A hierarchical model of visual processing simulates neural mechanisms underlying reflexive attention

Chloe Callahan-Flintoft

Hui Chen

Brad Wyble

czc213@psu.edu

(814) 867-2436

404 Moore Building

The Pennsylvania State University

University Park, PA, 16802, USA

Word Count: 13,707

Author Note: The current work's data was presented in posters at the Annual Psychonomic Society Meeting and the Annual Vision Sciences Society Meeting. The data from Experiment 1 were previously used in another publication (Callahan-Flintoft, C., & Wyble, B. (2017). *Vision research*, *140*, 106-119.)

Abstract

The visual system deploys attention reflexively in response to important stimuli to facilitate rapid selection. To better understand how a system that is distributed across a hierarchy of brain areas coordinates this rapid response, we have developed a computational model that simulates reflexive attention as a recurrent lock-on state between topographically aligned cortical areas. This model provokes key questions about reflexive attention that we study through the N2pc, an EEG component that indexes the timing and lateralization of stimulus-driven attention. A key finding from such research is that for two sequential targets sharing a location, the second elicits no N2pc, despite being easily perceivable. This suggests that attention locks-on to a target's location and carries forward in time to a second target without being redeployed. Here, we assess key properties of reflexive attention in four EEG experiments. First, the lock-on N2pc effect generalizes to feature target types. Second, the lock-on state is specific to a location, not just a visual hemifield. Third, reflexive attention decays as the interval between two targets increases, such that a second target will once-again elicit an N2pc at longer intervals. Finally, the lock-on state is not specific to a particular target type, which implies that reflexive attention is mediated using representations that are not stimulus-specific. These results provide important constraints on our understanding of visual attention. We incorporate those constraints in a formalized model of attention that elucidates the link between neural mechanisms and a key neural correlate of attention.

Introduction

A key role of visual attention is to select relevant information out of complex visual scenes for further processing. Our visual environment is constantly changing and information quickly shifts across the retina as our eyes saccade to sample information from different locations. It is therefore necessary that at least one form of attention would allow for rapid deployment to the location of task relevant information to extract that information before either the eyes or the stimulus moves. This rapid attentional mechanism is often referred to as transient or reflexive attention (Hopfinger & Mangun, 1998; Yeshurun & Carrasco, 2008), and it produces brief improvements in the ability to perceive stimuli at a cued location in experimental settings. This rapid deployment of attention also serves a crucial role in our ability to perceive by allowing current goals to influence the deployment of attention towards information that shares task-relevant attributes (Folk, Leber & Egeth 2002).

This reflexive form of attention is likely to be a crucial first stage of resource allocation within the visual system. It is extremely rapid, allowing it to be deployed within a single fixation, and yet it is also affected by goals. Thus reflexive attention is likely to serve as a crucial interface between bottom-up information and top-down goals. Moreover, the goal of understanding

reflexive attention has broad applicability to many areas of study in the field of visual cognition. For example, any paradigm in which two highly salient or task-relevant stimuli appear in rapid sequence is likely to be influenced by this form of attention. This would include, for example, the short-SOA conditions of visual cueing, attentional capture and attentional blink paradigms. In such cases, the second of two key stimuli will receive either the benefit of attention recruited by the first stimulus if they are in the same location, or exhibit reduced performance if they are in different locations (e.g. Mounts 2000). More importantly, reflexive attention presumably plays a role in allowing us to perceive a single stimulus when it is only briefly presented or fixated. Thus, understanding the mechanisms of this form of attention has broad applicability toward understanding visual cognition, broadly construed.

While there is a considerable body of research concerning reflexive attention, in terms of both its behavioral and electrophysiological correlates, its neural mechanisms remain unclear. The goal of this paper is to advance a computationally formalized theory of the neural mechanisms that direct the deployment and maintenance of reflexive visual attention. The theory is framed as a dynamic neural network that simulates the spatiotemporal evolution of activity across a hierarchical arrangement of spatiotopic neural sheets. A crucial advantage of this approach is that the model generates simulations of an EEG component that is thought to be related to attention (the N2pc; Eimer 1996; Luck & Hillyard, 1994). This allows findings from EEG experiments to be directly compared with the model's output, which opens new avenues for testing the model against human data.

The model being tested and developed here simulates the ability of attention to *lock-on* to a particular region of the visual field in response to a salient or task-relevant stimulus (Tan & Wyble 2015). This model explains a counterintuitive EEG finding when two targets are presented in sequence at the same location, which is that there is an N2pc to only the first target and not the second. This occurs despite the fact that subjects are actually more accurate in reporting the second target. The model explains this curious finding by simulating the N2pc as the neural correlate of localizing a target. This process of localization only needs to happen once for two targets in the same location, and hence there is only an N2pc to the first target. In this paper, we advance this theoretical framework by addressing a series of EEG predictions that were pre-registered in Tan & Wyble (2015) with a set of four experiments. Finally, in a fifth experiment we will explore some inconsistencies in the results of the first four. In the General Discussion the theoretical implications of the results will be assessed and the computational model will be revised to accommodate the new results.

Behavioral and neural correlates of reflexive attention

When measured with behavior, stimulus evoked attention exhibits a rapid onset and also a rapid offset. By varying the time between the cue and target onset, experimenters have shown that the attentional benefit peaks around 100ms post cue onset and then quickly begins to decline (Gottlob, Cheal, & Lyon, 1999; Nakayama & Mackeben, 1989). This time course suggests that

this form of attention cannot be maintained for a longer duration, even when subjects know that the target is likely or even certain to follow at the same location.

Furthermore, this form of attention is spatially locked to the location of the triggering stimulus, even when subjects expect the target to appear elsewhere. This was shown by Kröse & Julesz (1989), who presented a cue in a ring of stimuli and had the target appear predominantly on the opposite side of the ring. Despite this probability manipulation, attention was still deployed at the location of the cue and not at the predominant location of the target. This finding is crucial, for it suggests that the reflexive form of attention is triggered by an appropriate stimulus and cannot be easily directed to occur at a different time or place by expectation or learning. However, this form of attention does exhibit one element of top-down control: goals affect the *types* of stimuli that will be able to trigger attention (Folk, Remington & Johnston, 1992; Folk, Remington & Wright, 1994). In these findings, only cues which match a target defining feature (e.g. color, motion, etc.) are able to strongly capture attention.

Stimulus-evoked forms of attention have also been studied using the electroencephalogram (EEG), specifically through an Event Related Potential (ERP) called the N2pc. The N2pc is a brief negativity that can be recorded from electrodes on the opposite side of the brain relative to the location of a visual target. The N2pc occurs approximately 200-300ms after the presentation of a lateral target. This component is most pronounced in posterior electrodes and has been linked to several putative aspects of attentional selection. The first of such theories was that the N2pc reflected distractor suppression, as the N2pc was only found when a target was presented among distractors and not alone (Luck & Hillyard, 1994). However, Eimer's (1996) simplified array, which consisted of a target and distractor on either side of the fixation, also elicited an N2pc, despite the fact that there were no distractors to suppress in the vicinity of the target. These results inspired the theory that the N2pc reflects target enhancement, rather than distractor suppression. Other studies have found that the amplitude of the N2pc is affected by the number of targets but only when participants needed to report target numerosity, which the authors interpret as an indication that the N2pc reflects target individuation (Mazza & Caramazza, 2011).

Other experiments have found that the N2pc is a summation of two components: the Nt which is elicited by a target and the Pd which is thought to be related to distractor suppression (Hickey, DiLollo, & McDonald, 2009). As these components are only seen in the subtraction of the contralateral and ipsilateral waveform it is often difficult to disentangle them. The current model offers a way to better understand the mechanisms underlying the N2pc which will help us to pull apart these co-occurring waveforms. Furthermore, the N2pc is often followed by a sustained lateralized negativity known as the contralateral delay activity (CDA) or sustained posterior contralateral negativity (SPCN) (Jolicœur, Brisson & Robitaille, 2008), which is known to scale with working memory load (Vogel & Machizawa, 2004). This CDA component is most often seen after a post N2pc positivity which occurs with and without backwards masking (Robitaille & Jolicœur, 2006; Töllner, Müller, Zehetleitner, 2012)

While the N2pc is a complex waveform that may include multiple mechanisms, research has found consistent properties across experiments. For example, the N2pc seems to be related to the *onset* of attention. We know that the N2pc duration is not increased by a longer-duration target (Brisson & Jolicouer, 2007). Furthermore, the N2pc is only present for the first of two sequential targets presented at the same location. However when the same two targets are presented on opposite sides of the visual field in sequence, two sequential N2pcs of opposite polarity are observed (Tan & Wyble, 2015). Thus it seems that whatever attentional mechanisms are involved in the processing of one or more targets at a given location, the N2pc is unaffected by the total duration or number of targets at that location.

The CGF model

Collectively, these behavioral and neural constraints were used to develop a computational model of attention that simulates a combination of excitatory and inhibitory mechanisms acting in concert to enhance processing of targets, and to inhibit surrounding information. The model, called the convergent gradient-filter model (CGF), simulates EEG data as virtual ERPs (vERPs; Craston, Wyble, Chennu & Bowman 2009) resulting from the summed excitatory and inhibitory currents feeding into a master attention map (a more detailed description of the CGF model can be found in Tan & Wyble, 2015). The benefit of this approach is that hypotheses can be made about underlying neural mechanisms and the model will link those mechanisms to very specific, testable patterns of electrophysiological recordings. Furthermore, predictions stemming from computationally formalized accounts impose a unidirectional scientific method on the interpretation of experimental results that is clearer about what exactly was predicted by the theory.

Methods. The model is a rate-coded neural model, using the equations of O'Reilly & Munakata (2001) to determine membrane potential based on a combination of excitatory, inhibitory and leak currents (1).

$$M_{i,j,t} = M_{i,j,t-1} + dt_{VM} \times [(Bias_j + Excite_{i,j,t-1}) \times (EE_j - M_{i,j,t-1}) + Inhib_{i,j,t-1} \times (EI_j - M_{i,j,t-1}) + Leak_j \times (EL_j - M_{i,j,t-1})] \quad (1)$$

In this equation the membrane potential (M) of the i^{th} neuron on the j^{th} layer of the hierarchy at time t is determined by the excitatory (*Excite*), inhibitory (*Inhib*) and leak currents (*Leak*) acting on that neuron at time *t-1*. These influences are scaled by the difference between the neuron's membrane potential at time *t-1* and the reversal potential for each current type (*EE*, *EI*, and *EL*, respectively). This input is then multiplied by the neuron's rate of change (dt_{VM}).

Architecture. The primate visual system is composed of a multi-tiered hierarchy of cortical maps, many of which are organized according to a shared spatiotopic index (Silver, Ress, & Heeger, 2005). The CGF model implements an abstraction of this hierarchy with only one map per tier,

while preserving the key properties of shared spatiotopy between maps and receptive fields that increase in size through the hierarchy. This abstraction should not be considered a specific assumption that any layer of the model reduces to one cortical area. The lock-on mechanisms described here would apply to a network of any complexity, provided that the neurons preserve the property of shared spatiotopy, which we view as key for rapid decision making and attentional selection in a distributed network of visual areas.

The first layer of the CGF model is the Early Visual layer (EV), which is analogous to V1. The EV is the first layer to receive stimulus input, in the form of excitation of EV neurons that correspond to the location of a target in the visual field. Excitation from activated neurons in the EV feeds forward to the late visual layer (LV), corresponding to areas such as inferior temporal cortex in the ventral pathway where neurons show spatiotopy with large receptive fields (DiCarlo & Maunsell, 2003). Only a single LV layer is necessary for these simulations, but the assumption is that a number of highly specialized LV's are involved in visual processing. The LV then projects to a spatiotopically organized attention map (AM) that spans the entire visual field. While the anatomical location of the AM is not as well prescribed as the EV and LV layers, such a map is likely to consist of one or more cortical areas within the parietal cortex. Such regions have shown spatiotopically organized activation during attention tasks (Silver, Ress, & Heeger, 2005). Neurons in the AM have still larger receptive fields than in LV or EV, allowing them to integrate information from a larger area of the visual field. Thus activation from a single location in the EV produces a broad bump of activation in the LV and an even broader bump in the AM. Then, once any neuron in the AM is sufficiently activated, that neuron triggers attentional enhancement at its corresponding location in the visual field by amplifying feed-forward excitation from corresponding spatiotopic coordinates in EV (Figure 1). There is also lateral inhibition between neurons in the AM such that when a neuron is over threshold, it will

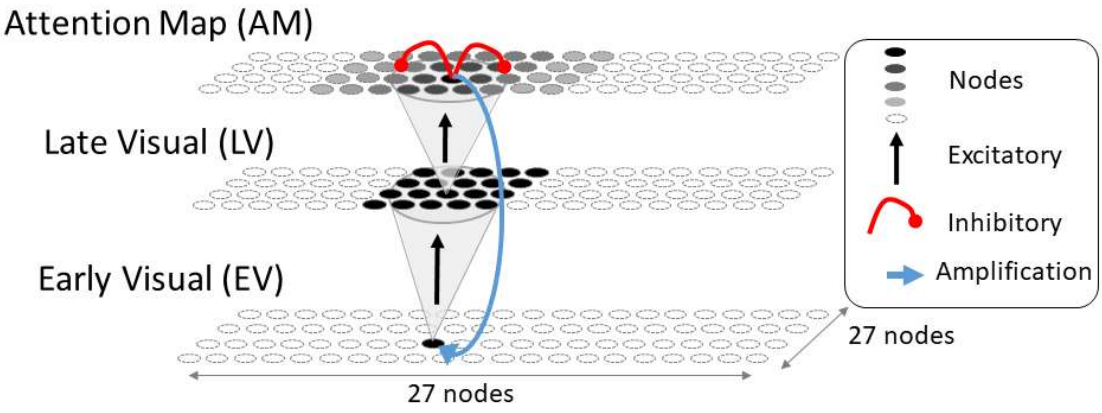

Figure 1. A schematic of the neural architecture of the CGF model. Excitation feeds forwards from the Early Visual (EV) layer to the Late Visual (LV) layer and then onto the Attention Map (AM). Once AM neurons cross threshold, attentional amplification is sent down to the corresponding EV neurons. Additionally, AM neurons above threshold project lateral inhibition onto neighboring neurons in the AM. Shades of gray indicate levels of neural excitation

send inhibitory signals to neighboring neurons in the AM. The consequence of this lateral inhibition is that while the activation of the AM in response to a target might be initially very broad, it is quickly narrowed such that only the specific location of a target is enhanced by the attentional amplification.

The lock-on state. Attending to a target at the level of the attention map can be broken apart into three different stages (Figure 2). The first stage is the initial localization of the target in which the AM shows a broad increase in activation around the target's location. This lack of specificity is due to the large receptive fields of the AM neurons which pool excitation from a wide region of the visual field. This initial localization stage continues as this region of the attention map continues to receive excitation and ends when AM neurons begin to cross their activation threshold. Once AM neurons begin to cross threshold, two things happen: they begin sending amplification to EV and laterally inhibiting neighboring AM neurons. The combination of these two mechanisms brings about the second stage of processing in the AM, referred to here as the lock-on state. Because of the lateral inhibition, neurons on the periphery of that initially broad bump of activation are suppressed, leaving active only the neuron(s) that directly correspond to the target coordinates available to receive the attentional amplification. This leaves the attention map with a spike of activation at the location of the target, surrounded by a deep well of inhibited activation to limit attention to the target location. We posit that this lock-on state is the neural origin of the behavioral phenomenon of reflexive attention. This activation profile is similar in some respects to dynamic field theory (DFT) models of working memory (Johnson, Spencer & Shöner, 2008; Johnson, Spence, Luck & Shöner, 2009). However this lock-on state is meant to model a rapid stimulus driven attentional deployment rather than a sustained

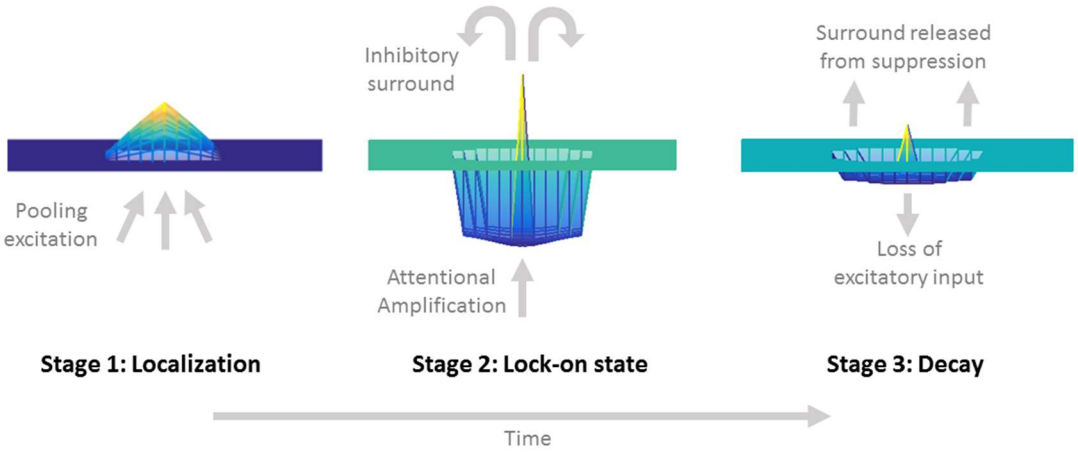

Figure 2. Three stages of target processing in the attention map (AM). Stage 1 is the localization of the target. Excitation is pooled from a broad area of neurons in the LV layer. Stage 2 is the lock-on state. AM neurons over threshold send down amplification in an excitatory loop while also projecting inhibition to surrounding neurons. Stage 3 is decay when there is no further information to feed excitation to the lock-on state, AM neurons begin to return to resting state membrane potential.

maintenance in working memory. It therefore establishes quickly but can only be maintained through the rapid presentation of new, relevant or salient information. Without the presentation of new information the lock-on state begins to dissipate, which is the third and final stage of the AM's response to a target. In the absence of excitation from the early and late visual layers, activated neurons in the AM return to their resting state membrane potential within a few hundred milliseconds. Once these neurons drop below threshold, they stop sending amplification and inhibiting the surrounding AM neurons. Without the incoming inhibition these surrounding neurons will also begin to return to baseline. In this way the lock-on state dissipates over several hundred milliseconds by losing its activation peak and its surrounding inhibitory well. The other important difference between the CGF and DFT models is that the CGF model describes a pattern of coherence organized across multiple layers of a hierarchy, rather than memory representations within a given layer. A key emergent property of hierarchical organization is that it ensures the spatial accuracy of the lock-on state. Because stimulus input is necessary to maintain the lock-on state, it is guaranteed to only be stable at the location of the target. Conversely, if this same activation profile was achieved solely through excitatory connections within the attention map (i.e. a 'bump' attractor as in the DFT theory), it would be likely that the attended location would drift away from the actual location of an important stimulus when a scene contains multiple stimuli at nearby locations.

Simulating the neural correlates of attention. The CGF model generates virtual event related potentials (vERPs) by summing synaptic activity across the AM. This is a coarse approximation based on our best understanding of the origin of EEG signals, which is that they result from a summation of postsynaptic potentials in cortical neurons that are anatomically oriented in the same direction within the cortex (Luck, 2014). Since the model lacks the physiological structure necessary for a complete forward model of ERP generation, we incorporated a polarity assumption, which is that excitatory synaptic potentials in cortical areas corresponding to the attention map produce a scalp negativity. To calculate the simulated N2pc at each time step, the AM is split down the vertical meridian and one side's vERP is subtracted from the other's much like how the contralateral and ipsilateral waveforms are subtracted from one another to produce a real N2pc.

Calculating the vERP in this way allows us to link the N2pc waveform to specific neural processes (Figure 3). The model attributes the initial negative deflection of the N2pc to the localization of a target, which involves a large mass of excitation in the vicinity of the target in the attention map. The abrupt positivity that typically brings the N2pc waveform back to baseline occurs when the lateral inhibition is triggered by the lock-on state and the surrounding neurons are suppressed. Importantly, the model proposes that the end of the N2pc component actually marks the *beginning* of the lock-on state, which is invisible to the EEG electrodes due to its narrow focus on one location. In the following section we will explain the evidence for this assertion.

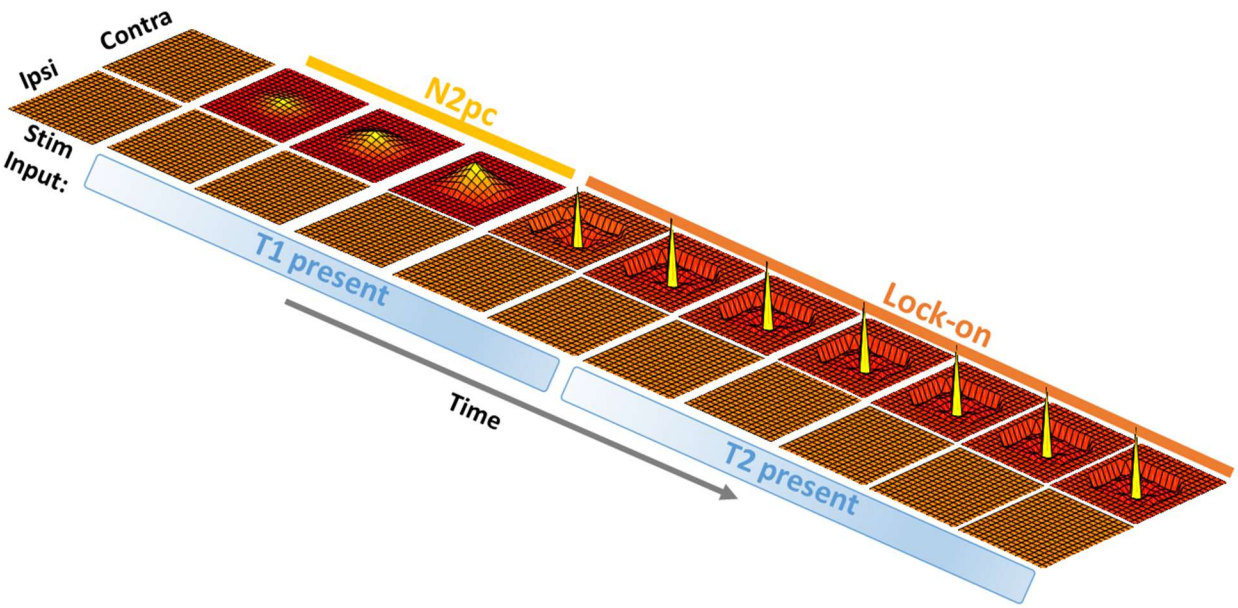

Figure 3. Progression of activity in the attention map to two consecutive targets. Blue boxes represent the stimulus input into the model through time. The half of the attention map contralateral to the target receives excitation resulting in the N2pc. The establishment of the lock-on state shuts off the N2pc as it returns the net difference in activation between contra and ipsilateral sides to baseline. The second target, occurring at the same location, inherits and extends the lock-on state but does not produce a second N2pc

Evidence for a lock-on state

Tan & Wyble (2015) developed and parameterized a preliminary version of the model with data from an experiment that sought to test whether a second target (T2) presented at lag-1 in the same location as the first (T1) would elicit a second N2pc. They presented participants with two RSVP streams (one on either side of fixation). Three conditions were used: T1 alone, T2 presented at lag-1 in the same stream as T1, and T2 presented at lag-1 in the opposite stream. The CGF model with its lock-on dynamic would predict that in the *Same Stream* condition, the lock on state would carry forward from T1 and T2, resulting in only a single N2pc. In this case, the excitatory input generated by T2 would continue to sustain the lock-on state established at that location by T1. However, if T2 appeared in the opposite RSVP stream, the model predicts a second N2pc of reverse polarity. This is because the T2 now has to establish a second lock-on

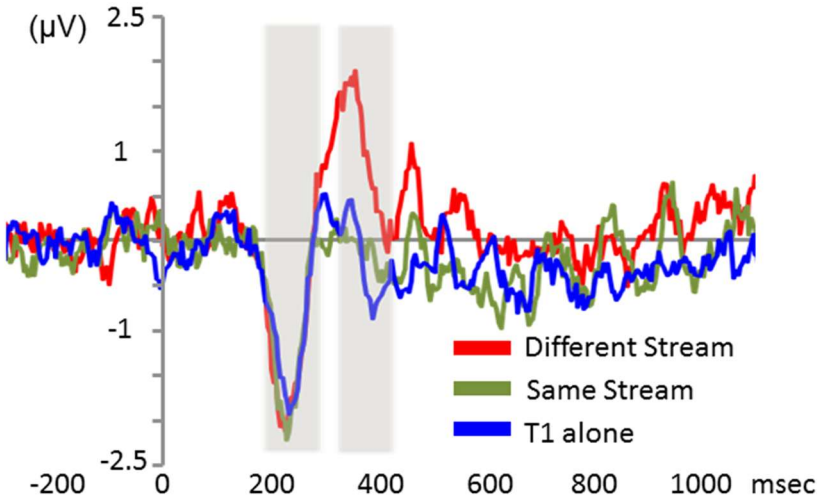

Figure 4. Difference waves (Contra-Ipsi) from P3/P4 electrodes of Tan & Wyble (2015). All conditions show a similar N2pc elicited by the first target. The Different Stream condition then shows a second N2pc of reverse polarity as the second target was presented in the opposite hemifield. Importantly the Same Stream condition shows no second N2pc to the second target.

state on the opposite side of the visual field, leading to a reversed pattern of voltages. The experimental result matched the findings of the vERPs, lending support to the idea of the lock-on state (Figure 4).

Moving Forward: addressing questions about the functional architecture of reflexive attention

The CGF model (Tan & Wyble 2015) is a mathematically and mechanistically formalized theory that prompts four specific questions about the functionality and architecture of reflexive visual attention. In this paper we use a series of EEG experiments to provide answers to those questions. The results will inform our understanding of attention by either reinforcing assumptions already inherent in the CGF model, or indicating which assumptions should be revisited. Three of these questions correspond to specific predictions that were pre-registered by publication in Tan & Wyble (2015).

*Question 1: Does the attentional lock-on state replicate for other kinds of targets? S*tudies have shown an N2pc in response to a wide variety of target types, such as color marked (Girelli & Luck, 1997; Kiss, Driver, & Eimer, 2009), shape singleton (Burra & Kerzel, 2013; Kiss, Jolicœur, Dell'Acqua & Eimer, 2008; Hickey, McDonald, & Theeuwes, 2006), and even complex line drawings (Nako, Wu, Smith & Eimer, 2014). The CGF model explains the N2pc as a consequence of attentional lock-on and therefore the crucial lock-on finding of Tan & Wyble 2015 (i.e. two sequential targets in the same location produce only one N2pc) should be obtained for different kinds of targets. Here we will contrast lock-on effects for two target types: feature targets (letters of a specified color) and orthographic targets (digits among letters, similar to what was used in Tan & Wyble 2015).

Question 2: Does the lock-on only occur at the location of the target or does it include a broader region in the vicinity of the target? A core assumption of the CGF model is that attention locks onto a specific location. Specifically, the CGF model predicts (Prediction 1, Tan & Wyble 2015) that a following target presented on the same side, but several degrees away from an earlier target would elicit a second N2pc as this second target would require a new lock-on state.

Question 3: Can a lock-on state be maintained over a long period of time or does it expire automatically? In the CGF model the lock-on state is a transient attractor state that requires continued target presence to be maintained. The lock-on state is therefore maintained when T2 is presented directly after T1. The model predicts (Prediction 2, Tan & Wyble 2015) that for a longer time interval between T1 and T2 (e.g. 500ms or more), the lock-on state will dissipate, as the neurons in the AM will return to baseline in the absence of input. Practically speaking this would mean that given sufficient time between the presentation of two targets the lock-on state will need to be re-established for the second target, resulting in a second N2pc. This would occur even if subjects can predict that the second target would always be presented in the same location as the first. Thus, a top-down expectation of the subject will be unable to maintain the deployment of reflexive attention in the absence of targets.

Question 4: *Is there a single attention map that is shared across target types?* Another key assumption of the CGF model is that there is a single superordinate attentional map that directs attention without regard to the type of target. Therefore, if one type of target (e.g. orthographic) establishes a lock-on state, a different kind of target (e.g. color-marked) presented sequentially should be able to maintain that lock-on state. This is Prediction 5 of Tan & Wyble (2015).

The results for each experiment will be presented in two parts. A confirmatory analysis will use exactly the same analysis methods and electrodes as were originally used in Tan & Wyble (2015). Following each confirmatory analysis is an exploratory analysis conducted over the entire set of electrodes using a cluster-mass approach (Groppe, Urbach, & Kutas, 2011). This second analysis allows scalp localization of statistically significant ERP effects while attempting to control for the family-wise error rate.

Experiment 1

The purpose of this experiment is to answer question 1: *Does the attentional lock-on state replicate for other kinds of targets?* An experiment was run in which targets were either letters of a specific color (feature target) or digits of any color (orthographic target). The comparisons of interest in this study were to see whether the lock-on state was similarly present for both target types by looking for the presence of a second N2pc.

Methods

Participants. Data from 36 subjects were collected for the first experiment. All subjects volunteered from the Pennsylvania State University psychology subject pool for this study. Participants all had normal or corrected-to-normal vision and were between the ages of 18 and

23 years old. Informed consent was obtained for each participant prior to each study in accordance with the IRB office of Penn State. Subjects were excluded for having too few usable trials per condition (threshold: 15) after discarding trials for inaccurate responses and EEG artifacts, leaving 25 subjects usable for analysis.

Stimuli & Apparatus. Participants sat in an electrically shielded room, 91cm away from the computer which had a 46cm CRT (1024 x 768, 60 Hz refresh rate). Stimuli were alphanumeric characters in size 55 Arial font (1.26°x 0.63° of visual angle), presented on a grey background using MATLAB 2012 with Psychophysics Toolbox-3 extension (Brainard, 1997). Feature targets were letters of a designated randomly selected target color (one out of five possible) whereas orthographic targets and distractors were presented in the remaining four colors with the constraint that the same color was not presented twice in the same frame. The target color was counterbalanced across participants, but for each participant the target color was consistent throughout the experiment (e.g., one participant was always looking for digits and red letters). The colors used were: green (0, 255, 0), red (255, 0, 0), blue (0, 0, 255), cyan (0, 225, 225) and yellow (225, 225, 0). The digits used as orthographic targets were: 2,3,4,5,6,7,8,9. Distractors and feature targets were drawn from the following collection of letters: A,B,C,D,E,F,G,H,J,K,L,N,P,Q,R,T,U,V,X,Y. Letters that resembled digits, such as 'O', 'S', 'Z' and 'I', were excluded as well as wider letters such as 'M' and 'W'. Target identities were randomly selected on each trial.

Procedure. On every trial participants were presented with two RSVP streams placed three degrees from a fixation cross, which remained throughout the trial (Figure 5). Both streams were updated simultaneously with a 150ms stimulus onset asynchrony (SOA) and no inter-stimulus-interval. Note that this SOA is slightly longer than that used in Tan & Wyble (2015). This change was implemented to increase accuracy, since analyses could only be done on trials in which subjects reported both targets successfully. Seven or eight distractors in each stream were presented before target onset and eight distractor pairs were presented afterwards. This number of pre- and post- target distractors was used in all experiments except Experiment 2. Every trial had one or two targets, the second of which (T2) always appeared in the same stream and immediately after the first (T1). No stimulus was repeated for at least two sequential frames. At the end of the trial, the fixation cross was replaced with either a period or a comma for 150ms. Two practice trials were used within the instruction block to demonstrate the task.

Targets were defined by feature (letters of a specific target color among letter distractors of the colors: green, red, blue, cyan, yellow, see Stimuli & Apparatus section for RGB values) or orthography (digit among letters—equally likely to appear in any of the four distractor colors). There were four conditions: feature target alone (Feat), orthographic target alone (Ortho), feature followed immediately by another feature target (Feat-Feat) and orthographic followed by another orthographic target (Ortho-Ortho). There were 280 trials distributed equally across the four conditions, intermixed within a single block. After every 20 trials, participants were given a self-paced break.

Instructions and responses. Participants were told that every trial would contain one or two targets that they would have to report, in the correct order, using the keyboard. They were also told that on some trials they would have to report which symbol had replaced the fixation cross. At the end of each trial, participants were asked for the first and second target, with the option of pressing Enter if they did not know it. On a third of trials participants were also asked to report the symbol (dot or comma) that appeared at fixation. This was done in order to encourage participants to remain engaged, with their eyes on fixation, for the entire trial and not just until the target(s) appeared. This technique was used to discourage eye movements but the actual elimination of trials with eye movements was done with EEG measures described below. Feedback was provided by showing the participants what targets in order were shown as well as what punctuation had replaced the fixation cross (dot or comma).

Participants were excluded from both behavioral and EEG analysis if they had fewer than 15 usable trials in any condition. Trials were considered usable if participants accurately reported

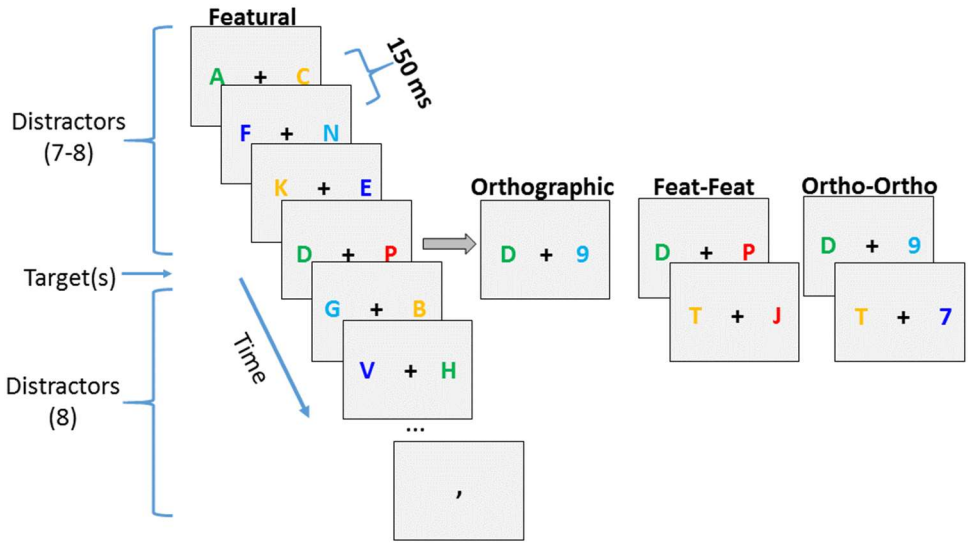

Figure 5. Paradigm for experiment 1. Participants were presented with two RSVP streams. There were 7-8 frames (150ms SOA) of distractors prior to target presentation and 8 frame post target. Targets were digits or letters of a designated color (in this depiction the target color is red, however for a given participant the target color could be any one of the five available colors).

the identity of all targets present, regardless of the order, and if the trial passed the EEG exclusion criteria described below. The average number of usable trials per condition was *M* = 33 (*SE* = 1.7), *M* = 38 (*SE* = 1.8), *M* = 35 (SE = 2.6), *M* = 42 (*SE* = 2.3) for Feat, Ortho, Feat-Feat, and Ortho-Ortho, respectively.

EEG Recordings. The same EEG system was used to collect data for all of the experiments presented here. A 32-channel sintered Ag/AgC1 electrode array mounted in an elastic cap according to the 10-20 system (FP1, FP2, F3, F4, FZ, F4, F8, FT7, FT8, FC3, FC4, FCz, T3, T4, C3, C4, Cz, TP7, TP8, CP3, CP4, CPz, T5, T6, P3, P4, Pz, O1, O2, and Oz) (QuikCap, Neuroscan Inc.) was used, with the tip of the nose used for a reference. The amplifier was a Neuroscan Synamps with a band pass filter of 0.05-100 Hz and a sampling rate of 500 Hz. The data was reduced offline to 250 Hz. Prior to the start of the experiment, the impedance for all electrodes was lowered to 5 kΩ or less. VEOG and HEOG electrodes were placed on the lower and upper orbital ridge of the left eye as well as the outer canthus of each eye respectively. All ERPs were time locked to the first target onset and a three second epoch (1 second prior to T1 onset, two seconds post) was used for each trial. Baseline activity from a window -200ms to 0ms relative to T1 onset was subtracted from each trial. If the average difference between HEOL and HEOR electrodes in a moving 32ms time window exceeded 20μV, the trial was judged as having contained a horizontal eye movement and was rejected. If the difference between VEOU and VEOL electrodes increases more than 100μV, the trials were marked as having contained a blink and was rejected. Additionally, if any channel exceeded an absolute value of 100μV during the epoched window, that trial was rejected. All artifact rejection and EEG analysis was performed with a combination of custom MATLAB 2012 script and EEGLab 13.3.2b functions (Delorme & Makeig, 2004).

Results

All of the following experiments were analyzed both with a confirmatory and exploratory approach. The confirmatory approach used the same methods and analysis parameters as Tan & Wyble (2015). To show that the first target in all conditions elicited an N2pc, an ANOVA was used to compare mean amplitude of contralateral vs. ipsilateral waveforms in the window of 200-300ms after first target onset for electrodes P3 and P4. To test for the existence of a second N2pc, a sliding-window permutation test was conducted using a MATLAB script to test whether there exists a significant difference between the difference wave (contra minus ipsi) of the single target condition (condition A) and the double target condition of interest (condition B) in electrodes P3/P4. The permutation analysis randomized the subtraction between condition A and condition B difference waveforms for each participant. This randomization was repeated 5,000 times, making a null hypothesis distribution for each 20ms time window from -1000ms to 2000ms relative to T1 onset. The true subtracted difference was then compared to the null distribution to provide a p-value at each time step. A time window was considered significant if at least 4 time steps in a row (80ms total time) had a p-value less than .05. Importantly there were no significant windows in the range of -1000 to 150ms relative to T1 onset for any of the

comparisons, demonstrating that this criterion was not prone to false alarms. This was exactly the analysis performed in Tan & Wyble (2015).

In addition to the confirmatory analysis, an exploratory analysis was also carried out using a cluster mass univariate approach with the Mass Univariate Toolbox (Groppe, Urbach, & Kutas, 2011). In this analysis a t-score was calculated for every time point at each electrode. T-scores that did not exceed a threshold set by an uncorrected p-value over 5% are eliminated. The remaining t-scores were clustered by spatial and temporal proximity and the t-scores are summed within clusters to calculate a cluster mass score. This process was repeated in a permutation test where cluster masses are compared against a distribution of the maximum cluster mass produced by each permutation. This method allows the exclusion of all clusters except those with a $p < .05$, and a control for family wise error (for further details see Groppe, Urbach, & Kutas, 2011). The time window for the cluster analysis was 0ms (T1 onset) to 200ms after the point at which the T2 N2pc (had it occurred) would have been expected based on the latency of the T1 N2pc. All electrodes were used in this analysis, however it is the difference between electrodes that is tested and so electrode pairs are presented in the figures.

Behavioral results. As the CGF model currently only simulates electrophysiology data, no specific predictions were made about accuracy. As a result, the behavioral accuracy of each condition is reported, but not analyzed. In Experiment 1 the mean accuracy rates and standard errors for each condition were, $M = .85$ ($SE = .02$) for Feat, $M = .85$ ($SE = .02$) for Ortho, $M = .59$ ($SE = .03$) for Feat-Feat, and $M = .68$ ($SE = .02$) for Ortho-Ortho. The accuracy for the two-target trials is the proportion of trials in which both targets were correctly reported, though not necessarily in the correct order.

Confirmatory EEG analysis. A two (contra vs. ipsi) by four (conditions) repeated measures ANOVA compared the mean amplitudes for the window 200-300ms post T1 onset between the contralateral and ipsilateral electrode of the P3/P4 pair to ensure that the first target in all conditions elicited an N2pc. There was a significant effect of laterality, $F(1,24) = 40$ $p < .001$, $\eta_p = .63$ and no effect of condition, $F(3,24) = .59$, $p = .62$, or interaction, $F(3,24) = 1.08$, $p = .36$. There was significant laterality in all conditions tested individually at $p < .001$.

Next a permutation analysis was run comparing conditions of like target type (i.e. Feat vs. Feat-Feat and Ortho vs. Ortho-Ortho) (Figure 6). For feature defined targets, when comparing the difference wave of the single versus two-target condition there were no significant time windows. However, for orthographic targets there was a significant difference between single and two target conditions during the window of 280-400ms post T1 onset which corresponds to time period just prior to when the second N2pc should have occurred, given that the T2 occurred exactly 150ms after the T1. Note that the orthographic target result is a slight deviation from the results of Tan & Wyble (2015) and the CGF model simulations, which will be discussed below.

Exploratory analysis. A cluster analysis showed no significant differences between the Feat and Feat-Feat conditions. There were, however, two significant clusters when comparing the Ortho to the Ortho-Ortho condition (Figure 7). The first started around 296ms post T1 onset in frontal-central electrodes and then extending into posterior electrodes until 328ms. The second started earlier in posterior electrodes at 356ms and finishes 400ms post T1 onset.

Discussion

The lack of difference between the Feat and Feat-Feat condition replicates the Tan & Wyble findings using color-marked targets instead of orthographic targets. Thus we show that EEG correlate of attentional lock on occurs for color marked stimuli and answer the question: *Does the attentional lock-on state replicate for other kinds of targets?* in the affirmative.

However, in the Ortho vs Ortho-Ortho comparison both the confirmatory and exploratory analyses showed a significant difference in the time window just before the second N2pc would have been predicted (~400ms). While the time course and magnitude do not match what would be expected of a full second N2pc, this is a deviation from the results of Tan & Wyble (2015)

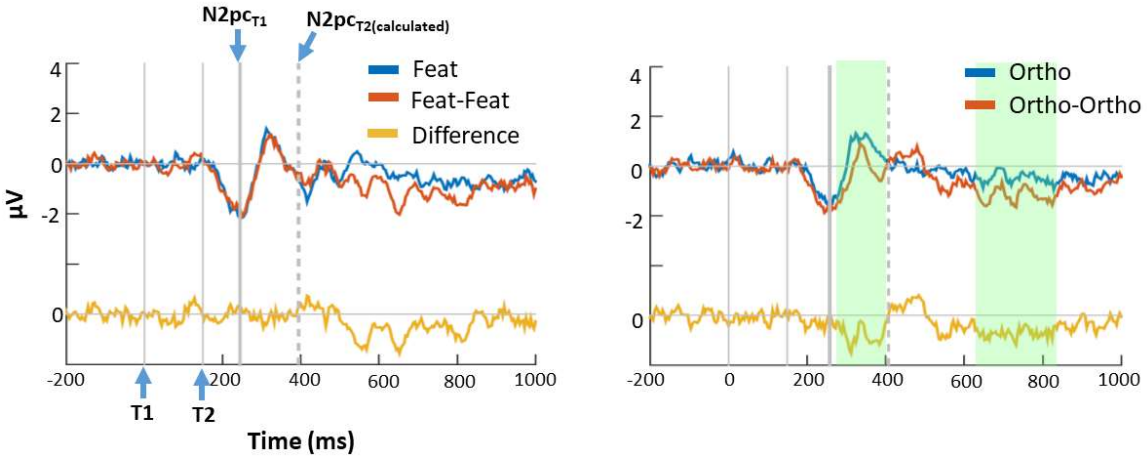

Figure 6. Difference waves from electrodes P3/P4 for each condition. The yellow waveform is the difference between the conditions' difference waves that are being compared. Green boxes show the windows of significant differences determined by the permutation analysis. The two vertical thin gray lines denote the time of T1 and T2 onset. The thick gray line is the point of 50% area under the curve for the N2pc elicited by T1. The dotted line is the predicted time of the N2pc to T2, calculated by adding the SOA to the time of T1's N2pc.

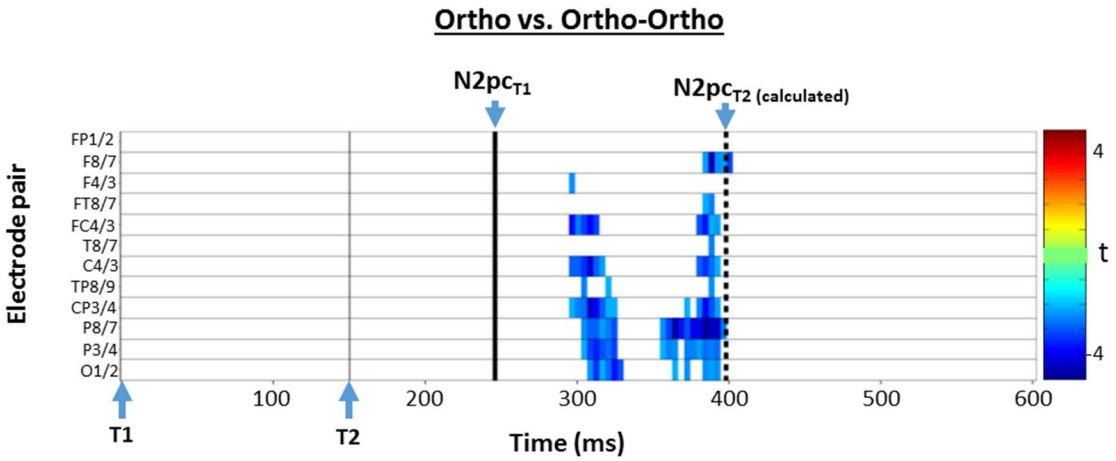

Figure 7. Cluster analysis comparing condition Ortho with condition Ortho-Ortho. Color bars denote significant time windows in a given electrode pair. Arrows mark the onset of T1 and T2. The thick vertical line is the time of 50% area under the curved for the first N2pc and the dotted line is the predicted time of the second N2pc.

who found essentially identical ERPs for the equivalent conditions, except that the SOA in Tan & Wyble was slightly shorter at 120ms. We will return to this point in the General Discussion.

Experiment 2

This experiment addressed question 2: *Does the lock-on only occur at the location of the target or does it include a broader region in the vicinity of the target?* by testing the CGF model's prediction that a second N2pc would be elicited by a second target presented on the same side but several degrees away from the T1.

Method

Participants. Twenty-six students participated in Experiment 2[1]. After filtering for accuracy and artifact rejections (see Experiment 1's method for details), 25 had the minimum number of trials to be included in analysis.

Stimuli & Apparatus. Distractors in this experiment were letters (R, L, C, P, F, K, B, G, Y, V, H, X, T, J, D, N) and targets were digits (2, 3, 5, 8, 9). All stimuli were black (RGB: [0 0 0]) presented in Arial font, size 78 (1.78°x 0.89° of visual angle). The target identities were selected randomly on each trial. Four RSVP streams were presented 3.26 degrees from fixation

[1] Two participants were added to this sample after the initial analysis to equate the number of participants in each experiment. The results did not qualitatively differ with this addition. The original results (from the n=23 sample) have been included in the supplemental as well as a Bayesian analysis to accompany the confirmatory ANOVA.

in a rectangle around the fixation cross (centered at corners of a rectangle 5.5° horizontal x 3.5° vertical). All monitor specifications are the same as in Experiment 1.

Procedure. In Experiment 2, participants were presented with 4 RSVP streams, creating a rectangle around fixation (Figure 8). For this experiment the SOA was 133ms. Distractors were black letters and targets were black digits. Four to seven sets of distractors were presented before the target frame and six to nine frames were presented afterwards so that each RSVP stream was fourteen characters long on every trial. In addition to a single target condition (Single, S) there were three two-target conditions. The T2 immediately followed the T1 either in the same stream (Same Location, SL), in a different stream but on the same side of the screen (Same Side, SS), or in the stream on the opposite side of the screen (Different Side, DS). At the end of the trial, participants were asked to report the targets that they saw in the order that they were presented. Trial types were intermixed. Participants completed 88 trials that were evenly distributed among all 4 conditions. The average number of usable trials per condition were: M = 53 (SE = 2.3), M = 61 (SE = 61), M = 35 (SE = 2.6), M = 51 (SE = 2.3), for Single, SL, SS, and DS conditions respectively.

Results

Behavioral results. The mean accuracy for each condition and standard error of Experiment 2 were: M = .86 (SE = .02) for Single, M = .83 (SE = .02) for Same Location, M = .46 (SE = .03) for Same Side, and M = .70 (SE = .03) for Different Side. As in Experiment 1, these are proportions of trials in which observers report both targets correctly, regardless of order.

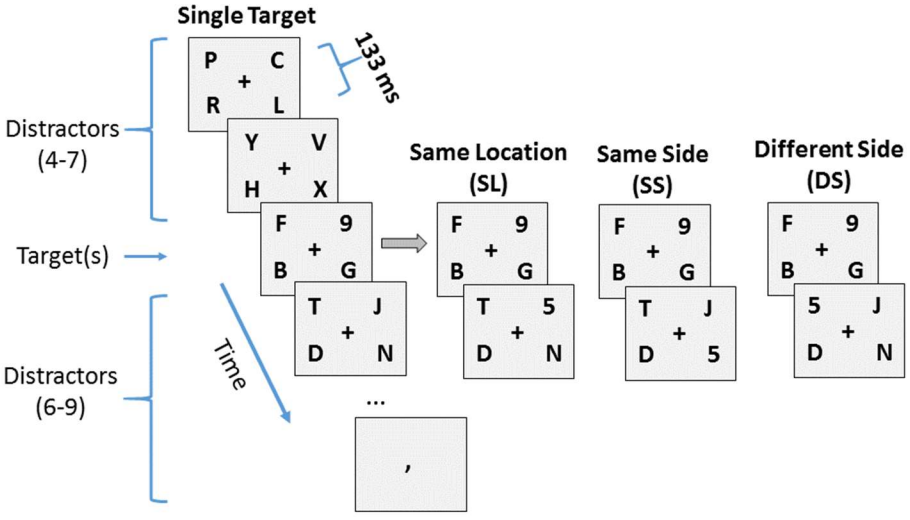

Figure 8. Paradigm for Experiment 2. Participants were presented with 4 RSVP streams. There were 4-7 pre-target frames of distractors (133ms SOA) and 6-9 frames post-target presentation. Targets were black digits among black letter distractors.

Confirmatory analysis. A two-way 2 (contra/ipsi) x 4 (conditions) repeated-measures ANOVA was run on the 200-300ms time window post T1 onset which showed a significant effect of laterality, $F(1,24) = 42$, $p < .001$, $\eta_p = .63$ in electrodes P3/P4. There was no significant effect of condition, $F(3,24) = 2.02$, $p = .12$, or interaction, $F(3,24) = 1.40$, p=.24. A comparison of contra and ipsi waveforms for each condition separately showed a significant effect of laterality in each, $p < .001$. A permutation analysis was run comparing the single target difference wave to each of the two target conditions (Figure 9). Comparing condition S (Single) with SL (Same Location) showed a significant window 500-640ms post T1 onset; S versus SS (Same Side) had a significant window from 400 to 500ms and 560 to 760ms; and S versus DS (Different Side) had a significant window from 340 to 440ms. An additional comparison that was important here was to compare SL versus SS since these conditions had similar numbers of targets. A permutation test revealed a significant difference between the two waveforms during the 420-480ms window.

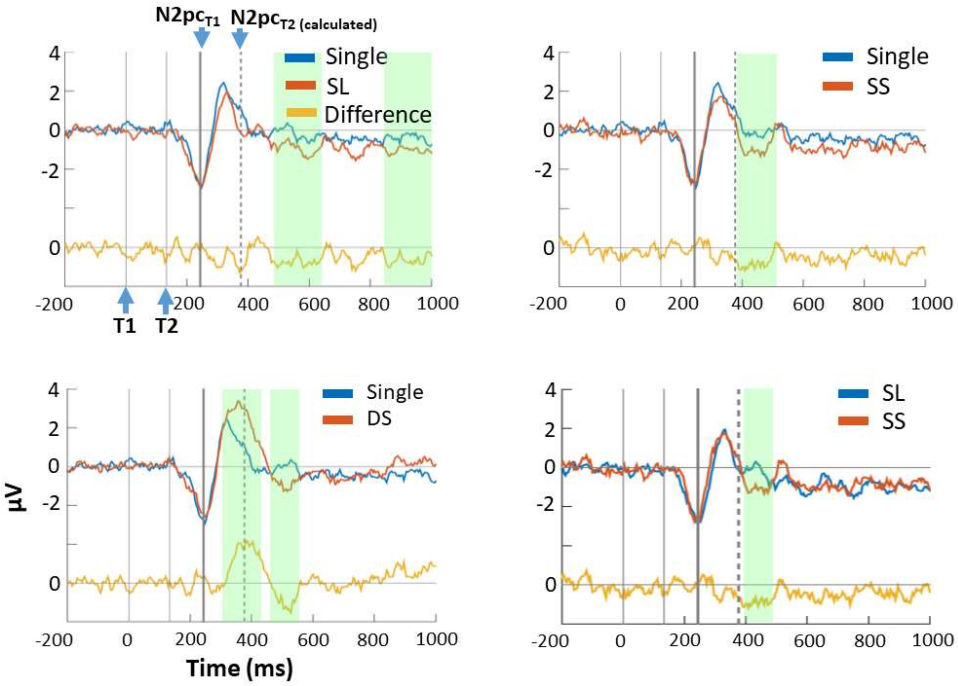

Figure 9. Difference wave comparison from Experiment 2. Green boxes mark the windows of significant differences from the permutation analysis. Vertical lines mark the onset time of T1 and T2 as well as the measured time of the first N2pc and the predicted time of the second N2pc.

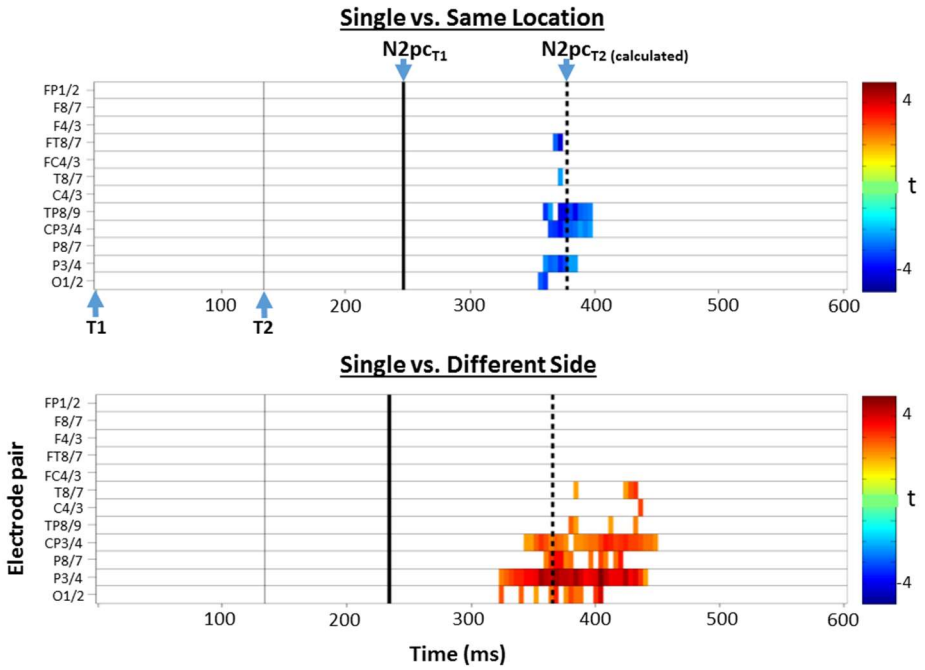

Figure 10. Cluster analysis comparing the Single condition to conditions Same Location and Different Side of experiment 2. Vertical lines mark the onset of T1 and T2 as well as the measured first N2pc and calculated time of the second N2pc.

Exploratory analysis. The cluster analysis revealed one significant difference between S and SL waveforms beginning in posterior electrodes at 356ms and spanning to central-and temporal-parietal electrodes until 396ms post T1 onset (Figure 10). A significant difference between S and DS started in posterior electrodes at 324ms and extended anteriorly until 448ms. There was no significant difference between S and SS waveforms in the cluster analysis or between SS and SL, contrary to what was found in the permutation confirmatory analysis

Discussion

The important comparisons for Experiment 2 were between S and SL, and SS and SL. In the CGF model, attention is deployed to the specific location of a target and does not need to be redeployed to a second target presented in quick succession at the same location as the first. Thus the CGF model predicts the absence of a second N2pc in response to T2 in the SL condition. As the lock-on state is location and not side specific, presenting a second target in the same side but in a different location (SS) would require the deployment of attention and establishment of a new lock-on state at the second location. When comparing the SS condition's difference wave to that of the S condition, we find a significant window between 400ms and 500ms post T1 onset. This window is slightly later than the expected time of the second N2pc (~375ms). One possible explanation for the delay is that instead of a second N2pc, it is an increased contralateral delayed activity (CDA) (Vogel & Machizawa, 2004) due to the additional working memory load in the SS condition. To rule out this possibility we compared the SS waveform to the SL waveform. Both of these conditions have the same number of targets and therefore should have similar CDA

amplitude. This comparison revealed a very similar time window of 420-500ms, suggesting the second negativity in the SS condition is not due to working memory load. One explanation for why this negativity is later than expected is that the waveforms elicited by T1 and T2 overlap, resulting in the N2pc to the second target being summed in part with the post-N2pc positivity seen in response to T1. This could result in a slightly delayed second N2pc. Another contributing factor could be that the second target was presented within the inhibitory surround of the first lock-on state. This late N2pc suggests the T2's ability to recruit attention was delayed in the SS condition. Thus question 2: *Does the lock-on only occur at the location of the target or does it include a broader region in the vicinity of the target?* is answered in favor of specific location, with the caveat that there may be a delay in the deployment of attention to the T2 in the same-side condition.

The comparison of the S condition with the SL condition in the confirmatory analysis is a replication of the results found in Tan & Wyble. While the permutation analysis showed a significant difference starting at 500ms this was much later than the expected second N2pc latency (~375ms) given the T1/T2 SOA and thus is likely to reflect a CDA. However the cluster mass analysis did find a significant window close to the time of the expected second N2pc. This window was only 40ms in duration and so did not pass the criterion of the permutation analysis. This brief, second negativity is not predicted by the model and is similar to that found in the Ortho vs. Ortho-Ortho comparison of Experiment 1. We return to this in the general discussion.

Experiment 3

This experiment addressed question 3: *Can a lock-on state be maintained over a long period of time or does it expire automatically?* by presenting T2 at various lags from T1.The CGF model predicts that, given enough time between T1 and T2 presentation, the lock-on state will dissipate and therefore will have to be re-established for T2, resulting in a second N2pc. Experiment 3 manipulated the lag at which T2 was presented to test this prediction.

Method

Participants. Twenty-six students participated in Experiment 3. After filtering for accuracy and artifact rejections (see Experiment 1's method for details), Twenty-five had the minimum number of trials to be included in analyses.

Stimuli & Apparatus. The same stimuli and monitor that were used in Experiment 1 were also used in Experiment 3 except all letters and digits were black (RGB values: [0 0 0]).

Procedure. The procedure for Experiment 3 was similar to Experiment 1 with a few exceptions (Figure 11). In Experiment 3 targets were orthographic (digits among letter distractors). Four conditions were used to test what duration between targets is required to elicit a second N2pc. Participants were either presented with only one target (Single), two sequential targets, the second immediately after the first (Lag 1), two targets with one frame of distractors in between (Lag 2), or two targets with three frames of distractors in between (Lag 4). As in

Experiments 1, the second target always appeared on the same side as the first. Each participant completed 280 trials, evenly distributed among conditions and randomly intermixed within block. The average count for accurate trials after artifact rejection was M = 52 (SE = 2.1), M = 46 (SE = 2.4), M = 40 (SE = 2.3), M = 43 (SE = 2.2) for Single, Lag-1, Lag-2 and Lag-4 conditions respectively.

Results

Behavioral results. The mean accuracy rates and standard errors for each condition in experiment 3 were: M = .92 (SE = .01) for Single, M = .81 (SE = .02) for Lag-1, M = .71 (SE = .03) for Lag-2, and M = .77 (SE = .02) for Lag-4.

Confirmatory analysis. A 2 (contra/ipsi) x 4 (conditions) repeated measures ANOVA was run to confirm the existence of a first N2pc in all conditions in the 200-300ms time window after T1. The ANOVA revealed a significant effect of laterality, $F(1,24) = 88$, $p < .001$, $\eta_p = .79$. There was no significant effect of condition, $F(3,24) = .10$, $p = .96$, however there was a significant laterality x condition interaction, $F(3,72) = 8.8$, $p < .001$. All conditions tested separately showed a significant effect of laterality, $p < .001$. A permutation analysis revealed a significant difference between Single and Lag-1 at 240-480ms post T1 onset (Figure 12). Between Single and Lag-2 there was a significant difference between 220-300ms and 500-800ms. Between Single and Lag-4 there was a significant difference between 820-920ms. As an additional comparison of conditions with equal number of targets, we also compared Lag-1 with Lag-2 and found a significant difference between 260-400ms and 500-620ms.

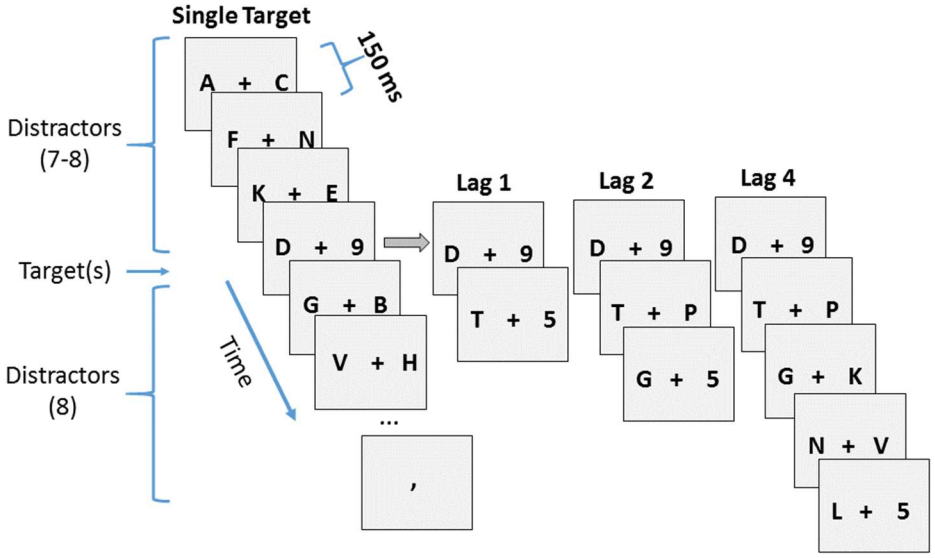

Figure 11. Paradigm for Experiment 3. Participants were presented with 2 RSVP streams. There were 7 or 8 frames of distractors (150ms SOA) prior to target presentation and 8 frames of distractors presented afterwards. Targets were black digits among black letter distractors. T2 was presented either at lag-1, 2, or 4.

Exploratory analysis. The cluster analysis showed a significant difference between Single and Lag-1 in the windows of 216-276ms and 296-388ms predominantly in the posterior electrode but extending frontally as well (Figure 13). There was a significant difference in posterior electrodes between Single and Lag-2 between the 480ms and 600ms. Between Single and Lag-4 there was a significant window between 792ms and 912ms, again predominantly in posterior electrode but extending into frontal electrodes as well.

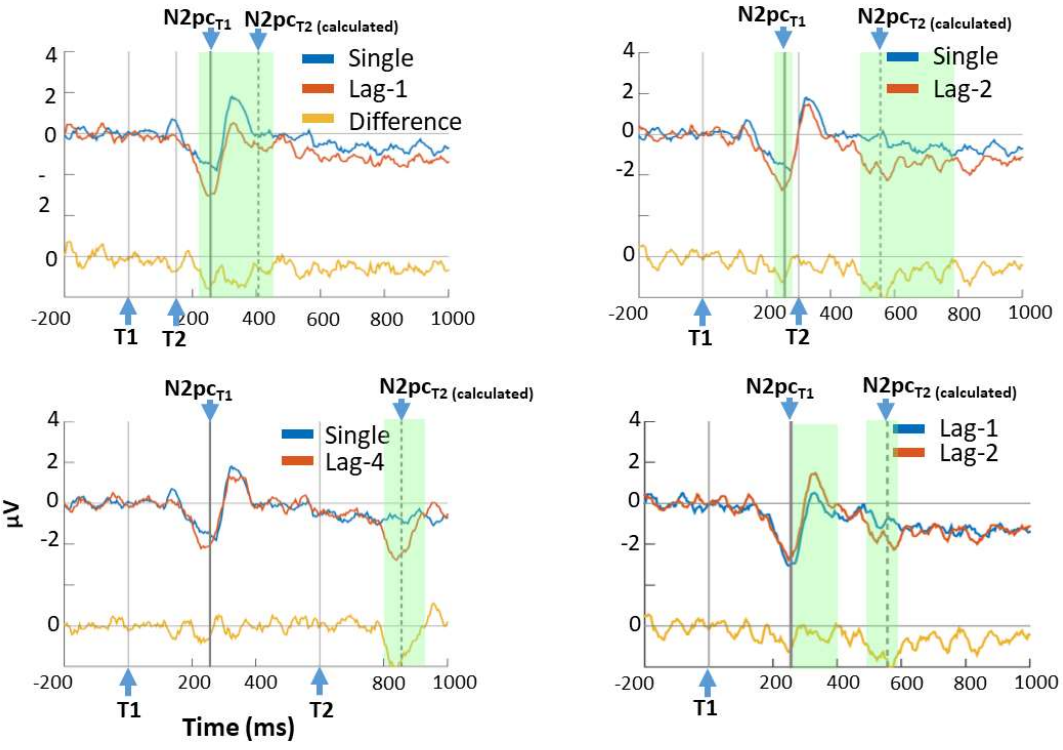

Figure 12. Difference wave comparisons for Experiment 3. Note that T2 onset and the predicted time of the N2pc elicited by T2 is marked explicitly in each graph as they change with T2 lag.

Discussion

The CGF model predicts that over time, the lock-on state will dissipate automatically, even if subjects expect a second target in the same location. More specifically, when T2 is presented after a sufficiently long interval, it will elicit a full second N2pc even if subjects expected the second target to be in that location. This prediction was clearly supported since a second target at Lag-4 produced a second N2pc at exactly the predicted time window. Both the permutation test and the cluster analysis supported this finding. Similarly, a significant negativity was found in the Lag-2 condition compared to the Single Condition. As the permutation analysis revealed a time window much broader than the typical N2pc, we also ran a comparison between Lag-2 and Lag-1, as the two conditions have the same number of targets and should help to rule out the possibility that the second negativity is increased or prolonged due to a CDA. The comparison

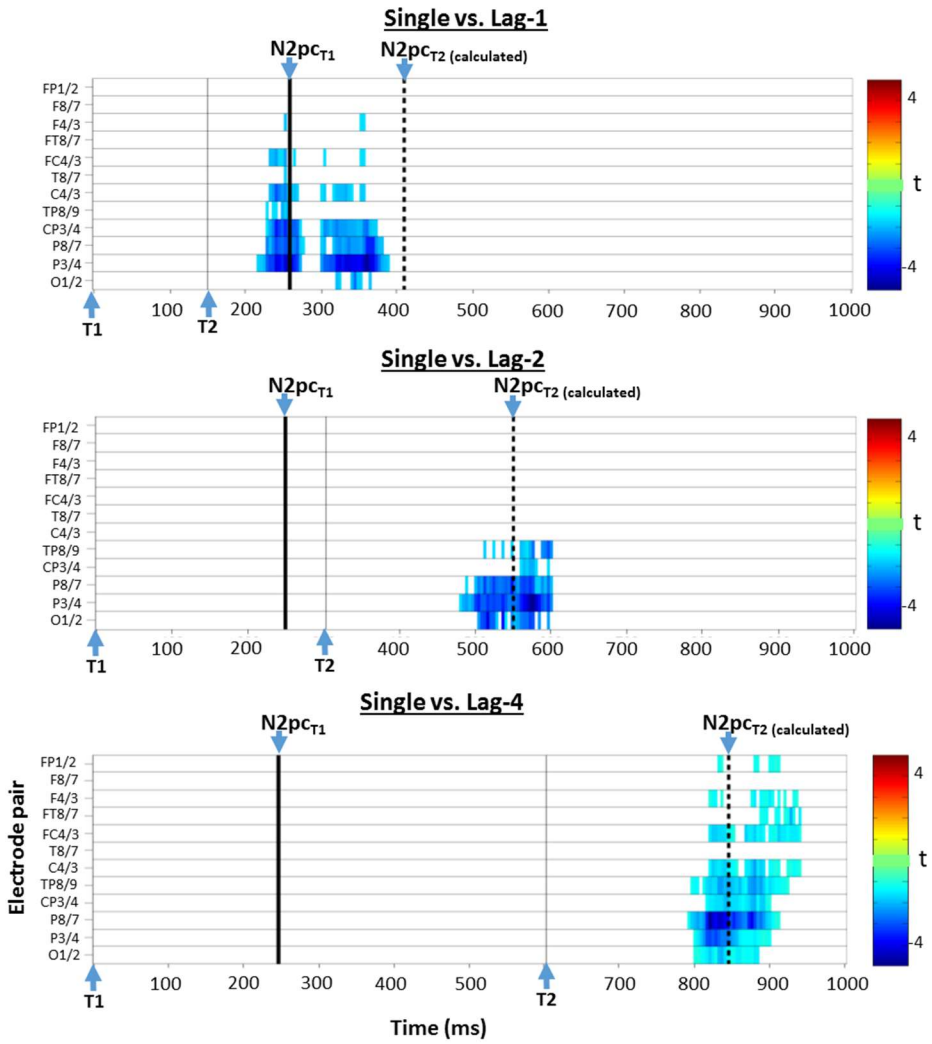

Figure 13. Cluster analysis comparing the Single condition with conditions Lag-1, Lag-2, and Lag-4 of experiment 3. Note that the presentation time of T2 and thus the time of its predicted N2pc are different for each comparison, as indicated by the vertical lines.

revealed a more focal difference at 500-620ms, which corresponds to the N2pc's expected latency.

Thus question 3: *Can a lock-on state be maintained over a long period of time or does it expire automatically?* is answered in favor of expiring automatically. Not only do these results confirm one of the CGF model's predictions, they also give us some insight into the nature of reflexive attention. In all conditions of this experiment, the second target was presented in the same stream as the first. That means that on every trial, participants knew the location of the second target once they had deployed attention to the first. An advantageous strategy should have been to keep attention at the T1's location for the next 600ms to ensure the second target was encoded. However, as the model predicted, the data suggest that the attentional lock-on state cannot be maintained through top-down control in the absence of relevant stimuli. Despite the fact that the second target's location was entirely predictable, attention needed to be redeployed to targets presented at later lags. So while the general prediction was validated, the results indicate that the lock-on state requires re-establishment at lag-2, which is earlier than the version of the CGF model published in Tan & Wyble had predicted. This will be discussed further in the general discussion.

Also, once again, there was the unexpected finding that a second target in the same location produces a slight additional negativity. The exploratory analysis broke this apart into two clusters. The first cluster shows that the first N2pc in the Lag-1 case is of greater amplitude than that to a single target. One could interpret this as two sequential targets eliciting a deeper N2pc. However the Lag-2 condition has a similarly deeper first N2pc compared to the single condition when the second target is presented after the ERP occurs (which means the presence of the T2 could not have directly influenced the N2pc). It is therefore unclear why the first N2pc is so much smaller in the Single condition than in the other conditions, but it does not seem to be influenced by a second target. The second cluster shows a significant negativity in the Lag-1 condition that is slightly earlier than the expected second N2pc. This result is not predicted by the CGF model but is consistent with the finding of both the first and second experiment. We will return to these issues in the discussion.

Experiment 4

Experiment 4 addressed question 4: *Is there a single attention map that is shared across target types?* by presenting subjects with two types of targets: featural (target colored letters among heterogeneously colored letters) and orthographic category (digits among letters). The CGF model predicts that because there is a unitary attention map, no second N2pc will be elicited even when sequential targets are of completely different types.

Method

Participants. Forty-two students participated in Experiment 4, ages 18-23 with normal or corrected-to-normal vision. All participants signed an informed consent form prior to the start of the experiment in accordance with the IRB office of Penn State. Twenty-five subjects passed exclusion criteria (see Experiment 1 for details) and were used in analysis.

Stimuli & Apparatus. The same apparatus and stimuli were used in Experiment 4 as in Experiment 1.

Procedure. A similar procedure was used in Experiment 4 as in Experiment 1 (Figure 14). Again, participants were looking for letters of a specified color (feature targets) or digits of any color (orthographic targets). Again, the target color remained constant throughout the experiment but was counterbalanced across participants. However in this experiment there were no single target conditions. Instead participants were presented with 4 different conditions: two sequential feature targets (Feat-Feat), two sequential orthographic targets (Ortho-Ortho), a feature target followed by an orthographic target (Feat-Ortho), or an orthographic target followed by a feature target (Ortho-Feat).

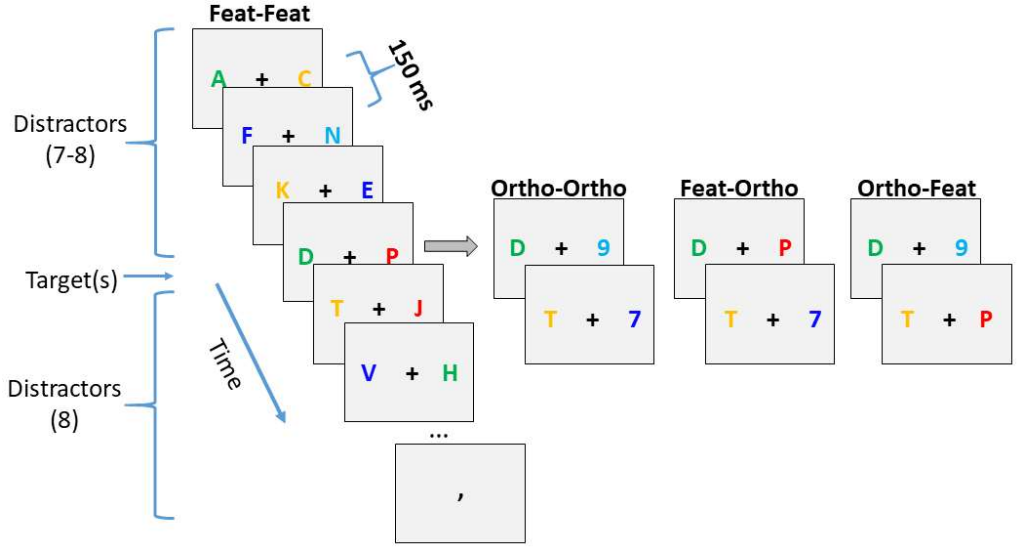

Figure 14. Paradigm for Experiment 4. Participants were presented with 2 RSVP streams. Targets were letters of a specified color or digits in any of the 4 remaining distractor colors. In this example, the target color is red but any of the 5 colors could be used as the target color for a given participant.

Each participant had a total of 272 trials equally divided among the four conditions. Trial types were intermixed within one block so participants did not know what type of targets to expect on any given trial. The average trial count per condition after inaccurate trials and those contain EEG artifacts were removed was M = 36 (SE = 2.4), M = 36 (SE = 2.3), M = 36 (SE = 2.2), M = 27 (SE = 1.9) for Feat-Feat, Ortho-Ortho, Feat-Ortho and Ortho-Feat conditions, respectively.

Results

Behavioral results. The mean accuracy rates and standard errors for each condition in experiment 4 were: M = .67 (SE = .03) for Feat-Feat, M = .64 (SE = .03) for Feat-Ortho, and M = .49 (SE = .02) for Ortho-Feat.

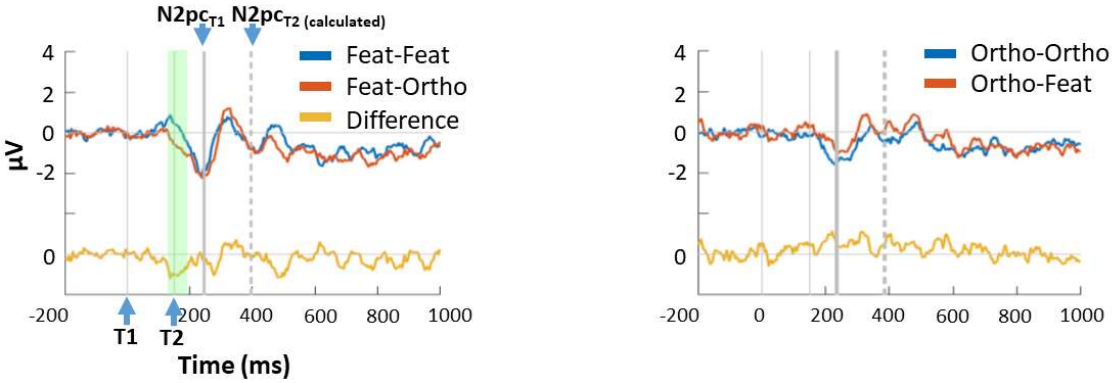

Figure 15. Difference wave from electrode pair P3/P4 for each condition in Experiment 4. The yellow line is the difference between the two waveforms being compared. Green boxes indicate windows of significant difference determined by the permutation analysis. Vertical lines show the timing of T1 and T2 presentation as well as the measured latency of the N2pc elicited by T1 and the calculated latency of T2's N2pc.

Confirmatory analysis. The two-way 2 (contra/ipsi) x 4 (conditions) repeated measures ANOVA in the 200-300ms window following T1 revealed a significant effect of laterality, $F(1,24) = 42$, $p < .001$, $\eta_p = .63$, but not of condition, $F(3,24) = 1.41$, $p = .25$. There was a significant interaction between laterality and condition, $F(3,24) = 3.53$, p .02. Testing each condition individually showed a significant effect of laterality in each condition at $p < .05$.

In the permutation analysis, Feat-Feat versus Feat-Ortho showed a significant difference during 120-180ms. whereas Ortho-Ortho versus Ortho-Feat showed no windows of significance (Figure 15).

Exploratory analysis. The exploratory analysis found no significant clusters when comparing the Ortho-Ortho and Ortho-Feat conditions, or the Feat-Feat and Feat-Ortho conditions.

Discussion

An assumption of the CGF model is that there is one universal attention map (either situated in a single cortical area or a network of interconnected areas) that responds to all target types which means that a lock-on state elicited by one type of target will carry forward to affect a subsequent target of a different type. However, if there exist separate attention maps for different target types (say, orthographic category and color-marked) then presenting two targets of different types sequentially would produce two N2pcs. This is because the second target would be activating a different attention map than the first, requiring a second localization stage to establish a second lock-on state. In Experiment 4, neither the comparison of Feat-Feat to Feat-Ortho nor that of Ortho-Ortho to Ortho-Feat showed any evidence of a second N2pc,which supports what the model would predict, given a unitary map of attention, or an interconnected network of attention-related areas.

Experiment 5

Both Experiments 1 and 3 showed a second early negativity to two sequential orthographic targets, which was not seen in the original Tan & Wyble results. One difference between the paradigms used here and that of Tan & Wyble is the slightly longer SOA of the current experiments (150 vs 120ms). It could be that the 30 ms difference in presentation rate is enough time for the lock-on state to begin to dissipate. This would require at least a partial re-establishment of the lock-on state for the second target, which could result in a weaker, earlier, second N2pc. To explore this possibility directly we ran an additional experiment comparing single and double target trials when stimuli were presented every 120 ms (as they were in Tan & Wyble) versus every 150ms (as they were in Experiments 1 and 3).

Method

Participants. Twenty-three students, ages 18-23 with normal or corrected-to-normal vision, participated in Experiment 5. All participants signed an informed consent form prior to the start of the experiment in accordance with the IRB office of Penn State. Twenty subjects passed exclusion criteria (see Experiment 1 for details) and were used in analysis.

Stimuli & Apparatus. The same apparatus and stimuli were used in Experiment 5 as in Experiment 3.

Procedure. A similar procedure was used in Experiment 5 as in Experiment 3. Experiment 5 had trials with one or two targets presented in RSVP streams with SOAs of 150 or 120ms. There were four conditions: Single-120, Single-150, Double-120, Double-150. Second targets were always presented at lag-1 in the same stream as the first. All trial types were intermixed within block.

Each participant had a total of 288 trials. Single target trials accounted for a third of the trials presented. As double target conditions required two targets to be accurately identified to be included, we presented twice as many double target trials than single to ensure enough usable trials. The average number of usable trials per condition were M = 31 (SE = 2.0), M = 35 (SE =

1.8), *M* = 50 (*SE* = 4.2), *M* = 59 (*SE* = 4.2), for Single-120, Single-150, Double-120, and Double-150 conditions, respectively. The larger number of usable trials in the double condition is a result of more of those trials being presented during the experiment.

Results

Behavioral results. The mean accuracy rates and standard errors for each condition in experiment

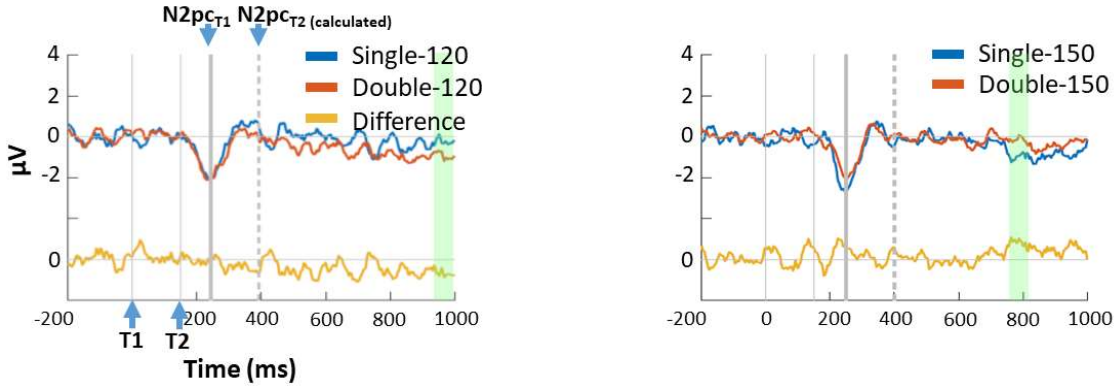

Figure 15. P3/P4 difference waves for each condition in Experiment 5. Highlighted in green are windows of significant difference determined by the permutation analysis. Difference between conditions is depicted in yellow. Vertical lines show the presentation time of T1 and T2 as well as the measured and calculated time of the first and second N2pc, respectively.

4 were: *M* = .84 (*SE* = .02) for Single-120, *M* = .90 (*SE* = .02) for Single-150, *M* = .65 (*SE* = .04) for Double-120, and *M* = .77 (*SE* = .04) for Double-150.

Confirmatory analysis. The two-way 2 (contra/ipsi) x 4 (conditions) repeated measures ANOVA in the 100-300ms window following T1 revealed a significant effect of laterality, $F(1,19) = 23$, $p < .001$, $\eta_p = .55$, but not of condition, $F(3,19) = ,85$, $p = .47$, $\eta_p = 4.26$, or an interaction, $F(3,19) = .77$, $p = .51$. Each condition individually showed a significant effect of laterality, $p < .01$.

In the permutation analysis, Single-120 versus Double-120 showed a significant difference during 940-1000ms and Single-150 versus Double-150 versus Ortho-Feat showed a significant difference during 760-820ms (Figure 16).

Exploratory analysis. The exploratory analysis revealed no significant clusters when comparing single and double target conditions of either SOAs.

Discussion

The results reveal no significant second negativity at the expected time of the second target's N2pc with either the longer or shorter SOA. The inconsistent nature of the second negativity seen in the orthographic conditions of Experiments 1 and 3 but not in equivalent conditions of

Experiments 2 and 5, or Tan & Wyble (2015) suggests that a second orthographic target produces a weaker lock-on state that is harder to detect, and/or varies according to task difficulty.

These results will be discussed in further detail in the General Discussion.

General Discussion

The ultimate goal of understanding the neural mechanisms that underlie attentional processes is going to require a cumulative approach to science in which models are developed and incrementally evaluated against their own predictions. In these experiments, we build on a previously published model (CGF, Tan & Wyble, 2015) by testing its key predictions, and now will close the inferential loop by evaluating what we have learned and how the model should be changed.

At the heart of this investigation is the idea that the visual attention system briefly enters an attractor state when attending to a stimulus. This state, termed a *lock-on* state in the CGF model, explains how a highly specific location can be attended, despite the larger receptive fields at superordinate stages of the visual hierarchy. The advantage of this approach is that it provides a clearly specified account of how neural mechanisms interact across levels of the visual hierarchy to produce the correlates of reflexive attention. Moreover, it explains counterintuitive findings, such as the fact that it seems to last for only a brief period, even when it would be advantageous to extend its deployment (e.g. Nakayama & Mackeben, 1989). The other key advantage of this model-based approach is that it provides a means to link underlying neural mechanisms to observable EEG components. This represents a fundamental shift in our approach to understanding ERPs because it provides an intuition of what those neural correlates represent in terms of changes in neuronal membrane potentials.

Ruling out alternative theories of reflexive attention

These predictions provide support for the CGF model by ruling out alternatives. Here we explore four predictions and will review each in turn as to what alternative theories these results rule out:

Experiment 1 demonstrated that color-marked targets and orthographic targets are treated similarly by reflexive attention. While the results of Tan & Wyble supported the model's attribution of the N2pc as a neural correlate of localization of attention, one theory could have been that this lock-on state was unique to categorically specified targets. This alternative would predict that two sequential color-marked targets would either invoke two N2pcs or perhaps a single N2pc of twice the duration. However Experiment 1's results show that this is not the case and suggested that the same mechanisms used for triggering attention by categorical targets are also used by feature-marked targets.

Another alternative theory that is ruled out by these experiments is that reflexive attention is triggered toward a large region of the visual field rather than a specific location. In previous experiments, the correlate of the lock-on was demonstrated by contrasting same-location vs

different-side targets. Those findings could equally support an alternative account in which the N2pc indicates attention towards a side of the visual field, rather than a location. This theory would predict a single N2pc to two sequential targets presented in the same hemifield regardless of being in the same location. The results of Experiment 2 show that this is not the case as a second N2pc is only elicited when the second target is presented in a different location than the first. These results rule out any alternative theory of reflexive attention that does not require its deployment to be location specific. However the small size of the second N2pc suggests the need to modify the model, as discussed below.

While Experiment 2's results set explicit spatial constraints on any theory of reflexive attention, Experiment 3's results set a temporal constraint. One theory of reflexive attention could be that once it is deployed to a location, it can remain at that location in the absence of relevant stimuli. However in Experiment 3 participants knew exactly where to expect the second target (as it always appeared in the same location as the first) and yet the results show that the lock-on state still needed to be re-established at later lags. We attribute the re-establishment of the lock-on state at later lags to the amount of time that has elapsed between targets allowing the original lock-on state to dissipate. However another difference between lag-4 and lag-1 is the number of intervening distractors, perhaps causing the lag-4 target to experience additional interference from the additional distractors. This explanation would apply a temporal interpretation of the spatial proximity effect of distractors on the N2pc amplitude (Luck & Hillyard, 1994; Mazza, Turatto, & Caramazza, 2009). However, typically in RSVP, it is the distractors on either side of a target that provide the interference, through forward and backward masking (Robitaille & Jolicœur, 2006). Moreover behavioral work suggests that time is the influential factor in closing an attentional window as both blank space and intervening distractors can cause an attentional blink (Nieuwenstein, Potter, & Theeuwes, 2009). Furthermore, beyond lag-2, increasing either the amount of time or distractors between two targets in an AB paradigm leads to better performance for the T2, which suggests that the additional distractors are not increasing the difficulty of attending to the T2.An important next step for future work is to run this experiment without intervening distractors between T1 and T2 as the CGF model would predict the same results found here.

Finally, the model proposes a single unifying map of attention for all target types. An alternative assumption would have been that each target type has its own attention map. The prediction of such a theory would be that two sequential targets of mixed type would require the establishment of two lock-on states (one in each map) resulting in two N2pcs. Experiment 4's results showed that this was not the case as the lock-on state established by either a feature or orthographic target, could be carried over to a second target of a different type. While the ERPs show no effects of switching between target types, the behavioral data shows a modest accuracy cost when two sequential targets are of different types. The model in its current form enhances all processing at the location of a target in the EV rather than enhancing processing of just the target attribute in one of the LVs. Thus, the model has no explanation for costs when attending to two

sequential targets defined by different criteria. As such, the CGF model predicts that set-switching costs seen in the behavioral occur at subsequent stages of processing. For example it may be more difficult to encode two distinct types of targets into memory. While these results support the theory of a universal attention map, they do not rule out the possibility of additional maps entirely. One way to explore this further could be to use even more distinct target types. Faces, for example, are a stimuli type that the brain is highly specialized in. A consequence of this specialization could be an independent attention map which would result in two N2pcs to the presence of an orthographic and face target presented sequentially.

Failed predictions and model revisions

While the specific predictions in Tan & Wyble (2015) were validated by the data, there were some findings that did not completely align with the CGF model simulations. Such findings provide important insight into what architectural changes or parameter adjustments would bring the model into closer register with the actual system. To do this, we first consider what mechanisms may have contributed to the unexpected results. We then make the appropriate adjustments and test whether this change improved the match between simulated and real data. These parameter changes are potential solutions for the discrepancies seen here and should be explored further in future work. For these experiments, the two unexpected findings were the second negativity to two sequential orthographic targets (Experiments 1 and 3), and the delayed second N2pc to a T2 presented on the same side but not the same location as T1 in Experiment 2. These findings will be discussed in turn as well as how they guided revisions of the model.

Second negativity in a two-target condition. In two out of three experiments where double target conditions were compared to single target conditions, a second negativity was seen in the sequential presentation of two orthographic targets at the same location, which was not seen in Tan & Wyble (2015). This second negativity was smaller than the first N2pc and sometimes earlier than it should have been had it been time-locked to the T2. One contributing factor could have been the slightly longer SOA used in these experiments. To test this explicitly we ran an additional study with conditions that presented stimuli at either 120 or 150ms SOA. No second N2pc was observed for two sequential targets with either SOA. The unreliability of the second negativity combined with its smaller size and early latency suggests a partial lock-on phenomenon, as would occur if the initial lock-on state had partially decayed when the T2 input arrived. Another important constraint on possible model adjustments is that this second negativity was not seen with color-marked targets. Thus, any change to the model would have to effect orthographic sequential targets and not color-marked ones.

One difference between color-marked and orthographic target tasks is that there is only one feature type by which color-marked targets are identified (e.g. red) whereas there are multiple types by which to identify an orthographic target (e.g. 1- 9 are all members of the digit target set). The LV layer of the CGF model was envisioned to represent late stage visual processing neurons that discriminate between different kinds of stimuli. However, in the Tan & Wyble implementation, for the sake of parsimony there was only one LV map containing one neuron at each location, which meant that T1 and T2 would both activate the same neuron, even if those targets were different digits or letters. This implementation is plausible for the Feat-Feat conditions as a single feature (e.g. the color red) indicated the presence of a target, and thus that feature's representation would remain active when T2 followed T1 in the same location. In terms of simulation this means there is no interruption of LV activity during the transition from T1 to T2. Because of this, excitation to the AM remains steady throughout the presentation of T1 and T2, and there is no disruption of the lock-on state from one target to the next. However, in the case of orthographic targets, this singular LV map solution is less appropriate. It is known that the visual system contains neurons that have highly specific firing patterns across different shapes, and thus it is unrealistic to assume a single set of neurons that would be activated by "digits". A more accurate simulation would use different nodes to represent different digits. The new version of the model, shown in Figure 16 illustrates the architectural change.

This change, by itself, was not enough to make the lock-on state sufficiently fragile to reproduce the second, very small N2pc for the Ortho-Ortho conditions. Therefore, another change we introduced is to make the CGF neurons sensitive to onset of a stimulus. This reflects known properties of neurons in IT cortex of the macaque, where firing rates are increased in response to the onset of a preferred stimulus within the receptive field, but then decrease quickly despite the stimulus still being present (Chelazzi et al., 1998). To implement this change, the excitatory input from a stimulus was shortened so that only the first 100ms of a stimulus sends excitation into EV. These changes combined with an increase to the rate of change for each neuron (*dtvm*) and a decrease to the strength of lateral inhibition that neurons project on to their neighbors

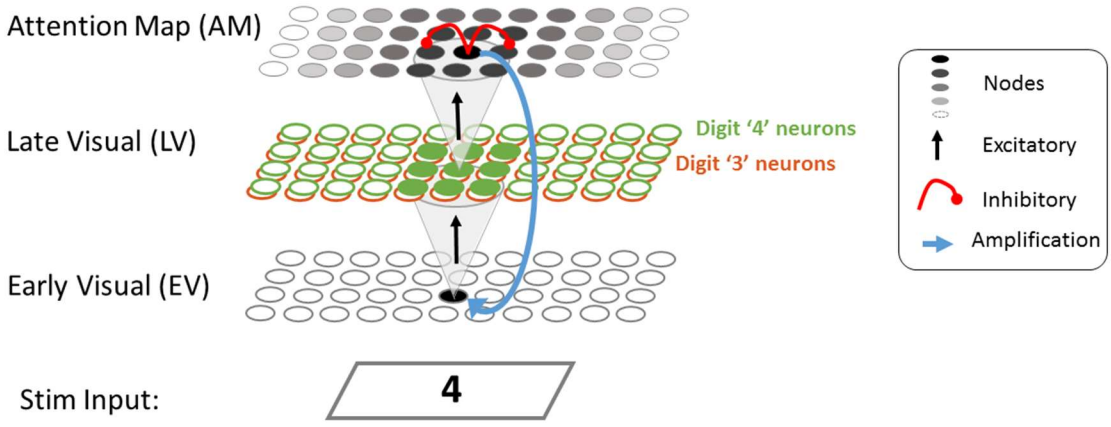

Figure 16. Schematic of the revised CGF model. The new version has a different pool of neurons that would be specifically tuned for different types of stimuli.

(*Inhib*), result in effectively reducing the rigidity of a lock-on state, making the system a bit more elastic. Consequently, the model produces a weaker, earlier N2pc to two sequential orthographic target but not when the targets are color defined (Figure 17). The weakness of the second negativity is such that it could be drown out by noise and so not reliably seen across experiments.

These changes in the lability of the attentional lock-on state produced another improvement in the model's simulations, which is that the original simulation predicted no second N2pc for a lag-2 T2, whereas Experiment 3 revealed a small N2pc at lag-2. These results suggest that the lock-on state collapses a bit more quickly than the original Tan & Wyble (2015) parameters allowed and these changes bring the model into closer register with the system.

Attention deployed later to second target in a different location from the first. In the second experiment we observed that while there was a second negativity in the Same Side condition to the second target, as predicted by the model, that N2pc was later and smaller than predicted. This could be due in part to the fact that the waveforms elicited by T1 overlap with those of T2. This would mean that the post-N2pc positivity seen in the single target condition would subtract from the N2pc elicited by a second target when presented on the same side and add to the N2pc of a second target presented on the opposite side. However another difference between Different Side and Same Side conditions of Experiment 2 is that while the stimuli were presented at equal eccentricity from fixation, the streams were organized such that there was a 5.5 degree horizontal separation between streams on opposite sides and a 3.5 degree vertical separation between streams on the same side. Therefore, the closer T2 in the SS condition was more likely to fall

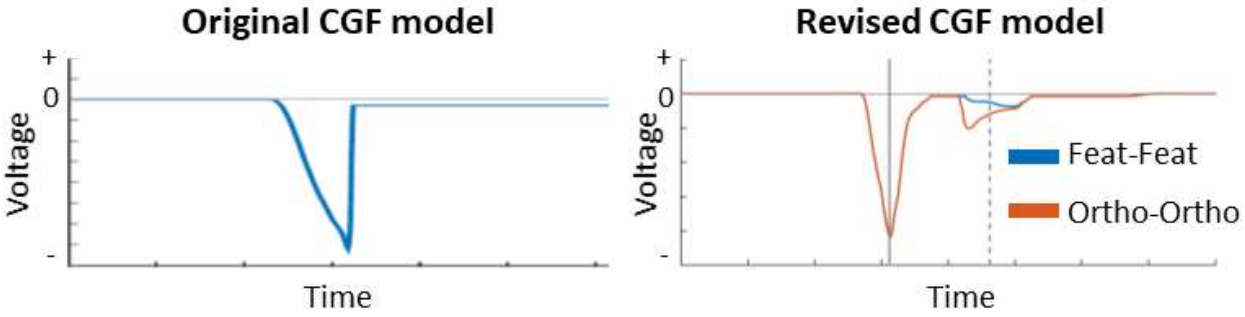

Figure 17. Simulated difference waves from the original and revised versions of the model. [Left] Simulated waveform from the original CGF model to two sequential targets in the same location. [Right] Because color marked targets (blue line) excite the same pool of neurons (neurons that are specifically tuned for the target color) there is no break in the excitation to the AM. However, because two different digits were used for T1 and T2 in the Ortho-Ortho condition (red line), two separate groups of neurons in the LV needed to be excited above threshold to pass excitation onto the AM. This causes a brief gap in the excitatory input to the AM and therefor some decay in the lock-on state, resulting in a small and early N2pc relative to the expected second N2pc time (dotted line).

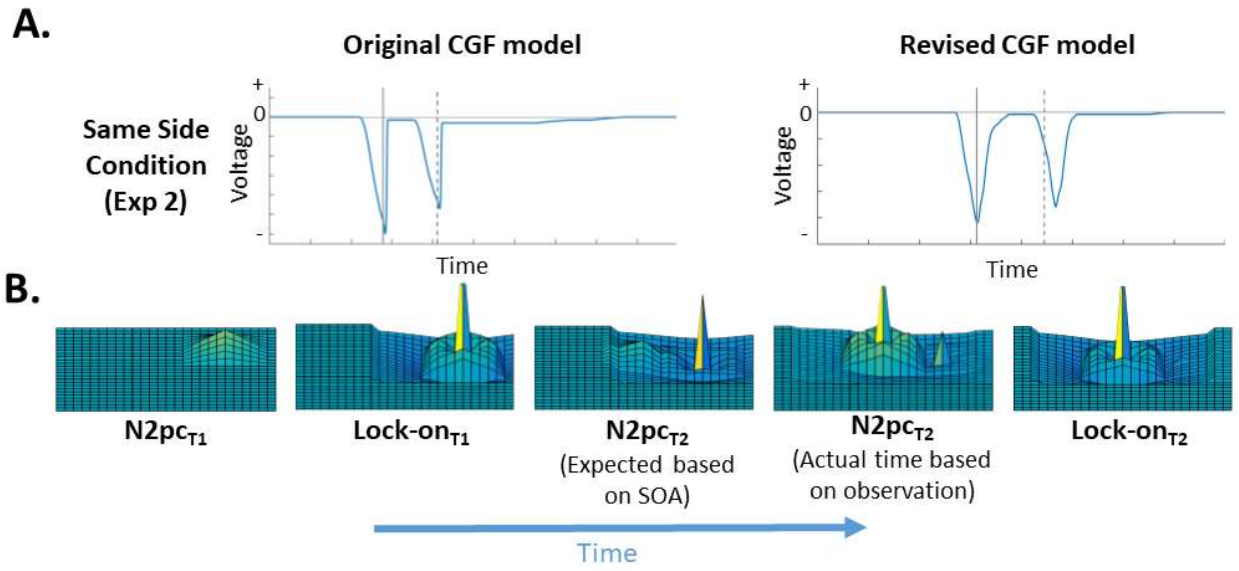

Figure 18. (A) Simulated difference waves from the original and revised model when presented with two sequential targets on the same side but not location of the visual field. The revised model has a larger range of lateral inhibition within the AM, which causes the rise of a spatially proximal T2's bump in the AM to be delayed, and thus delays the N2pc. (B) Surface plots of the AM activation through time in the revised model. Importantly at the time of the expected N2pc to T2, the excitation is being suppressed by the inhibition caused by T1's lock-on state. As T1's lock-on state decays, neurons fall below threshold and the inhibition is released, allowing for the delayed establishment of T2's lock-on state.

directly within the inhibited region produced by theT1's lock-on state. This inhibition would delay the formation of the localization bump within the AM, which would cause a latency shift of the N2pc and a reduction of its amplitude. Other studies have found similar behavioral effects of this inhibitory surround when presenting target 4° versus 6.8° (Dell'Acqua, Pascali, Jolicœur, & Sessa, 2003).

To allow the CGF model to simulate these results only one parameter change was needed. We increased the radius of inhibition in the AM, which had previously been set at 3.5 degrees, to 7.5 degrees. With this change, the model could simulate the data from Experiment 2 (Figure 18).

Shallow first N2pc to a single target in Experiment 3. Experiment 3 showed an abnormally small first N2pc to a single target compared to that of the Lag-1 condition, which is not predicted by the model. We believe this effect is an anomaly as it does not replicate in any of the other experiments presented here, nor even within experiment 3. Comparing the first N2pc for the Single condition with that of the Lag-2 condition shows a significant difference even though at the time of the N2pc, only one target has been presented to the visual system (see Figure 12). As only one condition in one experiment shows this shallow N2pc to a single target, we do not feel it warrants a change in the model to accommodate this result.

Future directions

Future directions for developing the CGF model will be to widen the breadth of findings the model can account for. For instance, as the model is currently structured, all targets are processed at the same speed resulting in equal N2pc latency for all target types. However recent work has showed that the N2pc latency is dependent on how attention is cued to a location (Jenkins, Grubert, & Eimer, 2018) as well as how a target is defined among distractors (Callahan-Flintoft & Wyble, 2017). Future work will look to investigate what neural mechanisms may drive these latency differences. Moreover, the model has no explicit simulation of differences in behavioral accuracy, and adding such mechanisms would tremendously expand the set of constraints that apply to the model.

Another direction that this work will explore is to explain the neural underpinnings of the Pd component, a contralateral positivity often associated with distractor suppression (Hickey, Di Lollo, & McDonald, 2009) or the termination of attention (Sawaki, Geng, & Luck, 2012). In future work we plan to use the CGF model to link this component to neural mechanisms, similarly to how the N2pc was explored here. Not only might this give new insight into what processes the Pd reflects but it might also enhance our understanding of the N2pc as the two components often appear within a single paradigm (Kiss et al., 2012). Linking these two components to underlying neural mechanisms will enhance our understanding of attentional selection and continue to guide future experimental work.

The revised version of the CGF model, along with scripts to simulate the data from these experiments can be found at [https://osf.io/zfh2j/?view_only=46b90b02a50742bab3e01b40c27801a1]. All code for running the experiments presented here as well as the analysis scripts and data files can be found at [https://osf.io/u5fnr/?view_only=c7d18747a7284778bfcd84766ca3d282]

Conclusion

The four experiments presented here answer four questions about the hypothesized attentional lock-on state of transient attention by testing predictions of the CGF model. The lock-on state was replicated for color-marked targets as well as those defined by orthographic category. The N2pc reflects attention to the location, not just the hemifield of visual attention. The neural mechanisms underlying reflexive attention at a given location decay over the course of several hundred milliseconds, even though subjects might expect an additional target at that location. Thus attention must be reallocated to a second target and this transience cannot be overridden by the knowledge that T2 will appear in the same location as T1. Finally, we found that an attentional lock-on state can be carried over to a second target presented at that location even when the second target is defined by a different criteria than the first (i.e. color-marked versus orthographic). This lends support to our assumption that different target types are processed within a universal attention map.

In addition to these supported predictions there were also a number of unexpected findings that informed adjustments of the model such as the specificity of neurons in the late visual layer and the range of lateral inhibition in the attention map.

Context paragraph

The current work is the culmination of an ongoing attempt to build a computationally formalized model of the neural underpinnings of reflexive visual attention. Earlier work demonstrated that visual targets trigger the deployment of attention in one location (Wyble, Potter, & Bowman, 2009). Subsequently we studied this effect using ERPs involving two sequential targets which revealed the core finding of a single N2pc to two targets. That finding led to the development of the first iteration of the CGF model (Tan & Wyble, 2015) and laid a roadmap for further investigation of the spatial and temporal dynamics of reflexive attention. This work builds on these previous findings by addressing key questions posed by the Tan & Wyble model and implementing a more advanced model that includes multiple kinds of targets and more accurately simulates the spatial and temporal extent of attentional deployment. The implications of this work extend to our lab's efforts to better understand visual cueing effects (Bay & Wyble 2014) and the attentional blink (Bowman & Wyble 2007; Wyble, Bowman & Nieuwenstein 2009; Swan & Wyble 2015) as well as other effects such as attentional capture and scenarios in which either eye movements or changing stimuli cause visual targets to be fleeting in nature.

Acknowledgement

This project was supported by NSF Grant 1734220 to Brad Wyble, and grants from National Natural Science Foundation of China (No. 31771201), Humanities and Social Sciences Foundation of the Ministry of Education of China (No.17YJA190001) to Hui Chen.

References

Bay, M., & Wyble, B. (2014). The benefit of attention is not diminished when distributed over two simultaneous cues. *Attention, Perception, & Psychophysics*, *76*(5), 1287-1297.

Bowman H., & Wyble B. (2007) The simultaneous type, serial token model of temporal attention and working memory. *Psychological Review*, 114(1):38-70

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial vision*, *10*, 433-436.

Brisson, B., & Jolicoeur, P. (2007). The N2pc component and stimulus duration. *Neuroreport*, *18*(11), 1163-1166.

Burra, N., & Kerzel, D. (2013). Attentional capture during visual search is attenuated by target predictability: evidence from the N2pc, Pd, and topographic segmentation. *Psychophysiology*, *50*(5), 422-430.

Cheal, M. L., Lyon, D. R., & Gottlob, L. R. (1994). A framework for understanding the allocation of attention in location-precued discrimination. *The Quarterly Journal of Experimental Psychology*, *47*(3), 699-739.

Chelazzi, L., Duncan, J., Miller, E. K., & Desimone, R. (1998). Responses of neurons in inferior temporal cortex during memory-guided visual search. *Journal of neurophysiology*, *80*(6), 2918-2940.

Craston, P., Wyble, B., Chennu, S., & Bowman, H. (2009). The attentional blink reveals serial working memory encoding: Evidence from virtual and human event-related potentials. *Journal of Cognitive Neuroscience*, 21(3), 550-566.

Dell'Acqua, R., Pascali, A., Jolicoeur, P., & Sessa, P. (2003). Four-dot masking produces the attentional blink. *Vision Research*, *43*(18), 1907-1913.

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of neuroscience methods*, *134*(1), 9-21.

DiCarlo, J. J., & Maunsell, J. H. (2003). Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. *Journal of neurophysiology*, *89*(6), 3264-3278.

Eimer, M. (1996). The N2pc component as an indicator of attentional selectivity. *Electroencephalography and clinical neurophysiology*, *99*(3), 225-234.

Folk, C. L., Leber, A. B., & Egeth, H. E. (2002). Made you blink! Contingent attentional capture produces a spatial blink. *Attention, Perception, & Psychophysics*, *64*(5), 741-753.

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology Human Perception and Performance*, *18*, 1030-1030.

Folk, C. L., Remington, R. W., & Wright, J. H. (1994). The structure of attentional control: Contingent attentional capture by apparent motion, abrupt onset, and color. *Journal of Experimental Psychology-Human Perception and Performance*, *20*(2), 317-329.

Girelli, M., & Luck, S. J. (1997). Are the same attentional mechanisms used to detect visual search targets defined by color, orientation, and motion?. *Journal of Cognitive Neuroscience*, *9*(2), 238-253.

Gottlob, L. R., Cheal, M., & Lyon, D. R. (1999). Time course of location-cuing effects with a probability manipulation. *The Journal of general psychology*, *126*(3), 261-270.

Groppe, D. M., Urbach, T. P., & Kutas, M. (2011). Mass univariate analysis of event-related brain potentials/fields I: A critical tutorial review. *Psychophysiology*, *48*(12), 1711-1725.

Hickey, C., Di Lollo, V., & McDonald, J. J. (2009). Electrophysiological indices of target and distractor processing in visual search. Journal of cognitive neuroscience, 21(4), 760-775.

Hickey, C., McDonald, J. J., & Theeuwes, J. (2006). Electrophysiological evidence of the capture of visual attention. *Journal of cognitive neuroscience*, *18*(4), 604-613.

Hopfinger, J. B., & Mangun, G. R. (1998). Reflexive attention modulates processing of visual stimuli in human extrastriate cortex. *Psychological science*, *9*(6), 441-447.

Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on pattern analysis and machine intelligence*, *20*(11), 1254-1259.

Jenkins, M., Grubert, A., & Eimer, M. (2018). The speed of voluntary and priority-driven shifts of visual attention. *Journal of experimental psychology: human perception and performance*, *44*(1), 27.

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

Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A dynamic neural field model of visual working memory and change detection. *Psychological science*, *20*(5), 568-577.

Kiss, M., Driver, J., & Eimer, M. (2009). Reward priority of visual target singletons modulates event-related potential signatures of attentional selection. *Psychological science*, *20*(2), 245-251.

Kiss, M., Grubert, A., Petersen, A., & Eimer, M. (2012). Attentional capture by salient distractors during visual search is determined by temporal task demands. *Journal of cognitive neuroscience*, *24*(3), 749-759.

Kiss, M., Jolicœur, P., Dell'Acqua, R., & Eimer, M. (2008). Attentional capture by visual singletons is mediated by top-down task set: New evidence from the N2pc component. *Psychophysiology*, *45*(6), 1013-1024.

Kröse, B. J., & Julesz, B. (1989). The control and speed of shifts of attention. *Vision research*, *29*(11), 1607-1619.

Luck, S. J. (2014). *An introduction to the event-related potential technique*. MIT press.

Luck, S. J., & Hillyard, S. A. (1994). Spatial filtering during visual search: evidence from human electrophysiology. *Journal of Experimental Psychology: Human Perception and Performance*, *20*(5), 1000.

Nieuwenstein, M. R., Potter, M. C., & Theeuwes, J. (2009). Unmasking the attentional blink. *Journal of Experimental Psychology: Human perception and performance*, *35*(1), 159.

Mazza, V., Turatto, M., & Caramazza, A. (2009). Attention selection, distractor suppression and N2pc. *cortex*, *45*(7), 879-890.

Mazza, V., & Caramazza, A. (2011). Temporal brain dynamics of multiple object processing: the flexibility of individuation. *PloS one*, *6*(2), e17453.

Mounts, J. R. (2000). Attentional capture by abrupt onsets and feature singletons produces inhibitory surrounds. *Attention, Perception, & Psychophysics*, *62*(7), 1485-1493.

Müller, H. J., & Rabbitt, P. M. (1989). Reflexive and voluntary orienting of visual attention: time course of activation and resistance to interruption. *Journal of Experimental psychology: Human perception and performance*, *15*(2), 315.

Nakayama, K., & Mackeben, M. (1989). Sustained and transient components of focal visual attention. *Vision research*, *29*(11), 1631-1647.

Nako, R., Wu, R., Smith, T. J., & Eimer, M. (2014). Item and category-based attentional control during search for real-world objects: Can you find the pants among the pans?. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1283.

Robitaille, N., & Jolicoeur, P. (2006). Fundamental properties of the N2pc as an index of spatial attention: effects of masking. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *60*(2), 101.

Tan, M., & Wyble, B. (2015). Understanding how visual attention locks on to a location: Toward a computational model of the N2pc component. *Psychophysiology*, *52*(2), 199-213.

Sawaki, R., Geng, J. J., & Luck, S. J. (2012). A common neural mechanism for preventing and terminating the allocation of attention. Journal of Neuroscience, 32(31), 10725-10736.

Silver, M. A., Ress, D., & Heeger, D. J. (2005). Topographic maps of visual spatial attention in human parietal cortex. *Journal of neurophysiology*, *94*(2), 1358-1371.

Töllner, T., Müller, H. J., & Zehetleitner, M. (2011). Top-down dimensional weight set determines the capture of visual attention: Evidence from the PCN component. *Cerebral Cortex*, *22*(7), 1554-1563.

Tsotsos, J. K., & Rothenstein, A. (2011). Computational models of visual attention. *Scholarpedia*, *6*(1), 6201.

Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*(6984), 748-751.

Wyble, B., Bowman, H., & Potter, M. C. (2009). Categorically defined targets trigger spatiotemporal visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(2), 324.

Yeshurun, Y., & Carrasco, M. (2008). The effects of transient attention on spatial resolution and the size of the attentional cue. *Perception & Psychophysics*, *70*(1), 104-113.