Task relevance moderates saccade velocities to spatially separated cues

Caitlyn M. McColeman (caitlyn_mccoleman@sfu.ca)
Department of Psychology, Simon Fraser University
8888 University Drive, Burnaby, BC V5A 1S6

Mark R. Blair (mark_blair@sfu.ca)
Cognitive Science Program & Department of Psychology, Simon Fraser University
8888 University Drive, Burnaby, BC V5A 1S6

Abstract

The study of eye movements has enjoyed a history of supporting theories of attention in different task settings by expanding our understanding of how people navigate tasks such as natural scene perception, reading and categorization. The theories and models that these data inform, however, are largely based on fixation patterns. Presently lacking is an understanding of how the eye movements preceding these fixations are affected by the task environment and if they change as a function of a shift in the state of knowledge. In an effort to close this gap, we report changing saccade velocities in two category learning experiments, evidencing the importance of understanding saccades in developing a stronger theory of the deployment of visual attention as it is influenced by higher level cognitive changes.

Keywords: Saccades; eye movements; category learning; visual attention; learning

The visual world contains an enormous amount of information. From when we wake up to check the time, to the familiar walk to the car, we make a remarkable number of eye movements to different elements of the environment in order to extract the information that is relevant to our goals at the time. However, we also make a large number of eye movements to elements of the environment that are irrelevant to our goals. For instance, the tree rustling outside the kitchen window is unrelated preparing a coffee, but we still opt to look there from time-to-time during the morning ritual. These eye movements, saccades, are punctuated by fixations. During fixations, the eyes pause while extracting information from the environment for further processing.

The saccade itself is not an effective source of information gathering, since visual perception is suppressed while the eye is in motion (Matin, 1974). However, the fixation is unable to gather information from another part of the environment without the preceding saccade to bring the eye to the target location. Since so many saccades are made each minute, small differences in the speed or accuracy of the movement can add up to important cumulative differences in completing a task or perceiving various parts of the visual environment. It is both the journey and the destination of the eye, on the scale of milliseconds, that indicate the processing underlying the deployment of visual attention.

The subtle properties of saccades have been explored in oculomotor learning tasks, wherein participants are trained to saccade toward a target. Work in non-human primates has shown that saccades quicken when they are deployed to a rewarded target location relative to alternative unrewarded locations (Takikawa, Kawagoe, Itoh, Nakahara & Hikosaka, 2002), implicating an important role of learning and reinforcement in programming and executing saccades.

In an anti-saccade task, when the participants' goal is to make a saccade to the mirror location of an onscreen cue, saccades tend to be longer and are more likely to miss the target location than when the task is simply to make a saccade to a cue (Hallett, 1977; Walker, Walker, Husain & Kennard, 2002). Recent work shows eye movements in the anti-saccade task can be further affected by factors such as drowsiness (Ahlstrom, Nyström *et al.* 2013), age and executive function (Mirsky, Heuer, Jafari, Kramer, Schenk, Viskontas, Miller, & Boxer, 2011), alcohol impairment (Roche & King, 2010) and anxiety (Cornwell, Mueller, Kaplan, Grillon, & Ernst, 2012). Through the variety of sampled participants, it is consistently found that executing a purely goal-driven saccade in the absence of a clear visual cue is more demanding, as is exhibited through slower eye movements. This finding also invites the possibility that the existence of a larger cognitive load - having to identify a cue and calculate a mirror target location to send the eyes to - enacts costs to the oculomotor system that is observed in sacrifices of saccade speed and accuracy relative to simply moving the eyes to an onscreen target. If a cognitively demanding decision does influence saccade velocities in a simple task, a natural next step is to explore how cognitive load may be reflected in saccadic properties in more challenging tasks.

Fixations have been explored more often than saccades in complicated task settings and are known to vary both as a function of the task environment and of participants' knowledge. Category learning tasks provide insight into the interplay between developing expertise in making category judgements and the corresponding trial-to-trial oculomotor activity. For instance, as knowledge of a category structure develops, participants' fixations to task relevant cues are longer than fixations to irrelevant cues (Blair, Watson, Walshe & Maj, 2009; Chen, Meier, Blair, Watson & Wood, 2013), and there are more fixations to relevant than irrelevant cues (McColeman, Barnes, Chen, Meier, Walshe & Blair, 2014; Rehder & Hoffman, 2005), demonstrating

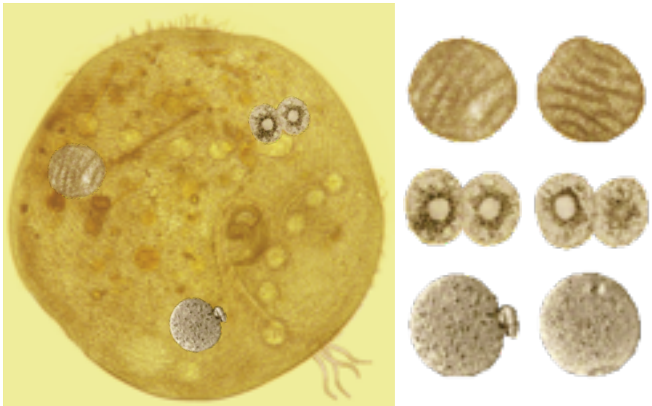

Figure 1. The stimuli used in the experiments. The features (right) could take on two possible values. They were pasted on the stimulus background (left).

flexibility in the mechanisms underlying fixations and responsiveness to higher level cognitive changes.

Understanding the influence of learning on fixations is important in understanding the allocation of attention in the context of problem solving, but there remains a question of what comes before those fixations. Saccades are an integral part of visual attention, and so in the following study, we explore two category learning tasks of varying complexity to uncover the influence of learning on saccade velocity. Over the course of these category learning experiments, participants must learn the abstract properties of feature values, and combine these to learn a category response rule.

Category learning tasks have been a significant part of cognitive psychology for decades, and various influences on high-level performance during categorization are well documented. For instance, single dimensional rules like the one used in Experiment 2 are typically much easier than two dimensional rules (Experiment 1) and yield higher accuracy scores (Maddox, Filoteo, Hejl & Ing, 2004). Since the high level performance differences between single- and two-dimensional rule tasks are well understood, the interpretation of new lower-level measures is not muddied by questions of higher-level phenomena in the same way that it would be in a novel learning paradigm.

Understanding how saccades change over the course of learning is invaluable in the pursuit of developing a full model of visual attention - be it in the context of categorization or with respect to learned visual tasks in general.

Experiment 1: Two Dimensional Categories

This experiment is a rule-based category learning task, wherein the participants learn to determine the value of two features to make a decision between four possible categories (Table 1). One additional (irrelevant) feature is presented as a distractor. The goal of this analysis is to examine the difference in saccade speeds between those directed to two relevant items and those targeting a single relevant distractor[1].

Methods

Stimuli were presented to the participants as a series of alien animal cells. The task was to sort them into different groups using the information conveyed by three features. There were 69 undergraduate students from Simon Fraser University's Research Participation Pool who received partial course credit for their participation. Participants were assigned to either a speed (n = 25) or an accuracy (n = 25) instruction condition prior to exclusion, where the speed or the accuracy of their responses was emphasized prior to the start of the experiment, respectively. Gaze quality criteria (>70% of trials with >75% of gaze points collected) identified 4 participants for exclusion, and an additional 15 participants were excluded for failing to reach a learning criterion of 12 consecutive correct responses.

The speed/accuracy manipulation was originally implemented to encourage a Speed/Accuracy Tradeoff and explore how eye movements may differ when speed (or accuracy) was prioritized. The experiment failed to elicit an effect of either reaction time or accuracy. Failure to find an effect of condition may have been a function of the weak instantiation of the manipulation, in that the condition was communicated only in the instructions and block breaks. Another possibility is that the participants intrinsically prioritized accuracy in order learn the categories regardless of their condition.

Stimuli and Category Structure Features on the alien cell were separated in space by 10.6° and each spanned 1.3°. Each could take on two possible values (Figure 1). The combination of two features was diagnostic of the category, while the third was irrelevant (Table 1).

Table 1: Experiment 1 Category Structure

Category	Feature 1	Feature 2	Feature 3
A	1	1	0 or 1
B	1	0	1 or 0
C	0	1	0 or 1
D	0	0	1 or 0

[1] The data from these experiments have been analyzed for other measures, and are reported in McColeman *et al.* (2014). Data are publicly available through Summit, Simon Fraser University's open access data repository (http://summit.sfu.ca/collection/94 under "Speed-Accuracy Trade-Offs in Category Learning").

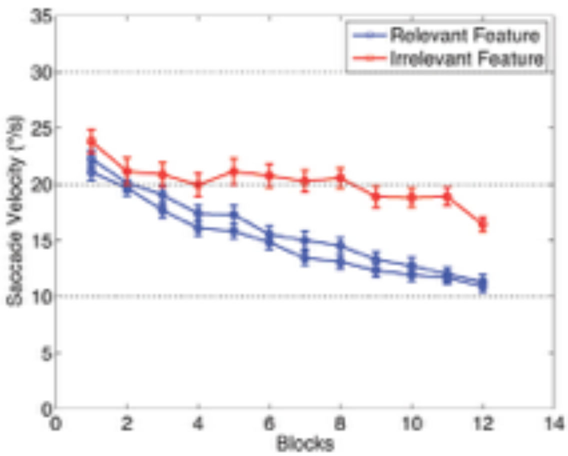

Figure 2. Saccade velocities for the Accuracy (solid line) and Speed (dashed line) conditions in Experiment 1. The velocity of saccades to relevant features is shown in blue, while the saccades the irrelevant features is plotted in red.

Procedure Participants are presented with 300 trials, shown in Figures 2 and 3 as twelve blocks of 25. During these trials, a Tobii X120 eye tracker recorded participants' eye movements, sampling at 120Hz with a spatial resolution of 0.5°. Fixations were defined using a modified dispersion algorithm (Salvucci & Goldberg, 2000), with a spatial threshold of 1.1° and a temporal threshold of 75ms. Saccade latency was defined as the difference in time between the end of one fixation and the beginning of the second; saccade amplitude is the angular distance between the spatial centroids of the gaze points that make up the first and second fixations. Saccade velocity is then calculated as amplitude divided by latency.

Each of the 300 trials began with a fixation cross, followed by the presentation of the stimulus upon which participants made a self-timed response. Feedback was provided with the the same stimulus, by showing the category label (A-D) in green font. If the participants' response was incorrect, then the label corresponding to their category choice was displayed in red. Participants were provided with breaks at the end of a block of trials.

Results

Because the primary interest of this study is the deployment of saccades in the context of categorization, those included in this analysis are only saccades that were made to one of the three possible features. Any saccade that ended with a fixation further than 150 pixels from the centre of a feature (8%) was dropped from the analyses independently from those reported here.

Saccade velocities (Figure 2) were predicted using a linear mixed effects regression model (LMER) including five predictors. The first was Condition (C), a between subjects predictor to identify the contribution of the speed/accuracy manipulation on the variance of saccade speeds. The first within subjects predictor was Block (B), where one block was 25 trials in the experiment. Estimates for the contribution of Block indicate how saccade speeds changed over the course of the experiment. Feature Relevance (R) was coded such that saccades ending on Feature 1 and 2 were assigned a value of 1 and saccades Feature 3 were assigned a value of 0. The Relevance predictor was meant to explore how the value of a fixated feature corresponds to the speed of the preceding saccade. An interaction between Block and Relevance (BxR) was included to investigate divergence in saccade velocities to relevant versus irrelevant features as the experiment unfolded. Fixation Order (FO) is a predictor meant to explore the influence of when a fixation occurs in the context of a trial on the speed of the saccade. The FO predictor can help in understanding how initial saccades may differ from later saccades. The Number of Fixations (NF) are included in the model to track repeated attempts to gather information from the stimulus in a trial (NF). For instance, there is a decrease in the number of fixations in each trial as the experiment runs its course: does this change the speed of the saccades that precede those few fixations? These predictors then form a model of predicting saccade velocity,

$$saccade\ velocity \sim \beta_0 + \beta_B + \beta_C + \beta_R + \beta_{BxR} + \beta_{FO} + \beta_{NF} + error$$

where β_0 is the intercept, and the remaining β values are coefficients for their respective predictors. Using the R Package "lme4" (Bates & Sarkar, 2007) to estimate the values of the coefficients, β_0=19.57 (t=31.57, SE=0.62), and the best predictors for saccade velocity are Block (β_B=-.96 t=-16.92, SE=0.06) suggesting a decrease in saccade velocity over blocks, and Relevance (β_R=1.24, t=5.31, SE=0.24) suggesting that saccades to irrelevant features are faster than saccades to relevant features. The remaining predictors, including the interaction between Block and Relevance (β_{BxR}=0.03, t=0.61, SE=0.05), Condition (β_C=0.17, t=0.22, SE=0.76), Fixation Order (β_{FO}=0.01, t=0.46, SE=0.01) and the Number of Fixations (β_{NF}=-0.01, t=1.19, SE=0.01) were poor predictors of saccade velocity.

Understanding how saccades relate to the fixations they precede is an important part of moving toward a richer account of visual attention in the context of learning. Fixation durations are known to be longer when the participant is fixating a relevant feature than an irrelevant feature (Chen, Meier *et al.,* 2013; Rehder & Hoffman, 2005), while we just showed that saccades are slower when they're directed to relevant features. It may be the case that there is a fixation duration/saccade speed trade-off, in that slower saccades use some of the resources that would be dedicated to processing during the fixation, but that seems unlikely given the findings just presented, and earlier work

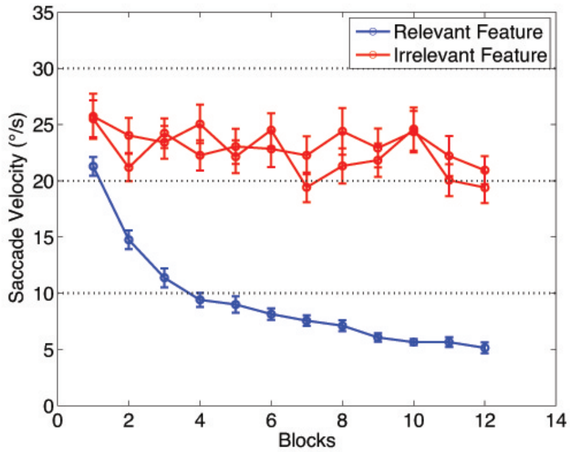

Figure 3. Saccade velocities for Experiment 2. The velocity of saccades to the relevant features is shown in blue, while the saccades to irrelevant features is plotted in red.

showing longer fixations to relevant features. Alternatively, slow saccades may precede longer fixations, in that additional cognitive processing occurs during the eyes' movements and the pause as information is gathered from the stimulus. Beyond that, it is possible that saccade velocities and fixation durations are informed by separable systems, and they display some independence.

To test these possibilities, a Spearman ranked correlation was conducted to investigate the relationship between saccade velocities and fixation durations. A negative correlation would indicate a trade-off between slow saccades and longer fixations, a positive correlation would suggest a common generator or a common motivator for slow saccades and long fixations while the absence of a correlation would suggest that both measures offer insight into separate processes. The test reveals a weak, but significant correlation between the two measures, $\rho_{(74178)}$=-0.01, p=0.004 although the very small ρ with the large sample size suggests little practical significance. Although the test indicates a relationship between the measures, it's clear that there is still a great deal of variance to account for, and that it is possible that the measures are reflecting different processes. It appears that both fixations and saccades are important to explore in developing a robust understanding of visual attention, at least in the context of a learning task.

Experiment 2: One Dimensional Categories

This is a simplified version of the category learning task, wherein only two categories are determined by a single feature. There are still three features on the stimulus, but two of them serve only as distractors (Table 2). In having a single feature, this experiment bridges more complicated learning tasks (like Experiment 1) with simpler target-following tasks that are more commonly employed in investigating saccades. The task still requires a decision, which in itself is more demanding than common cue-following saccade testing paradigms.

Methods

Unless otherwise noted, the stimuli (Figure 1), procedure and equipment are the same as Experiment 1. There were 67 participants, again were drawn from the Research Participation Pool. Three participants were excluded for failure to meet gaze quality criteria, and 7 failed to meet the learning criterion, leaving 28 participants in the accuracy condition, and 29 in the speed condition.

Results

Of the recorded saccades, 12% were not directed toward any of the three features and were excluded from analysis. A linear mixed effect regression model was built to explore the same variables as in Experiment 1 for their role in affecting saccade velocity: Condition (C), Block (B), Relevance (R), Block x Relevance (BxR), Fixation Order (FO) and Number of Fixations (NF). The only difference between the structure of this model and the one in Experiment 1 is the coding of Relevance (R). Since, in Experiment 2, Feature 1 is the only relevant feature, it alone is assigned an R value of 1. Features 2 and 3 are assigned an R value of 0.

The resulting model suggests that saccade velocities are influenced largely by Relevance (β_R=0.88, t=29.13, SE=0.03), Number of Fixations (β_{NF}=.01, t=11.62, SE=0.00), Block (β_B=-0.09, t=-11.01, SE=0.01), and the interaction between Block and Relevance (β_{BxR}=-0.04, t=5.85, SE=0.01), but negligibly by Condition (β_C=0.04, t=0.48, SE=0.09) and Order (β_{FO}=0.00, t=0.08, SE=0.00). The relatively strong influence of Relevance in predicting saccade velocity suggests that the task relevance of a saccade's target is affecting how quickly the it is executed.

A Spearman Rank correlation was conducted between saccade velocity and fixation duration to see if the two share a similar relationship to the underlying variables identified in the model above. The test failed to detect a relationship between the two measures, $\rho_{(37060)}$=-0.00, p=0.57. That is, changes in saccade velocities are either unrelated to changes in fixation durations, or there are underlying variables linking the two that have yet to be explored.

Table 2: Experiment 2 Category Structure

Category	Feature 1	Feature 2	Feature 3
A	1	1 or 0	0 or 1
B	0	0 or 1	1 or 0

Discussion

Through two experiments of varying complexity, we find that that saccades to irrelevant items are faster than saccades to goal-relevant targets. To complete the task in Experiment 1, participants had to use the information conveyed by two equally important stimulus features. Saccades to these features were slower than saccades directed to the irrelevant distractor. In Experiment 2, only a single feature was relevant for making a category decision, while there were two irrelevant distractors. In both experiments, saccades slow down over the course of learning, which may indicate a decrease in the proportion of quick, reflexive saccades relative to more purposeful eye movements.

These data begin to shed light on a more dynamic relationship between participants' developing understanding of task relevance and the rapidity of saccades. Critically, relevant items are learned over time, and as they are learned, saccades directed their way travel more slowly to their target. Also of note is that the saccades do slow over the course of learning. This is especially evident in Experiment 2, where the model identifies an important interaction between Feature Relevance and Block. The presence of this interaction provides statistical support for the divergence of saccade velocities based on the relevance of a feature.

Through inspection of Figure 3, it is apparent that the saccades to irrelevant features remain rather consistent over the course of learning, while saccades to the single relevant feature drastically slow down as the experiment progresses.

The importance of this divergence is twofold: for one, it provides evidence that the decrease in saccade velocity for the relevant items is not simply a function of fatigue. If that were the case, it would be expected that saccades to the irrelevant items would become slower too. Additionally, it suggests that increased knowledge of the category structure yields slower saccades to relevant features overall, but also incrementally as knowledge of the task increases (see McColeman *et al.* 2014 for higher level measures such as accuracy).

Experiment 1 paints a more complicated picture at first glance. Again, saccades do slow down over the course of learning, but saccades to both relevant and irrelevant items appear to decrease in velocity. Even so, the findings support an important influence of task relevance on saccade speed. The failure to elicit a strong divergence is likely due to the increased task complexity.

In Experiment 2, the participant simply has to execute a saccade to a pre-determined location. It is possible that the simplicity of this task evokes volitional saccades almost exclusively to feature 1, since few saccades will be directed to the known target after the rapid learning of the task structure. In Experiment 1, however, there is an extra step in that participants have to choose which of the two relevant features are their initial target.

As has been reported previously (Chen, Meier *et al.*, 2013) participants typically select a pattern of eye movements to deploy in the presence of multiple relevant features, and carry this pattern over a number of trials. Considering the feature relevance shown in Table 1, one participant may opt to employ a Feature 2-Feature 1-Feature 3 pattern of eye movements to extract information; another may gather information in a slightly different order: Feature 1-Feature 2-Feature 3. Both are equally good strategies, since Feature 1 and Feature 2 are of equal relevance and there is no reason to prefer one feature over the other. However, having to make this choice about the initial fixation may introduce some volatility into attentional processing, and invite a few more reflexive (fast) saccades than in the simpler Experiment 2. This study is an initial examination of the influence of task complexity, and further work will be necessary to flesh out the role of ordered fixations and choosing between multiple, equally good items to understand each factor's influence on saccade velocity.

It is important to note that these data differ somewhat from other work investigating saccade velocities. For instance, non-human primates saccade more quickly toward rewarded locations (Takikawa *et al.*, 2002; Chen, Hung, Quinet & Kosek, 2013); however, the findings shown here align with the previous contrast between volitional and reflexive eye movements, wherein consciously controlled saccades are understood to be slower than reflexive saccades (Godijn & Theeuwes, 2002; van Zoest, Donk & Theeuwes, 2004; Walker *et al.*, 2000). There are a number of differences in how tasks can be constructed to further investigate saccades and how they change over learning. It is possible that the primal reinforcement type of reward used more commonly in non-human primate saccade tasks differs from the more abstract reward that humans gather by learning new information or achieving task goals.

These early data investigating the change in saccade speed over learning provide motivation for further exploring the properties of saccades to elucidate complexities in the cognitive system. The evident dynamic interplay between high level learning and saccade velocity shown here offers a wide array of future research questions and possible practical applications. The increasing prevalence of gaze-based human computer interfacing in itself is an incentive to maximize the utility of gaze data, and to do so, understanding the relationship between the intention of an observer and recorded saccade speed is of critical import.

Acknowledgements

Thank you to the members of the Cognitive Science Laboratory for their ideas and contributions. Thank you especially to Joe Thompson and Jordan Barnes for their comments and Kim Meier for designing the stimuli used in the studies in addition to providing feedback on an early draft of this manuscript.

References

Ahlstrom, C., Nystrom, M., Holmqvist, K., Fors, C., Sandberg, D., Anund, A., Kecklund, G., & Akerstedt, T. (2013). Fit-for-duty test for estimation of drivers' sleepiness level: Eye movements improve the sleep/wake predictor. *Transportation Research Part C(26)*, 20-32.

Bates, D., Sarkar, D. (2007). lme4: Linear mixed-effects models using S4 classes. R package version 0.9975-12.

Blair, M. R., Watson, M. R., Walshe, R.C., & Maj, F. (2009). Extremely selective attention: Eye-tracking studies on dynamic attentional allocation to stimulus features. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*(5), 1196-1206.

Chen, L. L., Hung, L. Y., Quinet, J., & Kosek, K. (2013). Cognitive regulation of saccadic velocity by reward prospect. *The European Journal of Neuroscience, 38*(3), 2434–44.

Chen, L., Meier K.M., Blair, M.R., Watson, M.R., Wood, M.J.(2013). Temporal characteristics of overt attentional behaviour during category learning. *Attention Perception & Psychophysics, 75*(2), 244-256.

Cornwell, B., Mueller, S., Kaplan, R., Grillon, C. & Ernst, M. (2012). Anxiety, a benefit and detriment to cognition: Behavioral and magnetoencephalographic evidence from a mixed-saccade task. *Brain and Cognition, 78*(3), 257-67.

Godijn, R., & Theeuwes, J. (2002). Programming of endogenous and exogenous saccades: Evidence for a competitive integration model. *Journal of Experimental Psychology: Human Perception and Performance, 28*(5), 1039–1054.

Hallett, P. E. (1978). Primary and secondary saccades to goals defined by instructions. *Vision Research, 18*(10), 1279-96.

Matin, E. (1974). Saccadic suppression: A review and an analysis. *Psychological Bulletin, 81*(12), 899-917.

Maddox, W., Filoteo, J., Hejl, K., & Ing, D. (2004). Category number impacts rule-based but not information-integration category learning: Further evidence for dissociable category learning systems. *Journal of Experimental Psychology: Learning, Memory and Cognition, 30*(1), 227-245.

McColeman, C., Barnes, J., Chen, L., Meier, K., Walshe, R. & Blair, M. (2014). Learning induced changes in attentional allocation during categorization: A sizable catalog of attention change as measured by eye movements. *PLoS One, 9. doi: 10.1371/journal.pone. 0083302*Mirsky, J. B., Heuer, H. W., Jafari, A., Kramer, J., Schenk, A., Viskontas, I., Miller, B., & Boxer, A. (2011). Anti-saccade performance predicts executive function and brain structure in normal elders. *Cognitive and Behavioral Neurology, 24*(2), 50-58.

Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology, 51*(1), 1–41.

Roche, D. J. & King, A. C. (2010). Alcohol impairment of saccadic and smooth pursuit eye movements: impact of risk factors for alcohol dependence.*Psychopharmacology, 212,* 33-44.

Salvucci, D., Goldberg, J. (2000). Identifying fixations and saccades in eye-tracking protocols. *Proceedings of the Symposium on Eye Tracking Research and Applications*, 71-78.

Takikawa, Y., Kawagoe, R., Itoh, H., Nakahara, H. & Hikosaka, O. (2002). Modulation of saccadic eye movement by predicted reward outcome. *Experimental Brain Research, 142*, 284-91.

van Zoest, W., Donk, M., & Theeuwes, J. (2004). The Role of Stimulus-Driven and Goal-Driven Control in Saccadic *Visual Selection, 30*(4), 746–759.

Walker, R., Walker, D. G., Husain, M., & Kennard, C. (2000). Control of voluntary and reflexive saccades. *Experimental Brain Research, 130*(4), 540–4.