Forensic foraging of change detection in opponent strategies with a neural model of the interactions between temporal and prefrontal cortex

Matthew E. Phillips [a],*, Suhas E. Chelian [a], Peter Pirolli [b], Rajan Bhattacharyya [a]

[a] HRL Laboratories, LLC, Malibu, CA, United States
[b] Palo Alto Research Center, Palo Alto, CA, United States

Abstract

Change detection, and which information to attend to, are key research problems relevant to understanding adaptive behavior. Rational analyses of change detection have been developed in optimal foraging theory (McNamara & Houston, 1987; Stephens, 1987) and psychology (Gallistel, Krishan, Liu, Miller, & Latham, 2014). Information foraging models (Pirolli, 2007) have been developed to predict the optimal choice of information, and when to terminate collecting additional information. Optimal performance depend crucially on changing behavior when the world changes. For example, intelligence analysts track the dynamically changing strategies of adversaries and must choose information to weigh the utility of alternate actions regarding these agents. Neurobiologically-plausible descriptions of these behaviors have thus far been fragmentary (cf., Hayden, Pearson, & Platt, 2011). Here, behavioral and mechanistic differences in belief updating about an agent's strategy, and information foraging choices were investigated in a variant of the patch foraging task. Human subjects typically displayed non-normative updating and information foraging decisions which influenced future belief updating and foraging behaviors. To explore the biological basis of these behaviors, a neural model which forages, processes, and updates the beliefs of competing hypotheses was created. Strategy updating and information foraging behaviors were modeled in the temporal and prefrontal cor-

* Corresponding author.
  E-mail addresses: mephillips@hrl.com (M.E. Phillips), sechelian@hrl.com (S.E. Chelian), pirolli@parc.com (P. Pirolli), rbhattac@hrl.com (R. Bhattacharyya).

tices consistent with theories of human cognitive processing during adversarial game play. The neural model reproduced human behaviors observed in foraging, belief updating, and change detection. These results highlight the importance of updating strategy, foraging, and fatigue behaviors and provide a simulation framework for future studies to investigate other high-level cognitive behaviors.

Introduction: Change detection with information foraging

Change detection is a process critical to the survival of many organisms and crucial to human performance in many modern environments (e.g., investing in financial markets; intelligence analysis). Rational models of change detection have been developed in optimal foraging theory (e.g., McNamara & Houston, 1987; Stephens, 1987) for simplified, hypothetical environments. These models suggest that to survive in the natural world, animals need to actively explore the environment and need to do so in a way that maximizes the utility of their information foraging explorations. Several algorithms have been proposed to account for the reasoning processes involved in these behaviors. For example, a rank-order based model of attentional saliency was recently proposed to account for belief updating and decision-making behaviors in a dynamically-changing environment (Phillips et al., 2013b). Furthermore, anomaly detection, a related problem to change detection, can be addressed by several approaches including those that use hippocampal dynamics (Srinivasa & Chelian, 2007) and reticular inhibition (Phillips, Avery, Krichmar, & Bhattacharyya, 2013a). Another adaptation problem occurs in environments in which resources are encountered in ''patches,'' and foraging within patches has diminishing returns (e.g., fruit-bearing trees; herds of prey; information presented on a web page). Optimal foraging theory developed the *patch foraging model* (Charnov, 1976) to provide a rational analysis of how much time should be invested in processing individual patches, and this analysis has been extended to human information foraging behavior (Pirolli, 2007). In this paper, we present a neurocomputational model of an intelligence analysis task that involves both a patch foraging problem and a change detection problem. The task involves interacting with visualizations representing events occurring over time (a patch foraging problem) with the aim of discerning whether the tactics of an adversary have changed (a change detection problem).

Materials and methods

Neural model

The neural simulations performed in this work are based on the Leabra framework in the Emergent simulation software, described in (O'Reilly and Munakata, 2000). Emergent is a full-featured neural network simulator descended from PDP and PDP++. Emergent primarily supports complex, and multi-regional brain models (see Aisa, Mingus, & O'Reilly, 2008 for a detailed description).

Within the Emergent simulator, neurons summate excitatory, inhibitory, and leak conductances into a membrane potential that is thresholded and passed through a non-linear function to produce rate coded outputs. Layers in the model use ensembles of excitatory neurons with k-winner-take-all inhibition to simulate both populations efficiently. Learning is accomplished through a mixture of Hebbian and error driven learning that incorporates an adaptive threshold for firing. The model operates in 2 phases: a minus phase where only input stimuli are clamped onto input units (with no synaptic modification), and a plus phase where both input and output units are clamped (i.e., allowed to fire). Differential neural activations from the two phases drive error driven learning. Connections between layers follow a variety of topologies, including one-to-one, full, and geometric to create specific receptive fields.

Behavioral task

Human subjects (N = 123) and neural models (N = 40) played a variant of the patch foraging task adapted to a geospatial intelligence setting against a dynamic computer opponent. Subjects were told that the opponent's strategy would change exactly once during 30 trials. On each trial, subjects first estimated the likelihood of the opponent's strategy between 2 options (aggressive—likely to attack a point of interest or POI, and passive—less likely to attack the POI). In this task, an attack implicates an aggressive opponent strategy with a higher likelihood (although it is not constant across the 4 trial conditions), and a not-attack implicates a passive opponent (see Fig. 1, attack likelihoods averaged across all four trail conditions are 35% for passive and 65% for aggressive). After 1 non-attack and 9 attacks, for example, a subject might rate the computer opponent is 10% likely to be passive and 90% likely to be aggressive. After reporting strategy beliefs, subjects received intelligence layers or INTs that indicated the probability of an attack from the computer opponent. Subjects then chose to either divert or not divert against a potential attack. Next, the computer opponent would either attack the POI or not and subjects would receive payoff information based on their decision (e.g., not diverting when there was attack was a loss). At trials 10, 20, and 30 and only those trials, subjects were allowed to create a ''batch plot'' to examine the history of attacks in previous trials. After this batch plot, they could revise their beliefs of the computer opponent strategy and proceed with assessing the probability of an attack given intelligence layers on subsequent trials.