Phil. Trans. R. Soc. Lond. B (0000)
DOI: 10.1098/rstb0000.0000

Understanding decision making deficits in neurological conditions: Insights from models of natural action selection

Supplementary Appendix

Michael J. Frank, Anouk Scheres, and Scott J. Sherman

Departments of Psychology and Neurology, Program in Neuroscience
University of Arizona Tucson, AZ 85721

1. Appendix for Basic BG Model

This appendix outlines the implementational details for the "basic" BG model described in the main text (i.e, without the Orbitofrontal cortex.) For details of the OFC model, please see Frank and Claus (2006). The model code, written in PDP++, can be obtained by emailing the corresponding author at `mfrank@u.arizona.edu`. For animated video captures of model dynamics during response selection and learning, please see `www.u.arizona.edu/~mfrank/BGmodel_movies.html`

(*a*) *Implementational Details*

The model is implemented using the Leabra framework (O'Reilly & Munakata, 2000; O'Reilly, 2001). Leabra uses point neurons with excitatory, inhibitory, and leak conductances contributing to an integrated membrane potential, which is then thresholded and transformed via an $x/(x+1)$ sigmoidal function to produce a rate code output communicated to other units (discrete spiking can also be used, but produces noisier results). Each layer uses a k-winners-take-all (kWTA) function that computes an inhibitory conductance that keeps roughly the k most active units above firing threshold and keeps the rest below threshold.

The membrane potential V_m is updated as a function of ionic conductances g with reversal (driving) potentials E as follows:

$$\Delta V_m(t) = \tau \sum_c g_c(t)\overline{g_c}(E_c - V_m(t)) \qquad (1)$$

with 3 channels (c) corresponding to: e excitatory input; l leak current; and i inhibitory input. Following electrophysiological convention, the overall conductance is decomposed into a time-varying component $g_c(t)$ computed as a function of the dynamic state of the network, and a constant $\overline{g_c}$ that controls the relative influence of the different conductances. The equilibrium potential can be written in a simplified form by setting the excitatory driving potential (E_e) to 1 and the leak and inhibitory driving potentials (E_l and E_i) of 0:

$$V_m^\infty = \frac{g_e\overline{g_e}}{g_e\overline{g_e} + g_l\overline{g_l} + g_i\overline{g_i}} \qquad (2)$$

which shows that the neuron is computing a balance between excitation and the opposing forces of leak and inhibition. This equilibrium form of the equation can be understood in terms of a Bayesian decision making framework (O'Reilly & Munakata, 2000).

The excitatory net input/conductance $g_e(t)$ or η_j is computed as the proportion of open excitatory channels as a function of sending activations times the weight values:

$$\eta_j = g_e(t) = \langle x_i w_{ij} \rangle = \frac{1}{n}\sum_i x_i w_{ij} \qquad (3)$$

The inhibitory conductance is computed via the kWTA function described in the next section, and leak is a constant.

Activation communicated to other cells (y_j) is a thresholded (Θ) sigmoidal function of the membrane potential with gain parameter γ:

$$y_j(t) = \frac{1}{\left(1 + \frac{1}{\gamma[V_m(t) - \Theta]_+}\right)} \qquad (4)$$

where $[x]_+$ is a threshold function that returns 0 if $x < 0$ and x if $X > 0$. Note that if it returns 0, we assume $y_j(t) = 0$, to avoid dividing by 0. As it is, this function has a very sharp threshold, which interferes with graded learning learning mechanisms (e.g., gradient descent). To produce a less discontinuous deterministic function with a softer threshold, the function is convolved with a Gaussian noise kernel ($\mu = 0$, $\sigma = .005$), which reflects the intrinsic processing noise of biological neurons:

$$y_j^*(x) = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-z^2/(2\sigma^2)} y_j(z - x)dz \qquad (5)$$

where x represents the $[V_m(t) - \Theta]_+$ value, and $y_j^*(x)$ is the noise-convolved activation for that value.

†Author for correspondence mfrank@u.arizona.edu.

Param	Value	Param	Value	Param	Value	Param	Value	Param	Value	Param	Value
E_l	0.15	$\overline{g_l}$	0.10	E_i	0.15	$\overline{g_i}$	1.0	E_e	1.00	$\overline{g_e}$	1.0
V_{rest}	0.15	Θ	0.25	γ	600	k_{hebb}	.01	ϵ	.001		
Striatum (k=4)		$\overline{g_l}$	1.0*	Θ, +DA	0.32*	γ	2500*	γ, +DA	10000*	γ, -DA	300*
GPi		E_l	0.28*	$\overline{g_l}$	3.0*	V_{rest}	0.26*				
GPe		E_l	0.26*	$\overline{g_l}$	1.0*	V_{rest}	0.26*				
STN		E_l	0.2*	$\overline{g_l}$	1.0*	V_{rest}	0.25*				
Thal		$\overline{g_i}$	1.7*	$\overline{g_e}$	0.5*						
Premotor (k=3)		ϵ	1e-5*	k_{hebb}	1*	V_m noise	μ =.0015	V_m noise	σ=.0015		

Table 1. BG model parameters for model described in Frank (in press), which selects among 4 responses and includes the STN. First two rows indicate standard default parameters used in 100's of simulations with Leabra software; these parameters are used in the model except where noted with an * for specialized functions of the BG layers. Striatal units have a higher firing threshold θ and higher gain γ during DA bursts ("+DA"), and lower γ during DA dips, to simulate contrast enhancement and reduction (Frank, 2005). GP and STN units have higher than normal E_l, $\overline{g_l}$ and V_{rest}, leading to tonic baseline activity in the absence of synaptic input. Thal units have high $\overline{g_i}$ and low $\overline{g_e}$, enabling a default strong inhibition from BG output and only allowing top-down excitatory activity if disinhibited, thereby serving a gating function. Premotor units have Gaussian noise added to the membrane potential, learn with a slow learning rate via purely Hebbian learning. k (kWTA) parameters are shown for striatum and premotor areas, which have within-layer lateral inhibition.

(*b*) *Connectivity and Mechanics of the BG Model*

The network's job is to select one of four responses (R1-R4), depending on the task and the sensory input. At the beginning of each trial, incoming stimuli directly activate a response in premotor cortex (PMC). However, these direct connections are not strong enough to elicit a robust response in and of themselves; they also require bottom-up support from the thalamus. The job of the BG is to integrate stimulus input with the dominant response selected by PMC, and based on what it has learned in past experience, either facilitate (Go) or suppress (NoGo) that response.

Within the overall thalamocortical circuit, there are multiple parallel sub-loops that are isolated from each other, separately modulating the different responses. This allows for the BG to *selectively* gate one response, while continuing to suppress the other(s). The striatum is divided into two distributed subpopulations. The columns on the left are "Go" units for each of the potential responses, and have simulated D1 receptors. The columns on the right are "NoGo" units, and have simulated D2 receptors. The Go columns project only to the corresponding column in the GPi (direct pathway), and the NoGo columns to the GPe (indirect pathway). Both GPe columns inhibit the associated column in GPi, so that striatal Go and NoGo activity have opposing effects on GPi. Finally, each column in GPi tonically inhibits the associated column of the thalamus, which is reciprocally connected to premotor cortex. Thus, if Go activity is stronger than NoGo activity for R1, the left column of GPi will be inhibited, removing tonic inhibition (i.e. disinhibiting) of the corresponding thalamus unit, and facilitating its execution in premotor cortex.

The above parallel and convergent connectivity is supported by anatomical evidence discussed in Frank (2005). The network architecture simply supports the existence of connections, but how these ultimately influence behavior depends on their relative strengths. The network starts off with random weights and representations in both the BG and cortical layers are learned. Input to striatal units are initialized with random synaptic weights with a Gaussian distribution, with $\mu = 0.5$, $\sigma = 0.25$. Distributed activity within each striatal column enables different Go and NoGo representations to develop for various stimulus configurations during the course of training. Therefore whereas the different columns in the striatum represent Go and NoGo units for different responses, the different rows enable units with different initial random synaptic connectivity to become specialized to respond to particular stimulus-response conjunctions.

(*c*) *STN Connectivity with Other BG and Cortical Structures*

The STN was included in the model in accordance with known constraints on its connectivity in BG circuitry (Frank, in press). First, the STN forms part of the "hyperdirect" pathway, so-named because cortical activity targets the STN, which directly excites GPi, bypassing the striatum altogether (Nambu, Tokuno, Hamada, Kita, Imanishi, Akazawa, Ikeuchi, & Hasegawa, 2000). Thus initial activation of the STN by cortex leads to an initial excitatory drive on the already tonically active GPi, effectively making the latter structure more inhibitory on the thalamus, and therefore less likely to facilitate a response. Further, the STN gets increasingly excited with increasing cortical activity. Thus, if several competing responses are activated, the STN sends a stronger "Global NoGo" signal which allows the BG system to fully consider all possible options before sending a Go signal to facilitate the most adaptive one.

Second, the STN and GPe are reciprocally connected in a negative feedback loop, with the STN exciting the GPe and the GPe inhibiting the STN (Parent & Hazrati, 1995). The connections from STN to GPe are diffuse, and therefore are not likely to be involved in suppressing a specific response. Of the STN neurons that project to GPe, the vast majority also project to GPi (Sato, Parent, Levesque, & Parent, 2000). In the model, each STN neuron receives projections from two randomly selected GPe neurons. This was motivated by data showing that multiple GPe neurons converge on a single STN neuron (Karachi, Yelnik, Tande, Tremblay, Hirsch, & Francois, 2004). In contrast, each GPe neuron receives diffuse projections from all STN neurons (but with randomly different synaptic weights).

(*d*) *Inhibition Within and Between Layers*

Inhibition *between* layers (i.e for GABAergic projections between BG layers) is achieved via simple unit inhibition, where the inhibitory current g_i for the unit is determind from the net input of the sending unit.

For *within* layer lateral inhibition (used in Striatum and premotor cortex), Leabra uses a kWTA (k-Winners-Take-All) function to achieve inhibitory competition among units within each layer (area). The kWTA function computes a uniform level of inhibitory current for all units in the layer, such that the $k+1$th most excited unit within a layer is generally below its firing threshold, while the kth is typically above threshold. Activation dynamics similar to those produced by the kWTA function have been shown to result from simulated inhibitory interneurons that project both feedforward and feedback inhibition (O'Reilly & Munakata, 2000). Thus, although the kWTA function is somewhat biologically implausible in its implementation (e.g., requiring global information about activation states and using sorting mechanisms), it provides a computationally effective approximation to biologically plausible inhibitory dynamics.

kWTA is computed via a uniform level of inhibitory current for all units in the layer as follows:

$$g_i = g^{\Theta}_{k+1} + q(g^{\Theta}_k - g^{\Theta}_{k+1}) \tag{6}$$

where $0 < q < 1$ (.25 default used here) is a parameter for setting the inhibition between the upper bound of g^{Θ}_k and the lower bound of g^{Θ}_{k+1}. These boundary inhibition values are computed as a function of the level of inhibition necessary to keep a unit right at threshold:

$$g^{\Theta}_i = \frac{g^*_e \bar{g}_e(E_e - \Theta) + g_l \bar{g}_l(E_l - \Theta)}{\Theta - E_i} \tag{7}$$

where g^*_e is the excitatory net input.

Two versions of kWTA functions are typically used in Leabra. In the kWTA function used in the Striatum, g^{Θ}_k and g^{Θ}_{k+1} are set to the threshold inhibition value for the kth and $k+1$th most excited units, respectively. Thus, the inhibition is placed to allow k units to be above threshold, and the remainder below threshold.

The premotor cortex uses the *average-based* kWTA version, g^{Θ}_k is the average g^{Θ}_i value for the top k most excited units, and g^{Θ}_{k+1} is the average of g^{Θ}_i for the remaining $n-k$ units. This version allows for more flexibility in the actual number of units active depending on the nature of the activation distribution in the layer and the value of the q parameter (which is set to default value of .6). This flexibility is necessary for the premotor units to have differential levels of activity during settling (depending on whether or not a single response has been facilitated), and also allows greater activity in high-conflict trials.

(*e*) ***Learning***

Synaptic connection weights were trained using a reinforcement learning version of Leabra. The learning algorithm involves two phases, and is more biologically plausible than standard error backpropagation. In the *minus phase*, the network settles into activity states based on input stimuli and its synaptic weights, ultimately "choosing" a response. In the *plus phase*, the network resettles in the same manner, with the only difference being a change in simulated dopamine: an increase of SNc unit firing from 0.5 to 1.0 for correct responses, and a decrease to zero SNc firing for incorrect responses (Frank, 2005).

For learning, Leabra uses a combination of error-driven and Hebbian learning. The error-driven component is the symmetric midpoint version of the GeneRec algorithm (O'Reilly, 1996), which is functionally equivalent to the deterministic Boltzmann machine and contrastive Hebbian learning (CHL), computing a simple difference of a pre and postsynaptic activation product across these two phases. For Hebbian learning, Leabra uses essentially the same learning rule used in competitive learning or mixtures-of-Gaussians which can be seen as a variant of the Oja normalization (Oja, 1982). The error-driven and Hebbian learning components are combined additively at each connection to produce a net weight change.

The equation for the Hebbian weight change is:

$$\Delta_{hebb} w_{ij} = x^+_i y^+_j - y^+_j w_{ij} = y^+_j (x^+_i - w_{ij}) \tag{8}$$

and for error-driven learning using CHL:

$$\Delta_{err} w_{ij} = (x^+_i y^+_j) - (x^-_i y^-_j) \tag{9}$$

which is subject to a soft-weight bounding to keep within the $0-1$ range:

$$\Delta_{sberr} w_{ij} = [\Delta_{err}]_+ (1 - w_{ij}) + [\Delta_{err}]_- w_{ij} \tag{10}$$

The two terms are then combined additively with a normalized mixing constant k_{hebb}:

$$\Delta w_{ij} = \epsilon [k_{hebb}(\Delta_{hebb}) + (1 - k_{hebb})(\Delta_{sberr})] \tag{11}$$

2. Additional details for LC/NE simulations

The model used for these simulations was a standard BG model selecting among two responses, with additional LC modulation. In the standard model, there are 25 input units of which a subset represent particular stimulus cues. For example, in the LC simulations stimulus A was represented by the first column of input units, whereas stimulus B was represented by the fourth column of units. The Striatum and premotor cortex have to learn to facilitate the appropriate response associated with the combination of input features. Each response in premotor cortex is represented by a column of 2 units. This allows for a more distributed representation of motor activity, because activation of R1 involves high unit activity in both motor units of the same column.

As specified in the main paper, the effect of LC activation was to dynamically modulate the gain γ of the activation function of premotor units according to the following function.

$$\gamma = 20 + 600 * (LC_{act})^2 \tag{12}$$

where LC_{act} ranges from 0 to 1 and is the mean rate-coded activation of LC units, and γ affects the activation function in equation 4. To demonstrate the effects of tonic/phasic activation on noise in premotor units, the Gaussian distribution for noise in the membrane potential of these units was increased to μ =.0035 σ = .005. The stimulus input was also delayed at the beginning of each trial by approximately 30 cycles. This was accomplished by "soft-clamping" the input layer, so that input units did not activate immediately upon trial onset but instead needed to integrate activity with a slow membrane potential time constant τ = 0.1 and input gain = 0.05. These effects combine to produce more noisy activity in premotor units prior to stimulus processing, which is more realistic than standard conditions in which each trial is initiated with no activation in the network and input stimuli can then directly activate appropriate units in the rest of the network without

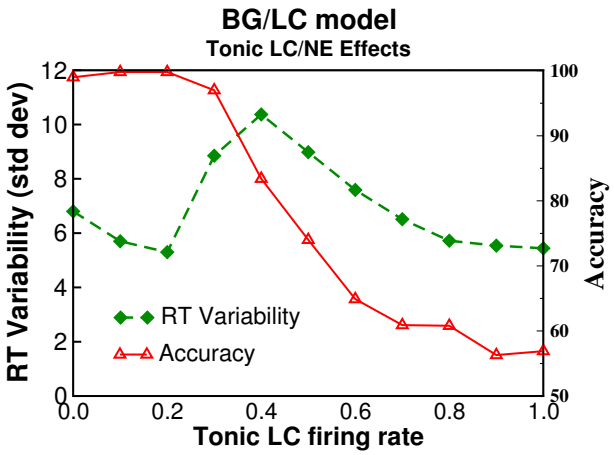

Figure 1. RT variability (standard deviation of processing cycles) and accuracy as a function of tonic LC firing rate (with no phasic response). Intermediate tonic LC levels are associated with high RT variability, while high tonic (supra-tonic) levels are associated with narrower distributions, with a cost in accuracy.

having to compete with ongoing noise. The effects in the paper show that phasic LC activity can therefore make the model more robust to noise, while tonic activity is associated with more sensitivity to noise.

Reaction times were calculated by determining when a particular response was facilitated by the BG. This amounts to recording from thalamic neurons, which facilitate the execution of responses upon being disinhibited by BG Go signals. RT's were calculated as the number of network processing cycles until one of the thalamic unit activity reach 0.9 (90% maximal value). The actual value does not much matter because thalamic neurons are completely silent until disinhibited, at which point they increase sharply in activity. Similar patterns were observed when recording from M1, but in that case it is less clear how to define when a response is actually executed, because M1 units become somewhat active from initial premotor activity prior to BG facilitation, so the thalamic units provide a more discrete RT readout.

Figure 1 shows RT variability and accuracy as a continuous function of LC tonic firing rate, in addition to the simple tonic (05) and supra-tonic (1.0) levels shown in the main paper. Overall, intermediate tonic levels are associated with high RT variability, due to noisy activation of competing responses. As tonic levels increase, noisy activity is likely to be facilitated, and RT variability decreases at the cost of increased error rates. This tradeoff demonstrates the need for an adaptive LC modulated gain.

3. References

Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and non-medicated Parkinsonism. *Journal of Cognitive Neuroscience*, *17*, 51–72.

Frank, M. J. (in press). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*.

Frank, M. J., & Claus, E. D. (2006). Anatomy of a decision: striato-orbitofrontal interactions in reinforcement learning, decision making, and reversal. *Psychol Rev*, *113*(2), 300–326.

Karachi, C., Yelnik, J., Tande, D., Tremblay, L., Hirsch, E. C., & Francois, C. (2004). The pallidosubthalamic projection: An anatomical substrate for nonmotor functions of the subthalamic nucleus in primates. *Movement Disorders*.

Nambu, A., Tokuno, H., Hamada, I., Kita, H., Imanishi, M., Akazawa, T., Ikeuchi, Y., & Hasegawa, N. (2000). Excitatory cortical inputs to pallidal neurons via the subthalamic nucleus in the monkey. *Journal of Neurophysiology*, *84*, 289–300.

Oja, E. (1982). A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, *15*, 267–273.

O'Reilly, R. C. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, *8*(5), 895–938.

O'Reilly, R. C. (2001). Generalization in interactive networks: The benefits of inhibitory competition and Hebbian learning. *Neural Computation*, *13*, 1199–1242.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

Parent, A., & Hazrati, L. (1995). Functional anatomy of the basal ganglia. II. the place of subthalamic nucleus and external pallidum in basal ganglia circuitry. *Brain Research Reviews*, *20*, 128–54.

Sato, F., Parent, M., Levesque, M., & Parent, A. (2000). Axonal branching pattern of neurons of the subthalamic nucleus in primates. *Journal of Comparative Neurology*, *424*, 142–52.