A COMPUTATIONAL MODEL OF EYE-MOVEMENTS IN CATEGORY LEARNING TASKS

by

R. Calen Walshe
B.A. (Hons), Simon Fraser University, 2009

THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF ARTS

in the
Department of Psychology
Faculty of Arts and Social Sciences

SIMON FRASER UNIVERSITY
Summer 2011

Approval

Name: R. Calen Walshe

Degree: Master of Arts (Department of Psychology)

Title of Thesis: A Computational Model of Eye-Movements in Category Learning Tasks

Examining Committee:

Chair: Dr. Kathleen Slaney
Associate Professor

Dr. Mark Blair
Senior Supervisor
Assistant Professor

Dr. Timothy Racine
Supervisor
Assistant Professor

Dr. Paul Tupper
Supervisor
Associate Professor

Dr. Nancy Hedberg
Associate Professor
Department of Linguistics

Date Defended/Approved: August 11, 2011

SFU SIMON FRASER UNIVERSITY
LIBRARY

Declaration of Partial Copyright Licence

STATEMENT OF ETHICS APPROVAL

The author, whose name appears on the title page of this work, has obtained, for the research described in this work, either:

(a) Human research ethics approval from the Simon Fraser University Office of Research Ethics,

or

(b) Advance approval of the animal care protocol from the University Animal Care Committee of Simon Fraser University;

or has conducted the research

(c) as a co-investigator, collaborator or research assistant in a research project approved in advance,

or

(d) as a member of a course approved in advance for minimal risk human research, by the Office of Research Ethics.

A copy of the approval letter has been filed at the Theses Office of the University Library at the time of submission of this thesis or project.

The original application for approval and letter of approval are filed with the relevant offices. Inquiries may be directed to those authorities.

Simon Fraser University Library
Simon Fraser University
Burnaby, BC, Canada

Abstract

Models of attention in category learning tasks have typically treated attention as a weighting of how influential a feature is to the correct classification of the overall stimulus. Attention shifting is frequently modelled as occurring after the trial is completed (Kruschke, 1992). Recent work has demonstrated in detail how learned attention develops during the course of a single trial. Currently, there is no model which can account for the dynamic attentional shifts that are identified by eye-tracking data. Additionally, research is many fields has identified the need to explore cognitive models that are based on a more naturalistic view of human behaviour. New mathematical techniques utilizing concepts from dynamical systems has greatly increased the tractability of developing such models. This thesis describes two category learning experiments and introduces a new computational model that produces a real-time simulation of eye-movements in these tasks. Human data is compared with the model output and the implications of this model to category learning and related fields is discussed.

Keywords: Category learning, Computational modelling, Attention, Dynamic field theory, Dynamical systems

Dedication

This project is dedicated to my wife Daniela who has always been there with love, support and a cup of tea when I needed them most.

Acknowledgements

I would like to thank the members of the Cognitive Science Lab who have continuously inspired me with their commitment and dedication. I would also like to thank my supervisors Drs. Mark Blair, Paul Tupper and Tim Racine for providing guidance throughout the development of this project.

Table of Contents

List of Tables

List of Figures

Introduction

One of the hallmarks of human cognitive capacities is our ability to classify objects from the physical environment into discrete categories. The capability to segregate that which is poisonous from that which is edible is surely a kind of categorization that would have aided our evolutionary ancestors. A key ability required for the correct classification of a stimulus is identifying which features of a class are those that are most likely to lead to correct classification and attending to those features accordingly. This remarkable ability is present both in early infancy when the capacity for categorization is first developing, and in experts who have developed intricate classification abilities. In a classic study of the sexing of chicks Biederman and Shiffrar (1987) documented the ability of experts to sex day old chicks with near perfect accuracy while performing over 1000 categorizations per hour. Their expertise was a result of their ability to identify and exploit the most informative features of the chicks sex while de-emphasising the irrelevant information. The ability to selectively attend has long been identified as forming a crucial component of human classification ability (Shepard, Hovland, & Jenkins, 1961).

One of the most well known formal models of human categorization is an exemplar model known as the Context Model. In exemplar models, it is assumed that a category consists of a collection of stored instances which collectively define the class. The context model produces the probability that a stimulus will be assigned a particular classification and there has been well documented evidence supporting the validity of the model (Busemeyer, Dewey, & Medin, 1984; Medin, Altom, Edelson, & Freko, 1982; Medin, Dewey, & Murphy, 1983). Nosofsky (1984, 1986 & 1988) extended the original model to cover continuous valued stimuli as well as to provide an explanation for the relationship between stimulus identification and stimulus classification. In the Generalized Context Model (GCM) a dimensional weight parameter was added that represented the amount of attention that was being applied to a stimulus dimension. The addition of the attention parameter unified explanations of identification and classification under the same explanatory framework by demonstrating that contradictory results can be explained by showing that different attention weights are being applied in the classification and identification tasks tasks (Logan, 2004). While the GCM includes an attention parameter that influences categorization, Nosofsky (1984) claimed that this parameter was selected by people in such a way so as to maximize the probability of correct classification and that the optimized attention weights are adjusted throughout learning. Although an assumption of optimized attention was made in the GCM it did not provide a computational mechanism to explain how these dimensional attention weights are learned.

Kruschke (1992) developed a model, ALCOVE (attention learning covering map), that utilizes all the same exemplar based representational machinery that is present in the GCM, but incorporated an error-driven learning mechanism that drives the association between exemplars and categories. In addition, ALCOVE eschews attention weight free parameters for a mechanisms which is capable of learning these weights on a trial by trial basis. ALCOVE's success at modeling numerous categorization results had a substantial impact on the field and stimulated the creation of a family of models that emphasize aspects of localized exemplar based categorization, dimensional attention parameters with error-driven learning assumptions (Ashby & Alfonso-Reese, 1998; Kruschke, 2001; Kruschke & Johansen, 1999; Love, Medin, & Gureckis, 2004).

A feature of these models is that they implement a variety of attention that could be called task-specific. Attention in these models is a learnt attentional distribution that is applied invariantly to all stimuli that are present in the task domain. For instance, when learning to classify bird; these models will learn attention weightings for beaks and feathers, but are not capable of defining an attention profile for individual birds, even when this may be required for classification or when it may increase the efficiency of attentional allocation (Blair, Watson, Walshe, & Maj, 2009). In Blair et al. (2009) the authors demonstrated attentional allocations that are contrary to those that are predicted by the exemplar models previously reviewed. The authors showed that fixation durations to the relevant and irrelevant features differed for stimuli belonging to distinct categories. The authors also reported unique temporal characteristics of fixations to stimuli of differing categories with the most relevant features to classification of a particular stimuli being fixated first. That this stimulus-specific attention is not predicted by many of the earlier exemplar based models is a clear shortcoming of their ability to account for this important aspect of human performance in categorization tasks.

Recent work using eye-tracking has made unique contributions to understanding the nature of attentional allocations under a variety of experimental conditions. Using head-mounted displays researchers have been able to study eye-movements as they develop during naturalistic tasks such as making tea or sandwiches and have demonstrated the importance of fixating on locations that will help to optimize performance on a task (Land, Mennie, & Rusted, 1999). For instance, when grasping the teapot the handle is likely going to be the most likely locus of information for successfully completing the task while the ability to flexibly monitor and shift gaze to a new location such as to the spout when pouring has been identified as crucial to optimal performance. This line of research has revealed the important role that that contextual factors have in determining both where eye-movements will be programmed to land, where they will go next as well as how the information in the field of vision will be processed (Rothkopf, Ballard, & Hayhoe, 2007). For instance, using VR simulations and eye-tracking, Triesch, Ballard, Hayhoe and Sullivan (2003) were able to control exactly the point at which a feature in a tracked object became task irrelevant and observe the subtle shifts in attentional processing that resulted. The dynamic, context sensitive nature of gaze allocation is also active in social interaction. In Foulsham, Cheng, Tracy, Henrich and Kingstone (2010) the authors demonstrated that when watching a video of a group decision making, task participants would fixate on high status individuals more often and for longer.

A number of computational models have been applied to the theoretical understanding of the nature of eye-movements and visual attention in naturalistic real-world tasks. Sprague and Ballard (2003) suggested the use of simulated agents as platforms to test cognitive theories and applied this approach to modelling attentional selection. The notion of using simulations to investigate theories of attention results from the observation that cognition can not be studied independently from an agent's physical state and nor are the physical states separate from the cognitive states under investigation (Clark & Toribio, 1994). In their model, the simulated environment creates a direct link between perception, action and attention and allows for the observation of the relationship between cognition and the environment. Another strength of the simulation methodology is that it enables direct comparison of both human and model data as both are presented with identical environments and produce qualitatively similar types of behavioural data.

In order to understand the temporal nature of eye-movements, it is important to understand what role eye-movements play in the overall pattern of cognition. One approach has been to treat fixations as information accessing queries and quantify the informational value that they represent. This approach has had numerous successes in modelling human information access in a wide variety of tasks such as visual search and concept learning (Nelson & Cottrell, 2007). Renninger, Verghese and Coughlan (2007) predicted eye-movements during a shape learning task by selecting the fixation site that yields the maximum information gain and extended this to predicting sequences of fixations (Renninger, Coughlan, Verghese, & Malik, 2005). In similar work Najemnik and Geisler (2005) developed a model to account for fixation sequences during visual search and found that human results were consistent with locations predicted by the model to maximize the likelihood of correctly identifying the location of the target. In concept learning Nelson and Cottrell (2007) used eye-movement data obtained from a study replicating a classic category learning result (Rehder & Hoffman, 2005). They used a Bayesian conceptual model combined with a sampling function that is capable of selecting a feature of the stimulus that maximizes the likelihood of correct classification. This model provided good quantitative fits to the prediction of eye-movements in both early and late learning. In addition to their success in fitting human behavioural results, support for an information maximization principle is also grounded in neurobiology. Lee and Yu (2000) have argued that the entropy encoded in the neuronal clusters in V1 can be used in a predictive model of saccade landing sites and Nakamura (2006) has shown that signals observed in dorsal premotor cortex of the monkey are consistent with an information-theoretic encoding. Clearly, there is no shortage of research exploring the link between selective attention and its role in successful performance of wide range of cognitive tasks.

The goal of this thesis it to further the understanding of the cognitive mechanisms responsible for attentional control in classification. A shortcoming that many previous models have is that they have neglected to describe is how attention is shifted moment to moment within a single trial. For instance, in most of the Bayesian and connectionist models introduced earlier the level at which the eye-movements or attentional parameters are modelled is at the level of the trial or the task. While the model is capable of fitting human data that reflects the ongoing trial by trial shifting of attention they are not capable of producing data that reflect attention shifts as they occur during the ongoing, continual acquisition of information that occurs in many learning environments. One way in which these category learning models could be modified to reflect the kind of attention observed in Blair et al. (2009) is to have attention update continuously after each successive saccade and have task goals be updated continuously as new information becomes available. While this would provide a higher resolution view of the role that attention has in category learning the resulting model would still be dissociated from important perceptual-motor influences reviewed earlier and would only result in a modest improvement over the discrete attention shifts that occur after each trial. Furthermore, since these models only calculate the value of the information contained at a target location in order to determine where an eye-movement should be placed, they are incapable of incorporating important aspects of human behaviour such as fixation duration differences.

The goal of the present project is to introduce a model which is capable of coherently integrating multiple levels of explanation into a description of human performance during categorization tasks. Insight into the processes of human decision making have demonstrated that cognitive phenomena are best understood when studied as an interacting group of phenomena rather than as isolated systems operating independently. Therefore, the rest of this thesis will be dedicated to developing and testing a new model of category learning that is in essence a real-time simulation of human eye-movements during a categorization task. The model involves the dynamic interaction of feature representation, vision, attention and motor pathways and produces data which is directly comparable to human eye-tracking data.

Dynamic field theory - Theory and Applications

An embodied view of cognition is one which emphasizes the close link between cognition, action and perception, describing them as inherently connected and mutually reliant. Dynamic systems theory is a framework from which these systems can be modelled and describe in a mathematically precise manner. One way in which dynamical systems are particularly useful is in describing the way in which small changes in the environment can have a massive non-linear impact on a systems behaviour. Observation of these kinds of non-linear phenomena in human perception is not a recent development (Necker, 1832). The necker cube is an example of a perceptual phenomenon in which a massive shift in the perception of the orientation of cube occurs despite an absence of any physical change in the orientation of the object. A number of theories have attributed this to the attractor dynamics of the low-level perceptual system (Kornmeier & Bach, 2004).

Applying the concepts from dynamical systems to cognition requires identifying what specific properties of dynamical systems resemble aspects of human cognition. The primary features of dynamical systems that make it such an appealing way to understand human behaviour is the concept of attractor states and their associated bifurcations, along with properties that result in the system moving from one attractor state to another. In dynamical systems terms a fixed-point attractor state occurs when the function describing the evolution of the dynamic system does not change under transformation. In a certain sense, these are the states that the dynamical system tends to move towards in the long run and states which the system can be said to prefer to stay in. Returning to the example of the necker cube, the attractor state would represent the state that the visual system is in when it has settled into one interpretation of the orientation of the cube. Crucial to the development of theories of cognition is the specification of how and why we move from one stable cognitive state to another. The analogy from dynamical systems is the concept of a bifurcation. Bifurcations occur when a dynamical system is placed into a state of instability and can lead to rapid large-scale changes in the overall organization of the system. The concept of bifurcation is especially interesting to the development of embodied cognitive theories. On the one hand, we are interested in understanding how cognitive states resist external influences and settle on a stable interpretation of the environment; however, the concept of a bifurcation explains how the system can respond flexibly by moving into a qualitatively new state when faced with environmental changes.

Dynamic field theory (DFT) developed out of an observation that the traditional dynamical systems approach inadequately addresses several key aspects of higher level cognition such as how representational states influence behaviour (Spencer, J & Schöner, 2003). To incorporate representational states into the dynamical systems perspective, DFT introduces the notion of an activation field which is defined over the metric dimension of the phenomena being modeled. For instance, a simple example of a dynamic field would be a one that represents the frequency of a tone. In this example the scale of the field would represent the tuning of the field in terms of what frequencies it is capable of responding to (see Figure 1). The formation of a peak at a particular location on the field would indicate that a tone of a particular frequency had been identified. Depending on the structure of the model under investigation the input to this field could be the result of input forces from the external environment, other fields internal to model or activation that the field itself is generating through self-excitation. These different sources of activation enumerate several of the possible attractor states that fields can find themselves in, such as a resting state, input-driven state and self-sustaining state (see Figure 2). The resting state of a field occurs in the absence of any activation either external or internal to the field itself, and can be thought of as an initial state of the field. Input-driven states occur when the field stabilizes to form a peak of activity in a particular location due to forces external to the field but will return to the resting-state when external input is removed. The self-sustaining states occur when the interaction between units of a field is sufficiently strong to maintain a localized peak of activity in the absence of external input. The ability to maintain this activity internally, within the field itself is a concept crucial to DFT's claims to providing a solution to the representational challenges of dynamical systems models of cognition (Johnson, Spencer, & Schoner, 2008).

Another strength of DFT is that by developing a unified set of mathematical concepts it is possible to implement a theory as a computational model enabling a strong link between developing cognitive theories and testing them through experimentation.

Figure 1. Frequency Field

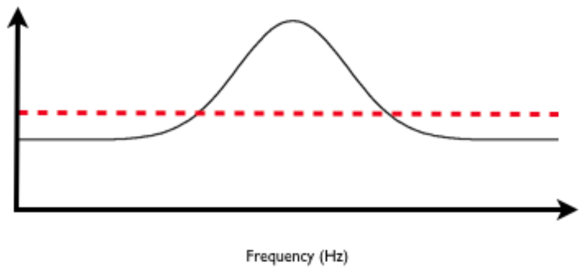

Note. A simple example of a neural field for the detection of the frequency of a tone. The peak of activation represents the frequency that the model is selecting for at that current point in time.

The DFT framework has a rich history of accounting for a diverse collection of experimental results from a variety of fields. In Schöner, Kopecz and Erlhagen (1997) the authors developed a theory of saccadic eye-movement planning that accounts for a number of aspects of human performance such as the variation in saccade averaging that occurs due to the spatial separation of targets, and in similar work Erlhagen and Schöner (2002) developed a model of movement preparation to account for the execution of motor plans. One area in which DFT models have had particular success is in the context of childhood development. Thelen, Schöner, Scheier, and Smith (2001) developed a theory of infant perseverative reaching in the A-not-B task. In this task a child repeatedly reaches for a toy which is consistently hidden at location A; even after the child watches the toy being hidden in a new location B, the child will perseverate in reaching for the toy at location A. Their model explains perseverative reaching as a result of the inability for infants at a certain stage of development to maintain a self-sustained peak of activity at the newly presented location, and also demonstrated how the parameter controlling this behaviour changes over the course of development leading to the disappearance of this behaviour. DFT has also succeeded in modelling the relationship between high-level cognition such as object identification and the mapping that those relationships have with motor surfaces. For instance, Faubel and Schöner (2008) situated a DFT model within an autonomous robot and had it learn to identify a large collection of objects in a small number of views and Spencer, Schneegans, and Hollingworth (2010) developed a model of the relationship between high-level visual working memory and saccade planning.

Figure 2. Dynamic Neural Field Attractor States

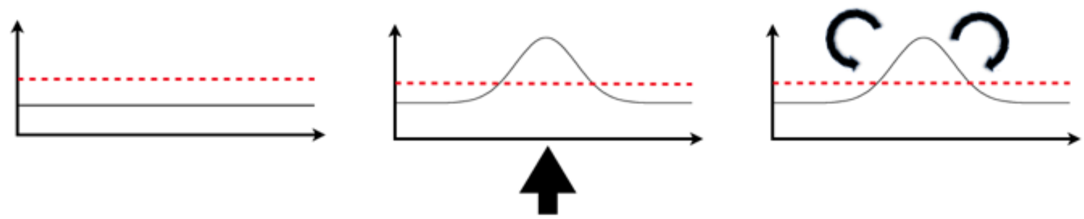

Note. The figure demonstrates three important states of dynamic neural fields. Throughout the course of a simulation any individual field will frequently transition between these three states. In a) the field is in a resting state, and is the state a field moves towards in the absence of input. b) occurs when external input is applied to their field, either from sensors connected to the external environment or from other fields internal to the model. In c) the field is in a state of self-excitation, which occurs when the activation dynamics of a single field are capable of sustaining activation in the absence of external input.

The goal of this project is to apply concepts and methodologies from dynamic field theory to the construction of a model capable of demonstrating a high degree of attentional flexibility. Currently there is no model which is capable of modelling eye-movements and attention at this level within the field of category learning. The proposed model would allow a far more detailed analysis of the relationship between attention and category learning.

Experiments

Two eye-tracking experiments were conducted in order to generate data that could be compared with model output to determine the validity of the overall approach. Experiment 1 is identical to the eye-tracking study conducted by Blair et al. (2009) except that a modification was made to the physical properties of the stimuli. This was done in order to attempt to provide the most direct comparison between human and model behaviour. In Experiment 2, a novel category structure was introduced and eye-tracking data was also collected.

Methods

Participants. In both experiments participants were from Simon Fraser University all who received course credit for participation in the study. There were 34 participants in Experiment 1 and 53 in Experiment 2. All students had either normal or corrected vision and no students reported colour blindness.

Stimuli and Design

Stimuli used in both studies were a three featured colour display with each colour presented horizontally across the midpoint of the screen. On each trial, the features varied between one of two distinct options, resulting is a total of eight possible stimuli. The colours used for Feature 1 varied between Red-Blue, Feature 2 varied between Green-Dark Blue and Feature 3 between Pink-Yellow (see Table 1 for stimuli combinations). The eight possible stimuli were presented many times throughout the experiment such that in every block of 24 trials the subjects would see an equal number of each stimuli. Each feature subtended 4° of visual angle and features were separated by 11°. While each feature was presented in the same location for each subject such that the varying colours did not appear at different locations, the feature location and relevance to categorization was counterbalanced across participants. Table 1 indicates how category membership for the stimuli is defined. A single row consisting of three colours represents a single stimuli. Next to each row the category that the stimuli belongs to is indicated. In both category structures the relevance of a feature to successful categorization varied between stimuli. By varying the relevance of features the goal is to elicit a variety of stimulus responsive attention (SRA) observed in Blair et al. (2009). In Experiment 1 it is possible to separate between the A and B categories by determining the value present in Feature 1. Once the information in this feature has been observed the value of either Feature 1 or Feature 2 will be relevant. In Experiment 2 an exception pattern is defined such that Feature 1 is relevant for the classification of half of the stimuli, but due to the category exception further information is required to distinguish between category B stimuli and the A exception.

The clear hierarchies that exist in these category structures are predicted to lead to attentional biases in the direction of features that have the highest relevance to categorization. Furthermore, due to the ordered nature of the relevance to classification it is predicted that there will be clear temporal patterns to the access of relevant features.

Table 1. Experiment 1 & 2 Category Structures

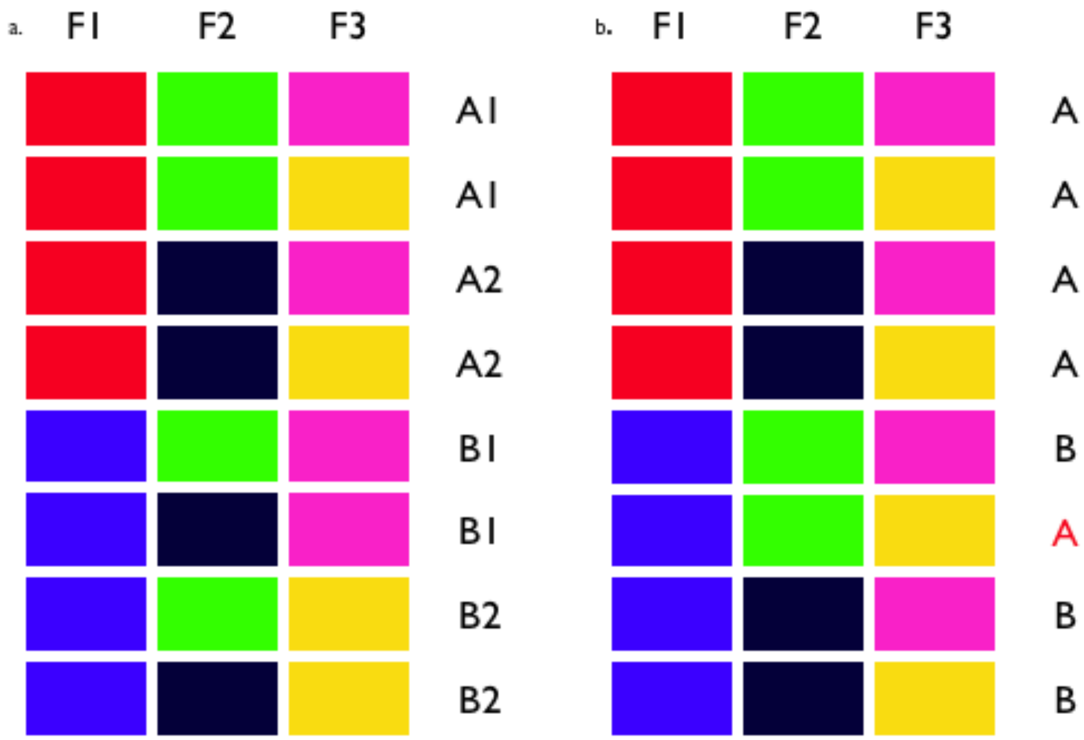

Note. Experiment 1 is indicated by category structure a. In Experiment 1, the relevant features for categorization of A1 and A2 are features 1 and features 2. In Experiment 2, Feature 1 is relevant for all categories. Due to the presence of the category exception, stimuli will require either 1, 2 or 3 features for correct classification.

Procedure

In both studies a trial consisted of fixation to a cross, stimulus presentation, making a categorization response and inspection of any feedback that was provided. The fixation cross was presented at a random position on the screen in order to minimize perceptual biases towards any particular initial position for the gaze such as at the central location of the display. Once the participants had detected and fixated at this location they pressed a button on the joystick that indicated their readiness to start the experiment. They were then presented with a stimulus and were required to determine which category it belongs to by studying its features. Participants were given as much time as they felt necessary in order to study the features. Their response was made by pressing a button on the joystick that corresponded to the selected category label. Once they made their selection the screen would briefly flash red if they had made an incorrect response or green if they had correctly classified the stimulus. The stimulus was maintained on the screen while feedback on their performance was provided by showing the category label they selected in the top left corner of the screen and the correct category label in the top right. They were also allowed to study the feedback portion of the trial for as long as they required. If a participant did not master the task by making 24 correct classifications in a row correct the experiment would terminate after the 200th trial. If the participant did reach the criterion of 24 in a row correct the experiment would move into a second phase for a further 72 trials. Phase 2 was identical to phase 1 in all aspects other than that feedback was removed for the remaining trials. After classification, participants would see a flash of grey for each trial and no information about the correct category was supplied in the upper right corner of the screen.

Equipment

A Tobii X120 eye-tracker sampling at a rate of 120Hz was used to record raw gaze data. In order to identify fixations a modified dispersion threshold algorithm was used with a threshold of 75 ms and 28 pixels (Salvucci & Goldberg, 2000). A fixation was considered to be located on a feature if it fell within 150 pixels to the left or right of the feature and 200 pixels above or below the feature.

Results

Experiment 1

In Experiment 1 17 out of 34 participants reached the learning criterion at an average of 111 trials. Accuracy among these participants was generally high with learners obtaining 71% accuracy before the criterion of 24 in a row correct and 91% after. Consistent with the increase in accuracy, a decrease in reaction time was also observed with post criterion reaction time dropping from 2715ms to 2329ms. The gaze data for 4 participants was excluded from the analysis, as the eye-tracker was able to collect less than 60% of their gaze throughout the experiment. In order to investigate patterns of attentional allocation that would indicate a stimulus-based profile of attention, an analysis of a trials mean total fixation duration to features grouped by category was conducted. Because the category structure separates between category A and B stimuli in terms of their relevant feature dimensions it was possible to collapse the analysis across categories. An ANOVA was conducted with feature and category included as within subjects variables. A Feature X Category interaction, $F(2,77) = 11.13$, $p < .01$ was observed confirming results from Blair et al. (2009) which showed category dependent fixation durations. Strangely, multiple comparisons showed that participants fixated Feature 1 less than Feature 2 or Feature 3 a finding which is at odds with the overall relevance of this feature and is also contrary to what was reported in Blair et al. (2009).

In order to investigate the temporal regularities present in the human data, we conducted a fixation probability analysis of the development of attentional allocation over the course of a trial. This analysis consists of separating a trial into equal size bins so that trials of different lengths can be compared. The probability that each feature will be fixated is then calculated for that bin. The shading on the line shows the standard error of the fixation probabilities for that bin. This analysis provides a detailed look at how attention shifts in response to the information that is gathered during the task. Fixation probability was calculated separately for each category, again with the As and the Bs grouped together due to their shared feature relevance. The results demonstrate an ordered temporal pattern to the allocation of attention (see Figure 4). On category A trials, Feature 1 initially has the highest probability of being fixated followed by an increase in Feature 2 fixation probability which is consistent with this Feature 2's relevance to classification. The pattern of allocation is not as clear for category B stimuli but is nevertheless consistent with participants ordering their fixations based on informativeness, with Feature 1 and Feature 3 being equally probable of receiving a fixation early in the trial and Feature 3 being more likely to receive fixation later in the trial.

Figure 3. Pre vs Post Criterion Accuracy for Experiment 1 and Experiment 2

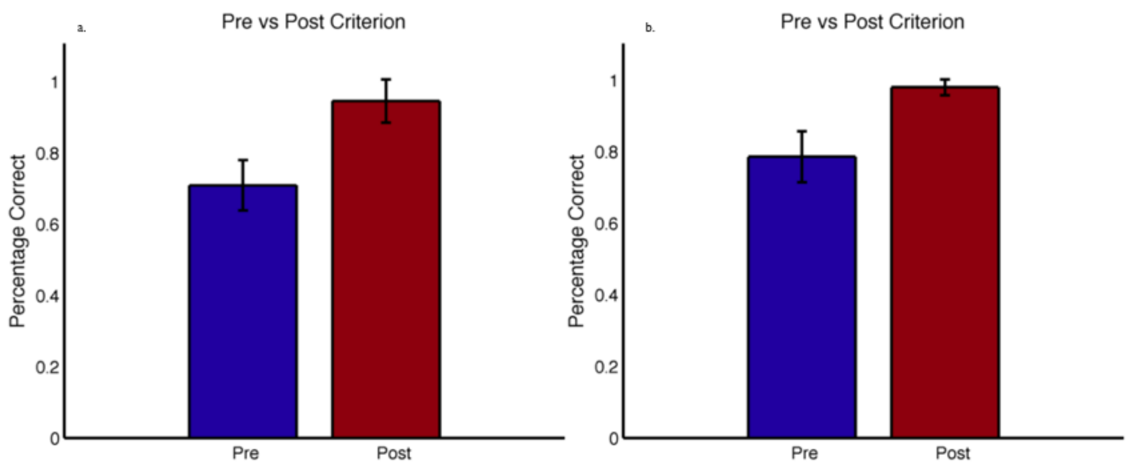

Note. Bars show the mean accuracy achieved by participants before and after reaching the criterion of 24 in a row correct. For a) mean accuracy was 71% and 91% and for b) the mean accuracy 83% and 91%.

While these results provide an indication that participants were utilizing category specific attentional allocations, the results were less robust than and contradict some results reported in the original Blair et al. (2009) study. In order to investigate possible reasons for this, a fixation probability analysis was conducted on stimulus locations to investigate the possibility of a location bias. Figure 6 shows the fixation probability to location averaged across all categories. The probability of fixating the central location is much higher at all stages of a trial regardless of category of feature relevance despite counter-balancing of location relevance and randomization of the initial gaze position. A central location bias would add a definite confound as it is difficult to determine whether a fixation to a feature was due to its spatial location or its feature relevance. Observing heat-maps of fixations over the course of learning provided further confirmation of a central location bias. Furthermore, heat-maps showed that some participants who had mastered the task appeared to look at nothing but the central location late in learning when only correct responses were being made. This leads to the speculation that participants may have been using peripheral vision to detect the value of distally located features.

Several possible improvements could be made to eliminate the effect of a location bias. If the speculation that the central location bias is due to the fact that distally located features can be peripherally detected, then reducing the features size may help in making discriminating features more difficult. The features used were also highly salient and designed to be maximally discriminable from each other. Designing stimuli that are more similar and difficult to distinguish may also help to make peripheral detection difficult. It will also be important to thoroughly pilot the effect that the salience due to stimulus colour and size has on participants ability to detect the stimuli.

Figure 4. Fixations to Features by Relevance

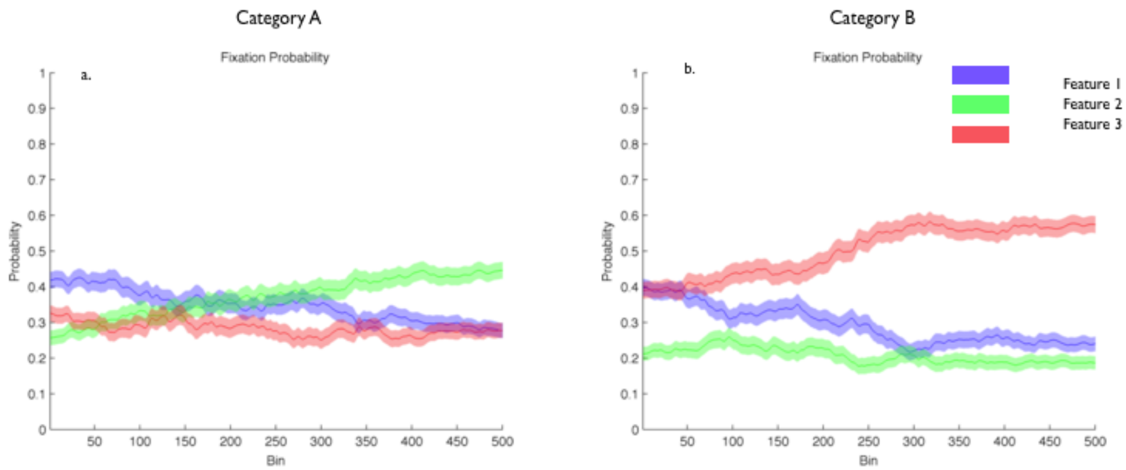

Figure 4. The lines show the probability of a feature being fixated at a particular point in a trial for Experiment 1. Data show the last 72 trials of the experiment after a participant had made 24 consecutive correct responses. Fixation probabilities demonstrate a temporal pattern in the allocation of attention reflective of the hierarchical nature of the category structure.

Experiment 2

In Experiment 2, the same approach to the analysis was extended from Experiment 1. Participants found this task difficult as only 20 out of 53 reached criterion and were included in the analysis. Figure 3 shows that the same pattern of pre-post criterion accuracy was observed in this task, with participants improving their accuracy from 83% to 91% and participants reached the learning criterion after an average of 129 trials.

A fixation duration analysis was conducted with a ANOVA using feature and category as within subjects variables. The levels of the category variable selected were category A, B and A-exception because of the prediction that each would require a unique pattern of attentional allocation. Contrary to what had been predicted no Feature X Category interaction was observed, $F(4,179) = .22$, $p = .92$. Also, the predicted Feature X Category interaction for mean total fixation duration was not confirmed. A fixation probability analysis for Experiment 2 was also conducted but no interaction effect was observed. As in Experiment 1 the perceptual bias to the central stimulus location was very large (see Figure 5).

The goal of these studies was to provide experimental results that would improve upon the ecological validity of (Blair et al., 2009) with respect to model comparisons. Arranging the stimulus in a row presents a convenient solution as this is the manner in which the model receives input. However, this spatial arrangement of features does represent a challenge due to strong biases that are present from activities such as reading that occur from left to right, or from a very strong bias towards fixating objects that are centrally located. While the results were not as robust as anticipated due to location confound, model results will be compared with predictions from the data observed in Blair et al. (2009) as well as the constrained results from Experiment 1 making the assumption that with the appropriate modifications the colour based stimuli are capable of eliciting stimulus responsive attention.

Figure 5. Fixation Probability to Features by Location

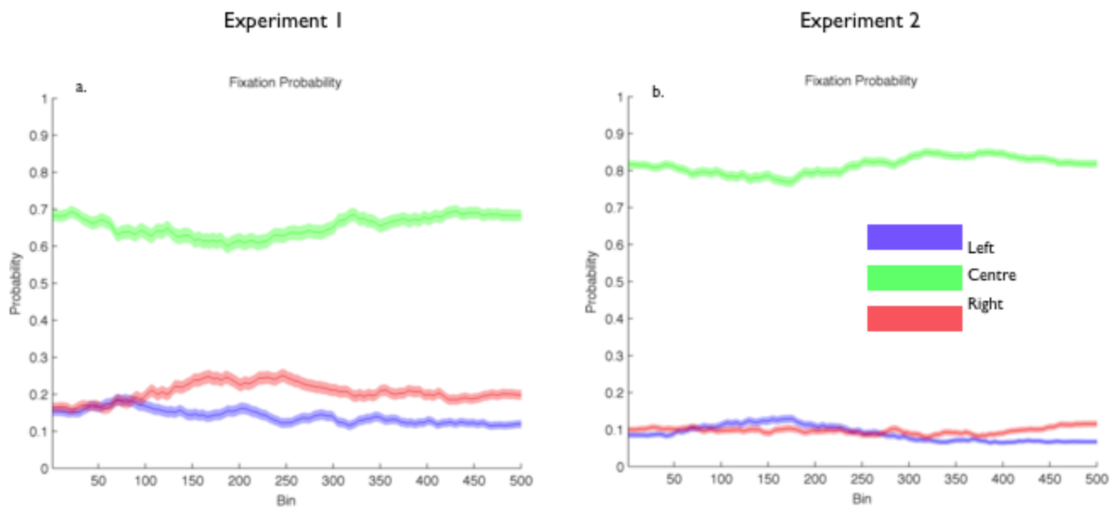

Note. This demonstrates the strong bias towards the central feature of all categories. The perceptual bias towards the central features adds significant noise to the results making between category attention results difficult to detect. The bias towards the central location is slightly stronger in Experiment 2 (b) than Experiment 1 (a).

Unnamed Computational Model (UCM)

Through the coupled interaction of several neural fields, UCM is capable of generating fixation data in the context of standard category learning tasks. The model utilizes a saccade timing mechanism developed in unpublished work by Spencer et al. (2010). The current model is unique in that no prior model has been able to simulate the temporal regularities that are present in the allocation of attention within the category learning paradigm. By measuring the activity on the saccade motor field as well as the resulting ballistic eye movement, temporal data that is directly analogous to a sequence of human fixation is generated. Furthermore, the fixations that are observed are a direct result of the low-level perceptual influences on attention as well as the top-down pressures from what has been learnt about the category structure. Figure 6 shows the basic structure of the model as well as the interactions between the components. The following section outlines a formal mathematical description of the generalized form of dynamic neural fields, as well as the equations describing the fields developed in this model.

Model Description

The general form of a dynamic neural field is given by (see Appendix of (Erlhagen & Schöner, 2002)):

$$\tau\dot{u}(x,t) = -u(x,t) + h + S(x,t)... \tag{1)}$$

Where $\dot{u}(x,t)$ represents the rate at which a particular field site x is changing at time t and is negatively proportional to the amplitude of the field site $u(x,t)$. $S(x,t)$ is the external input being applied to the field, and is generally provided in the form of a gaussian given by:

$$S(x,t) = \sum_i S_i(t) \exp\left\{ -\frac{(x - x_i)^2}{2\sigma^2} \right\} \tag{2)}$$

In all cases external input is combined with a parameter c which defines the input strength to the field. The parameter is omitted to increase the readability of the equations. A resting state, h, is also included and yields a stable attractor to the dynamical system described by 1) when no external input is applied.

The interaction of neurons within a field is modelled as local excitation and global inhibition. Roughly speaking, this results in neurons that are proximally located having greater excitatory force on each other than will neurons located further apart. The gradient of decline in excitation between field sites is given by:

$$w_u(x - x') = k_u \exp\left\{ -\frac{(x - x')^2}{2\sigma_u^2} \right\} \tag{3)}$$

In order to simulate the non-linear nature of neural activation a sigmoidal non-linearity is added to the interaction kernel:

$$\tau\dot{u}(x,t) = -u(x,t) + h + S(x,t) + \int w(x - x') f[u(x',t)]dx' \tag{4)}$$

The purpose of adding a sigmoidal non-linearity is to simulate the properties of biological neurons which generate activity only once they have reached a certain threshold.

Figure 6. Structure of the components of the simulator.

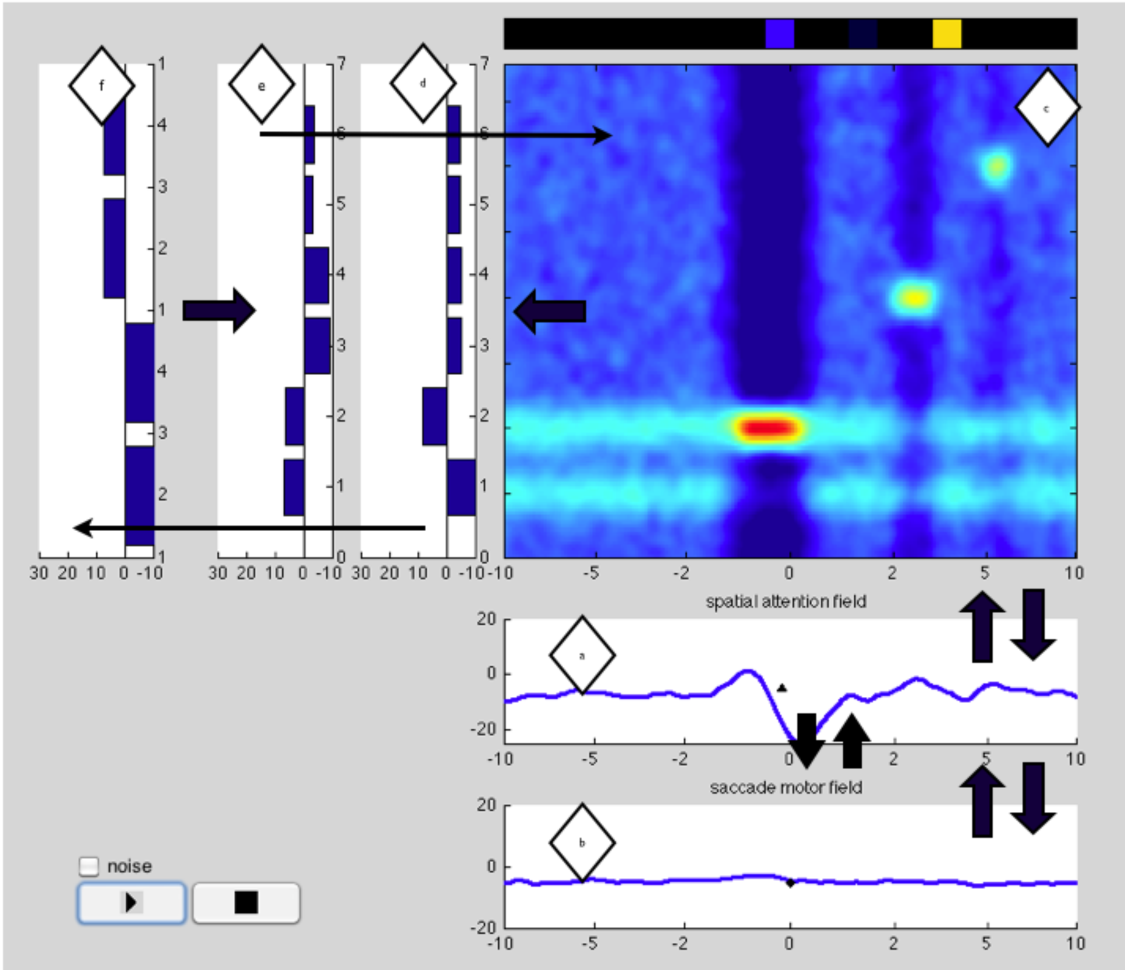

Note. Arrows indicate direction of activation between fields. See text for description of the fields. a) Attention Field b) Saccade Motor Field, c) Visual Field d) Feature neruons and e) Feature Hypothesis neurons and f) Categorization neurons

Attention field

The equation for the attention field (Figure 6, a) is given by:

$$\tau\dot{a}(x,t) = -a(x,t) + h + V(x,t) + S_{sacc}(x,t) + IOR(x,t) + INH(fovea,t) + P_{(preshape)}(x,t)... \\ + \int w(x-x')f[u(x',t)]dx' \quad 5)$$

The external input provided to this field arrives from several sources. $V(x,t)$ is an excitatory connection to the field that arrives from the visual field (see below for details). Field a also receives an excitatory connection from the saccade-motor field, and provides an additional boost to the spatial location that is currently being selected for an eye-movement. Inhibition of return is also incorporated into this model by applying an inhibitory influence to the attention field at the most recently fixated feature location (see (Satel, Wang, Trappenberg, & Klein, 2011) for evidence supporting this interpretation of IOR). In order to influence eye-movements away from the currently fixated location once feature information has been collected an inhibitory influence is applied to the field at the foveal location driven eye-movements away from the currently fixated location.

Saccade motor field

This field controls the programming of a saccadic location (Figure 6, b). When a certain threshold of activity has been reached the models field of view is shifted to centre at approximately the location indicated by saccade field's peak activity. The field dynamics are given by:

$$\tau \dot{s}(x,t) = -s(x,t) + h + a(x,t) + r(x,t) + \int w(x-x')f[u(x',t)]dx' \qquad 6)$$

where $a(x,t)$ is the excitatory input provided from the attention field and $r(x,t)$ is a global dampening of the field that results from a saccade target being selected. The strong global dampening is a feature of the saccade field that strongly pressures the field to select a single location as saccade target.

Visual Field

The dynamics in this field are unique and diverge slightly from the generalized description of field dynamics introduced earlier. The visual field v (Figure 6, c) is modelled as a two dimensional field where the y-axis represents a feature space (colour) and the x-axis represents a retinotopic spatial field of view. The modified field equation in two dimensions is then:

$$\tau\dot{v}(x,y,t) = -v(x,t) + h + a(x,t) + k(y,t)... \\ + \int\int w(x,x'-y,y')f[v(x',y',t)]dx'dy' \qquad 7)$$

with excitatory connections to dimension x of the field originating in the attention field and an external excitatory connection to the y dimension of the field from the feature hypothesis pathway. The excitatory activation propagating to the feature dimension of the visual field acts as a way of supplementing the activity of features that are currently represented.

Hebbian feature-learning pathway

The hebbian learning module (Figure 6, d,e,f) is responsible for extracting features that have been detected from the environment and learning associations between features and categories. One motivation for using a Hebbian-learning mechanism is to overcome challenges to the error-driven learning assumptions made by several influential models of category learning (Blair, Walshe, Barnes, & Chen, Submitted).

Features and categories are represented as collections of neurons as opposed to fields, due to the discrete nature of both the features and categories.

Although the same dynamic Hebbian processes could be applied to learn associations between two continuous valued neural fields, the discrete neural representation is simpler and no less valid for the scope of this project.

The dynamic activation for the neurons is very similar to the general neural field representation, although without the intra-field dynamics. The equation is:

$$\tau \dot{ft}_{feature}(t) = -ft_{feature}(t) + v\Big(t, y\big(ft_{location}\big)\Big)\ldots \qquad 8)$$
$$+ft_{inh}(t) + ft_{exc}(t)]$$

The term $v\big(t, y(ft_{location})\big)$ is provided as activation to a specific feature neuron as long as gaze has been directed to this feature on the visual field. $t_{location}$ indicates that the excitatory input is coming from a region surrounding $ft_{location}$ on the visual fields feature dimension. ft_{inh} describes an inhibitory connection between the two values of a feature. In these tasks the participants understand that the presence of one feature necessarily implies the absence of the feature paired at its location. In order to simulate this, an inhibitory influence from the detected feature is applied to the feature paired at the same location. This inhibitory force is applied as the negative of the activation present on the detected feature. Effectively, the feature nodes function to represent what features have been visually identified from the environment, passing this activation along to the category nodes via a matrix of associative connections between features and categories. In order to simplify the modelling of attentional learning, features could only be activated by visual detection before corrective feedback was applied. Simulation of model behaviour resulting from the feedback phase is left for future developments of the UCM.

The excitatory connection from $c_{categories}$ originates from the activity present on each category node, weighted by the learned connection strength between feature and category. The activation dynamics for the category nodes are similar:

$$\tau \dot{c}_{categories}(t) = -c_{categories}(t) + ft_{feature} + f_{category\_exc} + f_{category\_inh} \quad 9)$$

All feature nodes have weighted excitatory connections to each category node, and the input to the category node is defined as $ft_{feature}$.

In order to pass activation representing the the hypothesized values for which features are most likely to contain information necessary for classification a separate cluster of neurons is defined which receive excitatory input from the category nodes. The activation to these nodes is generated by multiplying the activity present on the category nodes by the matrix of learnt associations between categories and features. The equation for this field is:

$$\tau \dot{h}_{feature}(t) = -h_{feature}(t) + v\left(t, y\left(h_{location}\right)\right) \ldots \quad 10)$$
$$+h_{inh}(t) + h_{exc}(t)$$

An additional feature of these nodes is that pairs of hypotheses neurons share activation such that two nodes representing the same location in space are summed together to supplement their activation. The result of this is to have the hypothesis neurons provide the largest boost to the location, a combination of two features, whose sum of feature activation is the largest.

Dynamic Hebbian Learning

In order to learn associations between features and category, a dynamically updated version of a Hebbian learning rule was implemented:

$$\tau \dot{M}_i = \left(-M_i(j,t) + g\left(c(j,t)\right)\right) \cdot g\left(ft_i(t)\right) \quad \text{11)}$$

This process increases the connection strength between feature and category neurons that are co-activated. $M_i(j,t)$ indexes the weight between between feature node i and category node j. Weights between category and feature nodes are strengthened when, $f\left(ft_i(t)\right) > 0$ and $f\left(c(j,t)\right) > 0$ (Sandamirskaya & Schöner, 2010). This describes the learning rule that implements an associative mechanism such that the strength of connection between feature and category nodes is increased when positive activity is detected on both nodes.

Model Simulations

In order to compare model behaviour with human participants 15 simulations were conducted. All simulations were conducted on the category structure indicated in Table 1a and the spatial location of Feature 1 was counter-balanced to appear in all three locations. Two simulations were excluded from the analysis because the model did not learn the category structure by the 200th simulated trial. A trial for the model consisted of a starting position at one of two randomly selected stimulus locations between the features. The model was then free to fixate for 600 time steps (ts) before feedback was applied for 100 ts. This simplification was made in order to reduce the complexity of the modelling task while also maintaining important characteristics of behavioural experiments. During the feedback phase a boost was applied to the category node associated with the correct response, simulating the corrective feedback provided to human participants. The stimuli used were meant to roughly simulate human colour discriminations and had peak input responses separated by equal distances on the feature dimension of the visual field. One of the main strengths of this model is that it produces a type of output that is directly comparable to human participants. As the model itself produces observable fixations, there is no need to make indirect observations about the overt allocation because the model itself overtly allocates. Therefore, in order to compare this model with human behaviour, similar techniques that were used for the human data will be applied to model output.

The simulations demonstrate that the model is capable of producing a rich collection of data that can be compared with human participants. While the model does not incorporate a decision process which would initiate a self-directed response to be made, responses were identified post-hoc by identifying the category node that had the largest activation prior to receiving feedback on a given trial. Similarly to the humans, we identified a criterion of 24 in a row correct. Using this method the model learnt the category structure in a mean of 127 trials. In order to conduct an analysis of the models eye-movements to features, a fixation was identified as being located on a feature if was located with in 10 field units of the centre of the feature. Fixation durations were calculated as the number of time steps that occurred between the time a saccade landed on a feature and the time that the saccade moved away from the feature.

Temporal ordering of fixations were analyzed by conducting a fixation probability analysis as was conducted with human participants. The analysis revealed that the model is capable of deploying its attention in an analogous way to that observed in Experiment 1 and seen more clearly in Blair et al. (2009) (see Figure 7 for original results). The trials selected for this analysis were the last 72 trials the simulated experiment and all trials were in the post criterion phase after the model had responded 24 in a row correct.

Figure 8 shows that after learning the model is highly selective for Feature 1, the most diagnostic feature, then selecting either Feature 2 or Feature 3 depending on its relevance to correct classification. That is, for stimuli belonging to category A the model first fixates Feature 1 and then Feature 2, while on category B trials the model is most likely to fixate Feature 3 after first fixating Feature 1. This pattern of attentional allocation represents the optimal pattern of attentional allocation as it requires the fewest fixations in order to the classify the stimulus. Comparing the models post-learning behaviour with the first 72 trials of the simulation shows a striking

Figure 7. Allocation of attention within trials in Blair et al. 2009

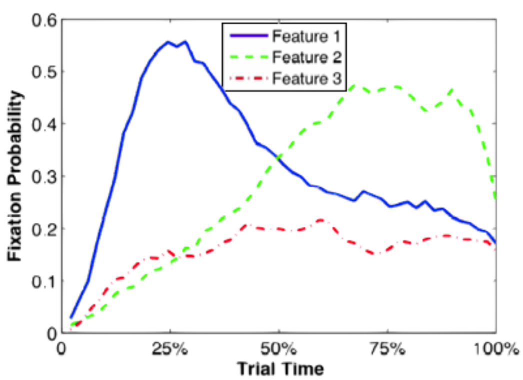

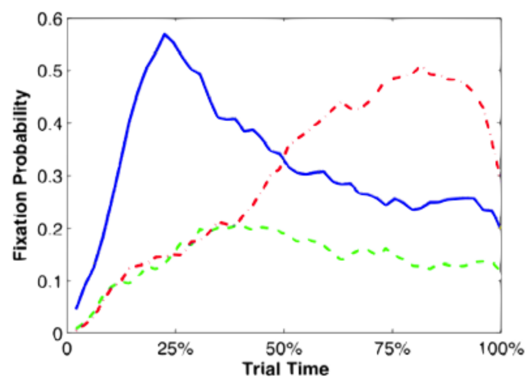

Note. Lines show probability of fixation to each feature. Data were averaged across the final 72 trials of Experiment 2 for Categories A1 and A2, shown in the left graph, and for B1 and B2, shown in the right graph. Model results show a very similar temporal pattern of attentional allocation.

Figure 8. Model Fixation Probabilities after Learning

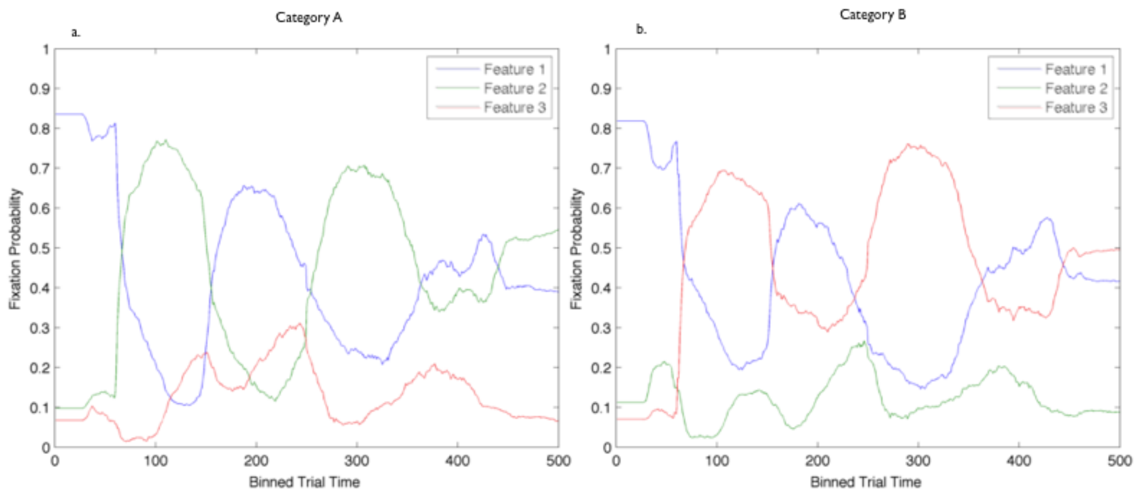

Note. In a) the model initially selects strongly for Feature 1 which is always diagnostic followed by Feature 2 which is diagnostic of this category. In b) Feature 1 is selected for first followed by an increase in the probability of Fixating Feature 3.

Figure 9. Model Fixation Probabilities for the First 72 Trials

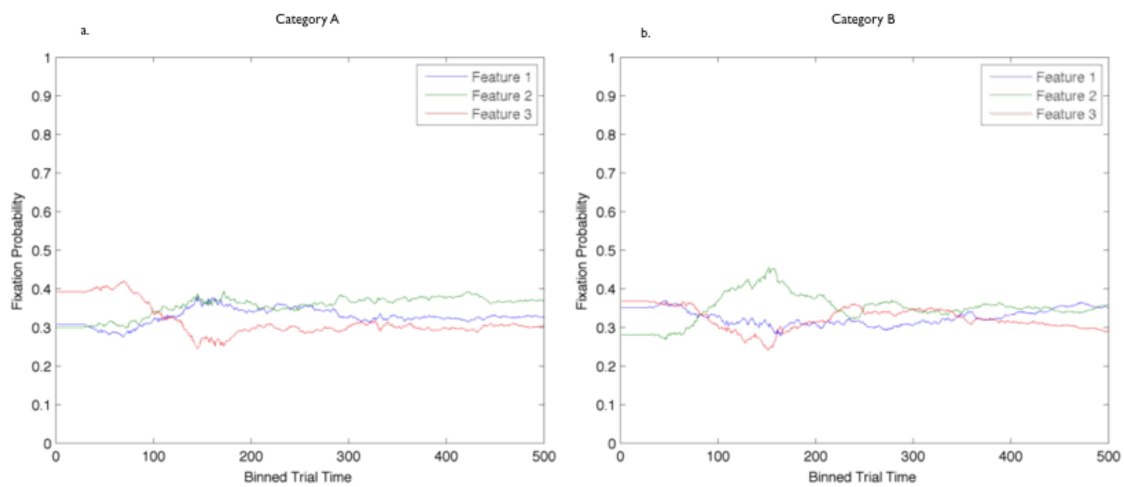

Note. There is almost no difference between a) and b) illustrating that the model does not show category differences in its temporal regularities before learning the category structure.

difference in that no strong temporal regularities tied to feature relevance are observed which is consistent with the models lack of knowledge about feature relevance at this stage in learning (see Figure 9). One difference between the model and human behaviour is that the model appears to oscillate between fixations to the relevant features. This is a result of the model not having any control over when it has acquired enough information to form a classification decision. The model is driven to fixate features until the trial time expires and therefore selects features that it has identified as most relevant on the current trial. In future extensions of the model a decision process will be implemented such that the model stops fixating and makes a response once it has detected that sufficient information has been gathered.

An analysis was conducted of mean fixation duration to stimulus features grouped by category (see Figure 10). In Experiment 1, no significant differences were observed; however, in Blair et al. (2009) strong category specific fixation duration differences were reported with the same category structure simulated in the model. The results from the simulations clearly replicate the findings from the original study. Feature 1 had relatively high fixation durations for both category A and B stimuli with Feature 1 and Feature 2 fixation durations being highly dependent on whether the Feature was relevant for classification. The results from the fixation duration analysis closely parallel results observed in the original study (see Figure 11).

Figure 10. Allocation of attention within a trial.

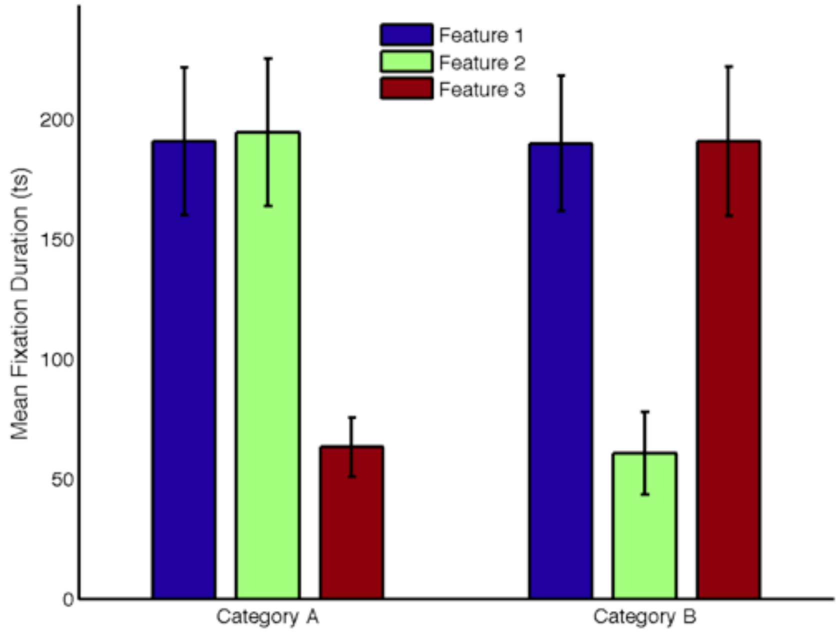

Note. Bars shows the mean total fixation duration that was allocated to a feature on a trial. This demonstrates that more time is spent on features that are relevant to classification. In Experiment 1, Feature 3 is irrelevant for Category A and Feature 2 is irrelevant for category B

Figure 11. Mean Fixation Duration to Features by Category (Blair et al. 2009)

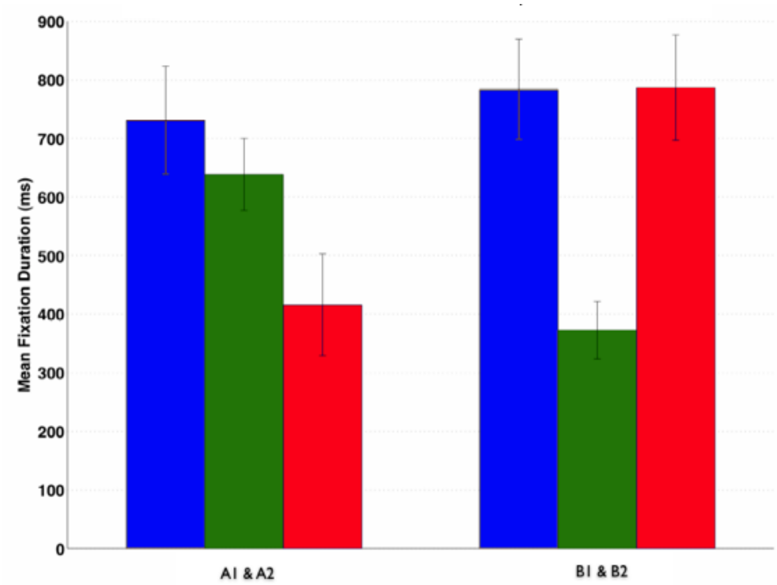

Note. Results reported in Blair, Watson, Walshe & Maj, 2009. (Used with permission) shows fixation durations to features on a trial are dependent on a features relevance to classifying a stimulus. Category structure is identical to that used in Experiment 1.

Discussion

The present work introduces a new model that is capable of predicting eye-movements in a category learning task. Early results indicate that the model deploys its attention in a manner such that it strongly selects for stimulus features that are relevant for classification in a particular context. The project utilizes techniques from Dynamic Field Theory to construct a model that incorporates perceptuo-motor as well as cognitive influences of behaviour into an integrated framework. The project has shown how a collection of softly-assembled components interacting in real-time can generate flexible attentional control in a way that has not been captured by other models in the field of category learning. Most models of attentional selection in these task have taken the perspective that attention to features is shifted once a trial is concluded and after a participant has made a decision. Studies from eye-tracking in category learning, as well as research in other fields, has shown that attentional deployment is far more subtle than is captured by the simple picture of dimensional relevance shifted over trials. Humans are constantly shifting the focus of their attention in response to environmental cues that they discover in an ongoing processes of information acquisition. The explanatory target of this model has been to show how that this behaviour can emerge from the dynamic interplay between the low-level perceptual motor systems and the higher-level learning structures that have here been implemented here as a dynamic Hebbian process. The model implements an attentional representation that is initially driven solely by the salience properties early in learning and over time, this representation becomes highly influenced by what the model has learnt about the category structure, much like human participants. The approach taken in this thesis can be seen as one implementation of an embodied cognitive agent. Embodiment is the perspective that cognitive states cannot be understood as being in a separate and independent relationship with the environment. Rather, it understands that there is a reciprocal relationship between the environment and cognition. In this model the dynamically changing scene that is being presented to the model is a direct result what the model has selected for its source of input, which in turn has a large scale impact on how it updates it's theory of its environment.

van Zoest, Hunt, & Kingstone (2010) have stressed the importance of understanding the specific time course of stimulus processing. They show that tasks in which response is measured early in processing may be heavily influenced by low-level, bottom-up processes reliant on stimulus salience while tasks in which the behaviour is allowed to develop more slowly can incorporate more complex information such as task-goals, context and prior knowledge. Category learning tasks of the type under investigation here are certainly of the latter variety. In these tasks participants have adequate time to explore various options and fully probe their prior experience in solving the task. While it is tempting to abstract away from the underlying stimulus properties, these aspects of a stimulus can have a great influence on task performance and can partially explaining fixations to high salience task-irrelevant features (Hickey, Van Zoest, & Theeuwes, 2010). A particular strength of the UCM is that its current behaviour is highly dependent on the stage in processing that it finds itself. Early in learning, the model is to a large degree driven by salience with a switch later in learning to knowledge driven allocational of attention. Within a single trial the importance of time to understanding the behaviour of the UCM is also clearly apparent as the attention that is applied to a stimulus feature is directly connected with what information the model is currently representing about the category structure.

Models such as the UCM that attempts to integrate multiple cognitive factors into an explanation of behaviour have unique benefits in that they naturally produce a rich set of data that can be compared to human participants. For instance, in the project currently under investigation we applied the model to the prediction of eye-movements in a category learning task. This type of measurement of the model was selected because of the availability of eye-tracking data that can be used to infer the hidden properties of attentional allocations in humans. However, in the UCM the saccadic eye-movement is a direct result of the ongoing activation present in the attentional field and can be directly observed. By modelling the attention-motor link in this dynamic manner allows observable manifestations of eye-movement behaviour such as fixation duration and saccade latency to be predicted from the underlying attentional control mechanisms. While this project focused solely on predictions that the model makes regarding the temporal nature of the eye-movements and fixation durations, future development of this model could be expanded into a framework to attempt to capture all aspects of eye-movement behaviour in these tasks such as saccade latency as well feature dependent fixation durations.

An assumption that is made in many popular category learning models, is that the shifting of attention is driven by errors that are detected on the current trial. Recent work has called this error driven assumption into question. Blair et al. (Submitted) used a novel measure of error-bias which measures the degree of attention shifting that occurs on correct and incorrect trials. The authors conducted simulations of two influential exemplar-based categorization models on numerous category structures and confirmed their prediction that the models would show a strong error-bias. Results from a number of experiments using these category structures demonstrated a different pattern of attentional shifting in actual human performance as the authors reported no specific difference in the error-bias for correct and incorrect trials. These results indicate that the error-driven assumptions that underlie numerous models of categorization may not be strongly founded in human performance. In the UCM the magnitude of the attentional shift is in no way connected with the model correctly or incorrectly classifying the stimulus. The learning mechanism that we implemented is a dynamic version of a Hebbian learning rule which simply associates the correct category node with the features that were observed on the current trial. This type of learning mechanism has been suggested by Blair et al. (Submitted) as alternative to the traditional error-driven approaches, and the UCM provides confirmatory evidence that such a mechanism produces attentional shifts analogously to human participants.

The capability for a model of category learning to produce a full range of eye-movement data is both new and exciting. The analysis of fixation probabilities represents a test case for the quality of the comparison that can be made between model output and human data. However, the current state of the model should be regarded to a large extent as a prototype for a model which extends more expansively into describing human behavioural processes. One aspect of the current model that was left undeveloped is a process that initiates a decision to be made and for feedback to begin. The lack of such a process limits the ability to make qualitative predictions about human behaviour as overt attentional data such as fixation durations, response times and total fixations are biased due to unrealistic trial lengths. Implementing a decision process would also allow the modeller to look much more closely at the appropriateness of the learning mechanism selected, as the models learning curve and response times can be directly compared to human responses. One candidate for such a mechanism would be to implement a decision node that initiates a response once it has reached a certain threshold. The output value of the decision node could be designed to represent both the amplitude of the category nodes as well as the uncertainty of the correct category. As the model fixates through the scene and gathers information about the category structure the decision node will reach threshold and the model can internally generate a signal that it has obtained sufficient information to respond.

One other aspect of the category learning task that has been neglected by current models as well as the UCM, is the feedback portion of category learning tasks. In Watson and Blair (2008) the authors showed a number of results that demonstrate the crucial importance of the feedback portion of the trial to learning the task in general. While the UCM does not assign any unique status to the feedback phase other than to activate the Hebbian learning processes, an exciting avenue for future development of this model would be to look into specific influences of post-response learning.

In addition to development components of the model that lead to a more accurate match with human performance, there is significant room for the development of a connection between the pathways outlined in the model and known neurological models of attentional control. There has been a plethora of research into understanding the neural mechanisms responsible for attention guided eye-movements that have incorporated both bottom up and top-down influences. For instance, a pathway that involves the superior-colliculus and pulvinar has been identified as possibly responsible for the maintenance of attention to features of the visual field as well as the generation of saccade targets (Berman & Wurtz, 2011). This pathway is very similar to the visuo-motor pathway that the UCM models as a visual, attention and saccade field (see Figure 4 a,b,c). One challenge to developing the link between the UCM and underlying neurobiology will be to show a link between a higher-level associative mechanism and the lower level selective attention mechanisms that are predominantly influenced by stimulus salience.

The main contribution of this thesis has been to make the first steps in developing a computational model capable of generate real-time eye-movement data during category learning tasks. By integrating a saccade generation system with an attentional system that is influenced by both top-down learned attention as well as stimulus salience the model is capable of simulating saccadic eye-movements that reflect human performance in high-level cognition. While the UCM is currently in an early phase of development, the connection between human eye-movement data and model output is encouraging. By developing the capacity for the model to more closely simulate human tasks the model will be in a unique position to describe aspects of human performance in category learning tasks that have not been previously attempted.

References

Ashby, F. G., & Alfonso-Reese, L. A. (1998). A neuropsychological theory of multiple systems in category learning. Psychological Review, 105(3), 442.

Berman, R. A., & Wurtz, R. H. (2011). Signals conveyed in the pulvinar pathway from superior colliculus to cortical area MT. The Journal of Neuroscience, 31(2), 373.

Biederman, I., & Shiffrar, M. M. (1987). Sexing day-old chicks: A case study and expert systems analysis of a difficult perceptual-learning task. Journal of Experimental Psychology: Learning, Memory, and Cognition, 13(4), 640.

Blair, M. R., Walshe, C., Barnes, J. I., & Chen, L. (Submitted). Rethinking the role of error in attentional learning.

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

Busemeyer, J. R., Dewey, G. I., & Medin, D. L. (1984). Evaluation of exemplar-based generalization and the abstraction of categorical information. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10(4), 638.

Clark, A., & Toribio, J. (1994). Doing without representing? Synthese, 101(3), 401-431.

Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. Psychological Review, 109(3), 545.

Faubel, C., & Schöner, G. (2008). Learning to recognize objects on the fly: A neurally based dynamic field approach. Neural Networks, 21(4), 562-576.

Foulsham, T., Cheng, J. T., Tracy, J. L., Henrich, J., & Kingstone, A. (2010). Gaze allocation in a dynamic situation: Effects of social status and speaking. Cognition.

Hickey, C., Van Zoest, W., & Theeuwes, J. (2010). The time course of exogenous and endogenous control of covert attention. Experimental Brain Research, 201(4), 789-796.

Johnson, J. S., Spencer, J. P., & Schoner, G. (2008). Moving to higher ground: The dynamic field theory and the dynamics of visual cognition. New Ideas in Psychology, 26(2), 227-251.

Kornmeier, J., & Bach, M. (2004). Early neural activity in Necker-cube reversal: Evidence for low-level processing of a gestalt phenomenon. Psychophysiology, 41(1), 1-8.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. Psychological Review, 99(1), 22.

Kruschke, J. K. (2001). The inverse base-rate effect is not explained by eliminative inference. Journal of Experimental Psychology: Learning, Memory, and Cognition, 27(6), 1385.

Kruschke, J. K., & Johansen, M. K. (1999). A model of probabilistic category learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 25(5), 1083.

Land, M. F., Mennie, N., & Rusted, J. (1999). Eye movements and the roles of vision in activities of daily living: making a cup of tea. Perception, 28(4), 1311-1328.

Lee, T. S., & Yu, S. X. (2000). An information-theoretic framework for understanding saccadic eye movements. Advances in neural information processing systems, 12, 834ñ840.

Logan, G. D. (2004). Cumulative progress in formal theories of attention. Annual Review of Psychology, 55, 207-234.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A Network Model of Category Learning. Psychological Review, 111(2), 309.

Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. Journal of Experimental Psychology: Learning, Memory, and Cognition, 8(1), 37.

Medin, D. L., Dewey, G. I., & Murphy, T. D. (1983). Relationships between item and category learning: Evidence that abstraction is not automatic. Journal of Experimental Psychology: Learning, Memory, and Cognition, 9(4), 607.

Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. Nature, 434(7031), 387-391.

Nakamura, K. (2006). Neural representation of information measure in the primate premotor cortex. Journal of neurophysiology, 96(1), 478.

Necker, L. (1832). Observations on some remarkable phenomena seen in Switzerland; and an optical phenomenon which occurs on viewing of a crystal or geometrical solid. Philosophical Magazine, 1, 329ñ337.

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

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10(1), 104.

Nosofsky, R. M. (1986). Attention, similarity, and the identificationñcategorization relationship. Journal of Experimental Psychology: General, 115(1), 39.

Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14(4), 700.

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

Renninger, L. W., Coughlan, J., Verghese, P., & Malik, J. (2005). An information maximization model of eye movements. Advances in neural information processing systems, 17, 1121-1128.

Renninger, L. W., Verghese, P., & Coughlan, J. (2007). Where to look next? Eye movements reduce local uncertainty. Journal of Vision, 7(3).

Rothkopf, C. A., Ballard, D. H., & Hayhoe, M. M. (2007). Task and context determine where you look. Journal of Vision, 7(14).

Salvucci, D. D., & Goldberg, J. H. (2000). Identifying Fixations and Saccades in Eye-Tracking Protocols. Paper presented at the Eye Tracking Research and Applications Symposium.

Sandamirskaya, Y., & Schöner, G. (2010). An embodied account of serial order: How instabilities drive sequence generation. Neural Networks, 23(10), 1164-1179.

Satel, J., Wang, Z., Trappenberg, T., & Klein, R. (2011). Modeling inhibition of return as short-term depression of early sensory input to the superior colliculus. Vision Research.

Schöner, G., Kopecz, K., & Erlhagen, W. (1997). The dynamic neural field theory of motor programming: Arm and eye movements. Advances in Psychology, 119, 271-310.

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. Psychological monographs.

Spencer, J., Schneegans, S., & Hollingworth, A. (2010). Dynamic interactions between visual working memory and saccade planning.

Spencer, J. P., & Schöner, G. (2003). Bridging the representational gap in the dynamic systems approach to development. Developmental Science, 6(4), 392-412.

Sprague, N., & Ballard, D. (2003). Eye movements for reward maximization. Advances in neural information processing systems, 16, 1419-1433.

Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. Behavioral and Brain Sciences, 24(01), 1-34.

Triesch, J., Ballard, D. H., Hayhoe, M. M., & Sullivan, B. T. (2003). What you see is what you need. Journal of Vision, 3(1).

van Zoest, W., Hunt, A. R., & Kingstone, A. (2010). Representations in Visual Cognition. Current Directions in Psychological Science, 19(2), 116.

Watson, M. & Blair, M. (2008).Attentional Allocation During Feedback: Eyetracking Adventures on the Other Side of the Response. In B.C. Love, K.McRae, & V. M. Sloutsky (Eds.) Proceedings of the 30th Annual Conference of the Cognitive Science Society (pp. 345-350). Austin, TX: Cognitive Science Society.