A neural model of the development of expertise

Travis DeWolf (tdewolf@uwaterloo.ca)
Chris Eliasmith (celiasmith@uwaterloo.ca)
Centre for Theoretical Neuroscience, University of Waterloo
Waterloo, ON N2J 3G4

Abstract

The ability to develop expertise through practice is a hallmark of biological systems, for both cognitive and motor based skills. At first, animals exhibit high variability and perform slowly, reliant on feedback signals constantly evaluating performance. With practice, the system develops a proficiency and consistency in skill execution, reflected in an increase in the associated cortical area (Pascual-Leone & Nguyet, 1995). Here we present a neural model of this expertise development. In the model, initial attempts at performing a task are based on generalizing previously learned control signals, which we refer to generically as 'actions', stored in the cortex. The basal ganglia evaluates these actions and modulates their contributions to the output signal, creating a novel action that performs the desired task. With repeated performance, the cortex learns to generate this action on its own, eventually developing an explicit representation of the action that can be called directly. This transference allows the system to more quickly and consistently execute the task, reflecting development of expertise. We present simulation results matching both behavioral and single cell spiking data.

Keywords: expert actions; neural engineering framework; automaticity; basal ganglia

Introduction

The development of expertise is a feature of biological systems that allows skills commonly employed to be executed more accurately and with greater efficiency, while taxing neural resources less. This is a critical feature of neural systems, required for the range of behaviors which animals are capable of displaying by enabling the system to perform tasks of increasing complexity as expertise develops. Learning expertise is also known as developing *automaticity*.

Consider a chess player, first learning the rules of the game. When analyzing the board and trying to visualize the game several steps ahead, the player must consider how each of the pieces may move, slowly trying to piece together how the game may evolve. With practice, commonly employed moves will become salient, and imagining the modified chess board will become easier, allowing the player to visualize the board several turns in the future. As expertise develops, this ability will grow and the results of more potential moves a number of turns into the future will be easily drawn up, freeing the player to focus on deciding which action leads to the best scenario. Chess experts are those players who have developed the ability to visualize the board up to 12-15 turns into the future.

This skill is gained by taking basic cognitive actions (visualizing the state of the board if a single piece is moved), combining them to make more complex actions (visualizing the state of the board 5 turns from now), developing new 'basic' actions to represent commonly used complex actions (series of movements that tend to follow each other), and using these new actions as building blocks to create even more complex actions (stringing together series of moves to envision the board 10 steps ahead).

Here we present a biologically plausible spiking neuron model of expertise development that includes spike timing based learning, specific anatomical mappings and computational models of the involved neural structures to control generic high-dimensional state spaces. Previous models have not been able to capture pre and post learning time differences (Gupta & Noelle, 2007), or use single artificial neurons to represent entire actions and targets requiring different structural implementations for each task (Ashby, Ennis, & Spiering, 2007). The model we present avoids these limitations, and provides novel functional hypotheses about the neuro-anatomical areas involved. For validation, we present simulation results capturing phenomena from the behavioral level to the level of single cell spiking data on a motor control task.

Automaticity and expertise

Automaticity is the ability to perform a given task proficiently, without requiring conscious effort. This is developed by repeating a task over and over, until the relevant neural systems have learned to automatically execute the action without explicit guidance. When learning continues and performance reaches a threshold level the system is said to have become an expert.

In (Ashby et al., 2007), the SPEED model of automaticity was introduced. They proposed that automaticity develops through two pathways: A fast loop through direct cortico-cortico connections, and a slow loop that passes through the basal ganglia. Learned actions are stored in the cortex, and quickly generate output in response to task information - this is the fast loop. When actions stored in the cortex are not sufficient for immediately completing the desired task, the slow loop is involved. The slow loop is a cortico-basal ganglia-cortico pathway, where the basal ganglia uses a feedback error signal to converge upon a solution, and projects it to the cortex. As a solution is found, the action is learned in the cortex, such that it can be called up through the fast loop in the future. Thus the slow loop becomes unnecessary and the neural resources are freed to be allocated elsewhere.

In the same paper, Ashby et al also present a simple neural model, with one spiking neuron population (representing the input sensory cortex), and a basal ganglia represented by three non-spiking nodes, where the number of pathways is changed for each task, relative to the number of actions possible. Here, we present a fully spiking model capable of handling high di-

mensional input and system feedback signals, detailing how a complex basal ganglia model could learn to generate the correct response by generalizing previous knowledge. Additionally, we extend the model of the cortex and implement in the context of a system that is capable of developing increasingly complicated actions.

Neural coding

The spiking neuron implementation of this model is based on the Neural Engineering Framework (NEF; (Eliasmith & Anderson, 2004)). The NEF is a general method for converting high-level algorithms into realistic spiking neuron models, a type of 'neural compiler' that calculates a connection weight matrix to perform a desired function. Two basic principles of the NEF are that a) groups of neurons form distributed representations of vectors, and b) connections between groups of neurons specify a computation to be performed on those vectors. In essence, the NEF provides a method for analytically solving for the synaptic connection weights that will efficiently compute any given function.

To represent a vector using a group of neurons, the NEF generalizes the idea of preferred direction vectors. Each neuron in a group has a randomly chosen vector $\mathbf{e}$ for which it will fire most strongly. In particular, the amount of current J flowing into the neuron is the dot product of the preferred vector e with the represented value $\mathbf{x}$, times the neuron's gain α, plus the background current J_{bias} (Eq. 1).

While Eq. 1 lets us convert $\mathbf{x}$ into neural activity, we can also do the opposite by computing $\mathbf{d}$ via Eq. 2. This produces a set of linear decoding weights that can be multiplied by the activity of each neuron in the group, giving the optimal least-squares linear estimate of $\mathbf{x}$. This allows us to estimate the value currently being represented by a set of neurons, given their spiking pattern.

Most crucially, we can use $\mathbf{d}$ to calculate the synaptic connection weights that will compute particular operations. To compute a linear operation where one group of neurons represents $\mathbf{x}$ and a second group should represent $\mathbf{Mx}$, where $\mathbf{M}$ is an arbitrary matrix, we set the connection weights between neuron i in the first group and neuron j in the second group to ω_{ij} as per Eq. 3. For non-linear operations, we need to compute the values of $\mathbf{d}$ via Eq. 4.

$$J = \alpha\, \mathbf{e} \cdot \mathbf{x} + J_{bias} \tag{1}$$

$$\mathbf{d} = \Gamma^{-1}\Upsilon,\ \Gamma_y = \int a_i\, a_j\, dx,\ \Upsilon_j = \int a_j\, \mathbf{x}\, dx \tag{2}$$

$$\omega_y = \alpha_j\, \mathbf{e}_j\, \mathbf{M}\, \mathbf{d}_i \tag{3}$$

$$\mathbf{d}^{f(x)} = \Gamma^{-1}\Upsilon,\ \Gamma_y = \int a_i\, a_j\, dx,\ \Upsilon_j = \int a_j\, f(\mathbf{x})\, dx \tag{4}$$

This approach allows us to convert high-level algorithms written in terms of vectors and computations on those vectors into detailed spiking neuron models.

The model

The anatomical mapping of the model presented here is based on the Neural Optimal Control Hierarchy (NOCH) framework, proposed in (DeWolf & Eliasmith, 2011). While broadly consistent with NOCH, the present model significantly extends the NOCH framework by detailing the forms of adaptation occurring in the basal ganglia and cortex. By building in the context of the NOCH framework, the development of expertise can also be examined in terms of its interaction with the rest of the system, forcing generalizability to be considered and providing insight into the kinds of signals processed. A simplified diagram of the model is shown in Figure 1.

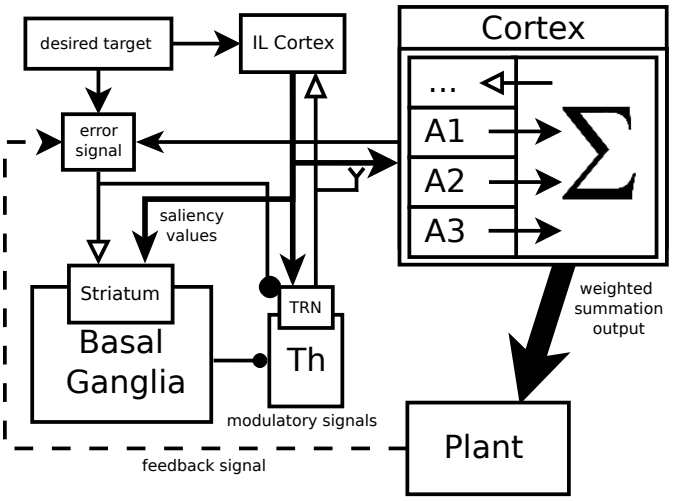

Figure 1: Th: Thalamus; TRN: Thalamic reticular nucleus; IL cortex: Infra-limbic cortex; A1-A3: Learned actions; Plant: The system being controlled. Solid arrows are regular excitatory connections, hollow arrows represent learning signals, forked connections are modulatory, circular connections are inhibitory, dashed connections represent sensory feedback.

When a goal is presented to the system, the set of weights stored in the infralimbic (IL) cortex to apply to the actions stored in the cortex (A1, A2, A3 in Figure 1) are projected to both the cortex and the basal ganglia. In the cortex, these values specify the contribution of each of the actions to the outgoing signal, which is calculated as a weighted summation and sent out to the plant in the 'fast loop'. In the basal ganglia 'slow loop' these weights are treated as saliency values for each action, providing an initial guess at the relevancy of each action to achieving the specified goal.

System feedback, relating the state of the plant after the outgoing signal has been applied, is used to generate a cortical error signal, representing the difference between the system state and the goal state. The available cortical actions also project into this neural circuit calculating feedback error, allowing the circuit to determine which action weights to increase or decrease to reduce the error. The information

is then projected into the striatum, specifying how to modify the saliencies of each of the actions. This error signal drives rapid adaption in basal ganglia through a reinforcement-like learning process (Stewart, Bekolay, & Eliasmith, 2012), taking advantage of the high learning rate in the striatum (Cragg, Rice, & Greenfield, 1997) to attempt to reduce the error to zero.

The difference between the weights determined by the basal ganglia and output by the IL cortex is calculated in the thalamus. This difference is projected both to the cortical action set to supplement the signal received from the IL cortex, and to the IL cortex directly to be used as a training signal for learning. The learning rate in the IL cortex is several orders of magnitude less than that of the striatum, so the basal ganglia is able to explore the effect of different action weightings on the system without causing system instability. Only those weights that minimize the error are held long enough to influence the IL cortex.

Simultaneously, another cortical population learns to produce the outgoing action, using an error modulated learning rule (MacNeil & Eliasmith, 2011). This newly learned action is subsequently added to the set of actions available in cortex (i.e. action 'A4' would be added to the cortical set in Figure 1). Consequently, the action created through slow loop modulation to minimize error is consolidated in the cortex, and in the future preferentially activated upon presentation of the same goal, allowing the goal to be achieved utilizing fewer system resources.

In this way expert actions for novel goal states can be developed by generalizing previously learned motor actions and learning to generate the action specified by the basal ganglia slow loop automatically in a faster cortical-cortical loop.

The basal ganglia

The basal ganglia model used here, shown in Figure 2, is a spiking neuron implementation of the computational model proposed in (Gurney, Prescott, & Redgrave, 2001), that has been extended to operate in high-dimensional vector spaces. This implementation based on neuro-anatomy performs a type of soft winner-take-all (WTA) functionality, identifying among its input signals those with the highest utilities. It features a unique efficiency of execution and scalability, and additionally matches experimental timing data (Stewart et al., 2012). WTA functionality has direct implications for the process of action selection, in which the basal ganglia is commonly implicated (Barto, 1994; Doya, 2000).

In past work, this basal ganglia model has been used with the intent of selecting a single winner from a set of input signals (Stewart et al., 2012; Eliasmith et al., 2012), requiring further post-basal ganglia neural circuity. Here, however, we wish to take advantage of its nonlinear character to modulate a set of learned actions to determine how they best combine to complete the desired task. This necessarily involves having more than one non-zero value output from the basal ganglia, which, when input saliency values are very similar, is the natural behavior of the circuit.

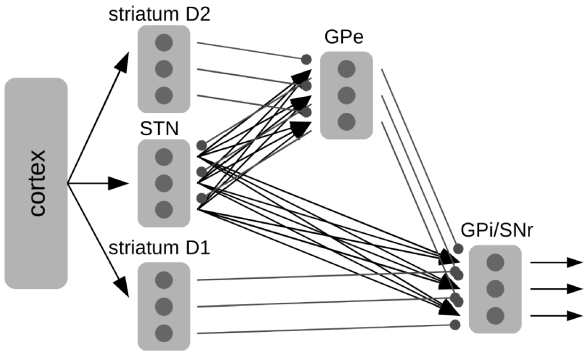

Figure 2: STN: Subthalamic nucleus; GPe: Globus pallidus external; GPi: Globus pallidus internal; SNr: Substantia nigra pars reticulata. Note that each area is a population of 100 neurons. Taken from (Stewart et al., 2010).

To appropriately do this, the population of neurons that the cortex projects to, representing the striatum, receives a feedback error signal from the cortex and action saliency values are modified based on their contribution to the current system error. If the error signal is high in a particular dimension, actions that contribute to that dimension have their saliency reduced. The system does this across all dimensions, attempting to drive the error to zero.

Continuing the chess example, if there is a set of cortical actions that represents moves the player can take at different points in time, then by choosing a subset the player can envision the resulting state of the board. Evaluating this state, it might be decided that it would be best for one of the pieces to be in a different position; forming a feedback error signal. The saliency of actions moving this piece in the desired direction would then be increased, until the output from the basal ganglia specifies a different set of actions, whose consequences can then be evaluated.

Eventually, the system will converge upon a combination of cortical actions that minimizes the feedback error signal. The next time the same goal is presented, the slow loop will generate the same set of modulatory output values to again minimize the error.

The cortex, transference, and the TRN

There are two parts of the cortex where learning occurs that drives the ability to quickly perform actions, and to acquire increasingly complicated actions: The IL cortex and the site where actions are stored.

Being able to reproduce the movement generated by the basal ganglia reduces to generating the same modulatory values to drive the cortical actions. A region in the medial prefrontal cortex, known as the infralimbic (IL) cortex, is an area of the brain that is associated with the subcortical suppression of activity and has been implicated as an executive controller of habits and behavioral strategies (Smith, Virkud, Deisseroth, & Graybiel, 2012; Daw, Niv, & Dayan, 2006). In

this model the IL cortex is responsible for learning the correct weights to apply given a specific target, such that the cortex can quickly drive the system to the target, without requiring the guidance of the slower subcortical loop.

To do this, the output of the basal ganglia modulates the weights of the cortical actions, so that it can explore the state space to converge on a correct solution, and slowly drive the IL cortex to represent these values. As the IL cortex learns these weights, the guidance of the basal ganglia is required less and less and its contribution goes to zero with the feedback error. In this way, the ability to generate the correct action is transfered from the slower subcortical loop to faster direct cortical control.

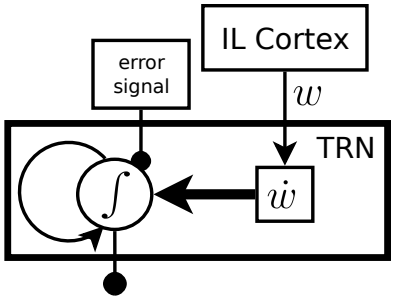

Figure 3: A diagram of the thalamic reticular nucleus (TRN). IL: Infralimbic cortex. *w* are the weights output by the IL cortex. Solid arrows are regular excitatory connections, circular connections are inhibitory. See text for details.

The thalamic reticular nucleus (TRN) plays a key role in this process. The TRN lies between the thalamus and cortex, and has been thought to prevent information from passing from the thalamus to the cortex through its inhibitory projections onto thalamic neurons. Additionally, it has been shown to be responsive to changes in firing pattern from its input (Guillery, Feig, & Lozsadi, 1998). A diagram of the TRN model presented here is shown in Figure 3. Projections from the IL cortex (w) are monitored for change above a certain threshold (i.e. $|\dot{w}| > thresh$). If there is a significant change in w, which occurs when a goal is first presented or changed but not during gradual learning, then an integrator circuit is strongly excited and begins to fire. These neurons send a constant inhibitory output to the thalamus, suppressing basal ganglia output. Due to the recurrent connection these neurons self-excite, and the inhibition of basal ganglia output will continue unless the error signal, conveying the discrepancy between the system state and the desired state, is high enough to suppress their activity. In essence, the TRN acts as a mechanism for waiting to see where the cortex drives the system before letting the basal ganglia exert control.

In order to develop the ability to perform increasingly complicated actions, the cortex contains feedback connections from the outgoing action such that it is driven to learn signals that are repeatedly output. With practice, cortical resources (shown as '...' in the Cortex block in Figure 1) will be allocated to develop a representation of the novel outgoing motor command. Because the basal ganglia has an upper bound on the number of actions that it can select at one time, as representations develop and sets of complex actions become available, the system will become better able to move through action space.

As an example of these types of cortical learning, consider a chess player who has arrived at an effective sequence of moves for a particular situation. The first several times this is repeated, effort is still required to 're-deduce' the sequence, but with practice it becomes easier and easier to produce the steps. This is the IL cortex learning the correct modulatory values and requiring less and less corrective guidance from the basal ganglia. At this point, the sequence of steps requires little attention to generate, but looking past those steps proves difficult, due to the basal ganglia's limit on selecting multiple actions. With enough practice, however, applying this sequence seems much more like only one step, reflecting its consolidation as a single complex action in the cortex. Now, the player is capable of stringing further movements on to the end of this series, and envision more steps into the future. With continued practice and development, the player will become able to string together longer chains of potential actions, and explore and evaluate consequences further and further into the action space.

Through these two types of cortical learning the system will develop the ability to quickly perform tasks and learn increasingly complicated actions.

Results

For validation of our model, we perform the same empirical data comparisons presented in (Ashby et al., 2007), looking at striatal neuron activity, the effect of outside interference in the basal ganglia on performance, and the response times profile throughout the course of developing expertise in a task. These comparisons test the model from the level of behavioral phenomena to single cell spiking data.

Striatal dropout with extended practice

In this task, (Carelli, Wolske, & West, 1997) train rats to press a lever when a tone is played. The rats were trained with 70 trials each day over 18 days, and their mean response time decreased throughout the training. Of interest here is the change in striatal neuron activity over the course of the experiment.

In this simulation the model is trained to respond with a specified weighting over a set of actions when an input signal (representing the audio tone of the experiment) is received. As the system learns the correct response and the cortex is trained the striatal neuron activity profile from the model reflects that of the single cell recording. Shown in Figure 4, the activity of both the monkey's striatal neurons (Figure 4 A) and the model's striatal neurons (Figure 4 B) shift their peak activity from before the lever pressed to after, before fading down to no activity at all.

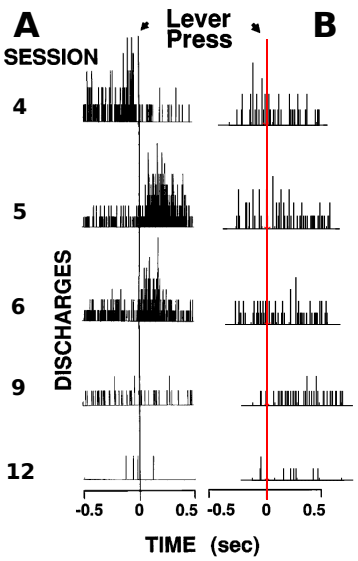

Figure 4: Striatal neuron activation throughout learning, experimental recordings on the left, model results on the right.

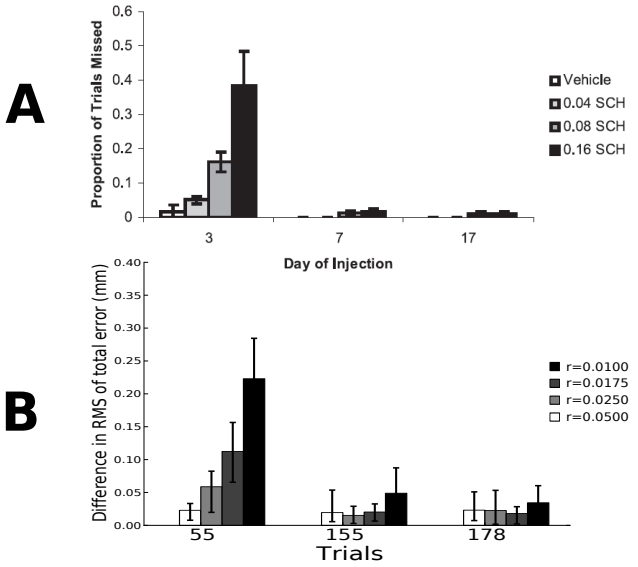

Figure 5: The effects of introducing D1 antagonist injections at different points during learning. A) Experimental results indicating the number of trials missed during a training session. B) Model results across 7 simulation trials showing the error introduced by decreasing the dopamine levels at different points during learning.

The role of dopamine in early vs late training

In this task, the model is compared against results from (Choi, Balsam, & Horvitz, 2005), where rats were trained to put their head in a compartment to receive a food reward when a tone sounded. Throughout the experiment, three different doses of a selective dopamine D1 antagonist were given to rats at different points of their training. The results show that the rats ability to learn the task is reduced proportional to the amount of D1 antagonist injected during the early days of training, but not at later stages in the experiment.

In (Ashby et al., 2007), injection of D1 antagonist is simulated reducing the dynamic range of the dopamine levels. Because the dopaminergic signal in the striatum is often proposed as an error signal for learning (Stewart et al., 2012), here, we model the D1 antagonist injection by reducing the range of values that the error signal may take on.

In the simulation the model was trained to move to a target location. The connection weights and simulation data were saved after every trial during normal learning, and then the system was loaded up at various points and driven with a reduced range on the striatal error signal. The results here (seen in Figure 5) show, analogous to the experimental results, that early in training the effect on the system performance is drastic, but is much reduced when the dopamine antagonist is introduced later in the learning process.

Response time profile

In this task, the response times (RT) profile of the model is compared to clinical data from (Nosofsky & Palmeri, 1997). Humans have a stereotyped response time (RT) profile over the course of learning a skill, where the mean RT decreases as a power function of the amount of practice. Analyzing the data from the model simulations (seen in Figure 6) used to generate the striatal dropout results, we can see that the response times profile throughout learning very clearly follows an exponential function, the same as humans.

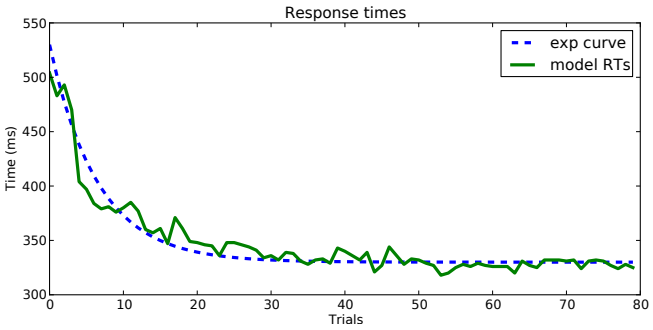

Figure 6: The response times profile of the model throughout learning, shown following a stereotypical exponential decay as seen in humans.

Discussion

The experimental results discussed above suggest that the basal ganglia is initially important to complete the task, but it becomes less involved in control of execution with further training. The results of the neural model of expertise development presented here match this data, in terns of shifting spiking activity and reducing it, sensitivity to dopamine antagonists, and exactly reproducing the classic exponential decay curve seen in human response times.

The presented model extends the model of automaticity proposed in (Ashby et al., 2007), incorporating both the infralimbic cortex, thalamic reticular nucleus, and a complex model of the basal ganglia. Several novel predictions are

made about the function and mechanisms used in these neural areas, such as the basal ganglia modulating the contributions of a set of cortical actions to outgoing control signals, and the neural circuit underlying the TRN. Additionally, it should be noted that the model is general enough to apply to learning automaticy in cognitive systems, as well as in motor systems.

There is, of course, much room for expansion of the presented model. Currently the model operates under the assumption of a static base set of cortical actions to build from. An interesting direction to explore would be to incorporate learning into these actions such that with continued exposure to a specific set of actions resources from previously learned populations are reallocated. Another interesting direction would be to develop a layered model, where the weighted output of cortical actions from one system projects to another, specifying the subsequent layer's goal. This is especially of interest in the context of the NOCH framework, which proposes a hierarchical abstraction of control where the premotor areas specify end-effector movement through 3D space, and the primary cortex is responsible for generating the correct muscle activation commands to carry out the movement on a limb (DeWolf & Eliasmith, 2010).

The model was built and simulated in Nengo (http://www.nengo.ca), using leaky-integrate and fire neurons, run in Ubuntu 12.04 on a Intel Core i7 Quad Core running at 3.2GHz with 14GB of ram.

Conclusion

We have presented a biologically plausible neural model of the development of expertise, featuring spiking neurons, high dimensional feedback and control signals, a complex model of the basal ganglia, and the ability to generalize previously learned data.

For validation of the model we compared simulation results to single cell recordings, neural systems level experimental interference studies, and behavioral data. These results demonstrate the potential of this model for providing further insight into the development of expertise in the brain. In particular we have shown that the model matches mammal performance not only in terms of behavior, but also in response to neurochemicals and the detailed spiking behavior of an individual neuron. Unlike Ashby's model, here we take into account a much wider range of brain areas, with all of them fully implemented in spiking neurons. We are aware of no other model at this level of complexity.

Acknowledgments

This work was supported by the Natural Sciences and Engineering Research Council of Canada.

References

Ashby, F. G., Ennis, J. M., & Spiering, B. J. (2007). A neurobiological theory of automaticity in perceptual categorization. *Psychological review*, *114*(3), 632–56.

Barto, A. (1994). Adaptive critics and the basal ganglia. *Models of information processing in the basal ganglia*, 215.

Carelli, R., Wolske, M., & West, M. (1997). Loss of lever press-related firing of rat striatal forelimb neurons after repeated sessions in a lever pressing task. *The Journal of neuroscience*, *17*(5), 1804–1814.

Choi, W., Balsam, P., & Horvitz, J. (2005). Extended habit training reduces dopamine mediation of appetitive response expression. *J Neurosci*, *25*(29), 6729–6733.

Cragg, S., Rice, M., & Greenfield, S. (1997). Heterogeneity of electrically evoked dopamine release and reuptake in substantia nigra, ventral tegmental area, and striatum. *Journal of neurophysiology*, *77*(2), 863–873.

Daw, N., Niv, Y., & Dayan, P. (2006). Actions, policies, values and the basal ganglia. *Recent breakthroughs in basal ganglia research*, 91–106.

DeWolf, T., & Eliasmith, C. (2010). *Noch: A framework for biologically plausible models of neural motor control.* Naples, FL.

DeWolf, T., & Eliasmith, C. (2011). The neural optimal control hierarchy for motor control. *J Neuro Eng*, *8*(6).

Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current opinion in neurobiology*, *10*(6), 732–739.

Eliasmith, C., & Anderson, C. (2004). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*.

Eliasmith, C., Stewart, T., Choo, X., Bekolay, T., DeWolf, T., & Rassmussen, D. (2012). A large-scale model of the functioning brain. *Science*, *338*, 1202–1205.

Guillery, R., Feig, S., & Lozsadi, D. (1998). Paying attention to the trn. *Trends in neurosciences*, *21*(1), 28–32.

Gupta, A., & Noelle, D. (2007). A dual-pathway neural network model of control relinquishment in motor skill learning. *Proceedings of IJCAI*, 405–410.

Gurney, K., Prescott, T., & Redgrave, P. (2001). A computational model of action selection in the bg. ii. analysis and simulation of behaviour. *Bio cybernetics*, *84*(6), 411–423.

MacNeil, D., & Eliasmith, C. (2011). Fine-tuning and the stability of recurrent neural networks. *PloS one*, *6*(9).

Nosofsky, R., & Palmeri, T. (1997). An exemplar-based random walk model of speeded classification. *Psych Review*, *104*(2), 266.

Pascual-Leone, A., & Nguyet, D. (1995). Modulation of muscle responses evoked by transcranial magnetic stimulation during the acquisition of new fine motor skills. *Journal of Neurophysiology*, *74*(3), 1037–1045.

Smith, K., Virkud, A., Deisseroth, K., & Graybiel, A. (2012). Reversible online control of habitual behavior by optogenetic perturbation of medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, *109*.

Stewart, T., Bekolay, T., & Eliasmith, C. (2012). Learning to select actions with spiking neurons in the Basal Ganglia. *Frontiers in neuroscience*, *6*, 2.

Stewart, T., Choo, X., & Eliasmith, C. (2010). *Dynamic behaviour of a spiking model of action selection in the basal ganglia.*