

The orbitofrontal cortex as a negative feedback control system: computational modeling and fMRI

Noah Zarr (noahnz@gmail.com)

Dept. of Psychological & Brain Sciences, Indiana University
1101 E Tenth St., Bloomington, IN 47405 USA

Joshua W Brown (jwbrown@indiana.edu)

Dept. of Psychological & Brain Sciences, Indiana University
1101 E Tenth St., Bloomington, IN 47405 USA

Abstract:

In this work we address two inter-related issues. First, the computational roles of the orbitofrontal cortex (OFC) and hippocampus in value-based decision-making have been unclear, with various proposed roles in value representation, cognitive maps, and prospection. Second, reinforcement learning models have been slow to adapt to more general problems in which the reward values of states may change over time, thus requiring different Q values for a given state at different times. We have developed a model of artificial general intelligence that treats much of the brain as a high dimensional control system in the framework of control theory. We show with computational modeling and combined fMRI and representational similarity analysis (RSA) that the model can autonomously learn to solve problems and provides a clear computational account of how a number of brain regions, particularly the OFC, interact to guide behavior to achieve arbitrary goals.

Keywords: goals; value-based decision-making; cognitive control; control theory; orbitofrontal cortex

Introduction

A significant limitation of both model-based and model free RL is that typically there is only a single ultimate goal. Q-values (Watkins & Dayan, 1992) are thus learned in order to maximize a single reward value. In contrast, real organisms will find differing reward values associated with different goals at different times and circumstances. This implies that goals will change over time, and re-learning Q-values with each goal change would be highly inefficient. Instead, a more flexible mechanism will dynamically assign values to various goals and then plan accordingly.

Methods

We developed the Goal-Oriented Learning and Sequential Action (GOLSA) model as a new approach to overcome the limitations of less flexible Q-values, while maintaining fidelity to known biological mechanisms and constraints such as localist learning laws. The model treats the brain as a high-dimensional control system. It drives behavior to maintain multiple and varying control theory set points of the agent's

state, including low level homeostatic and high level cognitive states. The model learns the structure of state transitions, then plans actions to arbitrary goals via a novel hill-climbing algorithm inspired conceptually by Dijkstra's algorithm (Dijkstra, 1959), and similar to that used in GPS navigation devices. The model provides a domain-general solution to the problem of solving problems and performs well.

The GOLSA model works by representing each possible state of the agent and environment in a network layer, with multiple layers each representing the same sets of states (Figure 1). The Goal Gradient layer is activated by an arbitrarily specified desired (Goal) state and spreads activation backward along possible state transitions represented as connections in the network. The Adjacent States layer receives input from a node representing the current state of the agent and environment and activates representations of all states that can be achieved with one state transition. The valid adjacent states then mask the Goal Gradient layer to yield the Desired Next State representation, which if achieved, will move the agent one step closer to the goal state. With that, the desired next state is mapped onto an action that is likely to effect the desired transition.

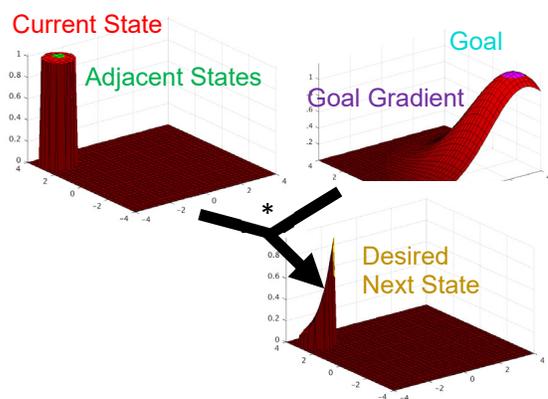


Figure 1: GOLSA model determines next desired state by hill climbing.

Results

Behaviorally, we found that the GOLSA model is able to learn to solve arbitrary problems (Figure 2).

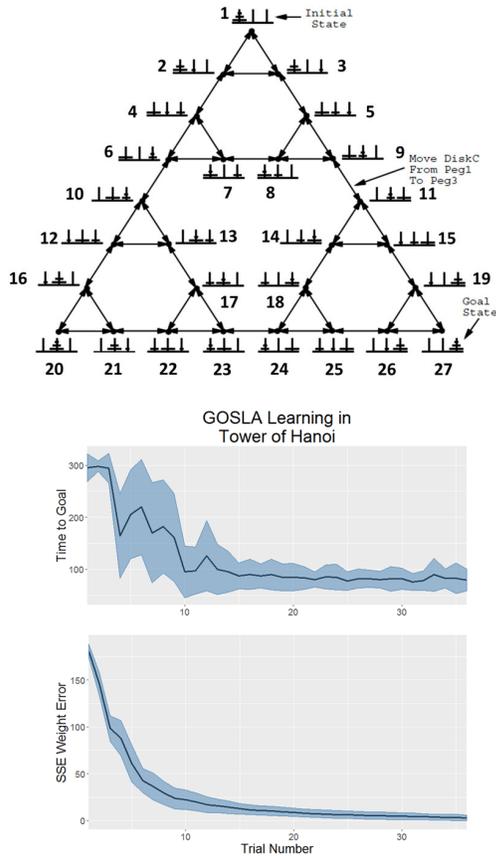


Figure 2: The GOLSA model learns to solve the Tower of Hanoi problem with arbitrary starting and desired ending states, with no hardwired task information.

To test the GOLSA model further, we use model-based fMRI with representational similarity analysis (RSA) (Kriegeskorte, 2008). We found that in addition to solving complex planning problems, the GOLSA model provides a novel computational account of network interactions of a number of brain regions involved in flexible action planning (Figure 3). The orbitofrontal cortex activity patterns match model components that represent both a cognitive map (Wilson, Takahashi, Schoenbaum, & Niv, 2014) and a flexible goal value representation (Schoenbaum, Takahashi, Liu, & McDannald, 2011), specifically matching the Goal and Goal Gradient layer activities. The hippocampus and striatum represent a conjunction of the current state and desired future state transitions (Buckner, 2010), which in the model is a necessary step toward selecting an appropriate action. The model and

RSA analyses account for specific roles of visual cortex, anterior inferior temporal cortex, and motor cortex as well.

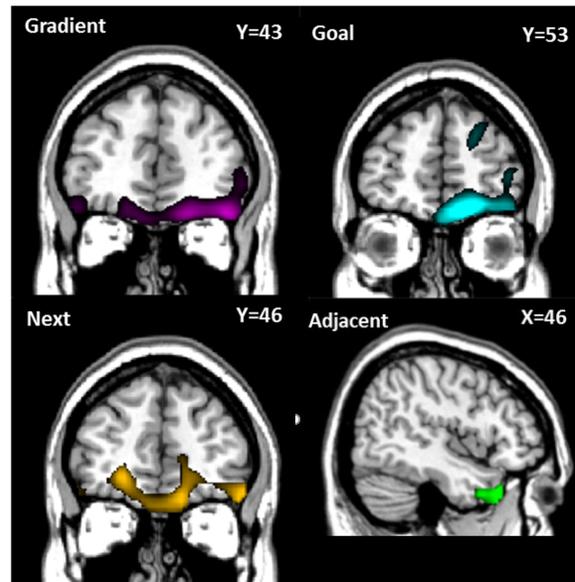


Figure 3: Representational Similarity Analysis of model layers vs. human subjects performing the same problem solving task.

Conclusion

Our results suggest a novel computational account of how the brain plans actions to solve problems, and how a number of brain regions perform interacting computational roles to such behavior. The orbitofrontal cortex represents a cognitive map of which state transitions are possible and also assigns value flexibly by activating the representation of whatever state is currently a desired goal state.

Acknowledgments

Supported by the Indiana University Imaging Research Facility.

References

- Buckner, R. L. (2010). The Role of the Hippocampus in Prediction and Imagination. *Annual Review of Psychology*, 61(1), 27–48.
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1), 269–271.
- Knoblock, C. A. (1990). Abstracting the Tower of Hanoi. In *IN WORKING NOTES OF AAAI-90 WORKSHOP ON AUTOMATIC GENERATION OF*

APPROXIMATIONS AND ABSTRACTIONS (pp. 13–23). AAAI.

Kriegeskorte, N. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*.

Schoenbaum, G., Takahashi, Y., Liu, T.-L., & McDannald, M. A. (2011). Does the orbitofrontal cortex signal value? *Annals of the New York Academy of Sciences*, 1239(1), 87–99. 210.x

Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3–4), 279–292.

Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal Cortex as a Cognitive Map of Task Space. *Neuron*, 81(2), 267–279.