Commentary

Biological Psychiatry

What is Computational Psychiatry Good For?

Michael Browning, Martin Paulus, and Quentin J.M. Huys

It is rare in the field of biological psychiatry for hypotheses to be definitively refuted. Rather, topics of investigation drift into and out of fashion, often driven by the initial excitement of technological innovation, followed by the necessary corrective of nuanced or underwhelming clinical results. A well-known example of this is the association between depression and abnormal function of the hypothalamic-pituitary-adrenal axis, as measured using the dexamethasone suppression test (DSST) (1). This observation led to a great deal of work investigating whether the association might help us identify useful subtypes of depression (2) or help us predict treatment response (3). As it turned out, the specificity and predictive value of the DSST was not thought to be of a level that would be useful clinically and the topic has gradually moved out of the spotlight. We are left in the familiar position of knowing that nonsuppression of the DSST is associated with depression, but not being sure how we can use this knowledge to help patients.

An understandable response to this situation is to try to develop better measures—if we had a more sophisticated version of the DSST, perhaps we would be able to realize our clinical aspirations. We make two suggestions in this commentary: First, that an alternative approach is to ask more precise questions, and second, that an underappreciated application of computational techniques is that they may help us to do so.

Computational Answers

Much of the ambition of computational psychiatry to date has been on the use of formal models to find hidden answers to interesting clinical questions. For example, a recent study asked whether computational modeling might help explain previously described electroencephalography and magnetic resonance imaging data from patients with schizophrenia (4). In this important study, Adams *et al.* (4) replicated the group differences between patients and control subjects using resting, mismatch negativity, and 40-Hz auditory steady-state response paradigms. They then applied a neural mass model to assess what changes at the microcircuit level might produce these effects, concluding that they could all be accounted for by reduced synaptic gain on pyramidal cells. In other words, the benefit of the modeling was that it linked electroencephalography and magnetic resonance imaging results to putative causal processes at the microcircuit level. In this example, the computational model is being used analogously to the DSST task described above—as a method of measuring a hidden but hopefully clinically important process. In the next section we consider how models may be useful in a different way: to ask, rather than answer, questions.

Computational Questions

A useful feature of computational models is that they do not just tell us what they know, but they also tell us what they do not know. This can be useful when we are deciding what questions we should answer. For example, turning to how depressed patients learn from rewarding experiences, a common finding across the literature is that, when presented with two choices, one more rewarding than the other (Figure 1A), patients with depression will select the more rewarding choice less consistently than people who are not depressed (5) (Figure 1C). One way of interpreting this finding is that reinforcement learning processes are disrupted in depressed patients. Relatively simple reinforcement learning algorithms (Figure 1B) can successfully emulate the choices of both patients and control subjects. The first type of algorithm that was used in this situation had 2 free parameters (6)—a learning rate controlling how quickly value learning occurs, and an inverse decision temperature controlling how much the model uses its learned values when selecting an action (Figure 1B–D). This algorithm attributed the behavior of depressed patients to a reduction in the inverse decision temperature. That is, patients were learning about the choices normally but were not selecting the best option because their decisions were less influenced by the values they had learned. Later, a second-generation algorithm was developed, in which the inverse decision temperature parameter was replaced by a reward sensitivity parameter, which controls how rewarding the reinforcer used in the task (e.g., points won, money) was judged to be. Using this model, the behavior of depressed patients was explained by a reduction in reward sensitivity (7). In other words, depressed patients selected the best option less frequently due to a specific reduction in how rewarding they found the reinforcer used in the task.

These 2 classes of model explain the same behavior in conceptually separate ways, as either an effect of how patients use information when they make decisions or as a consequence of reduced estimates of the reward value of the reinforcer. One way to arbitrate between these competing explanations might be to use a more sophisticated model in which both the inverse decision temperature and the reward sensitivity parameters are allowed to vary and ask which parameter differs in depressed patients relative to control subjects. The result of using this model is illustrated in Figure 1E: the model is unable to estimate reward sensitivity or inverse decision temperature as the 2 parameters are mathematically redundant (7) and therefore completely interchangeable. Thus, even though the parameters represent conceptually distinct hypotheses about the underlying cognitive processes, choice behavior on these simple tasks is not able to adjudicate between them. Here the important thing the model is telling us is that it does not know whether the choice behavior of patients is caused by changes in the inverse decision temperature and/or by changes in reward sensitivity.

https://doi.org/10.1016/j.biopsych.2022.08.030
ISSN: 0006-3223

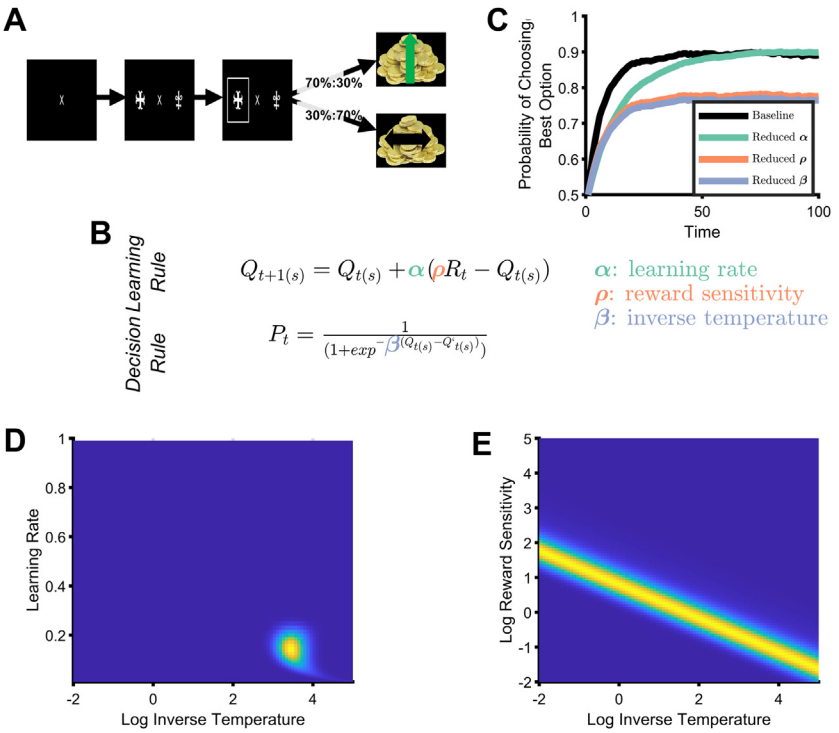

Figure 1. What computational models do not know. **(A)** A frequently used learning task in which participants must choose between 2 options, one of which is associated with a higher probability of reward than the other. A commonly observed finding is that depressed patients select the most rewarded option less consistently than nondepressed participants. **(B)** A simple computational model describes how participants may solve this task—Q_t, the value of an option, is learned using a simple updating process and is fed into a decision rule. The behavior of the model can be influenced by several different parameters, including learning rate (α), reward sensitivity (ρ), and inverse decision temperature (β). **(C)** An illustration of the effect of changing these parameters on choice in the task. The learning rate influences the rate at which the model reaches a behavioral plateau, and both the reward sensitivity and inverse decision temperature control the level of the plateau. The behavior of depressed patients can be captured by either reducing reward sensitivity or inverse decision temperature. **(D)** The posterior estimate of parameter values after fitting to a participant's choices, using a model in which learning rate and inverse decision temperature are allowed to vary (but reward sensitivity is fixed). Possible values of the inverse temperature parameter are represented along the x axis; possible values of learning rate are represented on the y axis. As can be seen, the parameters are precisely estimated, with learning rate lying between 0.1 and 0.2 and the log inverse temperature between 3.5 and 4. Yellow corresponds to the highest posterior probability, i.e., the most likely parameter value given the observed behavior. **(E)** The marginal posterior estimate of reward sensitivity and inverse temperature when all 3 parameters are allowed to vary and the model is fitted to the same data as **(D)** (for simplicity, the marginal probability of learning rate is not shown). Because the model is unable to discriminate between reward sensitivity and inverse decision temperature, it is unable to estimate either—any value of inverse temperature or reward sensitivity is possible. In other words, the model does not know whether participant choices are influenced by changes in reward sensitivity or decision temperature.

In this case, the model is useful not because it tells us what the answer is but because it tells us precisely what the question should be: Is reduced reward choice in depression caused by difficulty in using learned values when making decisions or by a lower value of reinforcers? Beyond this, the model suggests how these questions might be answered; if depression is associated with difficulty in using the value of options to make decisions, then the choices of depressed patients should be less consistent even when values of options are explicitly presented, and no learning is required [although Rutledge *et al.* (8) report evidence against this prediction]. Alternatively, if patients treat reinforcers as less rewarding, then this effect should be apparent even when no decisions are required [e.g., where response to a single rewarded stimuli is measured (9)]. To date, the literature is not consistent with a simple decision effect but is generally consistent with reduction in reward sensitivity as the most likely process associated with depression, raising interesting subsidiary questions about why this might occur (5).

In summary, computational models can be used to identify hidden processes, some of which might be useful for answering clinically interesting questions. But models also tell us when they are unable to discriminate between competing explanations, and when this occurs they are a useful way of framing the precise mechanistic questions we should be trying to answer to improve our measures.

The development of clinically useful measures from biological research requires us to ask questions that are sufficiently specific that they may be refuted. Computational models help us ask these questions.

Acknowledgments and Disclosures

MB is supported by the Oxford Health NIHR Biomedical Research Centre. QJMH is supported by the UCLH Biomedical Research Centre. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR, or the Department of Health. MP has received grants from the William K. Warren Foundation and National Institutes of Health Grant Nos. R01DA016663, P20DA027834, R01DA027797, R01DA018307, U01DA041089, and 1R01MH101453.

MB has received travel expenses from Lundbeck for attending conferences and has acted as a consultant for J&J, Novartis, P1vital Ltd., and CHDR and owns shares in P1vtial Products Ltd. QJMH has received options and consultancy fees from Aya Health and Alto Neuroscience and a research grant from Koa Health. MP has received royalties for an article about methamphetamine use disorder from UpToDate.

Article Information

From the Department of Psychiatry (MB), University of Oxford, and the Oxford Health NHS Trust (MB), Oxford; Division of Psychiatry (QJMH) and the Max Planck UCL Centre for Computational Psychiatry and Ageing Research and Wellcome Centre for Human Neuroimaging (QJMH), Queen Square Institute of Neurology, University College London, and the Camden and Islington NHS Foundation Trust (QJMH), London, United Kingdom; and

Laureate Institute for Brain Research (MP), University of Tulsa, Tulsa, Oklahoma.

Address correspondence to Michael Browning, D.Phil., M.B.B.S., M.R.C.P., M.R.C.Psych., at michael.browning@psych.ox.ac.uk.

Received Aug 30, 2022; accepted Aug 31, 2022.

References

1. Carroll BJ, Martin FI, Davies B (1968): Resistance to suppression by dexamethasone of plasma 11-O.H.C.S. levels in severe depressive illness. Br Med J 3:285–287.
2. Schatzberg AF, Rothschild AJ, Stahl JB, Bond TC, Rosenbaum AH, Lofgren SB, *et al.* (1983): The dexamethasone suppression test: Identification of subtypes of depression. Am J Psychiatry 140:88–91.
3. Rush AJ, Weissenburger J, Vasavada N, Orsulak PJ, Fairchild CJ (1988): Dexamethasone suppression test status does not predict differential response to nortriptyline versus amitriptyline. J Clin Psychopharmacol 8:421–425.
4. Adams RA, Pinotsis D, Tsirlis K, Unruh L, Mahajan A, Horas AM, *et al.* (2022): Computational modeling of electroencephalography and functional magnetic resonance imaging paradigms indicates a consistent loss of pyramidal cell synaptic gain in schizophrenia. Biol Psychiatry 91:202–215.
5. Huys QJM, Browning M (2022): A computational view on the nature of reward and value in anhedonia. Curr Top Behav Neurosci 58:421–441.
6. Kunisato Y, Okamoto Y, Ueda K, Onoda K, Okada G, Yoshimura S, *et al.* (2012): Effects of depression on reward-based decision making and variability of action in probabilistic learning. J Behav Ther Exp Psychiatry 43:1088–1094.
7. Huys QJM, Pizzagalli DA, Bogdan R, Dayan P (2013): Mapping anhedonia onto reinforcement learning: A behavioural meta-analysis. Biol Mood Anxiety Disord 3:12.
8. Rutledge RB, Moutoussis M, Smittenaar P, Zeidman P, Taylor T, Hrynkiewicz L, *et al.* (2017): Association of neural and emotional impacts of reward prediction errors with major depression. JAMA Psychiatry 74:790–797.
9. Knutson B, Bhanji JP, Cooney RE, Atlas LY, Gotlib IH (2008): Neural responses to monetary incentives in major depression. Biol Psychiatry 63:686–692.