Examples of Reinforcement Learning Models applied to Computational Psychiatry:

Application to Depression and Anxiety

Peggy Seriès, IANC
Informatics, University of Edinburgh, UK

pseries@inf.ed.ac.uk

CCN Lecture 11

A New Model for Mental Illness

Mental illness is the result of an impairment in prediction, due to having a distorted internal model of the world, possibly due to an impairment in learning.

Applications of RL models to Computational Psychiatry

- RL models have been used to model almost all psychiatric disorders.
- **idea**: disorder can be understood as impairment in learning/decision-making.

- Today:
  - Depression
  - Anxiety

Major Depression Disorder (MDD)

- High (lifetime) prevalence (esp. in developed countries)
  - USA: 16.2% (Kessler et al., 2003)
  - UK / Europe: 7-10% (Ayuso-Mateos et al., 2001)
  - Depression rates are rising (e.g. Mojtabai et al., 2016)

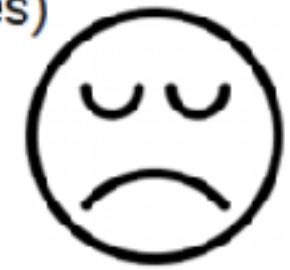

- High economic impact (Europe: €92 billion in 2010) (Olesen et al., 2012)
- People are suffering
  - Risk factor for suicide (Olfson et al., 2017)
    - And suicide rates are increasing
  - Cognitive Impairments (e.g. Snyder, 2013)
    - Attention, concentration, executive functioning, working memory, …
  - Impairments in Reinforcement Learning (Chen et al., 2015)

Major Depression Disorder (MDD)

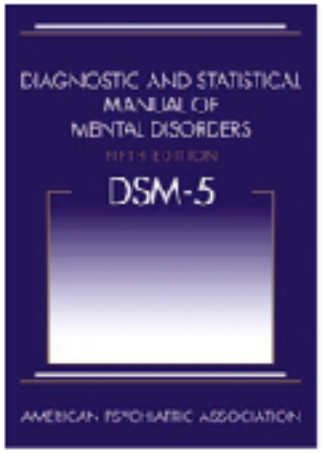

- Core symptoms:
  - **Depressed mood**
  - **Anhedonia** (inability to experience pleasure)
  - **Loss of energy**
  - Change in weight or appetite
  - Insomnia / Hypersomnia
  - Psychomotor agitation / retardation
  - Concentration difficulties
  - Suicidal thoughts / ideation

Major Depressive Disorder

Diagnostic Criteria

A. Five (or more) of the following symptoms have been present during the same 2-week period and represent a change from previous functioning; at least one of the symptoms is either (1) depressed mood or (2) loss of interest or pleasure.

Note: Do not include symptoms that are clearly attributable to another medical condition.

1. Depressed mood most of the day, nearly every day, as indicated by either subjective report (e.g., feels sad, empty, hopeless) or observation made by others (e.g., appears tearful). **(Note:** In children and adolescents, can be irritable mood.)
2. Markedly diminished interest or pleasure in all, or almost all, activities most of the day, nearly every day (as indicated by either subjective account or observation).
3. Significant weight loss when not dieting or weight gain (e.g., a change of more than 5% of body weight in a month), or decrease or increase in appetite nearly every day. **(Note:** In children, consider failure to make expected weight gain.)
4. Insomnia or hypersomnia nearly every day.
5. Psychomotor agitation or retardation nearly every day (observable by others, not merely subjective feelings of restlessness or being slowed down).
6. Fatigue or loss of energy nearly every day.
7. Feelings of worthlessness or excessive or inappropriate guilt (which may be delusional) nearly every day (not merely self-reproach or guilt about being sick).
8. Diminished ability to think or concentrate, or indecisiveness, nearly every day (either by subjective account or as observed by others).
9. Recurrent thoughts of death (not just fear of dying), recurrent suicidal ideation without a specific plan, or a suicide attempt or a specific plan for committing suicide.

Major Depression Disorder (MDD)

- Categorical view has little basis in biology?
  - Research moves towards dimensional view
- RDoC framework
  - Multiple levels of analysis
    - Neural circuitry, genes, behaviour
- Endophenotypes
  - Anhedonia inability to experience pleasure
  - Neuroticism

    a broad personality trait dimension representing the degree to which a person experiences the world as distressing, threatening, and unsafe.

Treatments

- Cognitive Behavioural Therapy (CBT)
- Antidepressant medication
  - Selective Serotonin Reuptake Inhibitors (SSRIs)
    - Primary first line treatment
  - Serotonin-Norepinephrine Reuptake Inhibitor (SNRIs)
  - Tricyclic Antidepressants (TCAs)
- Electroconvulsive therapy (ECT), Surgery
  - Very severe, treatment-resistant cases

Theories: Cognitive Theory

- **Aaron Beck 2008**

Depression results from **negative thinking patterns**, known as the cognitive triad:

1. Negative views of oneself
2. Negative views of the world
3. Negative views of the future

Cognitive distortions, such as overgeneralization and catastrophizing, reinforce depression.

Cognitive Behavioral Therapy (CBT) helps change these thought patterns

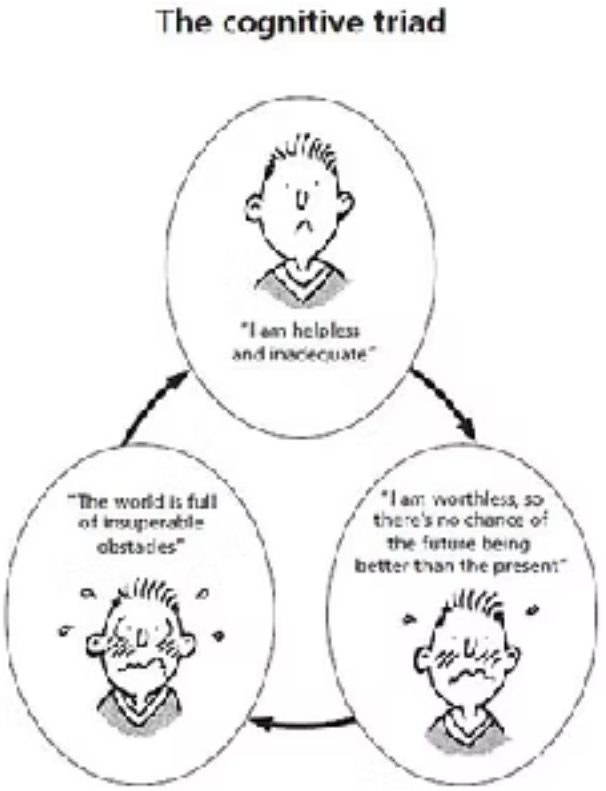

Theories: “Chemical imbalance”

Neurochemical (Monoamine) Theory

- linked to low levels of **serotonin, dopamine, and norepinephrine**—neurotransmitters that regulate mood, motivation, and pleasure.
- Antidepressants like **SSRIs (e.g., Prozac)** work by increasing serotonin levels.
- However, theory is incomplete since antidepressants take weeks to work, suggesting other mechanisms are involved.
- Despite popularity, serotonin theory of depression is highly debated
-

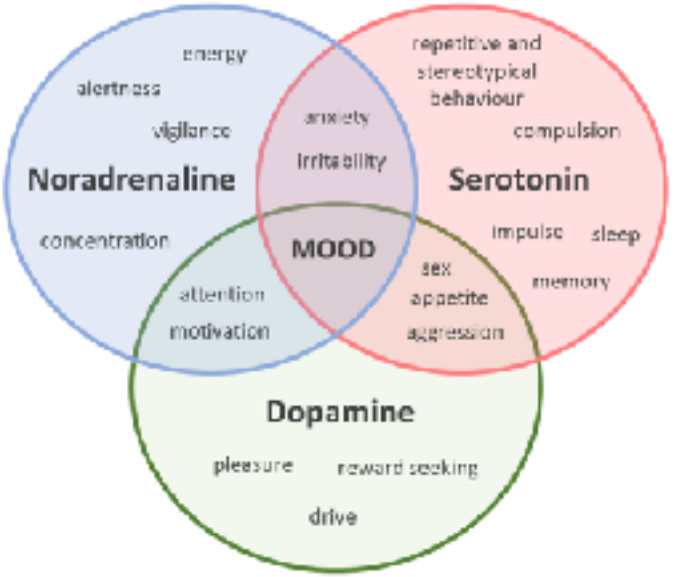

Theories: the role of stress

Neuroendocrine and stress theories

- Depression associated with **overactivity of the Hypothalamic-Pituitary-Adrenal (HPA) axis, leading to high cortisol levels (stress hormone)**.

- **Chronic stress** may damage brain regions like the hippocampus, impairing memory and emotion regulation.

- Stress could lead to deficit in reinforcement /**reward processing** and anhedonia (Pizzagalli, 2014).

- 70-80% of major depression events preceded by **major life event** (e.g., loss of a loved one, financial problems, trauma).

Depression and Reinforcement learning

- **Could depression be related to an impairment in learning and an underestimation of rewarding/positive outcomes?**

  That could potentially explain anhedonia, lack of motivation, negative biases etc..

Is Depression related to an impairment in RL?

- Depressed patients have shown deficits in RL (Chen et al 2015 for a review): important windows on the disease?

- Are impairments related to lower reward sensitivity, greater punishment sensitivity, or an inability to learn from reward (learning rate) or increased learning from punishment?

- All could explain a negative bias towards the processing of punishment

- but different implications for therapies (focus on experience quality of reward/punishment or focus on change in behaviour following reward or punishment)

Learning and Inference

- Learning: predict; control

  Δ weight α (learning rate) x (error) x (stimulus)

  - dopamine
    phasic prediction error for future reward
  - serotonin
    phasic prediction error for future punishment
  - acetylcholine
    expected uncertainty boosts learning
  - norepinephrine
    unexpected uncertainty boosts learning

Is Depression related to an impairment in RL?

- Signal detection task (Pizzagalli et al, *Biol Psychiatry*, 2004)

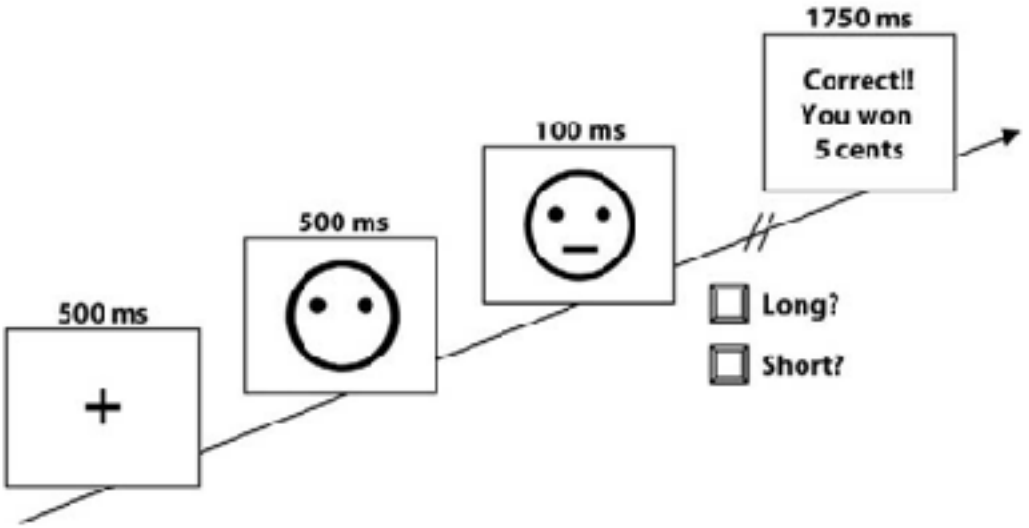

Figure 1. Schematic diagram of the task. After presentation of the mouth stimulus, subjects selected which type of mouth had been presented by pressing either the "z" or the "/" key (counterbalanced across subjects).

- One stimulus rewarded more often
  - (healthy) participants become biased towards it

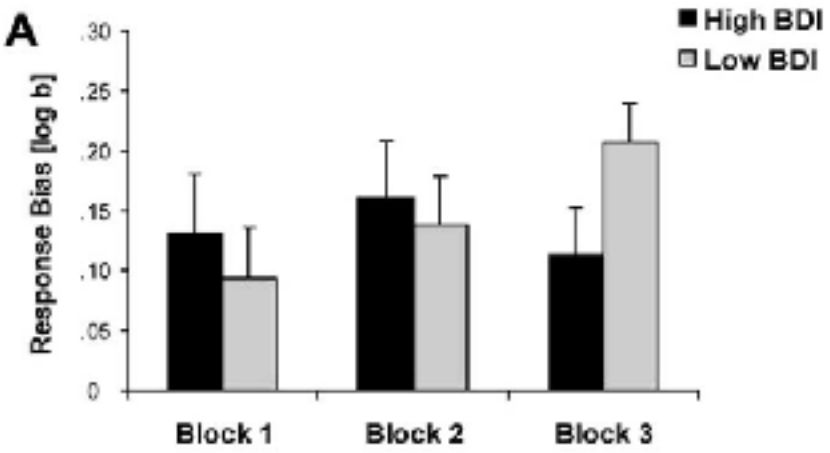

Is Depression related to an impairment in RL?

- Use probabilistic reinforcement learning models to fit the task
- Each option V_a and V_b is iteratively updated using:

$$V(t+1) = m \times V(t) + \varepsilon \times (\rho \times r(t) - V(t))$$

- Decide between the two options:

$$p(a \mid V, \theta) = \frac{1}{1 + \exp(-\beta \times (V_a - V_b))}$$

If learning from rewards is what is impaired in depression, what differs between controls and patients: **learning rate** (ϵ), **reward sensitivity** (ρ), or **stochasticity** in the response (β)?

Huys et al. Biology of Mood & Anxiety Disorders 2013, 3:12
http://www.biolmoodanxietydisord.com/content/3/1/12

Biology of Mood & Anxiety Disorders

RESEARCH Open Access

Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis

Quentin JM Huys[1,2,3,4*], Diego A Pizzagalli[5], Ryan Bogdan[6] and Peter Dayan[1]

Abstract

Background: Depression is characterised partly by blunted reactions to reward. However, tasks probing this deficiency have not distinguished insensitivity to reward from insensitivity to the prediction errors for reward that determine learning and are putatively reported by the phasic activity of dopamine neurons. We attempted to disentangle these factors with respect to anhedonia in the context of stress, Major Depressive Disorder (MDD), Bipolar Disorder (BPD) and a dopaminergic challenge.

Methods: Six behavioural datasets involving 392 experimental sessions were subjected to a model-based, Bayesian meta-analysis. Participants across all six studies performed a probabilistic reward task that used an asymmetric reinforcement schedule to assess reward learning. Healthy controls were tested under baseline conditions, stress or after receiving the dopamine D_2 agonist pramipexole. In addition, participants with current or past MDD or BPD were evaluated. Reinforcement learning models isolated the contributions of variation in reward sensitivity and learning rate.

A: the reward sensitivity

SCIENTIFIC REPORTS

OPEN

Major Depression Impairs the Use of Reward Values for Decision-Making

Samuel Rupprechter[1], Aistis Stankevicius[1], Quentin J. M. Huys[2,3], J. Douglas Steele[4] & Peggy Seriès[1]

Received: 2 November 2017
Accepted: 16 August 2018
Published online: 14 September 2018

Depression is a debilitating condition with a high prevalence. Depressed patients have been shown to be diminished in their ability to integrate their reinforcement history to adjust future behaviour during instrumental reward learning tasks. Here, we tested whether such impairments could also be observed in a Pavlovian conditioning task. We recruited and analysed 32 subjects, 15 with depression and 17 healthy controls, to study behavioural group differences in learning and decision-making. Participants had to estimate the probability of some fractal stimuli to be associated with a binary reward, based on a few passive observations. They then had to make a choice between one of the observed fractals and another target for which the reward probability was explicitly given. Computational modelling was used to succinctly describe participants' behaviour. Patients performed worse than controls at the task. Computational modelling revealed that this was caused by behavioural impairments during both learning and decision phases. Depressed subjects showed lower memory of observed rewards and had an impaired ability to use internal value estimations to guide decision-making in our task.

A: the beta of the softmax and memory parameter

JAMA Psychiatry | Original Investigation

Reinforcement Learning in Patients With Mood and Anxiety Disorders vs Control Individuals

A Systematic Review and Meta-analysis

Alexandra C. Pike, DPhil; Oliver J. Robinson, PhD

IMPORTANCE Computational psychiatry studies have investigated how reinforcement learning may be different in individuals with mood and anxiety disorders compared with control individuals, but results are inconsistent.

OBJECTIVE To assess whether there are consistent differences in reinforcement-learning parameters between patients with depression or anxiety and control individuals.

DATA SOURCES Web of Knowledge, PubMed, Embase, and Google Scholar searches were performed between November 15, 2019, and December 6, 2019, and repeated on December 3, 2020, and February 23, 2021, with keywords (*reinforcement learning*) AND (*computational* OR *model*) AND (*depression* OR *anxiety* OR *mood*).

STUDY SELECTION Studies were included if they fit reinforcement-learning models to human choice data from a cognitive task with rewards or punishments, had a case-control design including participants with mood and/or anxiety disorders and healthy control individuals, and included sufficient information about all parameters in the models.

DATA EXTRACTION AND SYNTHESIS Articles were assessed for inclusion according to MOOSE guidelines. Participant-level parameters were extracted from included articles, and a

Supplemental content

A: higher sensitivity to punishment, lower sensitivity to reward, the beta of the softmax

Major Depression Disorder (MDD): Learned Helplessness

- Psychological concept developed by Martin Seligman and coll. in 1960s.
- The state in which an individual, after experiencing repeated, negative events, which it is unable to escape from or avoid, stops trying to change their situation—even when opportunities to escape or improve things become available.
- **Animal model of depression**: mirrors loss of motivation, passivity, resignation, feelings of hopelessness, cognitive impairments: emotional distress.

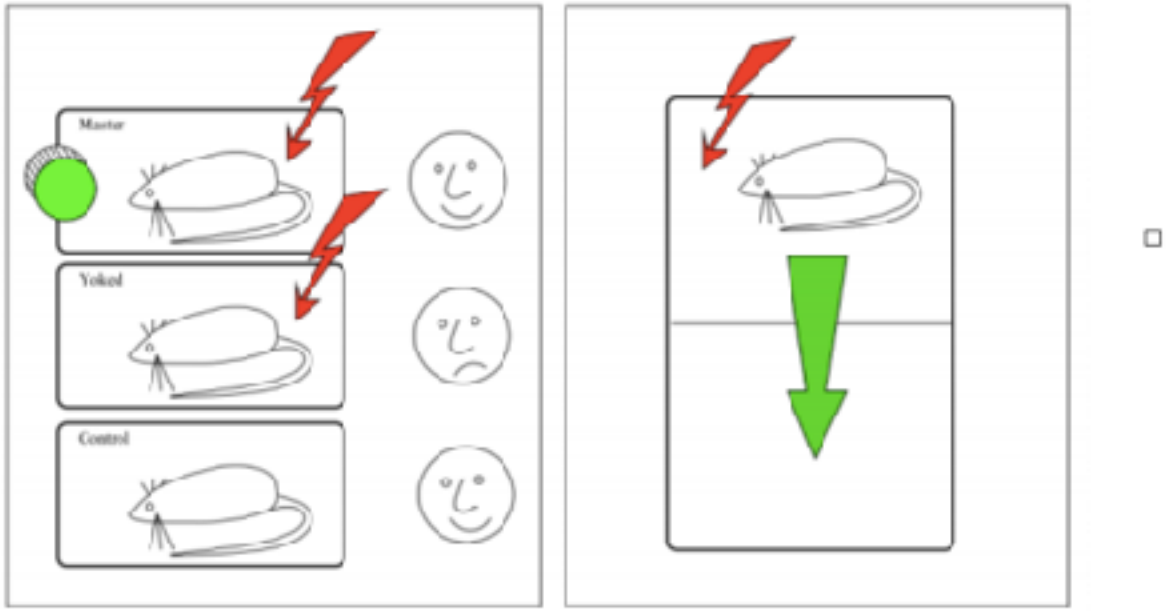

Not (only) about reward/values, but control?

Cognition
Volume 113, Issue 3, December 2009, Pages 314-328

A Bayesian formulation of behavioral control

Quentin J.M. Huys [a, b], Peter Dayan [a]

Show more

+ Add to Mendeley Share Cite

https://doi.org/10.1016/j.cognition.2009.01.008 Get rights and content

Abstract

Helplessness, a belief that the world is not subject to behavioral control, has long been central to our understanding of depression, and has influenced cognitive theories, animal models and behavioral treatments. However, despite its importance, there is no fully accepted definition of helplessness or behavioral control in psychology or psychiatry, and the formal treatments in engineering appear to capture only limited aspects of the intuitive concepts. Here, we formalize controllability in terms of characteristics of prior distributions over affectively charged environments. We explore the relevance of this notion of control to

Neuroscience & Biobehavioral Reviews
Volume 102, July 2019, Pages 371-381

Motivation and cognitive control in depression

Ivan Grahek [a], Amitai Shenhav [b], Sebastian Musslick [c], Ruth M. Krebs [d], Ernst H.W. Koster [a]

Show more

+ Add to Mendeley Share Cite

https://doi.org/10.1016/j.neubiorev.2019.04.011 Get rights and content

Highlights

- Review of motivational and cognitive impairments in depression.
- Impairments in key components of motivation contribute to cognitive control deficits in depression.
- Proposed framework focuses on the reduced value of control in depression.

"The environment is not controllable, so why bother."

Not (only) about reward/values, but the cost of effort?

Elevated Effort Cost Identified by Computational Modeling as a Distinctive Feature Explaining Multiple Behaviors in Patients With Depression

Fabien Vinckier, Claire Jaffre, Claire Gauthier, Sarah Smajda, Pierre Abdel-Ahad, Raphaël Le Bouc, Jean Daunizeau, Mylène Fefeu, Nicolas Borderies, Marion Plaze, Raphaël Gaillard, and Mathias Pessiglione

ABSTRACT

BACKGROUND: Motivational deficit is a core clinical manifestation of depression and a strong predictor of treatment failure. However, the underlying mechanisms, which cannot be accessed through conventional questionnaire-based scoring, remain largely unknown. According to decision theory, apathy could result either from biased subjective estimates (of action costs or outcomes) or from dysfunctional processes (in making decisions or allocating resources).

METHODS: Here, we combined a series of behavioral tasks with computational modeling to elucidate the motivational deficits of 36 patients with unipolar or bipolar depression under various treatments compared with 36 matched healthy control subjects.

RESULTS: The most striking feature, which was observed independent of medication across preference tasks (likeability ratings and binary decisions), performance tasks (physical and mental effort exertion), and instrumental learning tasks (updating choices to maximize outcomes), was an elevated sensitivity to effort cost. By contrast, sensitivity to action outcomes (reward and punishment) and task-specific processes were relatively spared.

CONCLUSIONS: These results highlight effort cost as a critical dimension that might explain multiple behavioral changes in patients with depression. More generally, they validate a test battery for computational phenotyping of motivational states, which could orientate toward specific medication or rehabilitation therapy, and thereby help pave the way for more personalized medicine in psychiatry.

https://doi.org/10.1016/j.bpsc.2022.07.011

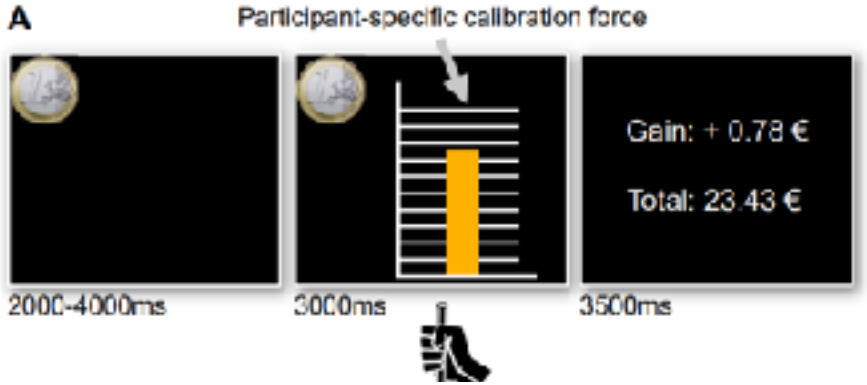

Anxiety

- **Different types**: Generalised anxiety disorder, Panic disorder, social anxiety disorder, specific phobia, agoraphobia, separation anxiety, selective mutism (OCD and PTSD no longer diagnosed under anxiety disorders).

- In UK, **6%** of people with a diagnosis of Generalised Anxiety
- More common in **young people** and **women**
- Covid made it worse: In 2022/23, in UK, an average of **37.1%** of women and **29.9%** of men reported high anxiety levels, a significant rise from 21.8% and 18.3%, respectively, between 2012 and 2015.

- Anxiety and depression are highly **comorbid**: about 50% diagnosed with one will also have the other one.

Table 1. Diagnostic Criteria for Generalized Anxiety Disorder

A. Excessive anxiety and worry (apprehensive expectation), occurring more days than not for at least 6 months, about a number of events or activities (such as work or school performance).

B. The individual finds it difficult to control the worry.

C. The anxiety and worry are associated with three (or more) of the following six symptoms (with at least some symptoms having been present for more days than not for the past 6 months):

Note: Only one item is required in children.

1. Restlessness or feeling keyed up or on edge.
2. Being easily fatigued.
3. Difficulty concentrating or mind going blank.
4. Irritability.
5. Muscle tension.
6. Sleep disturbance (difficulty falling or staying asleep, or restless, unsatisfying sleep).

D. The anxiety, worry, or physical symptoms cause clinically significant distress or impairment in social, occupational, or other important areas of functioning.

E. The disturbance is not attributable to the physiological effects of a substance (e.g., a drug of abuse, a medication) or another medical condition (e.g., hyperthyroidism).

F. The disturbance is not better explained by another mental disorder (e.g., anxiety or worry about having panic attacks in panic disorder, negative evaluation in social anxiety disorder [social phobia], contamination or other obsessions in obsessive-compulsive disorder, separation from attachment figures in separation anxiety disorder, reminders of traumatic events in posttraumatic stress disorder, gaining weight in anorexia nervosa, physical complaints in somatic symptom disorder, perceived appearance flaws in body dysmorphic disorder, having a serious illness in illness anxiety disorder, or the content of delusional beliefs in schizophrenia or delusional disorder).

Reprinted with permission from the American Psychiatric Association. Diagnostic and Statistical Manual of Mental Disorders. *5th ed. Washington, DC: American Psychiatric Association; 2013:222.*

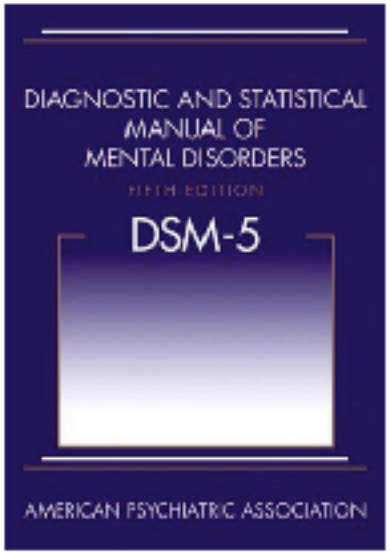

Anxiety and Reinforcement learning

- **Could anxiety be related to an overestimation of threat and prediction of negative outcomes?**
- **An inability to extinguish fear associations? An overgeneralisation of fearful associations?**
- **An inability to learn in changing environments?**

Research

JAMA Psychiatry | Original Investigation

Reinforcement Learning in Patients With Mood and Anxiety Disorders vs Control Individuals

A Systematic Review and Meta-analysis

Alexandra C. Pike, DPhil; Oliver J. Robinson, PhD

Supplemental content

IMPORTANCE Computational psychiatry studies have investigated how reinforcement learning may be different in individuals with mood and anxiety disorders compared with control individuals, but results are inconsistent.

OBJECTIVE To assess whether there are consistent differences in reinforcement-learning parameters between patients with depression or anxiety and control individuals.

DATA SOURCES Web of Knowledge, PubMed, Embase, and Google Scholar searches were performed between November 15, 2019, and December 6, 2019, and repeated on December 3, 2020, and February 23, 2021, with keywords (*reinforcement learning*) AND (*computational* OR *model*) AND (*depression* OR *anxiety* OR *mood*).

STUDY SELECTION Studies were included if they fit reinforcement-learning models to human choice data from a cognitive task with rewards or punishments, had a case-control design including participants with mood and/or anxiety disorders and healthy control individuals, and included sufficient information about all parameters in the models.

DATA EXTRACTION AND SYNTHESIS Articles were assessed for inclusion according to MOOSE guidelines. Participant-level parameters were extracted from included articles, and a

Higher sensitivity to punishment, lower sensitivity to reward, the beta of the softmax

However, the meta-analysis did not specifically differentiate between patients with depression and those with anxiety disorders

Anxiety

Anxious individuals have difficulty learning the causal statistics of aversive environments

Michael Browning[1], Timothy E Behrens[1], Gerhard Jocham[1,2], Jill X O'Reilly[1] & Sonia J Bishop[1,3]

Statistical regularities in the causal structure of the environment enable us to predict the probable outcomes of our actions. Environments differ in the extent to which action-outcome contingencies are stable or volatile. Difficulty in being able to use this information to optimally update outcome predictions might contribute to the decision-making difficulties seen in anxiety. We tested this using an aversive learning task manipulating environmental volatility. Human participants low in trait anxiety matched updating of their outcome predictions to the volatility of the current environment, as predicted by a Bayesian model. Individuals with high trait anxiety showed less ability to adjust updating of outcome expectancies between stable and volatile environments. This was linked to reduced sensitivity of the pupil dilatory response to volatility, potentially indicative of altered norepinephrinergic responsivity to changes in this aspect of environmental information.

nature neuroscience

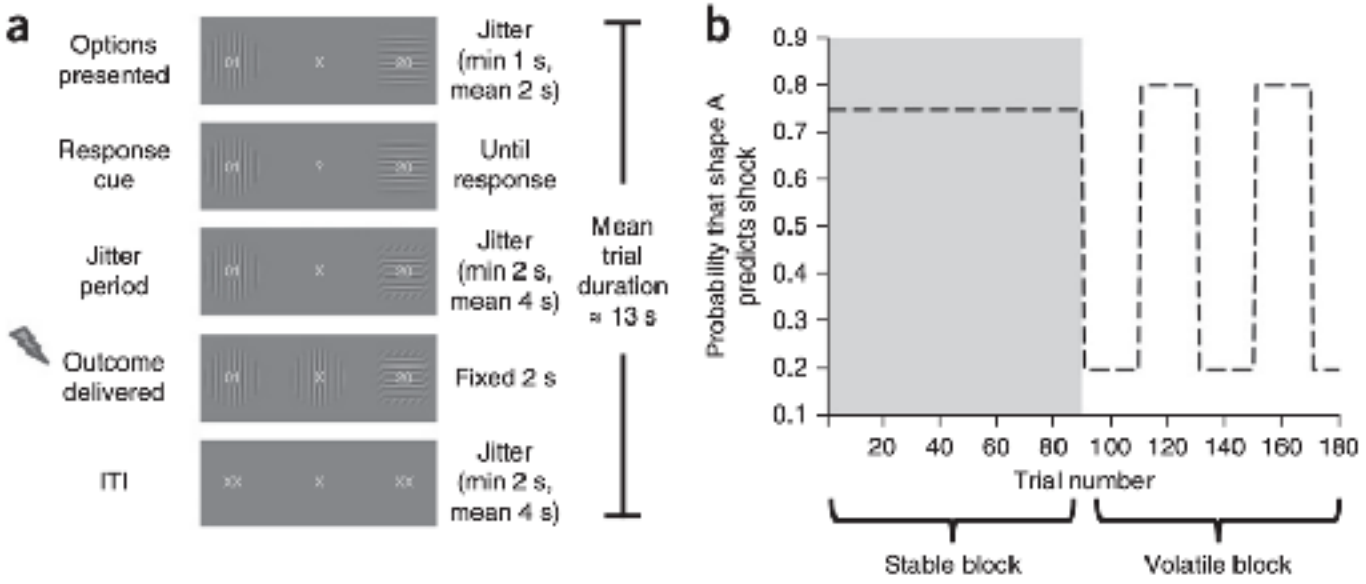

Overestimation of volatility

Participants need to predict whether they would receive an electric shock based on a given cue. Probability linking the cue to the outcome varied over time.

1. Two Conditions
    - Stable Condition: probability of receiving the shock given the cue remained relatively constant.
    - Volatile Condition: probability fluctuated frequently, requiring participants to continuously update their learning.

2. Reinforcement Learning Component:
    - Participants learn through trial and error how likely a cue was to lead to a shock.
    - Their learning was analyzed using computational models (RL and Bayesian).

Main Finding

- Low-anxiety individuals successfully adjusted their learning rates based on environmental volatility (i.e., they learned quickly when things were unpredictable and more slowly when they were stable).
- **High-anxiety individuals failed to properly adapt their learning rates**, meaning they either learned too rigidly or updated their expectations inconsistently, leading to maladaptive learning in uncertain environments.

Summary

- Psychiatric disorders viewed as deficits in learning and decision-making

- This makes RL tasks (e.g. gambling tasks) and modelling relevant: quantify differences in how values for states and actions are learned and used.

- Prominent models of addiction suggest that drug intake hijacks the learning processes (because dopamine surges interferes with the representation of prediction errors), hence leading to aberrant valuation of states leading to the drug.

- Impairments in valuation and decision-making is seen also in depression and anxiety, but not yet entirely clear how it relates to reward sensitivity, punishment sensitivity, learning rates, how values are used to guide decision-making, control, effort, estimation about volatility of the environment, model based vs model free.

- An explosion of hypothesis-driven quantitative research.

Readings

- Textbook Chapter 5 (Peter Dayan - Reinforcement Processes)
- Textbook Chapter 7 (Depression)
- Textbook Chapter 9 (David Reddish, Addiction)