Biological Psychiatry

Archival Report

Modeling Avoidance in Mood and Anxiety Disorders Using Reinforcement Learning

Anahit Mkrtchian, Jessica Aylward, Peter Dayan, Jonathan P. Roiser, and Oliver J. Robinson

ABSTRACT

BACKGROUND: Serious and debilitating symptoms of anxiety are the most common mental health problem worldwide, accounting for around 5% of all adult years lived with disability in the developed world. Avoidance behavior—avoiding social situations for fear of embarrassment, for instance—is a core feature of such anxiety. However, as for many other psychiatric symptoms the biological mechanisms underlying avoidance remain unclear.

METHODS: Reinforcement learning models provide formal and testable characterizations of the mechanisms of decision making; here, we examine avoidance in these terms. A total of 101 healthy participants and individuals with mood and anxiety disorders completed an approach-avoidance go/no-go task under stress induced by threat of unpredictable shock.

RESULTS: We show an increased reliance in the mood and anxiety group on a parameter of our reinforcement learning model that characterizes a prepotent (pavlovian) bias to withhold responding in the face of negative outcomes. This was particularly the case when the mood and anxiety group was under stress.

CONCLUSIONS: This formal description of avoidance within the reinforcement learning framework provides a new means of linking clinical symptoms with biophysically plausible models of neural circuitry and, as such, takes us closer to a mechanistic understanding of mood and anxiety disorders.

Keywords: Anxiety, Avoidance, Diathesis–stress, Pavlovian bias, Reinforcement learning, Threat of shock

http://dx.doi.org/10.1016/j.biopsych.2017.01.017

Avoidance is a core feature of anxiety (1,2) and plays a central role in psychological strategies for the treatment of anxiety (3), but its underlying neural and cognitive mechanisms are unknown. Avoidance can be adaptive: if an individual perceives a situation as stressful then it makes sense to avoid that stressor in the future. However, excessive avoidance can result in a pathological downward spiral. The more one avoids a situation, the less opportunity there is to learn that the situation is not as bad as feared, and a vicious cycle of avoidance and impaired extinction learning emerges, which in turn promotes further anxiety (1). For example, an individual who fears social embarrassment might ultimately end up housebound, avoiding all social interaction.

The diathesis-stress model of mood and anxiety disorders (4) proposes that maladaptive avoidance should be greatest during periods of environmental stress in vulnerable individuals. This idea has clear face validity and is supported by clinical anecdotes but is largely derived from retrospective, subjective self-reports. This is because quantifying avoidance under stress in an experimentally controlled yet ecologically valid manner in humans is methodologically challenging. In this study we address this challenge using 1) a translationally validated [i.e., comparable behavioral responses can be elicited across human and animal models (5)] threat-of-shock procedure to induce stress (6,7); 2) a cognitive task that has been shown to reliably index avoidance behavior in healthy individuals (1); and 3) a computationally precise method of defining of avoidance.

Specifically, we operationalize avoidance as a behavioral bias toward withholding action (no-go [i.e., inhibition]) in the face of potentially negative outcomes. This powerful prepotent reflexive (or pavlovian) bias has been observed consistently in humans and animals (8–11) and is so profound that it can disrupt instrumental goal-directed behavior (8–11). This is known as pavlovian-instrumental transfer (12), and we harness it here to measure the degree to which individuals rely on their prepotent avoidance biases. Given that both induced stress (13,14) and pathological anxiety have been associated with increased inhibitory control, it seems plausible that a combination of stress and anxiety will increase reliance on pavlovian inhibitory avoidance biases (15) [in contrast with depression alone, which might plausibly be associated with reduced reliance on pavlovian approach biases (16)].

Reinforcement learning algorithms can provide parameterizations of avoidance behavior that offer insight into both optimal behavior when set correctly (17) and to dysfunction and pathology when set incorrectly (18). Critically, reinforcement learning models enable us to parameterize the influence of pavlovian avoidance biases on task performance in a formal manner. A large body of work has applied these models to healthy humans (8–10) and they form the basis of human-level artificial intelligence (17), but to date they have not been applied to individuals with mood and anxiety disorders.

 ISSN: 0006-3223

We therefore tested individuals with mood and anxiety disorders and healthy individuals completing an approach-avoidance go/no-go task under stress, which was induced by threat of shock. Avoidance was defined and parameterized within a reinforcement learning framework. We predicted that the mood and anxiety group would show high reliance on avoidance bias and that this avoidance bias would be exacerbated by stress.

METHODS AND MATERIALS

Participants

All data, task scripts, and code to recreate the figures in this article are freely available online (https://figshare.com/articles/Avoidance_Anxiety_Materials/3860250). A total of 101 participants were included in the study. Healthy participants (n = 58 [originally n = 62 but 4 individuals were excluded because they failed to follow task instructions]; 36 men [62.1%]; age range = 18–57 years; mean ± SD age = 26.7 ± 7.1 years) and unmedicated individuals with pathological mood and anxiety symptoms (n = 43; 27 men [62.8%]; age range = 18–53 years; mean ± SD age = 28.8 ± 8.8 years) were recruited from online advertising and institutional subject databases. The primary difference between the groups in initial recruitment was that only the pathological group self-defined as experiencing distress from mood/anxiety symptoms. We recruited a mixed sample of anxiety and depression diagnoses because they are highly comorbid with overlapping symptoms and may not therefore represent truly distinct pathologies. Healthy participants responded to an advertisement asking for healthy individuals with no psychiatric symptoms. A phone screen confirmed no history of psychiatric, neurological, or substance use disorders. The mood and anxiety group responded to an advertisement for individuals suffering from low mood, anxious, or depressive symptoms. Following an initial phone screen, individuals who met criteria for mood or anxiety disorder symptomatology according to a face-to-face Mini-International Neuropsychiatric Interview (19) were included. According to the Mini-International Neuropsychiatric Interview, the majority of participants (n = 27) met criteria for both generalized anxiety disorder and major depressive disorder (MDD) (n = 9 with additional panic disorder), generalized anxiety disorder (n = 8; n = 3 with panic disorder, n = 1 with agoraphobia), panic disorder and MDD (n = 2), and MDD alone (n = 6; Supplemental Table S1). The average number of depressive episodes was 5 ± 7. The average duration of episodes was 7 ± 8 months (excluding one participant who reported a continuous episode since adolescence). Further details are provided in the Supplement.

Manipulation

State anxiety was induced via threat of unpredictable electric shocks delivered with two electrodes attached to the nondominant wrist using a Digitimer DS5 Constant Current Stimulator (Digitimer Ltd., Welwyn Garden City, United Kingdom). A highly unpleasant (but not painful) subjective shock level was established using a shock work-up procedure prior to testing. No more than five (to avoid habituation) shocks with gradual increasing shock level were administered. Participants rated each shock on a scale from 1 (barely felt) to 5 (unbearable). Shock level was matched at a level of four across participants. The experimental task was programmed in Psychtoolbox-3 (http://psychtoolbox.org) for MATLAB R2014b (version 8.4.0.1) (The MathWorks, Inc., Natick, MA), presented on a laptop and administered under alternating safe and threat blocks. During the safe block, the background color was blue and preceded by a 4000-ms message stating, “You are now safe from shock.” During the threat block, the background color was red and the message stating “Warning! You are now at risk of shock” was presented for 4000 ms. Participants were told that they might receive a shock only during the threat condition but that the shocks were not dependent on their performance. In practice, a single shock was delivered at a pseudorandom time point during one third of threat blocks (a total of four shocks across 480 trials). Note that it is the anticipation of these shocks, not the shocks themselves, that constitutes the manipulation (see the Supplement). At the end of each experimental task, participants retrospectively rated how anxious they felt during the safe and threat conditions on a 10-point Likert-type scale with responses ranging from 1 (not at all) to 10 (very much so).

Approach-Avoidance Task

The task was based on the design of a previous probabilistic go/no-go reinforcement learning task (10,20) modified to incorporate the threat manipulation. The prepotent pavlovian bias to a win is a go response (approach), and the prepotent pavlovian response to a loss is a no-go (avoid) response. As such, the task comprised four experimental conditions where action (go/no-go) was crossed with valence (reward/punishment): 1) go to win reward, 2) go to avoid losing (GA), 3) no-go to win reward (NGW), and 4) no-go to avoid losing. On each trial, participants were presented with one of four fractal cues per condition, followed by a target detection task and subsequently by a probabilistic outcome (Figure 1; more task detail in the Supplement).

Reinforcement Learning Models

Reinforcement learning modeling proceeded in the same way as described in a prior article (10). Briefly, we built seven parameterized reinforcement learning models to fit to the behavior of the subjects. All models were adapted Rescorla-Wagner models. We use the term “standard” to denote the six-parameter winning model from Guitart-Masip *et al.* (10) and either add or subtract parameters to test model fits for seven separate models (see Table 1 for a parameter specification summary).

Learning Models. All the models assigned a probability to each action a_t on trial t based on an action weight and the current stimulus. The action weights were constructed according to a simple Rescorla-Wagner–like update equation with a learning rate. Reinforcements were coded as +1 for a reward, –1 for a punishment, and 0 for no feedback. A sensitivity parameter determined the effective size of reinforcements for a subject. For the majority of models the sensitivity parameter could take on different values for the reward and punishment