A Neurocomputational Model of Posttraumatic Stress Disorder

Gregory P. Davis[1], Garrett E. Katz[2], Daniel Soranzo[3], Nathaniel Allen[3], Matthew J. Reinhard[3,4], Rodolphe J. Gentili[5,6,7], Michelle E. Costanzo[3], James A. Reggia[1]*

***Abstract*— Despite the well-defined behavioral criteria for posttraumatic stress disorder (PTSD), clinical care is complicated by the heterogeneity of biological factors underlying impairment. Eye movement tasks provide an opportunity to assess the relationships between aberrant neurobiological function and non-volitional performance metrics that are not dependent on self-report. A recent study using an emotional variant of the antisaccade task demonstrated attentional control biases that interfered with task performance in Veterans with PTSD. Here we present a neuroanatomically-inspired computational model based on gated attractor networks that is designed to replicate oculomotor behavior on an affective antisaccade task. The model includes the putative neural circuitry underlying fear response (amygdala) and top-down inhibitory control (prefrontal cortex), and is capable of generating testable predictions about the causal implications of changes in this circuitry on task performance and neural activation associated with PTSD. Calibrating the model with the results of behavioral and neuroimaging studies on patient populations yields a pattern of connectivity changes characterized by increased amygdala sensitivity and reduced top-down prefrontal control that is consistent with the fear conditioning model of PTSD. In addition, the model makes experimentally verifiable predictions about the consequences of increased prefrontal connectivity associated with cognitive reappraisal training.**

Keywords: **posttraumatic stress disorder, antisaccade task, inhibitory control deficits, attentional bias, cognitive control**

I. Introduction

PTSD is a neuropsychiatric illness clinically defined by behavioral symptoms such as intrusive thoughts, avoidance, and hyperarousal that persist following exposure to a traumatic event and result in functional impairment. According to the U.S. National Center for PTSD, estimates of PTSD prevalence range between 5-10% of the US population, and up to 30% in highly exposed populations such as combat veterans. Despite its well-defined criteria, the behavioral symptoms of PTSD are dependent on self-report and can involve heterogeneous biological factors [1]. The biological complexity of the disorder and lack of performance markers make it difficult to predict the appropriate treatment approach for a particular patient.

*Research supported by awards 18010252, VA24517P3981, and VA24517P3971 from the US Dept. of Veterans Affairs, NSF Award DGE-1632976, and ONR award N00014-19-1-2044.

[1]Department of Computer Science, University of Maryland, College Park, MD, USA [2]Department of Electrical Engineering and Computer Science, Syracuse University, Syracuse, NY, USA [3]WRIISC-DC, VA Medical Center, Washington DC, USA [4]Georgetown University Medical Center, Washington DC, USA [5]Department of Kinesiology, School of Public Health, University of Maryland, College Park, MD, USA [6]Neuroscience and Cognitive Science Program, University of Maryland, College Park, MD, USA [7]Maryland Robotics Center, University of Maryland, College Park, MD, USA Correspondence to: Gregory P. Davis, grpdavis@umd.edu

The fear conditioning model provides a rich theoretical architecture for understanding how trauma may cause neurobiological and behavioral abnormalities [2], [3]. According to this model, environmental triggers of fear responses become generalized to non-aversive contexts via associative learning. The primary neurobiological consequence of fear conditioning is increased amygdala activation [4]. Fear responses are regulated by reciprocal inhibition between the amygdala and ventromedial prefrontal cortex (vmPFC) [5]. Functional neuroimaging studies reveal a common pattern of amygdala hyperactivation and vmPFC hypoactivation in PTSD [6]. This indicates a key regulatory imbalance in fear conditioning neurocircuitry, but the specific neurobiological process in humans is still not clear [7]. Is the observed pattern a result of impaired prefrontal inhibition, amygdala sensitization, or a combination of both?

Greater understanding of the connection between dysregulated frontolimbic interactions and fear-related psychophysiology may inform therapeutic efforts to restore balance in function. Experimental paradigms challenge the regulatory processes that may be compromised in PTSD [2]. The antisaccade task has been commonly used to test inhibitory control by tracking eye movement reactions to stimuli [8]. This task requires recruitment of prefrontal regions to inhibit the pre-potent instinct to look towards a dominant target (i.e., a prosaccade) before executing the correct saccade away from a central fixation point. A recent study found that facial stimuli reliably interfered with performance on an antisaccade task more than neutral stimuli in PTSD [9]. Because the amygdala contributes to visual orientation towards emotionally salient stimuli such as human faces [10], these performance deficits may be explained by the frontolimbic regulatory imbalances characteristic of PTSD.

Computational modeling is one means of assessing the relationships between behavioral and neurobiological changes in psychiatric disorders. Here we introduce a neuroanatomically-inspired computational model that extends gated cortical networks [11], [12] by adding subcortical circuitry for visual and oculomotor processing in order to carry out cognitive operations implementing the antisaccade task. The model also includes a frontolimbic circuit modeled after the amygdala and vmPFC. We investigate how changes in the connectivity of this circuit impact task performance, and identify changes that account for the behavioral results of [9] and the characteristic BOLD responses demonstrated in fMRI studies [1], [6]. In addition, we investigate the impacts of increased prefrontal connectivity associated with cognitive reappraisal training [5], and propose a testable prediction that

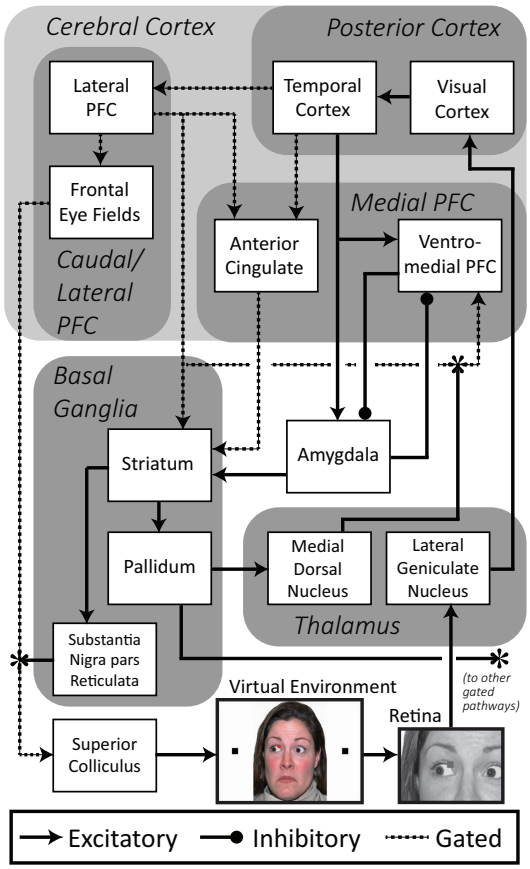

Fig. 1. High-level diagram of the neurocomputational model coupled with the antisaccade environment (bottom center). Connections ending in small circles are inhibitory, and dashed lines correspond to connections that are gated (asterisks) by the output structures of the basal ganglia: the substantia nigra pars reticulata gates motor output (frontal eye fields - superior colliculus, left), while the pallidum controls gating of cortical connectivity through the medial dorsal nucleus (lateral PFC - ventromedial PFC, center) and additional thalamic nuclei that are not modeled explicitly (remaining dashed connections, top). The gates on these remaining connections (e.g. temporal cortex - lateral PFC) are represented by an additional line exiting the pallidum (shown as a dangling asterisk on the bottom right). See text for details of model functionality.

such training attenuates antisaccade performance deficits in patients with PTSD by compensating for increased amygdala sensitivity with increased top-down prefrontal inhibition.

II. Methods

To simulate a participant carrying out the antisaccade task, we implemented a recurrent neural network embedded in a virtual environment, as shown in Figure 1. The network is composed of interacting components, and sensorimotor exchanges are implemented as pixel-level visual input and topographic eye movement motor commands. The task environment is modeled after [9]. The subject's goal is to generate eye movements in the opposite direction of facial stimuli presented between periods of rest. The time between stimulus presentation and the subsequent eye movement is measured as the antisaccade latency.

The task is carried out in multiple stages. First, visual processing in the ventral visual stream is monitored by the anterior cingulate cortex (ACC). Upon detection of a face in the temporal cortex (TC), the ACC emits a signal to the basal ganglia (BG) to proceed with the task. Next, based on the detected eye gaze direction, the lateral prefrontal cortex (LPFC) prepares an antisaccade in the frontal eye fields (FEF), which is emitted to the superior colliculus (SC), engaging the appropriate motor response. This flow of information is determined by the LPFC, which maintains goal-directed behavior via top-down control of gating operations in the basal ganglia (BG). Programmatic behavior is encoded in these regions as distributed activation states with gated transitions, and is learned using temporally asymmetric one-step learning [11], [12]. Together, the LPFC and BG cooperate to carry out the antisaccade task, mediating oculomotor behavior according to perceived environmental cues.

Visual input from the environment is processed by a cascade of layers from the retina to the TC, which is implemented as a convolutional neural network trained to recognize task-relevant visual features. TC activation is monitored by the ACC, which signals to the BG to proceed with the task upon face detection. The LPFC then prepares a contralateral saccade in the FEF based on the detected gaze direction, and the BG initiates the saccade via the SC.

Our central question concerns the contribution of amygdala hyperactivity and impaired top-down prefrontal inhibition to increased antisaccade latency. To address this question, the model includes reciprocally inhibitory connectivity between amygdala (AMY) ventromedial prefrontal cortex (vmPFC) regions [3]. AMY activation is driven by face detection in the TC, initiating a fear response to salient environmental stimuli that interferes with task execution [4], [10]. Strong AMY activation may overwhelm the vmPFC, diminishing top-down regulatory inhibition [2]. In this case, support is needed from prefrontal structures that are implicated in cognitive control, such as the LPFC [5].

The AMY interacts extensively with the basal ganglia, which is thought to contribute to its ability to rapidly interrupt and bias behavior under appropriate conditions [13]. In our model, AMY activation causes a shift in gating dynamics in the basal ganglia that activates the medial dorsal thalamus (MD), which is known to integrate cortical activity by synchronizing oscillatory rhythms [14]. MD activation opens a gate from the LPFC to the vmPFC, supporting the latter in providing time delayed top-down inhibition of the AMY. This ultimately releases the basal ganglia from interruption, allowing task execution to continue. The duration of the interruption, which produces measurable increases in antisaccade latency, depends on the dynamics of reciprocal inhibitory connectivity between the AMY and vmPFC.

This neurocircuitry allows us to explore the putative impairment of top-down inhibition of the AMY hyperactivity by the vmPFC and its impact on measurable response latencies. Experiments were conducted on instantiations of the model that served as simulated experimental subjects carrying out the antisaccade task (referred to as "subjects"). Response latency was measured as the number of model timesteps that elapsed between each face presentation and execution of

the corresponding saccade ($latency$), detected by monitoring the visual fixation point. BOLD signals were simulated for the AMY ($bold^{amy}$) and vmPFC ($bold^{vmpfc}$) by averaging net synaptic activity over the entire test [15]. Weight parameters in the emotional subnetwork of the model served as explanatory variables: w^{at} (TC to AMY), w^{vt} (TC to vmPFC), w^{av} (vmPFC to AMY), w^{va} (AMY to vmPFC) and w^{vl} (LPFC to vmPFC). We fix w^{vt} and w^{va} at 0.5 and 1.0, respectively, and vary w^{at}, w^{av}, and w^{vl} experimentally. These weights correspond to AMY sensitivity (w^{at}), strength of frontolimbic inhibition (w^{av}), and strength of prefrontal functional connectivity implicated in cognitive control of emotional regulation (w^{vl}) [5].

III. Results

To establish a baseline model, a healthy control group was simulated using weight parameters ($w^{at} = 0.5$, $w^{vt} = 0.5$, $w^{av} = 1.0$, $w^{va} = 1.0$, $w^{vl} = 1.5$). Note that the values of these parameters are expressed as positive values indicating connectivity strength, regardless of whether the connection is excitatory or inhibitory. A total of 20 subjects was generated using these parameters, and all 20 successfully produced antisaccades for 20 face presentations with a mean latency of 172.4 model timesteps. Because the mean latency of the control group in [9] was 440ms, one model timestep corresponds to 2.552ms of real time. Mean simulated AMY BOLD and vmPFC BOLD averaged over the test period were 247 and 884 units per second, respectively. Because fMRI experiments typically express results as the percentage of change in BOLD response, we left these in arbitrary units.

Our goal was to identify connectivity changes that explain the observed increases in antisaccade latency and changes in BOLD reported in neuroimaging studies. The mean latency of the experimental group in [9] was 633ms, and the results of [6] suggest that such a subject should demonstrate an increase in AMY BOLD, and a decrease in vmPFC BOLD. We assumed small but significant changes in BOLD signals of 5% to 259 (AMY) and 839 (vmPFC) units per second, respectively. We hypothesized that the expected changes in these response variables would be produced by an increase in AMY sensitivity (w^{at}), a decrease in top-down inhibitory strength (w^{av}), and a decrease in prefrontal functional connectivity (w^{vl}). In addition, we hypothesized that recovering prefrontal functional connectivity (w^{vl}) to baseline levels would recover some of the antisaccade performance deficit observed in PTSD, decreasing mean antisaccade latency.

To avoid confounds from the choice of initial parameters, we first conducted a search of the parameter space, varying w^{at}, w^{av}, and w^{vl} on an evenly spaced lattice in the ranges ([0.1, 1.0], [0.1, 2.0], [0.1, 2.0]) with an increment of 0.1 in each dimension. Each data point corresponded to a simulated subject (4000 total) that was tested on 5 face presentations. We filtered out subjects that failed to produce one or more antisaccades leaving 2134 subjects that successfully completed the task with non-zero task interference. These remaining subjects were used to generate a dataset of explanatory and dependent variables, shown in Figure 2. Although the BOLD response variables are in arbitrary units, these plots illustrate the impacts of changes in connectivity parameters: increased AMY sensitivity (w^{at}) tends to increase antisaccade latency and AMY BOLD, while increasing the strength of the LPFC-vmPFC-AMY pathway (w^{vl}, w^{av}) tends to decrease antisaccade latency.

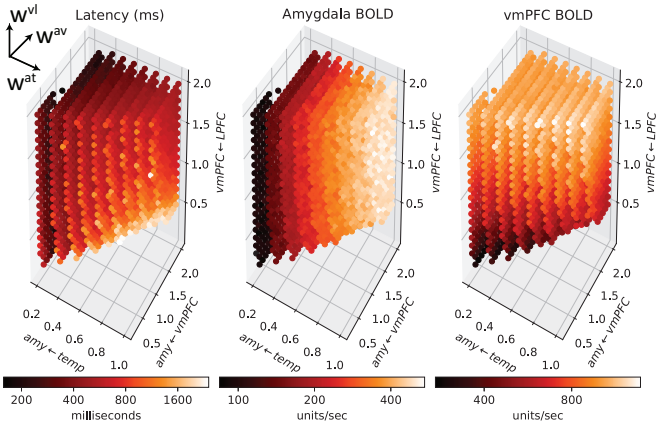

Fig. 2. Mean antisaccade latencies and BOLD responses for 2134 simulated subjects that successfully completed the task and demonstrated non-zero task interruption. Each axis corresponds to a variable connection weight (w^{at}, w^{av}, w^{vl}), and the color indicates the value of each response variable ($latency$, $bold^{amy}$, $bold^{vmpfc}$) in log-scale.

To quantify these relationships and eliminate arbitrary units, we generated a dataset of log-transformed fractional changes between each unique pair of subjects and performed ordinary least squares regression with fixed intercepts, yielding the coefficients shown in Table I. This model produces a reasonable fit to the data, and the signs of the coefficients corroborate the relationships illustrated in Figure 2. Results from [9] provided a target increase in antisaccade latency of 43.9% (633ms/440ms). Using the computed coefficients, we solved the following system of equations to identify the corresponding parameter changes:

$$\begin{bmatrix} 1.198 & -0.694 & -0.299 \\ 0.980 & 0.018 & 0.004 \\ 0.182 & -0.136 & 0.459 \end{bmatrix} \begin{bmatrix} w^{at'} \\ w^{av'} \\ w^{vl'} \end{bmatrix} = \begin{bmatrix} log(1.439) \\ log(1.05) \\ log(0.95) \end{bmatrix}$$

$$w^{at'} = 1.058, \quad w^{av'} = 0.721, \quad w^{vl'} = 0.794$$

This pattern of changes matches our hypothesis of increased AMY sensitivity, decreased top-down inhibition from the vmPFC, and decreased excitation from the LPFC to the vmPFC. We applied this pattern to our simulated control subjects ($hlth$) to produce parameters for simulated patients with PTSD ($ptsd$). To simulate cognitive reappraisal training, we generated a third set of subjects ($recov$) with recovered (healthy baseline) prefrontal connectivity:

$$\begin{bmatrix} w^{at}_{hlth} & w^{at}_{ptsd} & w^{at}_{recov} \\ w^{av}_{hlth} & w^{av}_{ptsd} & w^{av}_{recov} \\ w^{vl}_{hlth} & w^{vl}_{ptsd} & w^{vl}_{recov} \end{bmatrix} = \begin{bmatrix} 0.5 & 0.529 & 0.529 \\ 1.0 & 0.721 & 0.721 \\ 1.5 & 1.191 & 1.5 \end{bmatrix}$$

20 subjects were generated per group and and evaluated on 20 face presentations. Means and standard deviations are

reported in Table II, which includes net activation of the relevant regions. Statistical significance (Welch's t-test) is reported in Table III. The results show a reasonably close match to our predicted latency and BOLD responses for the PTSD group. Although the vmPFC BOLD response falls quantitatively short of the predicted 859 units/sec, the BOLD results match our qualitative predictions: AMY BOLD increased, while vmPFC BOLD decreased.

In addition, the results yield testable predictions about the impact of cognitive reappraisal training on antisaccade performance and BOLD measurements. If such training recovers baseline prefrontal connectivity, the model predicts that these patients would demonstrate a mean antisaccade latency of 587ms, no noticeable changes in AMY BOLD, and a significant increase in vmPFC BOLD. This somewhat paradoxical prediction can be explained by discrepancies between BOLD responses and neural activation: BOLD signals integrate excitatory and inhibitory synaptic activity. Despite increased AMY sensitivity, AMY activation in the model decreases by 11.7% after cognitive reappraisal training, while vmPFC activation increases by 5.9%. The increases in simulated AMY BOLD must be partially compensated by a net reduction in recurrent activation due to enhanced top-down inhibition from the vmPFC.

	w^{at}	w^{av}	w^{vl}	R^2
Latency (ms)	1.198	-0.694	-0.299	0.870
Amy BOLD	0.980	0.018	0.004	0.991
vmPFC BOLD	0.182	-0.136	0.459	0.947

TABLE I

LOG-TRANSFORMED LINEAR REGRESSION COEFFICIENTS

	Healthy	PTSD	Recovered
Latency (ms)	440 (9.4)	654 (19.9)	587 (16.8)
Amy BOLD	247 (9.65)	261 (7.16)	261 (9.90)
vmPFC BOLD	884 (13.8)	859 (11.7)	982 (17.8)
Amy Activity	9.81 (0.85)	21.51 (1.08)	19.0 (1.22)
vmPFC Activity	106.95 (4.87)	134.17 (3.08)	142.05 (4.81)

TABLE II

MEAN AND STD DEV OF MODEL PERFORMANCE ACROSS CONDITIONS

	ptsd/hlth	p	recov/ptsd	p
Latency (ms)	1.486	< 0.001	0.898	< 0.001
Amy BOLD	1.057	< 0.001	1.000	0.852
vmPFC BOLD	0.971	< 0.001	1.143	< 0.001
Amy Activ	2.193	< 0.001	0.883	< 0.001
vmPFC Activ	1.255	< 0.001	1.059	< 0.001

TABLE III

SIGNIFICANCE OF CHANGES ACROSS CONDITIONS

IV. DISCUSSION

A recent review identified the need to bridge the gap between neuroanatomical, functional, and behavioral changes in PTSD [7]. Our study begins to address this challenge by predicting changes in amygdala and prefrontal connectivity that may be tested empirically with methods such as optogenetic fMRI. Our model produces changes in simulated oculomotor behavior and BOLD responses that are consistent with empirical studies [6], [9]. Our experimental results support the hypothesis that elevated amygdala activation and diminished top-down prefrontal regulation significantly contribute to impaired behavioral control in patients with PTSD, consistent with the fear conditioning model [3].

Cognitive reappraisal training has been shown to increase functional connectivity within the prefrontal cortex and downregulate negative affect [5]. We used the model to explore the potential impacts of increased prefrontal connectivity and found a partial recovery of normal performance, as well as a significant increase in vmPFC BOLD. Although cognitive reappraisal training did not lead to changes in amygdala BOLD response, net activation of the amygdala decreased by about 11.7%, indicating that increased prefrontal connectivity can dampen amygdala response without producing noticeable changes in the BOLD signal. This result highlights potential methodological challenges involved with interpretation of neuroimaging data, and highlights the value of computational modeling in integrating neuroanatomical, functional neuroimaging, and behavioral evidence to understand the etiology of neuropsychiatric illnesses.

REFERENCES

[1] R. K. Pitman, *et al.*, "Biological studies of post-traumatic stress disorder," *Nature Reviews Neuroscience*, vol. 13, no. 11, p. 769, 2012.

[2] S. D. Norrholm and T. Jovanovic, "Fear processing, psychophysiology, and ptsd," *Harvard Review of Psychiatry*, vol. 26, no. 3, pp. 129–141, 2018.

[3] D. A. Ross, *et al.*, "An integrated neuroscience perspective on formulation and treatment planning for posttraumatic stress disorder," *JAMA Psychiatry*, vol. 74, no. 4, p. 407, Jan 2017.

[4] J. C. Erlich, *et al.*, "The role of the lateral amygdala in the retrieval and maintenance of fear-memories formed by repeated probabilistic reinforcement," *Frontiers in Behavioral Neuroscience*, vol. 6, 2012.

[5] K. N. Ochsner, *et al.*, "Functional imaging studies of emotion regulation," *Annals of the New York Academy of Sciences*, vol. 1251, no. 1, pp. E1–E24, 2012.

[6] A. Etkin and T. D. Wager, "Functional neuroimaging of anxiety," *American Journal of Psychiatry*, vol. 164, no. 10, pp. 1476–1488, 2007.

[7] R. J. Fenster, *et al.*, "Brain circuit dysfunction in post-traumatic stress disorder: from mouse to man," *Nature Reviews Neuroscience*, p. 1, 2018.

[8] S. B. Hutton and U. Ettinger, "The antisaccade task as a research tool in psychopathology: a critical review," *Psychophysiology*, vol. 43, no. 3, pp. 302–313, 2006.

[9] M. Reinhard, *et al.*, "Neuropsychological measurement of inhibitory control in posttraumatic stress disorder," *Journal of Clinical and Experimental Neuropsychology*, pp. 1–11, 2017.

[10] M. Gamer and C. Buchel, "Amygdala activation predicts gaze toward fearful eyes," *Journal of Neuroscience*, vol. 29, no. 28, pp. 9123–9126, 2009.

[11] G. E. Katz, *et al.*, "A programmable neural virtual machine based on a fast store-erase learning rule," *Neural Networks*, vol. 119, pp. 10–30, 2019.

[12] J. Sylvester and J. Reggia, "Engineering neural systems for high-level problem solving," *Neural Networks*, vol. 79, pp. 37–52, 2016.

[13] Y. T. Cho, *et al.*, "Cortico–amygdala–striatal circuits are organized as hierarchical subsystems through the primate amygdala," *Journal of Neuroscience*, vol. 33, no. 35, pp. 14 017–14 030, 2013.

[14] A. S. Mitchell and S. Chakraborty, "What does the mediodorsal thalamus do?" *Frontiers in Systems Neuroscience*, vol. 7, Aug 2013.

[15] R. Winder, *et al.*, "Functional connectivity in fmri," *NeuroImage*, vol. 34, no. 3, pp. 1093–1107, 2007.