

Expanding a Standard Theory of Action Selection to Produce a More Complete Model of Cognition

Brian W. Colder

The MITRE Corporation
bcolder@mitre.org

Abstract

A standard model of how brains produce natural cognition would provide a framework for organizing cognitive neuroscience research. A recent effort (Laird et al., in press) to build on consensus views of cognitive operations and produce a standard model of natural cognition started with common aspects of well-established cognitive architectures ACT-R, Sigma, and SOAR. The model captures scientific consensus on “how” the brain works, but it does not offer a coherent story for “why” the component modules (i.e., working memory, long-term memory, visual and motor areas) exist and interact in the ways described. This manuscript starts with background information on a well-cited theory of action selection, and extends that theory to a fuller explanation of decision-making, action and perception that includes a framework for the elements of cognition.

Need for a Standard Model

Our knowledge of how neural systems decide, plan, act and perceive is expanding rapidly with research at various levels from single cell recordings to high-level behavioral studies, yet it is currently very difficult to understand how many of these efforts are related. How do the findings of a neurophysiologist studying error-related EEG signal help a behavioral psychologist researching the effect of context on the perception of color? A Standard Model of the Mind offers the promise of collecting disparate elements of cognitive neuroscience theory into a single framework that provides a blueprint for how the elements work together.

A Consensus Mechanism for Selecting Actions

Neural systems allow animals to detect and react to features of their environment. When multiple actions are possible at any given time, an organism must choose from among the multiple alternatives (Redgrave et al., 1999). This section reviews research describing how activity in basal ganglia pathways may serve that purpose.

Research into how stimulus information influences decision making indicates that variations in frontal (Kim & Shadlen 1999; Romo *et al.* 2004) and parietal (Platt & Glimcher 1999; Sugrue *et al.* 2004) cortical activity can be linked to simultaneous neural network encoding of different actions. The affordance competition hypothesis (Cisek, 2007) proposes that distinct, parallel fronto-parietal neural networks encode potential actions and compete against each other for continued processing. This hypothesis suggests that multiple sources, including cortical-basal ganglia-thalamo-cortical loop activity, contribute activation that biases the competition between action networks and ultimately produces a single winning action (Cisek, 2007).

Prefrontal cortex (Alexander et al., 1986; Middleton and Strick, 2002), inferotemporal cortex (Middleton and Strick, 1996), and possibly other cortical regions participate in information loops that travel from cortex down to the basal ganglia, then to thalamus, and back to the same areas of cortex. These information loops have been well-studied, and are described in detail in Schroll and Hamker (2013). Information in cortical-basal ganglia-thalamo-cortical loops can travel multiple pathways through the basal ganglia. Some of these basal ganglia pathways produce excitatory input back to the cortex, while others result in inhibitory input. For more information on the anatomy of excitatory and inhibitory pathways through the basal ganglia, see Schroll and Hamker (2013).

Mink (1996) and Redgrave et al. (1999) proposed a role for cortico-basal ganglia-thalamo-cortical information loops in selecting a single action from among competing actions. A consensus may be forming (Friend and Kravitz, 2014; Stewart et al., 2010; Hazy et al, 2007; Gurney, 2001) around the theory that the interaction between the activity in excitatory and inhibitory pathways in cortico-basal ganglia-thalamo-cortical loops is responsible for selecting among mul-

multiple competing potential actions. The well-established notion is that the neural network activity in the information loop encoding the selected action is enhanced by relative increases in the activity in excitatory basal ganglia pathways, along with relative decreases in the activity of inhibitory basal ganglia pathways. These activity patterns cause an increase in the overall network activity of the neural information loop encoding the selected action, and that increased activity enables the selected action to be executed.

Action Selection in Leabra, ACT-R and Nengo

Cognitive architectures are theories of cognitive structure and function that are implemented in software, and can serve as building blocks for cognitive models. Three biologically-inspired cognitive architectures, Leabra, ACT-R and Nengo, have implemented methods of action selection that are patterned after the basal ganglia pathway interaction described above.

Leabra

The interaction between excitatory and inhibitory basal ganglia pathways is the modeled mechanism for action selection in the neural network-based cognitive architecture Leabra (O'Reilly et al., 2016). Leabra uses realistic simulations of spiking neurons to build models of cognition and behavior. The creators of Leabra are the authors of multiple papers describing the role of the BG in gating working memory (e.g. Hazy et al., 2007, O'Reilly and Frank, 2006), where they termed the excitatory and inhibitory pathways the "Go" and "No-Go" pathways, respectively. In Leabra models, Go and No-Go pathway activity in the basal ganglia gates both the performance of specific actions and the flow of information into working memory, which is maintained by neural network activity in prefrontal cortex (O'Reilly et al., 2016).

ACT-R

The ACT-R cognitive architecture models cognition at the functional level, using a set of specialized modules that process specific types of information (Anderson et al., 2004). Since the procedural module in ACT-R contains rules for action selection and execution, the role of the basal ganglia has been ascribed to the procedural module in ACT-R (Anderson et al., 2004). The time course of simulated activity in the procedural module predicts the measurement of blood oxygenation in the head of the caudate nucleus of the basal ganglia. (Anderson et al., 2008). However, the assignment of the procedural module to the basal ganglia may not fully explain known interactions between excitatory and inhibitory basal ganglia pathways (Stocco, 2017).

Nengo

The interaction between excitatory and inhibitory pathways through the basal ganglia is also modeled by the action selection mechanism in the Nengo cognitive architecture (Stewart et al., 2010). Nengo is a relatively new cognitive architecture that simulates large-scale neural systems, typically by combining many simulations of single neurons (Eliasmith, 2013). The architecture has been used to model the effects of drugs on working memory (Duggins et al., 2017), model the effect of changes in basal ganglia dopamine levels on speech (Senft et al., 2016), and create a large-scale model of the brain (Eliasmith et al., 2012). The basal ganglia modeled in Nengo chooses the highest utility action by suppressing inhibitory output for that action.

Action Selection in the Standard Model

The widespread adoption of the theory that action selection results from changes in the balance of excitatory and inhibitory basal ganglia pathways along with the incorporation of that theory into three widely used cognitive architectures suggests that these ideas should be a necessary part of any Standard Model of the Mind.

Expanding the Standard Model

Starting with this established theory of action selection, Colder (2015) described a hypothesis of natural neural system operation that is based on the manipulation of potential futures. According to this theory brains consider multiple potential futures for the organism. Each potential future includes some action by the organism, and the resulting expected sensory environment. The potential futures are instantiated in the brain by linked cortical activation in sensory and motor areas.

The key contribution to established theory from Colder (2015) is that the cortical neuronal networks in cortical-basal ganglia-thalamo-cortical loops do not instantiate only potential future actions - they also instantiate the expected outcome of those actions. Basal ganglia Go and No-Go pathways gate complete sensorimotor possible futures for the organism, and thus the basal ganglia is not just selecting potential actions, but also the desired future.

Directly following the action selection logic expressed by Mink (1996), Redgrave et al. (1999), Hazy et al. (2007) and others, Colder (2015) proposes that previously learned reward information stored in basal ganglia synapses acts on cortical-basal ganglia-thalamo-cortical feedback loops to disinhibit the cortical activation corresponding to a selected future. Disinhibition of the selected future allows its network activity to spread from more abstract cortical regions to primary motor and sensory areas, causing "realization" of the selected future. The motor portion of the selected future

is realized when the spread of that future's activation to primary motor areas results in an attempt to perform the action. Similarly, the expected sensation of the potential future is realized when that future's activation spreads to lower-level areas and shapes the way incoming sensory information is perceived, as described by the theory of predictive coding (Rao and Ballard, 1999).

This hypothesis places the established theory of the basal ganglia's role in action selection into a larger context that also provides a framework for relating empirical results in other areas of cognitive neuroscience research, such as perception, attention, and higher cognitive function (Colder, 2011). For instance, there is considerable evidence indicating that sensory expectations have a strong effect on our ongoing sensory perception (Nitschke et al., 2006; Bar, 2003). Similarly, control of attention is strongly tied to both normal sensory expectations, and violations of those expectations (Ruz and Lupianez, 2002). Also, activated representations of potential futures are theorized to play an important role in higher-level cognition (Barbey, 2009), and emotional prospection (Gilbert and Wilson, 2007).

A Consensus Description of Working Memory

The term "working memory" refers to information that is held in memory for a number of seconds and can be used to make decisions and guide actions. For many years there was a standard model describing the neural mechanisms underlying working memory, but in the last 10 years the story has changed somewhat (Esposito and Postle, 2015; Postle, 2006).

Early studies of working memory in non-human primates found that neurons in the pre-frontal cortex (PFC) increased activity during the delay period of tasks that required monkey to remember information in order to accomplish their task (Fuster and Alexander, 1971). These results combined with others showing that PFC lesions degrade working memory (Bauer and Fuster, 1976; Milner, 1963) led to the prevailing view that sustained activity in PFC neurons maintained stimulus representations and enabled working memory. However, recent evidence indicating that increased PFC activity during the delay period may be related to attention rather than memory (Lebedev et al., 2004) and that sustained activity firing in other cortical areas fits the memory maintenance profile (Gnadt and Andersen, 1988; Harrison and Tong, 2009) disrupts that hypothesis.

Regardless of the precise underlying neural mechanisms, the cognitive computational properties of working memories are well-known. Working memory can store a limited number of representations of recognizable stimuli (Baddeley and Hitch, 1974). Without intentional rehearsal, these representations are available to cognition for a short time,

on the order of seconds. Models of working memory typically assume that representations of current stimuli are established in working memory, where they are maintained for a limited time and may influence behavior in the same manner as current stimuli (Hazy et al., 2007). There are verbal and visuospatial components of working memory that operate independently of each other (Baddeley and Hitch, 1974). Along with the current sensory environment, the contents of working memory could also be filled by activated long-term memories (Fuster, 2006).

Future Creation Uses Working Memory and Current State

The future selection process described above and in Colder (2015) operates on many potential futures at once. According to this hypothesis this set of potential futures is created when the neural activity instantiating both the organism's current state and the current contents of working memory spreads to associated brain regions, creating a set of potential future states. This process is enabled by a lifetime of unsupervised learning that shapes cortical synapses so that the neural network activation that instantiates and sustains the current state, including working memory, spreads to potential actions and their expected outcomes that have been previously associated with similar current states and working memories. The set of potential actions and expected outcomes are potential future states that the organism considers for selection.

Colder (2015) suggests that these potential futures are encoded in the activity of cortical-basal ganglia-thalamo-cortical loops, and that selection of the most desirable future state is accomplished by the action of basal ganglia Go and No-Go pathways.

Future Realization Sets the Current State and Working Memory

As described above, Colder (2015) proposes that during the process of realization, the neural activity encoding a future that has been selected by the basal ganglia will migrate to primary sensory and motor cortices to generate action and influence perception. Upon realization, the current state of kinesthetic, proprioceptive and external sensory information can be placed into working memory for later use. While this (previously sensory) information is active in working memory, it is treated as an extension of the current state.

The description of basal ganglia activity selecting a desired state that, after realization, morphs into the current state which could become a working memory, bears some relation to the hypothesis of Hazy et al (2007). These authors theorized that excitatory (Go) and inhibitory (No-Go) pathways in the basal ganglia act as gates for information enter-

ing working memory. In the current formulation, the information selected by the Go and No-Go pathways will become the current state and potentially the contents of working memory after any changes introduced by realization.

The process of realization can introduce multiple sources of change, or error, from the initially selected “ideal” future to the ultimately realized current state and entries into working memory. One potential cause of a difference between the selected future and the final current state is the future’s level of abstraction. Initially created futures may not be very closely tied to reality. Redgrave (1999) described a hierarchical motor system, in which more abstract motor plans are selected by the BG, then those selected motor plans gain specificity and detail. Similarly, while the expected outcome of an action may be known in the abstract, the final details of the state that results from the realized sensory environment may add considerable information to the initial expectation.

Another source of error in realization is the selection of a future that is incompatible with the present. A clear example of this type of error comes when our expectations for the environment are wrong, such as when a driver begins to change lanes on a highway, only to hear a loud horn sound and discover another car was in their blind spot. In this case realization of the lane change was expected to produce a position in the adjacent lane, but those expectations were not realized, and the resulting environment and final state are different than the selected future.

Summary

Animals are constantly faced with decisions about the best course of action. All cognitive architectures include a method for deciding between multiple actions at a specific point in time. A standard model of the mind must include a method for action selection. The current manuscript lays out the neuroscience behind the widely-accepted view that action selection is executed by synapses in the basal ganglia controlling the activity of excitatory and inhibitory pathways in cortico-basal ganglia-thalamo-cortical loops. We also put forth a higher-level explanation for the mind that builds upon this theory of action selection, stating that rather than just selecting potential actions, the relative activity in excitatory and inhibitory basal ganglia pathways are selecting entire sensorimotor potential futures. We propose that neural systems create, select and attempt to enact these potential futures. This theory provides a unifying context for understanding empirical findings in studies of decision-making, action, perception and higher-level reasoning. Returning to the question posed in the beginning of the manuscript, the theory could provide a framework that integrates findings from the neurophysiologist with those of the behavioral psychologist, as neurophysiological signals can be

used to determine whether individuals were aware of color perception errors attributable to false expectations generated by contextual cues. Although no consensus has yet been established in support of this theory of future manipulation, it can serve as an inclusive organizing framework for a Standard Model of the Mind.

References

- Anderson J.R., Fincham J.M., Qin Y., Stocco A. (2008) A central circuit of the mind. *Trends in Cognitive Sciences*. 12:136–143.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., Qin, Y. (2004). An integrated theory of the mind. *Psychological Review* 111, (4). 1036–1060.
- Baddeley A., Hitch G.J.. Working memory. In: Bower G, editor. *Recent Advances in Learning and Motivation*. New York: Academic Press; 1974
- Bar, M. (2003). A cortical mechanism for triggering top-down facilitation in visual object recognition. *J. Cogn. Neurosci.* 15, 600–609.
- Barbey, A. K., Krueger, F., Grafman, J. (2009). Structured event complexes in the medial prefrontal cortex support counterfactual representations for future planning. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 364, 1291–1300.
- Bauer, R. H., Fuster, J. M. (1976). Delayed-matching and delayed-response deficit from cooling dorsolateral prefrontal cortex in monkeys. *J. Comp. Physiol. Psychol.* 90, 293–302. doi: 10.1037/h0087996
- Colder, B. (2015) The basal ganglia select the expected sensory input used for predictive coding. *Front. Comput. Neurosci.*, 23 September | <https://doi.org/10.3389/fncom.2015.00119>
- Colder, B. (2011). Emulation as an integrating principle for cognition. *Front. Hum. Neurosci.* 5:54. doi: 10.3389/fnhum.2011.00054
- Duggins P., Stewart T.C., Choo X., Eliasmith C. (2016) The Effects of Guanfacine and Phenylephrine on a Spiking Neuron Model of Working Memory. *Top Cogn Sci.* 2017 Jan;9(1):117-134. doi: 10.1111/tops.12247. Epub 2016 Dec 21
- D’Esposito, M., Postle, B. R. (2015). The cognitive neuroscience of working memory. *Annual Review of Psychology*, 66, 115–142. <http://doi.org/10.1146/annurev-psych-010814-015031>
- Eliasmith C. (2013) *How to build a brain: A neural architecture for biological cognition* Oxford University Press
- Eliasmith C., Stewart T.C., Choo X., Bekolay T., DeWolf T., Tang Y., Rasmussen D. (2012) A large-scale model of the functioning brain. *Science*. Nov 30;338(6111):1202-5. doi: 10.1126/science.1225266.
- Friend, D.M. Kravitz, A.V. (2014) Working together: basal ganglia pathways in action selection *Trends Neurosci.* Jun; 37(6): 301–303. doi: 10.1016/j.tins.2014.04.004
- Fuster, J. M. (2006). The cognit: a network model of cortical representation. *Int. J. Psychophysiol.* 60, 125–132.
- Fuster, J. M., Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science* 173, 652–654. doi: 10.1126/science.173.3997.652
- Gilbert, D. T., Wilson, T. D. (2007). Propection: experiencing the future. *Science* 317, 1351–1354.

- Gnadt, J. W., Andersen, R. A. (1988). Memory related motor planning activity in posterior parietal cortex of macaque. *Exp. Brain Res.* 70, 216–220.
- Gurney K., Prescott T.J., Redgrave P. (2001) A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biol Cybern.* Jun;84(6):401-10
- Harrison, S. A., Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458, 632–635. doi: 10.1038/nature07832
- Hazy, T. E., Frank, M. J., O’Reilly, R. C. (2007). Towards an executive without a homunculus: computational models of the prefrontal cortex/basal ganglia system. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 362, 1601–1613. doi: 10.1098/rstb.2007.
- Kim J.N., Shadlen M.N. (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci.* Feb;2(2):176-85.
- Laird, J. E., Lebiere, C., Rosenbloom, P. S. (In press). A Standard Model of the Mind: Toward a Common Computational Framework across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Magazine*.
- Lebedev M. A, Messinger A., Kralik J. D., Wise S. P. Representation of attended versus remembered locations in prefrontal cortex. *PLoS Biology.* 2004;2:1919–1935.
- Milner, B. (1963). Effects of different brain lesions on card sorting. *Arch. Neurol.* 9, 100–110. doi: 10.1001/archneur.1963.00460070100010
- Mink, J. W. (1996). The basal ganglia: focused selection and inhibition of competing motor programs. *Prog. Neurobiol.* 50, 381–425. doi: 10.1016/s0301-0082(96)00042-1
- Nitschke, J., Sarinopoulos, I., Mackiewicz, K., Schaefer, H., Davidson, R. (2006). Functional neuroanatomy of aversion and its anticipation. *Neuroimage* 29, 106–116.
- O’Reilly, R.C., Hazy, T.E. Herd, S.A. (2016) The Leabra Cognitive Architecture: How to Play 20 Principles with Nature and Win! In S. Chipman (ed) *Oxford Handbook of Cognitive Science*, Oxford: Oxford University Press
- O’Reilly, R.C., Frank, M.J. (2006). Making Working Memory Work: A Computational Model of Learning in the Frontal Cortex and Basal Ganglia. *Neural Computation*, 18, 283-328
- Platt M.L., Glimcher P.W. (1999) Neural correlates of decision variables in parietal cortex. *Nature*, Jul 15;400(6741):233-8.
- Postle, B. R. (2006). Working Memory as an Emergent Property of the Mind and Brain. *Neuroscience*, 139(1), 23–38. <http://doi.org/10.1016/j.neuroscience.2005.06.005>
- Rao, R. P., Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87. doi: 10.1038/4580
- Redgrave, P., Prescott, T. J., Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*. 89, 1009–1023. doi: 10.1016/s0306-4522(98)00319-4
- Romo R, Hernández A, Zainos A. (2004) Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron* 2004 Jan 8;41(1):165-73.
- Ruz, M., Lupiáñez, J. (2002). A review of attentional capture: on its automaticity and sensitivity to endogenous control. *Psicológica* 23, 283–309.
- Senft, V., Stewart, T.C., Bekolay, T., Eliasmith, C., Kröger, B.J. (2016) Reduction of dopamine in basal ganglia and its effects on syllable sequencing in speech: A computer simulation study, *Basal Ganglia*, Volume 6, Issue 1, Pages 7-17, ISSN 2210-5336, <http://dx.doi.org/10.1016/j.baga.2015.10.003>.
- Stewart T.C., Choo X., Eliasmith C. (2010). “Dynamic behaviour of a spiking model of action selection in the basal ganglia,” in *Proceedings of the 10th International Conference on Cognitive Modeling*, eds Salvucci D., Gunzelmann G., editors. (Philadelphia, PA: Drexel University), 235–240
- Stocco, A. (2017) A Biologically Plausible Action Selection System for Cognitive Architectures: Implications of Basal Ganglia Anatomy for Learning and Decision-Making Models. *Cogn Sci*, 2017 Jun 6. doi: 10.1111/cogs.12506.
- Sugrue L.P., Corrado G.S., Newsome W.T. (2004) Matching behavior and the representation of value in the parietal cortex. *Science*. Jun 18;304(5678):1782-7.