



HHS Public Access

Author manuscript

Nat Neurosci. Author manuscript; available in PMC 2015 April 23.

Published in final edited form as:

Nat Neurosci. 2011 February ; 14(2): 154–162. doi:10.1038/nn.2723.

FROM REINFORCEMENT LEARNING MODELS OF THE BASAL GANGLIA TO THE PATHOPHYSIOLOGY OF PSYCHIATRIC AND NEUROLOGICAL DISORDERS

Tiago V. Maia¹ and Michael J. Frank²

¹Department of Psychiatry, Columbia University, and New York State Psychiatric Institute, New York, New York, USA

²Department of Cognitive, Linguistic & Psychological Sciences, Department of Psychiatry and Human Behavior, and Brown Institute for Brain Science, Brown University, Providence, Rhode Island, USA

Abstract

Over the last decade and a half, reinforcement learning models have fostered an increasingly sophisticated understanding of the functions of dopamine and cortico-basal ganglia-thalamo-cortical (CBGTC) circuits. More recently, these models, and the insights that they afford, have started to be used to understand key aspects of several psychiatric and neurological disorders that involve disturbances of the dopaminergic system and CBGTC circuits. We review this approach and its existing and potential applications to Parkinson's disease, Tourette's syndrome, attention-deficit/hyperactivity disorder, addiction, schizophrenia, and preclinical animal models used to screen novel antipsychotic drugs. The approach's proven explanatory and predictive power bodes well for the continued growth of computational psychiatry and computational neurology.

INTRODUCTION

The limitations of the state-of-the-art in nosology in psychiatry have been much debated in the context of the development of the new edition of the *Diagnostic and Statistical Manual of Mental Disorders*^{1, 2}. There is widespread agreement that the current symptom-based system of classification must eventually be replaced with a system based on pathophysiology³. However, the current understanding of the neurobiology and genetics of psychiatric disorders remains too limited to form the backbone of nosology¹. This limited understanding is also reflected in the state-of-the-art in treatment, with most psychiatric medications having been found by serendipity, rather than through rational design. Neurology typically deals with disorders with better understood etiology (e.g., loss of dopaminergic neurons in Parkinson's disease), but even then, it is often unclear how these etiological processes produce complex patterns of symptoms, and why treatments can alleviate some deficits while exacerbating, or even causing, others^{4, 5}. Part of the problem is the complexity of the brain and mind, and the many levels of analysis that span the two. Computational models are a valuable tool to tame this complexity, because they foster a

mechanistic understanding that can span multiple levels of analysis and that can explain how changes to one component of the system (e.g., increases in striatal D2 receptor density) produce systems-level changes that translate to changes in behavior.

One area in which substantial progress has been made in integrating computational modeling and empirical research in neuroscience is that of reinforcement learning (RL)⁶. This approach has produced explicit models of the roles of dopamine and cortico-basal ganglia-thalamo-cortical (CBGTC) loops in learning about reinforcers (rewards and punishments) and in guiding behavior so as to acquire rewards and avoid punishments^{6, 7}. Existing models address a variety of functions of these circuits, including Pavlovian conditioning, instrumental conditioning, and their interactions; habits, goal-directed actions, and their interactions; and the inter-related issues of incentive salience, motivation, and vigor^{6, 8–11}.

Organizing behavior so as to obtain outcomes appropriate for the current motivational state (e.g., acquiring food if hungry) and so as to avoid harmful outcomes is crucial for survival and is therefore a central organizing principle of the nervous system. Not surprisingly, then, disturbances of the dopaminergic system and CBGTC circuits play a key role in several psychiatric and neurological disorders. RL models have recently started to be applied to these disorders, and they have been shown to have significant explanatory and predictive power^{12, 17}. The approach builds on an understanding of the computations that these circuits perform in healthy individuals, and investigates how pathophysiological processes alter these computations, producing symptoms. We therefore start by reviewing the computational neurobiology of the normal functioning of these circuits. We then discuss several disorders that have benefited or are ripe to benefit from the use of RL models. We close by discussing the future implications of this body of work for nosology and treatment.

In addition to conveying the specifics of how RL models provide insights into psychiatric and neurological disorders, we hope that this review will also help foster the development of the emerging disciplines of computational psychiatry and computational neurology. A powerful set of computational techniques can now be used to investigate pathophysiological processes and their relation to behavior (Figure 1), and these techniques are broadly applicable across model types and disorders. Some of these techniques have been used for many years; others are more recent. The concerted and widespread use of the full complement of these techniques, however, will likely flourish in the coming years, especially given the increasing sophistication in computational models of a variety of neurocognitive functions that are disturbed in several psychiatric and neurological disorders.

REINFORCEMENT LEARNING IN THE BRAIN

Dopamine and Prediction Errors

Dopamine neurons code reinforcement *prediction errors*^{6, 18}, a key signal in many RL models¹⁹. Prediction errors signal the difference between the observed and expected outcome: a positive prediction error signals that the outcome was better than expected, and a negative prediction error signals that the outcome was worse than expected. The magnitude of phasic dopamine-neuron bursts quantitatively represents positive prediction errors²⁰.

Whether phasic reductions in dopaminergic neuron firing quantitatively represent negative prediction errors is more controversial, because the low tonic firing rate of dopamine neurons implies that variations in such reductions are somewhat limited. The duration of the pause in dopaminergic neuron firing, however, seems to represent negative prediction errors quantitatively²¹. This asymmetric coding between positive and negative prediction errors – burst magnitude for positive and pause duration for negative – may be justified biologically, both because it permits lower tonic firing, which is advantageous metabolically, and in terms of the postsynaptic effects of these signals on D1 and D2 receptors⁷, as discussed below. An alternative or complementary hypothesis is that negative prediction errors may be coded by serotonin²². This hypothesis, however, remains to be adequately tested⁶.

Although the majority of dopamine neurons burst to positive prediction errors, smaller proportions of these neurons also burst under other conditions²³. For example, some dopamine neurons burst not only to positive events and stimuli that predict positive events, but also to negative events and stimuli that predict negative events^{24, 25}. There have been some attempts at reconciling these findings with RL^{26, 27}, but this is an area in need of additional research. Nonetheless, the predominant function of dopamine bursts is to code positive prediction errors²³. Supporting this notion, phasic optogenetic stimulation of dopamine neurons induces a subsequent preference for the place in which such stimulation occurred²⁸, just as if reward had been delivered at that place.

The Basal Ganglia and Action Selection

Prediction errors are used to learn the values (or ‘goodness’) of states (stimuli or situations) and/or of state-action pairs. These values are then used to select optimal actions¹⁹. A model of this process that has been used to account for many behavioral and neural findings^{6, 29, 30} is the actor-critic³¹. The actor-critic view of action selection in the brain suggests that the cortex represents the current state and the basal ganglia (BG) implement two computational modules: the critic, which learns state values and may be implemented in the ventral striatum^{6, 30} (possibly together with the amygdala and orbitofrontal cortex⁶); and the actor, which learns stimulus-response (S-R) associations and may be implemented in the dorsal striatum^{6, 30}. The critic and the actor both use the prediction errors signaled by dopamine to update their estimates (of state values and of S-R strengths, respectively). The mapping of the actor-critic to the BG is consistent with the view that the BG play a key role in S-R learning³², but with different portions of the striatum – which are involved in distinct, parallel CBGTC loops³³ – playing distinct roles³⁴.

A related but slightly different view suggests that the BG are responsible not for generating actions, but for arbitrating between multiple actions that may be under consideration in cortex, by facilitating the most appropriate action while suppressing competing actions^{35–37}. In other words, whereas according to the actor-critic view, the BG are fully responsible for selecting the action on the basis of the current state alone, this alternative view suggests that the cortex itself initially generates candidate actions (e.g., on the basis of the frequency with which they have previously been executed in the current state³⁷), and the BG then arbitrate between these actions (likely on the basis of their learned reinforcement probabilities^{37, 38}), to facilitate (‘gate’) the best one. The commonalities between this view and the actor-critic

one substantially outweigh their differences, however, so we will not delve into this distinction further.

The BG anatomy consists of direct, indirect, and hyperdirect pathways from cortex to BG output structures^{39–41} (Figure 2a). Neurocomputational models^{37, 42–46} have refined verbal theories^{35, 39–41} of the role of these pathways in action selection. An influential account, originally proposed to explain the pathophysiology of several neurological disorders^{39, 40}, suggests that the direct pathway provides focused facilitation of the appropriate action(s) for the current state, whereas the indirect pathway suppresses actions that are inappropriate for that state. Although the original version of this account was based on a somewhat oversimplified view of BG anatomy and function, recent findings strongly support this distribution of function between the direct (or ‘Go’) and indirect (or ‘NoGo’) pathways^{47–49}. This distribution of function finds formal expression in the basal ganglia Go/NoGo (BG-GNG) model³⁷ (Figure 2b), in which the probability that a given action is selected is proportional to the difference between the Go and NoGo activity for that action in the current state. Consistent with this scheme, electrophysiological findings demonstrate that the positive and negative values of actions are represented in distinct striatal populations, with greater activity in the neurons that represent the positive value of an action predicting selection of that action, and greater activity in the neurons that represent the negative value of an action predicting selection of an alternative action³⁸. The BG-GNG model further shows how aspects of BG anatomy not considered in the original account subserve other aspects of BG function⁷. For example, the model suggests that the hyperdirect (or ‘Global NoGo’) pathway provides global inhibition of all actions during the early stages of processing, particularly in situations of high conflict (i.e., when multiple actions are strongly activated simultaneously), to prevent premature, suboptimal responding⁴⁵.

The BG-GNG model also shows how the direct and indirect pathways can learn which actions to facilitate and suppress in each state, respectively, using the prediction errors conveyed by dopamine³⁷. The direct and indirect pathways predominantly express D1 and D2 receptors, respectively⁵⁰. In the model, when an action is followed by a dopamine burst, the corticostriatal synapses in the direct pathway into active Go units (which encode the conjunction between the state and the action that was executed) are strengthened via D1-dependent long-term potentiation (LTP), and the corticostriatal synapses in the indirect pathway into active NoGo units for the executed action are weakened via D2-dependent long-term depression (LTD). When an action is followed by a dopamine dip, the reverse occurs. These dual effects of dopamine on D1- and D2-mediated plasticity have been supported by empirical evidence⁵¹, as has the model prediction that the direct and indirect pathways mediate learning from positive and negative outcomes, respectively⁴⁸.

The dynamics of dopamine effects on D1 and D2 receptors may also explain why burst magnitude and pause duration code for positive and negative prediction errors, respectively⁷. D1 and D2 receptors have relatively low and high affinity for dopamine, respectively⁵². D1 stimulation is therefore hypothesized to depend on phasic dopamine bursts, with larger bursts producing greater stimulation. Burst magnitude therefore plays a key role in D1-mediated LTP due to positive prediction errors. D2 receptors, in contrast, are hypothesized to be stimulated tonically by baseline dopamine levels. The effect of pauses in dopaminergic

neuron firing on D2 receptors therefore depends on dopamine reuptake, with longer pauses allowing greater reuptake, and therefore a larger dip in dopamine concentration. Pause duration therefore plays a key role in D2-mediated LTP due to negative prediction errors. LTD mechanisms are also consistent with a key role for magnitudes and durations in coding positive and negative prediction errors, respectively. Positive prediction errors stimulate D2 receptors directly, producing LTD in the indirect pathway. Negative prediction errors may not strongly affect D1 receptors, because D1 receptors may not be significantly stimulated by tonic dopamine. In the BG-GNG model, negative prediction errors instead produce LTD in the direct pathway indirectly, via their effects on D2 receptors (dependent on reuptake and therefore pause duration), and subsequent inhibition of the direct pathway by the indirect pathway, leading to activity-dependent LTD in the direct pathway.

In addition to the role of phasic dopamine in learning, tonic dopamine increases excitability in the direct (Go) pathway and decreases excitability in the indirect (NoGo) pathway, because D1 receptors are excitatory (at least for neurons receiving strong concomitant glutamatergic input⁵³) and D2 receptors are inhibitory⁵⁴. Increases in tonic dopamine therefore produce a Go bias, whereas decreases produce a NoGo bias. Simulations using the BG-GNG model show that these biases exert strong effects on action selection and reaction times^{12, 37, 55}, with tonic dopamine promoting the execution and speed of actions (particularly actions with greater positive differences between their previously learned Go and NoGo associations). Dopamine therefore modulates not only learning, but also the expression of prior learning.

Although multiple parallel (albeit interacting⁵⁶) CBGTC loops course through the BG³³, a common division is into sensorimotor, associative, and limbic loops, which connect to, respectively, sensorimotor cortical areas, dorsolateral prefrontal cortex (DLPFC), and orbitofrontal cortex (OFC)/anterior cingulate cortex (ACC)⁵⁷. The BG-GNG model has been applied not only to motor action selection, involving the sensorimotor loop, but also to the selection of cognitive ‘actions,’ particularly working memory updating^{58, 59}, likely involving the associative loop. The idea in the latter case is that Go signals facilitate the gating of a stimulus into working memory, whereas NoGo signals prevent such gating (because, for example, the stimulus is not relevant for the task). Consistent with this idea, BG damage interferes with the ability to selectively gate only task-relevant stimuli into working memory⁶⁰.

The limbic loop, which involves the ventral striatum, may implement the critic, as noted above, learning the values of states. It is tempting to speculate that the direct and indirect pathways learn the positive and negative values of states, respectively. Consistent with this idea, in a conditioned place preference task—which involves learning the values of places (states) and depends on the nucleus accumbens⁶¹ – learning a positive value for a place depended on neurotransmission within the direct, but not the indirect, pathway⁴⁸. Furthermore, cocaine-induced conditioned place preference is increased by optogenetic activation of the direct pathway in the nucleus accumbens during learning, and it is decreased by activation of the indirect pathway⁶². The OFC has been hypothesized to implement working memory for state reinforcement values, helping to guide action selection when these values change rapidly, among other scenarios⁶³. If so, the limbic loop may also

implement Go and NoGo for ‘actions’ that determine when the values maintained in OFC should be updated so as to ensure adaptive, flexible behavior.

CLINICAL APPLICATIONS

The ideas and models discussed above, when applied to both motor and cognitive domains, can explain a variety of findings across several disorders. Given dopamine’s central role in RL, we focus on disorders with strong dopaminergic involvement: Parkinson’s disease (PD), Tourette’s syndrome (TS), attention-deficit/hyperactivity disorder (ADHD), drug addiction, and schizophrenia. Of course these disorders often also involve non-dopaminergic disturbances. For example, TS involves abnormalities in striatal interneurons^{64, 65}, ADHD involves noradrenergic abnormalities⁶⁶, etc. Furthermore, in some cases, the dopaminergic disturbances themselves may be caused by abnormalities in upstream processes. For example, schizophrenia involves increased mesolimbic and decreased mesocortical dopamine^{67, 68}, both of which may be caused, at least in part, by dysregulated cortical control of dopamine neurons due to NMDA-receptor abnormalities⁶⁹. (NMDA abnormalities may in fact directly contribute to many of the symptoms of schizophrenia by disrupting the stability of cortical attractors⁷⁰.) Comprehensive models of the pathophysiologies of these disorders and their relation to symptoms will ultimately have to integrate these various abnormalities. As we will see, however, even a more limited focus on the small number of principles articulated above concerning the computational functions of dopamine and CBGTC circuits provides substantial leverage to understand multiple aspects of these disorders.

Parkinson’s Disease

Dopaminergic cell death in PD results in reduced striatal dopamine, thereby producing an exaggerated tendency for NoGo^{12, 37, 71}. Simulations using the BG-GNG model demonstrate that this tendency explains not only the hypokinetic symptoms of PD but also a variety of cognitive deficits that accompany this disorder^{12, 37, 72}. For example, as noted above, the model suggests that Go and NoGo signals in the associative loop respectively facilitate and prevent working memory updating^{58, 59}. The hyperexcitable NoGo pathway in PD should therefore produce a deficit in working memory updating, while simultaneously producing increased resistance to distractors. Furthermore, 1-dopa and dopamine agonists should reverse these effects. Empirical studies confirmed these predictions^{73, 74}.

Subjects’ tendency to learn better from positive or from negative feedback (Go and NoGo learning, respectively) can be assessed using the probabilistic selection task⁷² (Figure 3). Healthy controls are equally good at learning to obtain positive outcomes (Go learning) and avoid negative outcomes (NoGo learning) in this task^{5, 72} (although there are individual differences in Go versus NoGo learning among healthy controls, which are predicted by genetic variations affecting dopamine function in the direct and indirect pathways^{47, 75}). Consistent with the predictions of the BG-GNG model, unmedicated Parkinson’s patients are better at NoGo learning than at Go learning⁷². Medication reverses these biases: medicated patients are better at Go than at NoGo learning, and are worse at NoGo learning than unmedicated patients or controls^{5, 72}. These medication effects were also predicted by the model, under the assumption that dopaminergic medications reduce dopamine dips

during negative prediction errors (because dopaminergic medications result in continued occupation of postsynaptic dopamine receptors during pauses in firing of dopamine neurons)⁷². Such blunting of negative prediction errors reduces learning from negative outcomes, producing the deficit in NoGo learning. Similar findings for both unmedicated and medicated patients have been obtained using other tasks that also assess the degree to which subjects learn more from positive or negative outcomes⁷⁶⁻⁷⁹. These asymmetries in learning from positive and negative outcomes may also have clinical implications. For example, the medication-induced tendency to learn more from positive than from negative outcomes may help explain why medication induces pathological gambling in a subset of Parkinson's patients⁸⁰.

The BG-GNG model also suggests that the subthalamic nucleus (STN) - the key node in the hyperdirect pathway - provides a 'Global NoGo' signal that transiently inhibits all actions during action selection⁴⁵. Activation of the STN, and therefore the Global NoGo signal, is dynamically modulated by the amount of response conflict (see also ref. 43). The Global NoGo signal therefore plays an especially important role in situations of high conflict, in which actions with relatively similar reinforcement histories are being considered. In such situations, the Global NoGo inhibition provides time for the best action to win the competition, preventing premature, suboptimal actions from being facilitated. The model therefore predicted that disruption of STN processing - e.g., by deep brain stimulation (DBS) - would disrupt subjects' ability to slow down in such high-conflict situations, resulting in faster but suboptimal responses⁵. This prediction was confirmed experimentally in Parkinson's patients undergoing DBS of the STN (and the deficit resolved when DBS was turned off)⁵. Furthermore, as predicted by the model, dopaminergic medications and DBS had doubly dissociable effects: medications affected the asymmetry in learning from positive and negative outcomes but not the ability to slow down in high-conflict situations, and DBS had the opposite effects⁵. The adverse effects of medications and DBS on real-life behavior in a subset of Parkinson's patients may therefore be due to distinct mechanisms.

Tourette's Syndrome

TS is characterized by recurrent, stereotyped movements and vocalizations, known as tics. Tics have been hypothesized to reflect abnormal activation of subsets of striatal neurons that provide Go signals for the tic⁸¹. Evidence from clinical pharmacology in TS and from experimental work in animals suggests that tics may be due to excessive excitability and/or plasticity in the direct (Go) relative to the indirect (NoGo) pathway. First, D2 blockers, the standard pharmacological treatment for TS, boost the indirect pathway (because the D2 receptor is inhibitory). Second, administration of dopamine, amphetamine, or a combination of D1 and D2 agonists into the striatum - all of which simultaneously boost the direct pathway (via D1 receptors) and inhibit the indirect pathway (via D2 receptors) - causes stereotypies in animals^{82, 83}. Several findings suggest that the effects on the two pathways work synergistically to induce stereotypies; for example, stereotypies induced by striatal amphetamine administration are reduced by pretreatment with either D1 or D2 antagonists⁸².

Excessive Go relative to NoGo activity in TS could be a consequence of excessive dopamine or dopaminergic receptor sensitivity in the striatum. Indeed, TS has been associated with increases in dopamine release, dopaminergic innervation, and D2 receptors in the striatum (although the evidence for the latter two alterations is somewhat inconsistent across studies)^{68, 84}. All of these alterations would result in a boosted Go relative to NoGo pathway. Consistent with a possible excess of striatal dopamine, unmedicated TS patients learn better from rewards than from punishments⁷⁸. These biases are the opposite of those shown by unmedicated Parkinson's patients^{5, 72, 78}, mirroring the fact that the symptoms of these two disorders are also, to a limited extent, the opposite: PD is a hypokinetic disease, whereas TS is a hyperkinetic disease. In fact, medication inverts these learning biases, making Parkinson's patients medicated with l-dopa and dopamine agonists perform like unmedicated TS patients (learning better from rewards than from punishments), and making TS patients medicated with D2 blockers perform like unmedicated Parkinson's patients (learning better from punishments than from rewards)⁷⁸. The mechanistic explanation for the former finding has already been discussed; the latter finding is consistent with the results of simulations showing that D2 blockade increases excitability and plasticity of the indirect pathway, thereby promoting NoGo learning¹². This enhancement of NoGo learning by D2 blockade suggests that acute administration of D2 antagonists may be an effective adjunct for behavioral therapies that work by assigning negative value to tics (e.g., contingency management, in which tics are followed by punishment or the absence of tics is positively reinforced, or massed negative practice, in which tics become aversive due to fatigue⁸⁵). Whether acute D2 blockade would also be useful as an adjunct to habit reversal training - the best current behavioral treatment for TS - is unclear, because this procedure does not obviously involve aversive learning.

Excessive Go relative to NoGo activity in the motor CBGTC loop may also explain the premonitory urges that are a prominent feature of TS⁸⁶. These urges are hypothesized to be caused by abnormal activation in the supplementary motor area (SMA), because electrical stimulation of the SMA causes similar urges⁸⁷. Consistent with this idea, the SMA is active prior to tics⁸⁸, and SMA activation is greater with tics than with movements that mimic tics (and that are visually indistinguishable from tics) performed by healthy controls (Wang et al., submitted manuscript). The SMA is the primary target of the motor CBGTC loop³³, so abnormal SMA activation could be a consequence of excessive relative Go activity in that loop. Alternatively, or additionally, the abnormal SMA activation could be driven by the corticocortical projections between the state and the SMA (or preSMA), as in the BG-GNG model (see Figure 2b). In the model, these connections are learned via Hebbian mechanisms, so repeated gating of a tic by the BG in one or multiple states (initially due to excessive Go relative to NoGo activity) would strengthen the connections between those states and that tic's motor plan in the SMA. That SMA motor plan would then tend to become activated in those states, producing the urge (which could then be gated by the BG into an actual tic emission, but could also be prevented from doing so). This account, if correct, would explain important clinical features of TS. For example, it would explain the state-dependency of tics (i.e., the fact that tics do not occur equally frequently in all contexts). It would also explain why treatments that prevent tic performance, such as habit reversal training, over time result in reduction of the urges: repeated activation of the tic-eliciting

states without corresponding tic emission would produce Hebbian unlearning of the association between those states and the tics. More broadly, this account suggests that TS involves a vicious cycle: performing a tic strengthens the urges to perform that tic (through corticocortical Hebbian learning), which in turn increases the tendency to tic (through corticocortical activation of the tic motor plan in the SMA, which increases the likelihood of BG gating of that tic).

Attention-Deficit/Hyperactivity Disorder

ADHD is characterized by abnormal levels of inattention, hyperactivity, and/or impulsivity. The classical theory of ADHD is that it results from a primary deficit in inhibitory control, which causes several deficits in executive function⁸⁹. Another prominent theory is that ADHD results from excessive discounting of delayed rewards⁹⁰. ‘Multiple-pathway’ accounts suggest that executive dysfunction and excessive delay discounting both play a role⁹¹. ADHD seems to involve a hypofunctioning dopaminergic system^{13, 92}.

One RL theory suggests that tonic dopamine in the ventral striatum determines the discount factor (the degree to which future reinforcers are discounted relative to immediate ones) in an RL system that can look ahead because it includes an internal model of the environment⁹³. Reduced tonic dopamine in the ventral striatum in ADHD would produce a smaller discount factor, causing excessive discounting of delayed rewards⁹³. This idea seems consistent with some circumstantial evidence: Systemic administration of dopamine blockers increases delay discounting^{94, 95}, and increasing dopamine via administration of psychostimulants or selective dopamine reuptake inhibitors generally (though not always) decreases delay discounting^{94–97}. Given the lesion studies implicating the nucleus accumbens core in delay discounting⁹⁶, it seemed reasonable to hypothesize that the effects of the systemic dopaminergic manipulations might be mediated by the nucleus accumbens. However, dopamine depletion in the nucleus accumbens does not seem to produce excessive delay discounting⁹⁸, whereas dopamine depletion in the OFC does⁹⁹. Excessive delay discounting in ADHD may therefore be caused by low dopamine in the OFC.

Biophysically realistic neurocomputational models suggest that dopamine stabilizes representations in PFC¹⁰⁰. Given the key role of top-down biases from PFC in attention, cognitive control (including inhibitory control), and working memory^{101, 102}, the hypothesized low PFC dopamine in ADHD could underlie deficits in all of these executive functions. (This idea contrasts with the idea that the primary problem in ADHD is with inhibitory control, with problems in the remaining areas being secondary⁸⁹.) Cognitive deficits in ADHD need not be caused solely by low PFC dopamine, however: low striatal dopamine in the associative loops may cause, for example, reduced gating of working memory¹³. Conversely, PFC dysfunction need not cause solely executive dysfunction: the lateral PFC appears to be involved in the ability to choose delayed rewards^{103, 104} (as is the OFC), so a dysfunctional PFC could also contribute to excessive delay discounting. The relation between neurobiological abnormalities and cognitive and motivational deficits may therefore not be one-to-one, and may vary across patients.

Drug Addiction

The key role of fast, phasic-like changes in striatal dopamine in the reinforcing effects of drugs¹⁰⁵ makes addiction a natural candidate for RL modeling. An influential RL theory suggests that these fast increases function as positive prediction errors that occur every time the drug is received¹⁶. This effect of drugs contrasts with the effect of natural rewards, for which the prediction error becomes zero after the reward is expected. The recurring drug-induced positive prediction errors produce a boundless increase in the values of states (and/or actions) that lead to drug receipt, prompting compulsive drug use. This theory explains important features of addiction¹⁶, but a subsequent study disconfirmed one of its key predictions. The theory predicted that when drugs are used as the unconditioned stimulus (US), blocking should not occur. Blocking is a paradigm in which first a stimulus A is paired with a US, and subsequently simultaneous presentations of A and another stimulus B are paired with the US. Usually, no learning occurs for B¹⁰⁶, because A already predicts the US, so there is no prediction error to support new learning. If, however, drugs always produce a positive prediction error, then subjects should learn to associate B with the US - but they do not¹⁰⁷. This finding prompted a search for alternative RL accounts of addiction^{17, 108}.

One approach¹⁰⁸ moved from standard TD learning to average-reward RL¹⁰⁹. In average-reward RL, reinforcements are evaluated relative to an average reinforcement value R_t calculated using a slowly changing weighted average of past reinforcements. The 'effective' reinforcement at time t is therefore $r_t - R_t$, where r_t is the received reinforcement. Part of the motivation for this approach to addiction was to capture the decrease in sensitivity to natural rewards that long-term drug use induces. Intuitively, if drugs are extremely reinforcing, long-term drug use inflates R_t , making natural rewards less reinforcing. The model, however, added further to this effect by artificially inflating R_t even more with each drug use. The model showed decreased sensitivity to natural rewards following long-term drug use, and also showed blocking and other relevant effects¹⁰⁸.

The search for simple, single-factor RL theories of addiction will undoubtedly continue, but multiple aspects of RL likely play a role in addiction^{17, 110}. For example, chronic drug use induces functional and structural changes in several RL brain regions, such as the OFC and ventral striatum^{111, 112}, thereby further dysregulating RL processes and potentially contributing to the maintenance or aggravation of addiction. As another example, optogenetic findings in mice demonstrate that direct and indirect pathway stimulation during drug administration respectively increases and decreases the reinforcing effects of the drug⁶², suggesting that reduced indirect relative to direct pathway activity could be a risk factor for addiction. In fact, reduced indirect pathway activity would also explain the reduced sensitivity to negative outcomes that characterizes addiction.

Schizophrenia

Schizophrenia is characterized by positive symptoms (e.g., delusions and hallucinations), negative symptoms (e.g., anhedonia and avolition), and cognitive symptoms (e.g., disturbances in attention and cognitive control). Schizophrenia involves excessive dopamine and D2 receptors in the striatum, but reduced dopamine in PFC^{67, 68}.

One theory, based on the idea that dopamine signals incentive salience¹¹³, suggests that dysregulated dopaminergic firing in schizophrenia imbues percepts, thoughts, and memories with abnormal salience, and that such abnormal salience experiences underlie delusions and hallucinations¹¹⁴. Another theory suggests that psychosis results from abnormal prediction errors that produce inappropriate associations, causal attributions, and attentional salience¹¹⁵. Patients with psychosis do exhibit abnormal neural activity during prediction errors^{116, 117}, but a causal relation between these abnormalities and psychosis remains hypothetical.

Negative symptoms might conceivably reflect reduced reward sensitivity, but the evidence for this is mixed. Consistent with this hypothesis, schizophrenia patients exhibit reduced neural responses to positive prediction errors, with weaker putamen responses associated with greater avolition¹¹⁸. Schizophrenia patients also exhibit reduced Go learning^{15, 55, 119}, but these deficits do not seem to correlate with negative symptoms^{55, 119}. In these studies, negative symptoms were instead associated with indicators of PFC dysfunction^{15, 55}. Anhedonia, in particular, was associated with reduced uncertainty-driven exploration (in which alternative actions are explored in proportion to the uncertainty about their reinforcement statistics relative to the uncertainty about the reinforcement statistics of the currently preferred action)¹¹⁹. This may reflect the fact that the anhedonia assessment used partly relies on the frequency with which patients engage in pleasurable activities, which may depend on strategic processes such as exploration.

Reduced striatal responses to positive prediction errors¹¹⁸ and reduced Go learning^{15, 55, 119} in schizophrenia are suggestive of reduced phasic striatal dopamine. Increased tonic striatal dopamine, as in schizophrenia⁶⁷, may reduce phasic dopamine via inhibitory autoreceptors¹²⁰. In fact, schizophrenia patients exhibit an overall Go bias (consistent with increased tonic striatal dopamine) coupled with decreased Go learning (consistent with decreased phasic striatal dopamine)⁵⁵. In the associative loop, the tonic Go bias may produce excessive updating of PFC representations with irrelevant information. Reduced PFC dopamine likely adds additional lability to PFC representations¹⁰⁰. The resulting extreme lability of PFC representations may underlie cognitive symptoms and contribute to positive symptoms.

Preclinical Animal Models: Conditioned Avoidance and Antipsychotics

RL models can also shed light on preclinical animal models used to test novel medications. One example is the use of conditioned avoidance to screen antipsychotics¹²¹. In conditioned avoidance, a CS is followed by shock unless animals perform a certain *avoidance response* after the onset of the CS but before the shock. The avoidance response produces a transition from a state with negative value (in which shock is expected) to a state with neutral value (in which no shock is expected), so it elicits a positive prediction error^{29, 122}. These positive prediction errors are hypothesized to strengthen the S-R association between the CS and the avoidance response in an actor-critic architecture^{29, 122}. Consistent with this idea, lesions of the ventral striatum (expected to damage the critic) and of the nigrostriatal dopaminergic projection (expected to prevent delivery of prediction errors to the actor) disrupt avoidance learning^{123–125}. Disrupting dopaminergic signaling in the dorsal striatum only following

training (through, for example, nigrostriatal lesions or dorsal striatal D2 blockade) does not, however, disrupt avoidance performance^{123, 124, 126}. These findings are also consistent with the predictions of the actor-critic account, because dopaminergic signaling of prediction errors in the dorsal striatum is necessary for S-R learning, but not for S-R expression.

Unlike in the dorsal striatum, D2 blockade in the nucleus accumbens following training disrupts avoidance performance¹²⁶. In fact, the standard use of conditioned avoidance to test antipsychotics is to administer them following training. The standard finding is that low antipsychotic doses disrupt avoidance but not escape from ongoing shock^{121, 127}, an effect that is mediated by the nucleus accumbens¹²¹. Dopamine in the nucleus accumbens can modulate the activation of instrumental behavior^{128, 129}, possibly via the striato-nigrostriatal spirals that allow the nucleus accumbens to influence the dorsal striatum^{56, 129}. The disruption of avoidance responding by antipsychotics may therefore reflect a decrease in the activation of instrumental behavior. Escapes would be less affected because the immediacy of pain would activate an innate flight response. An analogous situation is found in the appetitive domain: systemic or intra-accumbens administration of low doses of dopamine antagonists disrupts instrumental lever-pressing for food without affecting food approach or consumption¹²⁸.

Other Disorders

This review focused on a small number of RL principles and showed how they can shed light on multiple disorders. Much more, in fact, could have been said about each disorder, even using only these RL principles. Other aspects of RL and related computational approaches, however, seem likely to also be relevant for psychiatric and neurological disorders. For example, models of the role of the OFC in reversal learning⁶³ may be relevant for obsessive-compulsive disorder, which involves prominent OFC disturbances¹³⁰ and neural abnormalities during reversal learning¹³¹; models of the role of serotonin in RL^{22, 132} may be relevant for the many disorders that involve serotonergic abnormalities; and models of Pavlovian conditioning and extinction^{133, 134} may be relevant for some anxiety disorders (and have in fact already been shown to explain surprising findings in fear conditioning in humans¹³⁵).

CONCLUSIONS

RL models have been used to explain a wealth of findings across several psychiatric and neurological disorders. Although disorders as seemingly disparate as PD, TS, ADHD, schizophrenia, and addiction might seem to have little in common, they all involve disturbances in dopamine and CBGTC loops. The work reviewed above demonstrates that a mechanistic, computationally grounded understanding of the functions of these circuits sheds important light on all of these disorders. This approach relates to the new NIMH Research Domain Criteria (RDoC) initiative¹³⁶, which calls for research that cuts across diagnostic criteria and focuses instead on neurocognitive domains and how they go awry in a variety of DSM-defined conditions. The work reviewed above exemplifies this strategy. This work also amply demonstrates the new level of theoretical sophistication that computational psychiatry and computational neurology bring to the venerable disciplines of

psychiatry and neurology. Such theoretical sophistication and depth is essential if we are to fulfill the promise of a neuroscience-based, mechanistically detailed approach to diagnosis and treatment, which many agree should characterize the psychiatry and neurology of the future.

ACKNOWLEDGMENTS

Preparation of this article was funded by a Research Associate Award from the New York State Psychiatric Institute and the Research Foundation for Mental Hygiene, by NIMH grant R01 MH080066, and by a grant from the Michael J. Fox Foundation for Parkinson's Research.

REFERENCES

1. Hyman SE. Can neuroscience be integrated into the DSM-V? *Nat. Rev. Neurosci.* 2007; 8:725–732. [PubMed: 17704814]
2. Kendell R, Jablensky A. Distinguishing between the validity and utility of psychiatric diagnoses. *Am. J. Psychiatry.* 2003; 160:4–12. [PubMed: 12505793]
3. Charney, DS., et al. Neuroscience research agenda to guide development of a pathophysiologically based classification system. In: Kupfer, DJ.; First, MB.; Regier, DA., editors. *A Research Agenda for DSM-V.* American Psychiatric Association; Washington DC: 2002. p. 31-83.
4. Cools R. Dopaminergic modulation of cognitive function-implications for L-DOPA treatment in Parkinson's disease. *Neurosci. Biobehav. Rev.* 2006; 30:1–23. [PubMed: 15935475]
5. Frank MJ, Samanta J, Moustafa AA, Sherman SJ. Hold your horses: impulsivity, deep brain stimulation, and medication in parkinsonism. *Science.* 2007; 318:1309–1312. [PubMed: 17962524]
6. Maia TV. Reinforcement learning, conditioning, and the brain: Successes and challenges. *Cogn. Affect. Behav. Neurosci.* 2009; 9:343–364. [PubMed: 19897789]
7. Cohen MX, Frank MJ. Neurocomputational models of basal ganglia function in learning, memory and choice. *Behav. Brain Res.* 2009; 199:141–156. [PubMed: 18950662]
8. Dayan P, Niv Y, Seymour B, Daw ND. The misbehavior of value and the discipline of the will. *Neural Netw.* 2006; 19:1153–1160. [PubMed: 16938432]
9. McClure SM, Daw ND, Montague PR. A computational substrate for incentive salience. *Trends Neurosci.* 2003; 26:423–428. [PubMed: 12900173]
10. Dayan P, Balleine BW. Reward, motivation, and reinforcement learning. *Neuron.* 2002; 36:285–298. [PubMed: 12383782]
11. Niv Y, Daw ND, Joel D, Dayan P. Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology (Berl).* 2007; 191:507–520. [PubMed: 17031711]
12. Wiecki TV, Frank MJ. Neurocomputational models of motor and cognitive deficits in Parkinson's disease. *Prog. Brain Res.* 2010; 183:275–297. [PubMed: 20696325]
13. Frank MJ, Santamaria A, O'Reilly RC, Willcutt E. Testing computational models of dopamine and noradrenaline dysfunction in attention deficit/hyperactivity disorder. *Neuropsychopharmacology.* 2007; 32:1583–1599. [PubMed: 17164816]
14. Frank MJ, Scheres A, Sherman SJ. Understanding decision-making deficits in neurological conditions: insights from models of natural action selection. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 2007; 362:1641–1654. [PubMed: 17428775]
15. Waltz JA, Frank MJ, Robinson BM, Gold JM. Selective reinforcement learning deficits in schizophrenia support predictions from computational models of striatal-cortical dysfunction. *Biol. Psychiatry.* 2007; 62:756–764. [PubMed: 17300757]
16. Redish AD. Addiction as a computational process gone awry. *Science.* 2004; 306:1944–1947. [PubMed: 15591205]
17. Dayan P. Dopamine, reinforcement learning, and addiction. *Pharmacopsychiatry.* 2009; 42(Suppl 1):S56–S65. [PubMed: 19434556]
18. Montague PR, Dayan P, Sejnowski TJ. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* 1996; 16:1936–1947. [PubMed: 8774460]

19. Sutton, RS.; Barto, AG. Reinforcement Learning: An Introduction. MIT Press; Cambridge, MA: 1998.
20. Bayer HM, Glimcher PW. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*. 2005; 47:129–141. [PubMed: 15996553]
21. Bayer HM, Lau B, Glimcher PW. Statistics of midbrain dopamine neuron spike trains in the awake primate. *J. Neurophysiol.* 2007; 98:1428–1439. [PubMed: 17615124]
22. Daw ND, Kakade S, Dayan P. Opponent interactions between serotonin and dopamine. *Neural Netw.* 2002; 15:603–616. [PubMed: 12371515]
23. Schultz W. Dopamine signals for reward value and risk: basic and recent data. *Behav Brain Fund.* 2010; 6:24.
24. Brischoux F, Chakraborty S, Brierley DI, Ungless MA. Phasic excitation of dopamine neurons in ventral VTA by noxious stimuli. *Proc. Natl. Acad. Sci. U.S.A.* 2009; 106:4894–4899. [PubMed: 19261850]
25. Matsumoto M, Hikosaka O. Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature*. 2009; 459:837–841. [PubMed: 19448610]
26. Frank MJ, Surmeier DJ. Do substantia nigra dopaminergic neurons differentiate between reward and punishment? *J. Mol. Cell Biol.* 2009; 1:15–16. [PubMed: 19770170]
27. Morris G, Schmidt R, Bergman H. Striatal action-learning based on dopamine concentration. *Exp. Brain Res.* 2010; 200:307–317. [PubMed: 19904530]
28. Tsai HC, et al. Phasic firing in dopaminergic neurons is sufficient for behavioral conditioning. *Science*. 2009; 324:1080–1084. [PubMed: 19389999]
29. Maia TV. Two-factor theory, the actor-critic model, and conditioned avoidance. *Learn. Behav.* 2010; 38:50–67. [PubMed: 20065349]
30. O'Doherty J, et al. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*. 2004; 304:452–454. [PubMed: 15087550]
31. Barto AG, Sutton RS, Anderson CW. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Trans. Syst. Man Cybern.* 1983; 13:835–846.
32. Packard MG, Knowlton BJ. Learning and memory functions of the basal ganglia. *Annu. Rev. Neurosci.* 2002; 25:563–593. [PubMed: 12052921]
33. Alexander GE, DeLong MR, Strick PL. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu. Rev. Neurosci.* 1986; 9:357–381. [PubMed: 3085570]
34. Yin HH, Knowlton BJ. The role of the basal ganglia in habit formation. *Nat. Rev. Neurosci.* 2006; 7:464–476. [PubMed: 16715055]
35. Mink JW. The basal ganglia: focused selection and inhibition of competing motor programs. *Prog. Neurobiol.* 1996; 50:381–425. [PubMed: 9004351]
36. Prescott, TJ.; Gurney, K.; Redgrave, P. Basal ganglia. In: Arbib, MA., editor. *The Handbook of Brain Theory and Neural Networks*. MIT Press; Cambridge, MA: 2003. p. 147-151.
37. Frank MJ. Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and non-medicated Parkinsonism. *J. Cogn. Neurosci.* 2005; 17:51–72. [PubMed: 15701239]
38. Samejima K, Ueda Y, Doya K, Kimura M. Representation of action-specific reward values in the striatum. *Science*. 2005; 310:1337–1340. [PubMed: 16311337]
39. Albin RL, Young AB, Penney JB. The functional anatomy of basal ganglia disorders. *Trends Neurosci.* 1989; 12:366–375. [PubMed: 2479133]
40. DeLong MR. Primate models of movement disorders of basal ganglia origin. *Trends Neurosci.* 1990; 13:281–285. [PubMed: 1695404]
41. Nambu A, Tokuno H, Takada M. Functional significance of the cortico-subthalamo-pallidal 'hyperdirect' pathway. *Neurosci. Res.* 2002; 43:111–117. [PubMed: 12067746]
42. Beiser DG, Hua SE, Houk JC. Network models of the basal ganglia. *Curr. Opin. Neurobiol.* 1997; 7:185–190. [PubMed: 9142759]
43. Bogacz R, Gurney K. The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput.* 2007; 19:442–477. [PubMed: 17206871]

44. Brown JW, Bullock D, Grossberg S. How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Netw.* 2004; 17:471–510. [PubMed: 15109680]
45. Frank MJ. Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural Netw.* 2006; 19:1120–1136. [PubMed: 16945502]
46. Gurney K, Prescott TJ, Redgrave P. A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour. *Biol. Cybern.* 2001; 84:411–423. [PubMed: 11417053]
47. Frank MJ, Fossella JA. Neurogenetics and pharmacology of learning, motivation, and cognition. *Neuropsychopharmacology.* 2010
48. Hikida T, Kimura K, Wada N, Funabiki K, Nakanishi S. Distinct roles of synaptic transmission in direct and indirect striatal pathways to reward and aversive behavior. *Neuron.* 2010; 66:896–907. [PubMed: 20620875]
49. Kravitz AV, et al. Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. *Nature.* 2010
50. Gerfen CR. Molecular effects of dopamine on striatal-projection pathways. *Trends Neurosci.* 2000; 23:S64–S70. [PubMed: 11052222]
51. Shen W, Flajolet M, Greengard P, Surmeier DJ. Dichotomous dopaminergic control of striatal synaptic plasticity. *Science.* 2008; 321:848–851. [PubMed: 18687967]
52. Creese I, Sibley DR, Hamblin MW, Leff SE. The classification of dopamine receptors: relationship to radioligand binding. *Annu. Rev. Neurosci.* 1983; 6:43–71. [PubMed: 6220666]
53. Hernandez-Lopez S, Bargas J, Surmeier DJ, Reyes A, Galarraga E. D1 receptor activation enhances evoked discharge in neostriatal medium spiny neurons by modulating an L-type Ca²⁺ conductance. *J. Neurosci.* 1997; 17:3334–3342. [PubMed: 9096166]
54. Hernandez-Lopez S, et al. D2 dopamine receptors in striatal medium spiny neurons reduce L-type Ca²⁺ currents and excitability via a novel PLC[β]₁-IP₃-calcineurin-signaling cascade. *J. Neurosci.* 2000; 20:8987–8995. [PubMed: 11124974]
55. Waltz JA, Frank MJ, Wiecki TV, Gold JM. Altered probabilistic learning and response biases in schizophrenia: Behavioral evidence and neurocomputational modeling. *Neuropsychology.* (in press).
56. Haber SN. The primate basal ganglia: parallel and integrative networks. *J. Chem. Neuroanat.* 2003; 26:317–330. [PubMed: 14729134]
57. Postuma RB, Dagher A. Basal ganglia functional connectivity based on a meta-analysis of 126 positron emission tomography and functional magnetic resonance imaging publications. *Cereb. Cortex.* 2006; 16:1508–1521. [PubMed: 16373457]
58. Frank MJ, Loughry B, O'Reilly RC. Interactions between frontal cortex and basal ganglia in working memory: a computational model. *Cogn. Affect. Behav. Neurosci.* 2001; 1:137–160. [PubMed: 12467110]
59. O'Reilly RC, Frank MJ. Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural Comput.* 2006; 18:283–328. [PubMed: 16378516]
60. Baier B, et al. Keeping memory clear and stable--The contribution of human basal ganglia and prefrontal cortex to working memory. *J. Neurosci.* 2010; 30:9788–9792. [PubMed: 20660261]
61. Tzschentke TM. Measuring reward with the conditioned place preference paradigm: a comprehensive review of drug effects, recent progress and new issues. *Prog. Neurobiol.* 1998; 56:613–672. [PubMed: 9871940]
62. Lobo MK, et al. Cell type-specific loss of BDNF signaling mimics optogenetic control of cocaine reward. *Science.* 2010; 330:385–390. [PubMed: 20947769]
63. Frank MJ, Claus ED. Anatomy of a decision: striato-orbitofrontal interactions in reinforcement learning, decision making, and reversal. *Psychol. Rev.* 2006; 113:300–326. [PubMed: 16637763]
64. Kalanithi PS, et al. Altered parvalbumin-positive neuron distribution in basal ganglia of individuals with Tourette syndrome. *Proc. Natl. Acad. Sci. U.S.A.* 2005; 102:13307–13312. [PubMed: 16131542]
65. Kataoka Y, et al. Decreased number of parvalbumin and cholinergic interneurons in the striatum of individuals with Tourette syndrome. *J. Comp. Neurol.* 2010; 518:277–291. [PubMed: 19941350]

66. Biederman J, Spencer T. Attention-deficit/hyperactivity disorder (ADHD) as a noradrenergic disorder. *Biol. Psychiatry*. 1999; 46:1234–1242. [PubMed: 10560028]
67. Guillin O, Abi-Dargham A, Laruelle M. Neurobiology of dopamine in schizophrenia. *Int. Rev. Neurobiol.* 2007; 78:1–39. [PubMed: 17349856]
68. Nikolaus S, Antke C, Muller HW. In vivo imaging of synaptic function in the central nervous system: II. Mental and affective disorders. *Behav. Brain Res.* 2009; 204:32–66. [PubMed: 19523495]
69. Stahl SM. Beyond the dopamine hypothesis to the NMDA glutamate receptor hypofunction hypothesis of schizophrenia. *CNS Spectr.* 2007; 12:265–268. [PubMed: 17426663]
70. Rolls ET, Loh M, Deco G, Winterer G. Computational models of schizophrenia and dopamine modulation in the prefrontal cortex. *Nat. Rev. Neurosci.* 2008; 9:696–709. [PubMed: 18714326]
71. Wichmann T, DeLong MR. Functional and pathophysiological models of the basal ganglia. *Curr. Opin. Neurobiol.* 1996; 6:751–758. [PubMed: 9000030]
72. Frank MJ, Seeberger LC, O'Reilly RC. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*. 2004; 306:1940–1943. [PubMed: 15528409]
73. Moustafa AA, Sherman SJ, Frank MJ. A dopaminergic basis for working memory, learning and attentional shifting in Parkinsonism. *Neuropsychologia*. 2008; 46:3144–3156. [PubMed: 18687347]
74. Cools R, Miyakawa A, Sheridan M, D'Esposito M. Enhanced frontal function in Parkinson's disease. *Brain*. 2010; 133:225–233. [PubMed: 19995871]
75. Frank MJ, Moustafa AA, Haughey HM, Curran T, Hutchison KE. Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proc. Natl. Acad. Sci. U.S.A.* 2007; 104:16311–16316. [PubMed: 17913879]
76. Bodi N, et al. Reward-learning and the novelty-seeking personality: a between- and within-subjects study of the effects of dopamine agonists on young Parkinson's patients. *Brain*. 2009; 132:2385–2395. [PubMed: 19416950]
77. Moustafa AA, Cohen MX, Sherman SJ, Frank MJ. A role for dopamine in temporal decision making and reward maximization in parkinsonism. *J. Neurosci.* 2008; 28:12294–12304. [PubMed: 19020023]
78. Palminteri S, et al. Pharmacological modulation of subliminal learning in Parkinson's and Tourette's syndromes. *Proc. Natl. Acad. Sci. USA*. 2009; 106:19179–19184. [PubMed: 19850878]
79. Voon V, et al. Mechanisms underlying dopamine-mediated reward bias in compulsive behaviors. *Neuron*. 2010; 65:135–142. [PubMed: 20152119]
80. Dagher A, Robbins TW. Personality, addiction, dopamine: insights from Parkinson's disease. *Neuron*. 2009; 61:502–510. [PubMed: 19249271]
81. Mink JW. Basal ganglia dysfunction in Tourette's syndrome: a new hypothesis. *Pediatr. Neurol.* 2001; 25:190–198. [PubMed: 11587872]
82. Delfs JM, Kelley AE. The role of D1 and D2 dopamine receptors in oral stereotypy induced by dopaminergic stimulation of the ventrolateral striatum. *Neuroscience*. 1990; 39:59–67. [PubMed: 1982467]
83. Walters JR, Bergstrom DA, Carlson JH, Chase TN, Braun AR. D1 dopamine receptor activation required for postsynaptic expression of D2 agonist effects. *Science*. 1987; 236:719–722. [PubMed: 2953072]
84. Steeves TD, Fox SH. Neurobiological basis of serotonin-dopamine antagonists in the treatment of Gilles de la Tourette syndrome. *Prog. Brain Res.* 2008; 172:495–513. [PubMed: 18772048]
85. Peterson AL, Azrin NH. Behavioral and pharmacological treatments for Tourette syndrome: A review. *Appl. Prev. Psychol.* 1993; 2:231–242.
86. Leckman JF, Walker DE, Cohen DJ. Premonitory urges in Tourette's syndrome. *Am. J. Psychiatry*. 1993; 150:98–102. [PubMed: 8417589]
87. Peterson, BS., et al. Neuroanatomical circuitry. In: Leckman, JF.; Cohen, DJ., editors. *Tourette's Syndrome: Tics, Obsessions, Compulsions. Developmental Psychopathology and Clinical Care.* John Wiley & Sons, Inc.; New York, NY: 1999. p. 230-259.

88. Bohlhalter S, et al. Neural correlates of tic generation in Tourette syndrome: an event-related functional MRI study. *Brain*. 2006; 129:2029–2037. [PubMed: 16520330]
89. Barkley RA. Behavioral inhibition, sustained attention, and executive functions: constructing a unifying theory of ADHD. *Psychol. Bull.* 1997; 121:65–94. [PubMed: 9000892]
90. Sagvolden T, Sergeant JA. Attention deficit/hyperactivity disorder--from brain dysfunctions to behaviour. *Behav. Brain Res.* 1998; 94:1–10. [PubMed: 9708834]
91. Sonuga-Barke EJ. Causal models of attention-deficit/hyperactivity disorder: from common simple deficits to multiple developmental pathways. *Biol. Psychiatry*. 2005; 57:1231–1238. [PubMed: 15949993]
92. Sagvolden T, Johansen EB, Aase H, Russell VA. A dynamic developmental theory of attention-deficit/hyperactivity disorder (ADHD) predominantly hyperactive/impulsive and combined subtypes. *Behav. Brain Sci.* 2005; 28:397–419. [PubMed: 16209748]
93. Smith AJ, Becker S, Kapur S. A computational model of the functional role of the ventral- striatal D2 receptor in the expression of previously acquired behaviors. *Neural Comput.* 2005; 17:361–395. [PubMed: 15720772]
94. van Gaalen MM, van Koten R, Schoffelmeeer AN, Vanderschuren LJ. Critical involvement of dopaminergic neurotransmission in impulsive decision making. *Biol. Psychiatry*. 2006; 60:66–73. [PubMed: 16125144]
95. Wade TR, de Wit H, Richards JB. Effects of dopaminergic drugs on delayed reward as a measure of impulsive behavior in rats. *Psychopharmacology (Berl)*. 2000; 150:90–101. [PubMed: 10867981]
96. Cardinal RN. Neural systems implicated in delayed and probabilistic reinforcement. *Neural Netw.* 2006; 19:1277–1301. [PubMed: 16938431]
97. Pattij T, Vanderschuren LJ. The neuropharmacology of impulsive behaviour. *Trends Pharmacol. Sci.* 2008; 29:192–199. [PubMed: 18304658]
98. Winstanley CA, Theobald DE, Dalley JW, Robbins TW. Interactions between serotonin and dopamine in the control of impulsive choice in rats: therapeutic implications for impulse control disorders. *Neuropsychopharmacology*. 2005; 30:669–682. [PubMed: 15688093]
99. Kheramin S, et al. Effects of orbital prefrontal cortex dopamine depletion on inter-temporal choice: a quantitative analysis. *Psychopharmacology (Berl)*. 2004; 175:206–214. [PubMed: 14991223]
100. Durstewitz D, Seamans JK, Sejnowski TJ. Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. *J. Neurophysiol.* 2000; 83:1733–1750. [PubMed: 10712493]
101. Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 2001; 24:167–202. [PubMed: 11283309]
102. Maia TV, Cleeremans A. Consciousness: converging insights from connectionist modeling and neuroscience. *Trends Cogn. Sci.* 2005; 9:397–404. [PubMed: 16005677]
103. Figner B, et al. Lateral prefrontal cortex and self-control in intertemporal choice. *Nat. Neurosci.* 2010; 13:538–539. [PubMed: 20348919]
104. McClure SM, Laibson DI, Loewenstein G, Cohen JD. Separate neural systems value immediate and delayed monetary rewards. *Science*. 2004; 306:503–507. [PubMed: 15486304]
105. Volkow ND, Fowler JS, Wang GJ, Swanson JM, Telang F. Dopamine in drug abuse and addiction: results of imaging studies and treatment implications. *Arch. Neurol.* 2007; 64:1575–1579. [PubMed: 17998440]
106. Kamin, LJ. Predictability, surprise, attention, and conditioning. In: Campbell, BA.; Church, RM., editors. *Punishment and Aversive Behavior*. Appleton-Century-Crofts; New York: 1969. p. 279-296.
107. Panlilio LV, Thorndike EB, Schindler CW. Blocking of conditioning to a cocaine-paired stimulus: testing the hypothesis that cocaine perpetually produces a signal of larger-than-expected reward. *Pharmacol. Biochem. Behav.* 2007; 86:774–777. [PubMed: 17445874]
108. Dezfouli A, et al. A neurocomputational model for cocaine addiction. *Neural Comput.* 2009; 21:2869–2893. [PubMed: 19635010]
109. Mahadevan S. Average reward reinforcement learning: Foundations, algorithms, and empirical results. *Mach. Learn.* 1996; 22:159–196.

110. Redish AD, Jensen S, Johnson A. A unified framework for addiction: vulnerabilities in the decision process. *Behav. Brain Sci.* 2008; 31:415–437. discussion 437–487. [PubMed: 18662461]
111. Koob GF, Volkow ND. Neurocircuitry of addiction. *Neuropsychopharmacology.* 2010; 35:217–238. [PubMed: 19710631]
112. Schoenbaum G, Roesch MR, Stalnaker TA. Orbitofrontal cortex, decision-making and drug addiction. *Trends Neurosci.* 2006; 29:116–124. [PubMed: 16406092]
113. Berridge KC, Robinson TE. What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Res. Brain Res. Rev.* 1998; 28:309–369. [PubMed: 9858756]
114. Kapur S. Psychosis as a state of aberrant salience: a framework linking biology, phenomenology, and pharmacology in schizophrenia. *Am. J. Psychiatry.* 2003; 160:13–23. [PubMed: 12505794]
115. Corlett PR, Honey GD, Fletcher PC. From prediction error to psychosis: ketamine as a pharmacological model of delusions. *J. Psychopharmacol.* 2007; 21:238–252. [PubMed: 17591652]
116. Corlett PR, et al. Disrupted prediction-error signal in psychosis: evidence for an associative account of delusions. *Brain.* 2007; 130:2387–2400. [PubMed: 17690132]
117. Murray GK, et al. Substantia nigra/ventral tegmental reward prediction error disruption in psychosis. *Mol. Psychiatry.* 2008; 13:239, 267–276. [PubMed: 17684497]
118. Waltz JA, et al. Patients with schizophrenia have a reduced neural response to both unpredictable and predictable primary reinforcers. *Neuropsychopharmacology.* 2009; 34:1567–1577. [PubMed: 19052540]
119. Strauss GP, et al. Deficits in positive reinforcement learning and uncertainty-driven exploration are associated with distinct aspects of negative symptoms in schizophrenia. *Biol. Psychiatry.* (in press).
120. Grace AA. The tonic/phasic model of dopamine system regulation and its implications for understanding alcohol and psychostimulant craving. *Addiction.* 2000; 95(Suppl 2):S119–S128. [PubMed: 11002907]
121. Wadenberg ML, Hicks PB. The conditioned avoidance response test re-evaluated: is it a sensitive test for the detection of potentially atypical antipsychotics? *Neurosci. Biobehav. Rev.* 1999; 23:851–862. [PubMed: 10541060]
122. Moutoussis M, Bentall RP, Williams J, Dayan P. A temporal difference account of avoidance learning. *Network.* 2008; 19:137–160. [PubMed: 18569725]
123. Fibiger HC, Phillips AG, Zis AP. Deficits in instrumental responding after 6-hydroxydopamine lesions of the nigro-neostriatal dopaminergic projection. *Pharmacol. Biochem. Behav.* 1974; 2:87–96. [PubMed: 4151316]
124. Zis AP, Fibiger HC, Phillips AG. Reversal by L-dopa of impaired learning due to destruction of the dopaminergic nigro-neostriatal projection. *Science.* 1974; 185:960–962. [PubMed: 4471715]
125. Fantin G, Bottechia D. Effect of nucleus accumbens destruction in rat. *Experientia.* 1984; 40:573–575. [PubMed: 6723931]
126. Wadenberg ML, Ericson E, Magnusson O, Ahlenius S. Suppression of conditioned avoidance behavior by the local application of (-)sulpiride into the ventral, but not the dorsal, striatum of the rat. *Biol. Psychiatry.* 1990; 28:297–307. [PubMed: 2144458]
127. Smith A, Li M, Becker S, Kapur S. A model of antipsychotic action in conditioned avoidance: a computational approach. *Neuropsychopharmacology.* 2004; 29:1040–1049. [PubMed: 14997176]
128. Salamone JD, Correa M. Motivational views of reinforcement: implications for understanding the behavioral functions of nucleus accumbens dopamine. *Behav. Brain Res.* 2002; 137:3–25. [PubMed: 12445713]
129. Yin HH, Ostlund SB, Balleine BW. Reward-guided learning beyond dopamine in the nucleus accumbens: the integrative functions of cortico-basal ganglia networks. *Eur. J Neurosci.* 2008; 28:1437–1448. [PubMed: 18793321]
130. Maia TV, Cooney RE, Peterson BS. The neural bases of obsessive-compulsive disorder in children and adults. *Dev. Psychopathol.* 2008; 20:1251–1283. [PubMed: 18838041]
131. Remijnse PL, et al. Reduced orbitofrontal-striatal activity on a reversal learning task in obsessive-compulsive disorder. *Arch. Gen. Psychiatry.* 2006; 63:1225–1236. [PubMed: 17088503]

132. Dayan P, Huys QJ. Serotonin in affective control. *Annu. Rev. Neurosci.* 2009; 32:95–126. [PubMed: 19400722]
133. Dayan P, Kakade S, Montague PR. Learning and selective attention. *Nat. Neurosci.* 2000; (3 Suppl):1218–1223. [PubMed: 11127841]
134. Gershman SJ, Blei DM, Niv Y. Context, learning, and extinction. *Psychol. Rev.* 2010; 117:197–209. [PubMed: 20063968]
135. Maia TV. Fear conditioning and social groups: statistics, not genetics. *Cogn. Sci.* 2009; 33:1232–1251. [PubMed: 21585503]
136. Insel TR, Cuthbert BN. Endophenotypes: bridging genomic complexity and disorder heterogeneity. *Biol. Psychiatry.* 2009; 66:988–989. [PubMed: 19900610]
137. Aron AR, Behrens TE, Smith S, Frank MJ, Poldrack RA. Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI. *J. Neurosci.* 2007; 27:3743–3752. [PubMed: 17409238]

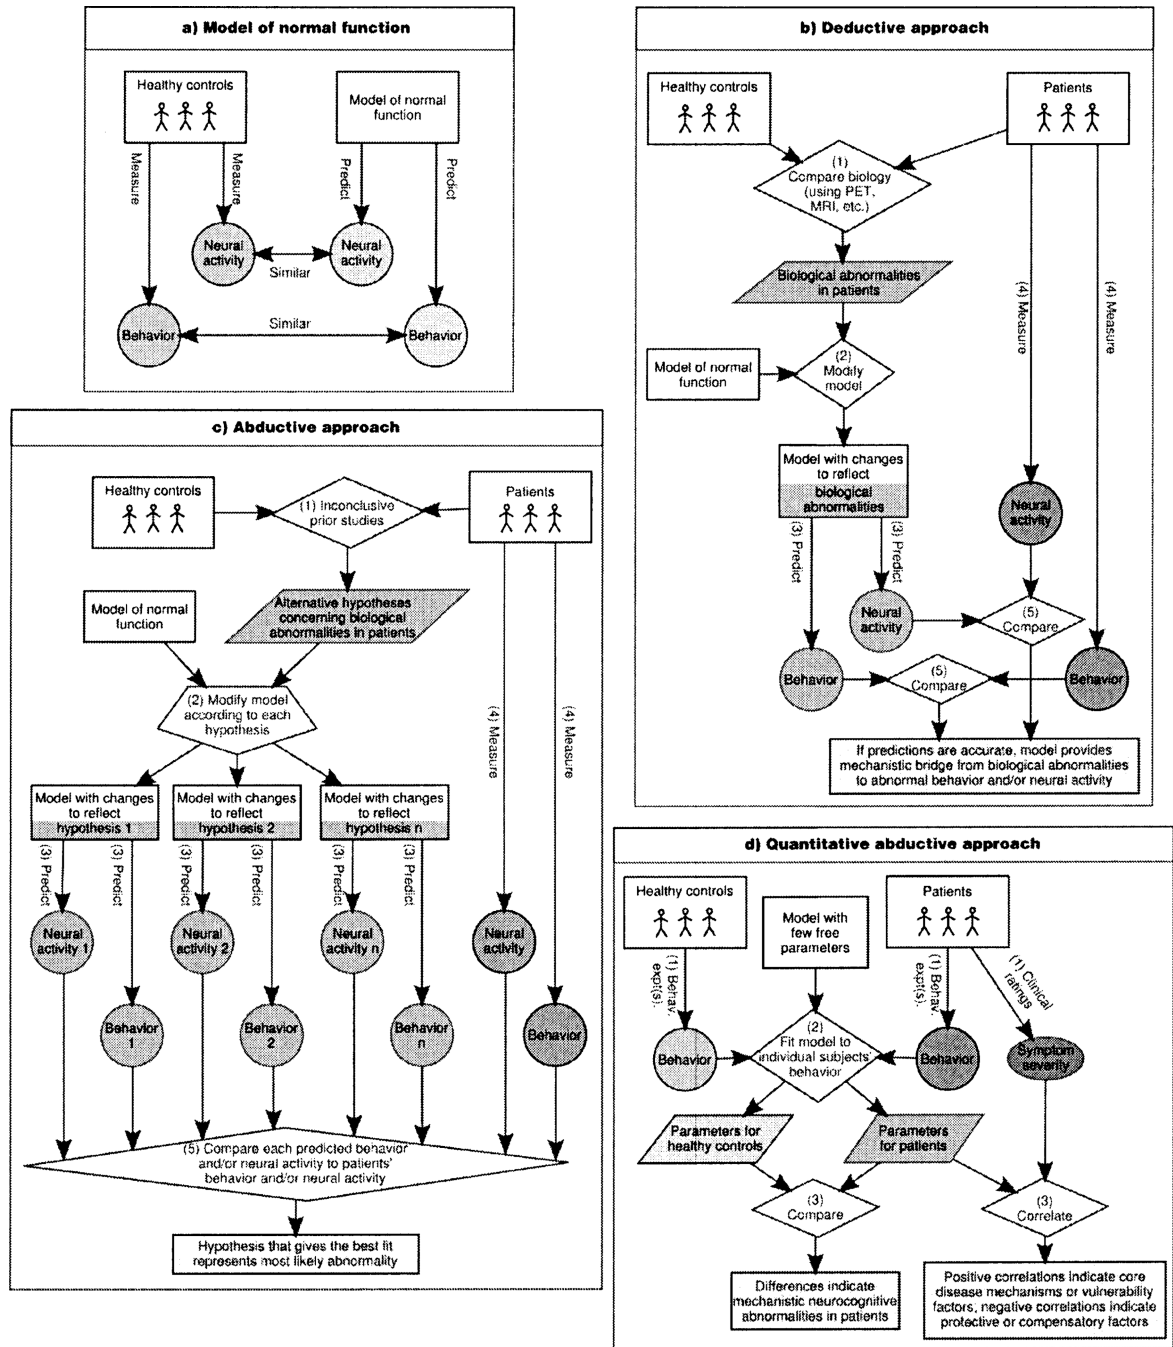
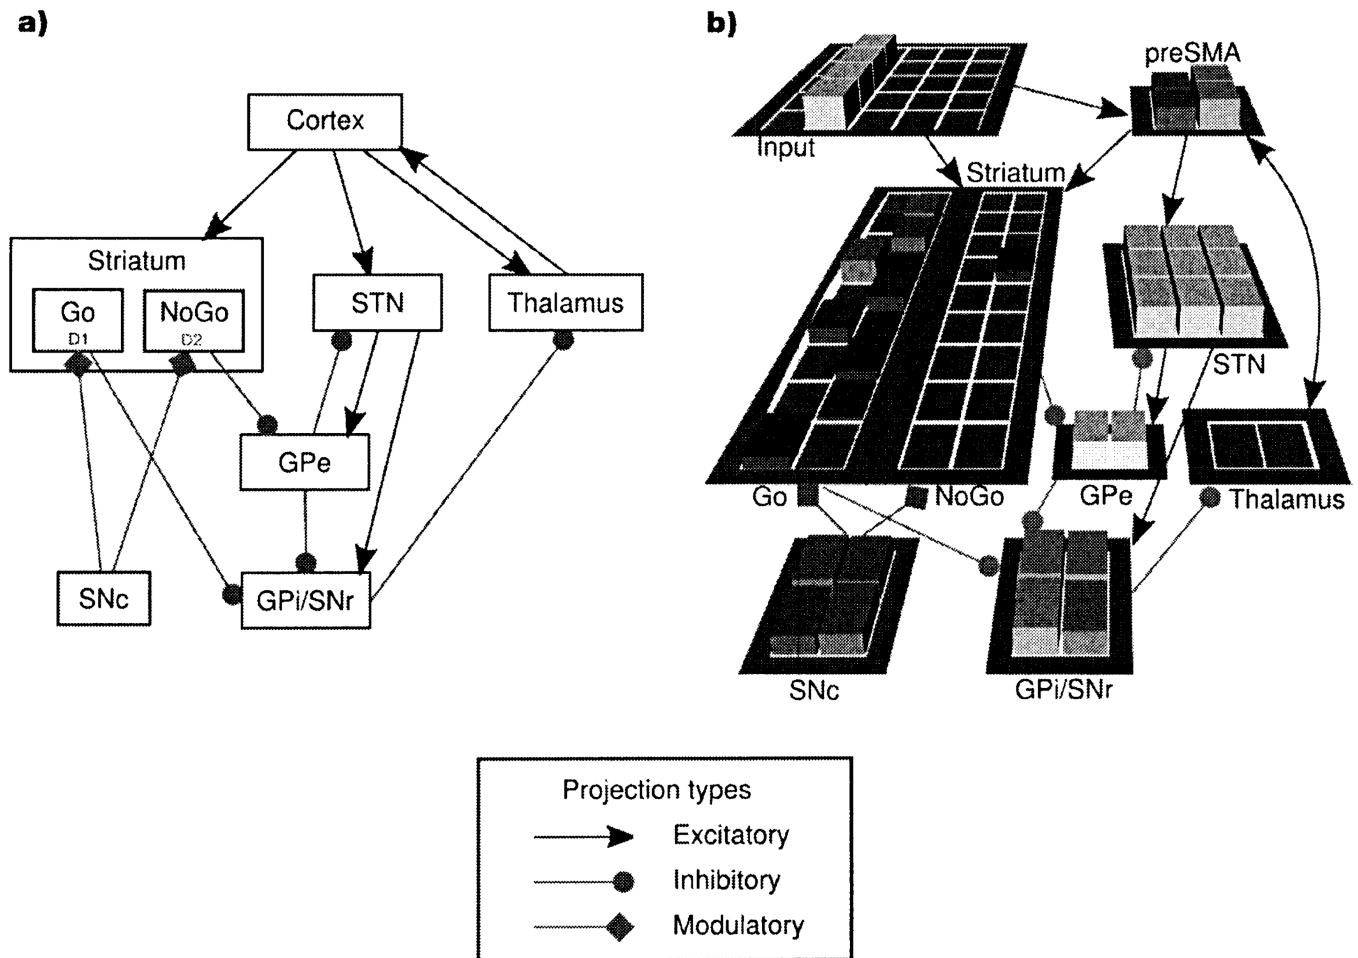


Figure 1. Principles of computational psychiatry and computational neurology, (a) The starting point in computational psychiatry and computational neurology is a model of normal function that captures key aspects of behavior and/or neural activity. Models at various levels of abstraction can be useful (e.g., algorithmic models from machine learning or neural models from computational cognitive neuroscience). Several complementary approaches can then be pursued, (b) With detailed neural models, pathophysiological processes can be simulated by making principled changes to the model that mimic biological alterations in the disorder

under consideration (e.g., alterations in striatal dopaminergic innervation or D2 receptor density). The systems-level and behavioral implications of these changes can then be explored using the model, leading to testable predictions. We call this approach ‘deductive,’ because the models are used to recreate the mechanistic link between causes (the biological abnormalities) and their consequences (abnormalities in behavior and/or systems-level neural activity). This approach can elucidate whether the observed biological abnormalities are sufficient to explain a behavioral phenotype. (c) A second approach involves using a model to try to infer the causes of the observed behavioral phenotype and/or of the observed alterations in neural activity. We call this approach ‘abductive,’ because it involves reasoning from consequences (the behavior or systems-level neural activity) to their possible causes (the underlying biological abnormalities). In this approach, alternative a priori hypotheses concerning possible biological abnormalities in a given disorder can be compared to determine which, if any, produce the same set of abnormalities in behavior and/or neural activity that is found in the disorder (Maia and Peterson, submitted manuscript), (d) A third approach, used more often with algorithmic than with neural models (largely because the former tend to have fewer parameters), involves fitting the model’s parameters to the behavior of individual subjects on a suitable task or set of tasks, and then determining if there are parameter differences between diseased and healthy subject groups, or correlations between parameters and disease severity. We call this approach ‘quantitative abductive,’ because it also involves reasoning from behavior to its mechanistic causes. A fourth, related approach (not shown graphically) also involves fitting a model’s parameters to subjects’ behavior, but the goal is to then calculate, on a trial-by-trial basis, each subject’s putative internal representation of the quantities calculated by the model (e.g., state values or prediction errors). These predicted internal representations are then used as regressors in functional imaging (e.g., fMRI, EEG), to find their neural correlates, which are then compared across the diseased and healthy groups. Each of these four approaches can also be adapted to study the effects of treatments (e.g., medication or neurosurgery). Furthermore, additional leverage can sometimes be gained by the synergistic use of different approaches and/or models at different levels of abstraction. For example, the deductive or abductive approaches are especially powerful with neural models, because these models embody mechanistic details that permit direct simulation of biological abnormalities. Such models, however, sometimes include too many parameters to make a quantitative abductive approach feasible. In some cases, a useful strategy is to construct an algorithmic model that includes parameters that reflect distinct mechanisms in the neural model (e.g., a Q-learning model with different learning rates for positive and negative prediction errors has been used to capture the prediction from a neural model of the basal ganglia that these two types of learning rely on distinct mechanisms⁷⁵). The neural model’s deductive predictions concerning how a disorder affects these parameters can then be verified using the quantitative abductive approach with the algorithmic model.

**Figure 2.**

(a) Anatomy of cortico-basal ganglia-thalamo-cortical loops. Striatal medium spiny neurons (MSNs) in the direct pathway ('Go' neurons) express mostly D1 receptors⁵⁰ and project directly to the globus pallidus internal segment (GPi) and the substantia nigra pars reticulata (SNr). [The GPi and the SNr have similar functions, so we treat them as a single complex (GPi/SNr).] Go neurons inhibit the GPi/SNr, which in turn results in disinhibition of the thalamus, thereby facilitating execution of the corresponding action. Striatal MSNs in the indirect pathway ('NoGo' neurons) express mostly D2 receptors⁵⁰ and project to the globus pallidus external segment (GPe), which in turn projects to the GPi/SNr. NoGo neurons produce a focused removal of the tonic inhibition of the GPe on the GPi/SNr, thereby disinhibiting the GPi/SNr, which in turn results in suppression of the corresponding action in the thalamus. Neurons in the subthalamic nucleus (STN) receive direct projections from the cortex in the hyperdirect pathway and project to both the GPe and GPi/SNr. The projections from the STN to the GPe and GPi/SNr are diffuse³⁵, so they are believed to modulate all actions rather than a specific action, **(b)** The basal ganglia Go/NoGo model^{37,45}. The synaptic connections in the model are consistent with the anatomical connections shown in (a). The model learns to map inputs, which represent the current stimuli and/or situation (i.e., the current state), to actions represented in the pre-supplementary motor area (preSMA). Corticocortical projections from the input layer to preSMA activate in preSMA

candidate actions that are appropriate for the current state. The basal ganglia then act to facilitate ('gate') the best action - i.e., the action with the best reinforcement history for the current state - while simultaneously suppressing the other actions (at the level of the thalamus). Distributed populations of Go and NoGo units represent, respectively, the positive and negative evidence for the candidate actions in the current state. Lateral inhibition between the Go and NoGo pathways ensures that the probability of selecting a given action is a function of the difference between the positive and negative evidence for that action. The positive and negative evidence for each action in each state is learned on the basis of past reinforcement history, through the actions of dopamine on D1 and D2 receptors in striatal Go and NoGo units, respectively (see text). The synaptic weights in the corticocortical projections from the input layer to preSMA are themselves learned, but through Hebbian mechanisms, thereby allowing these corticocortical projections to activate candidate actions in preSMA in proportion to their prior probability of being executed in the given state. Thus, the candidate actions generated by these corticocortical projections for a given state tend to be those that have previously often been gated by the basal ganglia in that state. When two or more actions become strongly activated in preSMA because they have similar reinforcement histories, this response conflict activates the STN via the hyperdirect pathway (consistent with the evidence for direct anatomical connections between the preSMA and the STN, and their co-activation in high-conflict situations¹³⁷). The STN then provides a 'Global NoGo' signal that prevents premature facilitation of suboptimal responses⁴⁵.

Training

ま み
A (80%) B (20%)

そ の
C (70%) D (30%)

ら や
E (60%) F (40%)

Testing

Go learning

ま > そ の ら や
A C D E F

NoGo learning

み < そ の ら や
B C D E F

Figure 3. The probabilistic selection task⁷². The probabilistic selection task is used to assess whether participants learn better from positive or negative outcomes. During training, in each trial participants are presented with one of the three pairs shown on top (AB, CD, and EF), and select one of the two stimuli. Feedback then indicates if the choice was correct or incorrect. The probabilities of each stimulus leading to correct feedback are indicated in the figure. Participants may learn to perform accurately during training (i.e., learn to select A, C, and E) by learning which stimulus in each pair is associated with positive feedback (Go learning), by learning which stimulus in each pair is associated with negative feedback (NoGo learning), or both. The test phase assesses the degree to which participants learned better from positive or from negative feedback. Participants are presented with novel pairs of stimuli consisting of either an A or a B paired with each of the other stimuli (C through F,

which on average had 50% probability of positive feedback during training). No feedback is provided during testing. If participants perform better on the pairs that include A than on those that include B, that indicates that they learned better to select the most positive stimulus (A) than to avoid the most negative stimulus (B), so they learn better from positive feedback (Go learning). If they perform better on the pairs that include B, they learn better from negative feedback (NoGo learning). Indeed, individual differences in neural responses to negative outcomes predict individual differences in performance on the pairs that include B (but not on those that include A)⁴⁷.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript