A Brain Inspired Cognitive Architecture for Autonomous Development

Tyler Streeter
Iowa State University, Ames, IA USA
VR Applications Center, Brainpower Labs LLC
www.tylerstreeter.net
tylerstreeter@gmail.com

James Oliver
Iowa State University, Ames, IA USA
VR Applications Center
www.vrac.iastate.edu/∼oliver
oliver@iastate.edu

Abstract

We present a novel cognitive architecture, Sapience, inspired by the high-level organization of the mammalian brain. This architecture, which provides a platform for autonomous development, combines a powerful set of computational elements, including a hierarchical learning Bayesian network, reinforcement learning, and artificial curiosity. This preliminary abstract gives a brief description of the architecture design and our plans for initial experimentation.

1. Introduction

We aim towards the long-term goal of engineering a general purpose "software brain" for autonomous development which could be applied to all kinds of real-time motor control tasks. How can this goal be achieved? Our general strategy is to design a cognitive architecture combining systems-level neuroscience with machine learning methods. We first take inspiration from the mammalian brain's high-level organization and function (but note that our focus is brain-inspired engineering, not biologically plausible brain models.) To make the architecture concrete, we then choose specific machine learning algorithms for every component, each providing a unique computational benefit to the overall system.

We use theoretical reinforcement learning (RL) as an organizing principle (Sutton and Barto, 1998). Thus, the primary objective is to learn complex motor control tasks defined as RL problems. We distinguish between two related objectives: to solve externally given RL tasks, and to improve an internal world model (which implicitly helps external RL). Corresponding to these objectives are two types of reinforcements: external (programmer-defined) reinforcements which define the task, and internal "curiosity rewards" for world model improvements, which encourage active sensory exploration.

2. Architecture

Our cognitive architecture (see figure), which we call "Sapience," represents a complete learning system which can be embodied within a real or simulated environment. It can handle arbitrary reinforcements and real-valued sensory and motor arrays. It is composed of five distinct components, functional abstractions of the brain's sensorimotor cortex, hippocampus, basal ganglia, cerebellum, and prefrontal cortex regions. It is designed to be practical: directly implementable in software, scalable with available computing resources, and generally applicable (e.g., for video games and robotics). The entire architecture has already been implemented in software, along with many test programs and analysis tools. No aspect described here is merely a hypothetical mechanism; every part has been instantiated explicitly in C++. Here we describe the architecture at a high-level, illustrating its overall organization and briefly introducing each component. More details will appear in the primary author's PhD thesis (in preparation).

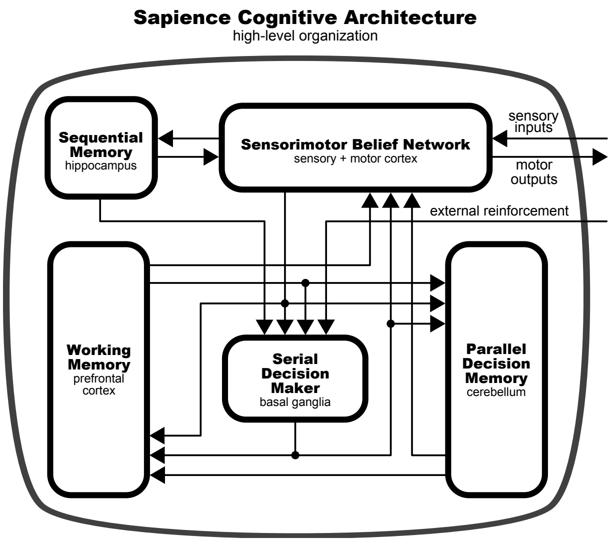

The **Sensorimotor Belief Network** is a probabilistic world model representing both sensory and motor patterns. It uses a "hierarchical empirical Bayesian network," combining practical belief propagation (Pearl, 1988) with unsupervised learning (kernel-based Maximum Entropy learning Rule, or kMER (Van Hulle, 2000)). We measure model improvements (for curiosity rewards) as the Jensen-Shannon divergence between prior and posterior distributions, a measure of information gain similar to that used in (Schmidhuber et al., 1995). The **Sequential Memory** is a neural sequence predictor which provides a prior distribution to the Bayesian hierarchy's root node. The **Serial Decision Maker** performs reinforcement learning with neural network-based value function and policy. It combines both external and internal (curiosity) reinforcements. Its action selection influences both motor outputs and working memory contents. The **Parallel Decision Memory** automates well-learned actions via supervised learning, freeing the RL process for trial-and-error learning, not repetitive actions, as described in (Peck et al., 2007). The **Working Memory** provides a set of general purpose memory cells whose contents can be actively stored and recalled, extending the set of RL choices to include motor *and* working memory actions, an arrangement inspired by the prefrontal cortex/basal ganglia/working memory (PBWM) model (Hazy et al., 2007).

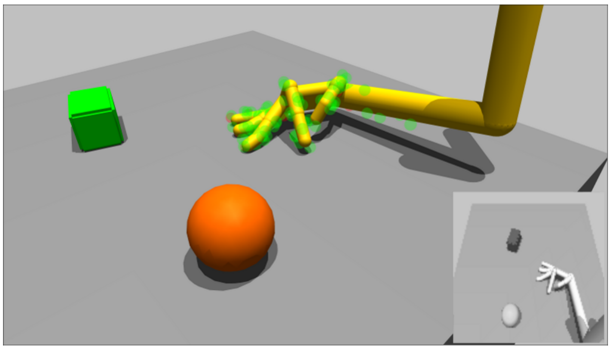

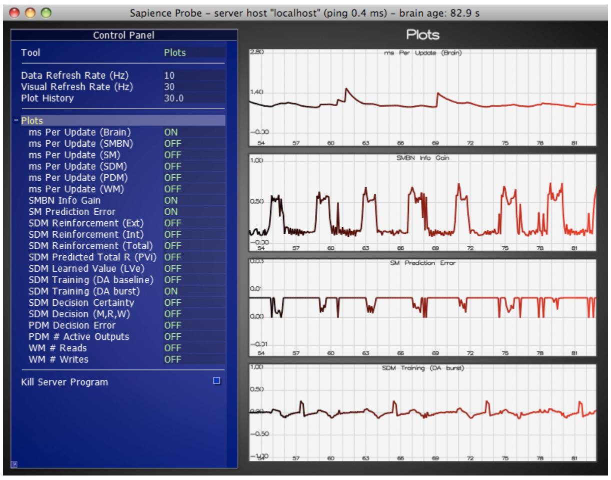

3. Experimental Platform

Currently much of our work is devoted to evaluating the architecture implementation. One of our experimental environments is a physically realistic arm and hand simulation (see figure) with tactile, visual, and proprioceptive sensory inputs and servo motor outputs. This test platform will be used for embodied motor learning, such as unsupervised learning of sensory/motor patterns, learning reaching tasks from scratch (no human guidance besides simple rewards), and open-ended curiosity-driven motor exploration.

For this kind of multi-component, multi-algorithm architecture, it is important to visualize internal data structures in real-time as learning progresses. Thus, we have created a probe tool (see figure), a client program which connects to a data server (embedded in the Sapience implementation) and generates plots of individual values and visual representations of 2D data arrays, neural networks, etc. Such a tool is crucial for observing and analyzing learning processes within our implementation as it drives a body in real or simulated environments.

Acknowledgements

This work is supported by AFOSR (FA9550-05-1-0384) and Brainpower Labs LLC. We are also grateful for the research infrastructure at the VRAC.

References

Hazy, T. E., Frank, M. J., and O'Reilly, R. C. (2007). Towards an executive without a homonculus: Computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society B*, 362:1601–1613.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann.

Peck, C., Streeter, T., and Kozloski, J. (2007). An integrated cerebro-cerebellar model demonstrating associative learning and motor control. In *Proceedings of the 10th Tamagawa-Riken Dynamic Brain Forum.*

Schmidhuber, J., Storck, J., and Hochreiter, J. (1995). Reinforcement driven information acquisition in non-deterministic environments. In *Proceedings of ICANN*, pages 159–164.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction.* MIT Press.

Van Hulle, M. M. (2000). *Faithful Representations and Topographic Maps: From Distortion- to Information-Based Self-Organization.* Wiley.