Cognitive, Affective, & Behavioral Neuroscience (2023) 23:716–726
https://doi.org/10.3758/s13415-023-01085-5

SPECIAL ISSUE/UNCERTAINTY

A computational model for learning from repeated traumatic experiences under uncertainty

Alfred P. Kaye[1,2] · Manasa G. Rao[3] · Alex C. Kwan[1,4] · Kerry J. Ressler[5,6] · John H. Krystal[1,2]

Accepted: 28 February 2023 / Published online: 10 May 2023

Abstract

Traumatic events can lead to lifelong, inflexible adaptations in threat perception and behavior, which characterize posttraumatic stress disorder (PTSD). This process involves associations between sensory cues and internal states of threat and then generalization of the threat responses to previously neutral cues. However, most formulations neglect adaptations to threat that are not specific to those associations. To incorporate nonassociative responses to threat, we propose a computational theory of PTSD based on adaptation to the frequency of traumatic events by using a reinforcement learning momentum model. Recent threat prediction errors generate momentum that influences subsequent threat perception in novel contexts. This model fits primary data acquired from a mouse model of PTSD, in which unpredictable footshocks in one context accelerate threat learning in a novel context. The theory is consistent with epidemiological data that show that PTSD incidence increases with the number of traumatic events, as well as the disproportionate impact of early life trauma. Because the theory proposes that PTSD relates to the average of recent threat prediction errors rather than the strength of a specific association, it makes novel predictions for the treatment of PTSD.

Keywords Stress sensitization · PTSD · Trauma · Uncertainty · Computational modeling · Bayesian · Reinforcement learning

Introduction

Computational psychiatry seeks to define psychiatric disorders in terms of fundamental algorithms for survival rather than only as pathological states (Bach & Dayan, 2017; Montague et al., 2012; Wang & Krystal, 2014). Quantitative models may allow personalization of mental health care, insight into the nature of the disorder, or predict the trajectory of symptoms (Galatzer-Levy et al., 2014; Galatzer-Levy et al., 2017; Saxe et al., 2017). For example, depression has been conceived as an adaptation to periods of low reward availability (Eldar et al., 2016). Similarly, hallucinations have been conceptualized as resulting from excessive weighting of prior expectations for auditory stimuli in a Bayesian model (Fletcher & Frith, 2009; Powers et al., 2017). One approach to describing a computational function of a neural system is using David Marr's three levels of analysis (Marr, 2010) (Fig. 1A), which seeks to map connections between computational goals, algorithmic procedures to achieve them, and the neurobiological substrate underlying these processes.

Posttraumatic stress disorder has a computational description that organizes theory and neurobiological data across Marr's three levels: associative fear learning (Fig. 1B) (Brown et al., 2018; Homan et al., 2019; Jovanovic et al., 2012; Ross et al., 2018; Yehuda et al., 1995). Threat learning models have been successfully applied to PTSD and underlie current conceptualizations of the disorder and treatment options (Foa, 2011; Levy & Schiller, 2021). PTSD is seen as an extreme outcome of associative fear learning, which in turn is a fundamental mechanism for predicting threats

✉ Alfred P. Kaye
alfred.kaye@yale.edu

1 Yale University Department of Psychiatry, New Haven, CT, USA

2 VA National Center for PTSD Clinical Neuroscience Division, West Haven, CT, USA

3 Mount Sinai Icahn School of Medicine, New York, NY, USA

4 Cornell University Meinig School of Biomedical Engineering, Ithaca, NY, USA

5 McLean Hospital, Division of Depression and Anxiety Disorder, Belmont, MA, USA

6 Harvard Medical School, Department of Psychiatry, Boston, MA, USA

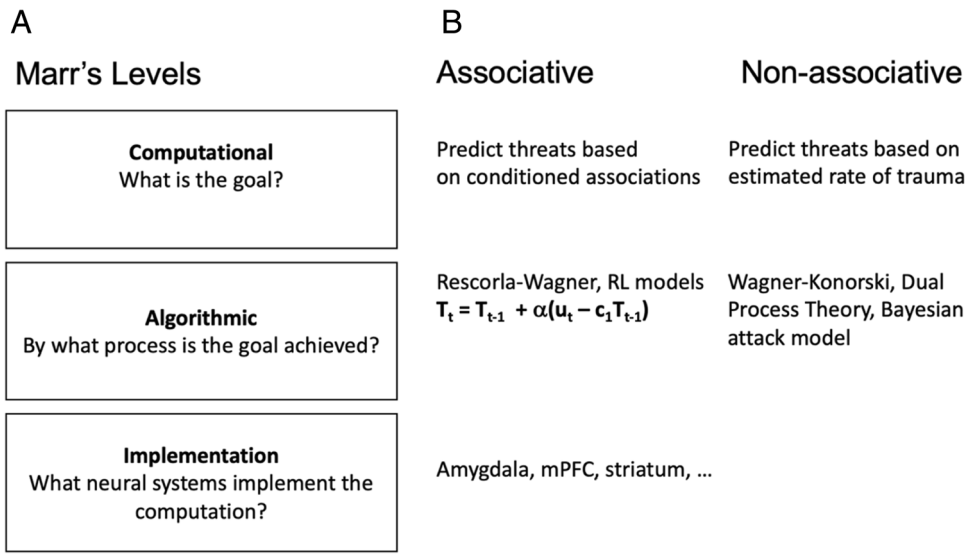

Fig. 1 David Marr's Levels of Analysis for computational neuroscience as applied to PTSD. **(A)** Definition of the three levels of analysis from reference (Eldar et al., 2016). **(B)** Application of those levels to associative learning (left) and non-associative learning (right) in PTSD. (left) Associative learning is a well-characterized system with a clear computational goal of ethological relevance (Computational), a mathematically defined formal model (Algorithm), and neural circuit mechanisms (Implementation). (right) Nonassociative learning is less well-understood but in this study conceptualizes linkages between nonassociative Bayesian estimation of trauma and RL models

based on previous experience (Rescorla & Wagner, 1972). In this model, PTSD occurs when life-threatening situations create potent associations between sensory reminders of the traumatic event and the emotional experience of fear (Levy & Schiller, 2021). The intensity of this association then motivates a person to avoid (Asmundson et al., 2004) future trauma cues, limits extinction of the fear memory (Izquierdo et al., 2006), and supports the subsequent formation of new fear memories via generalization and second-order conditioning (Beck & Sloan, 2012). This process can be described mathematically, enabling learning parameters to be precisely measured during new associative learning in a laboratory setting (Wagner, 1972). The precision with which associative learning can be controlled has enabled neurobiological studies into circuit mechanisms in both humans and animals (Maddox et al., 2019).

In contrast, nonassociative learning—i.e., increases (sensitization) or decreases (habituation) in response to a repeated stimulus (Thompson & Spencer, 1966)—is a prominent component of PTSD that lacks a formal algorithmic description (Fig. 1B). In humans, repeated traumatic events increase the probability of developing PTSD and may change the nature of the disorder (Almli et al., 2013; Khoury et al., 2010). However, the computational and neural processes that mediate this gradual increase in PTSD risk with trauma burden are not well understood. Core PTSD symptoms, such as hyperarousal, inherently involve an exaggerated response to sensory cues. Importantly, these cues need not be associated with the traumatic event to trigger the response (Morgan et al., 1995) but may result from sensitization of neuromodulatory systems ((Kelmendi & Southwick, 2018), but see also (Korem et al., 2022 for habituation responses in PTSD). Neurobiological studies in animals have shown that stress enhances both innate defensive behaviors (Li et al., 2018) and learning about unrelated fear cues (Rau et al., 2005). There are conceptual models of how habituation and sensitization occur (Dual Process Theory (Groves & Thompson, 1970); Wagner-Koniorsky Theory (Wagner, 1979)), which center the role of arousal in changing the response to a stimulus with repetition. However, these models lack the algorithmic detail and clear relation to survival value of Rescorla-Wagner and related reinforcement learning (RL) models (Sutton & Barto, 2018). This has limited the ability to parametrically manipulate and therefore understand nonassociative learning in PTSD patients and animal models.

These nonassociative adaptations are essential, because an organism must estimate the threat of violence to adapt to it. This process of estimation must necessarily involve information gathered across timescales, because threat may increase suddenly or over long periods (Pavluvík et al., 2015). Longer timescale estimation of threat involves integrating experience in disparate environments. To consider a concrete example: predator attacks are stereotyped events that have a significant probability of death (20% for mice exposed to an owl) (Ilany & Eilam, 2008). Life History Theory explains this by positing that stressful experiences in childhood provide information about organismal strategies

that will be adaptive in the adult environment (Frankenhuis & de Weerth, 2013). Childhood traumatic experiences have a strong impact on adult brain structure and function, which influence the development of PTSD (Callaghan & Tottenham, 2016). Previous approaches to computational modeling of PTSD have focused on defining changes in associative learning after traumatic experience (Brown et al., 2018; Homan et al., 2019; Jovanovic et al., 2012; Ross et al., 2018; Yehuda et al., 1995). PTSD is thus framed as a consequence of underlying mechanisms for predicting threat based on previous *associations*. In contrast, we were interested in defining how PTSD arises from the combination of associative learning and an agent's estimation of the frequency of traumatic events. An agent has to estimate the statistical relationship between experiences of trauma based on prior experience, which can be considered a Bayesian process.

We conceive of that Bayesian process as linking nonassociative and associative learning in simulations and in model-fitting to primary mouse behavioral data in a model of PTSD. First, we posit an ecological role for nonassociative learning in estimating the frequency of predator attacks (or other violence) across life history. In these simulations, we apply a Bayesian approach to understand how well an ideal agent could estimate predation risk from its own life history and consider the relationship to early life stress. Second, we acquired behavioral freezing data in a mouse model of PTSD, which is characterized by nonassociative sensitization to threat in a new context. We compare this behavioral data with classical RL models and with a recently developed RL momentum model (Eldar et al., 2016; Eldar & Niv, 2015; Rutledge et al., 2014; Trapp et al., 2018).

Results

This study used a combination of simulation and analysis of primary mouse stress data in a mouse behavioral model to understand PTSD behavior. A Bayesian model was used to understand, via simulation, how an agent can estimate the rate of traumatic attacks over its lifetime. The mouse behavioral model is stress-enhanced fear learning (SEFL), in which unpredictable footshocks in one context lead to long-lasting sensitization to another weak footshock in a new context (Supplementary Figure 1). Reinforcement learning models can be used to compare learning processes and to integrate nonassociative learning (about the frequency of threat) with associative learning (about the associations of threat) (Eldar et al., 2016; Eldar & Niv, 2015; Rutledge et al., 2014; Trapp et al., 2018). Further simulations of this model have implications for treatment and future research into the neurobiology of PTSD.

PTSD as trauma rate estimation (simulation)

If the probability of death is high, then the animal will experience few attacks before dying (Fig. 2B). In this information-poor environment, the animal must maximize the available information in estimating the rate of such attacks. In order to determine how well an ideal observer could do under such conditions, we constructed a simple probabilistic model with a fixed probability of attacks p_a and probability of dying per attack p_d at each time point (Fig. 2A).

Even with a flat prior, p_a could be estimated with a high degree of precision by the end of the lifetime of the animal (Fig. 2C). Variance in p_a decreases progressively over the lifetime of the animal (Fig. 2D). In contrast, inferring the probability of dying on each attack (p_d) with as much certainty (Fig. 2C; Supplementary Figure 1). This observation has an intuitive explanation; the organism knows it survived n attacks but not how many organisms did not survive. The limited estimability of p_d has further implications: each attack contributes to attack rate estimate ($\hat{p}_a$) and therefore the behavior of the agent. However, because individual attacks are of less defined lethality, the organismal response to them ought to be stereotyped.

The disproportionate impact of early life stress (ELS) on adult behavior (Pavluvík et al., 2015) is explained by the Bayesian trauma rate model. We evaluated the Bayesian trauma rate estimator in two scenarios with the same total number of traumatic events: one in which traumas occur early in life (ELS), and one in which they are spread across the lifespan (Fig. 3A). The number of attacks was equal in both cases. Variance in $\hat{p}_a$decreases with time in both models, as traumatic events reduce uncertainty in the true rate of violence (Fig. 3B). However, over the course of the lifespan, the ELS model shows a higher estimated rate of violence ($\hat{p}_a$). Thus, the increased response to ELS does not require specialized critical period mechanisms but instead arises naturally in a normative estimator of violence rate.

RL momentum model (behavior data and simulation)

In this section, we compare a recently proposed RL momentum model (Eldar et al., 2016; Eldar & Niv, 2015; Rutledge et al., 2014; Trapp et al., 2018) with a new data behavioral freezing behavior in a mouse PTSD model. Classical RL models, such as temporal difference learning (Fig. 4A), enable an organism to associate threatening experiences with the context in which they are experienced. However, threats in one context do not influence

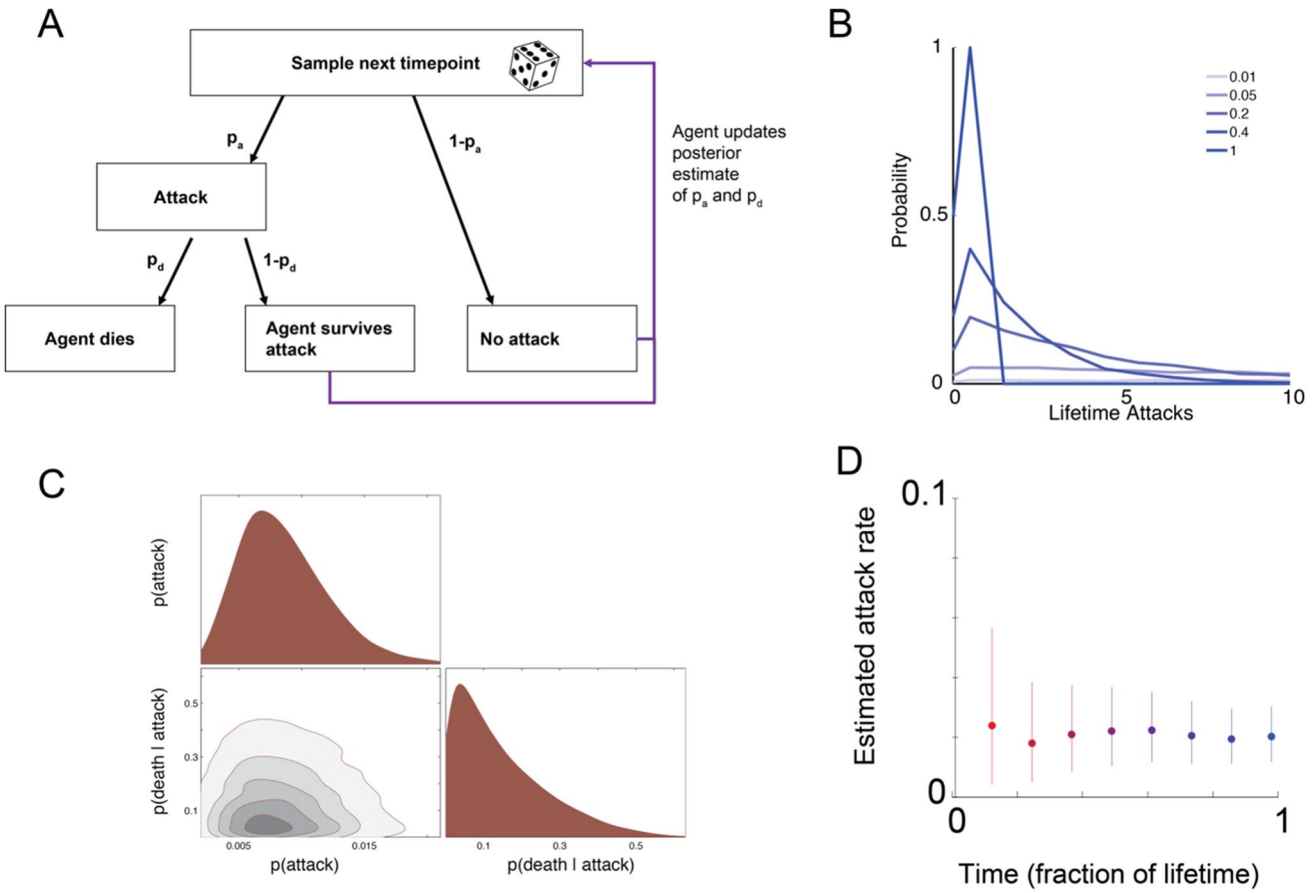

Fig. 2 Bayesian observer can measure the rate of traumatic attacks more easily than lethality. **(A)** Simplified doubly stochastic model of attacks (i.e., traumatic events). Attacks occur stochastically at each timepoint with a fixed probability p_a. Conditional on attacks occurring, agents die with probability p_d. If agents survive, they estimate the ongoing probability of attacks according to Bayes' rule. **(B)** Agents must estimate p_a and p_d in an information-poor environment. The number of attacks experienced by the typical agent is low, usually 3-5 during the course of a lifetime for $p_a = 0.2$, a typical value for the lethality of predator attacks (Rau et al., 2005). **(C)** The Bayesian estimator of p_a and p_d for a typical example sequence of attacks shows tight convergence for p_a and nonconvergence for p_d. The wider spread in p_d is indicative of the inability of the Bayesian agent to estimate the lethality of attacks. **(D)** As the agent continues over its lifetime (red to blue map), the estimate of p_a slowly narrows (vertical lines, 95% intervals). Greater time allows the agent to accumulate greater evidence about the true value of p_a

threats in another (Fig. 4A). In contrast, in the RL-momentum model, traumatic events occurring close in time but in unrelated environments contribute to a slowly varying momentum term (Fig. 4B). Momentum carries information about recent threats, allowing the agent to correctly assess risk in a changing environment. We determined how the ideal length of time for momentum to persist depends on how long threats persist using simulation (Fig. 4C). When attacks are uncorrelated in time, there is no advantage to momentum and the optimal momentum learning rate (ν) is zero. When attacks are correlated (Fig. 4C, light blue), a substantial improvement in threat estimation can be obtained by including the momentum parameter. The long timescale of optimal threat adaptation offers a potential explanation for the persistence of PTSD symptoms. If threat momentum, rather than the specific association with the initial traumatic event, were the source of PTSD symptoms, then this would have substantial implications for the understanding of PTSD.

To test this idea, we induced stress in a mouse model of PTSD (Stress-Enhanced Fear Learning; SEFL) and compared the performance of temporal difference learning (RL model) and a momentum model (RL momentum model) in explaining defensive behavior (Figure 5). In this model, mice receive unpredictable footshocks in one context (Context A) and then show sensitized threat responses to a single footshock in another context (Context B) later (Figure 5A, top). The RL momentum model fits the observed freezing behavior (Figure 5A, bottom) well, showing a disproportionate freezing response to the single footshock in a novel context (Figure 5B-C). This sensitized freezing behavior can be explained by

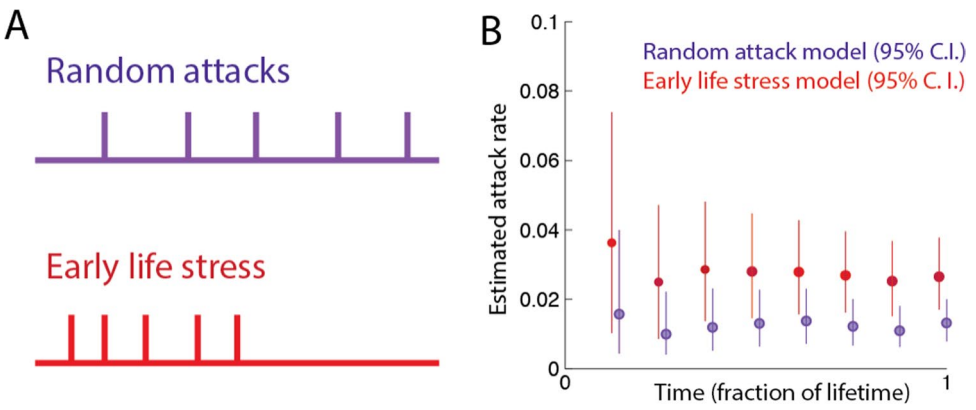

Fig. 3 Early life traumas have a disproportionate effect on the estimated attack rate. **(A)** Characteristic examples of two distributions of attack frequencies. In the random attack model (uniform attack probability), attacks are uniformly distributed across the lifespan. In the early life stress model, an identical number of attacks are uniformly distributed across the first half of the lifespan. **(B)** Two Bayesian agents (one for early life stress, one for random attack) posterior distributions for attack rate sequentially measured across the lifespan, for the random and early attack models. The discrepancy between estimated and true attack rate ($p_a = 0.01$) is greatest at the start of life due to a higher density of attacks in the early life stress model. Over the course of the lifespan, these two models converge.

the momentum term in the model, which links the threat prediction errors produced across contexts.

We compared Maximum Likelihood fits between the RL and RL momentum models (n = 18 unstressed, n = 17 stressed mice), using the Bayes Information Criteria (BIC; Fig. 5D). When the momentum learning parameter (ν) is zero, the two models are equivalent, but the the RL momentum model has a greater number of parameters (4 for RL momentum, 2 for RL model). Because the BIC penalizes the number of parameters, this produces model fits where the RL model is preferred (for unstressed mice, RL model was preferred in 17/18 animals). For stressed mice, however, the BIC strongly favored fits from the RL momentum model (14/17 animals). The RL momentum model predicts greater freezing in a novel context in stressed animals than the RL model, which accounts for the improved predictions over the RL model. This model also showed improved predictions compared with an

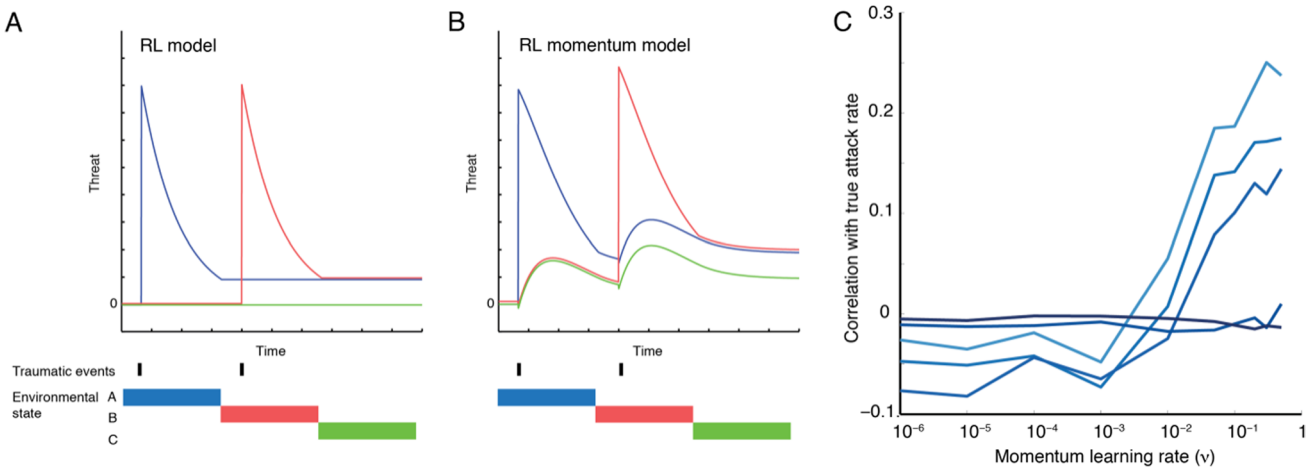

Fig. 4 Reinforcement learning with momentum allows improved estimation of autocorrelated attack rates. **(A)** Single traumatic events occur in different environmental states (contexts), leading to increased associated threat according to the RL model. **(B)** In the RL momentum model, the same series of attacks produces momentum which couples threat across contexts. Context C threat is due to momentum since the animal receives no footshocks in that state. **(C)** The momentum learning rate term of the RL momentum model enables extraction of information about fluctuating attack rates. Autoregressive attack rates were produced as shown in Figure 3 to produce 10,000 simulated attack sequences (light blue, highest autoregression to dark blue, lowest autoregression). All attacks occur in a different context. In the absence of momentum, the agent cannot extract information about fluctuations in underlying attack rate. With higher momentum, the agent can extract information about the underlying attack rate fluctuations.

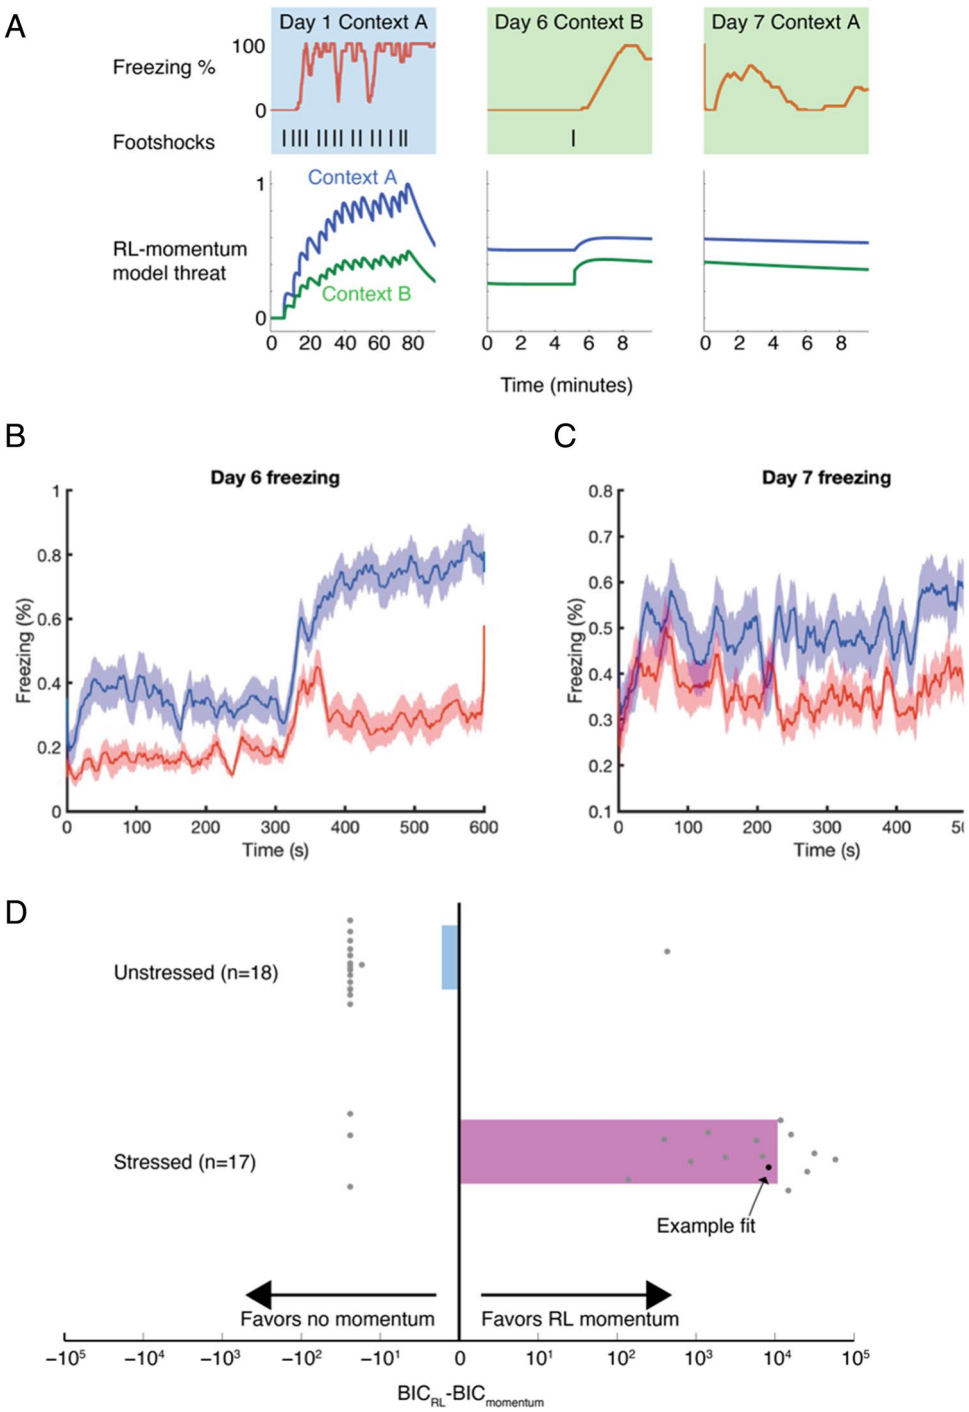

Fig. 5 RL momentum fits threat behavioral data in a mouse model of PTSD. **(A)** Example mouse behavioral data across three days of in the stress-enhanced fear learning model of PTSD (upper), along with RL momentum fit to behavioral data (lower). (upper left) Freezing across 90 minutes (red) of exposure to 15 unpredictable footshocks (black; 1mA, 1s). (upper center) Freezing across subsequent exposure to 1 uncued footshock in a new context. (upper right) Freezing during re-test in the new context (lower left) Threat according to maximum likelihood model fit of the RL momentum model (threat associated with context A – blue, context B- green) on day 1, (lower center) day 6, and (lower right) Day 7. Averaged freezing data across (n = 17 stressed, n = 18 controls) on Days 6 **(B)** and 7 **(C)**. **(D)** Model comparison between classic RL model and RL momentum model for SEFL mice (n = 17 stressed, n = 18 controls). Bayes information criterion (BIC) was calculated (see *Methods*) for maximum likelihood fits of the RL model and RL momentum model for either unstressed animals (0 shocks on Day 1) or stressed animals (15 shocks on Day 1). Difference in BIC between the two models is shown for individual animals (gray dots; black dot for example data from **(A)**), mean BIC difference per condition as bars (blue – unstressed, pink – stressed)

application of the Bayesian model of attack rate estimation (Supplementary Figure 2).

Implications for the treatment of PTSD (simulation)

The RL-momentum model of PTSD presents an additional learning mechanism by which PTSD symptoms may be ameliorated. In the classical RL model of PTSD, extinction learning (Fig. 6A) works to reduce PTSD by generating small prediction errors when the agent is reexposed to the traumatic context. This approach underlies evidence-based psychological therapies for PTSD, such as prolonged exposure and cognitive reprocessing therapy. The RL momentum model retains extinction of learned associations, but the threat prediction errors generated by extinction also generate negative momentum that reduces responses to novel threats (Fig. 6B). This model also offers a novel perspective on treatment failure of exposure therapy in PTSD.

Current learning-based accounts of this phenomenon posit that individuals may experience extinction renewal or extinction resistance, in which either extinction fails to occur or in which the extinction memory may be specific to the context in which it was generated (e.g., the therapy session). In contrast, the RL momentum proffers a simple explanation: unrelated mild stressors generate threat momentum, which increases threat associated with the original traumatic context (Fig. 6C). Similarly, an implication of this model is that exposure to novel threats independent of the traumatic context could reduce threat momentum. For example, an agent encountering an intense innate threat (e.g., standing on the side of a high cliff) without injury might experience a strong negative prediction error, which would reduce threat momentum for the same reason as exposure to a cue associated to a traumatic event.

Discussion

We formulated PTSD as a learning process directed at estimating both the rate of trauma and the specific associations with the trauma. We combined simulations of two distinct models with model fitting to primary behavioral data. The Bayesian formulation of this problem treated the agent experiencing trauma as an ideal observer. We found that the rate of traumatic events could be estimated well, but the lethality of traumatic events cannot be estimated even by an ideal observer. Early life trauma had disproportionate impact in this model even without specialized mechanisms for amplifying early life experience. We applied the reinforcement learning momentum model to PTSD and found that RL-momentum performs well when violence is clustered in time. The slower the change in trauma rate, the more momentum contributes to optimal learning from traumatic stress. This model also

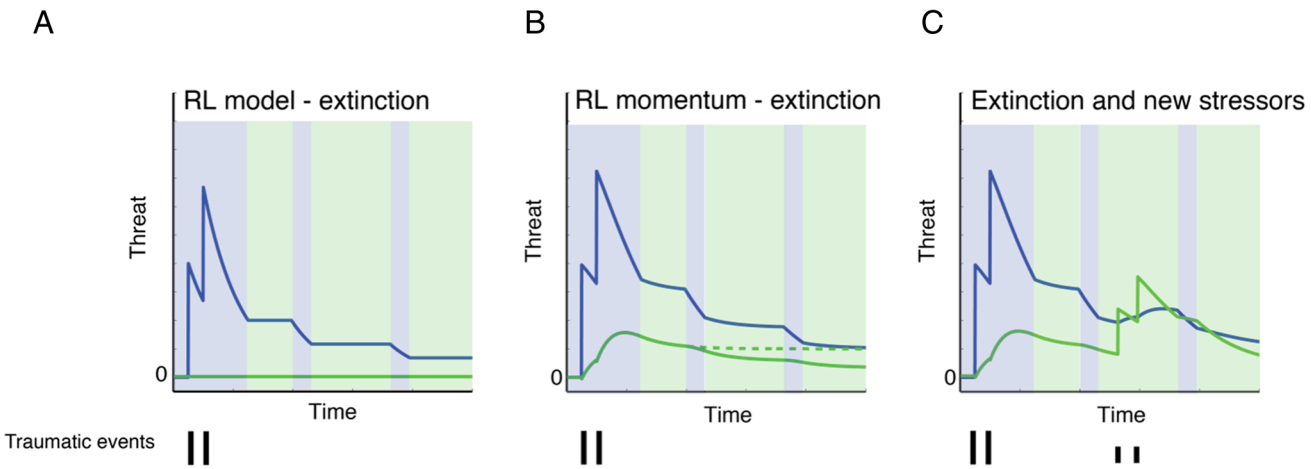

Fig. 6 RL momentum model offers a new perspective on mechanisms of extinction and symptom exacerbation in PTSD. **(A)** RL model: Two traumatic events in an initial context (context A; blue highlight) produce threat learning associated with that context (blue line) but no threat associated with a novel context (context B; green line) during exposure to that context (green highlights). Extinction occurs when exposure to the initial context A after the traumatic events causes threat prediction errors which decrease threat associated with context A (blue highlights, second and third exposures). **(B)** RL momentum model: Two traumatic events in initial context produce a momentum which increases threat in a novel context (green line). Reexposure to initial threat context (context A; blue highlights) reduces threat associated with context A (blue line) but also reduces threat momentum (green line). Green dotted line shows counterfactual threat momentum if no re-exposure to context A had occurred). **(C)** RL momentum model demonstrates a novel explanation for relapse during exposure therapy. Exposure to smaller stressors (small lines) in a novel context increases threat associated with context B (green line) but also, via the momentum term, increases threat associated with the initial traumatic context A (blue line)

offers a novel conceptualization of extinction learning and suggests that exposure to unassociated strong threats could affect threat momentum. Understanding the impact of innate danger on threat momentum requires further modeling and empirical investigation, because exposure to innate threat could lead to positive or negative changes in threat momentum.

Previous formal approaches to learning in PTSD have focused primarily on associative mechanisms. A simple model of a Bayesian observer experiencing potentially lethal attacks provided demonstrates several features that match the epidemiology of PTSD. For example, the question of why early life experiences lead to such profound effects on behavior is explained as a natural outcome of estimation bias induced by clustered trauma in early life. These obsevations may explain experimental observations of sensitization to new threats by previous stress are often used to model PTSD (Li et al., 2018; Morgan et al., 1995; Rau et al., 2005). We show that stress sensitization of threat, a model of PTSD, is well fit by the RL-momentum model. However, our ability to precisely fit the parameters of the RL-momentum model is limited by the binary nature of the stress in this dataset. Full validation and parameter-fitting for the RL-momentum model will require more precise manipulations of the sequence of threat prediction errors over time.

A further limitation of this study is that we did not consider parameter regimes that may give rise to habituation (decrease in response to repeated stimuli). Both sensitization and habituation can occur in the RL-momentum model, depending on chosen parameters (Eldar et al., 2016). In PTSD, habituation has recently been suggested as an outcome of repeated trauma (Stevens et al., 2018) and may relate to the numbing symptoms in PTSD. Habituation and sensitization have been thought of as separate processes that competitively modulate responses to repeated stimuli (Hopper et al., 2007). PTSD can involve both numbing and hyperarousal emotional reactions to stimuli after traumatic stress (Hopper et al., 2007; Krystal, 1971; Krystal, 1978). A more complete model of the impact of a sequence of threat prediction errors on subsequent emotional responses may explain this apparent contradiction.

Future progress in understanding the role of nonassociative learning in PTSD may depend on measuring the neural substrate of threat momentum (or estimated attack rate in the Bayesian model). Applying David Marr's three levels of analysis to nonassociative learning from threat (Fig. 1), we defined the computational problem ("predicting future threats based on a sequence of attacks") that must be solved. We compared two algorithms for accomplishing this goal: Bayesian attack rate estimation and RL momentum (Supplementary Figure 2). We find that the RL momentum model offers a formal mathematical approach at the implementation level, which explains clinical features of PTSD and behavior in a mouse model of PTSD. However, the implementation level of the RL momentum has not been identified.

Identifying PTSD with threat momentum may facilitate future neurobiological and translational studies of PTSD. Extensive work has shown that patients with PTSD have different learning rate parameters during fear and extinction learning than controls (Brown et al., 2018; Homan et al., 2019; Jovanovic et al., 2012; Ross et al., 2018; Yehuda et al., 1995). This study extends these findings by offering a model of how the sequence of threat prediction errors may generate other associative learning alterations in PTSD. The neurobiological correlates of threat momentum would be slowly varying summing functions of previous threat prediction errors, which sensitize defensive behaviors, such as neuromodulatory systems (Li et al., 2018) or molecular switches leading to persistent neural changes. Future extensions of this approach may link effects of arousal on learning rates (rather than overall threat) to averaged recent threat prediction errors, similar to Pearce-Hall learning (Pearce & Hall, 1980). Thus, the present study may facilitate future work linking nonassociative and associative mechanisms in PTSD. Such links are evident in behavioral and epidemiological data and have plausible biological mechanisms but have previously lacked a computational model to facilitate the design of future experiments.

Methods

In this study, mouse freezing behavior was induced by contextual fear conditioning after a stress manipulation that models sensitization effects of PTSD. This behavior was compared with predictions of reinforcement learning models to identify model-based descriptions of the stress-induced effects on contextual freezing. Simulations were used to model longer lifetime estimation of danger in an agent experiencing traumatic attacks over time. These simulations were then fit to the estimation of a Bayesian model for these attacks to better understand long lifetime estimation problems related to PTSD.

Animal behavior All procedures were performed in accordance with the ethical guidelines of the National Institutes of Health and were approved by the Institutional Animal Care & Use Committee of Yale University. Eight- to 12-week-old, C57Bl/6, male mice were stressed by using the Stress-Enhanced Fear Learning model, which has been shown to lead to long-lasting enhancement of fear and anxiety behaviors in both mice and rats (Rau et al., 2005; Sillivan et al., 2017). This model consists of 15 unpredictable footshocks (1 mA, 1 s) with random intershock intervals between 4 and 8 minutes. For contextual fear experiments, a second context (Context B) was used on Day 6, in a separate room

with different ambient auditory, visual, tactile, and olfactory characteristics. On Day 6, a single, 1-mA, 1-s shock was administered after 5 minutes, and then freezing was assessed for 5 more minutes. On Day 7, mice were returned to Context B for 10 minutes. MedAssociates boxes were used for all footshock experiments, and freezing was assessed as complete cessation of movement other than breathing (motion <18 a.u.) with automated VideoFreeze software.

Bayesian attack model At each time step, threat events (attacks) are binomially distributed with probability of attack $p_a = 0.01$ for 700 time steps (Fig. 2a). Deaths occur with probability $p_d = 0.2$ contingent on an attack occurring. Agents that die are not included in the analysis and in any comparisons attack rates are equalized to isolate the effect of estimators. The agent's estimate of p_a and p_d is derived from the sequence of attack observations ($\boldsymbol{x_t} = 0, 0, 1 \ldots 0$) according to Bayes' rule

$$p(p_a, p_d|\boldsymbol{x_t}) = \frac{p(\boldsymbol{x_t}|p_a, p_d)p(p_a, p_d)}{\int p(\boldsymbol{x_t}p_a, p_d)d\boldsymbol{x_t}} \tag{1}$$

Bayesian model fitting procedure The model is fit using a Markov Chain Monte Carlo sampler with a flat prior at time t = 0 with independent fitting performed at each time point for the events (attacks) up to that point. Specifically, an affine invariant ensemble MCMC sampler (MCMC Hammer, ref. (Akeret et al., 2013) toolbox for Matlab with 30 walkers was used to estimate the posterior for these two parameters (p_a and p_d), given the likelihood function

$$\begin{aligned} \ln\left(\left(\mathcal{L}\left(\boldsymbol{x_t}, p_a, p_d\right)\right)\right) = &\sum\nolimits_{attacks} \ln\left(p_a\left(1 - p_d\right)\right) \\ &+ \sum\nolimits_{non-attacks} \ln\left(1 - p_a\right) \\ &+ \sum\nolimits_{death} \ln\left(p_a * p_d\right) \end{aligned} \tag{2}$$

MCMC estimation is conducted with the 'gwmcmc' function in Matlab using the MCMC Hammer toolbox with 'burnin' .3 and 'stepsize' 2.

Autoregressive time series Autocorrelated attack rate time series were generated for an AR(Montague et al., 2012) autoregressive process

$$p_{a,t} = cp_{a,t-1} + \mathcal{N}(0,0.1), \tag{3}$$

where $p_{a,t}$ is the attack rate at time t, c is a constant equal to the correlation of successive time steps, and $\mathcal{N}(\mu, \sigma)$ is normally distributed noise with mean μ and standard deviation σ. Simulations used the arima function in Matlab. A total of 10,000 simulated lifetime attack rate time series were generated, then for each an agent's experienced attack time series was generated and the MCMC Hammer estimator was then used to progressively estimate attack rates as above. At each time step, the estimator was progressively estimated to generate the available estimate for an ideal Bayesian observer with information available to that time.

Reinforcement learning models In temporal difference learning, threat at time t in context c ($T_{c,t}$) is learned from a sequence of unconditioned stimuli (u_t) which produce prediction errors according to

$$T_{c,t} = T_{c,t-1} + \alpha\left(u_t - \gamma_1 T_{c,t-1}\right), \tag{4}$$

where α is a learning rate and γ_1 is a decay rate constant. Eq. 4 is referred to as RL model in the Results section, and describes the formation of associative threat learning. The addition of a momentum term (Eldar et al., 2016) allows prediction errors from different states to influence one another according to an RL momentum model

$$T_{c,t} = T_{c,t-1} + \alpha\left(u_t - \gamma_1 T_{c,t-1}\right) + fm_t, \tag{5}$$

where f is a scaling constant and m_t is the momentum at time t. This momentum term is defined by

$$m_t = m_{t-1} + \gamma_2 \sum\nolimits_{c=\{A,B,\ldots\}} \alpha\left(u_t - T_{c,t-1}\right) \tag{6}$$

in which the sum of decayed prediction errors across all contexts $c = \{A, B, \ldots\}$ with momentum decay constant γ_2. This can lead to either oscillatory behavior or slow summation of prediction errors across states depending on γ_2. Reinforcement learning models (RL – Eq. 4, RL with momentum – Eq. 5) were fit to smoothed freezing (sliding window, 15 s) on Days 1, 6, and 7. Inputs to the model were shock times and threat was fit for both Context A and Context B. Parameters for each model were fit using maximum likelihood estimation in Matlab, with threat variable rescaled to (0.1, 0.9) to match freezing probability. Maximum likelihood fit was compared by calculating the Bayes Information Criterion (BIC) for RL and RL with momentum models at the single animal level for both stressed and unstressed mice.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.3758/s13415-023-01085-5.

Declarations

Conflicts of interest No conflicts of interest are relevant to the current study. Dr. Kaye receives or has received research funding from Transcend Therapeutics and Freedom Biosciences. Dr. Krystal has received consulting fees from Aptinyx, Inc., Atai Life Sciences, AstraZeneca Pharmaceuticals, Biogen, Idec, MA, Biomedisyn Corporation, Bionomics, Limited (Australia), Boehringer Ingelheim International, Cadent Therapeutics, Inc., Clexio Bioscience, Ltd., COMPASS Pathways, Limited, United Kingdom, Concert Pharmaceuticals, Inc., Epiodyne, Inc., EpiVario, Inc., Greenwich Biosciences, Inc., Heptares Therapeutics, Limited (UK), Janssen Research & Development, Jazz

Pharmaceuticals, Inc., Otsuka America Pharmaceutical, Inc., Perception Neuroscience Holdings, Inc., Spring Care, Inc., Sunovion Pharmaceuticals, Inc., Takeda Industries, Taisho Pharmaceutical Co., Ltd, is on the board of directors at Freedom Biosciences, Inc.; he is a member of the scientific advisory board at Biohaven Pharmaceuticals, BioXcel Therapeutics, Inc. (Clinical Advisory Board), Cadent Therapeutics, Inc. (Clinical Advisory Board), Cerevel Therapeutics, LLC, Delix Therapeutics, Inc., EpiVario, Inc., Eisai, Inc., Jazz Pharmaceuticals, Inc., Novartis Pharmaceuticals Corporation, PsychoGenics, Inc., RBNC Therapeutics, Inc., Tempero Bio, Inc., Terran Biosciences, Inc.; he owns stock at Biohaven Pharmaceuticals, Sage Pharmaceuticals, and Spring Care, Inc. and owns stock options at Biohaven Pharmaceuticals Medical Sciences, EpiVario, Inc., Neumora Therapeutics, Inc., Terran Biosciences, and Inc., Tempero Bio, Inc. Dr. Krystal has also received free drug for research studies from Astra Zeneca (Saracatinib), Novartis (Mavoglurant), and Cerevel (CVL-751). Dr. Kwan is on the advisory board of Transcend Therapeutics, Freedom Biosciences, and Empyrean Neuroscience. Dr. Ressler has served on advisory boards or as a consultant for Bioxcel, Janssen, Takeda, and Verily, and he has received sponsored research support from Alkermes, Alto Neuroscience, BrainsWay, and Takeda.

References

Akeret, J., Seehars, S., Amara, A., Refregier, A., & Csillaghy, A. (2013). CosmoHammer: Cosmological parameter estimation with the MCMC hammer. *Astronomy and Computing, 2*, 27–39.

Almli, L. M., Mercer, K. B., Kerley, K., Feng, H., Bradley, B., Conneely, K. N., & Ressler, K. J. (2013). ADCYAP1R1 genotype associates with post-traumatic stress symptoms in highly traumatized African-American females. *American Journal of Medical Genetics Part B: Neuropsychiatric Genetics, 162*, 262–272.

Asmundson, G. J. G., Stapleton, J. A., & Taylor, S. (2004). Are avoidance and numbing distinct PTSD symptom clusters? *Journal of Traumatic Stress, 17*, 467–475.

Bach, D. R., & Dayan, P. (2017). Algorithms for survival: A comparative perspective on emotions. *Nature Reviews Neuroscience, 18*, 311–319.

Beck, J. G., & Sloan, D. M. (2012). *The Oxford handbook of traumatic stress disorders*. Oxford University Press.

Brown, V. M., Zhu, L., Wang, J. M., Frueh, B. C., King-Casas, B., & Chiu, P. H. (2018). Associability-modulated loss learning is increased in posttraumatic stress disorder. *eLife*, (*7*).

Callaghan, B. L., & Tottenham, N. (2016). The stress acceleration hypothesis: Effects of early-life adversity on emotion circuits and behavior. *Current Opinion in Behavioral Sciences, 7*, 76–81.

Eldar, E., & Niv, Y. (2015). Interaction between emotional state and learning underlies mood instability. *Nature Communications, 6*, 6149.

Eldar, E., Rutledge, R. B., Dolan, R. J., & Niv, Y. (2016). Mood as representation of momentum. *Trends in Cognitive Sciences, 20*, 15–24.

Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: A Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews Neuroscience, 10*, 48–58.

Foa, E. B. (2011). Prolonged exposure therapy: Past, present, and future. *Depression and Anxiety, 28*, 1043–1047.

Frankenhuis, W. E., & de Weerth, C. (2013). Does early-life exposure to stress shape or impair cognition? *Current Dir Psychology Science, 22*, 407–412.

Galatzer-Levy, I. R., Karstoft, K.-I., Statnikov, A., & Shalev, A. Y. (2014). Quantitative forecasting of PTSD from early trauma responses: A machine learning application. *Journal of Psychiatric Research, 59*, 68–76.

Galatzer-Levy, I. R., Ma, S., Statnikov, A., Yehuda, R., & Shalev, A. Y. (2017). Utilization of machine learning for prediction of post-traumatic stress: A re-examination of cortisol in the prediction and pathways to non-remitting PTSD. *Translational Psychiatry, 7*, e1070.

Groves, P. M., & Thompson, R. F. (1970). Habituation: A dual-process theory. *Psychological Review, 77*, 419–450.

Homan, P., Levy, I., Feltham, E., Gordon, C., Hu, J., Li, J., Pietrzak, R. H., Southwick, S., Krystal, J. H., Harpaz-Rotem, I., et al. (2019). Neural computations of threat in the aftermath of combat trauma. *Nature Neuroscience, 22*, 470.

Hopper, J. W., Frewen, P. A., van der Kolk, B. A., & Lanius, R. A. (2007). Neural correlates of reexperiencing, avoidance, and dissociation in PTSD: Symptom dimensions and emotion dysregulation in responses to script-driven trauma imagery. *Journal of Traumatic Stress, 20*, 713–725.

Ilany, A., & Eilam, D. (2008). Wait before running for your life: Defensive tactics of spiny mice (Acomys cahirinus) in evading barn owl (Tyto alba) attack. *Behavorial Ecology and Sociobiology, 62*, 923–933.

Izquierdo, A., Wellman, C. L., & Holmes, A. (2006). Brief uncontrollable stress causes dendritic retraction in infralimbic cortex and resistance to fear extinction in mice. *Journal of Neuroscience, 26*, 5733–5738.

Jovanovic, T., Kazama, A., Bachevalier, J., & Davis, M. (2012). Impaired safety signal learning may be a biomarker of PTSD. *Neuropharmacology, 62*, 695–704.

Kelmendi, B., & Southwick, S. M. (2018). Locus Coeruleus hyperactivity in posttraumatic stress disorder: Answers and questions. *Biological Psychiatry, 83*, 197–199.

Khoury, L., Tang, Y. L., Bradley, B., Cubells, J. F., & Ressler, K. J. (2010). Substance use, childhood traumatic experience, and posttraumatic stress disorder in an urban civilian population. *Depression and Anxiety, 27*, 1077–1086.

Korem, N., et al. (2022). Emotional numbing in PTSD is associated with lower amygdala reactivity to pain. *Neuropsychopharmacology, 1–9*.

Krystal, H. (1971). Trauma: Considerations of its intensity and chronicity. *International Psychiatry Clinics, 8*, 11–28.

Krystal, H. (1978). Trauma and affects. *The Psychoanalytic Study of the Child, 33*, 81–116.

Levy, I., & Schiller, D. (2021). Neural computations of threat. *Trends in Cognitive Sciences, 25*, 151–171.

Li, L., Feng, X., Zhou, Z., Zhang, H., Shi, Q., Lei, Z., Shen, P., Yang, Q., Zhao, B., Chen, S., et al. (2018). Stress accelerates defensive responses to looming in mice and involves a locus coeruleus-superior colliculus projection. *Current Biology, 28*, 859–871.e5.

Maddox, S. A., Hartmann, J., Ross, R. A., & Ressler, K. J. (2019). Deconstructing the gestalt: Mechanisms of fear, threat, and trauma memory encoding. *Neuron, 102*, 60–74.

Marr, D. (2010). *Vision: A computational investigation into the human representation and processing of visual information*. (MIT Press).

Montague, P. R., Dolan, R. J., Friston, K. J., & Dayan, P. (2012). Computational psychiatry. *Trends in Cognitive Sciences, 16*, 72–80.

Morgan, C. A., Grillon, C., Southwick, S. M., Davis, M., & Charney, D. S. (1995). Fear-potentiated startle in posttraumatic stress disorder. *Biological Psychiatry, 38*, 378–385.

Pavluvík, P., Poprach, K., Machar, I., Losík, J., Gouveia, A., & Tkadlec, E. (2015). Barn owl productivity response to variability of vole populations. *PLoS One, 10*.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review, 87*, 532–552.

Powers, A. R., Mathys, C., & Corlett, P. R. (2017). Pavlovian conditioning–induced hallucinations result from overweighting of perceptual priors. *Science, 357*, 596–600.

Rau, V., DeCola, J. P., & Fanselow, M. S. (2005). Stress-induced enhancement of fear learning: An animal model of posttraumatic stress disorder. *Neuroscience & Biobehavioral Reviews, 29*, 1207–1223.

Ross, M. C., Lenow, J. K., Kilts, C. D., & Cisler, J. M. (2018). Altered neural encoding of prediction errors in assault-related posttraumatic stress disorder. *Journal of Psychiatric Research, 103*, 83–90.

Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences United States of America, 111*, 12252–12257.

Saxe, G. N., Ma, S., Ren, J., & Aliferis, C. (2017). Machine learning methods to predict child posttraumatic stress: A proof of concept study. *BMC Psychiatry, 17*, 223.

Sillivan, S. E., Joseph, N. F., Jamieson, S., King, M. L., Chévere-Torres, I., Fuentes, I., Shumyatsky, G. P., Brantley, A. F., Rumbaugh, G., & Miller, C. A. (2017). Susceptibility and resilience to posttraumatic stress disorder-like behaviors in inbred mice. *Biology Psychiatry, 82*, 924–933.

Stevens, J., Ressler, K., & Jovanovic, T. (2018). T22. PTSD symptom profiles and amygdala function vary as a function of repeated trauma exposure: Numbing as a specific neurobiological phenotype. *Biological Psychiatry, 83*, S137.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. (MIT Press).

Thompson, R. F., & Spencer, W. A. (1966). Habituation: A model phenomenon for the study of neuronal substrates of behavior. *Psychological Review, 73*, 16–43.

Trapp, S., O'Doherty, J. P., & Schwabe, L. (2018). Stressful events as teaching signals for the brain. *Trends in Cognitive Science (Regul Ed.), 22*, 475–478.

Rescorla, R.A., & Wagner, A.R. (1972). A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In AH. Black & W.F. Prokasy (eds.), Classical conditioning II: Current research and theory (pp. 64-99) New York: Appleton-Century-Crofts.

Wagner, A. R. (1979). Habituation and memory. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation: A memorial volume for Jerry Konorski* (pp. 53–82). Lawrence Earlbaum Assoc.; Hillsdale.

Wang, X.-J., & Krystal, J. H. (2014). Computational psychiatry. *Neuron, 84*, 638–654.

Yehuda, R., Keefe, R. S. E., Harvey, P. D., Levengood, R. A., Gerber, D. K., Geni, J., & Siever, L. J. (1995). Learning and memory in combat veterans with posttraumatic stress disorder. *The American Journal of Psychiatry, 152*, 137–139.