An Engineering View of Cognitive Architectures

Steven B. Morphet

SRC (formerly Syracuse Research Corporation)
7502 Round Pond Road, Syracuse, NY 13212
morphet@srcinc.com

Abstract

Engineering principles when coupled with a clear set of goals/constraints/behaviors can help guide the development of effective cognitive architectures.

Position Statement

When designing a cognitive architecture, two important questions [among many] need to be answered. Specifically: (1) what components comprise a minimal cognitive substrate? (2) which level of the biological cognitive hierarchy should be focused on to effectively develop a cognitive architecture?

These questions are tightly coupled, and their answers are largely driven by the underlying implicit assumptions. For instance, is there an implication that the cognitive architecture will perform tasks that require human-level intelligence? This begs the question "what is trying to be achieved?" When a sufficiently detailed and accurate response to this question is formulated, one will be left with the set of necessary behaviors.

Taking inspiration from biology and given the necessary behaviors, it is possible to delineate a minimum set of biologically inspired modules that when coupled correctly will exhibit the desired behavioral properties.

It is difficult to know a priori what level of biological resolution will serve best – this is a function of the underlying task(s). But the answer to this question will also answer "at what level do we, as researchers, need to understand how the brain works (and what it does)?"

At a very low level of the resolution hierarchy, one can consider the neuron to be an atomic computation element. At a higher level, a neural cluster is a functional element, with capabilities including arithmetic or logical operations, correlations, coordinate transformations, etc. At an even higher level, the cortical column is a collection of neural clusters able to perform focus of attention, classification, etc.

When the best level of resolution is determined, there is still the question of how to emulate the lower levels. Fortunately, the concept of function equivalence comes to the rescue. Specifically, functional equivalence is defined as producing the same input/output behavior regardless of underlying implementation. So, for instance, it is not necessary to implement neural clusters as neurons; rather, neural clusters implemented as a lookup table with sufficient fidelity would be not only adequate, but indistinguishable from the neural implementation. The engineering benefit: efficient software emulations can turn intractable problems into tractable ones. Additionally, the right level of resolution helps dictate the functional modules.

Another important question regarding the design of a cognitive architecture is "how should the modules be 'wired' together?" To answer this question, one can examine the role of evolution. Looking at species development and the impact of evolution on human cognitive wet-ware, one can make an argument for (1) each brain area having a [relatively] fixed working function; (2) each function being implemented in overlapping neural structures. Giving these concepts some thought, they should be intuitively satisfying. Unfortunately the second is contrary to state-of-the-art engineering principles. That does not necessarily make it incorrect or bad; it does however present a challenge to the current thought processes in software design and reuse. The most prominent engineering implications of this insight are: (1) modularity in its classic form must be abandoned (or heavily revised); (2) assigning computational/cognitive roles to brain areas will require cross-domain modeling; (3) cross-domain uses must be considered at design time.

These are merely a few of the considerations that need to be addressed when designing a cognitive architecture. Many biological constraints and experimental results [from psychology] exist to help drive development. The optimal set of constraints and results is largely identified by the required goals and behaviors of the desired cognitive architecture.