A Neurocomputational Model of Automatic Sequence Production

Sebastien Helie[1], Jessica L. Roeder[2], Lauren Vucovich[2], Dennis Rünger[3], and F. Gregory Ashby[2]

Abstract

■ Most behaviors unfold in time and include a sequence of submovements or cognitive activities. In addition, most behaviors are automatic and repeated daily throughout life. Yet, relatively little is known about the neurobiology of automatic sequence production. Past research suggests a gradual transfer from the associative striatum to the sensorimotor striatum, but a number of more recent studies challenge this role of the BG in automatic sequence production. In this article, we propose a new neurocomputational model of automatic sequence production in which the main role of the BG is to train cortical–cortical connections within the premotor areas that are responsible for automatic sequence production. The new model is used to simulate four different data sets from human and nonhuman animals, including (1) behavioral data (e.g., RTs), (2) electrophysiology data (e.g., single-neuron recordings), (3) macrostructure data (e.g., TMS), and (4) neurological circuit data (e.g., inactivation studies). We conclude with a comparison of the new model with existing models of automatic sequence production and discuss a possible new role for the BG in automaticity and its implication for Parkinson's disease. ■

INTRODUCTION

Most behaviors unfold in time and include a sequence of submovements or cognitive activities. These behaviors are often automatic and repeated daily throughout life. Although sequence learning has received much attention (e.g., Keele, Ivry, Mayr, Hazeltine, & Heuer, 2003), relatively little is known about the neurobiology of automatic sequence production. Using single-cell recordings, Miyachi et al. (Miyachi, Hikosaka, & Lu, 2002; Miyachi, Hikosaka, Miyashita, Kárádi, & Rand, 1997) found task-sensitive neurons in the associative striatum during early sequence learning and in the sensorimotor striatum at a later stage. This result led to the proposal that automaticity in sequence production is characterized by a gradual transfer from the associative to the sensorimotor striatum (Nakahara, Doya, & Hikosaka, 2001). Specifically, a visual loop that connects pFC with the associative striatum and a motor loop that connects motor cortex with the sensorimotor striatum learn in parallel, but the visual loop learns more quickly than the motor loop. A coordinator located in the pre-SMA controls the contribution of each loop to the motor response and gradually transfers control from the visual loop to the more robust motor loop. The Nakahara et al. model accounts for many phenomena in early and late sequence production.

More recently, Desmurget and Turner (2010) injected muscimol into the internal segment of the globus pallidus (GPi) of monkeys that had received extensive practice in a sequence learning task to functionally disconnect the BG from frontal cortex. The results showed that, although movement kinematics were impaired, sequence knowledge was not, suggesting that the BG contribute to motor execution in automatic sequence production but are not necessary for the sequencing or storage of the overlearned sequence (Desmurget & Turner, 2010). Sequence production in the Nakahara et al. (2001) model always relies on the BG, so it could not easily account for this result.

The results of Desmurget and Turner, however, are consistent with a neurobiological model of automaticity in perceptual categorization called SPEED (Ashby, Ennis, & Spiering, 2007). Many parallels have been drawn between perceptual categorization and sequence learning, and the brain areas involved seem to overlap (Ashby, Alfonso-Reese, Turken, & Waldron, 1998). We propose a neurobiological model of automatic sequence production inspired by the SPEED model of automatic categorization (Ashby et al., 2007). The BG in the proposed model learn and produce individual responses in early sequence production and trigger Hebbian learning within premotor areas. The sequence is learned and stored within premotor areas so that automatic sequence production becomes purely cortical and no longer relies on the BG.

[1]Purdue University, [2]University of California, Santa Barbara, [3]University of Southern California

METHODS

Previous research shows that frontal cortex is connected to the BG by a number of closed loops (Middleton & Strick, 1997). In each case, the most prominent input to a given striatal channel is derived from the same area of cortex targeted by the channel's output projections (Strick, Dum, & Picard, 1995). This model focuses on the motor loop (Alexander, DeLong, & Strick, 1986), illustrated in Figure 1. The main cortical area contributing to the motor loop is the SMA, which receives massive projections from posterior parietal cortex (PPC) and projects to the putamen and primary motor cortex (M1; Flaherty & Graybiel, 1994; Lu, Preston, & Strick, 1994). Projections from PPC to SMA contain spatial information useful in designing motor plans essential for motor sequence production. The putamen, in turn, sends topographically organized projections to the ventrolateral part of the GPi. The GPi projects to the ventrolateral nucleus of the thalamus (VL), which projects back to the SMA (Akkal, Dum, & Strick, 2007), thereby closing the loop. The closed loop shown in Figure 1 receives visual/spatial input from PPC and outputs a motor response through M1. The next two sections present a conceptual description and a formal description of the neural architecture, respectively. Readers not interested in implementation details can focus on the conceptual description and skip the formal presentation.

Computational Model Architecture

The model architecture was inspired by the motor loop of Alexander et al. (1986), but the model assigns a new role to the BG in automatic sequence production. Specifically, the model differs from Nakahara et al. (2001) in that the BG do not store sequences but, instead, train cortical–cortical connections that mediate sequence production. This new role is supported not only by the

Desmurget and Turner (2010) data described above but also by evidence suggesting that synaptic plasticity in BG and cortex follow different learning rules (Doya, 2000; for a review, see Helie, Ell, & Ashby, 2015).

There is convincing evidence that synaptic plasticity at cortical–cortical synapses follows Hebbian-like learning rules (see, e.g., Doya, 2000). Ashby et al. (2007) further proposed that one role of the slower subcortical path through the BG is to learn to activate the correct post-synaptic target on the cortical–cortical path, which allows the appropriate cortical–cortical synapses in premotor cortex to be strengthened via Hebbian learning.[1] The cortical–cortical connections gradually become strong enough to depolarize the output neuron without help from the subcortical loop. Thus, according to this model, the development of motor skill automaticity is a gradual process in which control is passed from subcortical learning systems to cortical–cortical connections between separate units in posterior and premotor cortex (Hua & Houk, 1997).

Figure 2 shows the connectivity of the model for a four-element sequence. The model input is located in Layer V of the SMA (which projects primarily to the striatum), whereas the model output is located in Layer IV of the SMA (which receives projections from the thalamus; Gerfen & Bolam, 2010). In the SMA and putamen, inhibitory interneurons implement soft winner-takes-all dynamics. Synaptic plasticity between SMA Layer V and the putamen allows for learning stimulus → response associations using dopamine (DA)-mediated reinforcement learning (Ashby et al., 2007). In addition, all excitatory cortical–cortical synapses display synaptic plasticity using a Hebbian learning rule (Ashby et al., 2007).

Initially, the correct premotor target in SMA is activated via sensory input (e.g., from PPC). During this stage of learning, each motor response is triggered by a visual cue and processed through the BG. Closed loop projections from SMA through the putamen and back to SMA begin building connections from SMA Layer IV to Layer V

Figure 1. The "motor loop" used in the present model. The (visual/spatial) input to the model is located in the PPC, and the (motor) output of the model is located in the primary motor area.
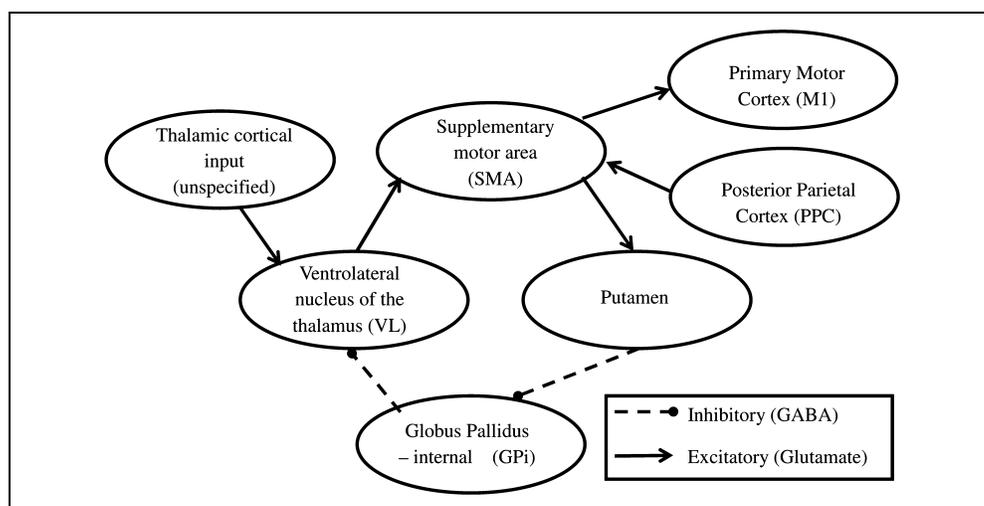
Figure 2. Neural architecture of the model. The BG loop connectivity in the figure is one-to-one, so activating Location 1 in SMA Layer V will produce Response 1 in SMA Layer IV. At the cortical level, connectivity from SMA Layer V toward SMA Layer IV implements automaticity, and these connections are learned using Hebbian learning. Automatic connections are learned slowly and allow for automatic responding without using the BG. Connectivity from SMA Layer IV toward SMA Layer V implements sequence knowledge. Sequence knowledge within the SMA is also learned using Hebbian learning. These connections allow for a response at time *t* to serve as a cue for

the response at time *t* + 1—thus removing the reliance on the visual cue. The example sequence shown is 1-2-3-4. Note that connectivity representing sequence knowledge is learned more quickly than connectivity allowing for automatic responding. When both sequence knowledge and automatic connectivity is established, the BG are no longer required to produce a response.
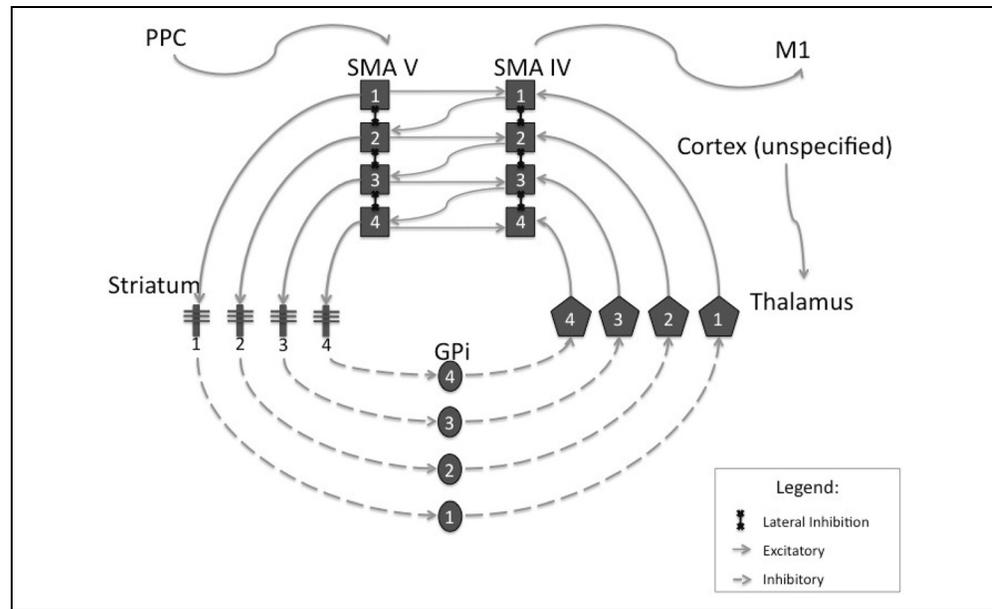
via Hebbian learning. This is because the PPC consistently activates the neurons representing the location of the stimulus on trial *t* + 1 (located in SMA Layer V) shortly after the response to trial *t* is initiated by the neurons located in SMA Layer IV. When these connections become strong enough, the model should be able to initiate the next motor response independently of visual input, as the preceding motor response will be a sufficient cue for the next response. Hebbian learning strengthens any active synapses, so if the BG can consistently activate the correct cortical units, Hebbian learning will gradually strengthen their interconnections. In the proposed model, these connections from SMA Layer IV to SMA Layer V encode the sequence information. It is important to note, however, that the proposed model does not learn the timing of events—only their ordering. For example, the model could learn to predict the next note in a piece of music, but not the note duration. Learning timing would likely require additional cognitive and neural mechanisms.

In addition to sequence knowledge, automatic sequence production is also learned via Hebbian modification. Automatic sequence production is encoded in the projections from SMA Layer V to SMA Layer IV. This is because the SMA unit initiating the response at time *t* in SMA Layer IV consistently fires following activation of the response location at time *t* in SMA Layer V. These connections allow activation of SMA Layer IV to be driven directly by SMA Layer V. Ultimately, the BG are no longer needed to produce the sequence. Instead, a subnetwork consisting of SMA units and their cortical–cortical connections becomes sufficient for the model to produce the correct motor sequence.

Neural Model

This section describes the equations used to implement the proposed model. The activation in each unit[2] was simulated using the Izhikevich (2007) model:

$$C\frac{dV(t)}{dt} = \beta + \gamma[V(t)-V_r][V(t)-V_t]-U(t)+I(t)+\varepsilon(t)$$
$$\frac{1}{\varphi}\frac{U(t)}{dt} = \kappa[V(t)-V_r]-U(t) \qquad (1)$$

where $V(t)$ is the intracellular voltage at time t, C is the membrane capacitance, β is related to the baseline firing rate, γ is related to the rheobase current, V_r is the resting membrane potential, V_t is the instantaneous threshold potential, $I(t)$ is the input at time t, and $\varepsilon(t) \sim N(0, \sigma)$ is randomly generated Gaussian noise at time t. $U(t)$ is an abstract regulatory term that is meant to describe slow recovery after an action potential. Like γ, κ is related to the neuron rheobase current, and ϕ is the recovery time. When $V(t)$ reaches 35 mV, a spike is recorded and $V(t)$ is reset to V_{reset}. In addition, $U(t)$ is reset to $U(t) + U_0$. The different parameters listed above are the only difference between the units in different regions of the model. The parameter values for each brain region are listed in Table 1. All these parameters were estimated using patch-clamp experiments and taken from Izhikevich (2007), unless stated otherwise.

The model of neural propagation focused on the temporal delays of spike propagation and the temporal smearing that occurs at the synapse. A standard solution is to use an α-function (e.g., Rall, 1967). Every time the

Table 1. Fixed Unit Parameter Values for Each Brain Region

Parameter	SMA	Putamen	GPi	VL
C	100	50	10	200
β	25	25	71	25
γ	0.7	1	0.7	1.6
V_r	−60	−80	−60	−65
V_t	−40	−25	−40	−63
σ	5	5	0	5
ϕ	0.03	1	–	0.01
κ	−2	−20	–	{0, 15}
V_{reset}	−50	−55	−50	−60
U_0	100	150	–	10

The same parameter values were used for SMA Layers IV and V. For the units in VL, if $V(t) < V_r$, $\kappa = 15$; otherwise, $\kappa = 0$.

presynaptic unit spikes, the following input is delivered to the postsynaptic unit (with spiking time $t = 0$):

$$\alpha(t) = \frac{t}{\lambda} e^{\left(1-\frac{t}{\lambda}\right)} \qquad (2)$$

In Equation 2, λ is scale-free and can be fixed to model any desired temporal delay (i.e., it does not correspond to any physical measure, but larger λ is related to longer delays). In the current model, $\lambda = 100$. If a second spike occurs before $\alpha(t)$ decays to zero, then a second α-function is added to the residual $\alpha(t)$.

The input to each region is described in following subsections. All the circuit parameters are listed in Table 2. These parameter values were based on biological constraints (e.g., excitatory vs. inhibitory) and prior experience with the model in our laboratory.

SMA Layer V

SMA Layer V is the input layer of the model. All SMA Layer V units were modeled after pyramidal neurons. The input to SMA unit i in Layer V is

$$I_{\text{SMA5},i}(t) = \text{PPC}_i(t) - \sum_{j \neq i} w^{5 \to 5} \alpha_{\text{SMA5},j}(t) + \varsigma \sum_{j} w_{j,i}^{4 \to 5} \alpha_{\text{SMA4},j}(t-1)$$

where $\text{PPC}_i(t)$ is the visual input simulated using a square function [$\text{PPC}_i(t) = 6000$, if the visual signal for cell i is "ON" and 0 otherwise]; $w^{5 \to 5}$ is the (fixed) inhibitory connection weight between SMA Layer V units; $\alpha_{\text{SMA5},j}(t)$ is the output α-function of SMA Layer V unit j at time t; $w_{j,i}^{4 \to 5}$ is the (learned) connection weight between units j and i in SMA Layers IV and V, respectively; $\alpha_{\text{SMA4},j}(t-1)$ is the

output α-function of SMA Layer IV unit j at time $t-1$; and ζ is a correction term reducing the input from SMA Layer IV after an error. Specifically, $\zeta = 0.5^{[\text{error\_duration} - \text{last\_error}]+}$, where error_duration is the duration of the error signal (in trials), last_error is the number of trials since the last error was made, and the function $[g(t)]^+$ equals $g(t)$ when $g(t) > 0$ and 0 when $g(t) \leq 0$. In words, the first term represents the visual input to the model, the second term represents lateral inhibition within SMA Layer V, and the last term represents the sequence knowledge learned by the model. In all the simulations, error_duration = 5.

Putamen

The putamen is the input structure of the BG in the motor loop of Alexander et al. (1986). All putamen units were based on medium spiny neurons. The proposed model focuses on automaticity, so it is assumed that the stimulus → response mapping is either already learned or trivial (e.g., the leftmost button corresponds to the leftmost location). As such, SMA Layer V units were preconnected to the putamen using a one-to-one mapping, and all SMA → putamen weights were equal. The input to putamen unit i is

$$I_{\text{put},i}(t) = w^{5 \to \text{put}} \alpha_{\text{SMA5},i}(t) - \sum_{j \neq i} w^{\text{put} \to \text{put}} \alpha_{\text{put},j}(t)$$

where $w^{5 \to \text{put}}$ is the fixed connection weight between SMA Layer V and the putamen, $\alpha_{\text{SMA5},i}(t)$ is the output α-function of SMA Layer V unit i at time t, $w^{\text{put} \to \text{put}}$ is the fixed inhibitory connection weight between the putamen units, and $\alpha_{\text{put},j}(t)$ is the output α-function of putamen unit j at time t.

Table 2. Fixed Circuit Parameters in the Model

Parameter	Value
WBASE	1
$w^{\text{GPi} \to \text{VL}}$	100
$w^{\text{VL} \to \text{SMA4}}$	80
$w^{\text{put} \to \text{put}}$	10
$w^{5 \to 5}$	10
$w^{4 \to 4}$	10
$w^{5 \to \text{put}}$	8
$w^{\text{put} \to \text{GPi}}$	0.7
$W_{\max}$	400
θ_{AMPA}	150
θ_{NMDA}	1500

WBASE = initial connection weights between the SMA layers (i.e., $w_{j,i}^{4 \to 5}$ and $w_{j,i}^{5 \to 4}$) before learning takes place. The remaining symbols are defined in the text.

In words, the first term represents the SMA Layer V input to the putamen and the second term represents lateral inhibition within the putamen.

Internal Segment of the Globus Pallidus

The GPi is the output structure of the BG in the motor loop of Alexander et al. (1986). GPi neurons have a high baseline firing rate, which is essential to the proposed network dynamics, but the exact firing pattern and timing is not because the present model will not be used to account for single-cell recording data in the GPi or abnormal GPi firing patterns (e.g., synchrony). Thus, we used a regular quadratic integrate-and-fire model. This simpler model is identical to Equation 1, except that only the top equation is evaluated (and the $U(t)$ term is removed). In the model, putamen units were connected to the GPi using a one-to-one mapping, and all the putamen → GPi weights were equal. The input to GPi unit i is

$$I_{\mathrm{GPi},i}(t) = -w^{\mathrm{put}\to\mathrm{GPi}}\alpha_{\mathrm{put},i}(t)$$

where $w^{\mathrm{put}\to\mathrm{GPi}}$ is the (fixed) inhibitory connection weight between the putamen and the GPi and $\alpha_{\mathrm{put},i}(t)$ is the output α-function of putamen unit i at time t.

Ventrolateral Nucleus of the Thalamus

VL is a relay structure between the BG and the SMA in the motor loop of Alexander et al. (1986). All VL units in the model were thalamocortical relay neurons. The patch-clamp parameters from Izhikevich (2007) were slightly tweaked to facilitate compatibility with the GPi unit described above (see Table 1). In the model, GPi units were connected to VL using a one-to-one mapping, and all the GPi → VL weights were equal. In addition, all VL neurons in the model receive weak excitatory input from nonspecified cortical areas for the duration of the stimulus presentation. The input to VL unit i is

$$I_{\mathrm{VL},i}(t) = -w^{\mathrm{GPi}\to\mathrm{VL}}\alpha_{\mathrm{GPi},i}(t) + E(t)$$

where $w^{\mathrm{GPi}\to\mathrm{VL}}$ is the fixed inhibitory connection weight between the GPi and VL, $\alpha_{GPi,i}(t)$ is the output α-function of GPi cell i at time t, and $E(t)$ is the nonspecific cortical input to thalamic cells simulated using a square-wave function. $E(t) = 300$ if the visual signal for cell i is "ON" and 0 otherwise.

SMA Layer IV

SMA Layer IV is the output layer of the model. All Layer IV SMA units in the model were pyramidal neurons

identical to the Layer V SMA units (e.g., same parameter values). In the model, VL units are connected to SMA Layer IV by one-to-one mapping, and all the VL → SMA weights are the same. The input to SMA cell i in Layer IV is

$$I_{\mathrm{SMA4},i}(t) = w^{\mathrm{VL}\to\mathrm{SMA4}}\alpha_{\mathrm{VL},i}(t) - \sum_{j\neq i} w^{4\to4}\alpha_{\mathrm{SMA4},j}(t)$$
$$+ \sum_{j} w_{j,i}^{5\to4}\alpha_{\mathrm{SMA5},j}(t)$$

where $w^{\mathrm{VL}\to\mathrm{SMA4}}$ is the fixed connection weight between VL and SMA Layer IV units, $\alpha_{\mathrm{VL},i}(t)$ is the output α-function of VL cell i at time t, $w^{4\to4}$ is the fixed inhibitory connection weight between SMA Layer IV units, $\alpha_{\mathrm{SMA4},j}(t)$ is the output α-function of SMA Layer IV unit j at time t, $w_{j,i}^{5\to4}$ is the (learned) connection weight between units j and i in SMA Layers V and IV, respectively, and $\alpha_{\mathrm{SMA5},j}(t)$ is the output α-function of SMA Layer V unit j at time t. The first term models subcortical input to the SMA, the second term models lateral inhibition within SMA Layer IV, and the last term models the automatic cortical–cortical sequence production process learned by the model.

Decision Process

The final component needed by the model is a decision rule. We used the following simple decision model: At every time step (msec), if $\mathrm{Max}_j\left[\alpha_{\mathrm{SMA4},j}(t)\right]\geq\tau$, set the model response to the most active unit in SMA Layer IV and set the RT to t. Otherwise, calculate another millisecond of model time. As described above, $\alpha_{\mathrm{SMA4},j}(t)$ is the output α-function of SMA Layer IV cell j at time t and τ is the response threshold. In all simulations, $\tau = 7.18$.

Synaptic Plasticity

Two sets of connections are learned: sequence knowledge from SMA Layers IV to V that allows prediction of the next element in the sequence and automatic response production procedures from SMA Layers V to IV that allow for purely cortical processing. Although both sets of connections are learned using the same Hebbian rule, the automatic response production procedures are learned more slowly than the sequence knowledge (i.e., with smaller learning rate parameter values). This is because the time between stimulus presentation and response production at the beginning of learning is typically longer than the time between a response and presentation of the next stimulus. Assuming decay on the synaptic trace (not modeled), sequence knowledge will be learned more quickly than automatic connections (but they converge with practice because of the rate limiting term; see below).

Sequence knowledge between SMA Layers IV and V is learned using

$$
\begin{aligned}
w_{i,j}^{4\to5}(n+1) = {} & w_{i,j}^{4\to5}(n) + \eta_{S,LTP} \int_t \alpha_{SMA4,i}(n-1,t)dt \\
& \left[\int_t \alpha_{SMA5,j}(n,t)dt - \theta_{NMDA}\right]^+ \left[w_{max} - w_{i,j}^{4\to5}(n)\right] \\
& - \eta_{S,LTD} \int_t \alpha_{SMA4,i}(n-1,t)dt \left[\theta_{NMDA} - \int_t \alpha_{SMA5,j}(n,t)dt\right]^+ \\
& \left[\int_t \alpha_{SMA5,j}(n,t)dt - \theta_{AMPA}\right]^+ w_{i,j}^{4\to5}(n) \quad (3)
\end{aligned}
$$

where $w_{i,j}^{4\to5}(n)$ is the connection weight between SMA Layer IV unit i and SMA Layer V unit j on trial n; $\alpha_{SMA4,i}(n-1,t)$ is the output α-function of SMA Layer IV unit i at time t on trial $n-1$; $\alpha_{SMA5,j}(n,t)$ is the output α-function of SMA Layer V unit j at time t on trial n; $\eta_{S,LTP}$ and $\eta_{S,LTD}$ are the learning rates for long-term potentiation (LTP) and long-term depression (LTD) of sequence knowledge, respectively; θ_{NMDA} denotes the threshold for strong activation of the NMDA receptors and θ_{AMPA} denotes the activation threshold of the AMPA receptors (with $\theta_{NMDA} > \theta_{AMPA}$); w_{max} is the maximum allowable weight value; and the function $[g(t)]^+$ equals $g(t)$ when $g(t) > 0$ and 0 when $g(t) \le 0$.

Note that $\int_t \alpha_{SMA4,i}(n-1,t)dt$ represents the integrated presynaptic activation in SMA Layer V on trial n whereas $\int_t \alpha_{SMA5,j}(n,t)dt$ represents the integrated postsynaptic activation in SMA Layer V on trial n. Hence, the first term of Equation 3 describes conditions under which LTP occurs (i.e., strong pre- and postsynaptic activation), whereas the second term describes conditions under which LTD occurs (i.e., strong presynaptic activation with weak postsynaptic activation). This model assumes that the change in synaptic strength is proportional to the product of the pre- and postsynaptic activations (and the final rate limiting term that prevents the strength of the synapse from exceeding w_{max}).

Automatic response production connectivity between SMA Layers V and IV is learned using the same learning rule as sequence knowledge:

$$
\begin{aligned}
w_{i,j}^{5\to4}(n+1) = {} & w_{i,j}^{5\to4}(n) + \eta_{A,LTP} \int_t \alpha_{SMA5,i}(n,t)dt \\
& \left[\int_t \alpha_{SMA4,j}(n,t)dt - \theta_{NMDA}\right]^+ \left[w_{max} - w_{i,j}^{5\to4}(n)\right] \\
& - \eta_{A,LTD} \int_t \alpha_{SMA5,i}(n,t)dt \left[\theta_{NMDA} - \int_t \alpha_{SMA4,j}(n,t)dt\right]^+ \\
& \left[\int_t \alpha_{SMA4,j}(n,t)dt - \theta_{AMPA}\right]^+ w_{i,j}^{5\to4}(n) \quad (4)
\end{aligned}
$$

where $w_{i,j}^{5\to4}(n)$ is the connection weight between SMA Layer V unit i and SMA Layer IV unit j on trial n; $\eta_{A,LTP}$

and $\eta_{A,LTD}$ are the learning rates for LTP and LTD of automatic response production, respectively (these are smaller than the corresponding values in Equation 3); and the remaining symbols are as described above. Again, the first term of Equation 4 describes conditions under which LTP occurs (i.e., strong pre- and postsynaptic activation), whereas the second term describes conditions under which LTD occurs (i.e., strong presynaptic activation with weak postsynaptic activation).

RESULTS

The proposed model was evaluated using numerical simulations of four different types of data from human and nonhuman animals: behavioral RTs, single-unit recordings, TMS, and neurological circuit data (i.e., inactivation studies). In all simulations, six units were included in each brain region, and voltage was estimated in every unit for each millisecond of a 3000-msec trial (Equation 1). On each trial, the first 1000 msec were burnout to stabilize the network, and the stimulus was presented 400 msec after the burnout period. The response and RT were determined as described in the Decision Process subsection above. The output of SMA Layer IV calculated after a response had been output by the model in trial t served as input to SMA Layer V on trial $t + 1$.[3] At the end of each trial, the learning rules were applied, and the membrane potential of each cell was reset. Each condition in each simulation was repeated 100 times, and the results are the average of the simulations.

The goal of the simulations is to test the neural architecture of the model—not parameter optimization. Hence, only five free parameters were varied across all of the simulations: the four learning rate parameters of Equations 3 and 4 were set to different values for modeling data from humans and monkeys (one value per parameter for humans and another for monkeys), and the NMDA threshold was increased to simulate TMS application. The fixed circuit parameters are shown in Table 2.

The Discrete Sequence Production Task

In the discrete sequence production (DSP) task, participants see a stimulus and press a button corresponding to the stimulus location. Similar to the serial RT task (SRTT), stimulus locations sometimes follow a predetermined sequence, so sequence learning is associated with RT speed-up. Unlike the SRTT, participants are allowed to produce a response before the stimulus appears on the screen, producing negative RTs. Matsuzaka, Picard, and Strick (2007) reported results from an extensive DSP task in which monkeys were trained with either pseudorandom or repeating sequences. The repeating sequences always included three separate button presses. Some sequences were practiced for more than 2 years. RTs for the repeating sequences substantially declined with practice, and the monkeys consistently learned to perform these sequences

without visual cues (i.e., predictively). Specifically, after a sequence had been practiced for a median of 22 weeks, the monkeys made more than 90% of their responses predictively.

Behavioral Data

Figure 3A shows median RTs in the DSP task for random (circles) and repeating (pluses) sequences for each of more than 400 experimental sessions for monkey GE (Matsuzaka et al., 2007). The dotted line in Figure 3A denotes stimulus onset time (RT = 0), so all responses below this line were predictive. Note that RTs for both repeating and random sequences decreased with practice, but only responses from repeated sequences became pre-

Figure 3. Median RTs in the discrete sequence production task. (A) Median RTs from monkey GE (Matsuzaka et al., 2006). (B) Mean median RTs from 100 simulations of the model. In both panels, O represents random sequences, and + represents repeating sequences.

dictive. When locations were random, RTs decreased by about 400 msec over the course of training, but RTs for the repeating sequence decreased by about 900 msec.

We simulated the DSP task[4] for 365,000 trials with random locations and 290,000 trials with the same repeated sequence used in Matsuzaka et al. (2007). The key results from the monkey data are that RTs decrease for both random and repeated sequences, and RTs for the repeated sequence become negative with extensive practice. The simulated median RTs are shown in Figure 3B. As can be seen, simulated RTs decreased by about 400 msec for the random locations and by about 850 msec for repeating sequences. Similar to the monkey data, the responses became predictive with extended practice, but only for repeating sequences, reproducing the key results from the monkey data. However, the monkey learning curve is somewhat steeper than the simulation learning curve; the monkey learning curve is more similar to a power function, whereas the model learning curve is more similar to an exponential function. This is not surprising, as Brown and Heathcote (2003) showed that averaging a number of exponential learning processes may result in an overall power function RT learning curve. In the proposed model, the focus is on the transition between sequence learning and automatic sequence production (which happens late in learning), and the simulation includes only one learning process (hence the exponential shape). The monkey likely used more than one learning process, hence the power function. Yet, although this early learning stage was not the focus of the model, the proposed model captured the data very well and the numerical fit of the model to the monkey data is $r^2 = .876$. Hence, the proposed model was successful at reproducing an exhaustive behavioral data set from the DSP task.

Single-cell Recordings in the SMA

Another important test of the model is whether its SMA units are qualitatively consistent with spiking behavior recorded from real SMA neurons during sequence production. Shima and Tanji (2000) recorded from SMA neurons in a three-element sequence learning task. They found that many task-related neurons showed transition selectivity. Specifically, these neurons fired during the interval between two specific sequence elements and remained silent during other movement transitions.

Although the task used by Shima and Tanji (2000) differed slightly from the DSP task, these differences are irrelevant for our purpose because the model does not simulate movements. Therefore, we observed the firing patterns of SMA output (Layer IV) units in the proposed model at the end of the DSP simulations. Figure 4 shows the firing pattern for a complete cycle through the learned sequence of a typical simulated SMA Layer IV neuron responsible for producing a response in the third location. As can be seen, the firing pattern of the model's SMA unit corresponds well with the pattern of firing reported by

Figure 4. Simulated cell activation of the SMA Layer IV unit representing the response in Position 3 during the DSP task. A complete cycle through the sequence is presented, and the response at the third position is indicated by the vertical dashed line.

Shima and Tanji.[5] Specifically, this neuron fires during the specific transition between the second and third elements (with response "3" being output at the vertical dashed line) and is silent for the remainder of the sequence. Other simulated SMA units show similar patterns of activity for other transitions. Figure 4 shows that the SMA, the key area in the model, has a transitional role that is consistent with the role ascribed to existing single-unit recording SMA data in a similar task.

Using TMS to Assess Sequence Representation in Cortex

Verwey, Lammens, and van Honk (2002) used TMS to explore the role of the SMA in the DSP task. Human participants were trained on two discrete six-element sequences, and repetitive TMS or sham stimulation was applied after the first training block. Results are shown in Figure 5A. Note that rTMS of the SMA interfered with the normal RT speed-up that occurs with practice. Verwey et al. concluded that after moderate practice, SMA mediates the execution of simple motor sequences.

The Verwey et al. results and interpretation are consistent with the role of the SMA in the proposed model, as learned sequences are represented within the SMA. To test this formally, we simulated the model's performance in the Verwey et al. experiment.[6] After the first block of training, we applied rTMS to the model using the TMS model of Huang, Rothwell, Chen, Lu, and Chuang (2011). According to this model, the effect of rTMS depends on differences in the pattern of Ca^{2+} influx through postsynaptic NMDA receptors that are induced by different stimulation frequencies. Accordingly, we simulated the rTMS pulses by raising the NMDA threshold in the model to $\theta_{NMDA} = 2500$, which made it more difficult for the post-

synaptic activation to rise above the NMDA threshold and trigger LTP. This caused the connections between the input units of the SMA and the output units to not be strengthened as much, so little decrease in RT was
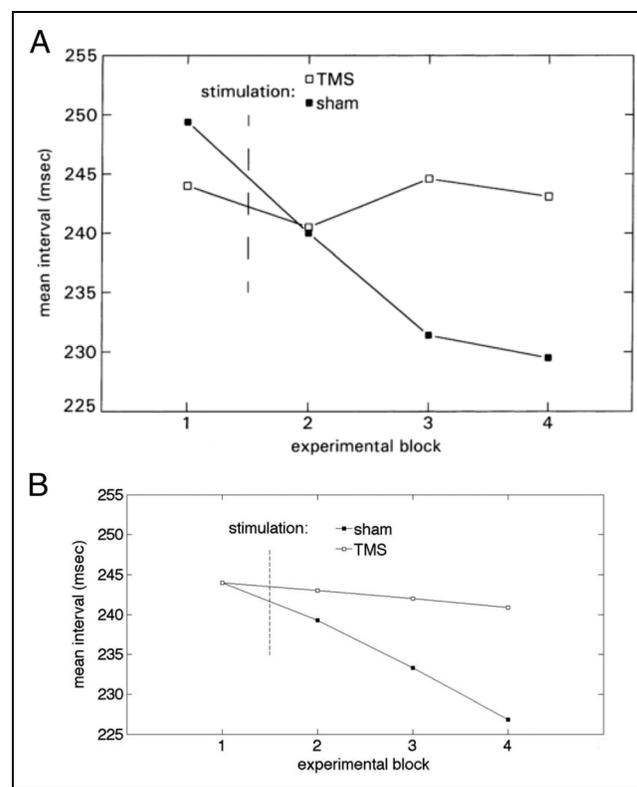
Figure 5. The effect of rTMS applied to the SMA in sequence learning. (A) Results collected by Verwey et al. (2002). (Reprinted with permission of Elsevier.) (B) Mean simulated results obtained with the proposed model (100 simulations).

observed. We simulated the sham condition by leaving the NMDA threshold at its baseline value in the model.

The simulated results are shown in Figure 5B. As can be seen, simulated rTMS applied to the SMA impaired normal learning, because the connections between the output and input layers of the SMA were not strengthened as much. As a result, the RT speed-up was diminished relative to the sham condition, as in the human data, reproducing the qualitative pattern of results obtained by Verwey et al. (2002). The numerical fit was also very good ($r^2 = .858$). This supports the hypothesis that the role of the SMA in the model is consistent with the role of the SMA in human sequence learning.

The Role of the BG in Automatic Sequence Production

Desmurget and Turner (2010) sought to test the hypothesis that sequence production becomes independent of the BG after overtraining. To test this hypothesis, monkeys were trained using both a random and a repeating sequence, and muscimol was injected into their GPi late in learning. Inactivating the GPi with muscimol functionally disconnects the BG from the thalamus and cortex. If automatic sequence production relies on the BG, then the muscimol injection should disrupt performance with the overtrained sequence. Interestingly, automatic sequence production was not disrupted by the muscimol injections. Instead, Desmurget and Turner showed that the monkeys had a similarly modest slowdown of RTs for both random and overlearned sequences, suggesting that automatic sequence production is not dependent on the BG.

The goal of this simulation is to show that, as in Desmurget and Turner's data, the model is not significantly impaired by inactivation of the GPi late in learning for either repeated or random sequences. This is because the model can account for continued sequence production

that is independent of the BG. In an overtrained sequence, the sequence representation is contained entirely within the SMA; the SMA Layer IV units have strong connections to the next input SMA Layer V unit in the sequence (motor cue), and the model is able to continue to respond without a functioning BG or external stimulation because of the strong connectivity between SMA Layer V (input) and SMA Layer IV (output) units (automatic response production).

To illustrate this explanation, the proposed model was trained for 200,000 trials with the same sequence used in Desmurget and Turner's experiment or a random sequence using the monkey learning rates described in Note 4. After 150,000 trials, the connection weight between the GPi and the thalamus was reduced by 90% to simulate the muscimol injection. The results are shown in Figure 6. As can be seen, the model responded similarly to the monkeys for both the random and overlearned sequences. In both cases, the lesion did not affect sequence production. The model can account for the lack of any appreciable effect of the muscimol injections on random sequences because with extensive practice the stimulus–response associations become automatic. The direct cortical–cortical connections from SMA Layer V to SMA Layer IV become strong enough that the BG are no longer necessary to generate a response, and therefore, any damage to subcortical structures will have minimal or no detrimental effect on RT (as shown by the full line). In addition, the muscimol injection has little effect on over-learned sequences (dashed line) because most of the speed-up is caused by cortical knowledge (i.e., the sequence knowledge encoded between SMA Layers IV and V). The model makes this prediction because the sequence can be produced independently of the BG loop after extensive training. Hence, inactivation of other regions in the BG loop (e.g., the sensorimotor striatum) should have a similar effect. Future work is needed to test this prediction.

Figure 6. Mean RT of 100 model simulations in sequence learning with a lesion of the GPi. The full line shows a late lesion with a random sequence and the dashed line shows a late lesion in a repeating sequence (as in Desmurget & Turner, 2010). The dotted line shows the predicted effect of an early lesion with a repeated sequence.

Table 3. Parameter Sensitivity Analysis (r^2)

Parameter	Matsuzaka et al. (2006)	Verwey et al. (2002)	Desmurget and Turner (2010)
$\eta_{A,LTP}$ + 10% (Equation 4)	.9977	.8421	.9885
$\eta_{A,LTP}$ − 10% (Equation 4)	.9976	.8508	.9920
$\eta_{S,LTP}$ + 10% (Equation 3)	.9976	.8491	.9927
$\eta_{S,LTP}$ − 10% (Equation 3)	.9971	.8457	.9944
$\eta_{A,LTD}$ + 10% (Equation 4)	.9977	.8631	.9876
$\eta_{A,LTD}$ − 10% (Equation 4)	.9981	.8336	.9923
$\eta_{S,LTD}$ + 10% (Equation 3)	.9981	.8242	.9884
$\eta_{S,LTD}$ − 10% (Equation 3)	.9979	.8560	.9889
θ_{NMDA} + 10% (Equations 3, 4)	.9977	.8341	.9920
θ_{NMDA} − 10% (Equations 3, 4)	.9978	.8310	.9885

Each simulation was repeated 100 times for each parameter.

The model also makes an interesting prediction that is currently untested: Although there is no effect of lesioning the GPi on repeating sequence production after extensive training, there should be a substantial effect of lesioning the GPi on sequence learning during early training. Because two-factor learning is slow, it takes significant training for cortex to become independent of the BG. Therefore, when the connections from subcortical structures are lesioned early in training, the projections from SMA Layer V to Layer IV are still weak, and it will take more time to depolarize the postsynaptic output unit. This will slow RTs, with the degree of slowing proportional to the amount of training that occurs before the lesion. This will also slow the strengthening of the predictive connections from SMA Layer IV to the SMA Layer V unit responsive to the next element of the sequence (resulting in impaired sequence learning). The simulated results of an early GPi lesion with a repeated sequence are also shown in Figure 6 (after 50,000 trials; dotted line). Note that the slowdown (about 75 msec) caused by the simulated lesion is more pronounced in early training, and responses never become predictive, even after 200,000 trials of practice. With no GPi lesion, responses become predictive after about 90,000 trials. Future work is needed to test this prediction. Thus, in addition to accounting for the failure of GPi inactivation to affect late sequence production (Desmurget & Turner, 2010), the model makes new (and untested) predictions.

Parameter Sensitivity Analysis

As mentioned above, the learning parameters (from Equations 3–4) were set to reasonable values based on previous work in our laboratory but were not optimized. The model's performance is relatively robust to deviations in exact numerical values, so all of the predictions derived in this article follow from the general architecture of the model and depend only minimally on our ability to find the optimal set of parameter values.

To substantiate these claims formally, we completed a sensitivity analysis for each simulation included in the article. The analysis proceeded as follows. For each of the five learning parameters in Equations 3–4, we successively changed the parameter estimate from the value used to generate the predictions shown in the figures by ±10%. After each change, we reran the model 100 times in the same conditions and calculated the r^2 between the averages of the new and old simulations. The results are shown in Table 3. As can be seen, the r^2 values are very high for all parameters in all simulations. The lowest r^2 is .824, and the average r^2 is .944. Thus, the proposed model is robust to changes in parameter values, confirming that our simulation results rely on the proposed new architecture and not the specific values assigned to the parameters.

DISCUSSION

In the proposed model, sequence knowledge is encoded within the SMA, and automatic sequence production is cortical (Hua & Houk, 1997). The BG activate the appropriate motor plan in SMA, which produces a response early in training and triggers Hebbian learning mechanisms between cortical areas that allow sequence learning and eventual automatic response production. These two mechanisms (i.e., Hebbian sequence learning and Hebbian automatic connectivity) will learn to automatize any simple sequence of cortical cell firing and were the only mechanisms required to account for all the data included in this article. As such, the proposed model shares many similarities with mechanisms responsible for habit learning (for a detailed review, see Ashby, Turner, & Horvitz, 2010).

The new model was used to simulate monkey behavioral and electrophysiological data collected in DSP tasks

(Matsuzaka et al., 2007; Shima & Tanji, 2000), human DSP data with rTMS application to the SMA (Verwey et al., 2002), and monkey DSP data following inactivation of the GPi (Desmurget & Turner, 2010). In addition to simulating many data sets, the model also makes an interesting and novel prediction about how the amount of training before GPi inactivation (or lesion) will affect sequence production.

Comparison with Existing Models of Automatic Sequence Production

The most prominent existing model of automatic sequence production is the Nakahara et al. (2001) model. In this model, automatic sequence production is gradually transferred from a visual or associative loop through the BG to a motor loop via a coordinator presumed to reside in the pre-SMA. Our model overlaps with the Nakahara et al. model in that they both include a prominent role for the motor loop and only use the direct pathway through the BG. However, the proposed model does not include the visual loop and instead includes a cortical circuit that allows for automatic responding. This is because the proposed new model focuses on the transition between early sequence production and automatic sequence production—not initial sequence learning.

There is now abundant evidence that human learning and memory are mediated by multiple systems (Squire, 2004). This evidence comes from many different tasks, including sequence learning (Keele et al., 2003; Clegg, DiGirolamo, & Keele, 1998). For example, Keele et al. (2003) proposed that sequence learning is controlled by the interplay of two systems. The first system requires selective attention, is driven by perceptually abstract stimulus or response features, and is mediated by a network that includes occipital, temporal, prefrontal, and lateral premotor cortices. The second system is driven by perceptually raw stimulus or response features and is mediated by a neural network that includes the SMA, M1, and PPC. The Nakahara et al. (2001) model maps nicely onto this theoretical viewpoint. Their visual loop, which is centered cortically in the dorsolateral pFC and uses working memory, is conceptually similar to the first Keele et al. (2003) system, and their motor loop, which is centered cortically in the SMA, is conceptually similar to the second Keele et al. system.

Although we accept the evidence for dual systems of sequence learning and we broadly endorse the Keele et al. (2003) and Nakahara et al. (2001) description of those systems, our goal was to focus more on how sequence production becomes automatized and less on initial learning. Thus, our model lacks an executive, pFC-based system and a detailed model of the BG and therefore may be an incomplete description of early sequence learning. Our model could be viewed as a more neurobiologically detailed version of the Nakahara et al. (2001) motor loop (and of the Keele et al. second sys-

tem) and one that is focused on accounting for automaticity. In Nakahara et al., sequence production that has reached the motor loop is automatic, whereas in our model, the motor loop through the BG is a transitional stage and sequence production truly becomes automatic only after transfer to cortex is complete. The GPi inactivation data of Desmurget and Turner (2010) are consistent with this transfer to cortex and inconsistent with the hypothesis that automatic behaviors are stored in the BG. For example, several fMRI studies have reported that BG activation increases with extensive practice in perceptual categorization but that the correlation between this activation and response accuracy decreases with practice (Helie, Roeder, & Ashby, 2010). This suggests that, although the BG are still active during automatic behaviors, they may not be responsible for mediating those behaviors (Muhammad, Wallis, & Miller, 2006). Other lines of evidence come from reports that some Parkinson's disease (PD) patients who have BG dysfunction are nevertheless able to emit an automatic motor response when presented with a familiar visual cue (kicking a ball), despite difficulties in initiating novel voluntary movements (Asmus, Huber, Gasser, & Schöls, 2008). A variety of other lines of evidence from several cognitive domains are consistent with this view (Helie et al., 2015; Ashby et al., 2010).

It is important to note that Miyachi et al. (1997) did report that inactivation of the sensorimotor striatum impaired the production of well-learned sequences in the 2 × 5 task. This seems at odds with the results of Desmurget and Turner (2010), but there are several possible explanations for these opposing results. One is that the 2 × 5 task is more cognitively demanding because of its hierarchical nature and its dependence on working memory, so participants may not be able to automatize this task to the same extent as the simpler DSP task or the SRTT. Hence, the motor loop may mediate the highest level of automaticity that is possible in the 2 × 5 task. Another possibility is that the increased task complexity only delayed the development of cortical automaticity and that longer training would have eventually allowed participants to reach cortical automaticity. Future research will be needed to disentangle this discrepancy.

Comparison with Cognitive Models of Sequence Learning

Although very few computational models have attempted to account for automatic sequence production, initial sequence learning has received more attention from cognitive modelers (for a review, see Cleeremans & Dienes, 2008). One of the earliest models of the SRTT was Cleeremans' (1993) Dual Simple Recurrent Network, which is a composite of two simple recurrent networks (Elman, 1990) that separately represent explicit and implicit knowledge. Simple recurrent networks use "contextual"

units (a running average of hidden unit activations) to supplement the input and learn sequences.

More recently, the CLARION cognitive architecture has been used to simulate the emergence of explicit knowledge in the SRTT (Sun, Slusarz, & Terry, 2005). The CLARION model includes two feedforward connectionist networks—one that models implicit processing using distributed representations and one that models explicit processing using localist representations. In both networks, the connection weights map the sequence element at time $t - 1$ to the sequence element at time t. One unique feature of CLARION models is the inclusion of bottom–up learning of explicit rules. Hence, the CLARION simulations focused on simulating the process of extracting explicit knowledge from implicit knowledge. Sun and colleagues were able to account for individual differences in performance by using different thresholds for learning new explicit rules, thus controlling the amount of explicit knowledge.

A third model, called TELECAST (Helie, Proulx, & Lefebvre, 2011), uses a hybrid connectionist/Bayesian network that implicitly learns contingencies between attended events, and the contingencies are then used to build a Bayesian network (both the structure and probabilities) that represents explicit causal knowledge. The sequence knowledge is redundantly represented in the connection weights (implicit) and the Bayesian network (explicit). Importantly, TELECAST was the first computational model of sequence learning that does not learn via backpropagation.

Among the approaches listed above, the model proposed here is closest to TELECAST. Both models learn contingencies, regardless of whether they are correct or not, using Hebbian learning. As such, both models make the implicit assumption that anything that is repeated will eventually be learned (as in habit learning). Even so, TELECAST's focus is on rule discovery in humans, whereas our model does not address the issue of implicit versus explicit sequence knowledge. Instead, its focus is on neurobiological detail and the later stages of learning. Knowledge discovery typically happens early in learning (if at all; Sun et al., 2005), and TELECAST includes no neurobiological detail. Even so, both models share a common emphasis on Hebbian learning, and they share overlap at the functional level, which suggests that an extension of the model proposed here may be able to account for some of the same early learning phenomena as TELECAST.

Sequence Learning in PD

The emphasis of our model and the Nakahara et al. (2001) model on the BG suggests that sequence learning research with PD patients should be informative about the role of the BG in sequence learning. Although PD is characterized by DA imbalances in the BG, sequence learning research with PD patients has yielded widely inconsistent results, with some studies finding an impairment while others finding normal performance (Clark, Lum, & Ullman, 2014). Clark and colleagues performed a meta-analysis of 27 studies that used the SRTT (a total of 505 PD patients and 460 neurologically intact controls). They found a statistically significant medium-sized effect of PD on the SRTT with a moderate to high study heterogeneity. To assess the source of heterogeneity, Clark et al. also used metaregression analyses to assess the role of several common independent variables, but none of the tested variables alone was able to account for the heterogeneity. However, some of the tested variables had an effect when the participants were trained under dual-task conditions. Specifically, PD patients' deficits were larger when the sequence was longer, the number of full coverage was larger (i.e., full sequence repetition), or the PD symptoms were more severe (Clark et al., 2014). In contrast, complex sequences yielded a smaller effect of PD. Note that all the above effects were only present under dual-task training conditions.

In our model, PD would be modeled by reduced DA levels in the BG and SMA. Helie, Paul, and Ashby (2012) argue that the role of DA in the BG is to serve as a reinforcement signal whereas the role of DA in cortex is to increase the signal-to-noise ratio. Although the role of DA was not explicitly simulated in this article, a PD version of the new model would be impaired in learning the stimulus–response assignments (because this is a BG function), and sequence learning and learning automatic connectivity in the SMA would be slowed down (because of a lower signal-to-noise ratio in cortex).

The PD results reviewed by Clark et al. (2014) are consistent with some of the predictions made by the new model. First, because the signal-to-noise ratio would be reduced in the SMA, the simulated PD learning rate would be slower than normal controls. Hence, differences would be larger for longer sequences (because more learning is required) and would increase with additional training. In addition, more severe PD would be simulated by an even smaller signal-to-noise ratio, which would also increase the difference between PD and controls. These three predictions are consistent with the effect of sequence length, number of coverage, and disease severity reviewed by Clark et al. However, the results from PD patients are only apparent in dual-task conditions, and the new model does not include a mechanism for dual-tasking. Likewise, Clark et al. found that the effect of PD was smaller for more complex sequences, but the current model focused on nonhuman sequence learning and only includes mechanisms to learn simple, first-order sequences. Future development will be required to account for learning of more complex tasks and dual tasking.

A New Role for the BG in Automaticity

A necessary feature of any reinforcement learning training signal is high temporal resolution. If the first response is correct, then DA must be released into the synapses

quickly, before the critical traces disappear. After the correct synapses are strengthened, excess DA must be quickly cleared from the synapse. If it is not and the next response is incorrect, then the residual DA will strengthen the synapses that caused the incorrect response, thereby undoing the beneficial learning that occurred following correct responses. Within the striatum, DA reuptake is fast and allows for trial-and-error learning, whereas in cortex, DA reuptake is slow and prevents trial-and-error learning (Ashby et al., 2007). For example, the delivery of a single food pellet to a hungry rat increases DA levels in pFC for approximately 30 min (Feenstra & Botterblom, 1996). Because cortex is unlikely to use DA as a trial-by-trial reinforcement signal, cortical–cortical connections will be strengthened regardless of whether or not the resulting behavior is appropriate. Hence, the BG, which can use DA as a reinforcement signal, are critical for initial trial-and-error learning to ensure that the appropriate behavior is learned. Once the cortical–cortical synapses are built, the BG are no longer required to produce the appropriate behavior, and the behavior is automatic (for a detailed argument, see Helie et al., 2015).

Although DA plays a role in the model in building cortical–cortical synapses, reinforcement learning is notably missing. This is because the role of three-factor learning in the striatum in simple sequence learning may be primarily to train stimulus–response associations (as in Ashby et al., 2007), and this process is not the focus of the current model. Instead, this model is focused on the development of automaticity in sequence learning, and stimulus–response associations in the above simulations were fixed (or prelearned). To model the learning of stimulus–response associations, the connections between SMA Layer V and the striatum would initially be weak, and the appropriate connections would be strengthened via three-factor learning (e.g., as in Ashby et al., 2007). Future work will add this early form of learning to the current model.

Future Work and Limitations

Although the proposed model accounts for many phenomena, like all models it is incomplete. First, the visual loop should be added to model early learning. Second, chunking may play an important role in automatic sequence production, and the new model, in its current form, does not have a chunking mechanism (e.g., Verwey et al., 2002). Third, many behavioral results from the human literature on the SRTT have not been simulated (e.g., Clegg et al., 1998). For example, the new model suggests that cues are initially visual, but that the spatial motor plan eventually becomes a cue for the next location. More work is needed to test this prediction. Finally, Matsusaka et al. (2006) found evidence of sequence representation in M1 after years of training in the DSP task. The current version of the model includes no M1. Future work should correct

this shortcoming and model the transition of sequence knowledge from SMA to M1.

Notes

1. The other role of the path through the BG is to learn stimulus–response associations using dopamine-mediated reinforcement learning (e.g., Ashby et al., 2007), but this role is not simulated in this article because the focus of the model is on automatic response learning and production.
2. We use the term "unit" rather than "neuron" to allow for the possibility that each unit in the model represents a group of tightly connected neurons that fire together.
3. For example, if the response threshold in trial t was reached after 2300 msec, the output of SMA Layer IV in the last $3000 - 2300 = 700$ msec was input to SMA Layer V for the first 700 msec of trial $t + 1$.
4. The simulated monkey learning rates were set to $\eta_{A,LTP} = 7 \times 10^{-15}$, $\eta_{A,LTD} = 8 \times 10^{-14}$, $\eta_{S,LTP} = 1 \times 10^{-12}$, and $\eta_{S,LTD} = 5 \times 10^{-11}$.
5. Note that we do not numerically fit the Shima and Tanji single-cell recording data because the timing of the tasks differed.
6. The simulated human learning rates were set to $\eta_{A,LTP} = 1.1 \times 10^{-13}$, $\eta_{A,LTD} = 1.01 \times 10^{-12}$, $\eta_{S,LTP} = 1.55 \times 10^{-12}$, and $\eta_{S,LTD} = 5 \times 10^{-12}$. This increases the learning rate of the model, consistent with humans learning sequences more quickly than monkeys.

REFERENCES

Akkal, D., Dum, R. P., & Strick, P. L. (2007). Supplementary motor area and presupplementary motor area: Targets of basal ganglia and cerebellar output. *Journal of Neuroscience, 27,* 10659–10673.

Alexander, G., DeLong, M., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience, 9,* 357–381.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review, 105,* 442–481.

Ashby, F. G., Ennis, J. M., & Spiering, B. J. (2007). A neurobiological theory of automaticity in perceptual categorization. *Psychological Review, 114,* 632–656.

Ashby, F. G., Turner, B. O., & Horvitz, J. C. (2010). Cortical and basal ganglia contributions to habit learning and automaticity. *Trends in Cognitive Science, 14,* 191–232.

Asmus, F., Huber, H., Gasser, T., & Schöls, L. (2008). Kick and rush: Paradoxical kinesia in Parkinson disease. *Neurology, 71,* 695.

Brown, S., & Heathcote, A. (2003). Averaging learning curves across and within participants. *Behavior Research Methods, Instruments, & Computers, 35,* 11–21.

Clark, G. M., Lum, J. A., & Ullman, M. T. (2014). A meta-analysis and meta-regression of serial reaction time task performance in Parkinson's disease. *Neuropsychology, 28,* 945–958.

Cleeremans, A. (1993). Attention and awareness in sequence learning. In *Proceedings of the 15th Annual Meeting of the Cognitive Science Society* (pp. 330–335). Hillsdale, NJ: Erlbaum.

Cleeremans, A., & Dienes, Z. (2008). Computational of implicit learning. In R. Sun (Ed.), *The Cambridge handbook of computational psychology* (pp. 396–421). New York: Cambridge University Press.

Clegg, B. A., DiGirolamo, G. J., & Keele, S. W. (1998). Sequence learning. *Trends in Cognitive Sciences, 2,* 275–281.

Desmurget, M., & Turner, R. S. (2010). Motor sequences and the basal ganglia: Kinematics, not habits. *Journal of Neuroscience, 30,* 7685–7690.

Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology, 10,* 1–19.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science, 14,* 179–211.

Feenstra, M. G., & Botterblom, M. H. (1996). Rapid sampling of extracellular dopamine in the rat prefrontal cortex during food consumption, handling and exposure to novelty. *Brain Research, 742,* 17–24.

Flaherty, A. W., & Graybiel, A. M. (1994). Input–output organization of the sensorimotor striatum in the squirrel monkey. *Journal of Neuroscience, 14,* 599–610.

Gerfen, C. R., & Bolam, J. P. (2010). The neuroanatomical organization of the basal ganglia. In H. Steiner & K. Y. Tseng (Eds.), *Handbook of basal ganglia: Structure and function* (pp. 3–28). London: Elsevier.

Helie, S., Ell, S. W., & Ashby, F. G. (2015). Learning robust cortico-frontal associations with the basal ganglia: An integrative review. *Cortex, 64,* 123–135.

Helie, S., Paul, E. J., & Ashby, F. G. (2012). A neurocomputational account of cognitive deficits in Parkinson's disease. *Neuropsychologia, 50,* 2290–2302.

Helie, S., Proulx, R., & Lefebvre, B. (2011). Bottom–up learning of explicit knowledge using a Bayesian algorithm and a new Hebbian learning rule. *Neural Networks, 24,* 219–232.

Helie, S., Roeder, J. L., & Ashby, F. G. (2010). Evidence for cortical automaticity in rule-based categorization. *Journal of Neuroscience, 30,* 14225–14234.

Hua, S. E., & Houk, J. C. (1997). Cerebellar guidance of premotor network development and sensorimotor learning. *Learning & Memory, 4,* 63–76.

Huang, Y. Z., Rothwell, J. C., Chen, R. S., Lu, C. S., & Chuang, W. L. (2011). The theoretical model of theta burst form of repetitive transcranial magnetic stimulation. *Clinical Neurophysiology, 122,* 1011–1018.

Izhikevich, E. M. (2007). *Dynamical systems in neuroscience: The geometry of excitability and bursting*. Cambridge, MA: MIT Press.

Keele, S. W., Ivry, R., Mayr, U., Hazeltine, E., & Heuer, H. (2003). The cognitive and neural architecture of sequence representation. *Psychological Review, 110,* 316–339.

Lu, M.-T., Preston, J. B., & Strick, P. L. (1994). Interconnections between the prefrontal cortex and the premotor areas in the frontal lobe. *The Journal of Comparative Neurology, 341,* 375–392.

Matsuzaka, Y., Picard, N., & Strick, P. L. (2007). Skill representation in the primary motor cortex after long-term practice. *Journal of Neurophysiology, 97,* 1819–1832.

Middleton, F. A., & Strick, P. L. (1997). New concepts about the organization of basal ganglia output. *Advances in Neurology, 74,* 57–68.

Miyachi, S., Hikosaka, O., & Lu, X. (2002). Differential activation of monkey striatal neurons in the early and late stages of procedural learning. *Experimental Brain Research, 146,* 122–126.

Miyachi, S., Hikosaka, O., Miyashita, K., Kárádi, Z., & Rand, M. K. (1997). Differential roles of monkey striatum in learning of sequential hand movement. *Experimental Brain Research, 115,* 1–5.

Muhammad, R., Wallis, J., & Miller, E. (2006). A comparison of abstract rules in the prefrontal cortex, premotor cortex, inferior temporal cortex, and striatum. *Journal of Cognitive Neuroscience, 18,* 974–989.

Nakahara, H., Doya, K., & Hikosaka, O. (2001). Parallel cortico-basal ganglia mechanisms for acquisition and execution of visuomotor sequences—A computational approach. *Journal of Cognitive Neuroscience, 13,* 626–647.

Rall, W. (1967). Distinguishing theoretical synaptic potentials computed for different soma-dendritic distributions of synaptic input. *Journal of Neurophysiology, 30,* 1138–1168.

Shima, K., & Tanji, J. (2000). Neuronal activity in the supplementary and presupplementary motor areas for temporal organization of multiple movements. *Journal of Neurophysiology, 84,* 2148–2160.

Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory, 82,* 171–177.

Strick, P. L., Dum, R. P., & Picard, N. (1995). Macro-organization of the circuits connecting the basal ganglia with the cortical motor areas. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 117–130). Cambridge, MA: MIT Press.

Sun, R., Slusarz, P., & Terry, C. (2005). The interaction of the explicit and the implicit in skill learning: A dual-process approach. *Psychological Review, 112,* 159–192.

Verwey, W. B., Lammens, R., & van Honk, J. (2002). On the role of the SMA in the discrete sequence production task: A TMS study. *Neuropsychologia, 40,* 1268–1276.