cortex xxx (2017) 1–11

Available online at www.sciencedirect.com

ScienceDirect

Journal homepage: www.elsevier.com/locate/cortex

Special issue: Review

Understanding active sampling strategies: Empirical approaches and implications for attention and decision research

Jacqueline Gottlieb [a,b,*]

[a] *Department of Neuroscience, Columbia University, USA*
[b] *Zuckerman Mind Brain Behavior Institute, Columbia University, USA*

ARTICLE INFO

Article history:
Received 15 March 2017
Reviewed 6 May 2017
Revised 1 August 2017
Accepted 14 August 2017
Published online xxx

Keywords:
Attention
Information seeking
Decision making
Eye movements
Active sensing

ABSTRACT

In natural behavior we actively gather information using attention and active sensing behaviors (such as shifts of gaze) to sample relevant cues. However, while attention and decision making are naturally coordinated, in the laboratory they have been dissociated. Attention is studied independently of the actions it serves. Conversely, decision theories make the simplifying assumption that the relevant information is given, and do not attempt to describe how the decision maker may learn and implement active sampling policies. In this paper I review recent studies that address questions of attentional learning, cue validity and information seeking in humans and non-human primates. These studies suggest that learning a sampling policy involves large scale interactions between networks of attention and valuation, which implement these policies based on reward maximization, uncertainty reduction and the intrinsic utility of cognitive states. I discuss the importance of using such paradigms for formalizing the role of attention, as well as devising more realistic theories of decision making that capture a broader range of empirical observations.

1. Introduction

The oculomotor system of humans and non-human primates holds a privileged status in neuroscience research. Motivated by the relative simplicity of the eye motor plant, the relative ease of measuring eye movements in the laboratory and the high degree of similarity between humans and non-human primates, scores of investigations have examined saccades – the rapid shifts of gaze that primates use to scan visual scenes – and the neural pathways involved in their generation.

However, while these studies have elucidated many of the sensorimotor mechanisms involved in saccades, progress stalled in explaining the cognitive aspects of saccades and attention – specifically, how the brain selects task-relevant cues. Behavioral evidence makes it clear that gaze is under strong task-related control – with humans deploying gaze very selectively to stimuli that are relevant to their immediate

* *Corresponding author*. Department of Neuroscience, Columbia University, 1051 Riverside Drive, Kolb Research Annex, Rm. 569, New York, NY 10032, USA.
E-mail address: jg2141@columbia.edu.
http://dx.doi.org/10.1016/j.cortex.2017.08.019

actions with minimal influence from salient distractors (Yarbus, 1967; Tatler, Hayhoe, Land, & Ballard, 2011). However, computational models of gaze allocation are based primarily on bottom-up saliency (Berg, Boehnke, Marino, Munoz, & Itti, 2009; White, Berg, Marino, Itti, &, 2017) with many fewer attempts to model task-related control (Navalpakkam & Itti, 2005; Tatler et al., 2011).

This gap in our understanding is particularly vexing for neurobiological investigations of oculomotor structures implicated in the selection of targets for attention or gaze, which include the superior colliculus, the frontal eye field (FEF) and the lateral intraparietal area (LIP) (Thompson & Bichot, 2005; Bisley & Goldberg, 2010; Krauzlis, Lovejoy, & Zénon, 2013). While abundant evidence shows *that* neurons in these areas encode top-down visual selection – selectively signaling the locations of task-relevant visual stimuli – we have little insight into *how* these responses arise. How do target selection neurons "know" which target to select? What is the computational definition of a "task-relevant" cue? While several lines of research have linked target selection responses in LIP with simple decisions based on perceptual evidence or rewards (Sugrue, Corrado, & Newsome, 2005; Kable & Glimcher, 2009; Hanks & Summerfield, 2017), these studies have yet to consider the unique information sampling nature of gaze (Gottlieb, Hayhoe, Hikosaka, & Rangel, 2014) and leave persistent unresolved questions about the selection process encoded by the cells (Maunsell, 2004; Gottlieb, 2012).

In this article I argue that, to understand task-related control, we must acknowledge the essential role of attention and gaze in sampling information – the fact that, in natural conditions, gaze and attention implement an *active sensing policy* that is coordinated with the decision maker's beliefs, goals or actions. The informational – or epistemic – nature of saccades and attention is recognized by theoretical frameworks- such as predictive coding, which emphasize the imperative of minimizing surprise or free energy (Friston, 2010; Friston & Ao, 2012; Schwartenbeck, Fitzgerald, Dolan, & Friston, 2013; Friston et al., 2015) or expanded reinforcement learning theories (Iigaya, Story, Kurth-Nelson, Dolan, & Dayan, 2016) – but we have scant empirical data that can constrain or refine these theories.

I this paper I review the few studies that have addressed questions concerning saccades, attention and information sampling, with a focus on behavioral paradigms that can probe the logic of active sensing policies and the key current findings regarding these policies in humans and monkeys. I argue that, although these approaches are relatively new to the field, developing them is essential for expanding our current understanding of both attention and decision making and bringing about a closer integration of research on these topics.

2. What does an observing decision entail?

Because the questions we are about to consider are relatively unfamiliar in the study of oculomotor control, it is useful to start by considering the computations that may be entailed by an active sampling policy. Active sampling is an ubiquitous aspect of natural behavior, and a core building block of the perception–action cycle: when reaching an intersection we *look* at the traffic (or a traffic sign, or a traffic light) to decide whether to stop of proceed and, before deciding whether to reach for the peanut butter jar we *look* at the jar. Understanding active sampling, therefore, requires us to consider two related decisions: the selection of a task-relevant cue, and the decision of which action to take based on that cue.

Sequential decisions of this kind are typically analyzed (e.g., in reinforcement learning frameworks) using a decision chain such as that illustrated in Fig. 1A, which specifies a sequence of states that the decision maker expects to traverse in a task, and the probabilistic actions and transitions that are possible from each state. In the case of a pedestrian reaching an intersection (Fig. 1A), the chain may start with the decision of whether to look at the traffic light or a cloud, followed by the decision of whether to stop or proceed followed by the observation of an outcome (e.g., staying safe, operationalized as a reward probability).

Our concern is with the first decision in this chain – the determination of which stimulus to sample – and the diagram in Fig. 1A illustrates three key points about this step: it depends on prior knowledge of the task structure, it may be guided by both expected rewards and the prospect of resolving uncertainty, and it requires the agent to estimate the desirability of the available cues in advance of the full sensory discrimination. Let us consider each feature in turn.

2.1. *Model-based selection*

One of the most important features of active sampling policies is that, like other types of decisions, they depend on prior knowledge of the task structure. This knowledge is embodied in a task model such as that shown in Fig. 1A, which specifies the states and actions involved in a task, as well as the relation between stimuli and subsequent states. It is only based on this knowledge that the agent can estimate the probability (or uncertainty) of competing actions, the *meaning* of sensory cues, and the information that the cues may bring about future states (e.g., that the colors of the traffic light are associated with crossing or waiting). This implies a hierarchical process whereby prior knowledge of the task structure organizes local sampling strategies. In other words, we need to know what we are doing in order to know what to sample. As we will see in the following sections, the role of hierarchical learning in task-related saccade and attention control is an important topic for further investigation.

2.2. *Dependence on reward and uncertainty*

A second critical feature is that, in the context of a task model, there are two possible mechanisms for distinguishing between informative and uninformative cues: the reward expectations associated with a cue, and the prospect that a cue will alter the decision maker's beliefs about future states (Sullivan, Johnson, Rothkopf, Ballard, & Hayhoe, 2012; Johnson, Sullivan, Hayhoe, & Ballard, 2014).

In conditions where the decision maker can act based on the sampled information – so called *instrumental* sampling paradigms – an informative cue is *by definition* one that signals the more desirable action, and thus the reliability of a cue is closely correlated with the chance of success in the task. In

Please cite this article in press as: Gottlieb, J., Understanding active sampling strategies: Empirical approaches and implications for attention and decision research, Cortex (2017), http://dx.doi.org/10.1016/j.cortex.2017.08.019