Title

A dynamic neural field model of self-regulated eye movements during category learning

Permalink

https://escholarship.org/uc/item/1t77n0fz

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

Authors

Barnes, Jordan I
Blair, Mark R
Tupper, Paul F
et al.

Publication Date

2015

Peer reviewed

A dynamic neural field model of self-regulated eye movements during category learning

Jordan I. Barnes (jordanb@sfu.ca)
Mark R. Blair (mblair@sfu.ca)
Department of Psychology, Simon Fraser University
8888 University Drive, Burnaby, BC V5A 1S6 Canada

Paul F. Tupper (pft3@sfu.ca)
Department of Mathematics, Simon Fraser University
8888 University Drive, Burnaby, BC V5A 1S6 Canada

R. Calen Walshe (r.c.walshe@sms.ed.ac.uk)
Department of Psychology, The University of Edinburgh
7 George Square, Edinburgh, EH8 9JZ Scotland

Abstract

Computational models of category learning and attention have historically focused on capturing trial and experiment level interactions between attention and decision. However, evidence has been accumulating that suggests that the moment-to-moment attentional dynamics of an individual affects both their immediate decision-making processes as well as their overall learning performance. To extend the scope of these formal theories requires a modeling approach that can index fine-grained decision-making at millisecond time scales. Here we implement a model of eye movements during category learning using concepts from Dynamic Neural Field Theory research. Our model uses a combination of timing signals, spatial competition and Hebbian association to simultaneously account for a number of foundational attentional efficiency results from eye tracking and category learning. We report the results of fitting this model to accuracy, fixation probabilities, fixation counts and fixation duration data in 42 subjects from a standard category learning experiment.

Keywords: attention, eye-tracking, dynamic field theory, cognitive modeling, category learning.

Introduction

Attention and learning are intrinsically related (Shepard, Hovland & Jenkins, 1961; Kruschke, 2011). When we learn a new skill, we learn to selectively attend to information relevant to that skill. As research into the fundamental neurophysiological principles of attention and learning progresses however, the historical psychological methods for studying attention and learning, and those implied by modern approaches in computational neuroscience, have thus far failed to converge. In the category learning literature for instance, where some of the clearest formalizations of the relationships between learning and attention have been developed, 'attention' is simply a weight on a feature dimension indicating the degree of importance it is assigned when making category decisions. Mathematical learning techniques that adjust these weights trial-by-trial over the course of an experiment have worked relatively well in fitting human data (Kruschke, 1992).

These approaches have been used in explanations for phenomena such as: attentional blocking and highlighting (Kruschke, Kappenman, & Hetrick, 2005), rapid attentional shifting (Kruschke & Johansen, 1999), and cluster driven attention (Love, Medin, & Gureckis, 2004). However, the biological basis for these techniques has not always been clear. What is becoming more clear, is that the brain is using a complex set of systems to perform attentional functions that may only look like those described by attention weights at a trial-level (Gottlieb, 2012).

As a first step toward bridging the gap between attention weights at a trial-level and the ongoing attentional decisions made within a trial, a recent model called *RLAttn*, standing for Reinforcement Learning of Attention (Barnes, McColeman, Stepanova, Blair & Walshe, 2014), used temporal difference error to learn transition probabilities between micro-information states, which get traversed by eye movements as part of arriving at a decision. This model showed the possibility of using aggregated eye movements to index attention in ways that correspond to classical descriptions of learned attention in terms of bias to particular feature dimensions. RLAttn is a first step, but it must be admitted that the target of a saccade is not just a function of the relative values of different information states. The spatial arrangement (Lipinkski, Spencer & Samuelson, 2010) and perceptual salience of visual information (McColeman, & Blair, 2013) for instance, affects these decisions as well. A more fundamental account of these systems guided by advancements in neuroscience is needed. For this reason, we have developed a new model that couples together the dynamically operating cognitive mechanisms needed to interface sensori-motor signals with learned attention.

Decision timing

When the relationships between category decisions and other measures like reaction time and ocular motor fixation durations are looked at closely, it quickly becomes clear that there is quite a bit happening beneath the level of the category decision. For instance, Blair and Watson (2008)

show that category learning performance can be well predicted from the duration of time that participants spent looking at the stimulus during the feedback phase in just the first 10 trials. Furthermore, individuals develop patterned fixation orders, and fixation durations tend to decrease over the course of the experiment, especially to relevant information (Rehder & Hoffman, 2005; Meier & Blair, 2013; Chen, Meier, Blair, Watson & Wood, 2013).

What then are the determinants of these sub-decisions regarding where and when to look, and how do they affect the trajectory of learning? These are some of the core questions we are trying to investigate with this model. Only by having a model that evolves its behaviour through learning can you begin to say how the factors influencing these decisions relate to one another. In a review of the cognitive and neural factors that influence decision timings for instance, Wittman and Paulus (2008) argue that timings are strongly related to attention and are influenced by a number of specialized neural networks. Inspired by ideas like these, we incorporate neurons for both decision responses and for saccade initiation mechanisms, whose activations grow with time and are reset when the relevant task is performed. One can think of these units as implementing a kind of impatience. Like Wittman and Paulus who suggest that timing is highly individual specific, Ghafurian and Reitter (2014) have implicated impatience as being a distinct trait between individuals. This work has shown that individuals maintain a constant offset in timing decisions, above and beyond factors like risk aversion, in tasks where individuals need to estimate optimal response times in order to maximize reward.

Dynamic Neural Field Theory

Models that attempt to describe processes as they unfold in time are naturally described in the language of differential equations. There are many models of the temporal characteristics of eye movements (see Nuthmann, Smith, Engbert, & Henderson, 2010, for example) but only a few are described purely in dynamic terms (see Perone & Spencer, 2013, or Schneegans, Spencer, Schöner, Hwang & Hollingworth, 2014). We opt for a modeling framework that looks beyond just the timing of a decision however, to incorporate a level of spatial dynamics that contributes to the actual decisions made. One framework that does this is called Dynamic Neural Field Theory (DNFT, or just DFT), and is built on some of the known properties of neural populations (Georgopoulos, Schwartz, & Kettner, 1986; Erlhagen & Schöner, 2002). The elements of these neural fields are defined by the entire population of neurons, weighted by their individual tuning preferences for the given element. These preferences are typically modeled by a standard receptive field across the population which has a locally excitatory and globally inhibitory difference of Gaussians (also known as 'Mexican hat') approximation (see Kopecz & Schöner, 1995, for an example of this kind of lateral interaction in a similar context).

Abstract fields with these properties can develop self-sustaining peaks of activity even in the absence of exogenous input, allowing for a kind of memory (Amari,

1977). The general form of the 2-dimensional field equation, like those used here, for a field u, is defined by:

$$\tau \dot{u}(x,y) = -u + h + S(x,y) + \int\int w(x-x', y-y') f[u(x',y',t)]dx'dy'$$

where τ scales the time, x and y represent points in space, h is a resting level, S is exogenous input, w specifies the interaction kernel of the field and f is a sigmoidal thresholding function. In this framework, any sufficiently activated neuron will contribute excitatory or inhibitory activation to the abstract field according to a distance along a metric of representation governed by the receptive field of the neuron.

Model

The model presented here is designed to work with similar kinds of task constraints as human subjects given the same experiment. A trial begins with the model looking at the center of a screen in the same manner that a human begins a trial with a central fixation cross. It then makes simulated eye movements to stimulus features that compete for attention. Each eye movement registers foveated information in a kind of visual working memory. This activates associated categories, which then affect the decisions about where to look next and when to look at feedback or end the trial. Experiment instructions specifying a particular set of features and responses are similarly coded into the model.

The model contains many dynamically interacting parts, as shown in Figure 1. To begin with, the visual field (A) has 2 spatial dimensions defined in retinotopic coordinates with the fovea always at the center. A 3rd dimension on this field contains the feature values for a particular trial that serve as input to a feature detection layer of neurons once a particular location is foveated; a particular color is only defined for one location for the whole experiment. Feature neurons (D) tuned to these specific locations of the color dimension of the input are connected to a layer of category neurons through a synaptic weight matrix. As a feature neuron is activated, it activates categories as a function of the level of its activation and the weight (F) on the synapse.
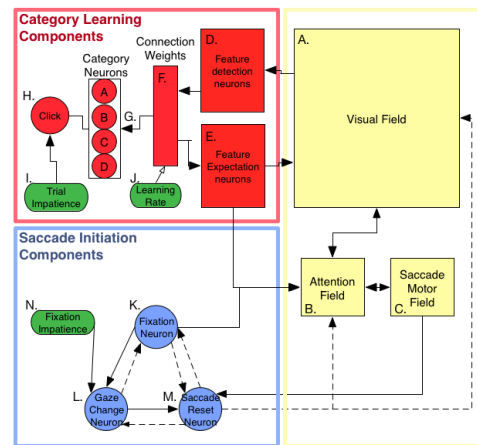

Figure 1: Schematic structure of the 3 sub-sections of the model: spatial representation, category learning, saccade initiation. Green shapes represent free parameters. Dashed lines indicate inhibitory connections.

A small gain factor is also applied that attenuates category activation according to the consistency of a feature with particular categories. This is done because there is often non-diagnostic information that is correlated across all categories.

The mutually inhibitory category neurons (G) in the model play a central role in the prioritization of attention and timing of decisions. These category neurons feed into a decision neuron (H) which, if activated strongly enough, will make a category decision, either initiating a feedback phase or, if very highly activated end the trial without looking at feedback. This decision neuron is excited by exponentially increasing input, whose rate of growth is set by the *trial-impatience* free parameter (I). The larger this number, the faster the model will tend to make a decision on a particular trial. Toward the start of an experiment when categories are not very strongly activated, the timing of category decisions is largely determined by the impatience of the model pushing the neuron over a threshold, but as categories begin to get selectively activated over the course of learning, category feedback pushes this neuron across threshold faster and response times gradually decrease. The same processes generalize to the feedback phase.

While this is happening, activity from the category neurons recurrently feeds back through the synaptic weight matrix and into a layer of feature expectation neurons (E). As particular categories grow in activation over the course of the trial, the outstanding features associated with that category become activated. Gain is again applied for this direction of the flow of activity, biasing the model to saccade to highly relevant features. Inhibitory input from the feature detection neurons simultaneously reduces the expectation of already viewed information, acting as a kind of inhibition of return on a time scale set by the memory for items in working memory.

The feature expectation neurons do two things. First, they magnify inputs along the 3rd dimension of the visual field. This has the effect of acting as a kind of salience boost to that information. Second, the feature expectation neurons are connected to specific locations on the spatial attention field (B). This divergence models the distinction between salience driven attention and the relative weightings that might be calculated there, and task level attention which could much more strongly drive voluntary shifts in attention.

The spatial attention field is modeled as a 2-dimensional dynamic neural field which represents the elements of the stimulus in spatiotopic, or world-centric, coordinates as opposed to the retinotopic coordinates of the visual field. The extant point of fixation receives an excitatory boost from a fixation neuron (K), which plays a role akin to neurons at the rostral pole of the superior colliculus. The intermediate layers of this brain region are known to topographically organize the selection of saccadic end-points (Robinson, 1972). Here, neurons coding for amplitude and direction of the eye movement compete in the initiation of saccades (Munoz & Wurtz, 1993). A gaze

change neuron (L) acts as the target for the exponentially increasing *fixation-impatience* parameter (N). This input ensures that the model will make new fixations. This gaze change neuron inhibits the fixation neuron, consequently modifying the balance of competition on the attention field, while also exciting the saccade motor field (C) and a 3rd and final neuron, the saccade initiation neuron (M). When this neuron crosses a threshold, the location of maximum activation on the saccade motor field is selected as a saccadic end-point and a saccade is initiated. This re-orients the visual field, which is suppressed during the saccade and resets the fixation impatience timer. Oscillatory fixation dynamics can be generated a number of different ways, for instance by treating the fixation neuron's resting level as a separate state variable sensitive to the changes in activation of the fixation neuron (Perone & Spencer, 2013), but we chose an impatience parameter as a way of separating the contribution of an individual's processing speed from their gaze strategy, which may cause saccades prior to finishing processing (Kiani, Hanks & Shadlen, 2008).

Learning occurs after the model has made a category decision and a feedback phase is entered. When the model looks at a particular location in space, where a feedback "button" is located, additional input to the correct category is provided. For every moment that the model spends in the feedback phase, Hebbian co-activation type association strengthens (J) the weights between the active features and the correct category. This is modulated with every fixation, in that the activation of a feature detecting neuron is boosted by fixational input above and beyond the attractor defined by the self-sustaining level of the feature detecting neurons, effectively increasing its association with the category active category in the feedback phase.

Method

Eye tracking and category learning data was obtained from the publicly available Meier and Blair (2013) dataset[1]. Full technical details can be found in the original paper. The experiment required participants to sort images of fabricated microorganisms defined by 3 organelle features (see Figure 2) into 1 of 4 categories (A1, A2, B1, B2) on each trial. Feedback indicating the subject's choice and the correct category was provided after each decision. We use only participants in the equal category base rate condition (as opposed to a condition where some categories appear more frequently) of this experiment, who had at least 70% of their gaze collected, and who reached a criterion of 24 trials in row correct (n=42). From these subjects, we only look at the first 360 trials of this data for each subject.

The category structure, shown in Table 1, is defined such that feature 1 is more informative than either feature 2 or feature 3, in that its value determines which of feature 2 or 3 is relevant for a particular trial. Optimal attention then, requires participants fixate feature 1 first.

---

[1] http://summit.sfu.ca/item/12715

Feature 1	Feature 2	Feature 3	Category
0	0	0/1	A1
0	1	0/1	A2
1	0/1	0	B1
1	0/1	1	B2

Table 1. Feature 1 is relevant for all categories. Feature 2 is relevant only for the A categories, and Feature 3 is relevant for only the B categories.
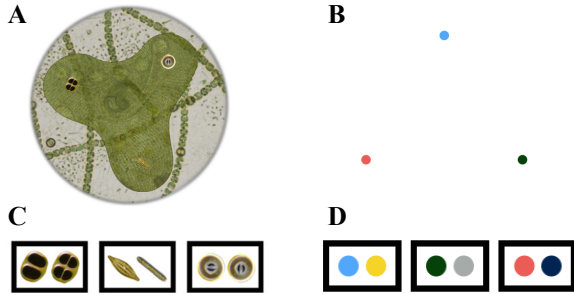

Figure 2: A: Example stimulus from Meier and Blair (2013). The three features of interest subtend approximately 1.7° of visual angle each, and are spaced roughly 10.6° apart. B: Three equidistant color features presented to the model. C/D: The two feature options for each location.

While the stimuli from this experiment are more complex than can be currently processed by the model, we have elsewhere employed the simple color stimuli suitable for use with the model (Barnes & Blair, 2014).

Simulations

To find out how well the model could simulate the human data we employed a grid search algorithm, varying our 3 free parameters (learning rate, trial impatience and fixation impatience) in looking for the best fitting version of the model for each individual. Model fits were judged based on the difference in mean and slope of 4 aspects of the human data (accuracy, trial fixation counts, probability of fixating the irrelevant feature and fixation durations). A single simulation of a single participant can take several hours, consequently, the model was just run several times at each level of a fairly coarse grid. Each subject's data was then fit with the point on the grid that minimized the weighted least squares error.

In what follows, we assess the quality of our fitted model in two ways. For each measure we first look at how the model performs on a representative individual in the population. This representative individual is the one with the median weighted squared error; so there are roughly equal numbers of subjects that we fit better or worse than this representative individual. We do this such that the reader can get a sense of the variability of the model under a single set of parameters. We next look at the distribution of the all the subjects and the corresponding fitted models for all the subjects.
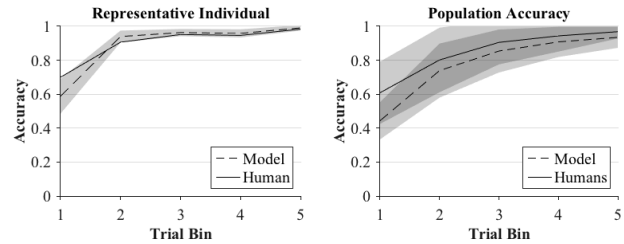

Figure 3: Accuracy learning curves. Grey bars indicate population standard deviation. On left is the fit of the model for the representative subject and on the right is the average behaviour over all subjects and fits.

Accuracy

We considered the fit of the accuracy, that is the fraction of the categories that were correctly guessed, as an important benchmark. As seen in Figure 3, the model performed well in fitting the learning performance data.

Probability of fixating irrelevant feature

Here we report a measure of attentional optimization characterized by the averaged binary probability of looking at an irrelevant feature on a particular trial. As seen in Figure 4, while the model captures the decreasing probability of fixating irrelevant information well across the whole experiment, the model always fixates all features at the beginning, whereas human do not. There are several non-exclusive possibilities that might explain this. The first possibility we considered is that the eye tracker may be losing track of the eyes at times during the experiment (due to blinks, or head turns, for instance) possibly depressing the initial probabilities of fixating irrelevant information. Upon further investigation, controlling for gaze loss at the individual level, this did not appear to account for the reduce probabilities. Further, Rehder and Hoffman (2005) report similar initial fixation probabilities for the features in their experiment. It is possible that not looking at all the features might be a strategic choice that human participants make in order to test simple rules first (possibly due to the assumption that the task itself could be quite simple). The simple rule first hypothesis is common to category learning models (e.g. Love, Medin & Guereckis, 2004; Nelson & Cottrell, 2007) and occurs naturally in models like RLAttn, where guessing the category may have as high an action selection probability as fixations to features until incorrect answers have a chance to punish this behaviour. Finally, it may be the case that precise fixations are not totally necessary in most category learning experiments (see Tatler, Hayhoe, Land & Ballard, 2011 and Coren, 1986). In any event, as currently constructed, the model presented here is too fastidious to produce the human data on this score.

Fixation Counts

Another common form of attentional optimization is to reduce the overall number of fixations (to features) per trial over the course of the experiment (McColeman et. al, 2014). Figure 5 shows human and model data. The model captures the overall trend of decreasing fixation count across the
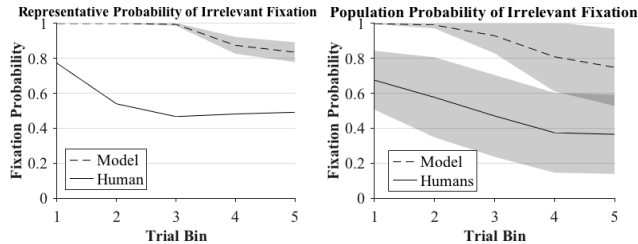
Figure 4: The probability of fixating irrelevant feature over the course of the experiment.

Figure 6: The mean fixation durations to all features, irrespective of relevance, over the course of the experiment.
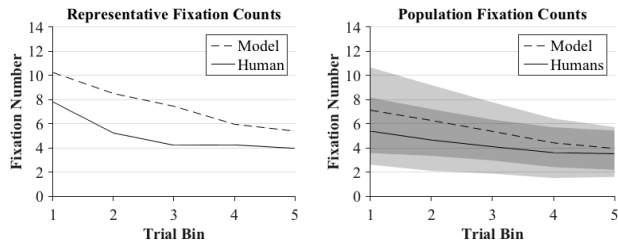

Figure 5: The total fixation counts to all features over the course of the experiment. The standard deviation for the sample of model runs that best fit the representative individual was very small.

course of the experiment, but once again overestimates the number of fixations at the beginning of the experiment.

Fixation Durations

We believe that fixation durations are an important measure for models to fit going forward as they reflect a number of important cognitive processes such as scene perception (Walshe & Nutthman, 2013), fixation type (Ballard, Hayhoe, Pook & Rao, 1997), and information relevance (Chen et. al, 2013) among others. Figure 6 shows the human data and model simulations. On average the model fits human data well. The model shows a modest but similar decline in durations as the human subjects over the course of the experiment, within a similar magnitude.

Discussion

Modeling eye movements during category learning is a recent trend (Nelson & Cottrell, 2009; Barnes et al,. 2014). There are disparate motivations for moving in this direction. Recent work has shown that learning methods based on error reduction alone are insufficient to explain human eye movement data in category learning, pointing to a need for new kinds of models that might learn with different methods (Blair, Walshe, Barnes & Chen, 2011; McColeman et. al, 2014). Additionally, work by Ballard, Hayhoe, Pook and Rao, (1997) on the relationship between working memory and skill acquisition, emphasizes the importance of "just-in-time" decision-making at an *embodied* time scale operating at around 1/3 - 2 sec. The motivating intuition is that there should be physical actions that line up with the time scales of information processing required for *variable binding*. In this view, instructions about where to look next can be
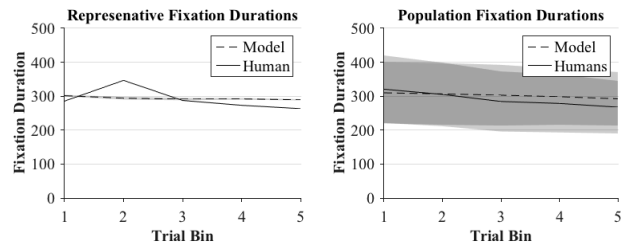
thought of as being stored "in the world" as opposed to in the head, reducing the representational demands on the computational cognitive system.

While ours is a much different model, the decisions it makes about where to look next are contingent on serially accessed parts of the visual world. In general, the category learning paradigm is well-suited for analysis at the embodied time scale because moment-to-moment decisions can reflect subtle manipulations in the category structure (Meier & Blair, 2013). Ultimately, the mechanisms that make gaze fixations relevant to overall learning are only just starting to be understood. The model of just-in-time gaze learning advances the idea that attentional optimization is a natural consequence of quickly dropping fixationally bound variables from working memory (Ballard, Hayhoe, Pook and Rao, 1997). A recent DNFT model of infant gaze behaviours (Perone & Spencer, 2013), with many similarities to our model, showed how the durations of individual fixations can work to modulate an infant's well-documented familiarity to novelty bias transition at around 8-10 weeks of age (Wetherford & Cohen, 1972): the idea being that longer fixations leave a larger Hebbian association in long term memory which combines with inhibition from working memory to make parafoveal information relatively more salient. Only models that allow for moment-to-moment changes in attention could feasibly model these kinds of emergent differences over the course of learning.

We believe that this line of research has the potential to show how complex behaviours can emerge from the interactions of a simple set of parameters, in our case just learning rate, fixation impatience and trial impatience, over the course of learning. Not only does the model presented here scale its learning through its overall gaze time but it also scales what it looks at it by what it knows. To our knowledge, no other model has attempted to simultaneously fit such a wide array of behavioural measures. An important future direction for these modeling efforts is to rigorously test this approach on other kinds of category structures, timing constraints and stimulus types. With the introduction of this model, we hope to provoke efforts that seek to explain many kinds of cognitive and sensori-motor behaviours simultaneously.

Acknowledgments

References

Amari, S.-I. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol Cybern, 27(2)*, 77–87.

Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. (1997). Deictic codes for the embodiment of cognition. *Behav Brain Sci, 20*(4), 723–42.

Barnes, J.I., & Blair, M.R. (2014). The Influence of Space and Relevance on Eye Movement Distributions. *Northwest Conference on Cognition and Memory.*

Barnes, J.I., McColeman, C.M., Stepanova, E., Blair, M.R. & Walshe, R.C. (2014). RLAttn: An actor-critic model of eye movements during category learning. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 1892-1897).

Barnes, J.I., Walshe, R.C., Tupper, P.F., & Blair, M.R. (2013) A dynamic neural field model of eye movements during category learning tasks. *Learning to Attend, Attending to Learn: Neurological, Behavioural, and Computational Perspectives*. Retrieved from: http://cslab-sfu.ca/wp-content/uploads/2013/11/AttnLearning-Workshop-Tempus-Poster.pdf

Blair, M.R., Walshe, C., Barnes, J.I., & Chen, L. (2011). Rethinking the role of error in attentional learning. In L. Carlson, C. Hölscher, & T. Shipley (Eds.), *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society* (pp. 1649-1655).

Chen, L., Meier, K. M., Blair, M. R., Watson, M. R., & Wood, M. J. (2013). Temporal characteristics of overt attentional behavior during category learning. *Atten Percept Psycho, 75*(2), 244–56.

Coren, S. (1986). An efferent component in the visual perception of direction and extent. *Psychol Rev, 93*, 391–410.

Erlhagen, W., & Schöner, G. (2002). Dynamic Field Theory of Movement Preparation. *Psychol Rev, 109(3),* 545-572.

Georgopoulos, A. P., Schwartz, A. B., Kettner, R. E. (1986). Neural population coding of movement direction. *Science*, 233, 1416–1419.

Ghafurian, M., & Reitter, D. (2014). Impatience, Risk Propensity and Rationality in Timing Games. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), Proceedings of the 35th Annual Conference of the Cognitive Science Society (pp. 2841-2846).

Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron, 76*(2), 281–95.

Hikosaka, O., Takikawa, Y., & Kawagoe, R. (2000). Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiol Rev, 80*(3), 953–78.

Kiani, R., Hanks, T. D., & Shadlen, M. N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci, 28*(12), 3017–29.

Kopecz, K., & Schöner, G. (1995). Saccadic motor planning by integrating visual information and pre-information on neural dynamic fields. *Biol Cybern, 73*, 49–60.

Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychol Rev, 99*(1), 22–44.

Kruschke, J. K. (2011). Models of Attentional Learning. In E. M. Pothos & A. J. Wills (Eds.), Formal Approaches in Categorization,(pp.120–152).

Kruschke, J. K., & Johansen, M. K. (1999). A model of probabilistic category learning. *J Exp Psychol Learn, 25*(5), 1083–119.

Kruschke, J. K., Kappenman, E. S., & Hetrick, W. P. (2005). Eye Gaze and Individual Differences Consistent With Learned Attention in Associative Blocking and Highlighting. *J Exp Psychol Learn, 31*(5), 830–845.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: a network model of category learning. *Psychol Rev, 111*(2), 309–32.

McColeman, C. M., Barnes, J. I., Chen, L., Meier, K. M., Walshe, R. C., & Blair, M. R. (2014). Learning-Induced Changes in Attentional Allocation during Categorization: A Sizable Catalog of Attention Change as Measured by Eye Movements. *PLoS ONE, 9*(1).

McColeman, C., & Blair, M. (2013). The Influence of Salient Distractors over the Course of a Category Learning Task. *J Vision, 13*(9), 506-506.

Meier, K. M., & Blair, M. R. (2013). Waiting and weighting: Information sampling is a balance between efficiency and error-reduction. *Cognition, 126*(2),319–25.

Munoz, D. P., & Wurtz, R. H. (1993). Fixation cells in monkey superior colliculus. II. Reversible activation and deactivation. *J Neurophysiol, 70*(2), 576–89.

Nelson, J. D., & Cottrell, G. W. (2007). A probabilistic model of eye movements in concept formation. *Neurocomputing, 70*(13-15), 2256–2272.

Nuthmann, A., Smith, T. J., Engbert, R., & Henderson, J. M. (2010). CRISP: a computational model of fixation durations in scene viewing. *Psychol Rev, 117*(2), 382–405.

Perone, S., & Spencer, J. P. (2013). Autonomy in action: linking the act of looking to memory formation in infancy via dynamic neural fields. *Cognitive Sc, 37*(1), 1–60.

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

Robinson, D. (1972). Eye movements evoked by collicular stimulation in the alert monkey.*Vision Res*,12, 1795–1808.

Schneegans, S., Spencer, J. P., Schöner, G., Hwang, S., & Hollingworth, A. (2014). Dynamic interactions between visual working memory and saccade target selection. *J Vision, 14*, 1–23.

Shepard, R. N., Hovland, C. I., & Jenkins, H. (1961). Learning and memorization of classifications. *Psychol Monogr-Gen A, 75(13),*1-42.

Tatler, B. W., Hayhoe, M. M., Land, M. F., & Ballard, D. H. (2011). Eye guidance in natural vision: reinterpreting salience. *J Vision, 11*(5), 1–23.

Walshe, R.C., & Nuthmann, A. (2014). Asymmetrical control of fixation durations in scene viewing. *Vision Res*, *100*, 38-46.

Watson, M. R., & Blair, M. R. (2008). Attentional allocation during feedback: Eyetracking adventures on the other side of the response. In *Proceedings of the 30th Annual Meeting of the Cognitive Science Society* (pp. 345-350).

Wetherford, M. J., & Cohen, L. B. (1973). Articles Developmental Changes in Infant Visual Preferences for Novelty and Familiarity. *Child Dev, 44*(3), 416–424.