Psychological Review
2007, Vol. 114, No. 3, 632–656

0033-295X/07/$12.00 DOI: 10.1037/0033-295X.114.3.632

A Neurobiological Theory of Automaticity in Perceptual Categorization

F. Gregory Ashby, John M. Ennis, and Brian J. Spiering
University of California, Santa Barbara

A biologically detailed computational model is described of how categorization judgments become automatic in tasks that depend on procedural learning. The model assumes 2 neural pathways from sensory association cortex to the premotor area that mediates response selection. A longer and slower path projects to the premotor area via the striatum, globus pallidus, and thalamus. A faster, purely cortical path projects directly to the premotor area. The model assumes that the subcortical path has greater neural plasticity because of a dopamine-mediated learning signal from the substantia nigra. In contrast, the cortical-cortical path learns more slowly via (dopamine independent) Hebbian learning. Because of its greater plasticity, early performance is dominated by the subcortical path, but the development of automaticity is characterized by a transfer of control to the faster cortical-cortical projection. The model, called SPEED (Subcortical Pathways Enable Expertise Development), includes differential equations that describe activation in the relevant brain areas and difference equations that describe the 2- and 3-factor learning. A variety of simulations are described, showing that the model accounts for some classic single-cell recording and behavioral results.

Keywords: striatum, skill learning, habit learning, dopamine, expertise

Humans have a remarkable ability to categorize a huge number of objects instantly and effortlessly. For example, when viewing a German Shepard, we immediately respond *dog* rather than *wolf*, even though such a categorization might require integrating perceptual information about the shape and size of the ears, the length, coarseness, and color of the hair, the size of the body, and many other perceptual features. Yet, this all happens instantly, with little or no awareness of any specific features. How can such judgments be made so easily and quickly? And what is their neural basis? The present article proposes answers to these questions.

Within cognitive psychology, the modern standard for determining that a behavior has become automatic is if it can be performed in parallel and without attention (Schneider & Shiffrin, 1977). Interest in the neural basis of automaticity dates back at least to Sherrington (1906), who proposed that automatic behaviors become reflexive and that by chaining strings of simple reflexes together, complex automatic behaviors could be produced. Sherrington's ideas led to the theory that dominated the 20th century: Novel behaviors require attention and flexible thinking and therefore are dependent on cortex, whereas automatic behaviors require neither of these and so are not mediated primarily by cortex. Instead, it has long been assumed that automatic behaviors are primarily mediated by subcortical structures. For example, in his classic and influential article entitled "In search of the engram," Lashley (1950) wrote that "it has been widely held that although memory traces are at first formed in the cerebral cortex, they are finally reduced or transferred by long practice to subcortical levels" (p. 466). This view is still widely held. For example, in a recent review of the prefrontal cortex (PFC), Fuster (2001) wrote that "routine, automatic, or overlearned behavioral sequences, however complex, do not engage the PFC and may be entirely organized in subcortical structures" (p. 323). In the present article, we argue exactly the opposite position—namely, that novel behaviors are mediated primarily by subcortical structures (i.e., the basal ganglia), whereas control of these behaviors is passed to cortex once automaticity is attained.

To our knowledge, this is the first article to propose a neurobiologically detailed theory of categorization automaticity. There is a large literature on how expertise[1] with a category affects the representation of its members in visual cortex (e.g., Gauthier & Tarr, 1997; Humphreys & Forde, 2001; Joseph, 2001). However, categorization is the ability to respond differently to objects in separate classes or categories, and thus categorization requires linking a percept to an action. There is good reason to believe that this linkage is not encoded in visual cortex (see, e.g., Ashby & Spiering, 2004). For example, when the category labels of two stimuli are switched (from *good* to *bad* , and vice versa), the firing properties of cells in the inferotemporal cortex of monkeys that are sensitive to those stimuli do not change (Rolls, Judge, & Sanghera, 1977). More recent studies have found similar null results with traditional categorization tasks (Freedman, Riesenhuber, Poggio, & Miller, 2003; Op de Beeck, Wagemans, & Vogels, 2001; Sigala, 2004; Thomas, Van Hulle, & Vogels, 2001; Vogels, 1999). In each

F. Gregory Ashby, John M. Ennis, and Brian J. Spiering, Department of Psychology, University of California, Santa Barbara.

This research was supported in part by Public Health Service Grant MH3760-2. We thank Jon Horvitz for his helpful suggestions. All of the software used in the simulations reported in this article is available online at http://www.psych.ucsb.edu/%7Eashby/numrout.htm

Correspondence concerning this article should be addressed to F. Gregory Ashby, Department of Psychology, University of California, Santa Barbara, CA 93106. E-mail: ashby@psych.ucsb.edu

[1] In this article, we use the terms *automaticity* and *expertise* interchangeably. In both cases, we simply mean a state that is acquired following extensive experience with a particular task. Expertise is defined similarly in much of the literature. However, it is important to note that in some fields (but not here), expertise also connotes some extra unique training (i.e., beyond mere experience).

case, single-cell recordings showed that the firing properties of cells in inferotemporal cortex did not change as monkeys learned to classify visual objects into one of two categories. The cells showed sensitivity to specific visual images, but category training did not make them more likely to respond to other stimuli in the same category, or less likely to respond to stimuli belonging to the contrasting category.

For these reasons, the theory of categorization automaticity that we propose is only partly motivated by the large and compelling literature on object representation in visual cortex. Instead, it was primarily motivated by the equally large literature on category learning (for a recent review, see, e.g., Ashby & Maddox, 2005).

Early theories of categorization virtually all assumed that humans have a single category-learning system that they use to learn all types of categories (for an exception, see Brooks, 1978). The dominant theories of this era were prototype theory (e.g., Homa, Sterling, & Trepel, 1981; Posner & Keele, 1968, 1970; Reed, 1972; Rosch, 1973, 1975; Smith & Minda, 1998), exemplar theory (Brooks, 1978; Estes, 1986, 1994; Hintzman, 1986; Lamberts, 2000; Medin & Schaffer, 1978; Nosofsky, 1986), and decision bound theory (Ashby & Gott, 1988; Ashby & Townsend, 1986; Maddox & Ashby, 1993). More recently, however, the category-learning field has been dominated by theories that assume that the learning of different types of category structures is mediated by different systems (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Brooks, 1978; Erickson & Kruschke, 1998; Nosofsky, Palmeri, & McKinley, 1994). Whereas most multiple systems theorists agree that one system is explicit and another is implicit, there is disagreement about the nature of the implicit system. Some argue for an exemplar-based system (e.g., Erikson & Kruschke, 1998; Nosofsky et al., 1994), some for a perceptual representation system (e.g., Reber, Stark, & Squire, 1998), and others for a procedural learning-based system (e.g., Ashby et al., 1998; Ashby, Ell, & Waldron, 2003; Ashby & Waldron, 1999). A likely possibility is that each proposal has some validity, and multiple implicit category-learning systems may exist (Ashby & O'Brien, 2005).

The theory of automaticity developed in the present article was motivated by a category-learning theory called COVIS (Ashby et al., 1998; Ashby & Valentin, 2005; Ashby & Waldron, 1999). Briefly, COVIS postulates two systems that compete throughout learning—a frontal-based explicit system that uses logical reasoning and depends on working memory and executive attention and a basal ganglia-mediated implicit system that uses procedural learning. A wide variety of category structures are learned in the procedural learning system, but learning occurs in a slow incremental fashion and is highly dependent on reliable and immediate feedback. In contrast, a fairly small set of category structures are learned quickly in the explicit rule-based system—specifically, those structures that can be learned via an explicit reasoning process. Tasks that require participants to learn such structures are called *rule-based category-learning tasks*. Frequently, the rule that maximizes accuracy (i.e., the optimal strategy) is easy to describe verbally. In the most common applications, only one stimulus dimension is relevant, and the participant's task is to discover this relevant dimension and then to map the different dimensional values to the relevant categories (e.g., as in the Wisconsin Card Sorting Test).

However, there are many category structures that cannot be learned in the explicit system. An important example occurs in *information-integration tasks*, in which learning requires that participants integrate perceptual information across two or more noncommensurable stimulus dimensions at some predecisional stage (Ashby & Gott, 1988). Perceptual integration could take many forms—from computing a weighted linear combination of the dimensional values to treating the stimulus as a Gestalt. Typically, the optimal strategy in information-integration tasks is difficult or impossible to describe verbally (which makes it difficult to discover via logical reasoning). Real-world examples of information-integration tasks are common. For example, deciding whether an x-ray shows a tumor not only involves some explicit reasoning but also shares many properties with information-integration tasks. For example, years of training are required, and expert radiologists are only partially successful at describing their categorization strategies.

The present article specifically addresses the development of automaticity in tasks in which the procedural learning system dominates. We begin by reviewing COVIS and especially its procedural learning system. We then describe a new model called SPEED (subcortical pathways enable expertise development), which extends COVIS to expert or automatic behaviors. Next, we describe a number of empirical tests of this new model and consider how it relates to existing models of expertise. Finally, we close with some general discussion and conclusions.

The Procedural Learning System of COVIS

Overview

Figure 1 shows the procedural learning system of COVIS (Ashby et al., 1998; Ashby & Waldron, 1999). The key structure is the striatum, which is a major input region of the basal ganglia

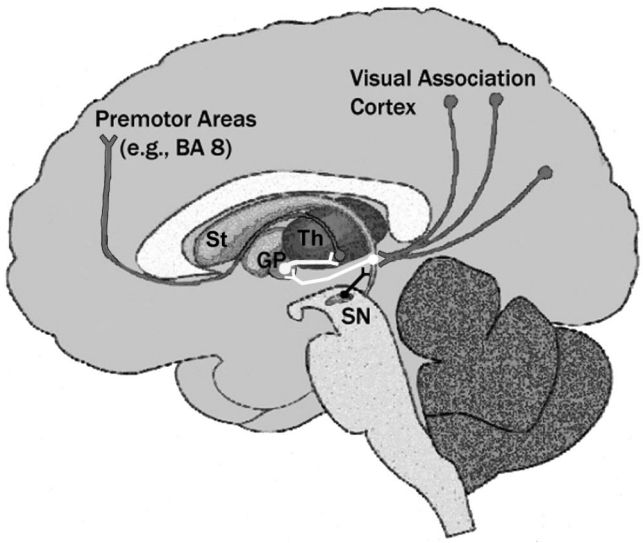

Figure 1. A schematic illustrating the procedural-learning system of COVIS. Gray projections are excitatory (glutamate), white projections are inhibitory (GABA), and black projections are dopaminergic; circles denote cell bodies, and forks denote axon terminals; SN = substantia nigra pars compacta; GP = internal segment of the globus pallidus; Th = thalamus; St = striatum (the striatal region shown here is the caudate nucleus); and BA 8 = Brodmann's Area 8.

and comprises a number of structures, including the caudate nucleus and putamen. For visual categories, the critical structure within the striatum is the caudate nucleus because (in primates) all of extrastriate visual cortex projects directly to the body and tail of the caudate. More important, these projections are characterized by massive convergence, with about 10,000 visual cortical cells converging on each caudate (medium spiny) cell (Wilson, 1995). COVIS assumes that, through a procedural learning process, each striatal unit associates an abstract motor program with a large group of visual cortical cells (i.e., all that project strongly to it).

The medium spiny cells in the body and tail of the caudate send projections to a variety of prefrontal and premotor cortical areas. There are two synapses on this pathway. The first synapse on the principal path is in the internal segment of the globus pallidus (and substantia nigra pars reticulata), which is a major output structure within the basal ganglia. The second synapse is in the thalamus, primarily in the ventral anterior nucleus, pars magnocellualaris (VAmc). The primary cortical projection from VAmc is to premotor areas and specifically to Brodmann Area 8 and the so-called supplementary eye fields (Shook, Schlag-Rey, & Schlag, 1991).

Figure 2 shows a close-up view of a synapse between the axon of a pyramidal cell originating in visual cortex and the dendrite of a medium spiny cell in the striatum. Note that glutamate projections from visual cortex and dopamine projections from the substantia nigra (pars compacta) both synapse on the dendritic spines of striatal medium spiny cells (e.g., DiFiglia, Pasik, & Pasik, 1978; Freund, Powell, & Smith, 1984; Smiley, Levey, Ciliax, & Goldman-Rakic, 1994). COVIS assumes these synapses are the critical site of learning in the procedural system.

Much is now known about how these cortical-striatal synapses are strengthened (e.g., via long-term potentiation [LTP]) and weakened (e.g., via long-term depression [LTD]). Specifically, the best available evidence indicates that three factors are required to strengthen synapses of the type shown in Figure 2: (a) strong presynaptic activation, (b) strong postsynaptic activation, and (c) dopamine release (e.g., Arbuthnott, Ingham, & Wickens, 2000; Calabresi, Pisani, Mercuri, & Bernardi, 1996; Nairn, Hemmings, Walaas, & Greengard, 1988; Pessin et al., 1994; Wickens, 1990, 1993). The first two factors guarantee that postsynaptic *N*-methyl-*D*-aspartate (NMDA) receptors are activated. The NMDA receptor is a glutamate receptor with a high threshold for activation that plays a critical role in LTP (e.g., Malenka & Nicoli, 1999). A high threshold is unlikely to be met by noise, so Factors 1 and 2 mean that only synapses driven by sensory cortical cells that respond to the stimulus are likely to be strengthened.

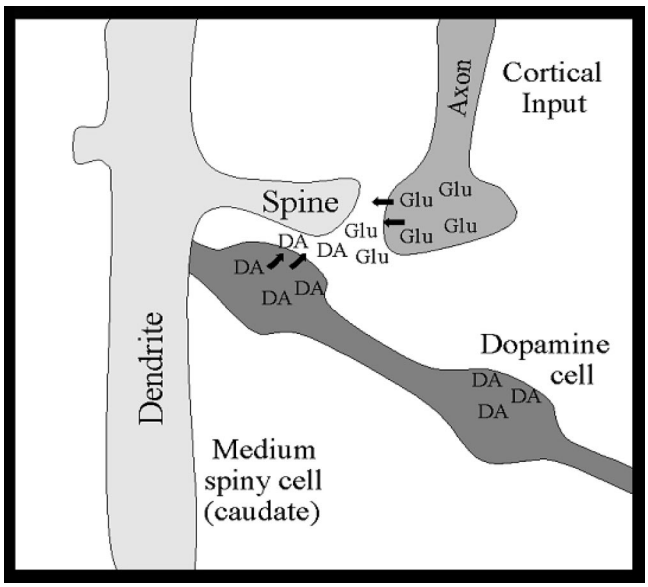

Figure 2. Close-up of a cortical-striatal synapse. The axon of a visual cortical cell synapses on the spine of a medium spiny cell in the caudate nucleus. A dopamine cell from the substantia nigra pars compacta releases dopamine into the cortical-striatal synaptic junction. Glu = glutamate; DA = dopamine.

The third factor, dopamine, is widely thought to serve as a reward-mediated training signal (e.g., Miller, Sanghera, & German, 1981; Montague, Dayan, & Sejnowski, 1996; Wickens, 1993). For example, a large literature shows that dopamine firing increases above baseline following unexpected reward and is depressed below baseline following the omission of an expected reward (e.g., Hollerman & Schultz, 1998; Mirenowicz & Schultz, 1994; Schultz, 2002). In addition, there is good evidence that elevated dopamine levels facilitate the strengthening of synapses (i.e., LTP) at which postsynaptic NMDA receptors have been activated (e.g., Kerr & Wickens, 2001). In summary, according to COVIS, synapses in the procedural learning system between visual cortical cells and medium spiny cells in the caudate are strengthened if the visual cortical cell responds strongly to the presented stimulus (Factors 1 and 2) and the participant is rewarded for responding correctly (Factor 3).

However, if either of Factors 2 or 3 is missing, then it is thought that the strength of the synapse will weaken (i.e., LTD will occur; e.g., Arbuthnott et al., 2000; Calabresi et al., 1996). Note that there are several ways this could happen. One is if the participant responds incorrectly (so Factor 3 is missing), and another is if the visual cortical cell responds only weakly to the presented stimulus. Thus, this model of LTD predicts that any synapse responsible for the participant emitting an incorrect response will be weakened, as will any synapse that is driven by a cell in visual cortex that does not respond strongly to the presented stimulus. The combination of this with the three-factor model of LTP produces a powerful learning algorithm.

The three-factor model of LTP is appealing, but a serious timing problem must be solved before it can operate effectively. The problem is that shortly after the stimulus is presented, the cortical-striatal synapse will be activated, but the dopamine release must necessarily occur some time later because dopamine release follows the reward, which follows the response, which follows stimulus presentation. Fortunately, because the spines are somewhat separated from the bulk of the intracellular medium, it takes several more seconds to reset the membrane potential in the spines than in the main cellular compartments (Gamble & Koch, 1987; MacDermott et al., 1986). Thus, so long as the reward is delivered within a few seconds of the response, a trace (e.g., calcium-dependent protein kinase; Hemmings, Walaas, Ouimet, & Greengard, 1987; Pessin et al., 1994) will still exist in the critical spines that were responsible for eliciting the behavior that earned the reward, and so the correct synapses will be strengthened. Note that an obvious and exceptionally strong prediction of this model is that if the feedback is delayed more than a few seconds, then learning

should be severely disrupted in information-integration tasks. This prediction was supported by Maddox, Ashby, and Bohil (2003), who found that delays as short as 2.5 s severely interfered with information-integration learning, but delays as long as 10 s had no effect on learning in rule-based tasks of equal difficulty.

Empirical Support for COVIS

A wide variety of evidence supports COVIS as a model of category learning, at least in rule-based and information-integration tasks. This evidence comes from single-cell recording studies, animal lesion experiments, studies with various neuropsychological patient groups, neuroimaging studies, and traditional cognitive behavioral experiments. A thorough review of all these data is beyond the scope of this article (for recent reviews, see Ashby & Maddox, 2005; Ashby & Valentin, 2005; Maddox & Ashby, 2004; Maddox & Filoteo, 2005). Instead, we briefly review some of the evidence implicating the basal ganglia in category learning (for a complete review, see Ashby & Ennis, 2006) and some behavioral studies in which a priori predictions of the model were tested.

First, lesion studies in rats and monkeys support the hypothesis that the tail of the caudate nucleus[2] is both necessary and sufficient for visual discrimination learning. Necessity was supported in a series of studies showing that lesions of the tail of the caudate impair the ability of animals to learn visual discriminations (e.g., McDonald & White, 1993, 1994; Packard, Hirsch, & White, 1989; Packard & McGaugh, 1992). Sufficiency, which is always more difficult to establish, was supported in a series of studies in which all pathways out of visual cortex were lesioned, except those into the tail of the caudate (e.g., projections into PFC were lesioned by Eacott & Gaffan, 1991, and Gaffan & Eacott, 1995; projections to the hippocampus and amygdala were lesioned by Gaffan & Harrison, 1987). Critically, none of these lesions affected visual discrimination learning. Another related line of work showed that visual discrimination learning is not mediated by medial temporal lobe structures (McDonald & White, 1993, 1994; Packard et al., 1989). Later in the present article, we model some monkey single-unit recording data that also strongly implicates the striatum in category learning.

Human neuropsychological and neuroimaging studies support these same conclusions. For our purposes, the two most relevant neuropsychological conditions are Parkinson's disease (PD), in which basal ganglia dopamine levels are depleted, and Huntington's disease, in which the medium spiny cells of the caudate nucleus die. Both groups are impaired in information-integration category learning (Filoteo, Maddox, & Davis, 2001a; Filoteo, Maddox, Salmon, & Song, 2005; Knowlton, Mangels, & Squire, 1996). As an important control, Filoteo, Maddox, and Davis (2001b) reported normal performance by amnesiacs in a difficult multiday information-integration task with nonlinearly separable categories that required hundreds of training trials. In the functional magnetic resonance imaging (fMRI) domain, Seger and Cincotta (2002) trained participants in an information-integration task before scanning and reported significant striatal and lateral occipital activation during performance of the task. Nomura et al. (2007) had participants learn either rule-based or information-integration category structures during fMRI. Successful rule-based learning was associated with increased activation in frontal areas, including the anterior medial temporal lobes, whereas successful information-integration learning was associated with increased activation in the body of the caudate nucleus.

There have also been a number of cognitive behavioral studies in which specific parameter-free a priori predictions made by COVIS have been tested. One set of experiments tested COVIS predictions about the nature and timing of trial-by-trial feedback about response accuracy. In particular, COVIS predicts that, because the explicit system has access to working memory and executive attention, it should be relatively unaffected by changes in the timing and form of the feedback signal. In contrast, as described above, COVIS predicts that effective learning is possible in the procedural system only if feedback is provided immediately following the response. Several studies support these predictions. First, in the absence of any trial-by-trial feedback about response accuracy, people can learn some rule-based categories, but there is no evidence that they can learn information-integration categories (Ashby, Queller, & Berretty, 1999). Second, observational training (in which the category label is shown before stimulus presentation) is as effective as traditional feedback training with rule-based categories, but with information-integration categories, feedback training is significantly more effective than observational training (Ashby, Maddox, & Bohil, 2002). Third, as mentioned above, information-integration category learning is impaired if the feedback signal is delayed by as little as 2.5 s after the response, whereas delays as long as 10 s had no effect on rule-based category learning (Maddox et al., 2003).

Another series of studies tested the fundamental assumption of COVIS that information-integration categorization uses procedural learning, whereas rule-based category learning does not. In the first study, neither switching hands on the response keys nor switching the keys interfered with rule-based category learning, whereas with information-integration categories, switching hands on the response keys caused no interference, but switching the locations of the response keys caused a significant decrease in accuracy (Ashby et al., 2003). Thus, it appears that response locations are learned in information-integration categorization but not specific motor programs. This hypothesis was supported in a separate study by Maddox, Bohil, and Ing (2004). These information-integration results essentially replicate results found with traditional procedural learning tasks (Willingham, Wells, Farrell, & Stemwedel, 2000), and therefore they provide direct evidence of procedural learning in perceptual categorization.

A third series of studies tested the COVIS prediction that working memory and executive attention are critical for rule-based category learning but not for information-integration category learning. First, several studies reported that a simultaneous secondary task requiring working memory and executive attention had a massive detrimental effect on the ability of participants to learn simple rule-based categories, but it had no significant effect on learning information-integration categories (Waldron & Ashby, 2001; Zeithamova & Maddox, 2006).

Second, Maddox, Ashby, Ing, and Pickering (2004) tested the COVIS prediction that feedback processing requires attention and

[2] In rats, the caudate and putamen are not distinct. So in the rat studies, the lesions were to an area of the striatum homologous to the primate tail of the caudate.

effort in rule-based category learning but not in information-integration category learning. Specifically, they showed that limiting the amount of time available to process the feedback signal significantly interfered with rule-based category learning, but it had no effect on information-integration learning.

The results suggesting that working memory and executive attention play at most a minor role in information-integration category learning argue against a major role for the PFC in these tasks. Even so, this view is challenged by some well-known single-unit recording results of Freedman, Riesenhuber, Poggio, and Miller (2001, 2002; Freedman et al., 2003), who reported neurons in lateral PFC that show category-specific responding in a task in which monkeys were taught to categorize computer-generated images as dogs or cats. On each trial, the monkey was shown two successive images, separated by a brief delay. If the two images were from the same category, then the monkey's task was to release a lever, whereas if the images were from different categories, then the monkey was to keep holding the lever. Single-unit recordings revealed many neurons in lateral PFC that seemed to respond to the category membership of the visual stimulus; that is, each of these cells responded almost equally to almost all members of one category and showed little or no response to members of the contrasting category.

There are at least two interpretations of the Freedman et al. (2001, 2002, 2003) results that are consistent with all other results described in the present article. One possibility is that the monkeys discovered a simple rule that allowed them to respond correctly. If so, then the Freedman et al. results are consistent with all the other results implicating the PFC in rule learning and rule use. A second possibility, however, is that the monkeys learned to categorize the images using some similarity-based system and that the results of each categorization judgment were loaded into PFC-based working memory (e.g., as in the model of Ashby, Ell, Valentin, & Casale, 2005). It is important to note, however, that if a similarity-based process was used, then it is not likely to have been the procedural learning system of COVIS. This is because in the Freedman et al. design, when the stimuli were in the category "dogs," sometimes the monkeys released the lever and sometimes they held it. The same was true for the category "cats.' Thus, there was no consistent mapping from category to response. As mentioned above, several studies have shown that a consistent category–response mapping is required for efficient information-integration category learning (Ashby et al., 2003; Maddox et al., 2004).

SPEED: A New Model of Automaticity

A Neural Account of Automaticity

One way to develop a neural model of automatic categorization, at least in information-integration tasks, might be to begin with the COVIS procedural learning system. Unfortunately, however, there are many reasons to believe that the neural mechanisms and pathways that mediate category learning are different from the neural structures that mediate automatic responses to highly learned categories. Thus, COVIS is inappropriate as a model of expertise. Later in the present article, we specifically model several data sets that convincingly demonstrate the inability of COVIS to model highly overlearned behaviors. For now, though, we simply describe one suggestive result. In particular, many neuropsychological groups that are impaired in category learning (e.g., frontal patients and PD patients) do not lose old, familiar categories (e.g., fruits and tools). Similarly, there is no evidence that people who lose a familiar category (i.e., who develop a category-specific agnosia) develop any general category-learning deficits (although we know of no studies that directly address this issue).

One way to begin developing a model of categorization expertise is to ask what neural events must transpire when a categorization response is made automatically. First, if the categorization stimulus is visual, then it is natural to assume there must be activation in some visual cortical region. And if a motor response is made, then there presumably must eventually be activation in some region of premotor or motor cortex. The theoretical difficulty is in determining how the signal makes its way from visual cortex to premotor cortex. COVIS assumes that the critical path is from visual cortex to the basal ganglia, then to the thalamus, and finally to premotor cortex. But there are also direct cortical-cortical projections from visual association areas to premotor cortex (e.g., Heimer, 1995). Because these involve only one synapse, they could be faster than the subcortical path proposed by COVIS (which includes at least four synapses).

COVIS postulates that procedural learning depends on the longer subcortical path because of the reward-mediated training signal provided by the dopamine input from the substantia nigra. This three-factor learning causes strengthening of synapses that were active during correct responses and weakening of synapses that were active during errors. As it happens, there is also a prominent (and well-known) dopamine projection into frontal cortex from the ventral tegmental area (VTA), and the best available evidence indicates that the environmental conditions that cause VTA dopamine cells to fire are essentially the same as the conditions that cause substantia nigra dopamine cells to fire (Schultz, 1998; Schultz, Dayan, & Montague, 1997). In other words, VTA dopamine cells will also increase their firing following an unexpected reward, and their firing will be depressed below baseline following an unexpected failure to receive a reward. So an obvious question is why is the subcortical pathway needed? Why is category learning not mediated by three-factor learning at cortical-cortical synapses?

A necessary feature of any reward-mediated training signal is high-temporal resolution. If the first response is correct, then dopamine must be released into the relevant synapses quickly, before the critical traces disappear. But after the correct synapses have been strengthened, it is also essential that excess dopamine be quickly cleared from the synapse. If it is not, and the next response is an error, then the residual dopamine will strengthen inappropriate synapses—namely, those responsible for producing the incorrect response. This would undo the beneficial learning that occurred following correct responses and prevent discrimination learning.

Within the body and tail of the caudate nucleus, dopamine reuptake is exceptionally fast (e.g., Cragg, Rice, & Greenfield, 1997), and many researchers have proposed the same role for striatal dopamine in three-factor learning that we propose here (e.g., Miller et al., 1981; Montague et al., 1996; Wickens, 1993). In contrast, in frontal cortex, dopamine reuptake is much slower (for reviews, see, e.g., Brannon & Roth, 1983; Seamans & Yang, 2004). For example, the delivery of a single food pellet to a hungry

rat increases dopamine levels in PFC for approximately 30 mins (Feenstra & Botterblom, 1996). We contend that this poor temporal resolution effectively rules out dopamine as a reward-mediated training signal in frontal cortex.[3]

To be clear, we are not arguing that dopamine does not facilitate LTP at frontal cortical synapses. In fact, the biochemical cascade initiated by dopamine in the medium spiny cells of the caudate that eventually promotes LTP will likely have the same effect in frontal cortical glutamate cells. Thus, if the first response is correct, then frontal cortical dopamine levels should rise, and the synapses responsible for eliciting the correct response should be strengthened. However, the elevated cortical dopamine levels that last for many minutes should cause all synapses that are subsequently active to be strengthened, whether their activation elicited a correct response or an error. From a computational perspective, this means that learning in frontal cortex is governed by classical two-factor, or Hebbian learning; that is, a synapse is strengthened anytime there is both pre- and postsynaptic activation. With Hebbian learning, active synapses are strengthened on trials when both correct and incorrect responses are made, and as a result, cortical-cortical projections governed by Hebbian learning, by themselves, are incapable of discrimination learning.

When a stimulus is presented early in learning, activation from the visual association units that encode the perceptual representation of that stimulus will presumably be propagated to all relevant premotor units. If somehow the correct postsynaptic premotor target was activated more strongly than the incorrect targets, then Hebbian learning would strengthen the correct synapse more than the incorrect synapses (because the product of the pre- and postsynaptic activations would be greatest at the correct synapse). We propose that this is precisely the major role of the subcortical path through the striatum. Specifically, via three-factor (reward-mediated) learning, we propose that the slower subcortical path activates the correct postsynaptic target on the cortical path, which allows the appropriate cortical-cortical synapses in the premotor area to then be strengthened via classical (two-factor) Hebbian learning. In this way, control is gradually passed from the slower path through the basal ganglia to the faster cortical-cortical path. Thus, according to this model, the development of categorization automaticity is a gradual process via which control is passed from the subcortical category-learning systems to cortical-cortical connections from sensory association areas directly to premotor cortex. To our knowledge, this idea was first proposed by Ashby et al. (1998, p. 453). We call the resulting theory, which is developed here for the first time, SPEED.

The basic idea underlying SPEED—that pathways through the basal ganglia facilitate the development of more permanent cortical-cortical projections—is similar to proposals that the hippocampus facilitates the development of more permanent episodic and semantic memories in cortex. For example, McClelland, McNaughton, and O'Reilly (1995) argued that "the hippocampal memory system . . . can be viewed . . . as the teacher of the neocortical processing system" (p. 424). The major difference, of course, is that the hippocampal memory system is thought to mediate the consolidation of declarative memories (e.g., Squire, Stark, & Clark, 2004), whereas SPEED applies to tasks that are mediated by procedural memory. Declarative memory systems clearly have their role in category learning (Ashby & O'Brien, 2005), although a wide variety of evidence suggests that in information-integration tasks, the basal ganglia play a more important role than the hippocampus and other medial temporal lobe structures (for reviews of this evidence, see Ashby & Ennis, 2006; Ashby & Maddox, 2005).

Although there is good evidence that the striatum helps mediate procedural learning, it is important to note that many models of executive function also posit a key role for the striatum. For example, many models assume the striatum serves to gate inputs into working memory and/or to facilitate switching of executive attention (e.g., Ashby et al., 2005; Beiser & Houk, 1998; Braver & Cohen, 2000; Frank, Loughry, & O'Reilly, 2001; Humphries, Stewart, & Gurney, 2006; Monchi, Taylor, & Dagher, 2000; O'Reilly, Braver, & Cohen, 1999; O'Reilly & Frank, 2006). Critically, however, these models generally focus on anterior regions of the dorsal striatum, such as the head of the caudate nucleus, because of their reciprocal connections with PFC. In contrast, SPEED focuses on posterior regions of the dorsal striatum (e.g., body and tail of the caudate nucleus) because of their pronounced cortical input from sensory association areas. Thus, SPEED does not contradict any of these models of executive function.

Figure 3 shows a schematic of SPEED for the case in which there are two contrasting visual categories. Even when it is specified at this crude level, SPEED could explain some seemingly perplexing results. For example, it explains why diseases of the basal ganglia that impair information-integration category learning (e.g., PD, Huntington's disease) do not cause people to lose old familiar categories. In addition, it might provide a neurobiological rationale for the well-known speed-up in categorization response time (RT) that occurs as people gain experience with the categories (e.g., Nosofsky & Palmeri, 1997). A rigorous test of the theory, though, requires specifying the model in more detail. In the next section, we derive a computational version of SPEED that makes precise quantitative predictions in a wide variety of experimental paradigms.

Computational Modeling

As specified so far, the theory identifies relevant brain areas and their interconnections. To develop a computational version of SPEED requires writing a set of equations that describe the neural activation in each of the brain regions shown in Figure 3. This must be done even for the subcortical component of SPEED that was motivated by the procedural learning system of COVIS because the previous computational versions of COVIS have only been developed at more abstract levels (i.e., not at the level of neural networks). In addition, because this is a model of learning, we must also specify exactly how learning is implemented in the network. The main challenge is to select an appropriate level at which to direct the equations. If too global a level is chosen, then the resulting model will lack biological plausibility, and, as a result, it will be unable to account for neuroscience data. If the

[3] Although frontal-cortical dopamine appears to operate too slowly to serve as a reward-mediated training signal, this slow action makes possible a number of other hypothesized benefits. For example, a wide variety of evidence suggests that modest increases in frontal dopamine levels facilitate working memory (Ashby, Valentin, & Turken, 2002; Williams & Goldman-Rakic, 1995) and creative problem solving (Ashby, Isen, & Turken, 1999).

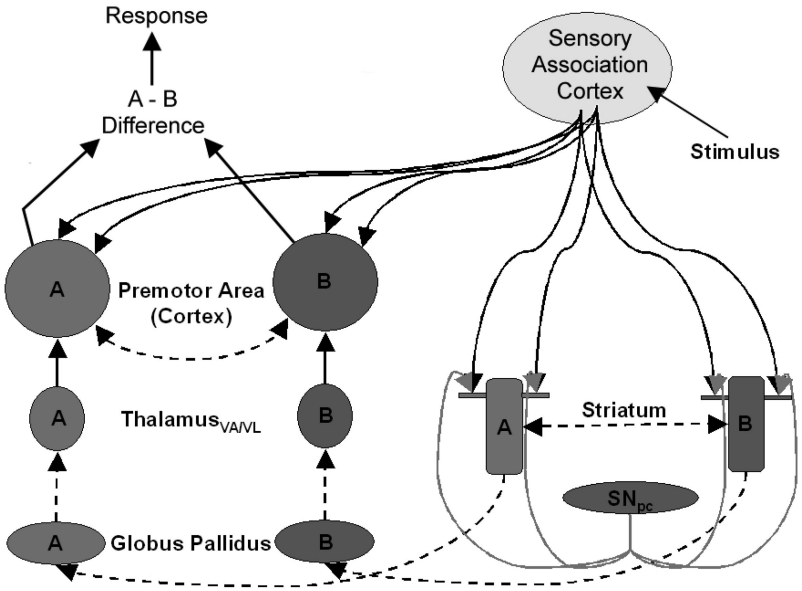

Figure 3. Schematic illustrating the subcortical pathways enable expertise development (SPEED) model of expertise in the case of two contrasting categories, A and B. Solid black lines denote excitatory (glutamate) projections, dashed lines denote inhibitory (GABAergic) projections, and solid gray lines denote dopamine projections.

model is too detailed, then it is likely to be too complex to account for human behavioral data.

Briefly, our approach is to model the functional interconnections between brain regions (as shown in Figure 3) and two key biophysical properties that are shared by all neurons: (a) saturation—every neuron has a maximum firing rate and (b) decay—if all inputs to a neuron cease, then the activation in that cell will decay to some baseline firing level. In our applications, these biophysical properties, together with the hypothesized interconnections, have been sufficient for modeling single-cell recording data, neuroimaging data, and human behavioral data (Ashby et al., 2005; Ashby & Valentin, 2006), but as the modeling efforts focus more heavily on neuroscience data, it may be necessary to increase the complexity of the models by trying to incorporate other, more subtle biophysical properties. For a thorough discussion of all the computational methods that were used in the present article, see Ashby and Valentin (2006). The next section describes the differential equations that result from following these methods. These equations describe how neural activation changes dynamically after stimulus presentation for each of the brain regions shown in Figure 3.

Activation Equations

To begin, we model sensory cortex as an ordered array of up to 10,000 units, each tuned to a different stimulus. We assume that each unit responds maximally when its preferred stimulus is presented and that its response decreases as a Gaussian function of the distance in stimulus space between the stimulus preferred by that unit and the presented stimulus.[4] For the present applications, it suffices to assume an exceedingly simple model of sensory cortex in which the activation of each unit is either off (with activation 0) or equal to some positive constant value during the duration of stimulus presentation. Specifically, we assume that when a stimulus is presented, the activation in sensory cortical unit K at time t is given by

$$I_K(t) = \frac{1}{\alpha} e^{-\frac{d(K,stimulus)^2}{2\alpha^2}}, \quad (1)$$

where α is a constant, and $d(K, stimulus)$ is the distance (in stimulus space) between the stimulus preferred by unit K and the presented stimulus. Equation 1, which is an example of a radial basis function (e.g., Buhmann, 2003), is a popular method for modeling the receptive fields of sensory units, in both models of categorization (e.g., Kruschke, 1992) and models of other tasks (e.g., Joo-Er, Wu, Lu, & Toh, 2002; Oglesby & Mason, 1991; Riesenhuber & Poggio, 1999; Rosenblum, Yacoob, & Davis, 1996).

We assume that the activation in striatal unit J at time t, denoted $S_J(t)$, is determined by the weighted sum of activations in all visual cortical cells that project to it and by lateral inhibition from other medium spiny cells:

$$\frac{dS_J(t)}{dt} = \left[\sum_K w_{K,J}(n) I_K(t) \right] [1 - S_J(t)] - \beta_S S_M(t) - \gamma_S [S_J(t) - S_{base}] + \sigma_S \varepsilon(t) S_J(t)[1 - S_J(t)], \quad (2)$$

where β_S, γ_S, S_{base}, and $\breve{\sigma}_S$ are constants, $M \neq J$, $I_K(t)$ is the input from visual cortical unit K at time t, $w_{K,J}(n)$ is the strength of the synapse between cortical unit K and striatal cell J on trial n, and $\varepsilon(t)$ is white noise.

[4] We make no other assumptions about the structure of sensory cortex.

An intuitive description of the derivative on the left is that it equals the change in activation in the striatal cell. Excitatory inputs to the cell increase this activation and hence will appear on the right as a positive term, and inhibitory inputs decrease the activation and so appear on the right as a negative term. The $[1 - S_J(t)]$ terms on the right model saturation. Because of these terms, neither the excitatory activation in the spines nor the noise can drive activation above 1 (which we have set as the arbitrary upper limit on activation). Similarly, the $S_J(t)$ in the last term on the right prevents noise from driving activation below 0. The second term on the right is a standard model of lateral inhibition (e.g., Usher & McClelland, 2001). The third term on the right models decay. If all inputs are zero, then this term guarantees that activation decays back to baseline (i.e., S_{base}). The last term in Equation 2 models noise.[5] Equation 2 assumes there are only two medium spiny cells (J and M).

Following this same approach (i.e., see Ashby & Valentin, 2006, for details on the approach), it is straightforward to write the other relevant differential equations. Activation in the globus pallidus at time t, denoted by $G_J(t)$, is described by

$$\frac{dG_J(t)}{dt} = -\alpha_G S_J(t) G_J(t) - \beta_G[G_J(t) - G_{base}], \quad (3)$$

where α_G, β_G, and G_{base} are constants. The first term models the inhibitory input from the striatum, and the second term ensures that in the absence of any other inputs, activation will decay to baseline G_{base}.

Similarly, activation in the thalamus at time t is given by

$$\frac{dT_J(t)}{dt} = -\alpha_T G_J(t) T_J(t) - \beta_T[T_J(t) - T_{base}], \quad (4)$$

where α_T, β_T, and T_{base} are constants. The first term models the inhibitory input from the globus pallidus, and the second term models decay to T_{base}. The ventral anterior and ventral lateral thalamic nuclei of the thalamus also receive a variety of excitatory inputs (e.g., cerebellum, PFC). We model these via T_{base}, which we set to a value that is higher than the true spontaneous firing rate. For our purposes, the most important of these excitatory inputs is from PFC (e.g., Anderson & DeVito, 1987). The PFC input is critical because striatal firing, by itself, can never trigger a motor response. When the striatum fires, it disinhibits the thalamus (i.e., by reducing pallidal inhibition), but it does not excite the thalamus. For this reason, random sensory stimuli that are encountered as one moves through the world could cause the striatum to fire, but this firing will typically not elicit an unintended motor response. In a category-learning task, instructions from an experimenter about how to respond could cause the PFC input to thalamus to increase, thereby priming the relevant response goals (e.g., "press the button on the left if the stimulus is a member of category A"). Because of the tonic inhibition from the globus pallidus, however, this PFC input is not enough to trigger a response. Instead, the striatum must first inhibit the globus pallidus, an event that would allow the thalamus to trigger one of the primed motor response goals.

Finally, activation in premotor cortex at time t is given by

$$\frac{dE_J(t)}{dt} = \left[\alpha_E T_J(t) + \sum_K v_{K,J}(n) I_K(t)\right][1 - E_J(t)] - \beta_E E_K(t) - \gamma_E[E_J(t) - E_{base}] + \sigma_E \varepsilon(t) E_J(t)[1 - E_J(t)], \quad (5)$$

where α_E, β_E, γ_E, σ_E, and E_{base} are constants; $I_K(t)$ again represents the input from visual cortical unit K at time t; $v_{K,J}(n)$ is the strength of the synapse between visual cortical unit K and premotor unit J on trial n; and $\varepsilon(t)$ is white noise. The second term on the right models lateral inhibition (mediated by GABAergic interneurons) in the same way as in Equation 2.

In tasks with only one response (e.g., Applications 2 and 3 below), we assume a response is initiated when the integrated activation in the premotor unit first exceeds a threshold τ. When there are two possible responses, A and B, (Applications 1 and 4), evidence suggests that cortical units in premotor areas are sensitive to the cumulated difference in evidence favoring the two alternatives (e.g., Shadlen & Newsome, 2001):

$$\Delta_{A,B}(t) = \int [E_A(t) - E_B(t)]dt. \quad (6)$$

In line with these results, we assume that in such tasks, response A is given when $\Delta_{A,B}(t)$ first exceeds τ, and response B is given when this integral is first less than $-\tau$. These response selection assumptions define a diffusion process (Ashby, 2000; Ratcliff, 1978) in which the drift rate is determined by the neural pathways postulated by SPEED.

Learning Equations

Although, of course, we expect some neural plasticity to occur at every synapse, the two critical synapses for automaticity are between the cells in sensory association cortex and the spines of the striatal medium spiny cells and between sensory association cortex and premotor cortex. A critical assumption of the theory proposed here is that these two synapses are modified according to qualitatively different rules. In particular, for the reasons described above, we assume that learning within the striatum is mediated by the three-factor learning rule, whereas in cortex, only two factors are required.

Two-Factor Learning. The two-factor cortical-cortical learning is also commonly referred to as *Hebbian learning*. In traditional Hebbian learning models, each synapse is strengthened by an amount proportional to the product of the pre- and postsynaptic activations. However, because postsynaptic NMDA receptor activation is required for LTP, we modify the traditional model to assume that all active synapses are strengthened in which the postsynaptic activation exceeds the NMDA receptor threshold, regardless of whether the preceding response was correct or incorrect. We also assume that active synapses in which the postsynaptic activation falls below the NMDA threshold are weakened.

Let $v_{K,J}(n)$ denote the strength of the synapse on trial n between visual cortical unit K and premotor unit J. Then, we model two-factor learning via the difference equation:

[5] The saturation and decay multipliers guarantee that the noise variance is greatest when activation levels are intermediate. This property of the model mimics the variance of a binomial distribution, which is greatest when the probability $p = .5$. In a binomial distribution (and in our noise model), when p (or activation) is near 0 or 1, there is very little variance because of the nearby hard limit.

$$v_{K,J}(n+1) = v_{K,J}(n) + \alpha_v \sum_t I_k(t) \left[\sum_t E_J(t) - \theta_{NMDA} \right]^+ \times [1 - v_{K,J}(n)] - \beta_v \sum_t I_k(t) \left[\theta_{NMDA} - \sum_t E_J(t) \right]^+ v_{K,J}(n), \quad (7)$$

where $\sum_t I_K(t)$ is the total activation over the course of the trial in sensory cortical unit K; $\sum_t E_J(t)$ is the total activation in premotor unit J; and α_v, β_v, and θ_{NMDA} are constants. θ_{NMDA} denotes the activation threshold of the NMDA receptor. The function $[f(t)]^+$ equals $f(t)$ when $f(t) > 0$, and 0 when $f(t) \leq 0$. The second (positive) term describes the conditions under which LTP occurs (premotor activation above the threshold for NMDA receptor activation) and the third (negative) term describes conditions that produce LTD (premotor activation below NMDA threshold). Note that this model assumes that the change in synaptic strength is proportional to the product of the pre- and postsynaptic activations (and the final rate limiting term that prevents the strength of the synapse from exceeding 1).

Three-Factor Learning. The three factors thought to be necessary to strengthen cortical-striatal synapses are (a) strong presynaptic activation, (b) strong postsynaptic activation, and (c) dopamine levels above baseline. According to this model, synapses between cells in sensory association cortex and medium spiny cells in the striatum are strengthened if the cortical cell responds strongly to the presented stimulus (Factors 1 and 2 are present) and the participant is rewarded for responding correctly (Factor 3). However, the strength of the synapse will weaken if the participant emits an incorrect response (Factor 3 is missing), or if the synapse is driven by a cell in sensory cortex that does not fire strongly to the stimulus (Factors 1 and 2 are missing). We also assume that cortical-striatal synaptic strength slowly decays if dopamine remains at baseline levels for long periods of time (which allows a kind of slow forgetting).

Let $w_{K,J}(n)$ denote the strength of the synapse on trial n between sensory cortical unit K and striatal unit J. We model three-factor learning as follows:

$$w_{K,J}(n+1) = w_{K,J}(n) + \alpha_w \sum_t I_K(t) \left[\sum_t S_J(t) - \theta_{NMDA} \right]^+ [D(n) - D_{base}]^+ [1 - w_{K,J}(n)] - \beta_w \sum_t I_K(t) \left[\sum_t S_J(t) - \theta_{NMDA} \right]^+ [D_{base} - D(n)]^+ w_{K,J}(n) - \gamma_w \sum_t I_K(t) \left[\theta_{NMDA} - \sum_t S_J(t) \right]^+ w_{K,J}(n) - \phi_w \left(1 - \frac{1 - [D(n) - D_{base}]^+}{1 - D_{base}} \right) w_{K,J}(n), \quad (8)$$

where $\sum_t S_J(t)$ is the total activation in striatal unit J; θ_{NMDA} again denotes the activation threshold of the NMDA receptor; D_{base} is the baseline firing rate of dopamine cells; $D(n)$ is the amount of dopamine released following feedback on trial n; and α_w, β_w, γ_w, and ϕ_w are constants. The second line describes the conditions under which LTP occurs (striatal activation above the threshold for NMDA receptor activation and dopamine above baseline), and lines three and four describe conditions that produce LTD (striatal activation above the NMDA threshold but dopamine below baseline, striatal activation below NMDA threshold). The last line models a slow decay in synaptic strength that occurs when dopamine stays at baseline levels for extended time periods.

In all simulations, we set the initial cortical-striatal weights $w_{i,J}(0)$ to be strong enough to cause one or more striatal units to fire when a novel stimulus is first presented. Without this property, postsynaptic activation would not exceed the NMDA receptor threshold, and an initial response would never occur (i.e., because LTP would be precluded). Fortunately, there is strong evidence supporting this assumption. Caan, Perrett, and Rolls (1984) recorded from single units in the tail of the caudate nucleus in monkeys as they passively viewed novel visual stimuli. The initial presentation of a stimulus (i.e., that was within the receptive field of the medium spiny cell) elicited a vigorous caudate response, which quickly habituated upon repeated presentations of the same stimulus. Because no response was required of the animals in this experiment, and no rewards were given, the learning component of SPEED predicts this habituation.

This model of three-factor learning requires that we specify the amount of dopamine released on every trial in response to the feedback signal. Although there are a number of powerful models of dopamine release (e.g., Brown, Bullock, & Grossberg, 1999), these models are quite complicated, in part because the midbrain dopamine areas receive a diverse array of excitatory and inhibitory inputs, and different classes of dopamine receptors have different effects. For our purposes, however, most of this complexity can be ignored. For example, we have no need to predict dynamic changes in dopamine release. In fact, the only critical variable is the amount of dopamine released to the feedback signal on each trial. Toward this end, the key empirical results are (e.g., Schultz et al., 1997; Tobler, Dickinson, & Schultz, 2003) (a) midbrain dopamine cells have a high spontaneous firing rate; (b) dopamine release increases above baseline following unexpected reward, and the more unexpected the reward, the greater the release; and (c) dopamine release decreases below baseline following unexpected absence of reward, and the more unexpected the absence, the greater the decrease.

A simple model that incorporates these properties is as follows: If the response on trial n was correct, then the amount of dopamine released equals

$$D(n) = D_{base} + [1 - P(C)](1 - D_{base}); \quad (9)$$

whereas if the response on trial n was incorrect, then,

$$D(n) = D_{base} - P(C)D_{base}, \quad (10)$$

where $P(C)$ is the current probability of a correct response. For practical implementation, we estimated $P(C)$ as the proportion correct over the previous 50 trials. Note that according to this simple model, dopamine release is greatest following a correct response when accuracy is low—that is, when the reward is unexpected. As accuracy increases and the participant is better able to predict when a reward will occur, dopamine release decreases to the point that after 50 correct responses in a row, another correct response leaves dopamine at baseline. The error model is similar. When errors are unexpected (i.e., accuracy is high), an error depresses dopamine release well below baseline, whereas expected errors have little (or no) effect on dopamine release.

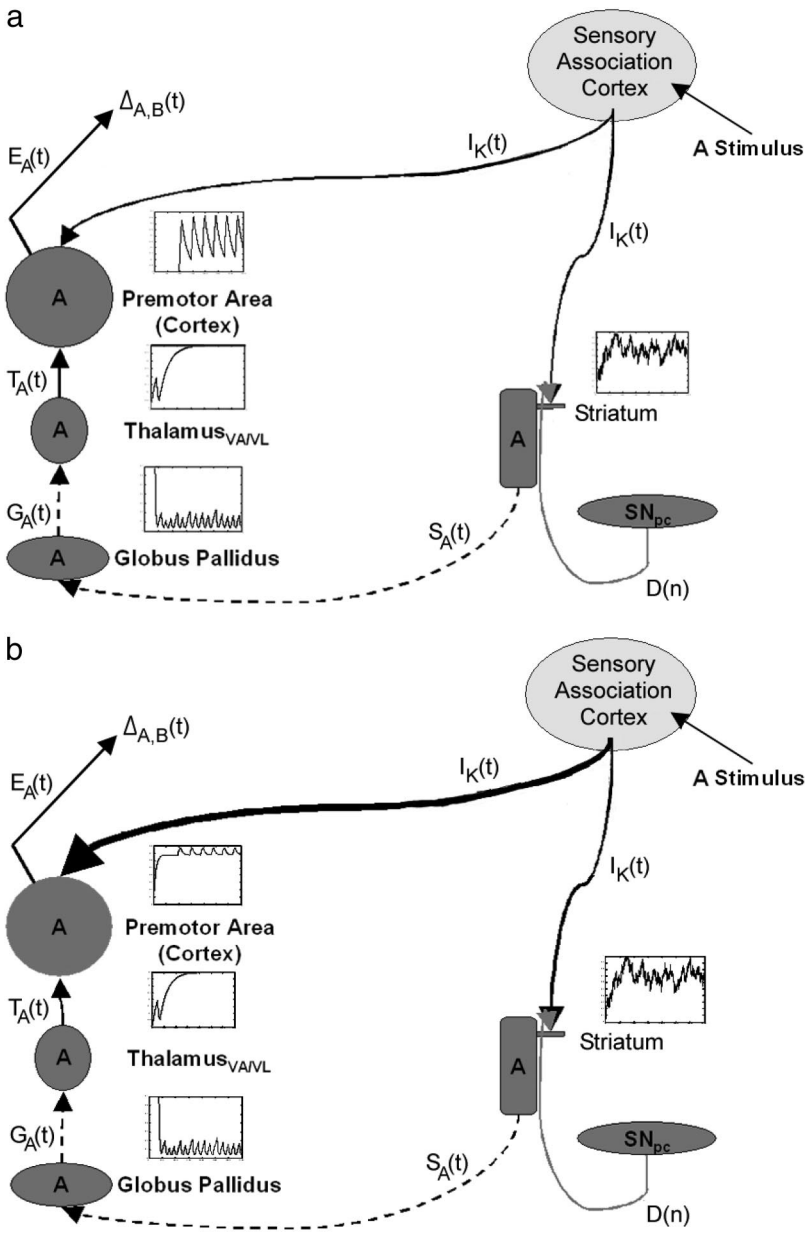

Figure 4. Schematics illustrating the subcortical pathways enable expertise development (SPEED) model of expertise on a trial when a stimulus from Category A is presented. Figure 4a depicts a trial early in learning. The graphs beside each brain area show the simulated solutions of the differential equations that define the model (see the Appendix). Figure 4b is the same as 4a, except later in training, after expertise has developed. Solid black lines denote excitatory (glutamate) projections, dashed lines denote inhibitory (GABAergic) projections, and solid gray lines denote dopamine projections.

Applying the Model

Figure 3 illustrates how SPEED might be applied to an information-integration task with two categories, A and B. Because there are two categories, the model includes two units in all brain regions except sensory cortex, which includes 10,000 units. Each medium spiny cell includes 10,000 spines—one for each unit in sensory cortex. Each unit in sensory cortex projects to a spine on each medium spiny cell with an initial synaptic strength that is randomly determined. Similarly, every sensory cortical unit projects to its own synapse on each premotor unit, and the initial strength of these 20,000 synapses is set to zero.[6]

Figure 4a shows the pathways in SPEED leading to an A response on a trial early in training when the stimulus is an

[6] There is no need to make the initial cortical-cortical synaptic strengths random because these synapses do not control early performance.

exemplar from Category A. The plots beside the striatum, globus pallidus, thalamus, and premotor units are simulated single trials derived from the differential equations that describe activation in each brain region. Note that presentation of the stimulus causes activation in striatal unit A to increase. The striatum sends an inhibitory projection to globus pallidus, so striatal activation inhibits the high spontaneous firing rate of pallidal units. Globus pallidus tonically inhibits the thalamus, so pallidal deactivation releases the thalamus from tonic inhibition, and, as a result, thalamic activation increases (presumably because of tonic excitatory input from PFC). Finally, thalamus excites cortex, so thalamic activation excites the premotor unit enough to trigger a response. Note that the strength of the cortical-cortical synapses is not great enough to allow sensory cortex to drive activation in the premotor unit high enough to trigger a response. As a result, SPEED responds correctly, but with a long RT.

Figure 4b is the same as Figure 4a, except for a trial much later in training, after expertise has been established. Note that the activation in each subcortical unit is essentially unchanged from earlier in training. However, now activation in the premotor unit is great enough to trigger a response before the striatal activation has even finished increasing. This is because Hebbian learning has progressed enough for visual cortical activation to be sufficient to drive the correct premotor unit high enough above the incorrect unit to initiate a response. Now the model is both accurate and fast.

Empirical Tests of SPEED

SPEED predicts that during early information-integration learning, the COVIS procedural and explicit systems will both be active. Thus, during early learning, SPEED predicts that neuroimaging studies should report activation in a great many brain areas, including visual cortex, much of the striatum (e.g., head, body, and tail of the caudate), globus pallidus, several thalamic nuclei (e.g., medial dorsal, ventral anterior, ventral lateral), premotor and motor cortices, PFC, anterior cingulate, and perhaps the hippocampus and other medial temporal lobe structures (e.g., Ashby & Valentin, 2005). As automaticity develops, however, activation in many of these regions is predicted to decrease substantially, with the exception of visual, premotor, and motor cortices (although activation in these regions could narrow).

We know of no studies that tested these predictions in an information-integration category-learning task (although the early learning predictions were generally supported by Seger & Cincotta, 2002, and by Nomura et al., 2007). However, several neuroimaging studies have examined changes in neural activations during the course of more traditional procedural learning tasks, in which participants were required to execute a fixed sequence of finger movements. To guarantee automaticity, participants in these studies either completed weeks of daily practice (Karni et al., 1995; Lehéricy et al., 2005) or practiced until they showed no interference from a simultaneous dual task (Poldrack et al., 2005; Wu, Kansaku, & Hallett, 2004). In general, the predictions of SPEED were supported in these studies. For example, after automaticity was achieved, decreased activation (relative to initial learning) was reported in cingulate, premotor, parietal, and prefrontal cortices, as well as in the caudate nucleus (Lehéricy et al., 2005; Poldrack et al., 2005; Wu et al., 2004). However, the reduced caudate activation reported by Lehéricy et al. (2005) was only in the associative/premotor regions of the caudate. These researchers reported increased activations in the sensorimotor regions of the putamen and globus pallidus. Pyramidal cells leaving motor cortex send collaterals to striatal neurons on the indirect pathway (Lei, Jiao, Del Mar, & Reiner, 2004), so SPEED predicts that these latter increases may reflect activity that occurred after response selection was complete. Also, in contrast to each of these studies, which reported no cortical regions that increased activity with practice, Karni et al. (1995) reported an increase in motor cortex activation in a study in which a high-field strength magnet was used and focused only on motor cortex.

Thus, the few relevant neuroimaging studies are generally consistent with the predictions of SPEED, but they are hardly conclusive. More rigorous tests can be achieved by examining the quantitative predictions of the model. This is the goal of the next four sections.

Application 1: Tactile Categorization

Romo, Merchant, and their colleagues reported a series of single-unit recording studies in which monkeys learned vibrotactile categories (Merchant, Zainos, Hernadez, Salinas, & Romo, 1997; Romo, Merchant, Ruiz, Crespo, & Zainos, 1995; Romo, Merchant, Zainos, & Hernadez, 1997). In these experiments, a rod vibrated against the monkey's finger at one of 10 different speeds. The monkeys were trained to push one button if one of the five low-speed vibrations occurred and to press a different button if they received one of the five high-speed vibrations. After extended feedback training, the monkeys reliably learned these categories.

Following training, the monkeys completed an additional session during which single-unit recordings were collected from the putamen and the premotor cortex. The putamen is a major structure within the striatum that receives direct projections from somatosensory cortex (e.g., Heimer, 1995). Thus, according to COVIS, the putamen would play the same role for tactile category learning that the caudate plays for visual category learning.

Within the putamen, the responses of 695 cells were characterized in detail. Of these, 196 responded to the onset of the vibrating rod, regardless of category membership, 258 responded to the monkey's arm movement, regardless of response, and 165 responded to the category membership of the stimulus. In the premotor cortex, 191 neurons were studied. Of these, 104 responded to category membership. Figure 5 shows recordings from two category membership neurons from the putamen during categorization of each of the 10 stimuli and response histograms that summarize the responses of the 104 categorization cells that were studied in the premotor cortex. Note that the putamen cell corresponding to the upper left panel of Figure 5 fires to every exemplar in the low-speed category but not to any exemplars in the high-speed category. The cell corresponding to the upper right panel shows the opposite firing pattern. Furthermore, these cells did not respond when the same stimuli were passively presented to the monkeys (i.e., when no categorization response was required), or when the monkeys made the same response to visual stimuli (Merchant et al., 1997). This is strong evidence that the categorical responses shown in Figure 5 developed as a result of category learning. The bottom half of Figure 5 shows that similar category-specific cells appear in premotor cortex. This is perhaps less surprising because, presumably, such cells must exist in some

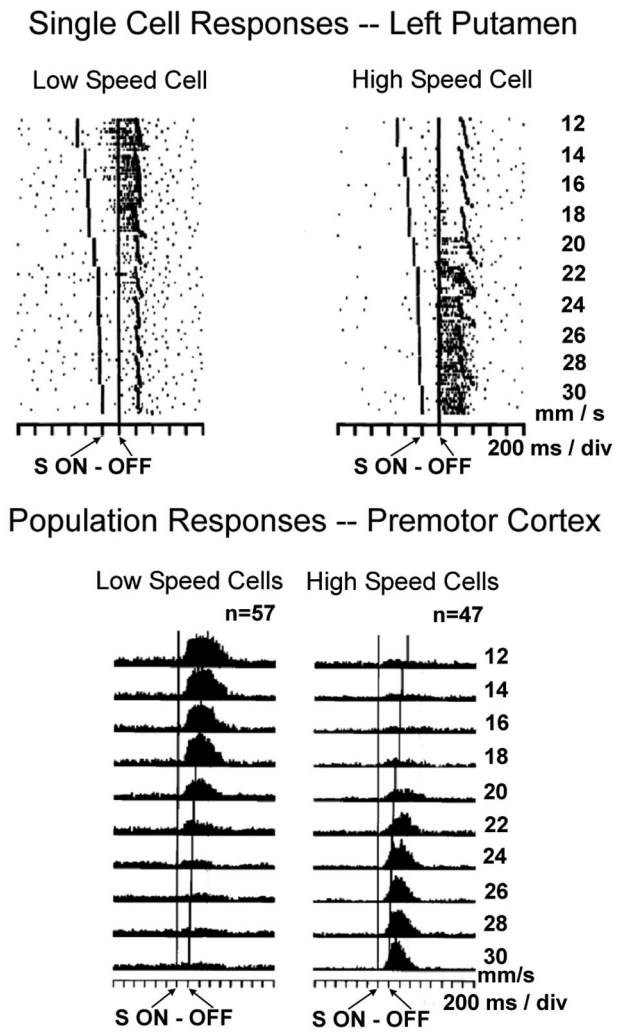

Figure 5. Single-cell responses from left putamen (Merchant et al., 1997; From "Functional Properties of Primate Putamen Neurons During the Categorization of Tactile Stimuli," by H. Merchant, A. Zainos, A. Hernandez, E. Salinas, and R. Romo, 1997, *Journal of Neurophysiology, 77,* p. 1143. Copyright 1997 by the American Physiological Society. Reprinted with permission.) and population responses from premotor cortex (Romo et al., 1997; From "Categorical Perception of Somesthetic Stimuli: Psychophysical Measurements Correlated With Neuronal Events in Primate Medial Premotor Cortex," by R. Romo, H. Merchant, A. Zainos, and A. Hernandez, 1997, *Cerebral Cortex, 7,* p. 320. Copyright 1997, by the Oxford University Press. Reprinted with permission.) to each of 10 stimuli in a tactile category-learning experiment. Note that cells in both regions develop category-specific responses.

motor area, given that the monkeys made reliable motor responses to the two categories. However, other results suggest that the similarity between the striatal and premotor responses that are shown in Figure 5 are causal rather than correlational (Brasted & Wise, 2004).

These data, together with the animal lesion results described earlier, support the fundamental assumption of COVIS that the striatum is a key site of learning in perceptual categorization. The lesion results support the hypothesis that the striatum is both necessary and sufficient for discrimination learning of the type required during information-integration categorization (e.g., Eacott & Gaffan, 1991; Gaffan & Eacott, 1995; Gaffan & Harrison, 1987; McDonald & White, 1993, 1994; Packard et al., 1989; Packard & McGaugh, 1992). The data from Romo, Merchant and their colleagues (Merchant et al., 1997; Roma et al., 1995, 1997) suggest that medium spiny cells in the striatum can become associated with a categorization response, presumably via a process of reward-mediated learning.

To simulate the performance of SPEED in this experiment, we constructed a version of the model like the one illustrated in Figure 3, except with a one-dimensional array of sensory cortical cells (rather than two-dimensional). Specifically, each stimulus triggered maximum activation in 1 of 100 somatosensory cells and triggered less activation in nearby cells. Because there were two contrasting categories, there were two units in every other brain region. Both striatal (i.e., putamen) units had 100 spines—one for each cell in sensory cortex. Thus, each sensory cortical cell projected to its own spine on both striatal units. The activations of all units were determined by the equations given above. A response was initiated when the integrated difference between the two premotor units first exceeded threshold. The cortical-cortical and cortical-striatal synaptic strengths were adjusted between trials according to the two- and three-factor learning equations given above. Details of this and all other simulations described in this article are given in the Appendix (including how parameters were estimated and spikes were generated).

Figure 6 shows the proportion of "low" responses (after training) given by the monkeys and by SPEED for each of the 10 vibration speeds. The solid vertical line shows the category bound. Note that the monkeys and the model both reliably learned the categories. Figure 7 shows simulated single-unit recordings from the striatal and premotor units of SPEED on trials when each of the

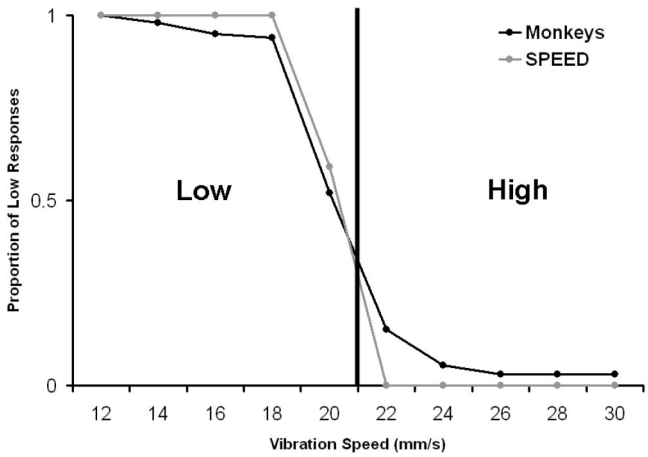

Figure 6. Proportion of "low" responses given to each of 10 stimuli in a tactile category-learning experiment (Merchant et al., 1997; From "Functional Properties of Primate Putamen Neurons During the Categorization of Tactile Stimuli," by H. Merchant, A. Zainos, A. Hernandez, E. Salinas, and R. Romo, 1997, *Journal of Neurophysiology, 77,* p. 1151. Copyright 1997 by the American Physiological Society. Reprinted with permission.) by monkeys (black line) and by SPEED (Subcortical Pathways Enable Expertise Development; gray line).

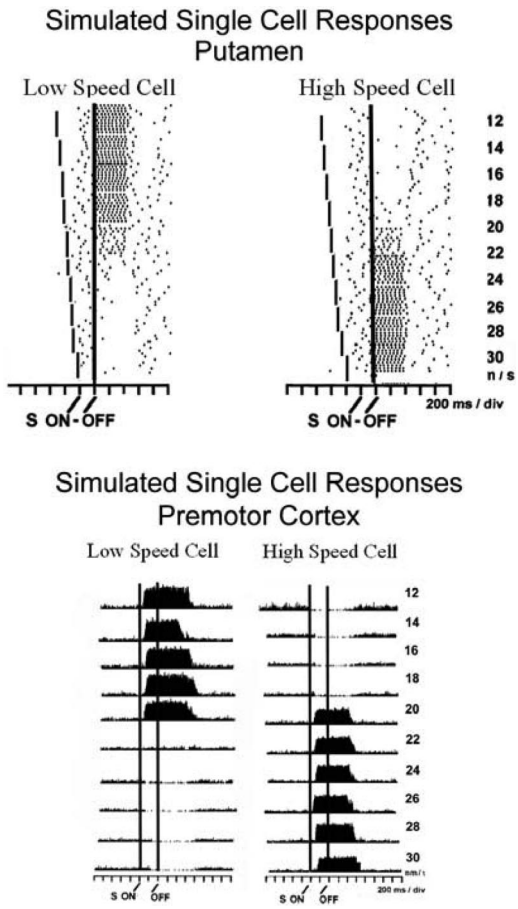

Figure 7. Single-cell responses from the striatum and premotor cortex of SPEED (Subcortical Pathways Enable Expertise Development) to each of 10 stimuli in the tactile category-learning experiment of Romo, Merchant, and their colleagues (Merchant et al., 1997; Romo et al., 1995; Romo, Merchant, Zainos, & Hernadez, 1997). Compare with the responses from monkeys shown in Figure 5.

10 vibrational speeds was presented. Not surprisingly, the SPEED units mimic the firing patterns of the monkeys.

Note that the SPEED units in both the putamen and premotor cortex exhibit a push-pull pattern of responding, in which high activity in one unit is associated with a suppression below baseline in the opposite unit. This phenomenon, which is due to lateral inhibition, can be seen in the Figure 5 putamen recordings of Merchant et al. (1997), but not in the Figure 5 premotor recordings. However, it is important to note that whereas the SPEED premotor simulations in Figure 7 are from a single unit, the Romo et al. (1997) premotor data shown in Figure 5 are collapsed across 104 neurons, which could obscure push-pull behavior in individual cells. This seems plausible because push-pull behavior has been reported in a cortical area closely associated with premotor cortex (i.e., lateral intraparietal cortex; Shadlen & Newsome, 2001).

Recall that the category-specific neurons in the putamen that are illustrated in Figure 5 did not respond when the same stimuli were passively presented to the monkeys (Merchant et al., 1997). The present version of SPEED, however, would not discriminate between passive and active experimental conditions, so its striatal units would respond the same in either condition. One possibility for correcting this shortcoming might be to model the cholinergic interneurons in the striatum known as the tonically active neurons (TANs). The TANs, which have a high spontaneous firing rate (hence their name), tonically inhibit the striatal medium spiny cells (Akins, Surmeier, & Kitai, 1990). Motivationally salient cues in the environment induce a pause in the firing of the TANs, thereby disinhibiting the medium spiny cells (Kimura, Rajkowski, & Evarts, 1984; Ravel, Legallet, & Apicella, 2003). So one possibility is that TANs maintain their high firing rate during passive viewing conditions, thereby preventing the medium spiny cells from firing in response to well-learned stimuli.

The TANs could also play an important role in recovery from extinction. In the present version of SPEED, removing the rewards shortly after a behavior has been learned would cause a reduction in the strength of the relevant cortical-striatal synapses (and therefore extinction). However, if the synaptic strengths became too small, then reacquisition of the behavior could be difficult once the rewards were reinstated. Striatal TANs receive a prominent (glutamatergic) input from cortex and thalamus, and these synapses display considerable synaptic plasticity (e.g., Suzuki, Miura, Nishimura, & Aosaki, 2001), which modulates the firing rate of TANs in response to changing reward conditions in the environment (e.g., Bar-Gad, Morris, & Bergman (2003). One intriguing possibility, therefore, is that during extinction, the TANs learn that the stimulus is no longer associated with reward before the synapses between cortex and the medium spiny cells have completely decayed. Once the TANs no longer pause during stimulus presentation, their high-spontaneous firing rate would prevent the medium spiny cells from responding, thereby hastening extinction. During reacquisition, the TANs might quickly learn that the stimulus was salient once again, which would cause them to pause during stimulus presentation, and allow the medium spiny cells to select a response. One particularly attractive feature of this model is that it could account for the well-documented phenomenon that reacquisition of a behavior following extinction is often faster than the initial acquisition.

The simulations described in this section provide support for the early learning component of SPEED (i.e., the procedural learning system of COVIS) as a model of category learning in monkeys—at both the single-unit and behavioral levels. The data from Romo, Merchant and their colleagues (Merchant et al., 1997; Romo et al., 1995, 1997), however, do not provide a test of the critical late-stage learning component of SPEED. This is because recordings were made at only a single time point in the learning sequence—shortly after the monkeys' categorization accuracy first reached asymptote. Testing the expertise component of SPEED requires data from tasks in which the participants were significantly overtrained. The next three empirical applications include this critical overtraining feature.

Application 2: Striatal Dropout With Extended Practice

Carelli, Wolske, and West (1997) trained rats to lever press to a tone. The rats completed 70 trials per day on each of 18 days. They learned to lever press reliably to the tone within just a few sessions, and their mean reaction time (RT) gradually improved over the course of training. Throughout the extended training period, Carelli et al. (1997) recorded from single units in the striatum.

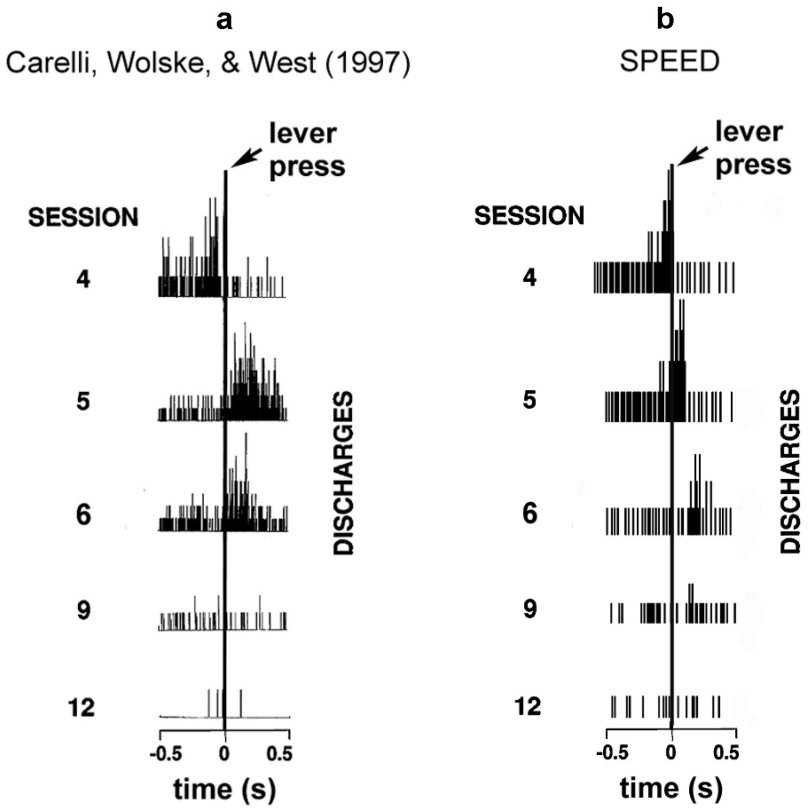

Figure 8. Figure 8a shows single-cell responses from the striatum of a rat in the instrumental learning task of Carelli, Wolske, and West (1997; From "Loss of Lever Press-Related Firing of Rat Striatal Forelimb Neurons After Repeated Sessions in a Lever Pressing Task," by R. M. Carelli, M. Wolske, and M. O. West, 1997, *Journal of Neuroscience, 17,* p. 1808. Copyright 1997 by the Society for Neuroscience. Reprinted with permission.). Figure 8b shows striatal responses from SPEED (Subcortical Pathways Enable Expertise Development) during Sessions 4, 5, 6, 9, and 12.

Spike histograms from one rat are shown in Figure 8a (ignore Figure 8b for now). Consistent with the data of Merchant et al. (1997), during Session 4, these striatal units fire a burst just before the rat lever presses. However, unlike Merchant et al. (1997), Carelli et al. (1997) continued recording for many days after the behavior was learned. Note that during Sessions 5 and 6, the same striatal units still fire bursts, but now these bursts come after the response has been made, and therefore they can play no role in response selection. In later sessions, presumably after automaticity is well established, the striatum ceases responding altogether; that is, neither the tone nor the response elicit any activity from the same striatal units that apparently controlled the response earlier in training.

The Figure 8 data, together with the other behavioral neuroscience data reviewed earlier, strongly suggest that novice and expert instrumental behaviors are controlled by different brain systems. In particular, the Carelli et al. (1997) results conclusively falsify any model that assigns a permanent role to the striatum in instrumental responding. For example, COVIS, by itself, could not account for these data.

The Carelli et al. (1997) experiment used only a single auditory cue, so in our simulations, SPEED was given only one sensory cortical unit, which was either active or not, depending on whether the stimulus was present. Similarly, because only one response was possible, SPEED was given only one unit in all other brain regions (striatum, internal segment of the globus pallidus, thalamus, premotor cortex). Initial responses were triggered by noise,[7] which sometimes happened in the presence of the stimulus. In this case, the dopamine levels were adjusted on the basis of the expectation of reward (see Equations 9 and 10), and the strength of the cortical-cortical and cortical-striatal synapses were adjusted according to the two- and three-factor learning equations (see Equations 7 and 8).

Like the rats, SPEED improved its accuracy and RT during each of the first four experimental sessions. Figure 8b shows the striatal responses of SPEED during Sessions 4, 5, 6, 9, and 12. As with the monkeys, note that the SPEED striatal unit fires bursts just before the lever press during Session 4. During Sessions 5 and 6, however, the striatal response begins occurring after the lever press. This is the time when the behavior is first fully controlled by the

[7] This is obviously a poor model of the behavior of the rat in the absence of a stimulus. A more realistic model, which is beyond the scope of this article, would have to account for a variety of complex states, such as motivation, arousal, and the drive to explore.

cortical-cortical projections. Even later in training, the SPEED striatal response disappears. After so much training, the reward no longer elicits dopamine release (because it is now expected), and without excess dopamine, the strength of the cortical-striatal synapses gradually decays back to baseline.[8]

The failure of an expected reward to elicit dopamine cell firing is less likely to reduce the strength of the cortical-cortical synapses in premotor cortex. This is because any event that causes VTA dopamine cells to fire will likely lead to increased dopamine levels in PFC for many minutes, and, therefore, during that time promote LTP at cortical-cortical synapses, even when the reward is expected. For example, cues that predict a rewarded session is about to begin might serve this purpose. In contrast, although such cues would also likely increase striatal dopamine levels, reuptake mechanisms would clear this dopamine from the cortical-striatal synapses before stimulus presentation.

Application 3: The Role of Dopamine in Early Versus Late Training

Choi, Balsam, and Horvitz (2005) reported the results of an experiment in which a task similar to the lever-pressing task of Carelli et al. (1997) was used. Hungry rats freely explored a chamber. At random times (with at least 30 s between trials), a tone sounded when a food pellet was dropped into a compartment attached to the chamber. Upon hearing this tone, every rat was trained to insert its head into the compartment to retrieve the food. Each rat completed 28 trials per day for a total of 17 days. Every rat was given one of four different drug injections on one of three different training days (so the design was 4 × 3). The drug was either a saline control or one of three different doses of a selective dopamine D1 antagonist (SCH 23390). Injections were administered (intraperitoneally) on Training Day 3, 7, or 17.

The relevant results are shown in Figure 9. The top panel shows the effects of the drug on total locomotor activity. Note that, as expected, the D1 antagonist reduced motor activity in a dose-dependent fashion and that the amount of this reduction did not depend on the injection day. The middle panel shows the effect of the drug on the instrumental response. On Day 3, the D1 antagonist significantly interfered with the expression of the learned behavior, and in a dose-dependent manner. For example, the rats given the largest dose failed to retrieve (i.e., in 10 s) approximately 40% of the rewards even though the control rats with the same amount of training retrieved virtually all of the available rewards. In contrast, on Days 7 and 17, the D1 antagonist had no effect on the performance of the learned behavior. Even the rats subjected to the highest dose managed to retrieve almost all rewards, despite their impaired motor abilities. Thus, dopamine D1 antagonists disrupt the expression of learned instrumental behaviors early in training, but not after the rats have had extensive practice at the task. As with the Carelli et al. (1997) data, these results are extremely problematic for any model that assumes instrumental learning is controlled by a single basal ganglia-dependent system.

When simulating the performance of SPEED in this task, we made the same architectural assumptions as in the Carelli et al. (1997) simulations (and used all the same parameter estimates). We modeled the dopamine antagonist by (linearly) reducing the dynamic range of the dopamine levels about baseline (see the Appendix for details). Thus, the 12 data points of Choi et al. (2005)

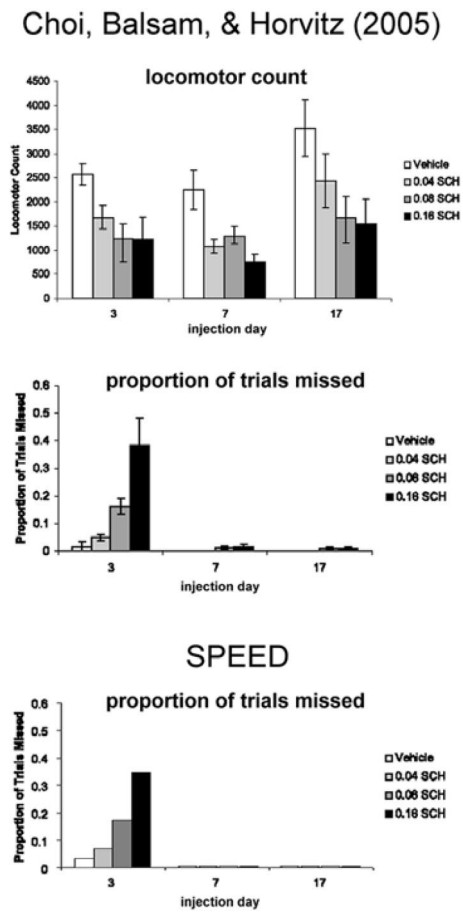

Figure 9. Locomotor count on Days 3, 7, and 17 for the rats in each of the four drug groups of the instrumental learning task of Choi, Balsam, and Horvitz (2005; top and middle panels, From "Extended Habit Training Reduces Dopamine Mediation of Appetitive Response Expression," by W. Y. Choi, P. D. Balsam, and J. C. Horvitz, 2005, *Journal of Neuroscience, 25,* p. 6730. Copyright 2005 by the Society for Neuroscience. Reprinted with permission.). The bottom two panels show the proportion of rewards missed by the animals (middle panel) and by SPEED (Subcortical Pathways Enable Expertise Development; bottom panel) on these same days and under these same conditions.

[8] It is well-known that when a reward becomes predictable, its presentation no longer elicits dopamine release. Instead, the dopamine release moves back to the cue that predicts the reward (e.g., Schultz, 2002). The dopamine release to the cue typically does not diminish, perhaps because the appearance of the cue is generally unpredictable. In the Carelli et al. (1997) experiment, however, the auditory tone was presented to the rat only when it stood on a piece of blue tape affixed to the floor of the cage. Because of this unusual design feature, the tone itself became predictable. So one possibility (suggested by J. Horvitz, personal communication, November 14, 2005) is that the tone eventually failed to elicit dopamine release, and it is only for this reason that the striatal response disappeared. SPEED predicts that in the absence of dopamine release, the striatal response will eventually disappear. Our simple model of dopamine release (see Equations 9 and 10) only models the amount of dopamine released to the feedback signal (i.e., the reward). For the present applications, this model appears sufficient. In tasks in which dopamine release to the cue predicting reward plays a prominent role, however, we would need to substitute a more sophisticated model of dopamine release for Equations 9 and 10.

were fit with three free parameters—the effective reduction in free dopamine that results from injection of the three doses of D1 antagonist.

Fits of SPEED to these data are shown in the bottom panel of Figure 9. Note that the fit is virtually perfect. In fact, after expertise has been established (i.e., after the behavior is controlled by the cortical-cortical projections), the model predicts no effects of dopamine antagonists regardless of the parameter values.

Application 4: Human Category Learning

As humans gain practice in virtually any skill, they naturally become faster and more accurate. In many laboratory studies of expertise or automaticity, asymptotic accuracy is perfect, and so the primary focus is on RT. The most widely replicated and best known empirical result in this area is that mean RT decreases as a power function of the amount of practice. Among many other examples, a power-function speed-up has been reported for cigar rolling (Crossman, 1959), proving geometry theorems (Neves & Anderson, 1981), and categorization of color patches (Nosofsky & Palmeri, 1997). In fact, by the early 1980s, the power-function speed-up was so well accepted that Newell and Rosenbloom (1981) proposed it as a scientific law.

According to the original formulation of the power law of speed-up, mean RT after N repetitions of a task should equal

$$\overline{RT_N} = T_{\min} + \alpha(N + n_0)^{-\beta}, \tag{11}$$

where T_{min} is the irreducible minimum RT, α is a constant, β is interpreted as the learning rate (usually between 0 and 1), and n_0 is interpreted as the number of trials of learning that occurred before the task began. Because SPEED has no prior learning, in our applications we consider a reduced version of the power law in which $n_0 = 0$. Not surprisingly, the power law of speed-up played a pivotal role in the development of a number of cognitive models of expertise (e.g., Anderson, 1983, 1993; Cohen, Dunbar, & McClelland, 1990; Logan, 1988; MacKay, 1982; Nosofsky & Palmeri, 1997; Rickard, 1997). Despite its ubiquity, however, it should be noted that there have been recent challenges to the universality of the law (Heathcote, Brown, & Mewhort, 2000; Rickard, 1997).

Nosofsky and Palmeri (1997) reported the results of a category-learning experiment that used information-integration categories in which the mean RTs for each individual participant were well described by the power law. In this experiment, 12 Munsell color patches that varied in saturation and brightness (but not hue) were equally divided into two information-integration categories (e.g., the categories were nonlinearly separable). Each participant completed 1,800 categorization trials (150 repetitions of each stimulus). By the end of training, all participants were responding with near perfect accuracy, so the dependent variable of primary interest was mean RT, which was computed for each consecutive 60-trial block. These mean RT profiles were well fit by the power law (i.e., the power law accounted for 88.3%, 74%, and 94% of the variance in the mean RTs of each participant, respectively).

We modeled these data with the same version of SPEED that was used to fit the data of Romo, Merchant, and their colleagues (Merchant et al., 1997; Roma et al., 1995, 1997) except in this application we included 10,000 sensory cortical units (see the Appendix for details). We let the model learn the categories 3,000 different times, using different initial random cortical-striatal synaptic strengths and different initial seeds for the noise generators. Figure 10 summarizes the performance of the model across these 3,000 simulations. Figure 10a shows the proportion of correct responses on each trial. Note that SPEED successfully learned the two categories, with accuracy becoming perfect after about 200 trials. Because the categories used by Nosofsky and Palmeri (1997) were nonlinearly separable, this application demonstrates that SPEED can learn nonlinearly separable categories.

Figure 10b shows the mean RTs (averaged across stimuli) predicted by SPEED on each trial of learning (ignore the bottom panel for now). We fit power and exponential functions to these mean RTs using the same procedures as in Nosofsky and Palmeri (1997; i.e., by grouping the mean RT data into 30 blocks). The power law gave an excellent account of the RT speed-up exhibited by the model (i.e., accounting for 99.51% of the variance). In contrast, an exponential function[9] gave a poorer account (i.e., accounting for 94.7% of the variance).

Figure 11 shows the RT density functions predicted by the model on a variety of different trials during the learning process. For example, the density function on Trial 77 was estimated by taking the RT on Trial 77 from each of the 3,000 replications and applying a Parzen (1962) kernel density estimator with a Gaussian kernel of width 10 ms. The RT density functions shown in Figure 11 display a number of properties that are commonly seen in empirical RT distributions. First, each of the distributions in Figure 11 is unimodal and skewed right, as are virtually all empirical RT distributions (e.g., Luce, 1986). Second, SPEED predicts a substantial decrease in RT variance with practice. Although the relation between RT variance and practice has not been as extensively studied as the relation between mean RT and practice, there is nevertheless much data supporting this prediction of SPEED (e.g., Logan, 1988; Rickard, 1997).

The unimodal nature of these distributions may, at first glance, seem counterintuitive. If the observable RT is sometimes determined by a slow subcortical path and sometimes by a fast cortical path, then one might expect the model to predict a bimodal RT distribution. In fact, SPEED predicts that, before automaticity has developed, both paths contribute to each response. The premotor unit that controls the response receives excitatory inputs from both the cortical and subcortical pathways. Initially almost all of this excitation comes from the subcortical path, and after automaticity has developed, it almost all comes from the cortical path. In between, however, both paths contribute. To show this, on each trial we can measure the amount of excitatory input to the critical premotor unit coming from the cortical and subcortical paths. The proportion of this excitation contributed by the subcortical path is illustrated in Figure 10c. The trials for which the RT densities in Figure 11 were estimated correspond to points on this curve that are equally spaced on the ordinate. Note that the Figure 10c proportion starts at 1 and gradually decreases to 0 after many trials. Thus, as expected, at the beginning of training, the subcortical path completely controls responding. Even so, the cortical path quickly begins to contribute, even before it can generate enough activation to cause a response on its own. The more the cortical path con-

[9] Both RT models had three free parameters. The best fitting version of the power law was $RT = 37 + 1339N^{-.20}$.

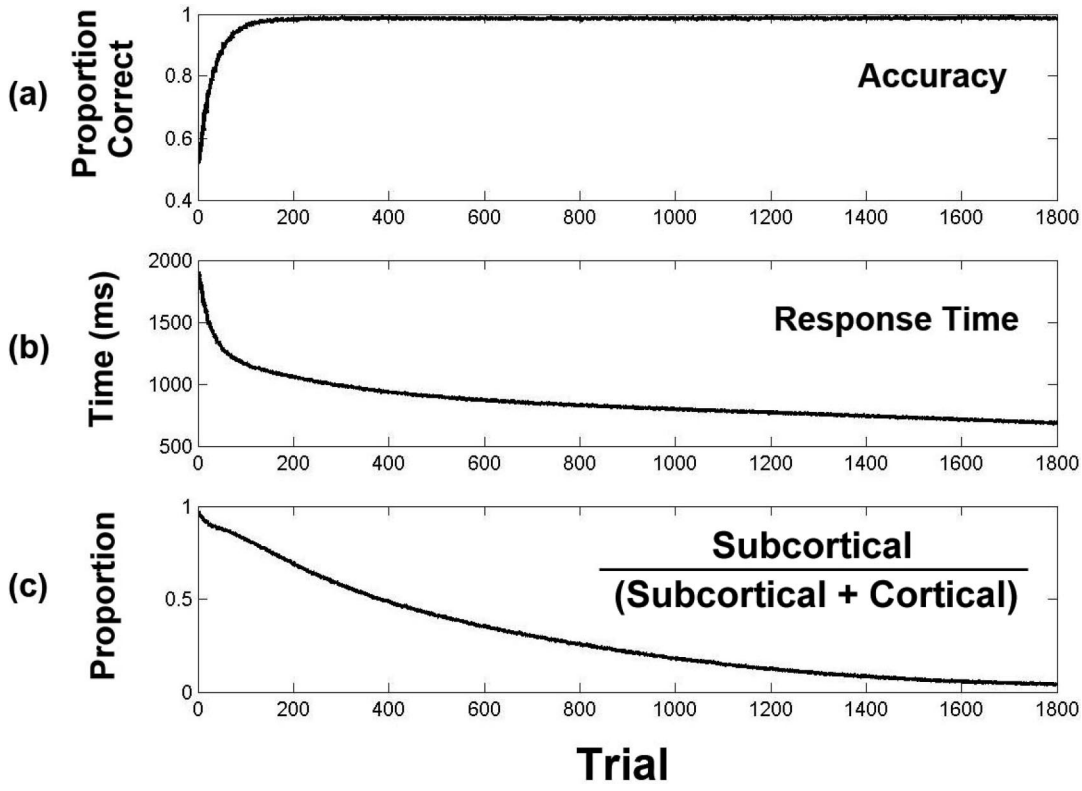

Figure 10. (a) Proportion correct and (b) mean response time on each trial of 3,000 SPEED (Subcortical Pathways Enable Expertise Development) simulations of the Nosofsky and Palmeri (1997) color category-learning experiment. Figure 10c shows the proportion of the total premotor activation (i.e., in the unit that controlled the response) that was generated by the subcortical path of SPEED in each simulation.

tributes, the faster the RT, so Figure 10c shows why SPEED predicts a continuous improvement in RT and why the RT distributions are unimodal. Figure 10 also shows that RT continues to decrease even after control has largely been passed to the cortical path. This is because the strength of the cortical-cortical synapses continues to increase (i.e., according to Equation 9) throughout training.

One thing Figure 10 does not show is that if the cortical path were lesioned, then SPEED would still show some RT improvement. This is because as the subcortical path learns, the critical cortical-striatal synaptic strengths increase, which allows sensory cortex to elicit a striatal response more quickly, which lowers RT. Thus, the RT improvement exhibited by SPEED is not due simply to passing control from a pathway with four synapses to a pathway with one synapse. Both paths themselves become faster because they both include a synapse that is gradually strengthened. This feature, together with the fact that it takes time for either pathway to generate enough postsynaptic activation in the premotor units to trigger a response, ensures that RT improvements in SPEED occur gradually (rather than suddenly) and that the RT distributions are unimodal.

This application shows that the predictions of SPEED are consistent with the power-law speed-up in RT, which is the most widely known result of the human expertise literature. In addition, Figure 11 shows that SPEED makes roughly the correct predictions about the effects of practice on RT variance and on the shape of the RT density function.

Relations to Other Models

Models of Expertise

Although we know of no other neurobiologically detailed models of categorization automaticity or expertise, there are a number of cognitive-based models. In this section, we briefly consider the relationship between SPEED and the more popular of these models.

Perhaps the most widely known cognitive model of expertise is Logan's (1988) instance theory of automaticity, which assumes that when a skilled behavior must be emitted, there is a race between accessing the procedures for computing the behavior and recalling a previous instance in which the behavior was successfully executed. With more and more practice, there are more and more stored instances of the behavior, so recall becomes faster and faster. SPEED is similar to the instance theory in that both models postulate two processes. In addition, in both cases, the process that dominates early performance is slow, whereas the process that dominates late (expert) performance speeds up greatly with practice. Rickard (1997) proposed an alternative version of the instance model, called the component power laws (CMPL) theory, which

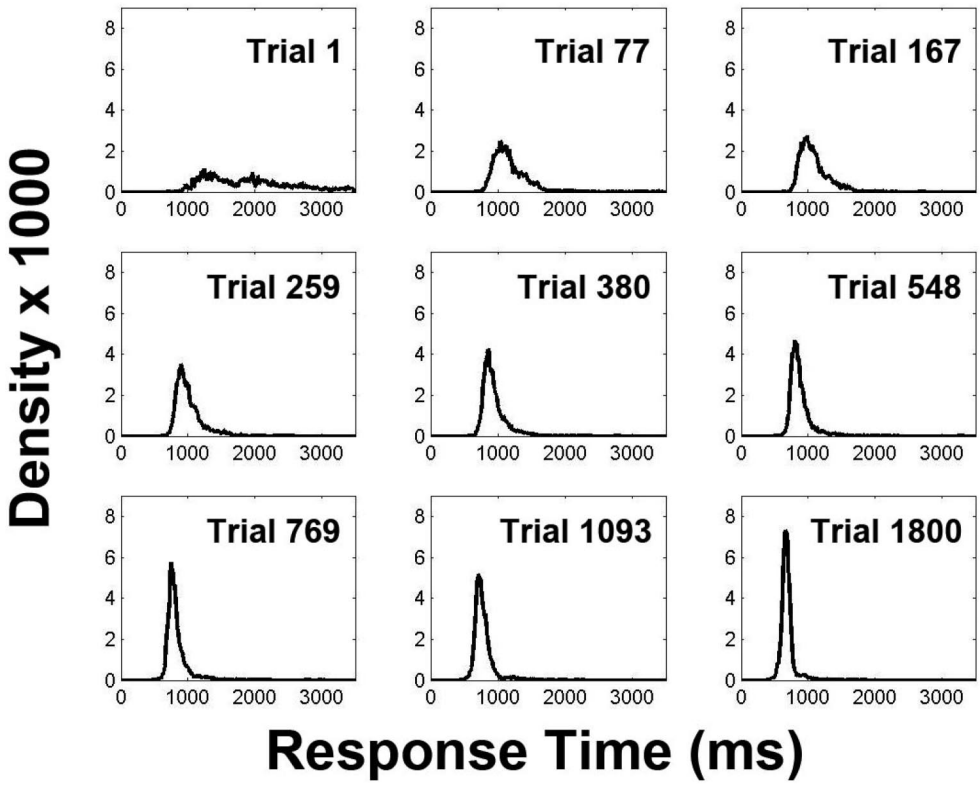

Figure 11. Predicted response time density functions estimated from 3,000 SPEED (Subcortical Pathways Enable Expertise Development) simulations of the Nosofsky and Palmeri (1997) color category-learning experiment.

assumes that on each trial, either the procedure or instance is executed, but not both. Thus, unlike Logan's (1988) model and SPEED, the CMPL model assumes no parallel race on each trial. Thus, SPEED is more similar to Logan's (1988) instance theory than to Rickard's (1997) CMPL model.

To our knowledge, the only model of categorization expertise that has been proposed is Nosofsky and Palmeri's (1997) exemplar-based random walk (EBRW) model. This model assumes that when a categorization stimulus is presented, the memory representations of previously seen exemplars from the two contrasting categories (call them A and B) are accessed, with the provision that exemplars that are more similar to the stimulus tend to be accessed more quickly. Response A is given as soon as the difference between the number of Category A and B exemplars accessed first exceeds a preset criterion. As expertise develops, there are more and more stored exemplars, so the criterion is reached more quickly. Thus, whereas SPEED proposes that the transition from novice to expert is characterized by a qualitative shift from one neural circuit to another, the EBRW assumes that the same process mediates all categorization performance, regardless of level of training.

The instance, CMPL, and EBRW models are purely cognitive models that make no neurobiological assumptions. Because they are expressed in such a different language from SPEED, detailed comparisons are difficult. At the minimum, some neural-based definition of instance and exemplar memory are needed. An obvious possibility is that the neural representations of both an instance and an exemplar memory are memory traces (e.g., of an episode) of the type consolidated by the hippocampus and other medial temporal lobe structures. For example, although there presently is no detailed neurobiological interpretation of exemplar theory, initial attempts to ground exemplar models of categorization in neurobiology have all assigned a key role to the hippocampus, at least during early learning (e.g., Pickering, 1997; Sakamoto & Love, 2004). According to this hypothesis, SPEED sharply disagrees with instance and exemplar-based models because of its focus on the basal ganglia as a key site of early learning in information-integration tasks, rather than the hippocampus. As mentioned above, this critical assumption of SPEED is supported by a variety of different experimental results (for a review, see, e.g., Ashby & Ennis, 2006). For example, there are numerous reports that basal ganglia disease patients are impaired in information-integration category learning (e.g., Filoteo et al., 2001a; Filoteo et al., 2005), whereas Filoteo et al. (2001b) reported that medial temporal lobe amnesiacs learn normally in information-integration categorization tasks.

It is important to note that we are not arguing that the hippocampus and other medial temporal lobe structures are not critical for normal category learning, only that they are not critical for learning in information-integration tasks. With other types of categories, medial temporal lobe structures might be vital for normal learning. For example, Ashby and O'Brien (2005) hypothesized

that the hippocampus may be critical in learning unstructured categories that contain only a few exemplars (e.g., my siblings; my personal numbers, such as zip code, phone number, and the like). Also, it has been hypothesized that the hippocampus is vital for long-term rule-based category learning or when learning complex rule-based categories (e.g., Ashby & Valentin, 2005; Nomura et al., 2007).

SPEED is more similar to exemplar models at the computational level than at the neurobiological level (e.g., they are both nonparametric classifiers; Ashby & Alfonso-Reese, 1995). Even so, the subcortical component of SPEED is not a neurobiological implementation of an exemplar model (Ashby & Waldron, 1999). It could be viewed, however, as a neurobiological version of other cognitive-based categorization models, including the striatal pattern classifier (Ashby & Waldron, 1999), Anderson's (1991) rational model, or the covering version of Kruschke's (1992) attentional learning covering map (ALCOVE) model (rather than to the more widely used exemplar-based version of ALCOVE). In each of these models, a low-resolution decision grid is associated with a high-resolution perceptual space, and response selection is based on what grid points are activated by the stimulus.

Models of Object Recognition

Riesenhuber and Poggio (1999) proposed a neurobiologically detailed model of object recognition that mimics the firing properties of cells in inferotemporal cortex when monkeys are categorizing objects as cats or dogs (Freedman et al., 2003). The model has a hierarchical structure in which cells from earlier visual areas tuned to simpler features feed forward to cells in higher visual areas that encode more complex and view-invariant objects. Despite the success of this model in mimicking single-cell firing data from categorization tasks, a number of results suggest that the Riesenhuber and Poggio (1999) model should not be viewed as a competitor to SPEED.

First, several studies have reported that following categorization training, cells in inferotemporal cortex showed enhanced sensitivity to diagnostic features compared with features that were irrelevant to the categorization judgment (Sigala, 2004; Sigala & Logothetis, 2002). Such changes are consistent with the widely held view that category learning is often associated with changes in the allocation of perceptual attention (Nosofsky, 1986). Second and most critical are the studies showing that categorization training did not make inferotemporal cortex neurons more likely to respond to other stimuli in the same category, or less likely to respond to stimuli belonging to the contrasting category (Freedman et al., 2003; Op de Beeck et al., 2001; Sigala, 2004; Thomas et al., 2001; Vogels, 1999). Third, Rolls et al. (1977) showed that the firing properties of cells in inferotemporal cortex did not change when the motor responses associated with category membership were switched (i.e., from "approach" to "avoid" and vice versa). For these reasons, the best evidence seems to suggest that inferotemporal cortex does not mediate the learning of new categories. Even so, this visual association area is crucial to the categorization process because it appears to encode a high-level representation of the visual stimulus. Cells in inferotemporal cortex project to the body and tail of the caudate nucleus, so the Riesenhuber and Poggio (1999) model could be incorporated into SPEED by using it to create a much more powerful model of the sensory input to the caudate and premotor units of SPEED.

Limitations

In its present formulation, SPEED has many limitations that prevent it from being considered a complete model of categorization expertise. First, it is meant only as a model in information-integration tasks. There is good evidence that other brain areas are recruited during category learning with different types of category structures. For example, a variety of evidence suggests that rule-based category learning depends on PFC (e.g., Kimberg, D'Esposito, & Farah, 1997; Monchi, Petrides, Petre, Worsley, & Dagher, 2001; Waldron & Ashby, 2001). In fact, a number of studies have shown that the application of a wide variety of abstract rules depends critically on the PFC (Asaad, Rainer, & Miller, 2000; Hoshi, Shima, & Tanji, 1998; Muhammad, Wallis, & Miller, 2006; Wallis, Anderson, & Miller, 2001; White & Wise, 1999).

An important generalization of SPEED will be to extend it to rule-based tasks. An obvious possibility would be to begin with the COVIS explicit system, which assumes that rule-based performance is largely mediated by a network that includes the PFC, the anterior cingulate, the hippocampus, and the head of the caudate nucleus (e.g., Ashby & Valentin, 2005). Then, as in SPEED, one could explore the possibility that the development of expertise coincides with a gradual transfer of control to cortical-cortical projections. Some neuroscience data support a model of this type. For example, Muhammad et al. (2006) recorded from single neurons in the PFC, head of the caudate, and premotor cortex while monkeys were applying rules. As predicted by COVIS, they found many cells in the PFC and caudate that fired selectively to a particular rule. However, after training the monkeys for 1 year, they also found many premotor cells that were rule selective, and, even more importantly, these cells responded, on average, about 100 ms before the PFC-rule selective cells. This result suggests a model in which the COVIS rule-based system trains cortical-cortical projections in much the same way as in SPEED.

SPEED must also be considered incomplete as a model of information-integration category learning. For example, we have greatly oversimplified the neuroanatomy of the basal ganglia, omitting striatal interneurons, striatal striosomes (i.e., patch compartments), as well as the indirect pathway out of the striatum (i.e., via the external segment of the globus pallidus and the subthalamic nucleus). This last omission is potentially serious because there are recent proposals that the indirect pathway plays a prominent role in learning what responses not to make when a stimulus is presented (Frank, 2005; Frank, Seeberger, & O'Reilly, 2004).

One possible consequence of these omissions is that the subcortical component of SPEED might oversimplify the process via which motor goals become associated with categories. In the present version of the model, the striatal units are each associated with a specific motor response goal before the start of category learning (e.g., press the left-response key). Thus, if the model were trained on a particular category and the response locations associated with each category were switched after training was complete, then the model would have to relearn the category responses from scratch. Such response switching does significantly interfere with the expression of information-integration category learning,

although interestingly it does not interfere with rule-based category learning (Ashby et al., 2003). Even so, it is natural to expect the recovery from response switching in information-integration tasks to be faster than the initial category learning (Ashby et al., 2003, did not look at recovery). If so, then it is likely that the striatal units become associated with specific motor goals at some point after the cortical-striatal synapse.

Conclusions

SPEED formalizes and tests a hypothesis of Ashby et al. (1998) that a primary function of the COVIS procedural learning system is to train faster cortical-cortical projections. In the nearly 10 years since COVIS was first proposed, a wide variety of evidence has been reported that supports the basic notion of COVIS that information-integration category learning is initially mediated largely within the striatum. Among the most convincing of such data are the single-unit recordings of Merchant et al. (1997), shown in the top panel of Figure 5, which show that category-specific responses develop in the striatum following initial training in tactile category learning. Figures 6 and 7 show that SPEED gives an excellent account of these data.

Equally important, there is an accumulating set of convincing data suggesting that COVIS is inadequate as a model of automatic categorization judgments. For example, basal ganglia disease patients, who are impaired in learning new information-integration categories, do not lose old familiar categories. In addition, although the early training data of Carelli et al. (1997) and Choi et al. (2005) are compatible with COVIS (i.e., up to the point at which accuracy asymptotes), the data collected in these studies after the animals had extensive training strongly contradict COVIS. Nevertheless, SPEED provides excellent accounts of both data sets, and it is perhaps these fits that provide the strongest overall evidence supporting SPEED as a general model of automaticity. Our last application showed that SPEED is also compatible with the most widely known behavioral phenomenon associated with human expertise.

In the present article, we fit SPEED to single-cell recording and behavioral data, and we modeled pharmacological treatments. Despite the diversity of these applications, they only scratch the surface of the possible empirical tests of SPEED. Formulating the theory as we did, in a neurobiologically detailed manner, allows tests of the model in a wide variety of experimental domains. For example, by adding a model of how neural activation is related to the fMRI blood oxygen level-dependent (BOLD) signal, SPEED could be used to generate predicted BOLD signals in each of the brain regions shown in Figure 3 and to make predictions about how these BOLD responses are changed by practice (e.g., by using the methods described by Ashby & Valentin, 2006). The model could also be generalized in a straightforward manner to make precise predictions about the categorization abilities of a variety of neuropsychological patient groups (e.g., PD, Huntington's disease). This diversity in possible applications is among the greatest strengths of SPEED and of the modeling approach that was used in its construction.

References

Akins, P. T., Surmeier, D. J., & Kitai, S. T. (1990, March 15). Muscarinic modulation of a transient K^+ conductance in rat neostriatal neurons. *Nature, 344,* 240–242.

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review, 98,* 409–429.

Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.

Anderson, M. E., & DeVito, J. L. (1987). An analysis of potentially converging inputs to the rostral ventral thalamic nuclei of the cat. *Experimental Brian Research, 68,* 260–276.

Arbuthnott, G. W., Ingham, C. A., & Wickens, J. R. (2000). Dopamine and synaptic plasticity in the neostriatum. *Journal of Anatomy, 196,* 587–596.

Asaad, W. F., Rainer, G., & Miller, E. K. (2000). Task-specific neural activity in the primate prefrontal cortex. *Journal of Neurophysiology, 84,* 451–459.

Ashby, F. G. (2000). A stochastic version of general recognition theory. *Journal of Mathematical Psychology, 44,* 310–329.

Ashby, F. G., & Alfonso-Reese, L. (1995). Categorization as probability density estimation. *Journal of Mathematical Psychology, 39,* 216–233.

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

Ashby, F. G., Ell, S. W., Valentin, V. V., & Casale, M. B. (2005). Frost: A distributed neurocomputational model of working memory maintenance. *Journal of Cognitive Neuroscience, 17,* 1728–1743.

Ashby, F. G., Ell, S. W., & Waldron, E. M. (2003). Procedural learning in perceptual categorization. *Memory & Cognition, 31,* 1114–1125.

Ashby, F. G., & Ennis, J. M. (2006). The role of the basal ganglia in category learning. *The Psychology of Learning and Motivation, 46,* 1–36.

Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 33–53.

Ashby, F. G., Isen, A. M., & Turken, A. U. (1999). A neuropsychological theory of positive affect and its influence on cognition. *Psychological Review, 106,* 529–550.

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology, 56,* 149–178.

Ashby, F. G., Maddox, W. T., & Bohil, C. J. (2002). Observational versus feedback training in rule-based and information-integration category learning. *Memory & Cognition, 30,* 665–676.

Ashby, F. G., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *TRENDS in Cognitive Science, 2,* 83–89.

Ashby, F. G., Queller, S., & Berretty, P. M. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics, 61,* 1178–1199.

Ashby, F. G., & Spiering, B. J. (2004). The neurobiology of category learning. *Behavior and Cognitive Neuroscience Reviews, 3,* 101–113.

Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review, 93,* 154–179.

Ashby, F. G., & Valentin, V. V. (2005). Multiple systems of perceptual category learning: Theory and cognitive tests. In H. Cohen & C. Lefebvre (Eds.), *Handbook of categorization in cognitive science* (pp. 547–572). New York: Elsevier.

Ashby, F. G., & Valentin, V. V. (2006). Computational cognitive neuroscience: Building and testing biologically plausible computational models of neuroscience, neuroimaging, and behavioral data. In M. J. Wenger (Ed.), *2004 Notre Dame series on quantitative methodology* (pp. 15–58). Hillsdale, NJ: Erlbaum.

Ashby, F. G., Valentin, V. V., & Turken, A. U. (2002). The effects of positive affect and arousal on working memory and executive attention: Neurobiology and computational models. In S. Moore & M. Oaksford (Eds.), *Emotional cognition: From brain to behaviour* (pp. 245–287). Amsterdam: John Benjamins.

Ashby, F. G., & Waldron, E. M. (1999). On the nature of implicit categorization. *Psychonomic Bulletin & Review, 6,* 363–378.

Bar-Gad, I., Morris, G., & Bergman, H. (2003). Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. *Progress in Neurobiology, 71,* 439–473.

Beiser, D. G., & Houk, J. C. (1998). Model of cortical-basal ganglionic processing: Encoding the serial order of sensory events. *Journal of Neurophysiology, 79,* 3168–3188.

Brannon, M. J., & Roth, R. H. (1983). Pharmacology of mesocortical dopamine neurons. *Pharmacological Reviews, 35,* 53–67.

Brasted, P. J., & Wise, S. P. (2004). Comparison of learning-related neuronal activity in the dorsal premotor cortex and striatum. *European Journal of Neuroscience, 19,* 721–740.

Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 713–737). Cambridge, MA: MIT Press.

Brooks, L. (1978). *Nonanalytic concept formation and memory for instances: Cognition and categorization.* Hillsdale, NJ: Erlbaum:

Brown, J., Bullock, D., & Grossberg, S. (1999). How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected rewarding cues. *Journal of Neuroscience, 19,* 10502–10511.

Buhmann, M. D. (2003). *Radial basis functions: Theory and implementations.* New York: Cambridge University Press.

Caan, W., Perrett, D. I., & Rolls, E. T. (1984). Responses of striatal neurons in the behaving monkey. 2. Visual processing in the caudal neostriatum. *Brain Research, 290,* 53–65.

Calabresi, P., Pisani, A., Mercuri, N. B., & Bernardi, G. (1996). The corticostriatal projection: From synaptic plasticity to dysfunctions of the basal ganglia. *Trends in Neurosciences, 19,* 19–24.

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

Choi, W. Y., Balsam, P. D., & Horvitz, J. C. (2005). Extended habit training reduces dopamine mediation of appetitive response expression. *The Journal of Neuroscience, 25,* 6729–6733.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review, 7,* 332–361.

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

Crossman, E. R. F. W. (1959). A theory of the acquisition of speed-skill. *Ergonomics, 2,* 153–166.

Difiglia, M., Pasik, T., & Pasik, P. (1978). A Golgi study of afferent fibers in the neostriatum of monkeys. *Brain Research, 152,* 341–347.

Eacott, M. J., & Gaffan, D. (1991). The role of monkey inferior parietal cortex in visual discrimination of identity and orientation of shapes. *Behavioural Brain Research, 46,* 95–98.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127,* 107–140.

Estes, W. K. (1986). Array models for category learning. *Cognitive Psychology, 18,* 500–549.

Estes, W. K. (1994). *Classification and cognition*. New York: Oxord University Press.

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

Filoteo, J. V., Maddox, W. T., & Davis, J. D. (2001a). A possible role of the striatum in linear and nonlinear categorization rule learning: Evidence from patients with Huntington's disease. *Behavioral Neuroscience, 115,* 786–798.

Filoteo, J. V., Maddox, W. T., & Davis, J. D. (2001b). Quantitative modeling of category learning in amnesic patients. *Journal of the International Neuropsychological Society, 7,* 1–19.

Filoteo, J. V., Maddox, W. T., Salmon, D. P., & Song, D. D. (2005). Information-integration category learning in patients with striatal dysfunction. *Neuropsychology, 19,* 212–222.

Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of the cognitive deficits in medicated and non-medicated Parkinsonism. *Journal of Cognitive Neuroscience, 17,* 51–72.

Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, & Behavioral Neuroscience, 1,* 137–160.

Frank, M. J., Seeberger, L., & O'Reilly, R. C. (2004, December 10). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science, 306,* 1940–1943.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2001, January 12). Categorical representations of visual stimuli in the primate prefrontal cortex. *Science, 291,* 312–316.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2002). Visual categorization and the primate prefrontal cortex: Neurophysiology and behavior. *Journal of Neurophysiology, 88,* 929–941.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2003). A comparison of primate prefrontal and inferior temporal cortices during visual categorization. *The Journal of Neuroscience, 23,* 5235–5246.

Freund, T. F., Powell, J. F., & Smith, A. D. (1984). Tyrosine hydroxylase-immunoreactive boutons in synaptic contact with identified striatonigral neurons, with particular reference to dendritic spine. *Neuroscience, 13,* 1189–1215.

Fuster, J. M. (2001). The prefrontal cortex—an update: Time is of the essence. *Neuron, 30,* 319–333.

Gaffan, D., & Eacott, M. J. (1995). Visual learning for an auditory secondary reinforcer by macaques is intact after uncinate fascicle section: Indirect evidence for the involvement of the corpus striatum. *European Journal of Neuroscience, 7,* 1866–1871.

Gaffan, D., & Harrison, S. (1987). Amygdalectomy and disconnection in visual learning for auditory secondary reinforcement by monkeys. *Journal of Neuroscience, 7,* 2285–2292.

Gamble, E., & Koch, C. (1987, June 5). The dynamics of free calcium in dendritic spines in response to repetitive synaptic input. *Science, 236,* 1311–1315.

Gauthier, I., & Tarr, M. J. (1997). Becoming a "Greeble" expert: Exploring mechanisms for face recognition. *Vision Research, 37,* 1673–1682.

Heathcote, A., Brown, S., & Mewhort, D. J. K. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review, 7,* 185–207.

Heimer, L. (1995). *Human brain and spinal cord*. New York: Springer-Verlag.

Hemmings, H. C., Walaas, S. I., Ouimet, C. C., & Greengard, P. (1987). Dopaminergic regulation of protein phosphorylation in the striatum: DARRP-32. *Trends in Neuroscience, 10,* 377–383.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review, 93,* 411–428.

Hollerman, J. R., & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience, 1,* 304–309.

Homa, D., Sterling, S., & Trepel, L. (1981). Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 418–439.

Hoshi, E., Shima, K., & Tanji, J. (1998). Task-dependent selectivity of movement-related neuronal activity in the primate prefrontal cortex. *Journal of Neurophysiology, 80,* 3392–3397.

Humphreys, G. W., & Forde, E. M. E. (2001). Hierarchies, similarity, and interactivity in object recognition: "Category-specific" neuropsychological deficits. *Behavioral and Brain Sciences, 24,* 453–509.

Humphries, M. D., Stewart, R. D., & Gurney, K. (2006). A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. *Journal of Neuroscience, 26,* 12921–12942.

Joo-Er, M., Wu, S., Lu., J., & Toh, H. L. (2002). Face recognition with radial basis function (RBF) neural networks. *IEEE Transactions on the Neural Networks, 13,* 697–710.

Joseph, J. E. (2001). Functional neuroimaging studies of category specificity in object recognition: A critical review and meta-analysis. *Cognitive, Affective, & Behavioral Neuroscience, 1,* 119–136.

Karni, A., Meyer, G., Jezzard, P., Adams, M. M., Turner, R., & Ungerleider, L. G. (1995, September 14). Functional MRI evidence for adult motor cortex plasticity during motor skill learning. *Nature, 377,* 155–158.

Kerr, J. N. D., & Wickens, J. R. (2001). Dopamine D-1/D-5 receptor activation is required for long-term potentiation in the rat neostriatum in vitro. *Journal of Neurophysiology, 85,* 117–124.

Kimberg, D. Y., D'Esposito, M., & Farah, M. J. (1997). Effects of bromocriptine on human subjects depends on working memory capacity. *NeuroReport, 8,* 3581–3585.

Kimura, M., Rajkowski, J., & Evarts, E. (1984). Tonically discharging putamen neurons exhibit set-dependent responses. *Proceedings of the National Academy of Sciences, USA, 81,* 4998–5001.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996, September 6). A neostriatal habit learning system in humans. *Science, 273,* 1399–1402.

Koch, C. (1999). *Biophysics of computation*. New York: Oxford University Press.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99,* 22–44.

Lamberts, K. (2000). Information-accumulation theory of speeded categorization. *Psychological Review, 107,* 227–260.

Lashley, K. S. (1950). In search of the engram. In *Society of Experimental Biology Symposium* (Vol. 4, pp. 454–480). Cambridge, England: Cambridge University Press.

Lehéricy, S., Benali, H., Van de Moortele, P. F., Pélégrini-Issac, M., Waechter, T., Ugurbil, K., & Doyon, J. (2005). Distinct basal ganglia territories are engaged in early and advanced motor sequence learning. *Proceedings of the National Academy of Sciences, USA, 102,* 12566–12571.

Lei, W., Jiao, Y., Del Mar, N., & Reiner, A. (2004). Evidence for differential cortical input to direct pathway versus indirect pathway striatal projection neurons in rats. *Journal of Neuroscience, 24,* 8289–8299.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492–527.

Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.

MacDermott, A. B., Mayer, M. L., Westbrook, G. L., Smith, S. J., & Barker, J. L. (1986, May 29). NMDA-receptor activation increases cytoplasmic calcium concentration in cultured spinal cord neurones. *Nature, 321,* 519–522.

MacKay, D. G. (1982). The problem of flexibility, fluency, and speed-accuracy trade-off in skilled behavior. *Psychological Review, 89,* 483–505.

Maddox, W. T., & Ashby, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics, 53,* 49–70.

Maddox, W. T., & Ashby, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioural Processes, 66,* 309–332.

Maddox, W. T., Ashby, F. G., & Bohil, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 650–662.

Maddox, W. T., Ashby, F. G., Ing, A. D., & Pickering, A. D. (2004). Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Memory & Cognition, 32,* 582–591.

Maddox, W. T., Bohil, C. J., & Ing, A. D. (2004). Evidence for a procedural learning-based system in perceptual category learning. *Psychonomic Bulletin & Review, 11,* 945–952.

Maddox, W. T., & Filoteo, J. V. (2005). The neuropsychology of perceptual category learning. In H. Cohen & C. Lefebvre (Eds.), *Handbook of categorization in cognitive science* (pp. 573–599). New York: Elsevier.

Malenka, R. C., & Nicoli, R. A. (1999, September 17). Long-term potentiation–A decade of progress? *Science, 285,* 1870–1874.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102,*419–457.

McDonald, R. J., & White, N. M. (1993). A triple dissociation of memory systems: Hippocampus, amygdala, and dorsal striatum. *Behavioral Neuroscience, 107,* 3–22.

McDonald, R. J., & White, N. M. (1994). Parallel information processing in the water maze: Evidence for independent memory systems involving dorsal striatum and hippocampus. *Behavioral and Neural Biology, 61,* 260–270.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85,* 207–238.

Merchant, H., Zainos, A., Hernadez, A., Salinas, E., & Romo, R. (1997). Functional properties of primate putamen neurons during the categorization of tactile stimuli. *Journal of Neurophysiology, 77,* 1132–1154.

Miller, J. D., Sanghera, M. K., & German, D. C. (1981). Mesencephalic dopaminergic unit activity in the behaviorally conditioned rat. *Life Sciences, 29,* 1255–1263.

Mirenowicz, J., & Schultz, W. (1994). Importance of unpredictability for reward responses in primate dopamine neurons. *Journal of Neurophysiology, 72,* 1024–1027.

Monchi, O., Petrides, M., Petre, V., Worsley, K., & Dagher, A. (2001). Wisconsin Card Sorting revisited: Distinct neural circuits participating in different stages of the task identified by event-related functional magnetic resonance imaging. *Journal of Neuroscience, 21,* 7733–7741.

Monchi, O., Taylor, J. G., & Dagher, A. (2000). A neural model of working memory processes in normal subjects, Parkinson's disease, and schizophrenia for fMRI design and predictions. *Neural Networks, 13,* 953–973.

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience, 16,* 1936–1947.

Muhammad, R., Wallis, J. D., & Miller, E. K. (2006). A comparison of abstract rules in the prefrontal cortex, premotor cortex, inferior temporal cortex, and striatum. *Journal of Cognitive Neuroscience, 18,* 1–16.

Nairn, A. C., Hemmings, H. C., Walaas, S. I., & Greengard, P. (1988). DARPP-32 and phosphatase inhibitor-1, two structurally related inhibitors of protein phosphatase-1, are both present in striatonigral neurons. *Journal of Neurochemistry, 50,* 257–262.

Neves, D. M., & Anderson, J. R. (1981). Knowledge compilation: Mechanisms for the automatization of cognitive skills. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 57–84). Hillsdale, NJ: Erlbaum.

Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 1–55). Hillsdale, NJ: Erlbaum.

Nomura, E. M., Maddox, W. T., Filoteo, J. V., Ing, A. D., Gitelman, D. R., Parrish, T. B., et al. (2007). Neural correlates of rule-based and information-integration visual category learning. *Cerebral Cortex, 17,* 37–43.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39–57.

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review, 104,* 266–300.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review, 101,* 53–79.

Oglesby, J., & Mason, J. S. (1991). Radial basis function networks for speaker recognition. *International Conference on Acoustics, Speech, and Signal Processing, 91,* 393–396.

Op de Beeck, H., Wagemans, J., & Vogels, R. (2001). Inferotemporal neurons represent low-dimensional configurations of paramterized shapes. *Nature Neuroscience, 4,* 1244–1252.

O'Reilly, R. C., Braver, T. S., & Cohen, J. D. (1999). A biologically based computational model of working memory. In A. Miyake & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control* (pp. 375–411). New York: Cambridge 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.

Packard, M. G., Hirsh, R., & White, N. M. (1989). Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: Evidence for multiple memory systems. *Journal of Neuroscience, 9,* 1465–1472.

Packard, M. G., & McGaugh, J. L. (1992). Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience, 106,* 439–446.

Parzen, E. (1962). On estimation of a probability density function and mode. *Annals of Mathematical Statistics, 33,* 1065–1076.

Pessin, M. S., Snyder, G. L., Halpain, S., Girault, J. A., Aperia, A., & Greengard, P. (1994). DARPP32/protein phosphatase1/Na + /K + AT-Pase system: A mechanism for bidirectional control of cell function. In K. Fuxe, L. F. Agnati, B. Bjelke, & D. Ottoson (Eds.), *Trophic regulation of basal ganglia* (pp. 43–58). New York: Elsevier.

Pickering, A. D. (1997). New approaches to the study of amnesic patients: What can a neurofunctional philosophy and neural network methods offer? *Memory, 5,* 255–300.

Pitt, M. A., Kim, W., Navarro, D. J., & Myung, J. I. (2006). Global model analysis by parameter space partitioning. *Psychological Review, 113,* 57–83.

Poldrack, R. A., Sabb, F., Foerde, K., Tom, S., Asarnow, R., Bookheimer, S., & Knowlton, B. J. (2005). The neural correlates of automaticity during motor skill learning. *Journal of Neuroscience, 25,* 5356–5364.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77,* 353–363.

Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology, 83,* 304–308.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85,* 59–108.

Ravel, S., Legallet, E., & Apicella, P. (2003). Responses of tonically active neurons in the monkey striatum discriminate between motivationally opposing stimuli. *Journal of Neuroscience, 23,* 8489–8497.

Reber, P. J., Stark, C. E. L., & Squire, L. R. (1998). Cortical areas supporting category learning identified using functional MRI. *Proceedings of the National Academy of Sciences, USA, 95,* 747–750.

Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology, 3,* 382–407.

Rickard, T. C. (1997). Bending the power law: A CMPL theory of strategy shifts and the automatization of cognitive skills. *Journal of Experimental Psychology: General, 126,* 288–311.

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience, 2,* 1019–1025.

Rolls, E. T., Judge, S. J., & Sanghera, M. (1977). Activity of neurons in the inferotemporal cortex of the alert monkey. *Brain Research, 130,* 229–238.

Romo, R., Merchant, H., Ruiz, S., Crespo, P., & Zainos, A. (1995). Neuronal activity of primate putamen during categorical perception of somaesthetic stimuli. *NeuroReport, 6,* 1013–1017.

Romo, R., Merchant, H., Zainos, A., & Hernandez, A. (1997). Categorical perception of somesthetic stimuli: Psychophysical measurements correlated with neuronal events in primate medial premotor cortex. *Cerebral Cortex, 7,* 317–326.

Rosch, E. (1973). Natural categories. *Cognitive Psychology, 4,* 328–350.

Rosch, E. (1975). Cognitive reference points. *Cognitive Psychology, 7,* 532–547.

Rosenblum, M., Yacoob, Y., & Davis, L. (1996). Human emotion recognition from motion using a radial basis function network architecture. *IEEE Transactions on the Neural Networks, 7,* 1121–1138.

Sakamoto, Y., & Love, B. C. (2004). Schematic influences on category learning and recognition memory. *Journal of Experimental Psychology: General, 133,* 534–553.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84,* 127–190.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology, 80,* 1–27.

Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron, 36,* 241–263.

Schultz, W., Dayan, P., & Montague, P. R. (1997, March 14). A neural substrate of prediction and reward. *Science, 275,* 1593–1599.

Seamans, J. K., & Yang, C. R. (2004). The principal features and mechanisms of dopamine modulation in the prefrontal cortex. *Progress in Neurobiology, 74,* 1–57.

Seger, C. A., & Cincotta, C. M. (2002). Striatal activity in concept learning. *Cognitive, Affective, & Behavioral Neuroscience, 2,* 149–161.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (Area LIP) of the rhesus monkey. *Journal of Neurophysiology, 86,* 1916–1936.

Sherrington, C. S. (1906). Observations on the scratch reflex in the spinal dog. *Journal of Physiology, 34,* 1–50.

Shook, B. L., Schlag-Rey, M., & Schlag, J. (1991). Primate supplementary eye field. II. Comparative aspects of connections with the thalamus, corpus striatum, and related forebrain nuclei. *Journal of Comparative Neurology, 307,* 562–583.

Sigala, N. (2004). Visual categorization and the inferior temporal cortex. *Behavioural Brain Research, 149,* 1–7.

Sigala, N., & Logothetis, N. K. (2002, January 17). Visual categorization shapes feature selectitivity in the primate temporal cortex. *Nature, 415,* 318–320.

Smiley, J. F., Levey, A. I., Ciliax, B. J., & Goldman-Rakic, P. S. (1994). D_1 dopamine receptor immunoreactivity in human and monkey cerebral cortex: Predominant and extrasynaptic localization in dendritic spines. *Proceedings of the National Academy of Sciences, USA, 91,* 5720–5724.

Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 24,* 1411–1430.

Squire, L. R., Stark, C. E. L., & Clark, R. (2004). The medial temporal lobe. *Annual Review of Neuroscience, 27,* 279–306.

Suzuki, T., Miura, M., Nishimura, K., & Aosaki, T. (2001). Dopamine-dependent synaptic plasticity in the striatal cholinergic interneurons. *Journal of Neuroscience, 21,* 6492–6501.

Thomas, E., Van Hulle, M. M., & Vogels, R. (2001). Encoding of categories by noncategory-specific neurons in the inferior temporal cortex. *Journal of Cognitive Neuroscience, 13,* 190–200.

Tobler, P. N., Dickinson, A., & Schultz, W. (2003). Coding of predicted reward omission by dopamine neurons in a conditioned inhibition paradigm. *Journal of Neuroscience, 23,* 10402–10410.

Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review, 108,* 550–592.

Vogels, R. (1999). Complex visual images by rhesus monkeys. Part 2: Single-cell study. *European Journal of Neuroscience, 11,* 1239–1255.

Waldron, E. M., & Ashby, F. G. (2001). The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychonomic Bulletin & Review, 8,* 168–176.

Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001, June 21). Single neurons in prefrontal cortex encode abstract rules. *Nature, 411,* 953–956.

White, I. M., & Wise, S. P. (1999). Rule-dependent neuronal activity in the prefrontal cortex. *Experimental Brain Research, 126,* 315–335.

Wickens, J. (1990). Striatal dopamine in motor activation and reward-mediated learning: Steps towards a unifying model. *Journal of Neural Transmission: General Section, 80,* 9–31.

Wickens, J. (1993). *A theory of the striatum*. New York: Pergamon Press.

Williams, G. V., & Goldman-Rakic, P. S. (1995). Modulation of memory fields by dopamine D1 receptors in prefrontal cortex. *Nature, 376,* 572–575.

Willingham, D. B., Wells, L. A., Farrell, J. M., & Stemwedel, M. E. (2000). Implicit motor sequence learning is represented in response locations. *Memory & Cognition, 28,* 366–375.

Wilson, C. J. (1995). The contribution of cortical neurons to the firing pattern of striatal spiny neurons. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 29–50). Cambridge, MA: Bradford.

Wu, T., Kansaku, K., & Hallett, M. (2004). How self-initiated memorized movements become automatic: A functional MRI study. *Journal of Neurophysiology, 91,* 1690–1698.

Zeithamova, D., & Maddox, W. T. (2006). Dual-task interference in perceptual category learning. *Memory & Cognition, 34,* 387–398.

Appendix

Simulation Methods

General Methods

In all applications, sample solutions[10] were estimated using Euler's method together with the constraint that the solutions were bounded below and above by 0 and 1, which are the asymptotic bounds of the true solutions. A dynamic integrate-and-fire method developed by Ashby and Valentin (2006) was used to model the flow of activation between brain regions. Briefly, this method integrates the activation in each region. When this integral exceeds a threshold, a square wave of activation is sent to the next region and the integral is reset to zero. One exception to this approach was that the activations from sensory cortex were assumed to be a constant multiple of the relevant synaptic strengths as long as the stimulus was present and 0 otherwise. This simplification was made to avoid the computation of 10,000 additional integrals at each time point and is effectively equivalent to assuming that sensory cortex has a constant rate of firing during the stimulus presentation.

Parameter Estimation

Our goal in the simulations was to determine whether SPEED is qualitatively consistent with the results of the various experiments. With single-cell recording data, it is likely that the responses of nearby cells differ quantitatively even if they show the same qualitative pattern of responding. Thus, our goal was not to provide precise quantitative fits to the various single-cell data. In addition, although we did not perform formal parameter space partitioning (Pitt, Kim, Navarro, & Myung, 2006), the qualitative predictions of SPEED were invariant under small changes to the parameter values. Specifically, for each parameter, there was a range of values under which the model would learn. Within this range, SPEED tended to make consistent qualitative predictions. For these reasons, we crudely searched the parameter space by hand.

Our general strategy was to first find parameter estimates that allowed SPEED to learn the Nosofsky and Palmeri (1997) categories (Application 4). In most cases, these estimates remained fixed for the other applications. For example, no new parameters were estimated in Application 1. However, task differences required that we adjust some parameter estimates in Applications 2 and 3. In particular, the experiments in Applications 1 and 4 used many stimuli, two alternative responses, and monkeys and humans as participants, respectively, whereas the experiments in Applications 2 and 3 used one stimulus, one response, and rats as participants. Because SPEED had only one premotor and striatal unit in Applications 2 and 3, but two each in Applications 1 and 4, for example, there was lateral inhibition between the SPEED premotor and striatal units in Applications 1 and 4 but not in Applications 2 and 3. This difference affected our estimates of the baseline activity levels in striatal and premotor areas.

The parameter estimates from all applications are listed in Table A1. Note that Applications 1 and 4 used identical estimates, as did Applications 2 and 3. Across these two types of tasks, the parameters that differed were (a) initial synaptic strengths from sensory cortex to striatum, (b) learning rates, (c) NMDA receptor thresholds, and (d) striatal and premotor baseline rates. These new values were estimated during Application 2 (i.e., Carelli et al., 1997) and used again in Application 3 (Choi et al., 2005). Thus, the only free

[10] ecause the differential equations listed in the text are stochastic, their solutions are stochastic processes. By sample solution, we mean a specific realization of the solution to a stochastic differential equation.

(*Appendix continues*)

Table A1

SPEED Parameter Estimates

Parameter	Applications 1 & 4	Applications 2 & 3
α	3	N/A
α_v	3.00×10^{-12}	2.00×10^{-13}
β_v	5.00×10^{-12}	2.00×10^{-13}
α_w	1.00×10^{-8}	4.50×10^{-11}
β_w	1.00×10^{-8}	3.00×10^{-11}
γ_w	1.00×10^{-8}	3.00×10^{-11}
φ_w	1.00×10^{-4}	3.00×10^{-7}
θ_{NMDA} (striatum)	800	300
θ_{NMDA} (premotor cortex)	400	500
D_{base}	.2	.2
S_{base}	.2	0
G_{base}	.7	.7
T_{base}	.4	.4
E_{base}	.2	0
β_S	.0085	N/A
γ_S	.004	.004
σ_S	.02	.02
α_G	.03	.03
β_G	.0025	.0025
α_T	.03	.03
β_T	.0025	.0025
α_E	.007	.007
β_E	.0085	N/A
γ_E	.004	.004
σ_E	.0125	.0125
τ	180	140

Note. SPEED = subcortical pathways enable expertise development; N/A = not applicable.

parameters in Application 3 corresponded to the effective reduction in dopamine that resulted from each of the three doses of the D1 antagonist. These estimates were 1 (vehicle), 0.95 (low dose), 0.86 (medium dose), and 0.5 (high dose).

Each application included a waiting period and a response deadline determined by the particular experiment featured in the application. In addition, there was a half-second period after the response was made to simulate the time during which feedback was given and processed. Inhibition was only active during stimulus presentation because the baseline firing rate models the result of inhibition and excitation from other cells during a neutral condition. In Applications 1 and 4, initial synaptic weights from sensory cortex to striatum were assumed to vary uniformly over the range (0.0002, 0.0002025). In Applications 2 and 3, the initial synaptic weight for the single synapse between sensory cortex and striatum was given the single number 0.0005. In all applications, initial synaptic weights from sensory cortex to premotor cortex were set to zero.

Sensory Unit Activation

As described in the text, Application 1 included 100 sensory cortical units, Applications 2 and 3 included one sensory cortical unit, and Application 4 included 10,000. In Application 1, each sensory cortical unit (somatosensory cortex) was assumed to have a preferred vibration speed, and these preferred speeds were equally spaced across the 100 units and spanned the range of presented stimuli (i.e., 12–30 mm/s). In Applications 2 and 3, the preferred stimulus of the single sensory cortical unit (auditory cortex) was the presented stimulus (i.e., tone). In Application 4, the 10,000 sensory cortical units (visual association cortex) were arranged in a two-dimensional rectangular grid, with axes corresponding to saturation and brightness. The rows and columns were equally spaced and covered the ranges of saturation and brightness used by Nosofsky and Palmeri (1997). The preferred stimulus of each unit was determined by its row and column. In all applications, the activation of each sensory cortical unit was determined by Equation 1. In each application, the stimulus was either feedback terminated or assumed to have a representation held in sensory cortex until feedback was processed.

Dopamine Levels

In all applications, the amount of dopamine released on each trial was determined by Equations 9 and 10. The probability of a correct response, $P(C)$, was estimated from SPEED's performance over the previous 50 trials. In Application 3, we assumed that the effect of the dopamine D1 antagonist was to restrict the dynamic range of the dopamine levels. For instance, without a drug effect, the dopamine levels varied between a maximal value of 1 and a minimal value of 0 about a baseline level of 0.2. With a high level of drug, dopamine levels varied between a maximal value of 0.45 and a minimal value of 0.1, still about a baseline level of 0.2.

Spike Trains

Finally, to produce the spike trains shown in Figure 7, we applied a standard integrate and fire model (e.g., Koch, 1999), as there was no advantage in this case to using a more complicated dynamic integrate and fire model (e.g., Ashby & Valentin, 2006).

Received March 23, 2006
Revision received December 15, 2006
Accepted December 18, 2006 ■