CellPress

Brain Networks for Exploration Decisions Utilizing Distinct Modeled Information Types during Contextual Learning

Jane X. Wang[1,*] and Joel L. Voss[1]
[1]Department of Medical Social Sciences, Ken & Ruth Davee Department of Neurology, and Interdepartmental Neuroscience Program, Feinberg School of Medicine, Northwestern University, Chicago, IL 60611, USA
*Correspondence: janewang@northwestern.edu
http://dx.doi.org/10.1016/j.neuron.2014.04.028

SUMMARY

Exploration permits acquisition of the most relevant information during learning. However, the specific information needed, the influences of this information on decision making, and the relevant neural mechanisms remain poorly understood. We modeled distinct information types available during contextual association learning and used model-based fMRI in conjunction with manipulation of exploratory decision making to identify neural activity associated with information-based decisions. We identified hippocampal-prefrontal contributions to advantageous decisions based on immediately available novel information, distinct from striatal contributions to advantageous decisions based on the sum total available (accumulated) information. Furthermore, network-level interactions among these regions during exploratory decision making were related to learning success. These findings link strategic exploration decisions during learning to quantifiable information and advance understanding of adaptive behavior by identifying the distinct and interactive nature of brain-network contributions to decisions based on distinct information types.

INTRODUCTION

Exploration behaviors during learning critically determine the information that is available and can be used to strategically acquire specific information needed to fill gaps in our memory/knowledge (Metcalfe and Jacobs, 2010). Exploration can thus determine what is learned, and learned information can, in turn, determine what will be explored. However crucial these mutual exploration-learning interactions are for memory success, little is known regarding their dynamics or neural mechanisms in humans.

Nonhuman animals can explore adaptively to improve learning. For instance, rodents sporadically exhibit iterative viewing of options at decision points during maze learning. This exploration pattern predicts learning success and effective generalization when the maze is subsequently altered (Tolman, 1948) and has been associated with hippocampal function (Buckner, 2010; Johnson and Redish, 2007). We have identified hippocampal-centered brain networks in humans associated with exploration behaviors that enhance learning, relative to receipt of the same stimuli but without active exploration (Voss et al., 2011a, 2011b). It is interesting that a specific exploration pattern that enhanced learning and hippocampal-prefrontal engagement was the revisiting of recently seen objects (Voss et al., 2011b), similar to the strategic exploration pattern observed in rodent maze learning. These findings implicate hippocampus and prefrontal cortex in online control of exploration (Buckner, 2010; Eichenbaum and Fortin, 2009; Wang et al., 2014), which could extend current functional accounts of these structures in advantageous decisions based on long-term memory (Buckner and Carroll, 2007; Schacter et al., 2012). In parallel research, dopamine-modulated pathways centered on the basal ganglia have been associated with strategic exploration during reinforcement learning and reward seeking (Hills, 2006; Pennartz et al., 2009), which could interact with hippocampus to support joint memory-reward influences on exploration (Shohamy and Adcock, 2010). However, further specification of the unique and interactive roles of hippocampus, prefrontal cortex, and basal ganglia in exploration will require measurement of the information that must be learned, so that the exploration decisions made to acquire this information can be isolated.

Indeed, it is an exceptional challenge to quantify the information on which individuals base exploration decisions during learning. Although it is possible to measure visual information for many stimuli (Beard and Ahumada, 1998), including entropy information relevant to novelty (Strange et al., 2005), this information does not necessarily drive exploration decisions. For instance, episodic learning is critically dependent on conceptual, gist, contextual, and other information types that are difficult to quantify. Moreover, current decision-making models, such as those for reinforcement learning, capitalize on the strong influence of reward on behavior to estimate internal decision variables (Frank and Claus, 2006), and in doing so conflate information available in the environment, information that is actually learned, and putative decision-making processes. Because available information cannot be isolated by these models (and, likewise, for many models of perceptual decisions), they do not permit isolation of the exploration decisions used to selectively acquire this information. Furthermore, existing decision-making

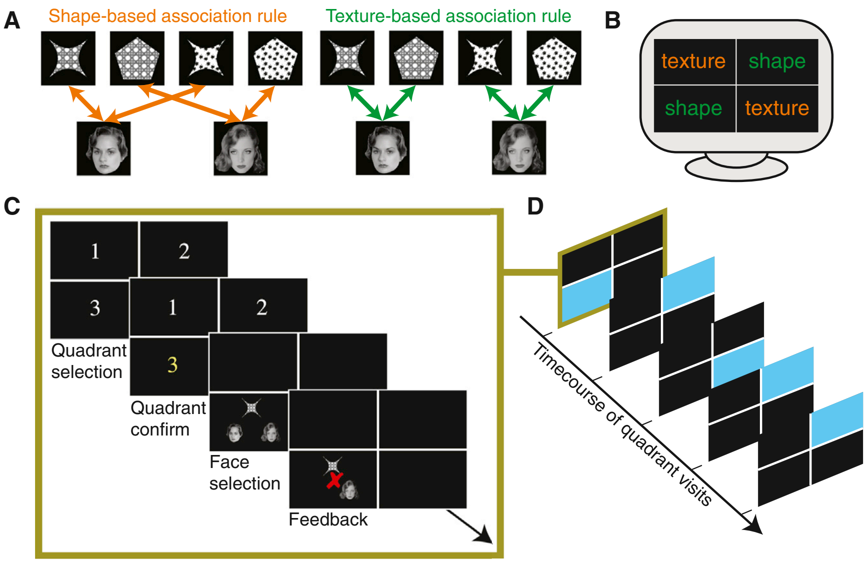

Figure 1. Contextual Object-Face Association Task

(A) Contextual associations were based on either shape or texture features of objects that served as cues. In shape quadrants, only shape (e.g., star-shaped versus pentagon-shaped) determined the correct object-face associations. In texture quadrants, only texture (e.g., white circles versus black dots) determined the correct object-face associations.

(B) Example configuration of quadrants, which varied for different blocks of the experiment, with two shape and two texture quadrants in each block. Subjects were not instructed regarding the salient feature in each quadrant but were required to learn contextual associations via feedback.

(C) Each trial involved highlight of the selected quadrant followed by presentation of the object cue and two faces, during which subjects attempted to select the target face. Trials concluded with feedback.

(D) Example quadrant sequence, with the quadrant selected for each trial highlighted in blue.

models generally account for learning of single parameters such as reward likelihood or perceptual identity (Ding and Gold, 2013). In contrast, episodic learning can require the integration of multiple information types over time (objects sampled within scenes, associations among sequentially presented items, etc.), thereby increasing the uncertainty of directly modeling decision-related variables.

To overcome these challenges, we adopted a blended modeling and experimental approach, whereby we modeled the information available during episodic learning and manipulated the ability to control exploration in order to isolate decisions based on modeled information. A contextual-association learning task required exploration of different contexts to identify contextual rules for item-item associations (similar to Badre et al., 2009). This allowed us to quantify contextual association information relevant for learning, based on extensions of optimal foraging theory that consider information as a finite resource that requires sampling (Hills, 2006; Pirolli and Card, 1999). Using a simple model with minimal assumptions, we quantified two aspects of information conceptualized as having distinct influences on learning and exploration (Frank et al., 2001; Johnson et al., 2012): (1) newly available information (NAI), which is the increase in available information provided when an event provides new information regarding contextual associations, and (2) accumulated available information (AAI), which is the total information previously encountered during exploration measured at any moment. To isolate exploration decisions, we manipulated the ability to actively explore using a condition in which subjects could control exploration (Active Learning) versus a condition in which the same information was passively studied (Passive Learning, as in Voss et al., 2011a, 2011b). This allowed us to isolate behavioral and neural correlates of exploration decisions based on modeled NAI and AAI using model-based fMRI in conjunction with comparisons between Active and Passive conditions.

We reasoned that neural activity associated with Active decisions based on NAI (relative to Passive exposure to NAI) would implicate regions in exploration decision making based on information that is immediately novel. Although prevailing accounts of hippocampal and prefrontal contributions to adaptive behavior emphasize long-term memory (Buckner and Carroll, 2007; Schacter et al., 2012), we found hippocampal and prefrontal involvement in NAI-based decisions, reflecting their role in the immediate use of novel information to support exploration decisions. In contrast, we identified regions of dorsal striatum associated with Active decisions based on AAI. This implicates dorsal striatum in exploration decisions based on accumulated information, substantiating theorized roles in strategic behavioral planning (Alexander et al., 1986; Martin, 1996) beyond involvement in slow learning of predictable stimulus-response associations (Packard and Knowlton, 2002). Finally, measures of background connectivity (Norman-Haignere et al., 2012) were analyzed to test putative network-level interactivity among these AAI-related and NAI-related regions in relation to advantageous exploration decisions. We found that greater interactivity predicted superior learning, indicating an important role for interplay of AAI- and NAI-related processing for advantageous exploration decisions.

RESULTS

Relationships among NAI and AAI, Exploration Strategies, and Learning

On each trial, an object and two faces were presented in one of four screen quadrants (Figure 1). The object had two features (shape and texture), and the quadrant determined the feature that was relevant for the object-face association. Subjects learned the correct object-face pairings; thus, the relevant feature for each quadrant based on feedback. We used the pattern of quadrant visits and object-face pairings to calculate NAI and AAI (Figure 2; Experimental Procedures). We first sought to identify effects of NAI and AAI on exploration choices and on learning success in the Active condition using the full sample