

Assembling Old Tricks for New Tasks: A Neural Model of Instructional Learning and Control

Tsung-Ren Huang, Thomas E. Hazy, Seth A. Herd,
and Randall C. O'Reilly

Abstract

■ We can learn from the wisdom of others to maximize success. However, it is unclear how humans take advice to flexibly adapt behavior. On the basis of data from neuroanatomy, neurophysiology, and neuroimaging, a biologically plausible model is developed to illustrate the neural mechanisms of learning from instructions. The model consists of two complementary learning pathways. The slow-learning parietal pathway carries out simple or habitual stimulus–response (S-R) mappings, whereas

the fast-learning hippocampal pathway implements novel S-R rules. Specifically, the hippocampus can rapidly encode arbitrary S-R associations, and stimulus-cued responses are later recalled into the basal ganglia-gated pFC to bias response selection in the premotor and motor cortices. The interactions between the two model learning pathways explain how instructions can override habits and how automaticity can be achieved through motor consolidation. ■

INTRODUCTION

We don't always need to learn things on our own, the hard way, from trial and error. By reading books and chatting with other people, we can immediately learn from the wisdom of others to avoid making mistakes. This way of learning from instruction/advice, unlike observational learning, does not require explicit demonstration and lies at the core of human communication and flexible behavior.

Given that instructional learning is arguably the most sophisticated form of learning seen in the animal kingdom, why does it appear so effortless in comparison with more basic, trial-and-error learning? How does instruction enable us to perform complex novel tasks perfectly on the first attempt? In this article, we address these questions with a biologically informed neural network model that can flexibly recombine old tricks for new tasks via mechanisms of hippocampal fast learning and prefrontal working memory.

Although instruction learning and following happen in almost every human laboratory experiment and in everyday life, learning from instruction has been much less explored from a mechanistic, computational modeling perspective, compared with other domains of learning, such as reinforcement learning. There has been some research in this area examining the role of instructional control in category learning (Noelle & Cottrell, 1996), probabilistic reward learning (Li, Delgado, & Phelps, 2011; Walsh & Anderson, 2011; Biele, Rieskamp, & Gonzalez, 2009; Doll, Jacobs, Sanfey, & Frank, 2009), and stimulus–response (S-R) learning (Cohen-Kadosh & Meiran, 2009; Wenke, Gaschler,

Nattkemper, & Frensch, 2009), instructional learning remains a poorly characterized phenomenon. Because different types of tasks likely engage different brain networks and mechanisms for encoding procedural knowledge in that particular domain, for simplicity we center our discussion on S-R instructions hereafter.

Multiple sources of neural data (e.g., neuroimaging and electrophysiology) show that implementation of novel S-R mappings involves several cortical and subcortical areas, such as premotor, prefrontal and posterior parietal cortices, hippocampus and striatum (Brass, Wenke, Spengler, & Waszak, 2009; Suzuki, 2007; Casey, Thomas, Davidson, Kunz, & Franzen, 2002). Although it remains unclear exactly how these brain areas coordinate to learn and implement instructions, the lateral pFC appears to be particularly important for instructional control (Hartstra, Kuhn, Verguts, & Brass, 2011; Ruge & Wolfensteller, 2010).

Several human and monkey studies on conditional motor learning (without the instruction component) provide insight into the functional roles of brain regions involved in processing S-R mappings. Whereas the hippocampal system underlies rapid acquisition of novel S-R associations (Brasted, Bussey, Murray, & Wise, 2005; Casey et al., 2002), the posterior parietal cortex subserves transformation of simple and well-learned S-R mappings (Grol, de Lange, Verstraten, Passingham, & Toni, 2006; Corbetta & Shulman, 2002). As for responses, motor preparation is carried out by the premotor cortex and other motor-related areas (Suzuki, 2007; Cavina-Pratesi et al., 2006), and the ACC is involved in monitoring response conflict (Brass et al., 2009; Botvinick, Cohen, & Carter, 2004).

On the basis of the neuroimaging and neurophysiological evidence mentioned above, a neural network model is developed to explore the dynamic interplay among various brain areas for instructional control. Mechanistically, the model explains how S-R instructions suppress habits for new conditional behavior and how automaticity can be achieved through motor consolidation. Next we outline the major model components and their connections in reference to relevant neuroanatomy.

Model Overview

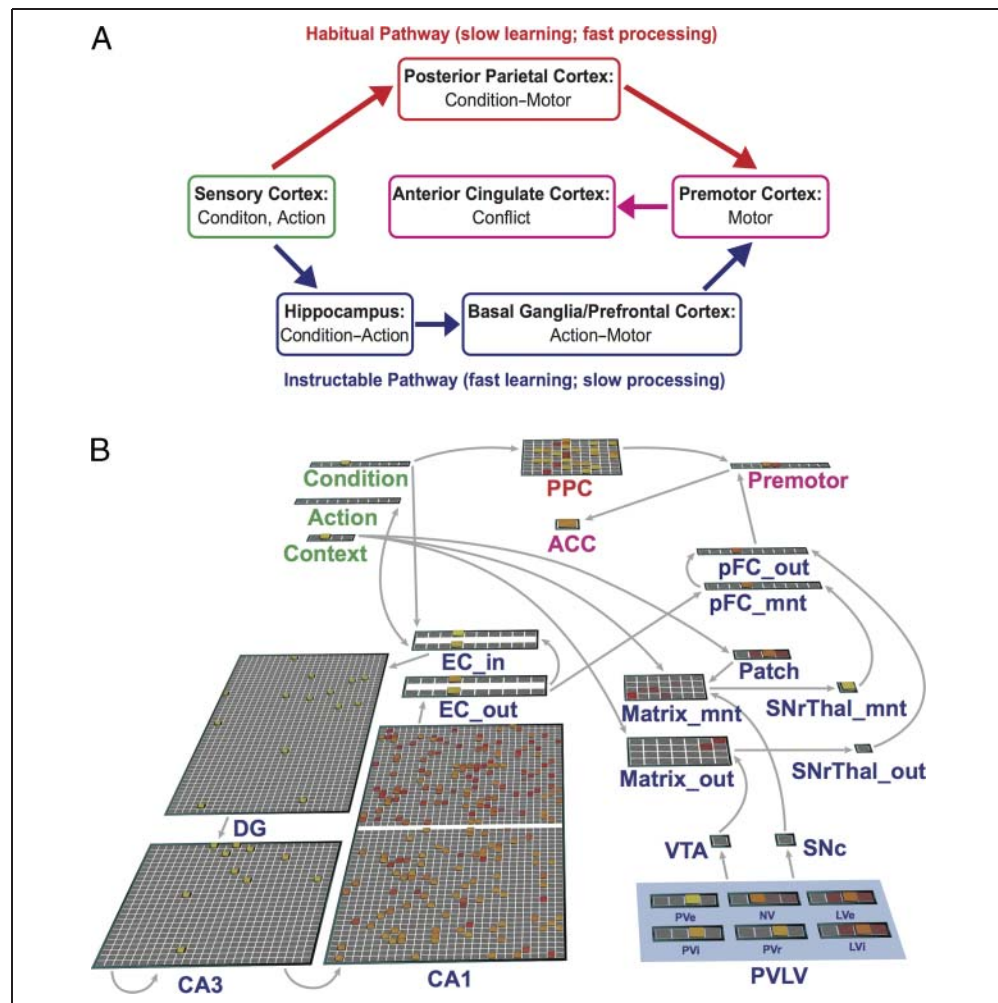
On a large scale, our model of instructional control consists of two complementary learning pathways, as shown in Figure 1. The fast-learning hippocampal–prefrontal pathway processes complex, novel S-R mappings like those learned from instructions, whereas the slow-learning parietal pathway processes simple, habitual S-R mappings. This overall architecture is consistent with findings about hippocampal and striatal contributions to initial learning and that neocortical learning is gradually shaped by experience (Pasupathy & Miller, 2005; for reviews, see Ashby, Turner, & Horvitz, 2010; McClelland, McNaughton,

& O'Reilly, 1995). Note, however, that with a focus on goal-directed control we model both the ventral and dorsomedial striatum (roughly, nucleus accumbens and caudate) but not dorsolateral striatum (putamen), which also governs habitual behavior (for reviews, see Ashby et al., 2010; Yin & Knowlton, 2006).

Our model helps to answer the following questions: Why is trial-and-error learning so arduous whereas instructed learning appears effortless? How do we successfully perform complex novel tasks on the first attempt? We reason that S-R instructions quickly assemble rather than slowly modify preexisting elements of perceptual and motor knowledge. For example, we can immediately follow the instruction: “press the left button when seeing a triangle; press the right button when seeing a square,” in which the action of button press is a preexisting motor skill and visual recognition of shapes is also an already learned perceptual ability. Note also that understanding the instruction requires a previously learned mapping from language (e.g., the verbal command of “press”) to actual behavior (e.g., the motor execution of “press”).

Our proposed model implements the aforementioned instructional control from neural to behavioral levels.

Figure 1. (A) The model macrocircuit. The hippocampal–prefrontal pathway processes newly instructed S-R rules, whereas the parietal pathway processes habitual S-R transformations. The motor signals in the premotor/motor cortices from the hippocampal–prefrontal pathway can train up the parietal pathway for automaticity. (B) The model microcircuit. The ACC becomes active whenever the parietal and hippocampal–prefrontal pathways suggest conflicting outputs in the premotor layer.



Unlike a single-purpose neural network that slowly re-wires the whole system to learn a new sensorimotor transformation, this general-purpose instructable model separates movement from plan representations and restricts plan updating to lie within the fast-learning hippocampus. Specifically, the pFC together with BG in the model learns a vocabulary of action items and the corresponding motor responses. The model hippocampus rapidly encodes S-R instructions as action episodes. It can retrieve a stimulus-dependent action command into the dorsolateral pFC as a goal for guiding subsequent behavior although sufficiently simple S-R rules, such as categorical rather than one-to-one mappings, may be directly maintained in the unmodeled ventrolateral pFC (Hartstra et al., 2011; Bunge, 2004) by verbal working memory instead of episodic memory. The model posterior parietal cortex alone can also carry out well-versed S-R transformations, and all the response plans outputted from the posterior parietal cortex (PPC) and pFC compete in the model premotor cortex. For all the simulated brain areas, below we discuss their anatomical connections and computational functions in the model.

Posterior Parietal Cortex

The PPC receives inputs from visual, auditory, and somatosensory cortices and outputs to motor areas, such as the premotor and primary motor cortices, consistent with the “how” or perception-for-action framework of Goodale and Milner (1992). After extensive learning, the PPC can process highly familiar S-R mappings in a fast, automatic, and unconscious manner (Rossetti et al., 2005; Schindler et al., 2004) and damage to the PPC in humans can produce a variety of sensorimotor deficits such as Apraxia (Andersen & Buneo, 2002). In the model, the PPC layer connects the input Condition layer with the output Premotor layer in the parietal learning pathway. It is essentially a hidden layer in a generic three-layer neural network that learns arbitrary input–output transformations.

Hippocampus

The entorhinal cortex provides the major interface for communication between the hippocampus and neocortex through perirhinal and parahippocampal cortices (Lavenex & Amaral, 2000). Within the hippocampus, there are two major pathways. The perforant pathway, originated from layer II of the entorhinal cortex, links dentate gyrus, Cornu Ammonis fields CA3, CA1, and subiculum by a long neuronal chain; the direct pathway projects from layer III of the entorhinal cortex directly to CA1, which in turn projects back to layer V of the entorhinal cortex via subiculum (Duvernoy, 2005).

Depending on the phase of the hippocampal theta rhythm, CA1 is driven mainly by either entorhinal or CA3 inputs for memory encoding and retrieval, respectively (Hasselmo, Bodelon, & Wyble, 2002). For goal-directed

decisions, memory retrieved in the hippocampus can be routed to the pFC via subicular, entorhinal, perirhinal, and parahippocampal cortices (Simons & Spiers, 2003; Goldman-Rakic, Selemon, & Schwartz, 1984).

In the model, we simulate both the perforant and direct pathways except subiculum, whose functional role in novelty processing (Lisman & Grace, 2005) is beyond the scope of this article. Computationally, the perforant pathway up to CA3 carries out pattern separation and completion (O’Reilly & McClelland, 1994) such that the model CA3 encodes sparse, conjunctive, and content-addressable memory representations of sensory inputs to the hippocampus, such as verbally instructed S-R associations; the direct pathway is a three-layer autoencoder network, which learns to perform a topographically organized identity mapping from the model EC_in (i.e., entorhinal superficial layers) via CA1 to EC_out (i.e., entorhinal deep layers) such that a recalled episode in CA3 can be reinstated into EC_out via the CA3-to-CA1 projection (Norman & O’Reilly, 2003).

Prefrontal Cortex and Basal Ganglia

Although the pFC comprises cytoarchitecturally distinct areas, as a whole it extensively interconnects with sensory, motor, and limbic systems. In particular, the dorsolateral pFC, implicated in working memory function, receives inputs from visual/auditory/somatosensory cortices, hippocampus (mainly via the orbitomedial pFC), and BG (via the thalamus) and outputs to motor areas such as the SMA, FEFs, and premotor cortex (Miller & Cohen, 2001).

In the model, we implement a pFC BG working memory (PBWM) system (Hazy, Frank, & O’Reilly, 2006, 2007; O’Reilly & Frank, 2006) to hold instructed actions recalled from the hippocampus. Specifically, the model pFC has two separate layers, pFC_mnt and pFC_out, to simulate neurons showing either tonic maintenance or phasic motor-related activities. The simplified model BG has distinct go and no-go neurons in every striatal matrix layer for deciding whether to maintain and/or output working memory contents. The pathway from Matrix_mnt to SNrThal_mnt controls pFC_mnt firing, whereas the pathway from Matrix_out to SNrThal_out controls pFC_out firing. The action units of the model EC_out layer projects to pFC_mnt, which in turn projects to pFC_out.

Dopamine (DA) neurons signal reward prediction errors, which allows learning of go versus no-go decisions of BG-gated working memory operations, such as outputting working memory contents to Premotor during performance but not during instruction memorization. Because appropriate go/no-go behaviors depend on the environmental context, different learning stages (see Methods for more details) are coded with localist representations in the model Context layer as a sensory, contextual input to the simulated matrixes of the striatum, namely Matrix_mnt and Matrix_out.

The reward-predictive firing properties of DA neurons in the substantia nigra pars compacta and ventral tegmental

area are simulated using a circuit that explains data on Pavlovian conditioning (Hazy, Frank, & O'Reilly, 2010; O'Reilly, Frank, Hazy, & Watz, 2007). Specifically, the simulated DA neurons increase firing for unexpected rewards and decrease firing for omission of expected rewards to respectively drive the go and no-go pathway neurons in the striatum. In the model simulation, the amount of reward supplied to this DA system is derived from the correctness of final motor responses.

Anterior Cingulate Cortex

The dorsal ACC interconnects with the PPC, dorsolateral pFC, and motor structures including premotor, supplementary motor, and primary motor areas. Thus, the dorsal ACC is in a position to evaluate or regulate executive control during both stimulus presentation and response selection, although it is engaged most strongly in monitoring and evaluating the outcomes of actions (Botvinick et al., 2004; Paus, 2001).

In the model, ACC unit receives uniform excitatory projections from all the Premotor units and has a high firing threshold such that it does not fire unless receiving inputs from more than two active units in the Premotor layer. In other words, the model ACC monitors existence of conflicting motor plans whereas the Premotor layer resolves conflicts, if any, during response selection through winner-take-all competition. Our simplified ACC does not project to any other model layers to further modulate

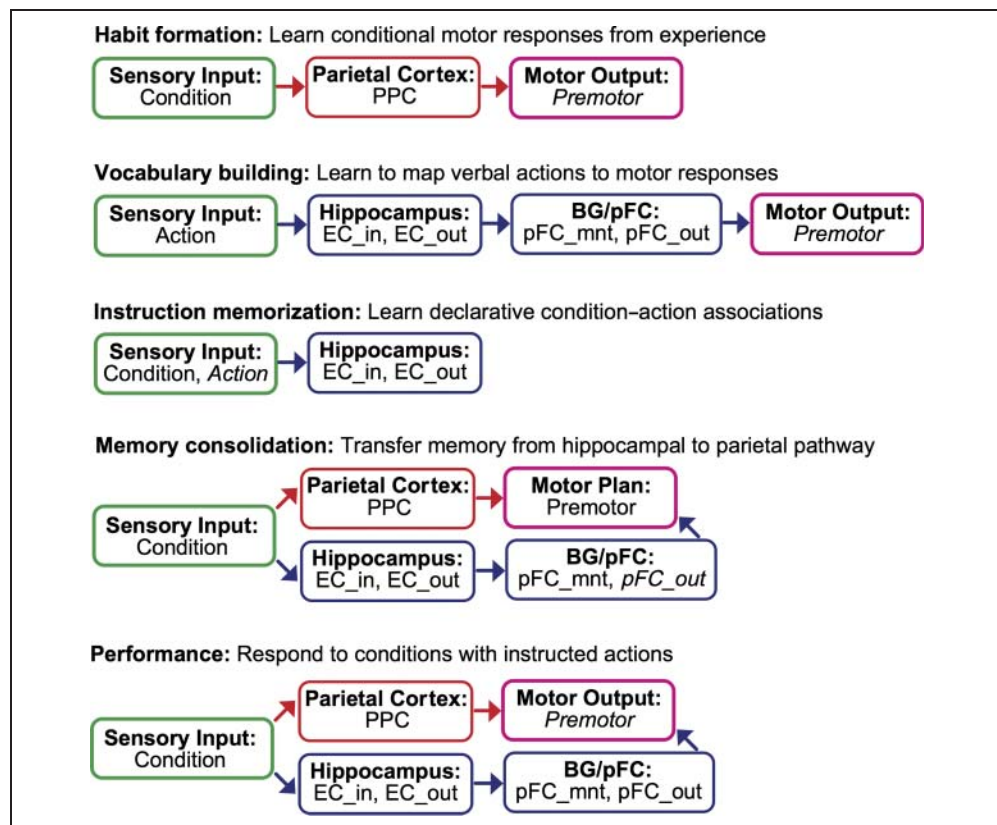
network dynamics, such as driving the basal forebrain cholinergic neurons to increase attention for more deliberate task processing (Krichmar, 2008).

METHODS

Our model was implemented using the Leabra framework for simulating biologically realistic neural networks at an intermediate level of detail. Leabra has been applied to explain a wide range of cognitive phenomena and their underlying neural mechanisms (O'Reilly, Hazy, & Herd, in press; O'Reilly & Munakata, 2000). Leabra combines locally computable mechanisms of error-driven and Hebbian associative learning to offer a biologically realistic way of training a hierarchical neural network in a supervised manner, which essentially modifies synaptic connectivity in a network to minimize any discrepancies between expectations (i.e., spontaneous network outputs in the Leabra minus phase) and actual outcomes (i.e., supplied teaching signals in the Leabra plus phase) in the network output layers.

Because conditional responses could be learned from instruction or from experience, we distinguished verbal action commands ("Action") from motor responses ("Motor") and supplied Condition–Motor, Action–Motor, and Condition–Action pairs as input–output teaching signals to respectively train the model PPC, BG/pFC, and hippocampus when simulating different types of learning (see the first three cases in Figure 2). Model performance

Figure 2. Five different simulation stages. Italicized labels denote the output layers that receive/provide teaching signals.



was evaluated based on these different teaching signals accordingly (see the italicized labels in Figure 2). For example, during instruction memorization, it was considered an error when the hippocampus failed to output the designated “Action” in response to a specific “Condition” input in a trial. Error-driven, supervised learning was then continued for each type of learning until no error was made for several epochs, where an epoch comprised 10 trials of different input–output pairs. For cortical consolidation, no external teaching signals were provided to the model. Instead, the hippocampal–prefrontal outputs served as internal teaching signals to correct spontaneous outputs from the model parietal cortex.

The proposed model used a rate-based coding scheme with a hybrid of localist and distributed representations to process information. Here the neural firing rates in the model were computed from excitatory conductances rather than membrane potentials to better capture the rates of neuronal spiking (O’Reilly et al., in press). Localist/sparse and distributed codes refer to whether certain information is represented by a relatively small or large number of neurons, respectively. Because we abstracted away sensory and motor processes during instructional control, it was difficult to manually specify a realistic pattern of distributed neural activities in the model input/output layers and thus localist representations were used. Otherwise, distributed coding was the default representation throughout the network.

For simplicity, the model input/output layers communicated with each other through one-to-one connections between the same localist representations. These interface layers, listed in Figure 2, included the inputs/outputs of the whole system (i.e., Condition, Action, and Premotor), the hippocampus (i.e., EC_in and EC_out), and the PBWM system (i.e., pFC_mnt and pFC_out). For example, one-to-one connections were set between Condition and the first unit group of EC_in, between Action and the second unit group of EC_in, between the second unit group of EC_out and pFC_mnt, between pFC_mnt and pFC_out, and between pFC_out and Premotor (see also Figure 1B).

As a whole, the proposed model is a synthesis of several specialized neural systems. Different brain regions simulated in the model share the same set of computational principles (e.g., lateral inhibition, bidirectional connectivity, and local learning) but become functionally specialized because of different parameterizations (e.g., differences in sparseness of neural codes, learning rate, and balance between error-driven and associative learning). The present work adopts the default structures and parameter values from previous studies that detail these specialized parameterizations for the neocortex (O’Reilly et al., in press), hippocampus (Norman & O’Reilly, 2003), mid-brain DA system (Hazy et al., 2010), and prefrontal-BG working memory (O’Reilly & Frank, 2006). For technical reference, the source code of our model and simulations are available for download at grey.colorado.edu/~tren/instruct/.

RESULTS

The model simulation results are presented below. Each simulation result shows the mean and the standard error of the mean of 20 independent model runs, each of which used randomly generated S-R mappings and initial weights for the network connections. In all the simulations, each of the Condition, Action, and Premotor layers used 10 localist units to respectively represent 10 encountered conditions, 10 verbal actions, and 10 motor outputs (see Methods for more details).

Each simulation ran through a specific set of stages that are described in Figure 2. Three of the stages, however, were common to all the simulations: The model was pre-trained with Action-to-Motor mappings (i.e., from verbal commands to real motor responses) during the vocabulary-building stage and then trained with Condition-to-Action mappings (i.e., S-R rules) during the instruction memorization stage. During the performance stage, it was tested with Condition-to-Motor mappings without any inputs from the Action layer.

Simulation 1: Instruction Following

To simulate basic instruction following, the proposed model was instructed with 10 novel pairs of S-R rules (e.g., if see S, then do R) and evaluated for its success in performing conditional actions (e.g., do R) when encountering a specific condition (e.g., see S). As shown in Figure 3A, the model quickly memorized an S-R rule in few trials during the instruction memorization stage, and without further practice it made no error in carrying out these instructions for response during the performance stage. Such a perfect performance was due to the contribution from the hippocampal–prefrontal learning pathway because lesioning the model posterior parietal cortex did not impair instruction following.

Simulation 2: Habit Suppression

S-R instructions can sometimes conflict with previously developed habits, as in S-R compatibility tasks. In such a case, newly instructed S-R associations stored in the hippocampus need to override prepotent responses promoted by old S-R habits (Casey et al., 2002). Although behaviorally people can make few or no error in following instructions, neurally they must resolve competition between new and old S-R mappings during response selection, as signaled by the dorsal ACC (Botvinick et al., 2004).

To simulate habit suppression, we first trained up the parietal habit pathway on 10 S-R rules before instructing the model on 10 newly remapped S-R rules, which recycled the same set of 10 stimuli and 10 responses in the old rules. Independently training the PPC and hippocampus was achieved by providing teaching signals to the model layers that interfaced with the parietal or

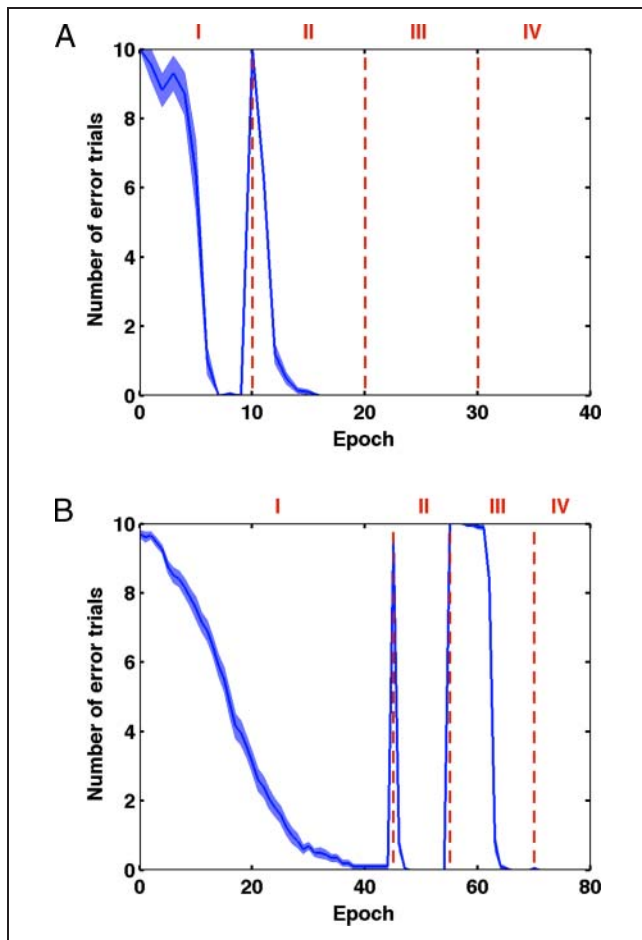


Figure 3. (A) Simulation of instruction-following. Stage I: Action-to-Motor vocabulary building; Stage II: Condition-to-Action instruction memorization; Stage III: Condition-to-Motor performance; Stage IV: Condition-to-Motor performance after PPC was lesioned. (B) Simulation of habits being temporarily suppressed by instructions. Stage I: Condition-to-Motor habit formation; Stage II: Action-to-Motor vocabulary building; Stage III: Condition-to-Action instruction memorization; Stage IV: Condition-to-Motor performance on the newly instructed actions.

hippocampal outputs: Premotor was an output layer and Action provided null inputs during the parietal training, whereas Action was an output layer and Premotor provided null inputs during the hippocampal training (see Figure 2). Such a treatment ensured proper credit assignment during error-driven learning.

As shown in Figure 3B, during the performance stage, the model successfully employed the newly instructed S-R rules without making any errors because the stronger outputs from the hippocampal–prefrontal pathways outweighed the weaker outputs from the parietal pathway in the Premotor layer, which resolved the competition between these two conflicting motor plans over time using a winner-take-all mechanism. During such a response selection process, the model ACC detected incompatible motor plans in the Premotor layer despite the model’s errorless performance (Figure 1B).

Simulation 3: Automaticity

Damage to the hippocampus in primates can lead to anterograde amnesia and temporally graded retrograde amnesia although novel S-R mappings can still be slowly acquired by the habit learning system in this circumstance (Bayley, Frascino, & Squire, 2005; Wise & Murray, 1999). It is thus believed that novel declarative memory formed in the fast-learning hippocampus is later consolidated in the slow-learning neocortex (McClelland et al., 1995; Alvarez & Squire, 1994). Although the BG is suggested to train up the neocortex during perceptual and motor consolidation (Ashby et al., 2010), we propose that the posterior hippocampus and the dorsolateral striatal circuit provide teaching signals of declarative and non-declarative S-R rules, respectively, to the neocortex for developing automaticity.

In the model, the hippocampal–prefrontal pathway is fast at learning but slow at processing, whereas the parietal pathway is slow at learning but fast at processing. Such

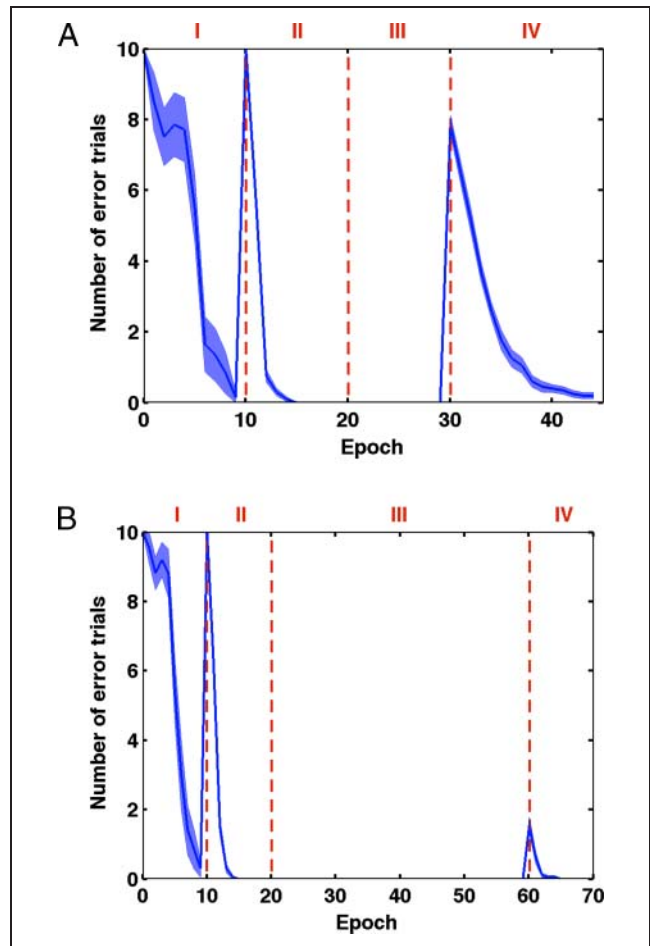


Figure 4. Simulation of cortical consolidation for automaticity. Stage I: Action-to-Motor vocabulary building; Stage II: Condition-to-Action instruction memorization; Stage III: Spontaneous hippocampal–prefrontal recall of Condition-to-Motor mappings (i.e., consolidation); Stage IV: Condition-to-Motor performance after the hippocampus and pFC were lesioned. (A) A shorter Stage III led to more errors in Stage IV. (B) A longer Stage III led to fewer errors in Stage IV.

a paradoxical property is because of a larger number of synaptic connections and more complex machinery involved in the hippocampal–prefrontal than the parietal pathway. As a result, when no teaching signal is supplied from the external environment to the premotor/motor cortices, a deliberately constructed motor plan that is transmitted from the hippocampal–prefrontal pathway to the model premotor cortex can serve as an internal teaching signal to the faster-responding parietal pathway to enable error-driven and/or supervised Hebbian learning. Gradually, the parietal pathway learns to conform its answers to those from the hippocampal–prefrontal pathway during this cortical consolidation process, and automaticity purely through parietal mediation can be achieved in the long run.

To simulate development of automaticity, a memory consolidation stage was introduced before the regular performance stage. During the consolidation stage, no explicit teaching signals were provided from the simulation environment to the model Premotor layer, and thus, the model performance was not evaluated. During the performance stage, all the model prefrontal and hippocampal layers were lesioned to see the pure parietal contributions to the model outputs. As shown in Figure 4, the degree of automaticity in terms of performance accuracy was positively correlated with the duration of consolidation.

DISCUSSION

We have presented a neural model capable of instructional control based on the relevant brain areas known to be involved. The model treats an instructed S-R task as a combinatorial generalization task in that novel combination of perceptual and motor knowledge is the primary learning problem faced by the system. This is a realistic simplification, in that people can only be instructed to perform tasks where they already know the basic task elements, and the instructions consist of a particular combination of these elements to be performed. Our model shows that these combinations can be easily memorized by the hippocampus and then implemented through top–down cognitive control via the pFC/BG system. This form of cognitive learning can be more rapid and flexible than perceptual-motor learning, which engages the slow-learning neocortex to develop task-specific representations.

Learning occurs in multiple parts of the architecture to support the ability to behave according to instructions. To comprehend the language of action instructions, the model needs a vocabulary-building stage during which the hippocampus learns to perform identity mapping for relaying information from the Action layer to the corresponding motor representations in pFC layers. Meanwhile, BG learns to open the execution gate for pFC to output a motor decision to the Premotor layer. During the instruction memorization stage, the hippocampus associates inputs from the Condition and Action layers and learns

each condition–action pair as an episodic pattern. Only after completing all these various forms of learning can the model follow instructions flawlessly during the performance stage—using mechanisms of pattern completion, the hippocampus recalls instructions about what action to perform based on retrieval cues from the Condition layer, and its downstream pFC either maintains a retrieved motor command in working memory when BG closes the execution gate or further triggers a motor decision in the Premotor layer when BG opens the execution gate.

Compared with reinforcement learning, instructed learning has several advantages. Reinforcement learning adapts behavior based on the consequences of actions, whereas instructed learning adapts behavior in accordance with instructed action rules. Hence, unlike the slow, retrospective process of trial and error in reinforcement learning, instructed learning tends to be fast, proactive, and errorless. In general, trial-and-error learning increases the likelihood of making the same mistakes in the future, and errorless learning avoids such a possibility of strengthening incorrect S-R associations through Hebbian mechanisms (McClelland, 2001). Note, however, that although errorless model performance is mainly because of perfect recalls of instructions from the hippocampus, we do not deny hippocampal involvement in trial-and-error learning, especially when task rules are explicit or deterministic. In fact, S-R studies that used trial-and-error procedures found hippocampal contributions to one-trial learning of correct responses (e.g., Brasted et al., 2005).

Our model differs from others in level of detail and in theorizing about how the pFC, BG, and the hippocampus interact with each other. For example, connectionist approaches (e.g., Noelle & Cottrell, 1995, 1996) and various reinforcement learning algorithms have been used to imply how instructions may interact with the experience system subserved by the BG (Li et al., 2011; Walsh & Anderson, 2011; Biele et al., 2009; Doll et al., 2009). However, the prefrontal and hippocampal mechanisms that mediate instructions remain underspecified in these studies compared with ours. Also, hippocampal memory and PBWM, respectively, underlies rule-based and habitual behavior in the bilinearity model (Dayan, 2007), but they work in tandem as the instructable pathway in our model where habitual S-R responses engage little or no prefrontal control.

Our model explains existing data and makes predictions. Note that postresponse feedback that defines correct S-R associations is functionally equivalent to prerresponse instructions, our model would predict impairment of learning and executing novel conditional responses in both the cases of S-R specification when either the pFC or hippocampus is damaged in humans, as shown by Petrides (1985, 1997). Moreover, whereas the BG plays a major role in habit learning and automaticity (Ashby et al., 2010; Yin & Knowlton, 2006), our model predicts hippocampal and prefrontal involvement in cortical consolidation of instructed one-to-one S-R mappings. Furthermore, because

the parietal pathway mainly supports habitual S-R execution, our model predicts that applying TMS onto the PPC will not disrupt learning and executing of instructed S-R mappings.

Finally, the model proposed here is being extended to address other important questions about instructional learning and control. For example, how does the brain handle hierarchical, probabilistic, or sequence rules other than simple S-R mappings? How does misleading instructions affect learning and performance of a given task? These more complex scenarios engage a larger brain network than the present work and require more detailed modeling of the pFC (Cole, Bagic, Kass, & Schneider, 2010; Bunge, 2004) as well as the BG and motor cortices.

Acknowledgments

We thank members of the CCN Lab at CU Boulder for helpful comments and discussion. This work was supported by ONR N00014-07-1-0651, ARL RCTA, NIH MH079485 and by the Intelligence Advanced Research Projects Activity (IARPA) via Department of the Interior (DOI) contract number D10PC20021. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI, or the U.S. Government.

Reprint requests should be sent to Tsung-Ren Huang, Department of Psychology and Neuroscience, University of Colorado Boulder, 345 UCB, Boulder, CO 80309, or via e-mail: tsungren.huang@colorado.edu.

REFERENCES

- Alvarez, P., & Squire, L. R. (1994). Memory consolidation and the medial temporal lobe: A simple network model. *Proceedings of the National Academy of Sciences, U.S.A.*, *91*, 7041–7045.
- Andersen, R. A., & Buneo, C. A. (2002). Intentional maps in posterior parietal cortex. *Annual Review of Neuroscience*, *25*, 189–220.
- Ashby, F. G., Turner, B. O., & Horvitz, J. C. (2010). Cortical and basal ganglia contributions to habit learning and automaticity. *Trends in Cognitive Science*, *14*, 191–232.
- Bayley, P. J., Frascino, J. C., & Squire, L. R. (2005). Robust habit learning in the absence of awareness and independent of the medial temporal lobe. *Nature*, *436*, 550–553.
- Biele, G., Rieskamp, J., & Gonzalez, R. (2009). Computational models for the combination of advice and individual learning. *Cognitive Science*, *33*, 206–242.
- Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: An update. *Trends in Cognitive Sciences*, *8*, 539–546.
- Brass, M., Wenke, D., Spengler, S., & Waszak, F. (2009). Neural correlates of overcoming interference from instructed and implemented stimulus–response associations. *Journal of Neuroscience*, *29*, 1766–1772.
- Brasted, P. J., Bussey, T. J., Murray, E. A., & Wise, S. P. (2005). Conditional motor learning in the nonspatial domain: Effects of errorless learning and the contribution of the fornix to one-trial learning. *Behavioral Neuroscience*, *119*, 662–676.
- Bunge, S. A. (2004). How we use rules to select actions: A review of evidence from cognitive neuroscience. *Cognitive, Affective, and Behavioral Neuroscience*, *4*, 564–579.
- Casey, B. J., Thomas, K. M., Davidson, M. C., Kunz, K., & Franzen, P. L. (2002). Dissociating striatal and hippocampal function developmentally with a stimulus–response compatibility task. *Journal of Neuroscience*, *22*, 8647–8652.
- Cavina-Pratesi, C., Valyear, K. F., Culham, J. C., Köhler, S., Obhi, S. S., Marzi, C. A., et al. (2006). Dissociating arbitrary stimulus–response mapping from movement planning during preparatory period: Evidence from event-related functional magnetic resonance imaging. *Journal of Neuroscience*, *26*, 2704–2713.
- Cohen-Kdoshay, O., & Meiran, N. (2009). The representation of instructions operates like a prepared reflex: Flanker compatibility effects that are found in the first trial following S-R instructions. *Experimental Psychology*, *56*, 128–133.
- Cole, M. W., Bagic, A., Kass, R., & Schneider, W. (2010). Prefrontal dynamics underlying rapid instructed task learning reverse with practice. *Journal of Neuroscience*, *30*, 14245–14254.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201–215.
- Dayan, P. (2007). Bilinearity, rules and prefrontal cortex. *Frontiers in Computational Neuroscience*, *1*, 1.
- Doll, B. B., Jacobs, W. J., Sanfey, A. G., & Frank, M. J. (2009). Instructional control of reinforcement learning: A behavioral and neurocomputational investigation. *Brain Research*, *1299*, 74–94.
- Duvernoy, H. M. (2005). *The human hippocampus: Functional anatomy, vascularization and serial sections with MRI* (3rd ed.). Berlin: Springer-Verlag.
- Goldman-Rakic, P. S., Selemon, L. D., & Schwartz, M. L. (1984). Dual pathways connecting the dorsolateral prefrontal cortex with the hippocampal formation and parahippocampal cortex in the rhesus monkey. *Neuroscience*, *12*, 719–743.
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neuroscience*, *15*, 20–25.
- Grol, M. J., de Lange, F. P., Verstraten, F. A. J., Passingham, R. E., & Toni, I. (2006). Cerebral changes during performance of overlearned arbitrary visuomotor associations. *Journal of Neuroscience*, *26*, 117–125.
- Hartstra, E., Kuhn, S., Verguts, T., & Brass, M. (2011). The implementation of verbal instructions: An fMRI study. *Human Brain Mapping*, *32*, 1811–1824.
- Hasselmo, M. E., Bodelon, C., & Wyble, B. P. (2002). A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation*, *14*, 793–817.
- Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2006). Banishing the homunculus: Making working memory work. *Neuroscience*, *139*, 105–118.
- Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2007). Towards an executive without a homunculus: Computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *362*, 1601–1613.
- Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2010). Neural mechanisms of acquired phasic dopamine responses in learning. *Neuroscience and Biobehavioral Reviews*, *34*, 701–720.
- Krichmar, J. L. (2008). The neuromodulatory system—A framework for survival and adaptive behavior in a challenging world. *Adaptive Behavior*, *16*, 385–399.

- Lavenex, P., & Amaral, D. G. (2000). Hippocampal-neocortical interaction: A hierarchy of associativity. *Hippocampus*, *10*, 420–430.
- Li, J., Delgado, M. R., & Phelps, E. A. (2011). How instructed knowledge modulates the neural systems of reward learning. *Proceedings of the National Academy of Sciences, U.S.A.*, *108*, 55–60.
- Lisman, J., & Grace, A. A. (2005). The hippocampal-VTA loop: Control of entry into long-term memory. *Neuron*, *46*, 703–713.
- McClelland, J. L. (2001). Failures to learn and their remediation: A Hebbian account. In J. L. McClelland & S. Siegler (Eds.), *Mechanisms of cognitive development: Behavioral and neural approaches* (pp. 97–121). Mahwah, NJ: Erlbaum.
- 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.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202.
- Noelle, D. C., & Cottrell, G. W. (1995). A connectionist model of instruction following. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 369–374). Hillsdale, NJ: Erlbaum.
- Noelle, D. C., & Cottrell, G. W. (1996). Modeling interference effects in instructed category learning. In G. W. Cottrell (Ed.), *Proceedings of the 18th Annual Conference of the Cognitive Science Society* (pp. 475–480). Hillsdale, NJ: Erlbaum.
- Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary learning systems approach. *Psychological Review*, *110*, 611–646.
- 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.
- O'Reilly, R. C., Frank, M. J., Hazy, T. E., & Watz, B. (2007). PVLV: The primary value and learned value Pavlovian learning algorithm. *Behavioral Neuroscience*, *121*, 31–49.
- O'Reilly, R. C., Hazy, T. E., & Herd, S. A. (in press). The Leabra cognitive architecture: How to play 20 principles with nature and win! *The Oxford Handbook of Cognitive Science*. Oxford: Oxford University Press.
- O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a tradeoff. *Hippocampus*, *4*, 661–682.
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.
- Pasupathy, A., & Miller, E. K. (2005). Different time courses for learning-related activity in the prefrontal cortex and striatum. *Nature*, *433*, 873–876.
- Paus, T. (2001). Primate anterior cingulate cortex: Where motor control, drive and cognition interface. *Nature Reviews Neuroscience*, *2*, 417–424.
- Petrides, M. (1985). Deficits on conditional associative-learning task after frontal- and temporal-lobe lesions in man. *Neuropsychologia*, *23*, 601–614.
- Petrides, M. (1997). Visuo-motor conditional associative learning after frontal and temporal lesions in the human brain. *Neuropsychologia*, *35*, 989–997.
- Rossetti, Y., Revol, P., McIntosh, R., Pisella, L., Rode, G., Danckert, J., et al. (2005). Visually guided reaching: Posterior parietal lesions cause a switch from fast visuomotor to slow cognitive control. *Neuropsychologia*, *43*, 162–177.
- Ruge, H., & Wolfensteller, U. (2010). Rapid formation of pragmatic rule representations in the human brain during instruction-based learning. *Cerebral Cortex*, *20*, 1656–1667.
- Schindler, I., Rice, N. J., McIntosh, R. D., Rossetti, Y., Vighetto, A., & Milner, A. D. (2004). Automatic avoidance of obstacles is a dorsal stream function: Evidence from optic ataxia. *Nature Neuroscience*, *7*, 779–784.
- Simons, J. S., & Spiers, H. (2003). Prefrontal and medial temporal lobe interactions in long-term memory. *Nature Reviews Neuroscience*, *4*, 637–648.
- Suzuki, W. A. (2007). Integrating associative learning signals across the brain. *Hippocampus*, *17*, 842–850.
- Walsh, M. M., & Anderson, J. R. (2011). Modulation of the feedback-related negativity by instruction and experience. *Proceedings of the National Academy of Sciences, U.S.A.*, *108*, 19048–19053.
- Wenke, D., Gaschler, R., Nattkemper, D., & Frensch, P. A. (2009). Strategic influences on implementing instructions for future actions. *Psychological Research*, *73*, 587–601.
- Wise, S. P., & Murray, E. A. (1999). Role of the hippocampal system in conditional motor learning: Mapping antecedents to action. *Hippocampus*, *9*, 101–117.
- Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience*, *7*, 464–476.