Cognitive Architecture Requirements for Achieving AGI

John E. Laird[*], *Robert E. Wray III*[**]

[*]Division of Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109-2121
[**]Soar Technology, Inc., 3600 Green Court, Suite 600, Ann Arbor, MI 48105

Abstract

We outline eight characteristics of the environments, tasks, and agents important for human-level intelligence. Treating these characteristics as influences on desired agent behavior, we then derive twelve requirements for general cognitive architectures. Cognitive-architecture designs that meet the requirements should support human-level behavior across a wide range of tasks, embedded in environment similar to the real world. Although requirements introduced here are hypothesized as necessary ones for human-level intelligence, our assumption is the list is not yet sufficient to guarantee the achievement of human-level intelligence when met. However, attempts to be explicit about influences and specific requirements may be more productive than direct comparison of architectural designs and features for communication and interaction about cognitive architectures.

Introduction

This paper explores requirements on cognitive architectures for artificial general intelligence. The goal of the analysis is to determine the requirements for cognitive architectures that support the full-range of human-level intelligent behavior. Although many different architectures have been proposed (and some built), understanding the relative strengths and weaknesses of these architectures and their unique contributions to the pursuit of human-level intelligence has proven elusive, whether via analytic comparison (e.g., Anderson & Lebiere, 2003; Jones and Wray, 2006) or empirical comparisons on task performance (e.g., Gluck & Pew. 2005).

However, as suggested by Cohen (1995), three influences determine an agent's behavior: the agent's structure, its environment, and tasks.[1] Given the diversity of environments and tasks, we are not attempting to create architectures that are necessarily the best or even sufficient for *all* possible environments and *all* possible tasks. We assume that agents exist in an environment and pursue tasks similar to those we find in the world we inhabit. The challenge is to take advantage of the structure of the environment and tasks in our architecture design, while avoiding optimizations that apply to only a subset of tasks. For specific problems, specialized architectures can be more appropriate (e.g., Deep Blue for chess, Campbell, Hoane & Hsu 2002).

---

[1] We use "task" for any type of problem, goal, drive, or reward that provides direction for agent behavior.

Figure 1 illustrates how the characteristics of the environment, tasks, and agent structure determine a set of requirements for a cognitive architecture. These requirements in turn are the basis for a specific architecture design. Cognitive architectures must provide a comprehensive computational story that puts all the pieces of intelligence together from end to end.
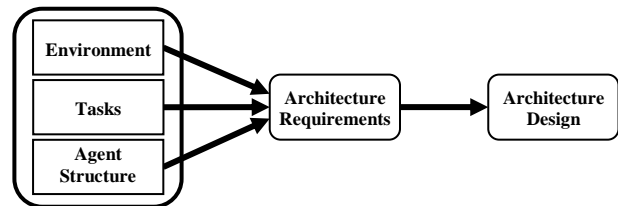

Figure 1: Influences on architecture design.

In practice, researchers have typically focused on communicating the architectural design of their systems and its performance on specific tasks rather than motivating the design via specific requirements. We propose to orient future discussion around requirements rather than specific designs. There will be two immediate benefits to this approach. First, it makes little sense to compare architectures (as works in progress) when they share few requirements. If one architecture attempts to satisfy a requirement that all decisions must be made in bounded time, whereas another is developed independent of that requirement, we would expect to see very different approaches that would be difficult, if not meaningless to compare. Being explicit about requirements will make it easier to see what "spaces" architectures are attempting to occupy – what environments and problems they are appropriate for. Secondly, because human-level intelligence is so broad, there is no existing list of necessary and sufficient requirements of AGI. This paper, drawing from our experience, proposes an initial list of these requirements. We expect it to be refined, extended, and corrected via interaction with other researchers.

We recognize this attempt is not novel. John Anderson took a step in this direction with the design of specific components of ACT-R using a rational analysis (Anderson, 1990). He determined optimal methods for primitive architectural functions, such as retrieving an item from declarative memory given the expected use of that memory in the future. This revolutionized his design process and led to significant advances in ACT, including the development of a new process for retrieving items from long-term declarative memory. Although rational analysis

is useful for designing the performance of specific components, it is difficult to apply to the specification of a complete cognitive architecture as it does not specify what components there should be, or how they combine together to provide general intelligent behavior.

The analysis also builds on previous descriptions of evaluation criteria for cognitive architectures (Langley, Laird & Rogers, 2009; Laird et al. 2009; Laird et al., 1996; Laird, 1991) and theories of cognition (Anderson & Lebiere, 2003; Newell, 1990). For example, previously identified criteria include a mixture of constraints on behavior (flexible behavior, real-time performance) architecture (support vast knowledge bases), and underlying technology (brain realization) (Newell, 1990, Anderson & Lebiere, 2003). We separate characteristics of the environment, tasks, agent structure, and the behavior of an agent, which are described below, and then use them to derive requirements for cognitive architectures.

Environment, Task, and Agent Characteristics

In this section, we list characteristics of environments, tasks, and agents that lead to requirements for architectures that support human-level intelligent agents. Some of these characteristics are obvious, or so ingrained in the literature that they are rarely made explicit, such as the existence of regularities at different time scales in the environment. Some of these are characteristics of one of the three components, independent of the others, but many of them are characteristics of interactions between two or even all three. The interactions are important because the characteristics of an environment are only important to the extent they influence the agent's ability to pursue its tasks.

C1. ENVIRONMENT IS COMPLEX WITH DIVERSE INTERACTING OBJECTS
The world is large and complex. Agents can usefully interpret the environment as if it consists of independent objects (together with materials that do not have object-like structure, such as air, water, and sand). There are many objects and the objects interact with each other (i.e., via physics). Objects have numerous diverse properties.

C2. ENVIRONMENT IS DYNAMIC
The agent's environment can change independently of the agent so that the agent does not determine the state of the environment and the agent must respond to the dynamics of the world. Because the world can change while an agent is reasoning, an agent must be able to respond quickly relative to the dynamics of the environment. Moreover, the dynamics of the environment are so complex that an agent cannot always accurately predict future states in detail.

C3. TASK-RELEVANT REGULARITIES EXIST AT MULTIPLE TIME SCALES
An environment, while it may be complex and dynamic, it is not arbitrary. The environment is governed by laws of interaction that are constant, often predictable, and lead to

recurrence and regularity that impact the agent's ability to achieve goals. Regularities exist at a variety of time scales.

C4. OTHER AGENTS IMPACT TASK PERFORMANCE
The agent is not alone, and must interact with other agents in pursuit of its goals. Other agents may help or hinder the agent's achievement of its tasks. The agent can communicate with the other agents to share knowledge, indicate intent, etc. In addition, some agents have similar structure and capabilities to the agent (similar perception, action, and mental capabilities), making it possible to learn from other agents by observing the methods they use for solving problems. This characteristic is a special case of C1, C2, and C3, but has sufficient impact on the structure of agents to warrant distinct enumeration.

C5. TASKS CAN BE COMPLEX, DIVERSE, AND NOVEL
A general, intelligent agent must be able to work on a diverse set of novel, complex tasks. Tasks can interact so that in some cases, achieving one task aids in achieving another, while in other cases, achieving one makes it more difficult to achieve another. Tasks can also vary in the time scales required to achieve them, where the agent must achieve some tasks at close to the timescale of relevant changes in the environment, while others tasks can require extended behavior over time.

C6. AGENT/ENVIRONMENT/TASK INTERACTIONS ARE COMPLEX AND LIMITED
There may be many regularities in the environment, but they are only relevant if they can be detected and influence the agent's ability to perform its tasks. Thus, an agent has sufficient sensory capabilities that it can detect (possibly only through extensive learning) task-relevant regularities in the environment. An agent also has many mechanisms for acting in the environment in order to pursue a task. Although sensing and action modalities can be extensive, they are limited. The environment is partially observable, both from inherent physical limits in the sensors and the size of the environment. Sensors have noise and can be occluded by objects, have limited range, etc. making the agent's perception of its environment incomplete and uncertain. The agent's actions must obey the physical limitations of the environment. For example, actions usually take time to execute and have limited extent.

C7. AGENT COMPUTATIONAL RESOURCES ARE LIMITED
The agent has physical limits on its computational resources relative to the dynamics of the environment. The agent is unable to perform arbitrary computation in the time it has available to respond to the environment. Thus, an agent has *bounded rationality* (Simon, 1969) and cannot achieve perfect rationality (or universal intelligence, Legg & Hutter, 2007) in sufficiently complex environments and tasks when it has large bodies of knowledge.

C8. AGENT EXISTENCE IS LONG-TERM AND CONTINUAL
The agent is always present in its environment and it needs to actively pursue core tasks (such as self-protection) related to its survival. The agent may act to position itself

so that the dynamics of the environment have little impact on it for extended times (e.g., hide in a protected area), but it has no guarantee that those efforts will be successful. Further, the agent has a long-term existence relative to its primitive interactions with its environment. Its activity extends indefinitely across multiple tasks, and possibly multiple instances of the same task.

Architectural Requirements

Based on the characteristics of environments, tasks, and agents presented in the previous section, we derive the following requirements for cognitive architectures related to knowledge acquisition, representation, and use. Our goal is to generate a list that is necessary, such that all human-level agents must meet these requirements and that it is sufficient, such that meeting these requirement guarantees human-level behavior. The requirements we derive do not include criteria related to how well they model human behavior, nor the ease with which humans can create, debug, maintain, or extend agents developed in cognitive architectures. We also have not included criteria related to properties of the theory underlying the architecture, such as parsimony (Cassimatis, Bello, & Langley, 2008).

R0.   FIXED STRUCTURE FOR ALL TASKS
An individual agent adapts to its environment not through changes in its architecture but through changes in knowledge. Architectures should not depend on parameters that are tuned to improve performance on a new task; although parameters can be useful for introducing variation across agents. Architectures also should not allow escape to a programming language for task-specific extensions. The rationale for this requirement is that environmental regularities exist [C3] at time scales that approach or exceed the life of the agent [C8] that are worth capturing in a fixed architecture.

R1.   REALIZE A SYMBOL SYSTEM
The consensus in AI and cognitive science is that in order to achieve human-level behavior, a system must support universal computation. Newell (1990) makes the case that symbol systems provide both sufficient and necessary means for achieving universal computation; that is, a symbol system is capable of producing a response for every computable function. Possibly most important, symbol systems provide flexibility. In particular, they provide the ability to manipulate a description of some object in the world "in the head" without having to manipulate the object in the real world. Symbol structures also provide arbitrary composability to match the combinatoric complexity and regularity of the environment [C1, C3]. Thus, structures encountered independently can be combined later to create novel structures never experienced together [C5]. This generative capability is what we do when we combine letters or sounds to make new words, and when we combine words to make new sentences, and so on. Symbol systems also allow us to accept instructions from another agent and then use those

instructions later to influence behavior (interpretation) – providing additional flexibility and more generality – so that not everything must be programmed into a symbol system beforehand. In addition, symbols are required for communication that does not cause the meaning to be directly experienced by the agent [C4]. For example, striking someone directly causes an experience in another agent, while a verbal threat involves the transmission of symbols that require interpretation.

Requiring that the agent realize a symbol system does not imply that symbolic processing must be implemented directly via some symbolic knowledge representation. Neural and connectionist models can obviously support human-level behavior. Rather, this requirement posits that such approaches must implement symbol systems to some degree (Barsalou, 2005).

R2.   REPRESENT AND EFFECTIVELY USE MODALITY-SPECIFIC KNOWLEDGE
Although pure symbol systems support universal computation, they rely on modality-independent methods for representing and reasoning to achieve universality and complete composability. However, complete composability is not always necessary. Modality-specific representations can support more efficient processing through regularities [C3] in sensory processing [C6]. For example, some representations and associated processes for visual input have qualitatively different computational properties for image operations. Examples include rotation and inversion, and detecting and reasoning about spatial relations. For some tasks [C5] given limited computational resources [C7], modality-specific representations are necessary for achieving maximal efficiency, especially in tasks that require real-time performance [C2].

R3.   REPRESENT AND EFFECTIVELY USE LARGE BODIES OF DIVERSE KNOWLEDGE
The agent must be able to represent and use large bodies of knowledge. This wealth of knowledge that arises from the complexity of the environment [C1] and its associated regularities [C3], the variety of tasks the agent must pursue [C5], its complex interaction with the environment [C6], and the agent's continual existence [C8]. This knowledge is diverse, including memories of experiences, facts and beliefs, skills, and knowledge about other agents [C4].

R4.   REPRESENT AND EFFECTIVELY USE KNOWLEDGE WITH DIFFERENT LEVELS OF GENERALITY
The agent must represent and use general knowledge that takes advantage of the environmental regularities [C3]. The agent must also be sensitive to details of its current situation and its relationship to its tasks. These details are ubiquitous in complex [C1], dynamic [C2] environments where the agent can have many tasks [C5].

R5.   REPRESENT AND EFFECTIVELY USE DIVERSE LEVELS OF KNOWLEDGE
An agent must be able to take advantage of whatever knowledge is available. For novel tasks and environments,

its knowledge is limited, and even for familiar tasks and environments, its knowledge may be incomplete, inconsistent, or incorrect. If there is extensive knowledge available for a task, the agent must be able to represent and effectively use it. There are regularities in the environment worth knowing [C3], the complexity of an agent's limited sensing of its environment [C6], the complexity of its environment and tasks [C5], and limits on its computational resources [C7]. Planning systems often fail on this requirement. They often have a required and fixed set of input knowledge (the task operators and a declarative description of the goal). Without this knowledge, they are unable to attempt the problem. Further, if additional knowledge is available (such as knowledge about the likelihood of an operator leading to the goal), the planner is often unable to use it to improve behavior.

R6. REPRESENT AND EFFECTIVELY USE BELIEFS INDEPENDENT OF CURRENT PERCEPTION

The agent must be able to represent and reason about situations and beliefs that differ from current perception. Perceptual information is insufficient because perception is limited [C6], the environment is dynamic [C2], and there are regularities in the environment worth remembering [C3] for task completion [C5]. Thus, the agent must be able to maintain history of prior situations as well as the ability to represent and reason about hypothetical situations, a necessary component of planning. An agent that satisfies this requirement can make decisions based not just on its current situation, but also on its memory of previous situations and its prediction of future situations.

R7. REPRESENT AND EFFECTIVELY USE RICH, HIERARCHICAL CONTROL KNOWLEDGE

The agent must have a rich representation for control, because the actions it can perform are complex [C6]. Because of the dynamics of the environment [C2], and the multiplicity of the tasks playing out at multiple time scales [C5], some actions may need to occur in rapid sequence while others may need to execute in parallel. To keep up with a rapidly changing environment [C2] with limited computational resources [C7], the agent must take advantage of the structure of regularities of the environment [C3], maximizing the generality of the knowledge it encodes because of the complexity and variability of the environment and the agent's tasks [C1, C5]. This often means organizing knowledge about actions hierarchically. The agent can then decompose some of its actions into sequences of simpler actions, using the context of higher-level actions to constrain choices and reduce the knowledge required to generate action.

R8. REPRESENT AND EFFECTIVELY USE META-COGNITIVE KNOWLEDGE

In addition to the different types of knowledge discussed above, it is sometime necessary for an agent represent and use knowledge about itself and about its own knowledge (meta-knowledge). An agent invariably faces novel tasks [C5] where its task knowledge and/or computational resources [C7] are insufficient to determine the appropriate behavior given the environmental complexity [C1], but in which there are regularities it can take advantage of [C3]. In these situations, an intelligent agent can detect its lack of task knowledge, and then use meta-knowledge to acquire new task knowledge. An agent can use other types of meta-cognitive knowledge to set its own goals and to direct future behavior in preparation for tasks, events, and situations that it expects to arise in the future. This is in response to the characteristics listed above as well as to the fact that the agent exists beyond a single task or problem [C8]. The exact range of necessary meta-cognitive knowledge is unclear – some appears to be necessary, but complete meta-cognitive knowledge is not required, at least in humans. Humans do not always know exactly what they know and often only discover what they know when they are put in a situation where that knowledge is useful.

R9. SUPPORT A SPECTRUM OF BOUNDED AND UNBOUNDED DELIBERATION

At one extreme, the agent must be able to react with bounded computation [C5] for tasks with time constraints close to those of the dynamics of the environment [C2]. It cannot reason or plan from first principles for all tasks because of inherent limits to its computational resources [C7]. At its most primitive level, the absolute time to respond must be bounded by the environmental dynamics for some subclass of its responses. Reactivity would be sufficient if the agent's knowledge of the environment and other agents was complete and correct and encoded for bounded access below the level of dynamics of the environment. However, in general, that is not possible because of the complexity of the environment [C1], the diversity of tasks [C5] and the limitations on environmental interaction [C6]. Moreover, at the other extreme, when there are sufficient computational resources available relative to the dynamics of the environment and task, the agent should have the ability to compose novel responses based on its knowledge that takes advantage of regularities in the tasks and environment [C3]. This composition is the basis for planning and it takes time, but allows the agent to integrate its diverse and potentially large bodies of knowledge for novel situations [R1-R8]. In between these two extremes, the agent must balance the tradeoff between deliberation and reaction based on its knowledge of the situation.

R10. SUPPORT DIVERSE, COMPREHENSIVE LEARNING

An agent with long-term existence [C8] requires different learning mechanisms when exposed to diverse environments [C1] and tasks [C5] having complex interactions [C6]. Learning takes advantage of regularities [C3], some of which can be extracted from a single situation in which all of the information is available at the same time, whereas in others; the information may be spread across time. Although general learning mechanisms exist, they are invariably biased toward specific types of knowledge that are available to the agent in different ways and often at different time scales. Moreover, a general cognitive architecture should be able to learn all the types

of task-specific knowledge it represents and uses, a property we call the *learning completeness principle.* A significant component of our own research is to explore what types of knowledge different learning mechanisms can contribute to achieving learning completeness.

R11. SUPPORT INCREMENTAL, ONLINE LEARNING

An agent with long-term existence [C8] that is in a complex active environment [C1, C2] with regularities [C3] must learn and modify its knowledge base so that it can take advantage of the environmental regularities [C3] when they are available. Once the experience has happened, it is gone. Only the information that the agent itself stores while it is behaving is available to guide its future behavior. This is not to suggest that an agent cannot recall prior situations and perform additional analysis at some future time [R6]; however, some primitive learning mechanism must store away the experience for that future, more deliberative learning. Moreover, the mechanisms for storing and retrieving those experiences must scale as more and more experiences are captured. Incremental learning incorporates experiences when they are experienced.

One implication of this requirement (together with the need for large bodies of knowledge) is that new knowledge must be acquired at low, bounded computational cost in real time; learning should not disrupt the agent's ongoing behavior by significantly slowing overall processing and negatively impacting its ability to react to its environment.

Summary

In this paper, we outlined characteristics of the environments, tasks, and agents important for human-level intelligence and, from these characteristics, derived requirements for general cognitive architectures. Architectural designs following from meeting the requirements should support human-level behavior across a wide range of tasks, embedded in environment similar to the real world. Figure 2 summarizes the analysis described above. The figure highlights the dense connectivity between characteristics and requirements – no single characteristic is solely responsible for any requirement and no characteristic influences only a single requirement.

Many characteristics are necessary to derive most of their associated requirements because eliminating a characteristic allows for extreme simplification. Simple environments [C1] require only simple agents. There is no need to have large bodies of knowledge, no need for rich representations of action, and limited need to learn. An agent that only pursues simple well-known tasks [C5], or has unlimited computation [C7] can be much simpler than one that supports agents and tasks, in an environments with these characteristics. At the extreme is the requirement for task-relevant regularities [C3], which has universal impact because only with environmental regularities are knowledge, reasoning, learning, and architecture useful.

Discussion & Conclusion

The requirements we derived (R0-R11) define a rough design envelope for underlying architectures. However, the role of knowledge in agent development complicates attempts to match the achievement of specific requirements with specific architectural components. Behavior in an agent is the result of the interaction between knowledge and architecture; some requirements may be achieved through general knowledge combined with multiple architectural components. For example, many cognitive architectures do not have explicit architectural support for planning. Not including such architectural support simplifies these architectures, but requires encoding of knowledge representation(s) and algorithms for planning using architectural primitives. Achieving a requirement directly with the architecture allows for a more efficient implementation. Achieving a requirement in knowledge usually leads to a simpler architecture while providing more flexibility and the possibility of improving the capability through learning. This tension is analogous to

	C1 Complex Environ.	C2 Dynamic Environ.	C3 Task Regularities	C4 Social Environ.	C5 Complex Tasks	C6 Limited Interaction	C7 Limited Computation	C8 Long-term existence
R0 Fixed structure			X					X
R1 Symbol system	X		X	X	X			
R2 Modularity knowledge		X	X		X	X	X	
R3 Large bodies knowledge	X		X	X	X	X		X
R4 Levels of generality	X	X	X		X			
R5 Levels of knowledge			X		X	X	X	
R6 Non-perceptual represent.		X	X		X	X		
R7 Rich action representations	X	X	X		X	X	X	
R8 Meta-cognitive knowledge	X		X		X		X	X
R9 Spectrum of deliberation	X	X	X		X	X	X	
R10 Comprehensive learning	X		X		X	X		X
R11 Incremental learning	X	X	X				X	X

Figure 2: Connections between environment, task, and agent characteristics (C1-C8) and requirements (R0-R11).

RISC vs. CISC trade-offs in traditional computer architecture.

Our own hypothesis is that significant bodies of knowledge in combination with the architecture are required for many of the cognitive capabilities needed to achieve human-level performance. Examples include natural language processing, logical thinking, qualitative reasoning, and multi-agent coordination. However, the requirements listed above do not address what knowledge is necessary to support such capabilities, or how that knowledge is acquired and encoded. Thus, even if we create architectures that satisfy all of the listed requirements, we will still fall short of creating human-level agents until we encode, or the systems learn on their own, the content required for higher-level knowledge-intensive capabilities.

Even when we restrict ourselves to considering the requirements within the context of cognitive architecture independent of knowledge, it is difficult to evaluate the sufficiency of these requirements by examination alone. Many of the requirements are qualitative and vague, making them difficult to apply to existing architectures. For example, how do we judge whether an architecture supports sufficient levels of generality in its knowledge representations, or sufficient representations of meta-cognitive knowledge, or sufficiently comprehensive learning mechanisms? Thus, an important goal for future research in human-level agents is to refine these requirements as we learn more about the capabilities that are necessary for human-level behavior.

The current list or requirements emphasizes necessity and may be missing some yet to be discovered requirements that are needed to guarantee human-level behavior. These requirements may arise from interactions among the existing characteristics (C1-C8) or they may arise because of the existence of yet additional characteristics of agents, tasks, and environments that are relevant to achieving human-level intelligence.

Our own hypothesis is that one of the best ways to refine and extend these sets of requirements and characteristics is to develop agents using cognitive architectures that test the sufficiency and necessity of all these and other possible characteristics and requirements on a variety of real-world tasks. One challenge is to find tasks and environments where all of these characteristics are active, and thus all of the requirements must be confronted. A second challenge is that the existence of an architecture that achieves a subset of these requirements, does not guarantee that such an architecture can be extended to achieve other requirements while maintaining satisfaction of the original set of requirements. Usually there are too many potential interactions between architectural components to guarantee such an incremental approach. It is for these reasons that our own research is inspired by studies of human psychology. We know that the human cognitive architecture is sufficient for generating the behavior we seek from our agents, and if we build systems that capture the core functionality of components of the human architecture, it is more likely that we will avoid dead ends in cognitive architecture development. Independent of what approach is used to develop a cognitive architecture, we propose that exploring how different architectures address (or do not address) these requirements, both theoretically and empirically, is our best chance to advance our knowledge of how cognitive architecture can support human-level intelligence.

References

Anderson, J. R. (1990). *The Adaptive Character of Thought,* Hillsdale, NJ: Erlbaum.

Anderson, J. R. & Lebiere, C. L. (2003). The Newell Test for a Theory of Cognition. *Behavioral & Brain Science* 26, 587-637.

Barsalou, L. W. (2008). Grounded Cognition. *Annual Review of Psychology, 59*, 617-645.

Campbell, M., Hoane, A. J., & Hsu, F. (2002). Deep Blue. *Artificial Intelligence*, 134(1-2) 57-83.

Cassimatis, N.L., Bello, P., & Langley, P. (2008). Ability, Parsimony and Breadth in Models of Higher-Order Cognition. *Cognitive Science*. 33(8), 1304-1322.

Cohen, P. R., (1995). *Empirical Methods for Artificial Intelligence,* Cambridge, MA: MIT Press.

Gluck, K., & Pew, R., eds. (2005). *Modeling Human Behavior with Integrated Cognitive Architectures: Comparison, Evaluation, and Validation*. Lawrence-Erlbaum Associates, Matawan, NJ.

Jones, R. M., & Wray, R. E. (2006). Comparative Analysis of Frameworks for Knowledge-Intensive Intelligent Agents. *AI Magazine* 27, 57-70.

Laird, J. E. (1991). Preface for Special Section on Integrated Cognitive Architectures. *SIGART Bulletin*, 2(12), 123.

Laird, J. E., Wray, R. E. III, Marinier, R. P. III, & Langley, P. (2009). Claims and Challenges in Evaluating Human-Level Intelligent Systems, *Proceedings of the Second Conference on Artificial General Intelligence.*

Laird, J. E., Pearson, D. J., Jones, R. M., & Wray, R. E. (1996). Dynamic Knowledge Integration During Plan Execution. *Papers from the 1996 AAAI Fall Symposium on Plan Execution: Problems and Issues*, 92-98. AAAI Press, Cambridge, MA.

Langley, P., Laird, J. E., & Rogers, S. (2009). Cognitive Architectures: Research Issues and Challenges. *Cognitive Systems Research* 10(2), 141-160.

Legg, S., & Hutter, M. (2007). Universal Intelligence: A Definition of Machine Intelligence. *Minds and Machines*.17(4).

Newell, A. (1990). *Unified Theories of Cognition,* Harvard University Press, Cambridge, Massachusetts.

Simon, H. A. (1969). *The Sciences of the Artificial (First Edition),* MIT Press.