, 2001; Kanan, Tong, Zhang, & Cottrell, 2008). This framework has been primarily

applied to viewing of static two-dimensional displays. Natural behavior on the other hand is embedded in a perception and action cycle where the human observer selects and acts on information in the world and must in turn deal with the dynamic consequences that those actions induce on the world, in addition to ongoing sensory events. It may be the case that vision has evolved in such a way that many stimuli that can capture attention or generate eye movements are relevant to survival. However, in complex interactive scenes and in the context of flexible task goals, there may be many visually salient stimuli that are irrelevant to the current goals and there must be some arbitration mechanism to evaluate what stimuli need to be attended to at the current moment. In this paper we focus on how this arbitration might be achieved in the presence of multiple potential targets for gaze.

Numerous experiments monitoring eye movements in natural tasks have shown that human gaze is tightly linked to ongoing task demands (Droll, Hayhoe, Triesch, & Sullivan, 2005; Hayhoe, Bensinger, & Ballard, 1998; Jovancevic & Hayhoe, 2009; Jovancevic-Misic, Hayhoe, & Sullivan, 2006; Land & Hayhoe, 2001; Land, Mennie, Rusted, 1999; Pelz, Hayhoe, & Loeber, 2001; Tatler, Hayhoe, Land, & Ballard, 2011). However, unlike the study of the bottom-up control of vision, there are few computational models of top-down control that have been proposed to explain gaze in natural tasks. In part this is because it is unclear how task structure should be represented, in contrast to the more straightforward image processing algorithms often found in bottom-up models. However, given the pervasive influence of task goals in gaze behavior, it is important to develop ‘a theory of tasks’ to understand how sensory information can be used to guide motor output towards some set of desired states. Theoretically, there are several ways one might approach this, in particular Sprague and Ballard (2003; Sprague, Ballard, & Robinson, 2007) have proposed a modular architecture for goal oriented visuomotor control and suggested that eye movements may be driven by two key parameters: reward and uncertainty. Within animal learning, the terms positive and negative reinforcement respectively refer to presenting a learner with an appetitive reward or withholding that reward. Similarly, positive and negative punishment refer to presenting or withholding an aversive stimulus. Note, however, that Sprague and Ballard describe these situations by using reward as a blanket term to refer to a numerical representation of an external learning signal that could be either numerically positive (appetitive), negative (aversive) or zero (neutral) and encompasses all of the above distinctions. This follows the naming tradition concerning Markov decision processes, an underlying mathematical framework for reinforcement learning,

that use the generic term reward function for a mapping between a state of the world and a learning signal. In this article we use reward in this general sense of a utility function for brevity although it lacks precision. Their model uses a set of context-specific task modules, which individually represent state variables and their uncertainty for their respective tasks. Over time these uncertainties increase and can only be reduced by obtaining a sensory measurement through an eye movement. By tracking the respective uncertainties of the state variables in each module and using the individual tasks’ rewards, one can compute an expected value of obtainable reward. If the expected value of reward for updating a particular task module is high, then gaze is allocated to update this module. A central premise of the model is that complex behavior can be broken down into a set of independent subtasks, and visual attention is allocated sequentially between these different tasks. Importantly the model allows flexible prioritization of visual tasks via reward weighting. Once the modules have been trained, their respective reward tables are normalized and each can be weighted (with the sum of weights across modules equaling one). The reward weighting on a module is proportional to its task priority and will directly influence how often that visual task receives new sensory information. On the face of it, this seems to describe the task selectivity of a wide range of natural behaviors and has the potential to guide our understanding of how gaze is allocated between competing task demands. While this model has been further developed for new visuomotor control scenarios (Nunez-Varela, Ravindran, & Wyatt, 2012, May; Rothkopf & Ballard, 2010, Sullivan, Johnson, Ballard, & Hayhoe, 2011), there has been little work addressing if and how the human visual system incorporates reward and uncertainty to control gaze in natural tasks. In this study our goal was to further the understanding of how these variables might be used by the visual system for eye movement control and provide behavioral observations for further modeling.

Reward and task incentive

Achievement of goals needs to be monitored via some sort of feedback. Reinforcement learning (RL) is a useful mathematical tool for function approximation in control problems (i.e., given an input with particular dynamics generate a desired output) and can allow a simulated agent to learn a variety of complex behaviors. It is also the case that many problems solved by animals may be cast in the reinforcement-learning framework. Additionally, there is considerable evidence that a variety of cortical and subcortical areas involved in the generation of

saccadic eye movements have activity that is correlated with reward and reward prediction. Neural firing in eye movement control regions in the cortex (lateral intraparietal area, frontal eye fields, supplementary eye fields, dorsolateral prefrontal cortex) has been shown to correlate with reward, although there is still much debate on the exact computation being carried out and what mathematical constructs best explain such activity (Deaner, Khera, & Platt, 2005; Dorris & Glimcher, 2004; Glimcher, 2003; Platt & Glimcher, 1999; Seo & Less, 2007; Stuphorn, Taylor, & Schall, 2000; Stuphorn & Schall, 2006; Sugrue, Corrado, & Newsome, 2004). LIP in particular has been subject to great study and it has been suggested it represents a combination of signals concerning knowledge and uncertainty regarding visual input, behavioral relevance, and action output to some degree (Gottlieb & Balan, 2010). The output from these cortical areas ultimately converge on the caudate nucleus in the basal ganglia including the caudate, putamen, and substantia nigra, all of which have neural activity correlated with reward. Study of the dopaminergic circuits in these areas has provided several proposed mechanisms for how such activity may be generated. This cortico-basal ganglia-superior collicular circuit appears to regulate the control of fixations and the timing of planned movements (Aggarwal, Hyland, & Wickens, 2012; Hikosaka, Nakamura, & Nakahara 2006, Lauwereyns et al., 2002; Watanabe, Lauwereyns, & Hikosaka, 2003). Additionally, behavioral studies have found evidence that human observers' search and gaze behavior can be sensitive to reward structure, although it is debated if humans optimize gaze behavior for reward collection (Stritzke, Trommershäuser, & Gegenfurtner, 2009; Hayhoe & Ballard, 2005; Jovancevic-Misic & Hayhoe, 2009; Navalpakkam, Koch, & Perona, 2009; Navalpakkam, Koch, Rangel, & Perona, 2010; Stritzke & Trommerhauser, 2007; Schutz, Trommesrhauser, & Gegenfurtner, 2012).

Given this evidence, it appears that there is a neural architecture in place that could be well suited to the learning and control of overt visual attention. However, it is worth noting that these types of studies have largely looked only at learning with a primary reinforcer, e.g., a monkey gets a sip of juice immediately after making a correct saccadic movement. While there are situations where visual stimuli might be directly rewarding (e.g., social situations), it is unlikely that representing primary reinforcers is the sole function of the circuitry described above (Deaner et al., 2005; Hayhoe & Ballard, 2005; Jovancevic-Misic & Hayhoe, 2009; Lee, Seo, & Jung, 2012; Sohn & Lee, 2006). Instead, during natural activity, animals must conduct a series of actions to gain a reward, e.g., scavenging for food. In human behavior, gaze changes

accrue information, which could confer secondary reward as steps are taken to achieve a desired behavioral goal state. In RL models such as the Sprague and Ballard model, reward contingencies for a particular state of the world incorporate both immediate gains as well as future discounted reward associated with future possible states. Theoretically, there has been a great deal of work examining how systems can use this kind of learning to acquire complex behavior that requires multiple steps (Sutton, 1988; Sutton & Barto, 1998). Experimental work examining operant conditioning and shaping have shown that the reward properties of an unconditioned stimulus can be associated with predictive stimuli far in advance of the final reward state (Schultz, 1998). This evidence suggests that visual attention and deployment of gaze may have the capacity to be learned over long time scales. This idea provides a possible low-level explanation for why natural vision appears to be dominated by task influences. Gaze strategies could be learned over time in a way that facilitates the collection of information that will eventually lead to a rewarding goal state for a task.

Humans have enormous flexibility in the selection of visual tasks and goal states. These task priorities can be altered both endogenously and exogenously. The concepts of motivation, utility, task priority and attention are deeply intertwined (Maunsell, 2006). In decision theory a utility structure of some sort is required to motivate why any decision would be made at all. While it is unclear exactly how humans select a subset of tasks from the wide variety that are available to be engaged in at any given point in time, it is likely that this prioritization is at least in part driven by external utility and internal utility (i.e., intrinsic motivation). As mentioned above, one way such flexibility might be implemented computationally is via a set of task weights that alter the expected rewards and costs for learned behaviors. From this perspective task priority and reward/cost weighting have a direct relationship.

In our driving study, all participants have had years of experience and presumably have learned the structure of rewards and punishments associated. As detailed below in our method section, we manipulate task priority via verbal instruction. We suggest that although this is an indirect manipulation, these instructions provide an implicit reward structure to participants that takes advantage of their prior experience and ability to set cognitive goals, which presumably have their own reward structure based around the instructions. While we do not directly estimate subject's subjective reward structure in this experiment, this type of interpretation is useful to understand how further research may proceed to unite high level behavioral findings with low level accounts of neural representation of reward.

Uncertainty

In addition to reward (in the general sense described above), sensory uncertainty also appears to be an important factor in human gaze behavior. Uncertainty can have a number of interpretations, since there are both external and internal sources of uncertainty. For example, uncertainty might be a consequence of inherent variability in the external stimulus or a consequence of information losses in the sensory apparatus as a consequence of low contrast or peripheral resolution losses, or other internal factors, such as memory decay or an outdated sensory signal. Some experiments have manipulated uncertainty about reward probability, and this has been shown to influence activity of both midbrain dopamine neurons and LIP neurons (Bromberg-Martin & Hikosaka, 2009; Peck, Jangraw, Suzuki, Efem, & Gottlieb, 2010). In our experiments below, we treat uncertainty as the variance of the probability distribution associated with a belief that the world is in a particular state given a set of visual observations over time.

Najemnik and Geisler (2005) used uncertainty measures for visibility in the peripheral retina to show that observers are similar to a Bayesian ideal observer that uses such uncertainty measures to optimize saccades in visual search. Renninger, Verghese, and Coughlan (2007) developed a model of eye movements that used an entropy measure of object features to predict how human observers foveate and compare complex visual shapes. They found that while human observers do not appear to foveate targets in a way that reduces global uncertainty, behavior is consistent with reducing local uncertainty. Senders (1980) also formulated an entropy-based uncertainty model of visual information accrual during driving. Using a set of estimates and parameters concerning how much data is present on a road, how fast this information can be absorbed, and how soon it is outdated or forgotten he was able to establish a relationship between the velocity a car traveled and information flux over time. Senders demonstrated that human drivers self-regulate their speed, contingent on information flow (e.g., if vision is disrupted via a blackout period they will slow down in a systematic manner). Additionally, it has been shown that visuomotor judgments can incorporate measures of exogenous uncertainty and often be carried out in a rational manner (Atkins et al., 2003; Graf, Warren, & Maloney, 2005; Schlicht & Schrater, 2007; Warren, Graf, Champion, & Maloney, 2012). However, similar to debates on optimality of reward-based behavior, it has been suggested that the human visual system may not incorporate uncertainty in an optimal fashion, and may be simply heuristic (Morvan & Maloney, 2012; Zhang, Morvan, & Maloney, 2010).

The goal of the present experiment is not to evaluate optimality, but to merely establish a link between reward and uncertainty in eye movement scheduling in a naturalistic task. In particular, we wanted to know if eye movements trade off between sources of uncertainty and intrinsic reward determined by task goals. To do this we set up a novel driving simulation environment where participants' eyes could be tracked while driving in a realistic three-dimensional environment with the experimental control available to vary task rewards and uncertainty. The rationale for the experimental manipulations derives from the structure of Sprague and Ballard's scheduling model and its underlying architecture for visuomotor processes. The model simplifies visual processing to consider only foveal vision and predicts where foveal attention is directed over time. A core idea in the model is that visual attention is deployed over time to a set of independent tasks that require different pieces of information and are serially updated by deployment of gaze to the relevant location in the scene. Each task is modeled as a control module (acquired via reinforcement learning) that receives an estimate of the state of the world and uses a reward mapping between this state and a set of actions available to 'vote' for a particular action, given that module's information about the world. For example, in Sprague and Ballard's simulations, an agent learned to navigate a sidewalk using visual information, for three modules including one module that was designed to keep the agent in the center of the sidewalk, one to avoid obstacles and another to 'pick up' certain objects. Each module maintains an estimate of its location in its state space, i.e., the dimensions of its sensory estimates (e.g., this sidewalk module's state space included the agent's distance to the sidewalk center and its relative angle to the center). This estimate of state uncertainty is computed using a Kalman filter, which provides a mean and variance associated with each module's sensory inputs. Estimates of reward and uncertainty can be combined together to form an expected value for each module. This can be used as a measure that allows a cross module comparison to determine, which single module has the most to gain if it receives new sensory information. If a module is not selected, its mean and variance estimates are propagated into the future without new sensory information and gradually become less accurate over time.

Given that the Sprague and Ballard model has the potential to explain some properties of natural gaze control, our goal in the current experiment was to demonstrate that the primary variables in the model, reward and uncertainty, were in fact significant determinants of the way gaze is allocated. We do not make a quantitative comparison between performance and the model predictions in this paper, but merely attempt to validate the approach.




Figure 1. Depiction of the driving simulator. (Left) View of the driving platform. (Right) Subject's view of the virtual environment in the simulator (subjects were presented with stereo image pairs). The white crosshair shows the subject's point of gaze on the speedometer. Neither the crosshair nor the eye image was visible to the subject.

Methods

Experimental design

Our primary goal in this experiment was to have participants perform a naturalistic task where there were several task goals, and uncertainty and reward could be manipulated. Given the theoretical work described above that links task goals for model reward parameters, it seems reasonable to assume that the task instructions given to subjects will effectively manipulate implicit reward. Subjects had two instructed tasks: (a) following a leader car at a constant distance and (b) maintaining a constant speed. The priority of the two tasks was varied by instruction. This strategy has the disadvantage that there is no objective specification of the reward. Explicit rewards, on the other hand, have the disadvantage that they are unlikely to reflect natural behavior, where primary rewards are typically not immediate, and may not be effective in a well-learned task such as driving. We therefore assume here that the task instructions function as a reward manipulation. We elaborate on these conditions in more detail below. To manipulate uncertainty, we chose to add noise to the car velocity. Our graphics environment incorporated a speedometer that gave subjects a constant analog readout of their speed. This speedometer could give noisy readings that fluctuate around a mean, adding uncertainty to subject's speed estimates and potentially increasing fixations, if subjects maintain a constant speed. Note, while it is possible to affect the speedometer reading only, we decided to

introduce noise into the car's gas pedal signal. This was chosen because drivers in pilot studies experienced difficulty when the speedometer's gauge visually fluctuated but did not influence the car's behavior. In our conditions with noise present (detailed below), noise from a uniform distribution was added to the subject's car velocity. In the Sprague and Ballard model, task modules are assumed to be independent and unique uncertainty distributions can be introduced to individual task modules. Our manipulation was intended to only affect the speed maintenance task. But, this speed variability inevitably also affected leader following performance. While noise made both speed maintenance and leader following more difficult, our results indicate that this manipulation primarily affected eye movements concerned with speed maintenance. In all conditions, the leader car drove at an average of 60.8 kph, $SD=7.2$, (37.8 mph, $SD=4.3$), close to the subject's instructed velocity of 40 mph but sufficiently varied to encourage engagement in both tasks. Given our manipulations of reward and uncertainty, our expectation based on the Sprague and Ballard model was that both high uncertainty and high reward would increase the priority of foveating a particular task-related object, if these factors are indeed important in controlling natural gaze behavior.

Driving platform overview

The virtual driving platform consisted of a stripped-down car cab interior with a seat, steering wheel, brake, and gas pedal (see Figure 1). The wheel and pedals were

connected to an analog-to-digital converter that allowed real-time positional measurement for control of the driving simulation software. The simulated car had an automatic transmission that required no interaction from the driver other than using the gas pedal. Additionally, a transducer speaker was mounted onto the car seat to provide vibration and sound that were proportional to the activity of the engine in the virtual environment.

Participants wore a head-mounted display (HMD), a NVIS nVisor SX111 (NVIS, Inc., Reston, VA) with $\sim 102^\circ \times 64^\circ$ binocular field of view, running at a resolution of 1280×1024 and updated at 60 hz. A Polhemus Fastrak motion tracking system (Polhemus, Inc., Colchester, VT) was mounted on the HMD and recorded subjects' head movements with 6 degrees of freedom and was used to update the image at 60 hz. The HMD display was rendered by in-house software running on a Dell Precision T7500 running Windows XP, using an Intel Xeon E5507, 4GB memory and a NVIDIA GeForce GTX 460 video card.

An Arrington Research ViewPoint EyeTracker (Arrington Research, Inc., Scottsdale, AZ), a dark pupil-based video tracking system, sampled eye position at 60 hz and in ideal conditions can track with a precision and accuracy of $\sim 1^\circ$. Subjects were calibrated to a nine-point grid at the beginning of each trial and the calibration was checked at the end of each trial to allow subsequent error measurements. Video records of the eye and scene camera were saved directly to disk in a custom QuickTime digital format, which also allows the data from our sensor arrays and simulation (e.g., position of objects in the world) to be saved as synchronized metadata on each video frame.

Once a subject was wearing the HMD and calibrated for the eye tracker, the HMD displayed the driving environment. From the driver's perspective they were inside a car with a dashboard, steering wheel etc., as show in [Figure 1](#). The virtual car that subjects drove is modeled by the Vortex physics software engine by CM Labs (CMLabs, Montreal, Quebec, Canada). This software application programming interface allows the simulated forces on three-dimensional objects and the transmission of a car, including gear ratios and revolutions per minute levels for shifting gears. The virtual environment was generated via the Tile Management Tool software package created by researchers at the University of Iowa's National Advanced Driving Simulator. The path driven was a continuous four-lane road in an urban setting without traffic signs or signals or any intersections. The driving path contained many static objects including buildings, plants, cars parked on the side of the road, and pedestrians on the sidewalk. The path was approximately 2 km long and subjects drove its entirety in one direction in about two minutes at ~ 50 -72 kph (30-45 mph). To increase the

realism of the driving environment, several dynamic object models were added. Using only two car models, a truck and a sedan, eight oncoming and nine outgoing cars (i.e., cars traveling in the same direction as the subject) were added to the environment. Additionally, there was a single red sports car in the environment that acted as a leader car. The cars were spaced out roughly uniformly throughout the subject's driving path and followed unique non-intersecting paths whose trajectories were captured from the paths of a human driver. Outgoing cars were arranged in a staggered fashion in the two lanes on the subject's side of the road. These cars remained fixed in position allowing the lead car and the subject's car to weave between cars in the left and right lanes. Because the driving environment is simulated, experiments can be conducted across subjects with identical routes and car paths allowing gross level control of the visual stimuli they observe. However, since subjects are in control of the exact speed and course of the drive, each had a unique visual experience and trajectory.

Procedure

Once participants had practiced in the environment and had been calibrated with the eye tracker, all were read a set of standardized instructions where they were told to drive in a lawful manner in an urban setting, to follow a lead car at a distance of two car lengths and to maintain a constant speed of 40 mph (~ 65 kph). Our experimental conditions manipulated task priority by varying the relative emphasis of these two tasks in our instructions. In the Follow Task condition, subjects were told that following a lead car at a distance of two car lengths was most important. In the Speed Task condition, subjects were told that maintaining a constant speed was most important. They were also informed that whichever task was not emphasized, should also be done but was less important. The subjects' primary goal was always reiterated before they initiated driving.

Each subject performed all four conditions, resulting in a 2×2 within subjects design (Speed and Follow, with and without noise). The order of the Speed and Follow conditions was counterbalanced. Within each of these conditions, the order of trials with or without noise was counterbalanced. To introduce uncertainty (+Noise conditions), uniform noise was added to the car's gas pedal command. The Vortex software simulation uses a value from 0 to 1 to indicate the degree to which the gas pedal is depressed. When the subject achieved a velocity above 36 kph the software then added a randomly chosen value between 0 and 0.5 every 200 ms to the car's gas command. This resulted in driving a car that appeared to have a mechanical problem that made it

quite challenging to maintain a constant speed. Although this type of noise is dependent on the subject's current speed, when the gas pedal was held constant the speedometer representation in the car would typically vary by ± 0 to 18 kph.

Subjects were not informed about the presence or absence of noise in the gas pedal signal. While we strove to make the car simulation realistic, there is a certain amount of learning that needs to occur to adapt to the car's dynamics and wearing the HMD. To ensure that participants were familiar with these driving dynamics in the simulation, prior to the experiment starting all participants were given a practice period of 5-10 minutes to drive the entire length of the driving course at least two and up to four times depending on their level of driving comfort. During the practice session other cars were absent and subjects were not exposed to the +Noise condition. Subjects were encouraged to try hard turns, accelerations and braking to make sure they were completely familiar with how the car responded. Once subjects confirmed that they were comfortable driving the car, they performed the main experimental conditions.

Thirty-four undergraduate participants with normal or corrected to normal vision from the University of Texas at Austin gave informed consent and took part in the experiment. Ten subjects' data were not used because they could not complete all four conditions either due to motion sickness or a poor eye tracking signal. Of the remaining 24 subjects, eight additional subjects' data were not considered further in the full analysis due to infrequent checking of the speedometer in any one condition. We used a criterion to exclude subjects who made less than seven fixations on the speedometer, as each driving condition was 98 s duration on average and the overall mean number of fixations on the speedometer across all conditions was 34 (13.5 *SD*). While all subjects were able to complete the conditions, these particular subjects were possibly using a different strategy for measuring speed (e.g., optic flow) rather than directly foveating the speedometer, making their behavior difficult to compare to the other subjects who were actively using the speedometer's information. While this type of behavior may be present in real world driving, due to our experimental manipulation we exclude such subjects to ensure that they were using a common source of visual information for speed. This yielded a set of data from 16 subjects for our analysis. Of these 16, there were nine males and seven females with a mean age of 19.6 years (2.6 *SD*), with a mean of 3.8 years driving experience (2 *SD*). Additionally, to avoid the inclusion of fixation behavior at the end and beginning of each trial, when subjects often were not actively driving or weren't in position with other cars, all data files were segmented to include only the portion where vehicle velocity was

above 54 kph. Note in the data and figures below, the terms uncertainty and noise are used interchangeably to describe the presence of the uniform noise added to the subject's car velocity.

Data collection and analysis

Subjects' eye position data were analyzed using an automated, in-house developed system. The eye signal was preprocessed using a median filter and a moving average over three frames to smooth the signal. A 60×60 pixel window, $\sim 2^\circ \times 2^\circ$ of visual angle, was centered around the location of the point of gaze on each frame and each pixel in the window returned a label for the type of object it contained. Each subject's eye position was measured against a calibration screen at the beginning and end of each trial. During the experiment, a data file was generated that contained readings from the car controls (steering and pedals), positions of the subject's head and eye as well as the positions of all of the objects in the environment. By using the head orientation and position of the subject along with the positions of all the objects in the driving environment our analysis tool can create a complete reconstruction of the experimental environment. Each data frame is analyzed for the types of objects present in the local area near fixation. Given the reconstruction, we use the projection of the eye position in three-dimensional to query the pixels in a local area for the identity of the object present in that pixel, this is conducted over the entire window and used to index the object categories present. Due to a technical limitation in how we recover object identity information from our three-dimensional models, the automated data can provide object labels only for the speedometer, the leader car, oncoming cars and outgoing cars. All other fixations on the road, buildings and other scenery were entered into a category termed 'other.'

The eye position data of subjects included in the analysis had a precision of $\sim 1^\circ$ - 2° , and global offsets were applied as needed to each subject's data as needed to ensure that the point of gaze signal was correctly lined up with the calibration targets. Eye movement data were initially segmented in two ways, segmentation via object labels and using the eye velocity signal with an adaptive velocity threshold algorithm (Rothkopf & Pelz, 2004). Both approaches have limitations, object label segmentation examines the data for sets of continuous object labels near the point of gaze that last at least two frames. Label segmentation is subject to noise in the eye position signal so if the point of gaze is near the border between two objects, multiple segments may be found. Additionally, it does not reveal fixations within an object category (i.e., the label never changes) and thus can




Figure 2. Example of gaze behavior from a single driving trial from the Speed Task condition. The horizontal axis displays time in seconds and the vertical corresponds to object category. Each rectangular chunk corresponds to a portion of time where the center of the subject's left eye was within $\sim 2^\circ$ of an object class. While subjects look at several object classes the most relevant behavior to our experimental manipulations is the switching behavior exhibited between looks on the speedometer and on the leader. Looks at multiple objects can occur simultaneously due to labeling of any object that enters the 60×60 pixel window centered around the fovea.

only determine looks, the beginning and ending a fixation or set of fixations on a distinct object class. The velocity analysis does not use label knowledge but is also subject to noise in the eye gaze signal. The velocity analysis marked fixations by looking for frame sequences of at least 50 ms in length where the eye velocity is below $35^\circ/\text{s}$ (although this parameter can change on the local noise level in a local window). After these segments were found, the object labels for a given segment were counted. To allow a liberal description of the image content near the fovea, all pixel labels in 60×60 pixel window were considered in our analysis. This means that multiple looks to different objects could occur simultaneously if they resided within the pixel window. These looks could be also have asynchronous start and end times due to the simulation's dynamics, e.g., a look at the leader may start at $t = 1$ and last until $t = 10$, and another car may also enter the pixel window for a look that begins at $t = 5$ until the car exits the window at $t = 8$. To have a look labeled as being on the 'other' category, all of the pixels had to be labeled as 'other.' Due to the nature of our simulation, eye position can remain relatively stable but the content of fixation can change. For example, the velocity segmentation may report a fixation of 1.5 s in length but that fixation could actually have 1 s of labels for one object and 500 ms for another due to objects moving in and out of the fixation window. While there are many ways to deal with this segmentation and labeling problem, we ultimately chose to categorize eye movement data via the label segmentation method due to the

simplicity of its assumptions, i.e., the content in the fixation-labeling window is of primary importance, and all subsequent graphs and analyses present looks segmented via object labels. Figure 2 shows an example of the type of segmented data yielded from this analysis for a single trial where the speed following task was emphasized. Note that subjects devote more time to the lead car regardless of task (see below).

Results

Driving behavior

We first examine whether subjects' performance reflected task instruction. Although all subjects were instructed to both follow a leader and maintain a constant speed, one might expect that subjects would exhibit superior performance in the task instructed as the most important. Table 1 displays subjects' mean speed and mean standard deviation in speed, as well as their mean distance to the leader and the mean standard deviation in that distance. It was hypothesized that subjects in the Speed Task conditions should have a speed closer to 40 mph (64.4 kph) and less variation in their speed. Similarly, it was expected that subjects in the Follow Task conditions should have less variation in their distance to the leader. Subjects were told to drive at a heuristic of two car lengths from the leader

| | Speed (kph) ¹ | Speed SD (kph) ² | Distance (m) ³ | Distance SD (m) ⁴ |
|--------------|--------------------------|-----------------------------|---------------------------|------------------------------|
| Follow | 61.29 (0.22) | 6.58 (0.28) | 22.43 (1.16) | 4.03 (0.35) |
| Follow+Noise | 60.92 (0.16) | 10.44 (0.36) | 26 (1.4) | 7.13 (1.78) |
| Speed | 61.63 (0.13) | 5.87 (0.28) | 20.79 (1.19) | 4.47 (0.32) |
| Speed+Noise | 61.3 (0.16) | 9.97 (0.35) | 27.2 (1.98) | 8.52 (0.92) |

Table 1. Summary of driving performance across conditions. For each condition we present the mean performance for following, leader distance and leader distance standard deviation, and for maintaining a constant speed, speed and speed standard deviation. SEM is presented in parentheses. Two-way repeated-measures ANOVA were performed for all measures including: (A) Mean leader distance, which showed no main effect of task on $F(1,15) = 0.02$, $p = 0.89$, a main effect of noise, $F(1,15) = 21.7$, $p = 3e-4$ and a marginally significant task and noise interaction, $F(1,15) = 3.44$, $p = 0.083$. (B) Mean standard deviation of leader distance which had no main effect of task, $F(1,15) = 2.69$, $p = 0.12$, a main effect of noise, $F(1,15) = 31.3$, $p = 5.08e-5$, and no significant interaction, $F(1,15) = 1.12$, $p = 0.31$. (C) Mean car velocity which had a main effect of task, $F(1,15) = 6.5$, $p = 0.02$, a main effect of noise, $F(1,15) = 5.4$, $p = 0.03$, and no significant interaction, $F(1,15) = 0.01$, $p = 0.94$. (D) Mean standard deviation in speed with a marginal main effect of task, $F(1,15) = 3.96$, $p = 0.065$, a significant main effect of noise, $F(1,15) = 138$, $p = 5.8e-9$ and no significant interaction, $F(1,15) = 0.24$, $p = 0.63$.

(~15 m as measured from the center of the lead car to the center of the subject's car, assuming cars of 5 m in length), and we cannot address how close subjects were to a correct distance.

Overall, subjects' driving behavior is roughly equivalent across our task manipulations. There was a significant main effect of task in mean speed and a marginally significant difference in speed variability. However, the effect sizes were rather small, 0.36 kph for mean speed and 0.59 kph for speed variability. Much larger effects were observed due to noise. While subjects have very similar driving performance across tasks, our eye movement analyses detailed below reveal more substantial differences.

Fixation proportion

There are several metrics available to capture how subject gaze behavior varies in each condition. In this section and all following sections, a two-way repeated-measures ANOVA was used to evaluate gaze behavior. One basic metric is the proportion of time spent looking at the various object categories in the driving world as shown in Figure 3.

This figure shows that much of the subjects' time driving is spent looking at the leader car and other cars. In addition there is a clear effect of task on the allocation of gaze between the Speedometer, Leader Car, and Oncoming Cars. Note that due to the method of look labeling, this downplays the amount of information coming from the road. In nearly all fixations, except those on the speedometer and some rare cases (e.g., looking at the sky), the eye is in a position to gather information about the road. However, because we cannot explicitly label the road in our analysis and do not have any specific hypotheses concerning looks to the road or to a specific part of the road, these types of fixations will not be addressed.

Most important for our manipulation are the differences between looks to the leader car and the

speedometer. Figure 4a shows the mean gaze behavior on the leader across conditions. Note that for this and all subsequent analyses, the graphs display between-subjects' standard errors, but all analyses were within subjects. A main effect of task was found where the proportion of leader looks was reduced from 84% to 63% by the instruction to emphasize speed control, $F(1,15) = 8.4$, $p = 1e-6$. Noise reduced the proportion of looks on the leader from 76% to 71%, $F(1,15) = 41.6$, $p = 0.01$. Additionally, there was a significant task and noise interaction, $F(1,15) = 10.33$, $p = 0.006$. This interaction was examined with paired t-tests and




Figure 3. Mean percentage of looks across all object categories. Look proportions are calculated for each category and controlled for amount of time present onscreen and then averaged across subjects. Note that proportions in the figure do not sum to one since looks can contain multiple labels across object classes. Error bars show ± 1 SEM.

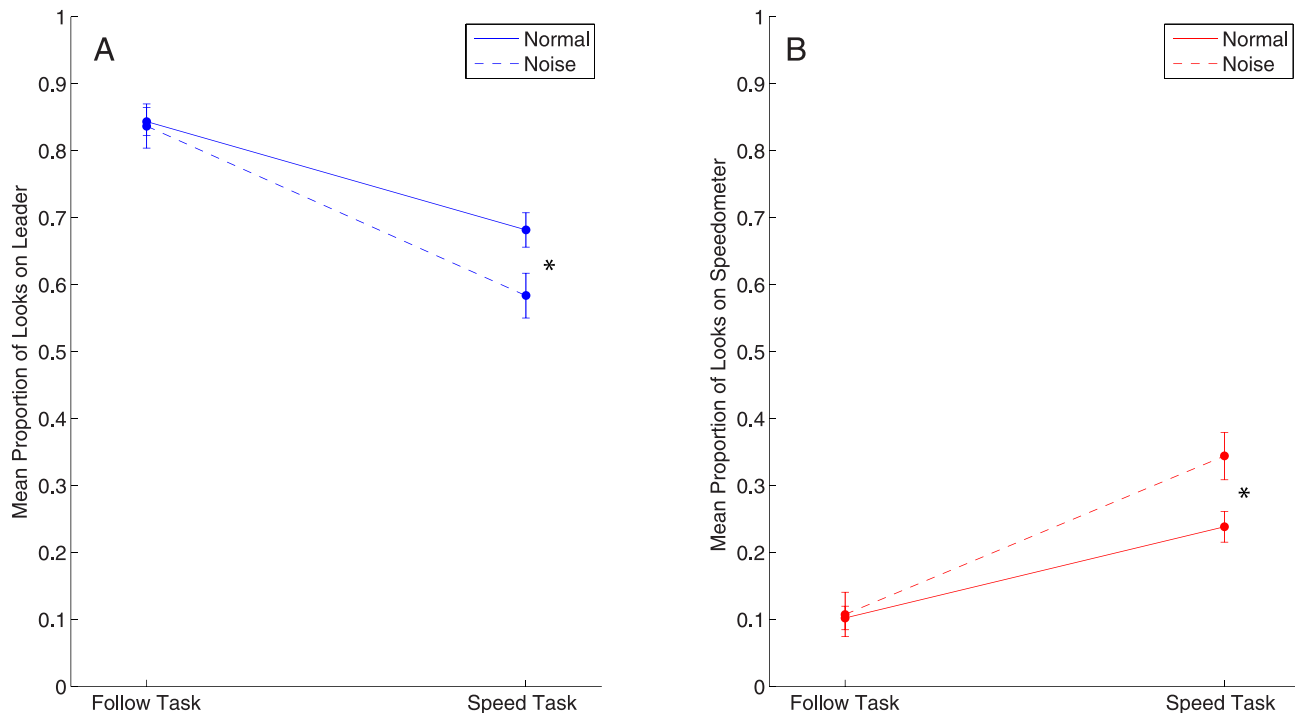


Figure 4. Mean look proportions to the leader car and speedometer in Follow and Speed Tasks. (A) Proportion of looks to the lead car in both tasks, with and without noise added. (B) Proportion of looks to the speedometer in both tasks, with and without noise added. Dashed lines show the Noise conditions. Note that a repeated-measures ANOVA was used for statistical analysis but between subjects data is plotted for ease of visualization. The asterisk indicates a statistically significant difference ($p < 0.05$ via paired t -test) between normal and noise conditions. Error bars show ± 1 SEM between subjects.

showed a significant effect in the Speed Task where noise decreased leader fixation proportion from 68% to 58%, $t(15) = 3.59$, $p = 0.0027$, but no such effect was present in the Follow Task conditions, $t(15) = 0.38$, $p = 0.71$. On average, subjects decreased look proportions on the leader by $\sim 10\%$ when the speed condition was emphasized and noise was present. This result suggests that subjects are using uncertainty information but only in the condition where the speedometer's task has high priority.

Figure 4b summarizes gaze proportion data for looks to the speedometer. There was main effect of task that increased the mean proportion of speedometer looks from 11% to 29% in the Speed Task conditions, $F(1,15) = 37.3$, $p = 2e-5$. Speedometer fixations also increased when noise was present, $F(1,15) = 13.3$, $p = 0.002$. There was also a significant noise and task interaction, $F(1,15) = 16.18$, $p = 0.001$. This derived from the Speed Task conditions, where noise increased looks from 24% to 34%, $t(15) = 4.8$, $p = 2.4e-4$, but no effect was present in the Follow Task conditions, $t(15) = 0.32$, $p = 0.75$. On average, subjects increased look proportions on the speedometer by $\sim 10\%$ in the Speed Task when noise was present. This suggests that when subjects are in the Speed Task+Noise condition, the extra proportion of look durations they devote towards the speedometer is subtracted from the leader task.

Look proportions on the speedometer and leader were analyzed for order effects. ANOVAs were conducted comparing fixation distributions per condition across groups of subjects who drove in the same condition order. Marginal effects of condition order were found on speedometer look proportions in the Follow Task, $F(3,12) = 2.97$, $p = 0.08$, and on leader look proportions also in the same task, $F(3,12) = 2.94$, $p = 0.08$. Further inspection found that for subjects who performed trials in the order Follow+Noise, Follow, Speed+Noise, Speed, there was a trend to look more frequently at the speedometer and less so at the leader Follow Task than other trial orderings. It is not entirely clear what this trend means and how order effects might contribute to the data we collected. However, due to the marginal significance of the ANOVA results, we proceeded with our analysis assuming that ordering effects were not prominent.

Look durations

Look durations on the leader and speedometer were also examined for the influence of task and noise. Figure 5a summarizes this data for looks to the leader. There was a main effect of task, where subjects decreased the mean duration of leader looks from

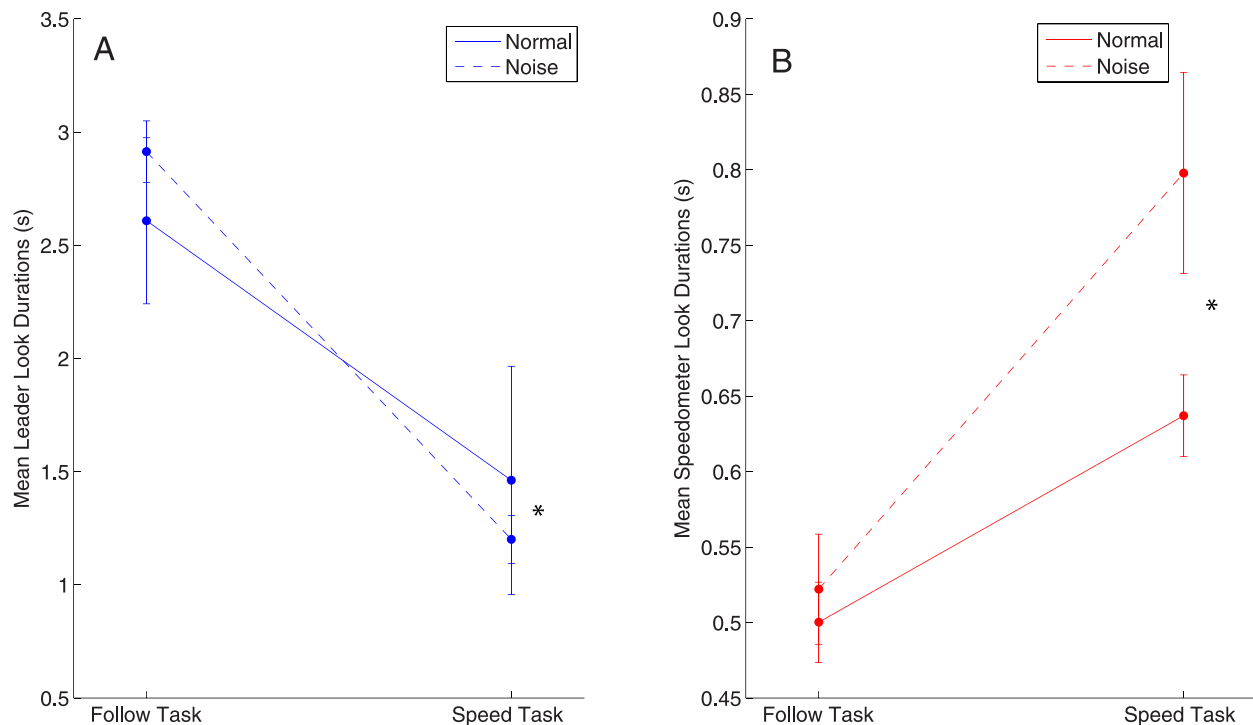


Figure 5. Mean look durations on the lead car and speedometer. (A) Average duration of looks to the lead car in both tasks, with and without noise added. (B) Average duration of looks to the speedometer in both tasks, with and without noise added. Dashed lines show the Noise conditions. Note that a repeated-measures ANOVA was used for statistical analysis but between subjects data is plotted for ease of visualization. The asterisk indicates a statistically significant difference ($p < 0.05$ via paired t -test) between normal and noise conditions. Error bars show ± 1 SEM between subjects.

2.76 s in the Follow Tasks to 1.33 s in the Speed Tasks, $F(1,15) = 11.52$, $p = 0.004$. There was no main effect of noise, $F(1,15) = 0.027$, $p = 0.87$. However, there was a significant interaction between task and noise presence, $F(1,15) = 5.32$, $p = 0.036$. Paired t -tests were used to examine the interaction. Compared to the Speed Task condition, on average subjects decreased look durations on the leader by 260 ms in the Speed+Noise Task, $t(15) = 2.14$, $p = 0.049$. In the Follow Tasks, there was no reliable effect, $t(15) = 1.37$, $p = 0.19$.

Figure 5b summarizes data for look durations on the speedometer. There was a main effect of task such that in the Speed Task conditions, look durations on the Speedometer increased from 0.51 s to 0.72 s over the Follow Task conditions, $F(1,15) = 19.7$, $p = 0.0004$. There was also main effect of noise increasing speedometer look durations from 0.57 s to 0.66 s when noise was present, $F(1,15) = 9.27$, $p = 0.008$. Additionally, there was a significant interaction between Noise and Task variables, $F(1,15) = 6.8$, $p = 0.02$. This interaction was examined with paired t -tests, and showed that speedometer looks in the Speed+Noise Task increased by an average of 160 ms compared to the Speed Task, $t(15) = 3.56$, $p = 0.003$, but no such effect was present in either Follow Task conditions, $t(15) = 0.65$, $p = 0.53$.

Interlook interval duration

The duration between successive looks on an object gives a metric for polling frequency, i.e., if a subject looks away from an object, how long does the subject wait on average before getting new sensory information about that same object with a new look? Interlook intervals were calculated for the leader and the speedometer and examined for effects of task and noise. Figure 6a summarizes this data for interlook intervals for the leader. There was a main effect of task where the interlook interval on the Leader increased from 0.46 s in the Follow Task conditions to 0.76 s in the Speed Task conditions, $F(1,15) = 33.12$, $p = 0.0001$. There was also a main effect of noise, where the presence of noise increased the interlook interval on the leader from 0.55 s to 0.67 s, $F(1,15) = 25.4$, $p = 0.0001$. There was also a significant interaction between these variables, $F(1,15) = 13.46$, $p = 0.002$. This resulted from the effect of noise in the Speed Task, where it increased interlook interval on the leader (0.66 s vs. 0.86 s; $t(15) = 5.11$, $p = 1.3e-4$), but no effect was present in either Follow Task condition, $t(15) = 1.53$, $p = 0.15$. On average, subjects increased look interlook interval durations on the leader by ~ 200 ms in the Speed Task when noise was present. This presumably reflects the




Figure 6. Mean interlook intervals for the leader car and speedometer. (A) Average interlook intervals for looks to the lead car in both tasks, with and without noise added. (B) Average interlook interval duration for looks to the speedometer in both tasks, with and without noise added. Dashed lines show the Noise conditions. Note that a repeated measures ANOVA was used for statistical analysis but between subjects data is plotted for ease of visualization. The asterisk indicates a statistically significant difference ($p < 0.05$ via paired t -test) between normal and noise conditions. Error bars show ± 1 SEM between subjects.

increasing demands of the Speed Task+Noise condition.

Figure 6b summarizes the interlook interval data for the speedometer. In the Speed Tasks mean interlook interval was much shorter than in the Follow Tasks, 1.92s versus 4.48s; $F(1,15) = 16.57$, $p = 0.001$. The effect of Noise was not significant, $F(1,15) = 0.82$, $p = 0.37$, nor was the interaction significant, $F(1,15) = 2.5$, $p = 0.13$. However, planned paired-sample t -tests revealed an effect of noise for the Speed Task conditions, 2.3 s versus 1.6 s; $t(15) = 3.56$, $p = 0.003$ and no effect of noise in the Follow Task conditions, $t(15) = 0.35$, $p = 0.73$. Thus on average, subjects decreased look interlook intervals on the speedometer by 0.72 s when Noise was present in the Speed Task.

Discussion

The goal of this research was to gain insight into how human eye movements actively gather information from a dynamic world. We investigated whether gaze deployment is controlled by both task priority and uncertainty in the face of competing potential targets, as in the Sprague and Ballard model, assuming that task priority can be interpreted as implicit reward. As

expected from previous work on the importance of task goals, we found that task emphasis was a primary factor in the allocation of gaze to the task-relevant location. The task that was prioritized led to increased probability of looking at task relevant targets, increased look durations, and reduced interlook intervals. We showed in addition, that the effect of task is modulated by sensory uncertainty. Subjects made more frequent and longer looks to the speedometer when speed noise was added. However, this effect of uncertainty was manifest only when the Speed task was emphasized as the task with high priority. This suggests that if uncertainty primarily affects one visually guided task, this uncertainty will be reduced via changes in gaze pattern only if the task has a high priority and uncertainty alone is not sufficient. It is worth noting that despite the reliable differences found in eye movements across conditions, there was surprisingly little difference in driving behavior due to task instruction. There are several reasons this may be the case. Subjects who had very infrequent looks to the speedometer were left out of our analysis that would make our analysis sample more homogenous. Additionally, our driving conditions were not designed to test the limits of subject performance. Our subjects were experienced drivers and it is possible that the tasks were not sufficiently challenging to depend critically on

the visual information afforded by the fixations. It appears, at least, that the relationship between such information and driving performance is more complex than we supposed.

While uncertainty and reward (both explicit reward and task priority as a proxy for implicit reward) have each been topics of recent research in human gaze behavior, the present study manipulated both these factors and showed that they jointly affect allocation of gaze between competing targets in the context of natural behavior. There have been other demonstrations of comparable task-based reward effects in natural behavior (Jovancevic-Misic & Hayhoe, 2009; Rothkopf, Ballard, & Hayhoe, 2007), but there has been no previous demonstration of the role of uncertainty in regulating gaze in the natural world. Our primary result showed that when confronted with a signal with perceptual uncertainty, gaze allocation changes only if this signal is associated with a task that has a sufficient amount of reward or behavioral relevance. This finding is difficult to reconcile with bottom up models of visual attention. Such models are not designed with task-based vision in mind and must be modified to incorporate ideas of reward or task relevance. One might argue that uncertainty in our experiment yields a more salient visual stimulus and thus drives visual attention. However, this does not explain why we observed an interaction where uncertainty only changed gaze behavior when associated with sufficient task priority. This result is generally consistent with the Sprague and Ballard modeling framework, although it is not clear if this exact interaction would be predicted, and this is a topic of our current research. This model is distinctive among other current models of visual attention in that it explicitly models the visual perception and action loop with the premise that an agent or organism in a dynamic world needs to allocate visual attention in a rational way that actively reduces uncertainty based upon reward driven task priorities. While it is important to understand early bottom-up processing of visual information, it is simplistic to approach the deployment of visual attention as a process of stimuli attracting attention. It is prudent to consider how visual information is selected for control purposes in naturalistic scenarios. Bottom-up models can be biased by knowledge of human gaze behavior to exhibit more ‘top-down like’ behavior towards particular visual features, but this sidesteps the point of why certain objects are useful to look at in the first place. Within the context of our experiment, one could model monitoring the speedometer and the leader car as two separate visual tasks. While the exact visual computations being carried out are speculative, following could be accomplished using the angle subtended by the leader car as a control signal (Andersen & Sauer, 2007).

Speed estimation has many potential cues, but in our manipulation it is plausible that subjects integrated the orientation of the speedometer needle over time to estimate its position. The introduction of uncertainty in speed could hinder this integration process and attempts to store the position in memory, explaining why subjects look longer and more frequently at the speedometer. Because the introduction of noise made performance worse across our task manipulations, it is not obvious why our results didn’t show a similar effect for fixations on the leader. It may be the case that since the speedometer immediately reflected multiple changes per second, uncertainty information was more readily available, whereas the fluctuation in distance to the leader involved gradual distance changes due to simulation’s physics model and distance changes due to variations in the leader car’s speed.

Within the study of behavior and eye movements in driving several control models have been suggested for particular behaviors, e.g., car following, lane following (Andersen & Sauer, 2007; Land & Horwood, 1995; Land & Lee, 1994; Salvucci & Gray, 2004). Salvucci and Taatgen (2008, 2011) have also presented a ‘multithreaded theory of cognition’ that is conceptually quite similar to the Sprague and Ballard scheduling model. Salvucci studied attentional allocation in driving using an implementation of his theory using the ACT-R architecture (Anderson, 1996). Recently, ACT-R has been modified to more flexibly incorporate reinforcement learning architectures and additionally the system uses psychological plausible parameters for memory to mimic internal uncertainty and recent simulations indicate this may be another viable way to model such reward and uncertainty interactions in visual attention (Anderson, 2007; Janssen, Brumby, & Garnett, 2012). Models that treat the deployment of gaze from a task-based level as a set of sensory-motor interactions in a dynamic environment have considerably more traction in explaining natural fixation behavior than other approaches. However, as mentioned in our [Introduction](#), a core problem is in building an adaptive task-based framework that captures such behavior. Additionally, the unification of bottom-up and top-down control within such a framework has not been addressed.

Our experiments relied on implicit manipulations of task related reward and while we obtained large task effects, such manipulations make quantitative analysis difficult compared to experiments using explicit reward. One potential strategy to mitigate this is to estimate the intrinsic reward weights used by individual subjects using a technique called inverse reinforcement learning which, under certain assumptions, can recover the implicit reward weights of a human actor given the set of behaviors they execute over time (Ng & Russell, 2000; Rothkopf & Ballard, [submitted](#); Rothkopf &

Dimitrakakis, 2011). This is a focus of our current research.

Concerning manipulations of uncertainty, it is useful to consider Senders' work, where he captured data with parametric variations in uncertainty in driving 'black-out' experiments. He found a systematic relationship between uncertainty and the ability of a driver to keep a car at a given velocity. This type of approach, although not always easy experimentally, would be helpful in illuminating subtleties in behavior beyond the simple assertion that human eye movements strategies can be shown to change under uncertainty. Additionally, there is enormous variety in types of uncertainty. A number of features in our experiments could be manipulated, e.g., reducing contrast by adding fog, rain, or a nighttime situation, or making other cars erratic instead of the subjects. Additionally, the statistics of such uncertainty distributions, in particular the statistics of task-related uncertainty, could be rigorously categorized instead of merely using one particular type of noise as in our experiments.

It is useful to bear in mind that Sprague and Ballard's perceptual arbitration algorithm was devised as a rational approach to reduce uncertainty within a reward based framework. It was not modeled directly on human behavior, nor has it been shown to be an optimal solution for this type of problem. In our experiments, we noted that the human drivers in the Follow Task did not make any extra looks on the speedometer when noise was present compared to when it was absent. Current work is under way to test if the model can capture the uncertainty and reward interaction we observed.

Conclusion

Our study has demonstrated a novel tradeoff between task priority and uncertainty in a naturalistic driving task. If task priority is interpreted as an implicit manipulation of reward structure, our data provide evidence for a reward and uncertainty based weighting scheme that could be incorporated into top-down models of vision. The active polling of task-related information in the world with sensory systems can be prioritized by a mechanism that only reduces uncertainty on visual features if they are associated with a sufficiently high priority task.

Acknowledgments

Thanks to Dmitry Kit and John Stone for their work on the automated gaze analysis software. This work

was supported by NIH Grants EY05729 and EY019174. Constantin A. Rothkopf was supported by "Bernstein Fokus: Neurotechnologie Frankfurt, FKZ 01GQ0840" and the EU-Project IM-CLeVeR, FP7-ICT-IP-231722.

Commercial relationships: none.

Corresponding author: Brian T. Sullivan.

Email: brians@ski.org.

Address: Smith-Kettlewell Eye Research Institute, San Francisco, CA, USA and Department of Psychology, University of Texas at Austin, Austin, Texas, USA.

References

- Aggarwal, M., Hyland, B. I., & Wickens, J. R. (2012). Neural control of dopamine neurotransmission: Implications for reinforcement learning. *European Journal of Neuroscience*, 35(7), 1115–1123.
- Andersen, G. J., & Sauer, C. W. (2007). Optical information for car following: The driving by visual angle (DVA) model. *Human factors: The journal of the human factors and ergonomics society*, 49(5), 878–896.
- Anderson, J. R. (1996). ACT: A simple theory of complex cognition. *American Psychologist*, 51(4), 355.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe? (Vol. 3)*. New York: Oxford University Press.
- Atkins, J. E., Jacobs, R. A., & Knill, D. C. (2003). Experience-dependent visual cue recalibration based on discrepancies between visual and haptic percepts. *Vision Research*, 43(25), 2603–2613.
- Bromberg-Martin, E. S., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron*, 63(1), 119.
- Bruce, N. D., & Tsotsos, J. K. (2009). Saliency, attention, and visual search: An information theoretic approach. *Journal of Vision*, 9(3), 5, 1–24, <http://www.journalofvision.org/content/9/3/5>, doi:10.1167/9.3.5. [PubMed] [Article]
- Deaner, R. O., Khera, A. V., & Platt, M. L. (2005). Monkeys pay per view: Adaptive valuation of social images by rhesus macaques. *Current Biology*, 15(6), 543–548.
- Dorris, M. C., & Glimcher, P. W. (2004). Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron*, 44(2), 365–378.

- Droll, J. A., Hayhoe, M. M., Triesch, J., & Sullivan, B. T. (2005). Task demands control acquisition and storage of visual information. *Journal of Experimental Psychology: Human Perception and Performance*, 31(6), 1416.
- Fecteau, J. H. (2007). Priming of pop-out depends upon the current goals of observers. *Journal of Vision*, 7(6):1, 1–11, <http://www.journalofvision.org/content/7/6/1>, doi:10.1167/7.6.1. [PubMed] [Article]
- Folk, C. L., Remington, R. W., & Wright, J. H. (1994). The structure of attentional control: Contingent attentional capture by apparent motion, abrupt onset, and color. *Journal of Experimental Psychology: Human Perception and Performance*, 20(2), 317.
- Folk, C. L., & Remington, R. (1998). Selectivity in distraction by irrelevant featural singletons: Evidence for two forms of attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, 24(3), 847.
- Forster, S., & Lavie, N. (2008). Failures to ignore entirely irrelevant distractors: The role of load. *Journal of Experimental Psychology: Applied*, 14(1), 73.
- Glimcher, P. W. (2003). The neurobiology of visual-saccadic decision making. *Annual Review of Neuroscience*, 26(1), 133–179.
- Gottlieb, J., & Balan, P. (2010). Attention as a decision in information space. *Trends in Cognitive Sciences*, 14(6), 240.
- Graf, E. W., Warren, P. A., & Maloney, L. T. (2005). Explicit estimation of visual uncertainty in human motion processing. *Vision Research*, 45(24), 3050–3059.
- Hayhoe, M., & Ballard, D. (2005). Eye movements in natural behavior. *Trends in Cognitive Sciences*, 9(4): 188–194.
- Hayhoe, M. M., Bensinger, D. G., & Ballard, D. H. (1998). Task constraints in visual working memory. *Vision Research*, 38(1), 125–137.
- Hikosaka, O., Nakamura, K., & Nakahara, H. (2006). Basal ganglia orient eyes to reward. *Journal of Neurophysiology*, 95(2), 567–584.
- Itti, L., & Baldi, P. (2006). Bayesian surprise attracts human attention. *Advances in neural information processing systems*, 18, 547.
- Itti, L., & Koch, C. (2001). Computational modeling of visual attention. *Nature Reviews Neuroscience*, 2(3), 194–203.
- Janssen, C. P., Brumby, D. P., & Garnett, R. (2012). Natural break points: The influence of priorities and cognitive and motor cues on dual-task interleaving. *Journal of Cognitive Engineering and Decision Making*, 6(1), 5–29.
- Jovancevic, J., Sullivan, B., & Hayhoe, M. (2006). Control of attention and gaze in complex environments. *Journal of Vision*, 6(12):9, 1431–1450, <http://www.journalofvision.org/content/6/12/9>, doi:10.1167/6.12.9. [PubMed] [Article]
- Jovancevic-Misic, J., & Hayhoe, M. (2009). Adaptive gaze control in natural environments. *The Journal of Neuroscience*, 29(19), 6234–6238.
- Kanan, C., Tong, M. H., Zhang, L., & Cottrell, G. W. (2009). SUN: Top-down saliency using natural statistics. *Visual Cognition*, 17(6-7), 979–1003.
- Knudsen, E. I. (2007). Fundamental components of attention. *Annual Review of Neuroscience*, 30, 57–78.
- Land, M. F., & Hayhoe, M. (2001). In what ways do eye movements contribute to everyday activities? *Vision research*, 41(25), 3559–3565.
- Land, M. F., & Horwood, J. (1995). Which parts of the road guide steering? *Nature*, 377, 339–340.
- Land, M. F., & Lee, D. N. (1994). Where we look when we steer. *Nature*, 369(6483), 742–744.
- Land, M., Mennie, N., & Rusted, J. (1999). The roles of vision and eye movements in the control of activities of daily living. *Perception*, 28(11), 1311–1328.
- Lauwereyns, J., Takikawa, Y., Kawagoe, R., Kobayashi, S., Koizumi, M., & Coe, B., et al. (2002). Feature-based anticipation of cues that predict reward in monkey caudate nucleus. *Neuron*, 33(3), 463–473.
- Leber, A. B., & Egeth, H. E. (2006). It's under control: Top-down search strategies can override attentional capture. *Psychonomic Bulletin & Review*, 13(1), 132–138.
- Lee, D., Seo, H., & Jung, M. W. (2012). Neural basis of reinforcement learning and decision making. *Annual Review of Neuroscience*, 35, 287–308.
- Lu, S., & Han, S. (2009). Attentional capture is contingent on the interaction between task demand and stimulus saliency. *Attention, Perception, & Psychophysics*, 71(5), 1015–1026.
- Maunsell, J. H. R. (2004). Neuronal representations of cognitive state: Reward or attention? *Trends in Cognitive Science*, 8, 261–265.
- Morvan, C., & Maloney, L. T. (2012). Human visual search does not maximize the post-saccadic probability of identifying targets. *PLoS Computational Biology*, 8(2), e1002342.

- Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, *434*(7031), 387–391.
- Navalpakkam, V., Koch, C., & Perona, P. (2009). Homo economicus in visual search. *Journal of Vision*, *9*(1):31, 1–16, <http://www.journalofvision.org/content/9/1/31>, doi:10.1167/9.1.31. [PubMed] [Article]
- Navalpakkam, V., Koch, C., Rangel, A., & Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences*, *107*(11), 5232–5237.
- Ng, A. Y., & Russell, S. (2000). Algorithms for inverse reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*, (pp. 663–670).
- Nunez-Varela, J., Ravindran, B., & Wyatt, J. L. (2012, May). Where do I look now? Gaze allocation during visually guided manipulation. In *Robotics and Automation (ICRA), 2012, IEEE International Conference*, pp. 4444–4449. IEEE.
- Peck, C. J., Jangraw, D. C., Suzuki, M., Efem, R., & Gottlieb, J. (2009). Reward modulates attention independently of action value in posterior parietal cortex. *The Journal of Neuroscience*, *29*(36), 11 182–11 191.
- Pelz, J., Hayhoe, M., & Loeber, R. (2001). The coordination of eye, head, and hand movements in a natural task. *Experimental Brain Research*, *139*(3), 266–277.
- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*(6741), 233–238.
- Renninger, L. W., Verghese, P., & Coughlan, J. (2007). Where to look next? Eye movements reduce local uncertainty. *Journal of Vision*, *7*(3):6, 1–17, <http://www.journalofvision.org/content/7/3/6>, doi:10.1167/7.3.6. [PubMed] [Article]
- Rothkopf, C. A., & Ballard, D. H. (Submitted). Modular inverse reinforcement learning for visuomotor behavior.
- Rothkopf, C., & Dimitrakakis, C. (2011). Preference elicitation and inverse reinforcement learning. *Machine learning and knowledge discovery in databases, Lecture Notes in Computer Science*, Vol. 6913, pp. 34–48. European Conference, ECML PKDD 2011, Athens, Greece, September 5–9, 2011 Proceedings, Part III. Berlin: Springer.
- Rothkopf, C. A., & Ballard, D. H. (2010). Credit assignment in multiple goal embodied visuomotor behavior. *Frontiers in Psychology*, *1*(173), 1–13, doi:10.3389/fpsyg.2010.00173.
- Rothkopf, C. A., Ballard, D. H., & Hayhoe, M. H. (2007). Task and scene context determines where you look. *Journal of Vision*, *7*(14):16, 1–20, <http://www.journalofvision.org/content/7/14/16>, doi:10.1167/7.14.16. [PubMed] [Article]
- Rothkopf, C. A., & Pelz, J. B. (2004). Head movement estimation for wearable eye tracker. In A. T. Duchowski & R. Vertegaal (eds.), *Proceedings of the 2004 Symposium on Eye Tracking Research & Applications*. (pp. 123–130). ACM.
- Salvucci, D. D., & Gray, R. (2004). A two-point visual control model of steering. *Perception-London*, *33*(10), 1233–1248.
- Salvucci, D. D., & Taatgen, N. A. (2008). Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review*, *115*(1), 101.
- Salvucci, D. D., & Taatgen, N. A. (2011). Toward a unified view of cognitive control. *Topics in Cognitive Science*, *3*(2), 227–230.
- Schlicht, E. J., & Schrater, P. R. (2007). Effects of visual uncertainty on grasping movements. *Experimental Brain Research*, *182*(1), 47–57.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*(1), 1–27.
- Schutz, A., Trommershauser, J., & Gegenfurtner, K. (2012). Dynamic integration of information about salience and value for saccadic eye movements. *Proceedings of the National Academy of Sciences*, *109*, 7547–7552.
- Senders, J. W. (1980). *Visual scanning processes*. Soest, the Netherlands: Drukkerij Neo Print.
- Seo, H., & Less, D. (2007). Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *Journal of Neuroscience*, *27*, 8366–8377.
- Sohn, J.-W., & Lee, D. (2006). Effects of reward expectancy on sequential eye movements in monkeys. *Neural Networks*, *19*, 1181–1191.
- Sprague, N., & Ballard, D. (2003). Eye movements for reward maximization. *Advances in Neural Information Processing Systems*, *16*, 1467.
- Sprague, N., Ballard, D., & Robinson, A. (2007). Modeling embodied visual behaviors. *ACM Transactions on Applied Perception (TAP)*, *4*(2), 11.
- Stritzke, M., & Trommershäuser, J. (2007). Eye movements during rapid pointing under risk. *Vision research*, *47*(15), 2000–2009.
- Stritzke, M., Trommershäuser, J., & Gegenfurtner, K.R. (2009). Effects of salience and reward information during saccadic decisions under risk.

- Journal of the Optical Society of America, 26, B1–B13.
- Stuphorn, V., Taylor, T. L., & Schall, J. D. (2000). Performance monitoring by the supplementary eye field. *Nature*, 408(6814), 857–860.
- Stuphorn, V. & Schall, J. D. (2006). Executive control of conterminating saccades by the supplementary eye field. *Nature Neuroscience*, 9(7), 925–931.
- Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, 304(5678), 1782–1787.
- Sullivan, B. T., Johnson, L. M., Ballard, D. H., & Hayhoe, M. M. (2011). A modular reinforcement learning model for human visumoto behavior in a driving task. *Proceedings of the AISB 2011 Symposium*, 33–40.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1), 9–44.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1, No. 1). Cambridge, MA: MIT Press.
- Tatler, B. W., Hayhoe, M. M., Land, M. F., & Ballard, D. H. (2011). Eye guidance in natural vision: Reinterpreting salience. *Journal of Vision*, 11(5):5, 1–23, <http://www.journalofvision.org/content/11/5/5>, doi:10.1167/11.5.5. [PubMed] [Article]
- Theeuwes, J. (2004). Top-down search strategies cannot override attentional capture. *Psychonomic Bulletin & Review*, 11(1), 65–70.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136.
- Warren, P. A., Graf, E. W., Champion, R. A., & Maloney, L. T. (2012). Visual extrapolation under risk: human observers estimate and compensate for exogenous uncertainty. *Proceedings of the Royal Society B: Biological Sciences*, 279(1736), 2171–2179.
- Watanabe, K., Lauwereyns, J., & Hikosaka, O. (2003). Neural correlates of rewarded and unrewarded eye movements in the primate caudate nucleus. *The Journal of Neuroscience*, 23(31), 10052–10057.
- 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.
- Wolfe, J. M. , & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, 5(6), 495–501
- Yantis, S., & Egeth, H. E. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, 25(3), 661.
- Zhang, H., Morvan, C., & Maloney, L. T. (2010). Gambling in the visual periphery: A conjoint-measurement analysis of human ability to judge visual uncertainty. *PLoS Computational Biology*, 6(12), e1001023.