

AFRL-RI-RS-TR-2008-177
Final Technical Report
June 2008



INTERACTING BRAIN MODULES FOR MEMORY: AN ADAPTIVE REPRESENTATIONS ARCHITECTURE

Rutgers University

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. V027

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

**The views and conclusions contained in this document are those of the authors
and should not be interpreted as necessarily representing the official policies,
either expressed or implied, of the Defense Advanced Research Projects
Agency or the U.S. Government.**

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2008-177 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/

WILMAR W. SIFRE
Work Unit Manager

/s/

JAMES A. COLLINS, Deputy Chief
Advanced Computing Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

Table of Contents

List of Figures	ii
Abstract	iii
Summary, Introduction & Methods	1
Results and Discussions	1
Cortico-hippocampal circuits (episodic memory, spatiotemporal relations)	1
Cortico-striatal circuits (intention selection and dopamine modulation)	4
Modulation of action contingencies via dopamine	6
Dopamine reward circuits (Intrinsic Reward and its Neural Basis)	8
Cerebellum	9
Conclusion	11
References	12
Acronyms List	15

List of Figures

Figure 1.1: Hippocampal Formation interaction with Cortex	2
Figure 1.2: Schematic of a single corticostriatal loop.....	4
Figure 1.3: Basal ganglia output channels.....	5
Figure 1.4: Schematic of main connections of the dopaminergic system.....	7
Figure 1.5: Schematic of the major connections of the cerebellum.....	10

Abstract

In our original proposal for the DARPA BICA project, our plan was to develop an innovative and novel biologically-based computational model of interacting brain modules for memory, using the adaptive representations architecture of Gluck & Myers (2001, Gateway to Memory: An Introduction to Neural Network Models of the Hippocampus and Learning, MIT Press). The approach was to begin with a connectionist-level architecture for the hippocampal region (medial temporal lobe) as a central system for creating optimal and adaptive stimulus representations, and then work outwards from the hippocampal region to the brain systems that it modulates, including the cerebellum, cerebral cortex, basal ganglia, as well as other structures which, themselves, reciprocally modulate the hippocampus (ventral tegmental area/VTA, medial septum of the basal forebrain). Ultimately, this would define a novel biologically-inspired and constrained architecture for the neural substrates of a broad range of learning and memory behaviors and capabilities.

Our direction changed toward the end of the first year as we saw the opportunity to collaborate with other BICA teams to create a completely new biologically-inspired architecture, called TOSCA. TOSCA was to be the basis for our collaborative Phase II submission to BICA. The TOSCA team included:

- Michigan (John Laird, Richard Lewis, Thad Polk, Doug Pearson (Three Penny))
- MIT (Cynthia Breazeal, Linda Smith (Indiana), Larry Barsalou (Emory))
- Dartmouth (Richard Granger, Carey Priebe (Johns Hopkins), Anna Tsao (Algotek))
- Harvard (Stephen Kosslyn, Giorgio Ganis, Bruce Draper (CSU))
- Rutgers (Mark Gluck)

Summary, Introduction & Methods

This report reflects the research we have done under BICA and represents a summary of the Rutgers contributions to the larger TOSCA architecture. The design of TOSCA starts at the brain system and circuits levels. In developing an initial version of TOSCA, our team chose to abstract away from much of the complexity of the brain. Many brain systems include multiple subsystems that are extremely complex in their own right (e.g., vision and hearing within sensory systems) and the sophisticated computational mechanisms underlying these systems. This is purely a tactical decision to get us started and we fully plan to greatly expand the systems and subsystems in TOSCA in the future. Our strategy is to include those neural systems that we consider most important in constructing an initial functional architecture that provides end to end behavior.

The Rutgers team had primary responsibility and/or significant contributions to three components: (1) Cortico-hippocampal circuits, (2) Cerebellum, and (3) Fronto-striatal/basal ganglia. Refer to the Michigan report for a broader overview and the summary of other components contributed by other team members.

These three components are described below in the results and discussion section. The final deliverable was a blue-print for an integrated architecture for cognition which, had it been continued, would have been our proposed work to be done under Phase II funding of BICA.

Results and Discussion

Cortico-hippocampal circuits (episodic memory, spatiotemporal relations)

Anatomical structure

As illustrated below in Figure 1.1, our network model of cortico-hippocampal circuits for learning and memory include modules corresponding to the dentate gyrus (DG), CA3 and CA1 fields of the hippocampus proper, and superficial and deep entorhinal cortex, which receive inputs from the perirhinal and parahippocampal cortices which, in turn, get projections from the rest of the brain.

Entorhinal Cortex (EC): The entorhinal cortex contains six layers that, for simplicity, can be divided into "superficial" (layers I-III) and "deep" (layers V-VI). The superficial layers receive highly-processed multimodal sensory input from neocortex (primarily via perirhinal and postrhinal cortex). Principal neurons in the superficial layers include pyramidal neurons (in layer III) and stellate cells (in layer II). The stellate cells project via the perforant path to DG and CA3, while the pyramidal cells project to CA1 (and subiculum). The superficial layers also contain a large number of GABAergic interneurons that exert a widespread inhibitory control over the output of principal cells. The deep layers receive input from CA1 (and subiculum) and project back to the same neocortical areas that provided input to the superficial layers. There is also a projection from deep to superficial EC that causes both excitation and feedforward inhibition (van Haeften et al., 2003). Pyramidal cells in the deep layers show graded persistent firing (over 5 minutes) which could allow for reverberating circuits (superficial EC to hippocampus to deep EC

to superficial EC) to maintain stimulus representations across short delays (Frank & Brown, 2003).

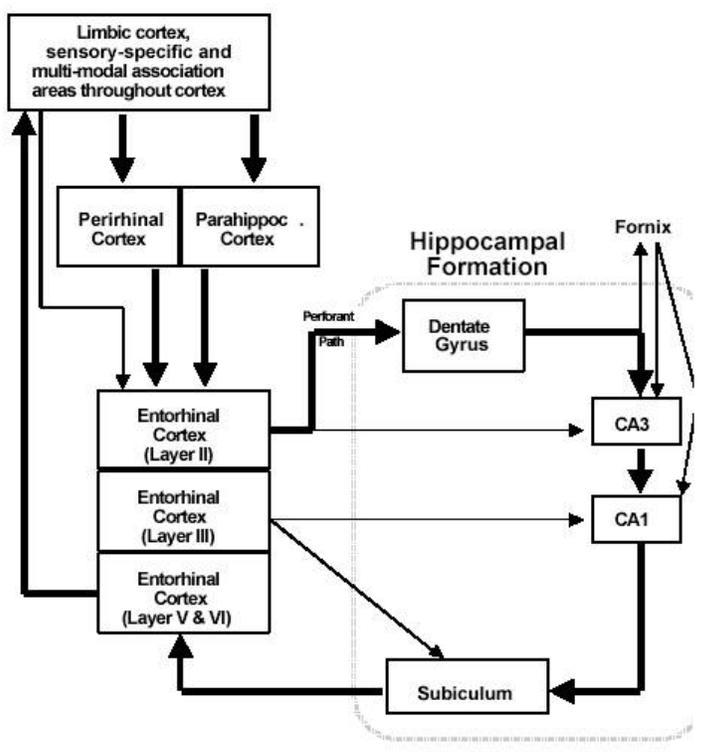


Figure 1.1: Hippocampal Formation interaction with Cortex

Hippocampal Formation: The hippocampus includes a DG layer, a CA3 layer, and a CA1 layer. Connections from DG to CA3 and from EC to CA1 are topologically organized. Each stellate neuron in EC contacts a subset of the possible postsynaptic targets in DG and in CA3. Each neuron in CA3 contacts a subset of the possible postsynaptic targets in CA3 and CA1.

Physiological operation

EC neurons receive external input representing highly pre-processed multimodal sensory information from cortex. They will be modulated by interneurons providing both feedback and feedforward inhibition; for simplicity. Strong inhibitory processes and local circuit feedback in the EC cause representational compression, implementing representational clustering function proposed by Myers et al. (1995). Deep EC neurons form the principal output of the hippocampal region back to cortex and also project to principal cells in superficial EC.

Computational function

In our implementation of TOSCA, we will follow the widely accepted hypothesis that the hippocampal region plays a critical role in the acquisition of new memories, both (1) rapidly-acquired memories for autobiographical events, sometimes collectively called **episodic memory** (e.g. Squire, 1987; Squire et al., 2004), as well as being critically involved in developing novel **adaptive stimulus representations** that are important both for episodic memories but also for

incrementally-acquired procedural memories which are otherwise mediated through the cerebellum and basal ganglia. As a starting point we plan to incorporate our previous neural network modeling of hippocampal region processing in the larger architecture (Gluck & Myers, 1993, 2001; Myers & Gluck, 1994). This model assumes that the hippocampal region develops new stimulus representations that encode contextual and stimulus-stimulus regularities. Specifically, we found that known features of the anatomy and physiology of EC (sparse activation of principal neurons, dense inhibition, and local plasticity mechanisms) give rise to the compression of redundant features in the input. This model accounted for data showing that latent inhibition and sensory preconditioning, which depend on compressing together the representations of conditioned stimulus and context and/or co-occurring cues, survive selective hippocampal lesion but are impaired after EC or broad hippocampal-region damage (Myers et al., 1995). We will adopt this same model in the initial version of TOSCA. We will also follow our previous modeling in assuming that the hippocampal layer forms a compact code for the whole situation in which the organism finds itself (what we call the 'ensemble'; Murnane, Phelps, & Malmberg, 1999). Such representations will form the basis of episodic memory in TOSCA, which are acquired in one or a few exposures and include information about the spatial and temporal context in which learning occurred (e.g. Meeter et al., 2004; Hasselmo & Eichenbaum, 2005; O'Reilly & Rudy, 2000), or on spatial and sequence learning, which may be animal analogues of human episodic learning (e.g. Lisman et al., 2005; Sharp, 1999; Tsodyks et al., 1996).

Systems

Interactions between the hippocampal system and other neural systems will play a crucial functional role in TOSCA. At the highest level, the hippocampal system will constantly be encoding and storing compressed representations of the current state (as represented in posterior cortex). When similar states are encountered in the future, they will activate the previously stored compressed representation, which will in turn reinstantiate information from the previously stored state in posterior cortex. Once this information is represented in posterior cortex, it can influence which actions/intentions are proposed and selected. Furthermore, we envision corticohippocampal loops in TOSCA storing and retrieving temporal sequences of events that have been experienced. Specifically, each event in a sequence could provide cues that lead to retrieval of the next event in the sequence. In this way, the hippocampal system could be used to replay a sequence of events from the past. Doing so could be potentially very valuable to the agent, because it would make it possible to plan ahead and predict likely future events that may improve its present decision making.

The interaction between the hippocampal system and anterior cortex could provide another crucial functionality for TOSCA. Recall that one critical assumption of the architecture is that it learns how and when to perform mental operations as well as motor actions. That is, the same learning algorithms will be used to reinforce rewarding actions, whether they are mental actions or motor actions. The initial design of TOSCA will exploit this strategy in order to learn how best to exploit its episodic memory system. For example, TOSCA should be able to learn when the mental act of attempting an episodic memory retrieval is likely to lead to long-term reward. Similarly, it should learn when episodic storage is called for. Indeed, the agent should even be able to learn what retrieval cues to set in posterior cortex in order to retrieve memories that are likely to help in deciding how to act. Put simply, TOSCA should be able to learn how to use its

episodic memory most effectively in addition to learning episodic memories themselves.

Cortico-striatal circuits (intention selection and dopamine modulation)

Anatomical structure

The basal ganglia (BG) are a set of interconnected, sub-cortical nuclei which form a complex network of loops integrating cortical, thalamic and brainstem information (Alexander et al 1986). There are three main pathways from the cortex, through the BG, and back to the cortex as illustrated in Figure 1.2. The striatum is the input nucleus of the direct pathway. It projects directly to the output nuclei of the BG, the globus pallidus interna (GPi) and substantia nigra pars reticulata (SNr). The output nuclei project back to the cortex via the thalamus, with the input returning to the same cortical module that provided the excitation to the striatum. The striatum also has a second pathway to the output nuclei, the indirect pathway. This two step inhibitory pathway provides delayed excitation to the same area of the output nuclei that the striatum inhibited via the direct pathway. The hyperdirect pathway provides a route for cortical excitation to be passed to the output nuclei of the BG.

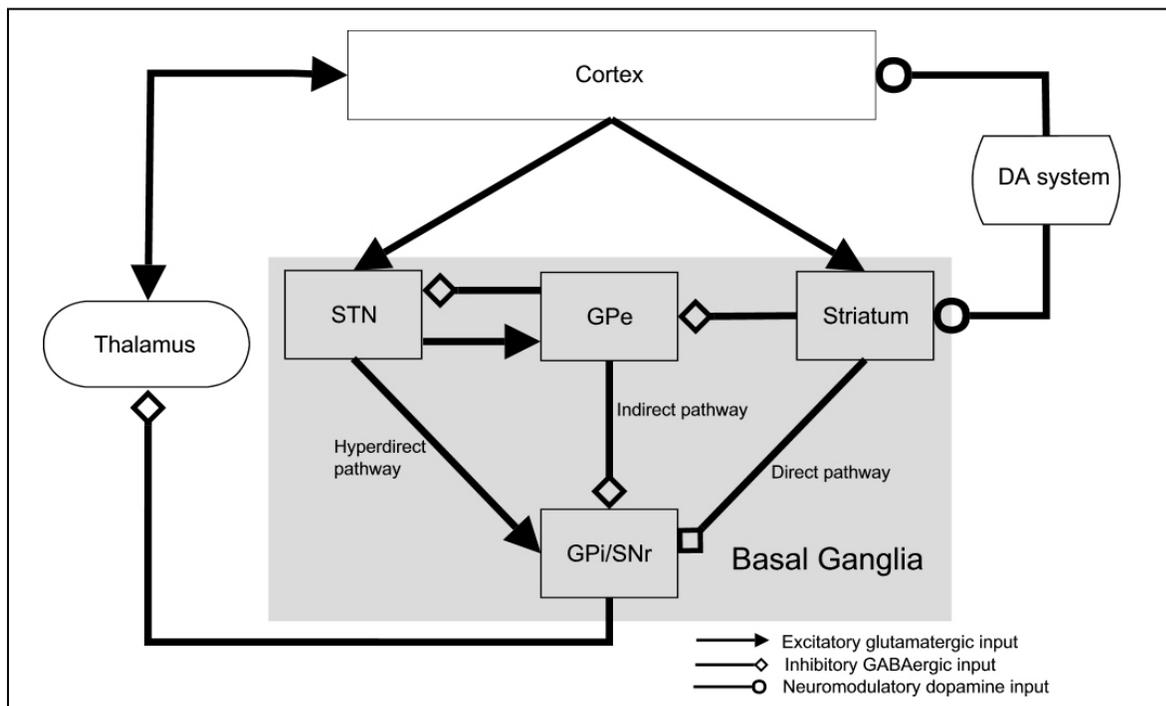


Figure 1.2: Schematic of a single corticostriatal loop.

Each loop through the basal ganglia originates in a specific cortical area and terminates in the same area. This provides a set of parallel loops through the basal ganglia as shown by the relationship of the output channels with specific cortical areas as seen in Figure 1.3. Communication between the channels occurs at the level of corticothalamic loops and cortico-cortical circuits.

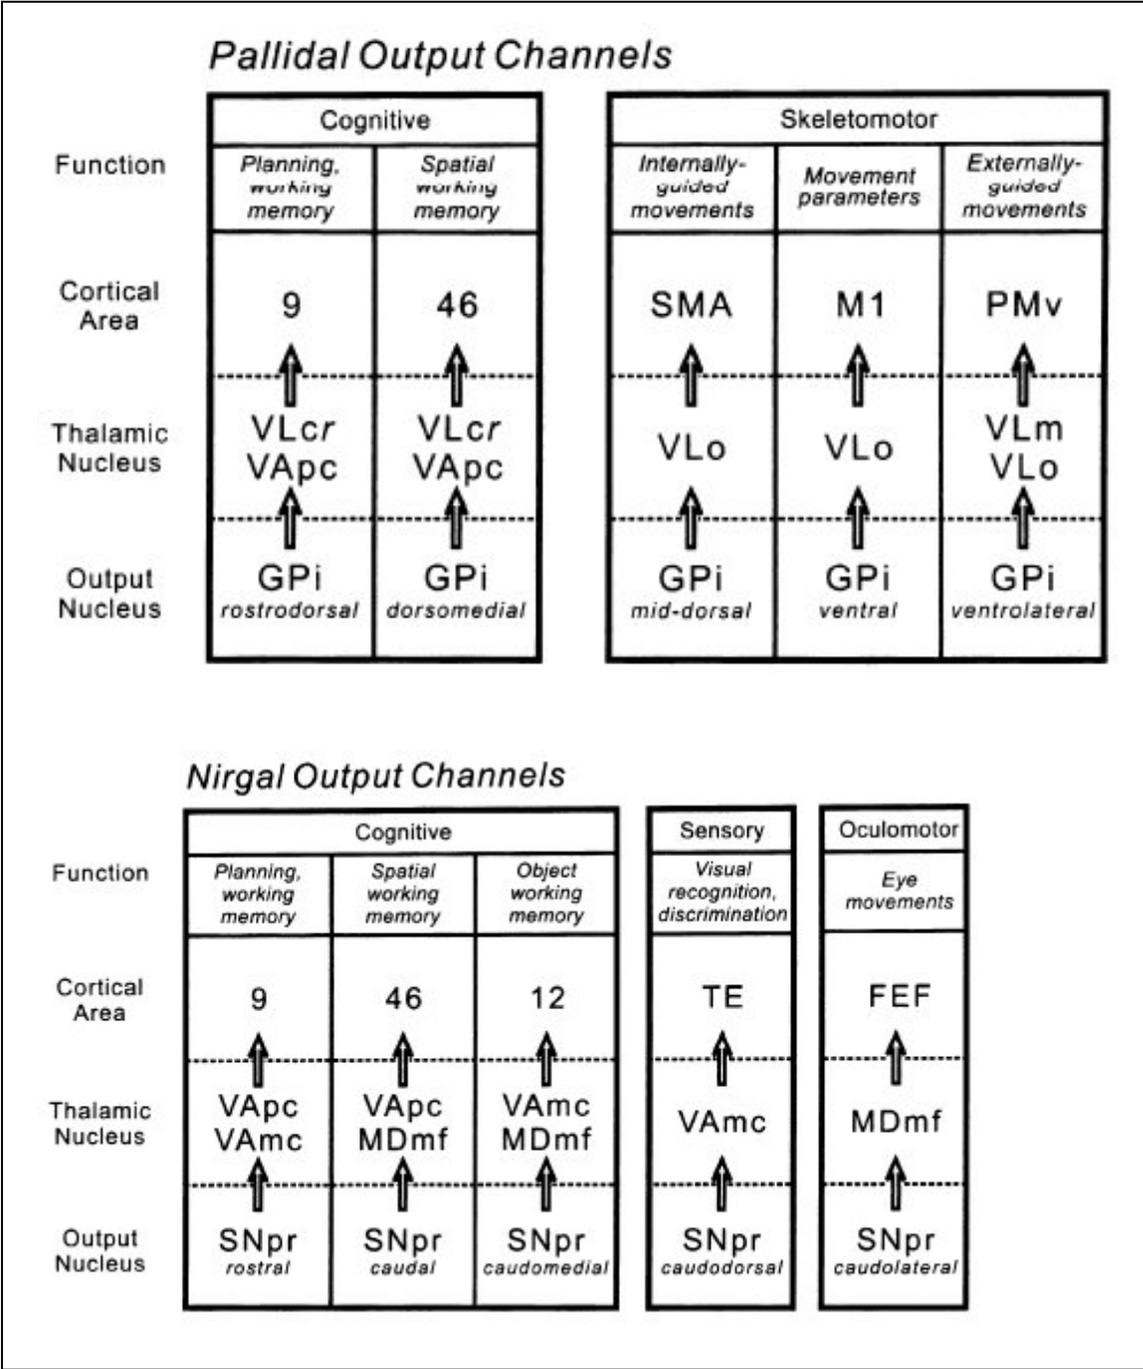


Figure 1.3: Basal ganglia output channels

Physiological operation

The cortical module proposes a number of contesting intentions. These are held in check by the tonic inhibitory output of the GPI/SNr acting via the thalamus. The striatum acts to decide amongst the competing intentions using information from past rewards obtained in similar environmental contexts. The three pathways provide mechanisms for intention selection, control the force of release of the intention and duration of release of the intention. The presence of multiple, parallel corticostriatal loops allows for the selection of multiple intentions in parallel.

Intentions that are mutually exclusive (e.g. reach for ball with left hand and scratch head with left hand) will be presented within the same corticostriatal loop, and will therefore be decided between. Intentions that can be executed in parallel (e.g. walk and talk) can be selected in parallel and thus executed simultaneously. The segregated corticostriatal loops interact at the cortical level, with the feeding of information generally from areas of more abstract intentions to more motor intentions. As an example, the first corticostriatal loop, communicating with areas in prefrontal cortex may decide that the medium term intention goal is to satisfy hunger. This decision will be passed back to the prefrontal cortex and forwarded to more motor planning areas. The next corticostriatal loop, originating from the motor planning areas, will decide that the current motor plan is to go to the cafeteria. This decision is then communicated back to the motor planning cortical area and forwarded to a shorter term motor planning area. This series of loops continues until the first action of the sequence is decided upon, perhaps rising from a chair. The medium term goal of hunger satiation remains. The actions needed to fulfill that goal are executed in sequence until the goal has been met and another medium term goal attains a higher priority and is therefore selected in the corticostriatal loop.

Derived computational functionality

We assume that a central function of corticostriatal circuits is action selection (or more accurately, intention selection). Specifically, the corticostriatal circuits in TOSCA will act as a winner-take-all network to mediate between mutually exclusive intentions. The main computation is performed at the level of the striatum where the intrinsic membrane properties of the principal neurons provide the capability to differentiate between the expected rewards from each of the competing intentions. When a rewarding (or aversive) event occurs, the intentions that led to the event will be strengthened (or weakened) within the striatum so that they are more (or less) likely to be selected the next time a similar environmental context is encountered.

Systems

As previously discussed, projections from posterior to anterior cortex can naturally encode associations between actions/intentions and features of the state that suggest that action. Multiple different, and potentially, conflicting intentions can be activated in parallel and it will often be necessary to select among conflicting actions. The neuroanatomy of corticostriatal circuits make them particularly well-suited to this function and interactions between cortex and basal ganglia will be crucial in doing so. Interactions between this system and the dopamine system will also be crucial for learning in TOSCA. Specifically, when an action leads to unexpected reward, the value of that action in the current state/context will be increased by potentiating the cortical associations between the state features and the action representation. The corticostriatal action-selection system will be sensitive to these values, so that when that action is proposed in similar states in the future, its probability of being selected will be higher.

Modulation of action contingencies via dopamine

Anatomical structure

Dopamine producing neurons are located in two midbrain nuclei, the ventral tegmental area (VTA) and the substantia nigra pars compacta (SNc). They receive excitatory input primarily from the pedunculopontine tegmental nucleus (PPTN), which conveys information about the occurrence of primarily rewarding events, and prefrontal cortex and inhibitory input from the

ventral striatum. They project to the prefrontal cortex and striatum where they fire in a phasic fashion to release dopamine in response to rewarding situations (Romo & Schultz 1990, Schultz 1996).

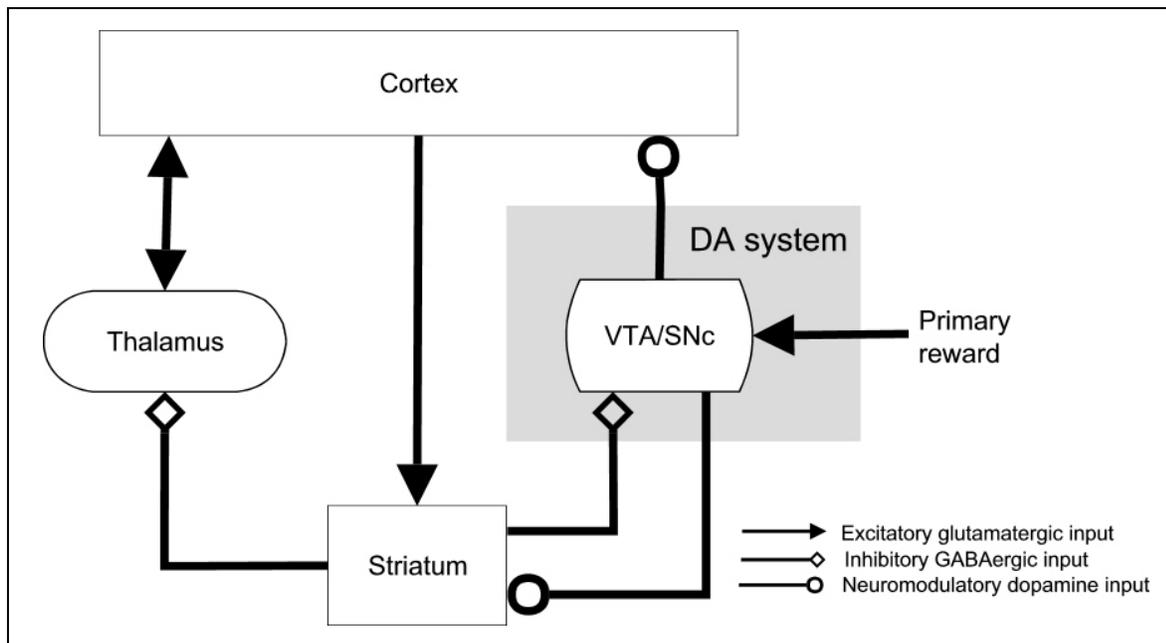


Figure 1.4: Schematic of main connections of the dopaminergic system

Physiological operation

The corticostriatal loops of the basal ganglia are the substrate for selecting between intentions. Learning of the correct intention in a given environmental context is under the control of the dopaminergic system shown in Figure 1.4. When a reward is encountered, the synaptic strengths in the corticostriatal circuits that were activated prior to the reward are increased. This makes it more likely that the same intention will be executed in a similar environmental context on future occasions. An unexpected (primary) reward elicits a phasic response in the dopaminergic neurons of the VTA/SNc. When a CS has been learned to reliably predict an upcoming reward, the time of response of the dopamine neurons shifts to coincide with the CS. These phasic releases of dopamine are utilized in the recipient structures to direct learning. The action of phasic dopamine signals is to increase synaptic strength by a 3-factor learning rule. In this rule, the relative timing of synaptic input, neuronal firing and dopamine pulse conspire to dictate the amount of learning from a single rewarding event.

Derived computational functionality

Dopamine neurons have long been associated with reward learning and rewarded behavior, partly because of clear evidence of their key role in drugs of addiction (DiChiara, 1999), and because they are among the best targets for self-stimulation. The observation that the activity of dopamine cells in the monkey midbrain in reward-learning tasks closely follows the form of a key training signal in reinforcement learning (the temporal difference prediction error), is an important backdrop for TOSCA. In particular, temporal difference based reinforcement learning

(RL) methods will serve to modulate state-action associations by potentiating associations between clusters in posterior cortex (representing complex internal state information) and clusters in anterior cortex (representing internal and external action intentions).

Systems

The dopamine system is tightly bound to the corticostriatal system, mediating learning in the prefrontal cortex and both divisions of the striatum. This system is also now known to provide neuromodulatory input to the hippocampal and thalamic systems.

Dopamine reward circuits (Intrinsic Reward and its Neural Basis)

Anatomical structure

Recent studies (Kakade & Dayan 2002, Dayan & Balleine 2002) have focused on the idea that dopamine not only plays a critical role in the extrinsic motivational control of behaviors aimed at harvesting explicit rewards, but also in the intrinsic motivational control of behaviors associated with novelty and exploration. For instance, salient, novel sensory stimuli inspire the same sort of phasic activity of dopamine cells as novel rewards (Schultz 1998, Horvitz et. al. 1997}. However, this activation extinguishes more or less quickly as the stimuli become familiar. This may underlie the fact that novelty itself has rewarding characteristics (Montague et al. 1996).

The novelty-based release of dopamine onto one of its major targets, the striatum, causes both general psychomotor activation (Hooks & Kalivas 1994) and also specific exploratory or seeking behaviors such as approach that cause animals to engage with those novel stimuli. Approach of this sort is a Pavlovian response---it is like a pre-wired action inspired by novelty (and also reward prediction). Theoretical treatments (Kakade & Dayan 2001, Kakade & Dayan 2002) have directly related the dopamine activity with mechanisms for controlling exploration in the RL literature such as exploration and shaping bonuses (Sutton, 1993, Dayan & Sejnowski 1996, Ng et. al. 1999) effectively completing the circle of interaction between computational, psychological and neural approaches. In TOSCA, we will explore a wider set of mechanisms by which animals control and benefit from exploration, using it to build sophisticated mechanisms for manipulating and exploiting novel environments. This wider set of mechanisms include the desire for mastery over one's environment and often leads to purposeful and sustained experimentation, as well as the motivation of an agent in a social setting to be liked by other agents (like-me) which leads to imitative behavior in social settings.

Various studies have also considered the neural basis of the assessment of novelty. Of particular relevance are two further neuromodulators, acetylcholine (ACh) and norepinephrine (NE), which are known to be involved in uncertainty and unexpectedness, and also to interact with the dopamine system. Theoretical treatments of these (Dayan & Yu 2003, Yu & Dayan 2002) focus on their roles in reporting specific sorts of uncertainty---uncertainty arising from ignorance (which is what should drive exploration) and uncertainty arising from environmental stochasticity (which should not). The difference between these forms of uncertainty is relative to models of the environment, which form a key component of any theory of novelty. The ideas on ACh and NE are in their infancy; there is scope for a productive interaction between our explorations via TOSCA and future experiments and theory on the drives and effects of NE and ACh.

Derived computational functionality

The intrinsic motivations listed above will serve as mechanisms for providing internal reward to the agent and this in turn will help direct the agent's behavior during exploration and play both in the presence and absence of externally specified tasks. These internal rewards will lead to the learning of useful mental and physical skills in the form of options or abstract actions that in turn will become available to the reinforcement learning system in TOSCA as actions. This will allow an incremental buildup of a hierarchy of useful cognitive and physical skills by the agent that would not be possible in the absence of intrinsic motivations.

Systems

The dopamine system is tightly bound to the corticostriatal system, mediating learning in the prefrontal cortex and both divisions of the striatum. This system is also now known to provide neuromodulatory input to the hippocampal and thalamic systems.

Cerebellum

Anatomical structure

The cerebellum can be subdivided into the cerebellar cortex and the deep cerebellar nuclei, which sit on top of the cerebellar peduncle. Figure 1.5 illustrates a schematic diagram of the major connections of the cerebellum. The largest subdivision of the cerebellar cortex in humans is the cerebrocerebellum which occupies most of the lateral cerebellar hemispheres and receives input from many areas of the cerebral cortex. The phylogenetically oldest part of the cerebellar cortex is the vestibulocerebellum, which comprises the caudal lobes. The third division is the spinocerebellum, which occupies the median and paramedian zone of the cerebellar hemispheres. The deep cerebellar nuclei are embedded within the white matter of the cerebellum. The connections between the cerebellum and other parts of the nervous system occur by way of three large pathways called the cerebellar peduncles. The middle cerebellar peduncle is an afferent pathway arising mainly in the pons and the superior cerebellar peduncle is an efferent pathway from the deep cerebellar nuclei to the thalamus.

The majority of cerebral cortical inputs to the cerebellum arise in the primary motor and premotor cortices of the frontal lobe, the primary and secondary somatic sensory cortices of the anterior parietal lobe and the secondary visual regions of the posterior parietal lobe. The cerebellum projects mainly to the upper motor neurons in the cerebral cortex via relay neurons in the thalamus.

Physiological operation

The cerebellum influences movements by modifying the activity patterns of the upper motor neurons. The primary function of the cerebellum is to detect the difference (or motor error) between an intended movement and the actual movement and, through its projections to the upper motor neurons, to reduce the error (Gluck et al 2001). These corrections can be made both during the course of a movement and as a form of motor learning when the correction is stored.

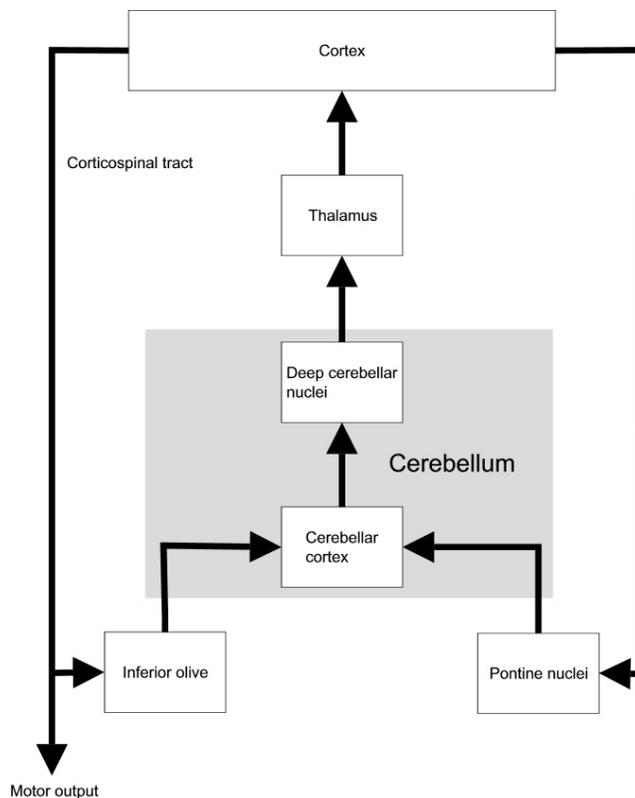


Figure 1.5: Schematic of the major connections of the cerebellum

From lesion studies it has been found that the cerebellar loop is critical for the performance of planned, voluntary, multi-joint movements. The activity of the cerebellum instructs the motor cortex in the direction, timing and force. For ballistic movements these instructions are based entirely on predictions about their outcome.

Derived computational functionality

In the TOSCA architecture, the cerebellum acts to store complex motor programs as they are learned. Individual movements, originally used as separate parts of a complex movement sequence, will be gradually compiled into motor programs in the cerebellum. These motor programs generate the appropriate motor sequences on demand and through supervised learning gradually make execution of the movement sequences smoother and better coordinated (Gluck et al 1994).

Systems

The cerebellum interacts primarily with the cerebral cortex. In early phases of motor learning, the motor programs will be simple and proposed by the cerebral cortex. When the intended action has been selected by the basal ganglia, the action will be executed by the primary motor cortex.

The cerebellum will receive information about the intended outcome of the action from the motor cortices and the outcome of execution of the action from the sensory cortices. The difference between the intention and the outcome will be used by the cerebellum for learning of the motor action. The next time the same motor action is proposed the cerebellum will have an influence on

the execution of the action and will use the error in the action execution to continue learning.

Conclusion

The previous section lays out the Rutgers University components of the broader team vision for TOSCA at the level of brain systems and circuits. It explores the physiology we are trying to capture in TOSCA as well as the low-level computation being performed in individual brain systems and in brain circuits. However, it is down at a level where it is often difficult to see how human-level behavior emerges from these components and their connections.

Two primary features of the design of TOSCA are its representational system and its control system. Learning permeates the operation of the TOSCA system: the system is continually learning and cannot help but learn, and thereby builds up representations from combinations of perception and prior knowledge, as well as building up control knowledge.

References

- Alexander GE, DeLong MR, Strick PL. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience* **9**, 357–38.
- Alexander GE, DeLong MR, Strick PL. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience* **9**, 357–38.
- Arbuthnott GW, MacLeod N, Rutherford A. The rat cortico-striatal pathway in vitro. *Journal of Physiology London* **367** (1985), 102P.
- Arbuthnott GW, MacLeod N, Rutherford A. The rat cortico-striatal pathway in vitro. *Journal of Physiology London* **367** (1985), 102P.
- Barsalou, L.W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101-140). Cambridge: Cambridge University Press.
- Barsalou, L.W. (1989). Intraconcept similarity and its implications for interconcept similarity. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 76-121). Cambridge: Cambridge University Press.
- Barsalou, L.W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A.C. Collins, S.E. Gathercole, & M.A. Conway (Eds.), *Theories of memory* (pp. 29-101). London: Lawrence Erlbaum Associates.
- Barsalou, L.W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 577-660.
- Barsalou, L.W. (2003a). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society of London: Biological Sciences*, *358*, 1177-1187.
- Barsalou, L.W. (2003b). Situated simulation in the human conceptual system. *Language and Cognitive Processes*, *18*, 513-562.
- Barsalou, L.W. (2005). Abstraction as dynamic interpretation in perceptual symbol systems. In L. Gershkoff-Stowe & D. Rakison (Eds.), *Building object categories* (389-431). Carnegie Symposium Series. Mahwah, NJ: Erlbaum.
- Barsalou, L.W., Niedenthal, P.M., Barbey, A., & Ruppert, J. (2003). Social embodiment. In B. Ross (Ed.), *The Psychology of Learning and Motivation*, Vol. 43 (pp. 43-92). San Diego: Academic Press.
- Barsalou, L.W., Simmons, W.K., Barbey, A.K., & Wilson, C.D. (2003). *Grounding conceptual knowledge in modality-specific systems. Trends in Cognitive Sciences*, *7*, 84-91.
- Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, *132*, 163-201.
- Damasio, A.R. (1989). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition*, *33*, 25-62.
- Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of Experimental Psychology: General*, *120*, 339-357.
- Gluck, M. A., Myers, C. E., & Thompson, R. F. (1994). A computational model of the cerebellum and motor-reflex learning. In S. Zornetzer, J. Davis, T. McKenna, & C. Lau (Editors). *An Introduction to Neural and Electronic Networks* (Second Edition). 91-80.
- Gluck, M.A., Allen, M.T., Myers, C.E., & Thompson, R.F. (2001) Cerebellar substrates for error-correction in motor-reflex conditioning. *Neurobiology of Learning and Memory*, *76*, 314-341.
- Kemp JM, Powell TPS. The corticostriate projection in the monkey. *Brain* **93** (1970), 525-546.
- Kemp JM, Powell TPS. The corticostriate projection in the monkey. *Brain* **93** (1970), 525-546.
- Kosslyn, S.M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.
- Kosslyn, S.M. (1994). *Image and brain*. Cambridge, MA: MIT Press.
- Martin, A. (2001). Functional neuroimaging of semantic memory. In R. Cabeza & A. Kingstone

- (Eds.), *Handbook of functional neuroimaging of cognition* (pp. 153-186). Cambridge, MA: MIT Press.
- McRae, K., de Sa, V.R., & Seidenberg, M.S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*, 99-130.
- Montague PR, Dayan P, Sejnowski TJ. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of Neuroscience*, **16** (1996), 1936–1947.
- Montague PR, Dayan P, Sejnowski TJ. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of Neuroscience*, **16** (1996), 1936–1947.
- Niedenthal, P.M., Barsalou, L.W., Winkielman, P., Krauth-Gruber, S., & Ric, F. (2005). Embodiment in attitudes, social perception, and emotion. *Personality and Social Psychology Review*, *9*, 184-211.
- Pulvermüller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, *22*, 253-336.
- Romo R, Schultz W. Dopamine neurons of the monkey midbrain: contingencies of responses to active touch during self-initiated arm movements. *The Journal of Neurophysiology* **63** (1990), 592-606.
- Romo R, Schultz W. Dopamine neurons of the monkey midbrain: contingencies of responses to active touch during self-initiated arm movements. *The Journal of Neurophysiology* **63** (1990), 592-606.
- Schultz W. Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey. *The Journal of Neurophysiology* **56** (1986), 1439-1461.
- Schultz W. Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey. *The Journal of Neurophysiology* **56** (1986), 1439-1461.
- Schyns, P.G., Goldstone, R.L., & Thibaut, J.P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, *21*, 1-54.
- Simmons, W.K., & Barsalou, L.W. (2003). The similarity-in-topography principle: Reconciling theories of conceptual deficits. *Cognitive Neuropsychology*, *20*, 451-486.
- Smith, L., & Gasser, M. (2005). The development of embodied cognition: Six lessons from babies. *Artificial Life*, *11*, 13-29.
- Sutton, R.S. (1990). Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. Proceedings of the Seventh International Conference on Machine Learning, pp. 216-224, Morgan Kaufmann.
- Sutton, R.S., Precup, D., Singh, S. (1999). Between MDPs and semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artificial Intelligence* *112*:181-211
- Thelen, E. (2000). Grounded in the world: Developmental origins of the embodied mind. *Infancy*, *1*, 3-30.
- Thompson-Schill, S.L. (2003). Neuroimaging studies of semantic memory: inferring “how” from “where”. *Neuropsychologia*, *41*, 280-292.
- D. E. Berlyne. A theory of human curiosity. *British Journal of Psychology*, *45*:180–191, 1954.
- D. E. Berlyne. *Conflict, Arousal, and Curiosity*. McGraw-Hill, N.Y., 1960.
- D. E. Berlyne. *Aesthetics and Psychobiology*. Appleton-Century-Crofts, N.Y., 1971.
- D. O. Hebb. *The Organization of Behavior*. Wiley, N.Y., 1949.
- G. Di Chiara. Drug addiction as dopamine-dependent associative learning disorder. *European Journal of Pharmacology*, *375*(1-3):13–30, 1999.
- G. Barto. Adaptive critics and the basal ganglia. In J. C. Houk, J. L. Davis, and D. G. Beiser, editors, *Models of Information Processing in the Basal Ganglia*, pages 215–232. MIT Press, Cambridge, MA, 1995.
- K. J. Friston, G. Tononi, G. N. Reeke, O. Sporns, and G. M. Edelman. Value-dependent selection

- in the brain: Simulation in a synthetic neural model. *Neuroscience*, 59:229–243, 1994.
- S. Kakade and P. Dayan. Dopamine bonuses. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 131–137. MIT Press, 2001.
- S. Kakade and P. Dayan. Dopamine: Generalization and bonuses. *Neural Networks*, 15:549–559, 2002.
- Ng, D. Harada, and S. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning*. Morgan Kaufmann, 1999.
- P. R. Montague, P. Dayan, and T. J. Sejnowski. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of Neuroscience*, 16:1936–1947, 1996.
- P. Dayan and B. W. Balleine. Reward, motivation and reinforcement learning. *Neuron*, 36:285–298, 2002.
- P. Dayan and T. J. Sejnowski. Exploration bonuses and dual control. *Machine Learning*, 25:5–22, 1996.
- P. Dayan and A. J. Yu. Uncertainty and learning. *IETE Journal of Research*, 49:171–182, 2003.
- S. Kakade and P. Dayan. Dopamine bonuses. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 131–137. MIT Press, 2001.
- S. Kakade and P. Dayan. Dopamine: Generalization and bonuses. *Neural Networks*, 15:549–559, 2002.
- M. S. Hooks and P. W. Kalivas. Involvement of dopamine and excitatory amino acid transmission in novelty- induced motor activity. *Journal of Pharmacology & Experimental Therapeutics*, 269:976–988, 1994.
- J. C. Horvitz, T. Stewart, and B. Jacobs. Burst activity of ventral tegmental dopamine neurons is elicited by sensory stimuli in the awake cat. *Brain Research*, 759:251–258, 1997.
- W. Schultz. Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80:1–27, 1998.
- W. Schultz, P. Dayan, and P. R. Montague. A neural substrate of prediction and reward. *Science*, 275:1593–1598, March 1997.
- R. S. Sutton, D. Precup, and S. Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112:181–211, 1999.
- J. Yu and P. Dayan. Acetylcholine in cortical inference. *Neural Networks*, 15:719–730, 2002.
- S. Singh, M. R. James, and M. R. Rudary. Predictive State Representations: A New Theory for Modeling Dynamical Systems In *Uncertainty in Artificial Intelligence: Proceedings of the Twentieth Conference (UAI)*, pages 512-519, 2004.
- S. Singh, A. G. Barto, and N. Chentanez. Intrinsically Motivated Reinforcement Learning. In *Proceedings of Advances in Neural Information Processing Systems 17*, 2005.

Acronym List

ACh – Acetylcholine
BG – Basal Ganglia
BICA – Biologically Inspired Computer Architecture
CA1 – Cornu Ammonis 1
CA3 – Cornu Ammonis 3
CS – Conditioned Stimulus
DG – Dentate Gyrus
EC – Entorhinal Cortex
GPe – Globus Pallidus externa
GPi – Globus Pallidus interna
NE – Norepinephrine
PPTN – PedunculoPontine Tegmental Nucleus
RL – Reinforcement Learning
SNc – Substantia Nigra pars compacta
SNr – Substantia Nigra pars reticulata
STN – Triangular Septal Nucleus
TOSCA –
VTA – Ventral Tegmental Area