The art of scaling up : a computational account on action selection in basal ganglia

Bhargav Teja Nallapu, Bapi Raju Surampudi, Nicolas Rougier

HAL Id: hal-01534459
https://hal.inria.fr/hal-01534459

Submitted on 7 Jun 2017

The art of scaling up : a computational account on action selection in basal ganglia

Bhargav Teja Nallapu*‡, Bapi Raju Surampudi*†, Nicolas P. Rougier‡§¶
*International Institute of Information Technology, Hyderabad
†School of Computer and Information Sciences, University of Hyderabad
‡Inria Bordeaux Sud-Ouest, Talence, France
§LaBRI, UMR 5800, CNRS, Bordeaux INP, Universit de Bordeaux, Talence, France
¶Institut des Maladies Neurodgnratives, UMR 5293, CNRS, Universit de Bordeaux, Bordeaux, France

Email : bhargav.teja-nallapu@inria.fr

Abstract—**What makes a computational neuronal model 'large scale' ? Is it the number of neurons modeled? Or the number of brain regions modeled in a network? Most of the higher cognitive processes span across co-ordinated activity in a network of different brain areas. However at the same time, the basic information transfer takes place at a single neuron level, together with multiple other neurons. We explore modeling a neural system involving some areas of cortex, the basal ganglia (BG) and thalamus for the process of decision making, using a large-scale neural engineering framework, Nengo. Early results tend to replicate the known neural activity patterns as found in the previous action selection model by Guthrie et al. 2013, besides operating with a larger neuronal populations. The power of converting algorithms to efficiently weighed neural networks in Nengo (Stewart et al. 2009 and Bekolay et al. 2013) is exploited in this work. Crucial aspects in a computational model, like parameter tuning and detailed neural implementations, while moving from a simplistic to large-scale model, are studied.**

I. Introduction

Action selection is one of the most vital processes for the survival of an organism.

The very importance of the process of action selection (and learning from it) in the survival of an organism, pushes for a need of understanding the process from an evolutionary, biological and physiological point of view.

Brain works by running complex dynamics forming complex connections within itself, specific to each behavioral repertoire. It is at the level of few (or many) neurons that the information transfer takes place, which drives all the processes like perception, memory, decision making, language or any other brain process. Ideally, it is that level of detail one wants to model any sub-system of brain or the whole model of brain itself. Because of the practical limitations like lack of detailed experimental data, computational challenges in studying the existing data forces major modeling efforts to employ simplistic neuron models and smaller population representations.

It is a positive sign that despite these limitations, large scale models still manage to emerge, if not at a full-human scale, but at least at a scale of a rodent's brain. A recent work in [4] proposes a thalamo-cortical network model of BG to reproduce Parkinsonian tremor. Quite interestingly, the model consists of four neuron types in all the subareas, modeled at the full spatial size and neuron numbers of the rat. However, depending on the functional dynamics one tries to model like in [5], [1], often simpler neuron (rate) models are employed to keep the computations easier and emphasize more on the network dynamics than on the detail of neuronal implementations.

We considered one of our previous works on a computational study of decision making and its dynamics using a thalamo-cortical BG network [6] inspired by [1] and replicated in [7]. It is a framework on the top of an optimal decision making model, accommodating representations of varying visual salience and delayed presentation of the stimuli, studying their effect on optimal decision making. The framework uses known network dynamics as in [1] and the model exhibits similar characteristics as compared to experimental work described in [8]. The framework is indeed robust to study the other factors mentioned, that affect optimal decision making - such as stimuli properties, presentation delays and different associations of reward in learning. Aforementioned model is still quite simplistic with as few as 72 neurons representing all of the cortico-thalamic basal ganglia network, involved in decision making. This is where we emphasize on modeling the same functional networks at a large scale.

The simple model has been successfully expanded to a large scale, both in terms of neuron populations and extending to a variety of tasks that the BG are known to be involved in. For example, we have attempted to build a similar large scale model of BG also demonstrating an arm reaching a target[1] (based on the model described in [9]), besides reward based decision making. Scaling up involved exploring various large scale neural modeling frameworks like DANA [10] and NEF[2] (Nengo). The model is scaled using DANA to as many as 66,000 neurons. With the large scale model, the functional properties and neural activity patterns have been found to be consistent with the simplistic model. Using the large scale neural modeling platform, Nengo, also yielded the same dynamics of the simplistic model but now the model

---

[1]https://www.youtube.com/watch?v=oIHfDabVqcg
[2]Neural Engineering Framework - http://nengo.ca/

is expanded to a massive scale of 300K neurons. The power of converting existing models to efficiently weighed neural networks in Nengo is exploited in this work. Additionally, *NEF* accounts for various biologically realistic properties of neurons which form the internal neural networks.

II. Methods

A. Task

The task used in this work is a probabilistic learning task that was described in [8] and used in [1], [7], [6]. Four target shapes are associated with different reward probabilities (see figure 1). Each trial (an independent task turn), any two of the four possible shapes are presented at two random positions (out of the four cardinal positions - up, right, down and left). By the end of trial period, a choice is made and the reward is given according to the reward probability associated with the chosen shape.

In a single independent trial, the cognitive decision (shape of the cue) and motor decision (direction of position) are independent of each other. However, ideally, the motor decision is expected to be in the direction of the cue shape that is chosen. On subsequent reward association of the cue decision, one should eventually select the direction in which the most rewarding cue is present.

B. Model

In [5], authors demonstrated an action selection mechanism in the cortico-basal ganglia loops based on a competition between the positive feedback, direct pathway through the striatum and the negative feedback, hyper-direct pathway through the sub-thalamic nucleus. In [1], authors investigated further how multiple level action selection could be performed by the basal ganglia, and the model has been extended in a manner consistent with known anatomy and electro-physiology of the basal ganglia in the monkey (see figure 2). This model allows a bidirectional information flow between loops such that during early trials, a direction can be selected randomly, irrespective of the cue positions. However, after repeated trials, the model is able to consistently make the motor decision, only after cognitive decision (indeed biased by the cognitive decision) towards the position of the more rewarding cue shape.

To begin with, we reproduce the thalamo-cortical basal ganglia network with similar characteristics as in [1]. On a different note, this effort also reminds how a model needs to be reproducible and less constrained on its parametric values. In the new model, we could successfully demonstrate the decision in a single trial. With dopamine modulated learning in the cognitive channels between cortex and striatum, the model reached a state where it consistently chooses the higher rewarding stimulus. We also developed the model to serve as a framework to accommodate several other complex scenarios that are more likely to happen in real life decision making.

The model in its simplest form is based on few implicit assumptions, which facilitate the demonstration of optimal action selection. One of them being the cortico-striatal learning only in the cognitive channel. That is, attributing reward only

to the shape of the cue chosen but not to the position it is shown. Secondly, in all the experiment scenarios, both the stimuli are presented simultaneously. Lastly, both the stimuli are represented by same visual salience in the experimental scenarios, which is an unlikely case with real-world choices. The model has been extended so as to account such assumptions, providing a framework to study factors affecting optimal decision making. However, to expand it to a larger scale, we chose to base on its basic form so that the larger model would provide better representations to challenge such assumptions.

One of the major limitations of the simple model, we propose, is the discrete representation of stimuli in the discussed brain areas. In every structure (thoughout this paper, we refer each brain region that is modelled as *structure*), a single neuron (in each of cognitive and motor channels) representing each cue shape stands too simplified a representation and limits us from analyzing cases with complex stimuli. In figure 3a, an ambiguous shape is shown, which resembles two of the known shapes but still uncertain to identify. Figure 3b shows a familiar shape but with added noise that makes it difficult to identify as a known shape. It is not just difficult, but impossible to represent such complex stimuli in the simple and basic model. Thus we begin to expand the model to a scale that is sufficient to provide more biologically plausible stimuli representations, thus demonstrating action selection. Subsequently, the large scale model poses new testable hypothesis in the context of stimuli representation.

III. Large scale model

With the idea of scaling up the model and to have a larger population of neurons representing the stimuli, we decided to implement the model with discussed dynamics using DANA and Cython [11] individually and then scaled up the number of neurons representing the cue shape inputs in important areas of cortex and striatum. DANA provides very intuitive way of modeling neural networks by using structures called Groups - to represent neuron groups and connections - various type of connectivities between the groups. It internally uses modified NumPy [12] structured arrays made suitable for intensive computations in the group fields, by making them a contiguous block of memory.

In the simplistic model, in a regular structure (excluding the associative structure), there is one neuron representing each stimulus (stimulus - cue shape in a cognitive channel, cue position in a motor channel) and an associative structure has a 2D array of neurons where number of rows and columns both equal to the number of stimuli. Within each structure, each neuron (in each channel of any loop) represents the average activity of a small population and is modeled as a single rate neuron with the equation:

$$\tau \frac{dm}{dt} = -m + I_s + I_{Ext} - T \qquad (1)$$

decay time constant of the synaptic input τ and threshold of the neuron T are set to respective constant values. m is the output of the neuron. I_s is the synaptic current. I_{Ext} is an
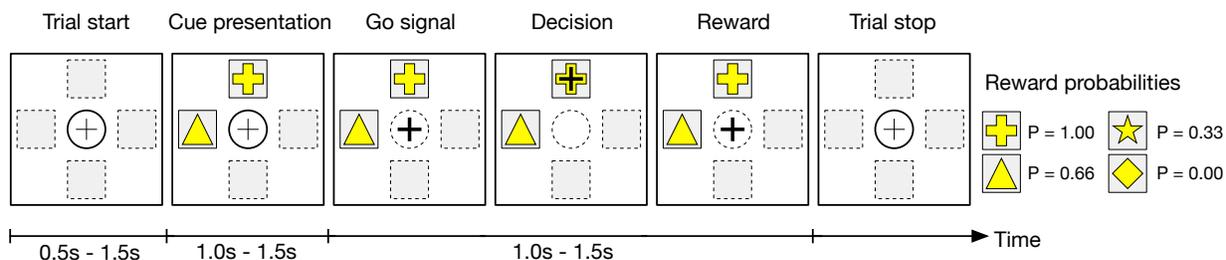
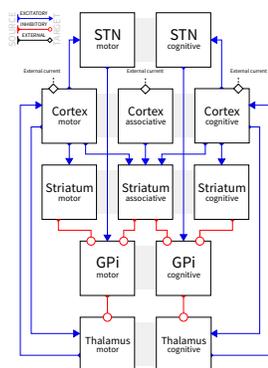Fig. 1. The two armed bandit task as described in [8], [1].

Fig. 2. Architecture of BG model from [1].

Fig. 3. Stimuli with added complexities (a) An ambiguous shape which resembles two of the known shapes (b) A noisy shape which looks familiar but not certainly known.
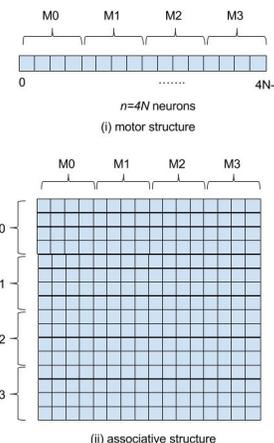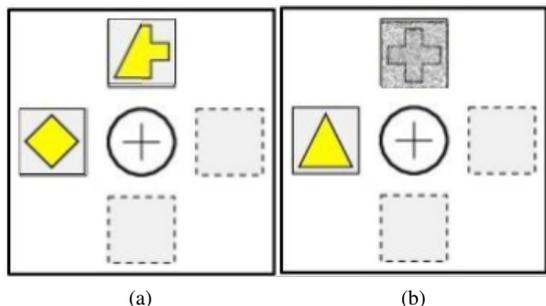

Fig. 4. (i) A motor structure with 4N neurons, with each N neurons representing a motor stimulus (cue position). A cognitive channel also has similar structure. (ii) An associative structure with 4Nx4N (16x16 in this case) neurons. All the columns of neurons represent 4 motor stimuli equally and all rows represent the 4 cognitive stimuli.

external input representing the sensory visual salience of the stimulus, is considered as the input to cortex in the model and it is a fixed magnitude to the single neuron representing that stimulus. The values of τ and T used are shown in Table I

In the expanded version, a population of neurons together represents all the stimuli. To begin with, we have a clear distinction in such population that each stimulus is represented by a subgroup of neurons and each subgroup with equal number of neurons. For example, in a cognitive channel of a structure, a population of 4N neurons (N>1) for 4 cue shapes, a subgroup of N neighboring neurons represent each cue shape (see figure 4). As for the synaptic input (I_{ext} in equation 1), a Gaussian input at the center of the subgroup for a particular cue shape is given.

It is also interesting to explore how the changes in such stimuli affect the dynamics of decision making. This complex and varying representation of stimuli is what we want to emphasize while beginning to working on a large scale model. Few examples can be seen in figure 5, where different possibilities of stimuli are represented. This kind of encoding also could address the complex stimuli presented in figure 3.

Between such structures, the connectivity is based on the assumption that the neurons in both the structures correspond to particular stimuli (shape or position) and only those subgroups of neurons are connected in the respective structures. It might be biologically implausible to assume such a distinct connectivity between the structures but it provides a good understanding of expanding this kind of model from simplistic implementations to complex highly populated networks. Using the above described structures and connections, a large computational model is developed, which exhibits the intrinsic functional properties of the simpler model mentioned. The type of connections, few corresponding to those between cortex and striatum, are listed in table II.

The model is large in terms of significantly higher number of neurons in each structure. Cortex particularly has larger population than the other structures. The population numbers of structures are listed in table III. Higher population in cortex

Structure	Threshold(T)	Decay (τ)
Cortex	-3	10
Striatum	0	10
Globus Pallidus interna (GPi)	10	10
Subthalamic Nucleus (STN)	-10	10
Thalamus	-40	10

TABLE I
PARAMETERS IN EACH STRUCTURE

Connection	Source	Target
One to One	cognitive cortex	cognitive striatum
One to One	motor cortex	motor striatum
Associative to Associative	associative cortex	associative striatum
Cognitive to Associative	cognitive cortex	associative striatum
Motor to Associative	motor cortex	associative striatum

TABLE II
CONNECTIONS BETWEEN STRUCTURES

allows complex stimulus representations. In the simpler model, the synaptic input corresponding to a stimulus is represented in the cortical structures as external current of fixed magnitude of 7 spike(sp)/s to the neuron representing that stimulus. In case of larger populations, for each stimulus, a Gaussian input (of mean magnitude of 14 sp/s and variance spreading over the subgroup of the cortical structure) is provided.
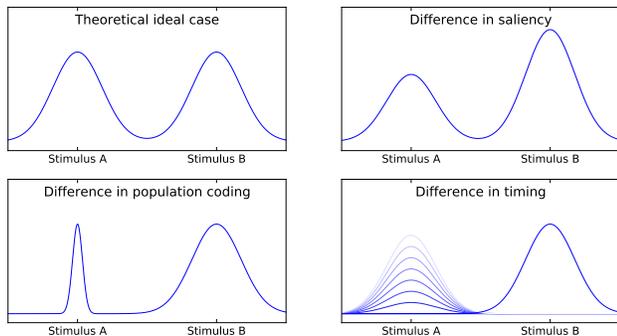

Fig. 5. Stimulus in the form of a Gaussian input over a subgroup of neurons and various scenarios of input representations. The inputs shown here are only to demonstrate the differences in their mean or variance, or in the timing (see *Difference in timing*), irrespective of any scale.

For instance, if the cortex has 1024 neurons, 256 neurons for each stimulus, and to activate a particular stimulus, a Gaussian input spreading over these 256 neurons, is given at the beginning of the trial. The mean cortical activity (over each sub-population), when the model is run for an independent trial, is found to be similar to that of simple model before

Structure	Channel			Total
	cognitive	motor	associative	
Cortex	256	256	65536	66048
Striatum	16	16	256	288
GPi	16	16	-	32
Thalamus	16	16	-	32
STN	16	16	-	32
				66432 neurons

TABLE III
NO. OF NEURONS IN EACH STRUCTURE

expanding. This is one of the key characteristics a model should have, when scaled up.

The model is also run for a session of 120 trials, with cortico-striatal learning. Although the time taken for the completion of session is long (around 40 minutes, whereas the simpler models run in few seconds), model reached a performance over 90%. Owing to this computational expense, it was not easy to run the model for 250 simulations to generate a mean performance over each of the 120 trials of all 250 sessions. The key challenge in computation lies with the structures like associative cortex, which is a 2D representation. Since the connections involving associative cortex are like the ones described in figure 6, the weight matrices between the structures span up to a scale of over thousands of rows and columns, which inevitably makes the computation expensive.
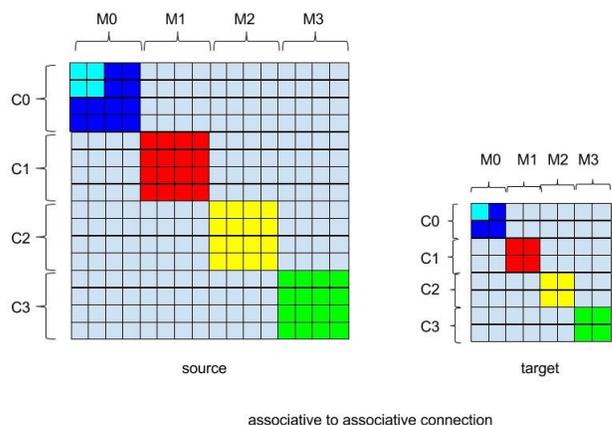

Fig. 6. Neurons corresponding to each stimulus in both the structures are connected. For a given stimulus (C_0, M_0), each neuron in target is connected to a block of neurons in source. For eg. neuron (0,0) in target is connected to the block [(0,0),(0,1),(1,0),(1,1)] in source. Similarly all the neurons that are connected are highlighted in colors.

This leaves an interesting question about the connection weights between two highly but unequally populated (like cortex and striatum). And to point out, even in this larger model, each neuron remains to be a mere computational unit, without accounting for any other of the biological neuron properties (spiking or population coded etc). Hence, we move further exploring another large scale neural modeling framework, Nengo which appears to solve this problem by effectively computing some local weights within a sub-population of neurons while providing ready-to-use spiking neurons to build the structures.

IV. THE *NEF*, NENGO AND RESULTS

The underlying principles of *NEF* are: representation, transformation and dynamics. The activity of a group of neurons is considered to be best represented by encoding of some underlying variable, e.g., any vector generally of smaller dimensionality. The *transformation* of information from one neural group to the other neural group. i.e., the neural group

representing **x** can be connected to other group of neurons representing a variable **y**, such that $\mathbf{y} = f(\mathbf{x})$.

Any desired function can be approximated with improved accuracy by increasing the number of neurons. The *dynamics* of NEF allows us to build recurrent neural networks. These networks can compute the time evolution of a given variable **x** of the form $\frac{d\mathbf{x}}{dt} = f(\mathbf{x}, \mathbf{u})$ for any control variable **u**. In a connection between two neurons, as the activity in the pre-synaptic neuron changes over time, **x(t)**, the value actually represented in the post-synaptic neuron over time will be **f(x(t)) × h(t)**, where **h(t)** represents the post-synaptic activity (e.g, $\mathbf{h(t)} = e^{-t/\tau}$, for t>0 with a time constant τ).

The structures involved are modeled as neuron ensembles, collection of several thousands of neurons. The basic pattern of connectivities are retained as listed in table II. When neuron ensembles in two different structures are connected, the individual spiking activity of each neuron is taken into account and a resultant activity of the ensemble is calculated. It is this resultant activity which is provided as the synaptic input to the target structure. This is a striking advantage in using Nengo. The individual neuron works on the basis of chosen neural dynamics, while the network dynamics can be modeled as required.

in the model dynamics that the position corresponding to the chosen shape only will be selected.

There is an array of neuron ensembles representing each stimulus and corresponding arrays are connected between structures. The effective activity of each of these arrays is taken as the resultant activity for the stimulus the ensemble array represents. Figure 8 displays a case where shape 0 is shown at position 3 and shape 1 is shown at position 2. Since shape 1 has higher synaptic weight, it should be chosen over shape 0. It can be seen that shape 1 is selected (label_1 - green in CTX_COG). Subsequently, for the motor decision, position 2 is selected (label_2 - orange in CTX_MOT).

For the models claiming to have network-induced dynamics, the choice of the neuronal dynamics should not make a significant difference. This has been demonstrated clearly in this work. The model developed is tested under similar conditions of action selection task, switching between direct computations and neural implementations like Sigmoidal and Integrate & Fire (IF). To start with, we used simple neurons that directly compute a function. Then the neurons are changed to detailed implementations, but to a lower scale. However, the switch between neuron implementations involved minimal tuning of parameters of the network as well as optimization of few framework parameters, provided there is computational power to handle the large scale of spiking or IF neurons.
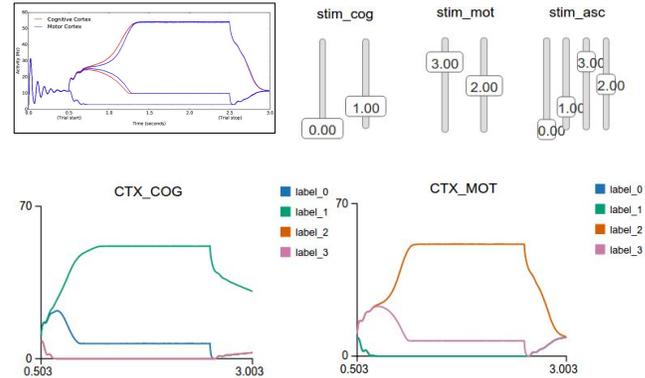
Fig. 8. For the combinations of cue shape and position, the best rewarding cue shape is selected (see CTX_COG) and the corresponding position (as per stim_cog and stim_mot represent shape 0 is shown at position 3 and shape 1 is shown at position 2). It can be seen that, at a point later than choosing shape, corresponding position is chosen (see CTX_MOT). *Inset* The decision in the original simplistic model
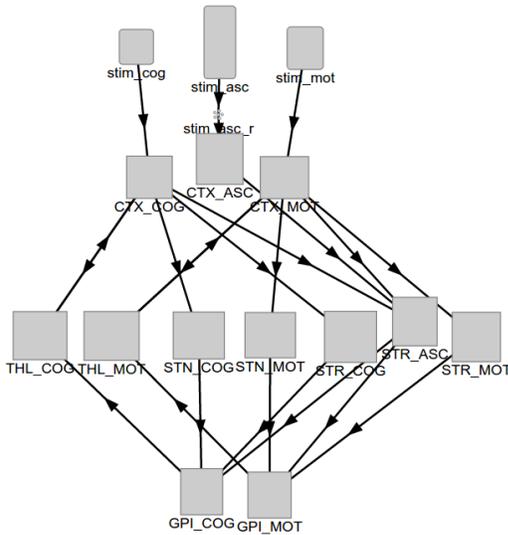

Fig. 7. BG network model implemented in nengo

The following simulations assume cue shapes are represented by [0, 1, 2, 3] in the cognitive channel of the cortex and the cardinal directions, where the shapes could be presented, are also represented by [0, 1, 2, 3] in the motor channel of the cortex. The cue shapes have a rewarding probability in the order 1>0>2>3. The synaptic weights that were learned in the previous model are applied. To ensure the optimal action is selected in a generic manner, we exchanged the weights of the first (0) and the second (1) stimuli. These weights ensure that the highest rewarding cue is always selected. It is also implicit

V. DISCUSSION

Although the population numbers in this work are comparable to that of experimental numbers of a rat's BG and some recent modeling work by Moren et al. 2015 [4] (see table IV), we are still far away from the scale of that of humans' (150K in rats where as 2.2 million neurons, only the globus pallidus and sub-thalamic nucleus combined [13]). As for results, the neural activity corresponding to synaptic weight-driven decision making and the ideal choice as per high reward, were consistent. We also believe, at a full range

simulations on the model, performance would match that from the previous simplistic model. In future, such a functional large scale network could become readily usable in other functional processes involving motor reaching or motor learning [14], modeling deficits in neuro-degenerative disorders such as Parkinsonianism [15] etc.

One pertinent question that arises is about the usefulness of a larger network, if the same qualitative behavior could be realized in a simpler, smaller network. In addition to being able to replicate the basic features of experimental observations as in the case of simple model, a large scale model is expected to provide additional advantages. Some of the advantages are: the flexibility to cope with sophisticated stimulus representations as opposed to simplistic ones; the degree of match with the details of known anatomical features, that is, the feasibility of a more biologically plausible architecture; and more importantly the ability to cope with insults / injuries to the system without degradation in performance, that is, graceful degradation. These additional benefits make it more appealing to try and scale up a simpler architecture. But when is a model sufficiently large? A model does not need to be as large as it can get but could be just large enough to answer questions like how do we conclude from the activity of a group of neurons that a stimulus is chosen? and how do we compare a weakly activated stimulus to a strongly activated one?

Model	Cortical level	Total
Simplistic model [1]	24	72
Model implemented	24000	387,200
Full-scale rat model [4]	180K	3 million

TABLE IV
No. OF NEURONS ACROSS MODELS

VI. CONCLUSION

In a model, when it comes to designing structures with bigger populations, it is quite challenging to come-up with appropriate connectivities among these structures. Assuming that an ideal large scale model takes the realistic neuronal population ratios between structures (like a rat cortex having considerably large populations than that in its striatum, as suggested by experimental data [16]), the connectivity patterns are difficult to assume and hence compute. In this research, an attempt was made by using some simple many-to-one connections in different sparsities. Using a sophisticated framework like Nengo provides benefit of taking care of internal computations among such sub-populations to propagate the resultant population activity through the inter-structural connections.

A general observation is that in neuroscience, there have been very few general principles governing systems. Also quite often, computational models are very rigid in terms of their specifications like having to specify range of values for various parameters, etc. Hence one should try to identify some generic characteristics of the structures involved in a foolproof manner, making them independent of the model details or the range of parameters used. This means model should be able to represent multiple scenarios over a range of parameters

used. In this work using Nengo, such configurations have been explored. Figure 7 displays how certain parameters can be isolated from the model to better tune them. Many key parameters involving reward-prediction errors, dopamine modulation etc, could be observed independently.

The framework provides various types of neuronal implementations in-built, allowing configuration of different neuron models in different structures easily. Because of the optimal weights computation within the neuron ensembles in Nengo, inter-structural connections carry network level dynamics in a very reasonable way.

The large scale models described here, could be extended to model the formation habits, expanding structures like cortex. It has been suggested that associative striatum plays a crucial role in goal-directed actions, contrary to sensori-motor that is part of the habitual system. New evidence suggests that automatic behaviors, such as habits, are stored in cortex and that the role of BG is to train cortex. In other words, BG learns quickly via reinforcement learning to activate the correct postsynaptic target in cortex, which leads to appropriate cortical Hebbian learning ([17], [18]). More evidence supporting this idea is provided by Piron et al. 2016 [19]. The scaled up cortex could provide different structured representations individually for stimulus representations, habitual learning etc, so that it could make decisions independent of BG after the acquisition of habits.

In most of the common tasks of action selection, identifying whether a stimulus is the same or different from a previous one involves coordination between the incoming sensory information and working memory [20]. Also, in one of the cases discussed about complex stimuli, when all the choices are not presented at once and new choices appear in the middle of decision making, one needs to model a strategy with which whether value computation of each choice is re-computed. Thus the role of working memory in decision making paradigms where the basal ganglia also play a crucial role, should be studied.

Thus, a comprehensive network with all the structures modeled, demonstrating optimal action selection, motor reaching (in the lines of demonstration in [9]), habitual learning, could be referred as true large scale model. Besides the mere high population of neurons, 'large-scale' should be with respect to the number of structures involved, different responsibilities each structure switches between and number of cognitive or biological processes simulated. Also, from a unified research perspective, when models of certain important pathways are developed computationally and made available on frameworks like Nengo, they are often reusable in the development of higher functional models using those pathways. Our two cents to reproducible science, brief pieces of code demonstrating simple as well as discussed large scale models are made public[3]. After all, science ought to be reproducible, replicable[4] and extendable.

References

[1] M. Guthrie, A. Leblois, A. Garenne, and T. Boraud, "Interaction between cognitive and motor cortico-basal ganglia loops during decision making: a computational study," *Journal of neurophysiology*, vol. 109, no. 12, 2013.

[2] T. C. Stewart, B. Tripp, and C. Eliasmith, "Python scripting in the nengo simulator," 2009.

[3] T. Bekolay, J. Bergstra, E. Hunsberger, T. DeWolf, T. C. Stewart, D. Rasmussen, X. Choo, A. R. Voelker, and C. Eliasmith, "Nengo: a python tool for building large-scale functional brain models," *Frontiers in neuroinformatics*, vol. 7, 2013.

[4] J. Moren, J. Igarashi, J. Yoshimoto, and K. Doya, "A full rat-scale model of the basal ganglia and thalamocortical network to reproduce parkinsonian tremor," *BMC Neuroscience*, vol. 16, no. 1, p. 1, 2015.

[5] A. Leblois, T. Boraud, W. Meissner, H. Bergman, and D. Hansel, "Competition between feedback loops underlies normal and pathological dynamics in the basal ganglia," *Journal of Neurosciences*, vol. 26, pp. 3567–3583, 2006.

[6] B. T. Nallapu and N. P. Rougier, "Dynamics of reward based decision making a computational study," in *International Conference on Artificial Neural Networks*, 2016.

[7] M. Topalidou and N. Rougier, "[re] interaction between cognitive and motor cortico-basal ganglia loops during decision making: a computational study," *ReScience*, vol. 1, no. 1, 2015.

[8] B. Pasquereau, A. Nadjar, D. Arkadir, E. Bezard, M. Goillandeau, B. Bioulac, C. E. Gross, and T. Boraud, "Shaping of Motor Responses by Incentive Values through the Basal Ganglia," *Journal of Neuroscience*, vol. 27, no. 5, 2007.

[9] S. R. C. Vignesh Muralidharan, "A scalable cortico-basal ganglia model to understand the neural dynamics of targeted reaching," *Computational Neuro Science*, 2016.

[10] N. P. Rougier and J. Fix, "Dana: Distributed numerical and adaptive modelling framework," *Network: Computation in Neural Systems*, vol. 23, no. 4, pp. 237–253, 2012. [Online]. Available: http://www.tandfonline.com/doi/full/10.3109/0954898X.2012.721573

[11] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn, and K. Smith, "Cython: The best of both worlds," *Computing in Science & Engineering*, vol. 13, no. 2, pp. 31–39, 2011.

[12] S. Van Der Walt, S. C. Colbert, and G. Varoquaux, "The numpy array: a structure for efficient numerical computation," *Computing in Science & Engineering*, vol. 13, no. 2, pp. 22–30, 2011.

[13] C. D. Hardman, J. M. Henderson, D. I. Finkelstein, M. K. Horne, G. Paxinos, and G. M. Halliday, "Comparison of the basal ganglia in rats, marmosets, macaques, baboons, and humans: volume and neuronal number for the output, internal relay, and striatal modulating nuclei," *Journal of Comparative Neurology*, vol. 445, no. 3, pp. 238–255, 2002.

[14] K. Magdoom, D. Subramanian, V. S. Chakravarthy, B. Ravindran, S.-i. Amari, and N. Meenakshisundaram, "Modeling basal ganglia for understanding parkinsonian reaching movements," *Neural Computation*, vol. 23, no. 2, pp. 477–516, 2011.

[15] V. Muralidharan, P. P. Balasubramani, V. S. Chakravarthy, S. J. Lewis, and A. A. Moustafa, "A computational model of altered gait patterns in parkinsons disease patients negotiating narrow doorways," *Frontiers in computational neuroscience*, vol. 7, 2015.

[16] G. Paxinos, *The Rat Nervous System.* Gulf Professional Publishing, 2004.

[17] F. G. Ashby and J. M. Ennis, "The role of the basal ganglia in category learning," *Psychology of Learning and Motivation*, vol. 46, pp. 1–36, 2006.

[18] F. G. Ashby, B. O. Turner, and J. C. Horvitz, "Cortical and basal ganglia contributions to habit learning and automaticity," *Trends in cognitive sciences*, vol. 14, no. 5, pp. 208–215, 2010.

[19] C. Piron, D. Kase, M. Topalidou, M. Goillandeau, H. Orignac, T.-H. N'Guyen, N. Rougier, and T. Boraud, "The globus pallidus pars interna in goal-oriented and routine behaviors: Resolving a long-standing paradox," *Movement Disorders*, 2016.

[20] A. Pooresmaeili, D. R. Bach, and R. J. Dolan, "The effect of visual salience on memory-based choices," *Journal of neurophysiology*, vol. 111, no. 3, pp. 481–487, 2014.