NeuroImage 172 (2018) 415–426

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Rule activation and ventromedial prefrontal engagement support accurate stopping in self-paced learning

Sean R. O'Bryan [a,*], Eric Walden [b], Michael J. Serra [a], Tyler Davis [a]

[a] Department of Psychological Sciences, Texas Tech University, Lubbock, TX 79409, USA
[b] Rawls College of Business, Texas Tech University, Lubbock, TX 79409, USA

ARTICLE INFO

Keywords:
Attention
Categorization
Decision making
fMRI
Representational similarity analysis

ABSTRACT

When weighing evidence for a decision, individuals are continually faced with the choice of whether to gather more information or act on what has already been learned. The present experiment employed a self-paced category learning task and fMRI to examine the neural mechanisms underlying stopping of information search and how they contribute to choice accuracy. Participants learned to classify triads of face, object, and scene cues into one of two categories using a rule based on one of the stimulus dimensions. After each trial, participants were given the option to explicitly solve the rule or continue learning. Representational similarity analysis (RSA) was used to examine activation of rule-relevant information on trials leading up to a decision to solve the rule. We found that activation of rule-relevant information increased leading up to participants' stopping decisions. Stopping was associated with widespread activation that included medial prefrontal cortex and visual association areas. Engagement of ventromedial prefrontal cortex (vmPFC) was associated with accurate stopping, and activation in this region was functionally coupled with signal in dorsolateral prefrontal cortex (dlPFC). Results suggest that activating rule information when deciding whether to stop an information search increases choice accuracy, and that the response profile of vmPFC during such decisions may provide an index of effective learning.

Introduction

Should I keep studying for my math test? Do I know enough about cars to pick out a good one? When gathering evidence for a decision, individuals are continually faced with the question: Have I learned enough yet? Learners must strike a compromise between collecting enough information to make accurate decisions while avoiding collecting redundant information and—consequently—wasting time and resources.

Currently, little is known about the neurobiological mechanisms that govern decisions about when to stop gathering new information. In the domain of value-based decision making, behavioral research has often focused on heuristics or stopping rules, such as take-the-best, that people employ when presented with cues of varying predictive value (Gigerenzer and Goldstein, 1996). The use of such strategies, however, can vary across participants, even in decision environments that encourage the use of a particular heuristic (Newell and Shanks, 2003; Newell et al., 2004). Thus recent research has begun to focus on participants' use of confidence thresholds for determining when stopping is appropriate, as opposed to application of specific rules per se (Svenson, 1992; Karelaia, 2006; Hausmann and Läge, 2008). In the present study, we test the neural mechanisms that contribute to stopping decisions during learning, and how activation of information associated with a choice evolves leading up to when a decision threshold is reached.

Neurobiologically, in a recent study that required participants to take or decline sequentially-presented stock options, stopping of information search was found to engage anterior cingulate, insula, and ventral striatum (Costa and Averbeck, 2015). Additionally, accumulated value and reward associated with stopping decisions in sequential sampling paradigms have been associated with activation in lateral orbitofrontal cortex, vmPFC, and the basal ganglia (Gluth et al., 2012; Costa and Averbeck, 2015). Although these results have shed light on the neural correlates of stopping in value-based choice, how they translate to stopping in learning contexts, such as when people make decisions about their mastery of new concepts, remains an open question.

Rule-based category learning provides an ideal context to examine the neural basis of stopping in learning because many real world concepts are associated with rules, and because the neural systems that support rule-based categorization are well-understood (for review, see Seger and

* Corresponding author. Texas Tech University, Department of Psychological Sciences, MS 2051, Lubbock, TX 79409, USA.
E-mail address: sean.r.obryan@ttu.edu (S.R. O'Bryan).

https://doi.org/10.1016/j.neuroimage.2018.01.084
Received 31 July 2017; Received in revised form 26 January 2018; Accepted 30 January 2018
Available online 2 February 2018
1053-8119/

Miller, 2010; Ashby and Maddox, 2011). Cognitively, rule-based category learning involves using hypothesis testing and selective attention to establish and focus on stimulus dimensions that are relevant for predicting category membership (Maddox and Ing, 2005; Zeithamova and Maddox, 2006). For example, when learning to distinguish between birds and mammals, people may learn to selectively attend to whether an organism has wings. Selective attention emerges over the course of learning to minimize prediction error, and has the effect of expanding the representation of dimensions that lead to successful categorization (Nosofsky, 1986; Kruschke, 1992; Folstein et al., 2013). Neurobiologically, rule-based category learning is thought to depend on executive cortico-striatal loops connecting the prefrontal cortex with the head of the caudate (Seger and Miller, 2010), with the ventral striatum playing a particularly important role in both initial rule acquisition and reversal learning in the event of a rule switch (Seger and Cincotta, 2006; Liu et al., 2015). The medial temporal lobes are thought to be involved in the long-term maintenance and retrieval of these category rules (Poldrack et al., 2001; Davis et al., 2012).

Currently, how the neurobiological systems involved in rule-based category learning contribute to decisions to stop learning is not clear because the vast majority of neuroimaging studies have employed fixed numbers of trials and not given participants leeway in decisions about whether to continue learning. However, it is possible to infer what mechanisms may underlie decisions to stop learning by incorporating predictions from recent neuroimaging research on stopping decisions in value-based choice and research on confidence in choice behavior more generally. In value-based decision making, the vmPFC has been shown to track subjective value (Tom et al., 2007; Bartra et al., 2013), and is sensitive to cost-benefit discrepancies among response options (Basten et al., 2010; Lim, O'Doherty and Rangel, 2011). These findings are consistent with results showing that the vmPFC tracks accumulated value in value-based stopping decisions. Recent findings suggest that vmPFC may also code general decision evidence or confidence associated with a choice, perhaps in parallel with value (De Martino et al., 2013; Barron et al., 2015; Lebreton et al., 2015). Indeed, a number of basic category learning tasks have found that the engagement of vmPFC is correlated with greater decision evidence for categorization choices (Grinbald et al., 2006; DeGutis & D'Esposito, 2007; Seger et al., 2015; Davis et al., 2017). Thus we expect that decision evidence/confidence signals from the vmPFC may contribute to subjective thresholds participants use when deciding whether they have learned enough information.

In addition to the vmPFC, we also expect the dorsolateral PFC (dlPFC) to contribute to stopping decisions. Recent research has suggested that a region of the posterior dlPFC may be involved in comparing accumulated perceptual information to decision criteria when making perceptual decisions (Heekeren et al., 2004; White et al., 2012). In terms of decisions to stop gathering new information, the dlPFC may monitor information from a number of regions to determine when stopping criteria have been reached, including confidence signals from the vmPFC. Indeed, several studies have observed increased functional connectivity between the dlPFC and vmPFC when participants make decisions that require weighing the subjective values of different choice options (Baumgartner et al., 2011; Rudorf and Hare, 2014). A similar coupling between dlPFC and vmPFC may support computing decision criteria for stopping in category learning.

As a complement to univariate activation, which measures the extent to which brain regions are engaged leading up to decisions to stop, multi-voxel pattern analysis (MVPA) may provide an additional window into how participants are processing category information leading up to stopping decisions. Recent studies have found, as a result of learning, similarities between activation patterns elicited for members of a category come to reflect how participants attend to the stimulus dimensions, such that items sharing values along a rule dimension come to elicit more similar activation patterns (e.g., Mack et al., 2013; Mack et al., 2016). Such changes in neural similarity spaces may track participants' decisions to stop learning. Specifically, we expect that as subjects selectively narrow their attention to a particular rule dimension leading up to a decision to stop gathering information, activation patterns associated with this dimension will become increasingly prominent in the underlying neural similarity space.

To test our predictions for engagement of the vmPFC/dlPFC and how rule-relevant information will be activated leading up to a stopping decision, we trained participants to categorize triads of visual stimuli using simple rules based on one of three binary stimulus dimensions (Fig. 1). The stimulus dimensions consisted of three distinct visual categories: faces, objects and scenes. Participants learned, using trial and error, which stimulus dimension was the rule dimension and was predictive of category membership. Each full task trial was comprised of three to four fixation-separated subcomponents: a learning trial, which involved categorizing the visual stimulus; the presentation of correct/incorrect feedback; a decision trial where participants chose whether to solve the category rule or to continue learning; and finally, if a solve response was made, a trial for selecting between the three possible category rules. The use of a Type I category structure (Shepard et al., 1961) allowed participants to solve rules rapidly, allowing us to robustly measure how activation of rule-relevant information evolved over many different individual decisions to stop learning and solve the rule. To further maximize our ability to detect subtle changes in attentional weighting that result from learning, face, object, and scene images were used as stimulus dimensions. These real-world categories exhibit strong properties for representational decoding (Haxby et al., 2001; O'Toole et al., 2005), and were localized within each subject using independent scans to create ROIs that were unbiased to any potential learning effects (Davis et al., 2014). For the purposes of this study, the representational analysis was focused on the activation of patterns associated with the rule dimension that participants eventually chose as the solution.

To successfully navigate our task, participants must first learn to selectively attend to the rule dimension to predict category membership. We hypothesized that the increasing selective attention to the rule dimension prior to stopping would result in multi-voxel patterns that become increasingly similar to the object class associated with the chosen rule dimension as participants' neared a choice to stop. Once attention is allocated to the potential rule dimension, it is necessary for the participant to monitor the evidence consistent with this category rule. We predicted that vmPFC would be involved in representing the confidence or decision evidence for an attended rule dimension, with stopping marked by greater vmPFC activation than decisions to continue learning. Finally, participants must make a decision to stop learning once this evidence has reached a criterion. We hypothesized that dlPFC would track confidence/decision evidence signals from the vmPFC to determine whether a stopping threshold had been reached.

To preview our findings, the representational analysis revealed a dynamic neural accumulation process whereby activation of multi-voxel patterns associated with the object class eventually chosen as the rule dimension increased over the trials leading up to stopping decisions. Compared with the choice to continue learning, participants engaged a widespread network of brain regions including medial PFC when choosing to stop; within this contrast, activation of vmPFC and object-selective cortex were positively correlated with participants' ability to solve rules accurately throughout the task. Moreover, we show that decisions to stop acquiring information and solve a rule are associated with enhanced functional connectivity between vmPFC and dlPFC.

Materials and methods

Participants

Twenty-five healthy, right-handed volunteers (ages 21–57, mean ± SEM = 27.32 ± 1.67, 17 women) were recruited through online newsletters and flyers posted at Texas Tech University. All subjects were included in the final analysis. All subjects provided written informed consent prior to participation, and were compensated $35 for a 1.5 h