Modeling the Effects of Post-Traumatic Stress on Hippocampal Volume

Briana M. Smith (brianam2@uw.edu)
Department of Psychology, University of Washington
Seattle, WA 98195 USA

Madison Chiu (mchiu44@uw.edu)
Department of Psychology, University of Washington
Seattle, WA 98195 USA

Yuxue Cher Yang (chery@uw.edu)
Department of Psychology, University of Washington
Seattle, WA 98195 USA

Catherine Sibert (sibert@uw.edu)
Department of Psychology, University of Washington
Seattle, WA 98195 USA

Andrea Stocco (stocco@uw.edu)
Department of Psychology, University of Washington
Seattle, WA 98195 USA

Abstract

Post-Traumatic Stress Disorder (PTSD) is a psychiatric disorder often characterized by the unwanted re-experiencing of a traumatic event through nightmares, flashbacks, and/or intrusive memories. This paper presents a neurocomputational model using the ACT-R cognitive architecture that simulates intrusive memory retrieval following a potentially traumatic event (PTE) and derives predictions about hippocampus volume observed in PTSD. Memory intrusions were captured in the ACT-R Bayesian framework by weighting the posterior probability with an emotional intensity term *I* to capture the degree to which an event was perceived as dangerous or traumatic. It is hypothesized that (1) Increasing the intensity *I* of a PTE will increase the odds of memory intrusions; and (2) Increased intrusions will result in a concurrent decrease in hippocampal size. A series of simulations were run and it was found that *I* had a significant effect on the probability of experiencing traumatic memory intrusions following a PTE. The model also found that *I* was a significant predictor of hippocampal volume reduction, where the mean and range of simulated volume loss match results of existing meta-analysis. The authors believe that this is the first model to both describe traumatic memory retrieval and provide a mechanistic account of changes in hippocampal volume, capturing one plausible link between PTSD and hippocampus size.

Keywords: Post-Traumatic Stress Disorder; Hippocampus; Amygdala; Declarative Memory; Long-Term Memory; ACT-R; Cognitive Architecture

Introduction

Post-Traumatic Stress Disorder (PTSD) is a psychiatric disorder that originates after experiencing or witnessing a traumatic event, such as rape, domestic violence, assault, a serious accident, or military combat. At the behavioral level, PTSD is characterized by persistent avoidance, alterations in mood, as well as cognitive distortions surrounding the trauma. One of the most characteristic and disruptive behavioral effects of PTSD, however, is the unwanted reexperiencing of the trauma through nightmares, flashbacks, and/or intrusive memories. Traumatic experiences evoke an emotional response that is accompanied by increased activation of subcortical areas such as the amygdala. Intrusive memories are thought to occur because of the simultaneous activation of the amygdala and hippocampus during memory encoding (Marks, Franklin, & Zoellner, 2018). At the subcortical level, PTSD is also characterized by a marked reduction in the volume of the hippocampus—a medial temporal lobe structure necessary for memory functioning. It is important to note that this change is primarily structural, and, although often remarkably apparent, decreased hippocampus size is not accompanied by a functional impairment in long-term memory performance (Karl et al., 2006).

The goal of this paper is to derive predictions about the changes in hippocampus volume observed in PTSD by using a neurocomputational model to simulate intrusive memories over time within an integrated cognitive architecture. The central idea of the model is that intrusive memories operate *within* the context of a general theory of declarative memory, specifically episodic memory. Within this framework, the persistent memory intrusions observed in PTSD can be seen not as a maladaptive response, but rather as the runaway process of an otherwise adaptive memory system.

As a memory is retrieved more frequently, its priority increases and its rate of decay decreases. A traumatic memory, however, tends to out-compete more contextually

appropriate memories due to the fact it was encoded in a highly emotional state. With each retrieval of the traumatic memory, disproportionately more resources are allocated to it, leading to the further preservation and growth of these unwanted memory intrusions. In this framework, it is proposed that the corresponding changes in hippocampal volume associated with PTSD can be explained as the natural result of a biological process to efficiently allocate resources to changing memory demands.

The model presented herein is framed within ACT-R's theory of declarative memory (Anderson, 2007). This choice was motivated by three reasons. First, ACT-R is the most commonly adopted cognitive architecture in psychology and the cognitive neurosciences (Kotseruba & Tsotsos, 2018). Second, ACT-R has a long and established history of application to brain sciences, making the process of drawing new inferences at the neural level easier and less tentative. Finally, ACT-R is based on a Bayesian framework, which provides an elegant foundation of declarative memory retrieval processes and can be easily extended to incorporate the proposed theory of memory retrieval according to their emotional intensity.

The Model

Before introducing the model from a neural and an algorithmic point of view, it is important to frame it within Anderson's analysis of human episodic memory in terms of "Rational Analysis" (Anderson, 1990), or, as it is called currently, Bayesian terms. Throughout this paper, this analysis will be referred to as a guiding principle to modify ACT-R and make inferences about its neural substrates.

In the Bayesian framework, a memory m's probability of being recalled in the presence of a context $Q = \{q_1, q_2, \ldots q_n\}$ reflects the memory's retrieval *need*, and is a Bayesian function of both the past history of m and the degree to which each contextual cue q predicts m. In Anderson's (1990) formulation, the retrieval need of a memory m in a context Q is expressed in terms of a memory's activation $A(m)$, a quantity that reflects its log posterior odds of being retrieved in the presence of Q. Following Bayes rule, the posterior odds can be separated into two different quantities, the prior odds and the likelihood odds:

$$\begin{aligned} A(m) &= \log [P(m|Q) / P(m/\neg Q)] \\ &= \log [P(m) / P(\neg m)] + \log [P(Q|m) / P(Q|\neg m)] \\ &= \log [P(m) / P(\neg m)] + \log \textstyle\prod_q [P(q|m) / P(q)] \\ &= \log [P(m) / P(\neg m)] + \textstyle\sum_q \log [P(q|m) / P(q)] \end{aligned} \quad (1)$$

In ACT-R, it is customary to give different names to the two quantities that make up the right-hand side of Eq. 1, referring to the prior odds as the base-level activation or $B(m)$, and to the likelihood odds as the spreading activation or $S(m)$.

A memory's base-level activation $B(m)$ increases with the frequency of its usage and decreases over time, reflecting the effects of frequency and recency. In ACT-R, each use of a memory m leaves a trace i, and each trace i decays exponentially over time with a decay rate d, which represents an individual-specific rate of forgetting (Sense et al. 2016). A single memory m is associated with multiple traces, each of which corresponds to a time during which m has been encoded, and re-encoded, or retrieved. Thus, the log odds of retrieving m correspond to the sum of the log odds of retrieving each of its individual decaying traces, and $B(m)$ can be expressed as:

$$B(m) = \textstyle\sum_i (t - t_i)^{-d} \quad (2)$$

Spreading activation $S(m)$, instead, can be interpreted in reference to semantic networks, in which memories are connected by associative links, and activation flows through the links to associated nodes in the network. In this case, the activated nodes represent the elements q in the context Q, and the links represent the degree of association or similarity between q and each memory's features. By means of spreading activation, the proper context can facilitate the retrieval of memories whose base-level activation would, otherwise, be too weak. The amount of spreading activation is proportional to the product between the strength of the link connecting q to m (indicated as $s_{q \to m}$) and an *attentional weight*. The weight is usually simplified as a single scalar quantity, W, divided over the number of active elements in the context, N:

$$S(m) = \textstyle\sum_q (W/N)\, s_{q \to m} \quad (3)$$

Different values of W in Eq. 3 alter the degree to which memory retrieval depends on spreading (and, therefore, contextual cues) vs. base-level activation (and, therefore, statistical priors).

ACT-R In the Context of Memory Consolidation

Although ACT-R has been described in many ways, it is useful, given the goal of this paper, to compare it to a prominent neural theory of memory consolidation, the Multiple Trace Theory (MTT: Moscovitch et al., 2005). The MTT assumes that episodic memories originate from distributed representations that span multiple cortical areas (Figure 1). During the encoding phase (red lines in Fig. 1), the different features of an event ($q_1 \ldots q_N$) are encoded by different cortical areas and bound together into a single association map in the hippocampus (as attributes $a_1 \ldots a_N$), through the multiple descending pathways that converge from the cortex through the dentate gyrus. MTT posits that the hippocampus is the permanent store of episodic memories, and that each encoding episode leaves a permanent trace. During retrieval, the hippocampus trace is temporarily re-activated (blue lines in Figure 1) and, through ascending pathways from the temporal lobe to the cortex, causes the re-activation of the original neurons. This reactivation, in turn, might be re-encoded as a second trace.

Base-level activation and spreading activation reflect, therefore, two distinct neural processes. Specifically,

base-level activation reflects processes that are internal to the hippocampal network, such as decay or interference due to accumulation of memory traces (Alvarez & Squire, 1994), while spreading activation reflects the mechanism by which cortical inputs might trigger contextual memory retrieval (Rolls & Treves, 1998).

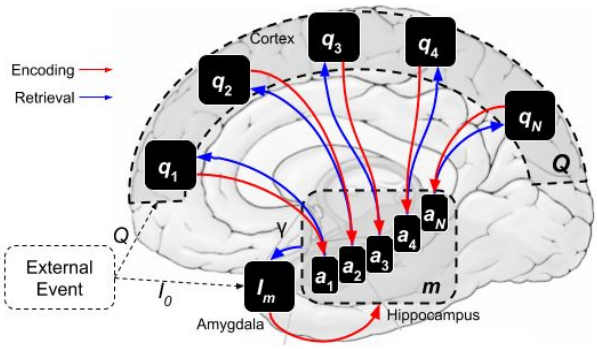

Figure 1: A neuroanatomical interpretation of the model presented herein.

Extending ACT-R to Include Trauma

It has been noted several times, even by Anderson himself (Anderson, 2007, Chapter 3), that one limitation of this approach is that it considers all memories as equally important. On the contrary, not all memories are. Memories of emotional events are thought to persist longer and be more readily available for retrieval than non-emotional memories because of the activation of the amygdala during memory encoding (Marks, Franklin, & Zoellner, 2018). Specifically, memories of events that incorporate threat or fear are of greater importance evolutionarily because they are often critical for survival (Ledoux, 1998). Although some authors have generalized this approach to all emotions (and advanced strong arguments), this paper will limit itself to the responses of the amygdala which are directly connected to PTSD and well understood in neurophysiological terms (Bryant, et al. 2008).

In a Bayesian framework, the concept of survival importance can be easily captured by weighting the posterior probability by an emotional impact term, referred to as *intensity*, $0 < I(m) < \infty$, which captures the degree to which an event was potentially dangerous or traumatic. The posterior odds now become:

$$\begin{aligned} A(m) &= \log [P(m|Q) / P(m/\neg Q)] \, [I(m)/\mathrm{I}(\neg m)] \\ &= B(m) + S(m) + \log I(m) - \log I(\neg m) \\ &= B(m) + S(m) + \log I(m) - k \end{aligned} \qquad (4)$$

The last passage was motivated by the consideration that, over a lifetime, $I(\neg m)$ would approach the mean traumatic value of all memories and thus could be considered a background constant k.

In summary, the proposed Bayesian framework suggests that traumatic events add a *constant bias* that makes a memory more likely to be retrieved, even in the absence of contextual cues and in proportion to the perceived intensity of the traumatic event. In biological terms, this perceived intensity bias can be interpreted as the contribution of the amygdala to hippocampal activation (Fig. 1). The amygdala is bidirectionally connected to the hippocampus and is known to play a key role in processing event salience (Anderson & Phelps, 2001), fear (LeDoux, 1998) and in boosting memory for stressful events (references). Importantly, and consistently with our interpretation, the amygdala is hyper-responsive in individuals suffering from PTSD (Shin, Rauch, & Pitman, 2006).

Deriving ACT-R Predictions For Hippocampus Size

The final step to test this theory consists of deriving predictions about hippocampus size from the augmented ACT-R framework. To calculate hippocampus size, the following analysis was adhered to. In general, it is known that the size of the hippocampus changes with experience. For instance, in a landmark study (Maguire, Woolett, & Spiers, 2006), cab drivers of London were shown to have larger hippocampus volume than the general population. Additionally, another study showed the volume of the hippocampus co-varies with the years of education (Noble et al., 2012). An accepted explanation for this effect is that the size of the hippocampus reflects the biological investment in storing memories that need to be re-used often (Wollet & Maguire, 2011).

An *efficient* memory storing system would encode cells so memories that need to be accessed more frequently use less resources (in neural terms, less cells or synapses) than memories that need to be accessed less often (Huffman, 1952). In the Bayesian terms described above, memories that are accessed more often have the highest *priors* and, in ACT-R terms, the higher *base-level activations*. Knowing the priors of memory utilization, the size of the hippocampus could then be approximated by a measure of the homogeneity of the distribution of the priors. Here, the long-term memory's *information entropy*, H, was utilized, i.e., the quantity (Shannon, 1948):

$$H = -\sum_m P(m) \log P(m) \qquad (5)$$

This quantity captures how much information is represented in declarative memory, once the different probabilities of each memory are taken into account. Consider, for example, the case of two London cab drivers who have memorized the same number of addresses but use them with different probabilities. For one driver, all addresses are equally likely to be retrieved, reflecting the fact that his clients are equally likely to request a ride to all of these locations. For the second driver, on the other hand, one single address is requested all the time, while all the others are seldom, if ever, requested by clients. Information entropy is high for the first driver because it is impossible to

predict which address will be requested by the next client. For the second driver, on the other hand, entropy is low, since one memory is highly predictable and all the other can be ignored. Biologically, the first driver needs to allocate more resources (hippocampal cells) to maintain all of these memories than the second, for whom a small number of cells could be used to encode the single memory that predicts most of the clients' rides in their daily routine.

In ACT-R, a memory's probability of being used, $P(m)$, is reflected in its base-level activation (Eq. 1) and, in this paper's specific model, in its intensity $I(m)$. Note the base-level activation $B(m)$ reflects the memory's prior *odds* rather than true probabilities. To translate them into probabilities, base-level activations were normalized across all memories into long-term memory (LTM):

$$P(m) = [B(m) + I(m)] / \sum_{i \in LTM} [B(i) + I(i)] \quad (6)$$

Hypotheses and Predictions

Given the theory outlined above, it is hypothesize that (1) Increasing the emotional intensity I of a potentially traumatic event (PTE) will increase the odds of the event memory being retrieved out of context, predicting intrusive memory occurrence observed clinically in patients with PTSD; and (2) Increased intrusion occurrence will result in a concurrent decrease in hippocampal size, driven by the altered landscape of memory recall priors, and thereby capturing the relationship between trauma and hippocampus size.

Methods

To test the hypothesis driving this experiment, a series of computational simulations were run. The following sections describe the details of the simulations

Memory Representations

The simulations described herein differ significantly from most ACT-R models because they focus on modeling episodic memories over extended durations (~6 months) rather than on specific tasks for very short times. Thus, they adopt a uniform memory representation for all memories instead of different, task-dependent structures. Specifically, all memories are vectors of $N = 8$ features. Each feature is given a randomly selected value, called an attribute, from a pool whose size is determined by a given parameter, A (not relevant for this study and thus not discussed). The attributes for all "normal" events are always selected from the same pool, which captures the common features found in one's daily environment. Attributes of PTEs are selected from a different pool of attributes, representing the unique extraordinary features associated with traumatic circumstances.

Model Behavior

The model performs routine behaviors following a perceive-retrieve-respond loop. The loop initiates when a new event occurs in the external world. The event is perceived by the model, and its features are held in sensory buffers that, together, form the current context Q (Fig. 1). The model responds to the current context by first setting a goal to resolve it. When the goal is set, the model retrieves the memory with the highest total amount of activation, $A(m)$. The retrieval process is influenced by three factors: (1) the base-level activation of the model's memories of previous events, $B(m)$, which determines the memories' *priors*; (2) The spreading activation from the current context $S(m)$, modulated by the model's executive attention W; and (3) The intensity $I(m)$ of previous events. This loop captures a simple decision-by-sampling strategy (Stewart, Chater, & Brown, 2012): Facing a new situation, the model responds by retrieving the most contextually appropriate situation faced in the past, balancing recency, frequency, and contextual cues through spreading activation. Once the memory is retrieved, the goal is resolved and a new memory is formed to encode the current event using the contextual cues $q_1, q_2 \ldots q_N$ as its attributes (Fig. 1).

Daily Event Distribution and Simulation Time Window

To model the accumulation of memories in a plausible manner, new events are presented to the model at a frequency that follows a gamma distribution and a realistic daily schedule. On average, the model is presented with approximately ~20 events per day. Events occur between 8:00 AM and midnight, with a peak probability at around noon. This event distribution was chosen to reflect the normal waking hours of a person, with a greater concentration of events during working hours (8:00AM-4:00PM). Each event's emotional intensity I was randomly selected from a uniform distribution between 0 and 2, so that their mean was equal to 1 (and thus the bias term k in Eq. 4 was equal to zero).

Each simulated run of the model lasted 160 consecutive days, starting 100 days before the occurrence of a traumatic event and extended 60 days after that. On the midnight of day zero, a PTE was generated and presented to the model. The intensity of the PTE was explicitly manipulated throughout the simulations, and given the values of $I_{PTE} = 1$ (control condition), 20, 40, 60. The model's time window extended to another 60 days after the PTE.

Dependent Variables

Two dependent variables are the focus of this study. The first is the probability of experiencing an intrusive memory during the day. This is defined as the probability that the model retrieves a memory of the PTE in response to a situation throughout the day. Note that, because the PTE's attributes are different from those of the daily events, its retrieval is always contextually inappropriate, and thus its recall qualifies as *intrusive*.

The other variable is the hippocampal volume reduction, which is measured as a percent change from a control condition. To get a suitable baseline, the average value of H (as a proxy for hippocampus size) over the last 10 days of the simulation (corresponding to days 50-60 after the PTE) was compared to the average value of H for the same period of a model run with an identical combination of parameters except $I = 1$.

Simulations

In addition to the intensity I of the traumatic event, a number of other parameters were manipulated parametrically. These parameters were derived from a recent review of the PTSD literature (Marks et al., 2018) and reflect idiographic factors that moderate the behavioral outcomes of traumatic stress. They include the vividness of memory re-experience γ; the vividness of sensory encoding, modeled as the size of attributes pool A; individual differences in working memory capacity W (see Eq. 3); the tendency to ruminate over the traumatic event R; and the potential overlap C between features of the traumatic event and attributes of daily situations. Although these parameters will not be discussed in this paper, they are summarized in Table 1 and were left in the analysis as they contribute to representative variability in the simulated results. To obtain stable estimates, the model was run 50 times for each of the 576 combinations of parameter values. estimates. In total, the simulations spanned 4,608,000 simulated days, and 103,330,000 simulated events.

Table 1: Model parameters manipulated in the simulations

Parameter	Meaning	Values
I	Intensity of PTE	1, 20, 40, 60
A	Size of attributes pool	6, 8
γ	Vividness of re-experience	0.80, 0.90, 0.95
W	Working memory	4, 8, 12
C	Similarity between PTE and daily events	0, 0.25, 0.5, 0.75
R	Number of rumination events in a day	0, 20

Results

Given the large number of simulations that were run, it is impossible to fully report the complete set of results. For the purpose of this paper, there are two aspects to concentrate on. First, as expected the model does indeed show worse clinical outcomes in response to more traumatic events. Figure 2, below, shows the daily incidence of traumatic memories. A 3x60 ANOVA, using emotional intensity I and the days after PTE as factors, revealed that I had a significant effect on the probability of experiencing traumatic memories in the days following a traumatic event [$F(2, 1295345) = 37{,}115.6$, $p < .0001$], with higher values of I corresponding to higher intrusion probability. Furthermore, I interacted significantly with the day [$F(118, 1295345) = 10.3$, $p < .0001$], resulting in different recovery trajectories (Figure 2).

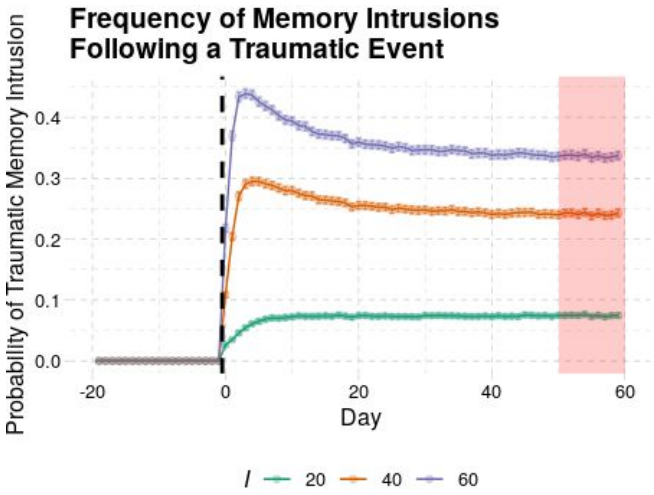

Figure 2: Predicted increase in memory intrusion following a PTE on Day 0 (black dashed line) as a function of emotional intensity I. The shaded red area marks the time interval in which the hippocampus volume was calculated.

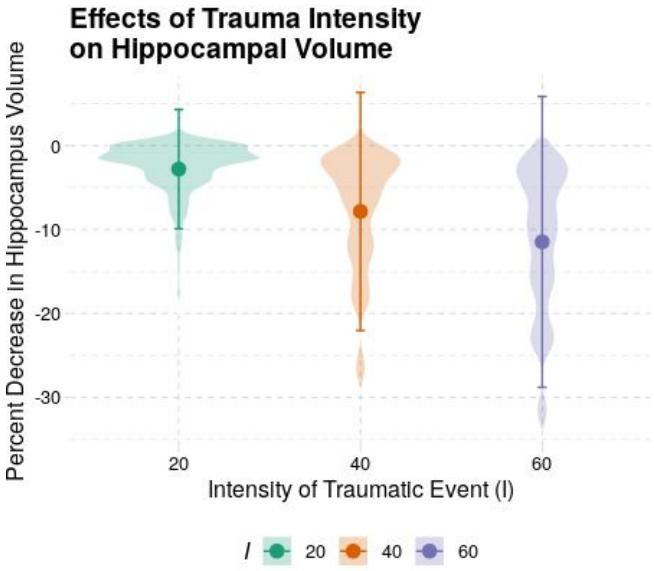

Figure 3: Effect of trauma intensity on hippocampal volume. The violin plots represent the distribution densities of model runs resulting in the corresponding decreases of hippocampal volumes. Solid circles and lines represent means +/- SD.

Having established that the model succeeds in capturing these signatures of PTSD, the results were further examined

to estimate the effects of traumatic stress on hippocampus size. It was observed, across all parameters, that there was general reduction of simulated hippocampus size, ranging from zero to 33.89% with a mean decrease of 7.35% [$t(21,599) = 140.83$, $p < .0001$]. Both mean and range match the results of existing meta-analysis. For example, in Smith's (2005) influential review of structural MRI studies, the range: 0 to 44% and the mean 6.9%. A second question was whether the severity of the reduction was predicted by the severity of trauma. To this end, the model found that the emotional intensity I was a significant predictor of hippocampal volume reduction [$F(2, 21594) = 774.7$, $p < .0001$], with the decrease in hippocampus size growing with greater values of I (all pairwise comparisons significant at $p < .0001$, Bonferroni corrected). This is shown in Figure 3, which shows the distributions of predicted decreases of hippocampal volumes in the simulations, visualized (as violin plots) separately for different values of intensity I.

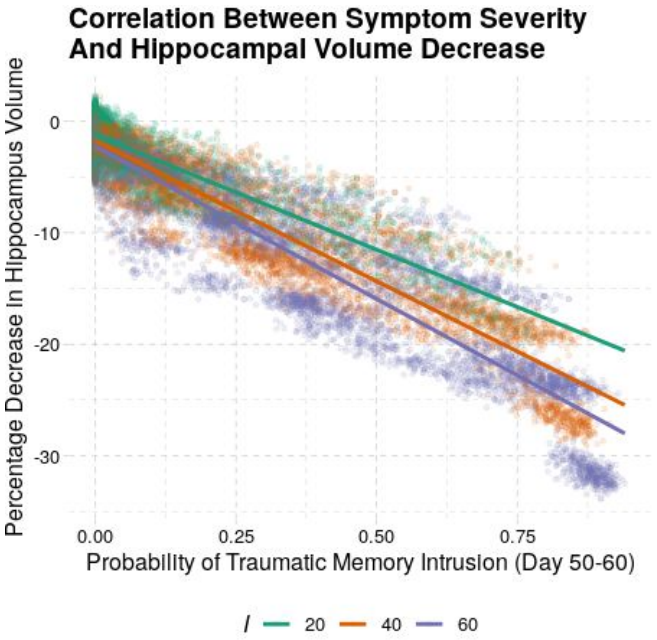

Figure 4: Correlation between the probability of traumatic memory intrusions in hippocampal volume for varying levels of trauma intensity. Each point represents a single run of the model; solid lines represent the mean regression line.

The final analysis investigated was whether or not the degree of hippocampus volume was correlated to the degree of symptom severity. This is important because, although symptom severity is clearly driven by the severity of the traumatic event, it also depends on other factors that were explicitly manipulated in the simulations (see Methods and Table 1). To do so, the mean daily probability of memory intrusions in the last 10 days of the simulations (red shaded area in Figure 2) and the corresponding percentage decrease in hippocampus were calculated for each run of the model. Three separate linear regressions, one for each level of I = 20, 40, and 60, were then computed In all cases, a significant linear regression was found [$I = 20$: $\beta = -20.55$, $t(7198) = -157.9$, $p < .0001$; $I = 40$: $\beta = -25.18$, $t(7198) = -225.6$, $p < .0001$; $I = 60$: $\beta = -27.35$, $t(7198) = 205.1$, $p < .0001$] as shown in Figure 4.

Discussion

This paper has presented a computational model that draws a link between the prevalence of intrusive memories and the changes in hippocampal volume observed in patients with PTSD. To the best of our knowledge, this is the first model to do so, and to provide a connection between Bayesian theories of memories and the underlying neurobiology. The results of this model are also consistent with estimates from the clinical and medical literature. As such, the model may shed light on a number of cognitive factors, such as traumatic memory activation, that contribute to neurophysiological changes associated with PTSD.

In the presentation of this computational model, there are a few obvious limitations. Although an effort was made to account for numerous idiographic factors (see Table 1), it was impossible to account for all various factors that have been deemed clinically important such as age, gender, duration of trauma, recurrence of trauma, comorbidity of other psychiatric disorders, presence and occurrence of other PTSD symptoms, and genetic predisposition. With that in mind, it is feasible that this model can be altered to account for some of these varying factors, as well as other individual differences not aforementioned. Something imperative to take into consideration for future improvements of this model would be, for example, the specific role of the stress hormone cortisol on hippocampal functioning.

These limitations notwithstanding, the model's success in capturing some behavioral and biological factors is encouraging. Theoretically, this model, along with the other research concerning PTSD and its perceived effects on the hippocampus and amygdala, could be used in the future to enhance clinical practice. Targeted, individualized treatments could be developed in which an individual's biological and behavioral measures are used to parametrize a computational model which is, in turn, used to predict long-term recovery trajectories under different medical options.

Acknowledgments

This work was supported by a scholarship from the University of Washington Institute for Neuroengineering (UWIN) to BMS, and partially supported by an award from the Defence Advanced Research Project Agency (DARPA, Grant No. FA8650-18-C-7826). All of the model code, simulation data, and analysis scripts can be found in the Cognition & Cortical Dynamics Laboratory's Github repository at: http://github.com/UWCCDL/PTSD.

References

Alvarez, P., & Squire, L. R. (1994). Memory consolidation and the medial temporal lobe: a simple network model. *PNAS, 91*(15), 7041-7045.

Anderson, A. K., & Phelps, E. A. (2001). Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature*, 411(6835), 305-309.

Anderson, J. R. (1990). *The adaptive character of thought*. Psychology Press.

Anderson, J. R. (2007). How can the human mind occur in the physical universe? *Oxford University Press*.

Astur, St Germain, Tolin, Ford, Russell, & Stevens. (2006). Hippocampus function predicts severity of post-traumatic stress disorder. *Cyberpsychology & Behavior:,* 9(2), 234-240.

Bryant, R., Kemp, A., Felmingham, K., Liddell, B., Olivieri, G., Peduto, A., . . . Williams, L. (2008). Enhanced amygdala and medial prefrontal activation during nonconscious processing of fear in posttraumatic stress disorder: An fMRI study. Human Brain Mapping, 29(5), 517-523.

Huffman, D. A. (1952). A method for the construction of minimum-redundancy codes. *Proceedings of the IRE, 40*(9), 1098-1101.

Karl, A., Schaefer, M., Malta, L. S., Dörfel, D., Rohleder, N., & Werner, A. (2006). A meta-analysis of structural brain abnormalities in PTSD. *Neuroscience & Biobehavioral Reviews, 30*(7), 1004-1031.

Kotseruba, I., & Tsotsos, J. K. (2018). 40 years of cognitive architectures: Core cognitive abilities and practical applications. *Artificial Intelligence Review*, 1-78.

Maguire, E. A., Woollett, K., & Spiers, H. J. (2006). London taxi drivers and bus drivers: a structural MRI and neuropsychological analysis. *Hippocampus, 16*(12), 1091-1101.

Marks, E. H., Franklin, A. R., & Zoellner, L. A. (2018). Can't get it out of my mind: A systematic review of predictors of intrusive memories of distressing events. *Psychological Bulletin, 144*(6), 584.

Moscovitch, M., Rosenbaum, R. S., Gilboa, A., Addis, D. R., Westmacott, R., Grady, C., ... & Nadel, L. (2005). Functional neuroanatomy of remote episodic, semantic and spatial memory: a unified account based on multiple trace theory. *Journal of anatomy, 207*(1), 35-66.

Noble, K. G., Grieve, S. M., Korgaonkar, M. S., Engelhardt, L. E., Griffith, E. Y., Williams, L. M., & Brickman, A. M. (2012). Hippocampal volume varies with educational attainment across the life-span. Frontiers in human neuroscience, 6, 307.

Rolls, E. T., Treves, A., & Rolls, E. T. (1998). *Neural networks and brain function* (Vol. 572). Oxford: Oxford university press.

Sense, F., Behrens, F., Meijer, R. R., & van Rijn, H. (2016). An individual's rate of forgetting is stable over time but differs across materials. *Topics in cognitive science*, *8*(1), 305-321.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell system technical journal*, 27(3), 379-423.

Shin, L. M., Rauch, S. L., & Pitman, R. K. (2006). Amygdala, medial prefrontal cortex, and hippocampal function in PTSD. *Annals of the New York Academy of Sciences*, 1071(1), 67-79.

Smith, M. E. (2005). Bilateral hippocampal volume reduction in adults with post‐traumatic stress disorder: A meta‐analysis of structural MRI studies. *Hippocampus, 15(*6), 798-807.

Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cognitive Psychology*, 53(1), 1-26.

Woollett, K., & Maguire, E. A. (2011). Acquiring "the Knowledge" of London's layout drives structural brain changes. *Current biology, 21*(24), 2109-2114.