



AFRL-RH-WP-TR-2008-0080

Developing a Complete and Effective ACT-R Architecture

**John R. Anderson
Christian Lebiere**

**Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213**

January 2008

Final Report for July 2005 to January 2008

**Approved for public release;
distribution is unlimited.**

**Air Force Research Laboratory
Human Effectiveness Directorate
Warfighter Interface Division
Cognitive Systems Branch
Wright-Patterson AFB OH 45433-7022**

NOTICE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 88th Air Base Wing Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

THIS REPORT HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

AFRL-RH-WP-TR-2008-0080

//SIGNED//
JOHN L. CAMP
Program Manager
Cognitive Systems Branch

//SIGNED//
DANIEL G. GODDARD
Chief, Warfighter Interface Division
Human Effectiveness Directorate
Air Force Research Laboratory

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) January 2008		2. REPORT TYPE Final		3. DATES COVERED (From - To) July 2005 - January 2008	
4. TITLE AND SUBTITLE Developing a Complete and Effective ACT-R Architecture				5a. CONTRACT NUMBER FA8650-05-C-7254	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) John R. Anderson, Christian Lebiere				5d. PROJECT NUMBER 63123F	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER DRPA0507	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University 5000 Forbes Avenue Pittsburgh PA 15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Materiel Command Air Force Research Laboratory Human Effectiveness Directorate Warfighter Interface Division Cognitive Systems Branch Wright-Patterson AFB OH 45433-7022				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RHCS	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RH-WP-TR-2008-0080	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES 88 th ABW/PA cleared on 04 September 2008, 88ABW-2008-0007					
14. ABSTRACT The Carnegie Mellon University team focused on extending their current cognitive architecture, ACT-R, to show how visual imagery, language, emotion and meta-cognition affect learning, memory, and reasoning. The architectural modules are associated with the processes of the basal ganglia, anterior cingulate, amygdala, and the motor, occipital, prefrontal, and parietal cortices.					
15. SUBJECT TERMS Cognitive Architecture, Collaboration, Decision Making, Human Performance, Modeling and Simulation, Perception					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 38	19a. NAME OF RESPONSIBLE PERSON John L. Camp
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED			19b. TELEPHONE NUMBER (include area code)

THIS PAGE LEFT INTENTIONALLY BLANK

TABLE OF CONTENTS

1.	Introduction.....	1
2.	Overview of SAL	2
3.	Cognitive and Neural Synergies between ACT-R and Leabra in the SAL Architecture.....	5
3.1	The Procedural-Declarative Distinction	5
3.2	Reinforcement Learning for Procedural Processing	6
3.3	Declarative Learning and Processing Mechanisms	6
3.4	Visual Mechanisms	7
4.	Initial Concrete Implementation of SAL	7
5.	Future Evolution of the SAL Architecture	12
5.1	Procedural.....	12
5.2	Declarative	13
5.3	Motor.....	14
5.4	Visual	14
6.	Conclusion	14
7.	Supplementary Reports on Research Activities	14
7.1	Report 1: Models of Algebra Learning.....	14
7.2	Report 2: Attentional Blink Model.....	17
7.3	Report 3: Additional ACT-R Research Thrusts.....	19
7.3.1	<i>Spatial modules and navigation (AFRL, NRL, Rice)</i>	<i>20</i>
7.3.2	<i>Situation Awareness/Multimodal Integration/Episodic Memory (Drexel, RPI, AFRL).....</i>	<i>20</i>
7.3.3	<i>Motor modules/Robotic Embodiments (Rice, PSU, NRL).....</i>	<i>20</i>
7.3.4	<i>Language/Ontologies (Michigan, Xerox PARC, Alion).....</i>	<i>20</i>
7.3.5	<i>Social Interactions/Theory of Mind (Alion, Drexel)</i>	<i>21</i>
7.4	Report 4: Proposed Challenge Task - Assembly Required.....	21
7.4.1	<i>Challenge Task</i>	<i>21</i>
7.4.2	<i>Ecological Validity</i>	<i>21</i>
7.4.3	<i>Cognitive Requirements.....</i>	<i>21</i>
7.4.4	<i>Robotic Platform.....</i>	<i>22</i>
7.4.5	<i>Test Instantiation</i>	<i>22</i>
7.4.6	<i>Civilian Applications</i>	<i>23</i>
7.4.7	<i>Military Applications.....</i>	<i>23</i>
7.4.8	<i>Potential Platforms.....</i>	<i>23</i>
7.5	Report 5: SAL Integration Protocol.....	23
7.5.1	<i>ACT-R Protocol</i>	<i>24</i>
7.5.2	<i>Leabra Protocol.....</i>	<i>25</i>
7.5.3	<i>Glue Script.....</i>	<i>25</i>
7.6	Report 6: Directions for Further SAL Integration	26
7.6.1	<i>Motivation.....</i>	<i>26</i>
7.6.2	<i>Principle</i>	<i>27</i>
7.6.3	<i>Mechanisms</i>	<i>28</i>
7.6.3.1	<i>Learning Similarities.....</i>	<i>28</i>
7.6.3.2	<i>Learning Chunk Subtypes</i>	<i>29</i>
7.6.3.3	<i>Learning Intermediate Types</i>	<i>30</i>
7.6.3.4	<i>Integration</i>	<i>31</i>
8.	Related Approaches.....	33
9.	References.....	34

List of Figures

Figure 1:	Tripartite Architecture of SAL.....	3
Figure 2:	Cognitive Module Architecture of SAL.....	5
Figure 3:	The SAL Demonstration Model Architecture.....	9
Figure 4:	A sample view of a room with single object in the virtual environment.....	10
Figure 5:	The state of the Leabra network as it finishes correctly identifying the armor in the scene above.....	10
Figure 6:	BOLD response in four cortical regions for simple and complex equations.....	17
Figure 7:	Example of activation rise and fall for 7 output cells.....	18
Figure 8:	Data and model results of the attentional blink experiment.....	19
Figure 9:	Bridging ACT-R and Leabra representations.....	29
Figure 10:	Switch from ACT-R type to Leabra network and back.....	30
Figure 11:	Intra-module Organization of SAL Architecture.....	31

1. Introduction

In keeping with the direction of phase 1, we teamed with the Leabra research group. This group was led by Prof. Randy O'Reilly at the University of Colorado and included Brad Aisa, Thomas Hazy and Seth Herd of the University of Colorado, David Jilk of eCortex, Jonathan Cohen of Princeton, David Noelle of Vanderbilt University, and Todd Braver of Washington University.

The goal of the BICA program is to produce flexible systems that capture the power of human cognition – systems that can both adapt to new environments and be tasked with new instructions without reprogramming. To achieve that goal we are bringing together ACT-R, a high functionality architecture that already performs well on the goal of taskability, with Leabra, a high neural-fidelity architecture, that already performs well on adapting to new environments. The following are examples of high-end applications of ACT-R:

1. A system that can take instruction for a new domain of mathematics and reproduce student behavior (including brain imaging data) with no special programming.
2. A MOUT (Military Operations in Urban Terrain) system that simulates soldiers navigating in urban terrain and executing USMC combat doctrine while remaining adaptive to environment and opponent actions.
3. A driving model that can drive a car (in a simulator) and predict the degradation on performance that occurs when devices like cell phones or GPS systems are introduced.

The following are examples of high-end applications of Leabra:

1. A model of the visual neural pathways supporting object recognition, that learns to recognize objects despite wide variations in size, location, orientation, and context, performing at state-of-the-art levels based solely on existing general-purpose Leabra learning mechanisms.
2. A model of the prefrontal cortex and basal ganglia that can learn a wide range of complex working memory and cognitive control tasks based on trial and error learning, without any task-specific pre-configuration.
3. A model of the hippocampus and neocortex that can account for learning and memory data from over 20 different experiments on rats and humans, accurately capturing the effects of hippocampal damage and the complex division of labor between these neural systems in learning.

Although they are at different levels of description, ACT-R and Leabra have deep compatibilities that enable them to be synthesized into a new system that we will call SAL (Synthesis of ACT-R and Leabra). We have already put components of the two architectures together, and they interact successfully to solve a few interesting problems. In our work in Phase II we hope to go beyond just piecing together the best of both systems. We will combine the insights of each system into the components of the new SAL architecture. We will also develop more powerful means of interaction among the components.

In this document, we begin with a broad overview of the principles and neural basis of this new SAL architecture and how it relates to the existing ACT-R and Leabra architectures. Then, we

describe a concrete instantiation of a SAL model in a simplified version of the “Egg Hunt” task, operating in the Unreal Tournament simulation engine. This model shows how SAL can do things that neither Leabra nor ACT-R can currently do by itself, demonstrating the promise of the integrated architecture. This example also shows how we can easily extend the functionality of SAL without reprogramming the system. Finally, we outline a trajectory for further development of the SAL architecture as we move into Phase II, highlighting some of the key scientific and technical challenges and payoffs. It should already be quite evident that this synthesis represents an important new development in the field of cognitive architectures, and we are only at the very early stages.

The following are the publications to emerge from this research:

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.
- Lebiere, C., Gonzalez, C., & Martin, M. (2007). Instance-based decision-making model of repeated binary choice. In *Proceedings of the 8th International Conference on Cognitive Modeling*. New York: Psychology Press.
- O'Reilly, R., Lebiere, C., & Jilk, D. (2008). An explicitly pluralistic hybrid cognitive architecture. To appear in *Journal of Theoretical and Experimental Artificial Intelligence (JETAI) Special Issue on Pluralism and the Future of Cognitive Science*.
- Taatgen, N.A., Juvina, I., Herd, S., Jilk, D. & Martens, S. (2007). Attentional blink: An internal traffic jam? *Proceedings of the Eight International Conference on Cognitive Modeling* (pp. 91-96). New York: Psychology Press.

2. Overview of SAL

The ACT-R and Leabra architectures are both characterized by the attempt to account for a wide range of cognitive and neural phenomena using a small and therefore strongly constrained set of computational primitives, as contrasted with the predominant “one-off” and “grab-bag” cognitive models in the field. These architectures have been focused on largely complementary domains: Leabra on the neural mechanisms subserving the processing of individual stimuli and short sequences thereof, and ACT-R on more abstract, longer time-scale controlled cognition unfolding over minutes. These architectures are each arguably the most successful in their domain at rigorously accounting for a wide range of cognitive and neural phenomena, with each model providing detailed accounts of hundreds of distinct types of data.

Given their independent success at describing the human cognitive system, it is reassuring, and quite remarkable, that they have arrived at very convergent views of the overall cognitive architecture. This convergence is particularly significant given that the Leabra architecture is derived from more bottom-up neuro-computational constraints about the kinds of processing different parameterizations of a common neural substrate can support, while ACT-R is derived more top-down based on regularities and constraints present in human cognitive performance. This independent convergence provides a strong basis for confidence in the veracity of the emerging SAL architecture.

Both the Leabra and ACT-R architectures can be described at the most abstract level in terms of complimentary systems that are specialized for cognitively and neurally dissociable forms of processing. These dissociable neural systems form the basis of the SAL architecture, and can be categorized most broadly in a tripartite architecture, as previously documented by the Leabra team (see Figure 1):

- The **posterior cortex**, which performs basic sensory (e.g., visual, auditory, and somatosensory) processing (in the occipital and inferotemporal lobes) and motor processing (in the parietal lobe, which interacts strongly with posterior frontal cortical motor areas). This area is also critical for encoding higher-level semantic and declarative knowledge about the world, including many aspects of language and reasoning (in higher level association cortex in both temporal and parietal lobes).
- The **prefrontal cortex**, which is necessary for active maintenance of information and executive control of cognitive processing, and interacts closely with the **basal ganglia**, which is specialized for action selection and learning about which actions lead to reward or punishment. This system is critical for procedural processing and learning.
- The **hippocampus**, which is responsible for rapid learning of new information, often of a declarative (verbally-mediated) form (e.g., the location of a given object in an environment, the name of someone you’ve just met, or a new fact such as “the capital of Pakistan is Karachi”).




Figure 1: Tripartite Architecture of SAL. Human cognition is conceptualized in terms of the computational properties of distinct brain areas; each specialized for different incompatible forms of learning (e.g., rapid learning in the hippocampus vs. slow learning in the cortex). Red arrows represent top-down cognitive control (which results from interactions between Frontal Cortex and Basal Ganglia), while black arrows represent standard neural communication.

This tripartite, neurally-focused architecture can be decomposed into separable cognitive modules, which corresponds very closely with the ACT-R architecture, as shown in Figure 2. This mapping of functional modules onto neural structures is only approximate, particularly in the case of the imaginal and declarative modules. In the case of the imaginal module, while the control and maintenance are believed to be in the prefrontal cortex, the actual imaginal transformations seem to be performed in the parietal cortex, which is part of posterior cortex. In the case of declarative memory, while the hippocampus is a critical component, as in Leabra, much of the cortex can store declarative memory as well, and the prefrontal cortex plays an important role in controlling encoding and retrieval operations.

Before elaborating the cognitive and neural synergies between the Leabra and ACT-R architectures as captured in SAL, we can illustrate the general operation of this system in the context of SAL performing the “Egg Hunt” task in the Unreal Tournament (UT) environment (more details are provided below). SAL first hears a command like “find the armor”, which initiates a search through a set of rooms until it finds the target object, at which point it takes possession of the object. In this scenario, the different modules play the following roles:

1. The **aural module** holds a representation of the spoken sentence for processing.
2. The **goal module** maintains a representation of the activity (“find”) and the object (“armor”) throughout the episode.
3. The **declarative module** is accessed both to retrieve knowledge of the layout of the rooms and to maintain knowledge of which rooms have been searched.
4. The **imaginal module** is used to maintain a representation of the current room and the locations that have already been examined in the room.
5. A **motor module** is used to request movements from room to room and to orient to various objects in the room.
6. A **visual module** is used to represent the visual scene and identify objects.
7. The **procedural module** steps the agent through the tasks of planning its moves, performing the actions, and recognizing when the task objective has been achieved. It learns to improve its performance in future attempts, based on the success and failure of these actions.




Figure 2: Cognitive Module Architecture of SAL (where the broad tripartite architecture has been subdivided into finer grained separable cognitive mechanisms).

3. Cognitive and Neural Synergies between ACT-R and Leabra in the SAL Architecture

3.1 The Procedural-Declarative Distinction

The most notable area of convergence between ACT-R and Leabra is in the broad division of the cognitive architecture into procedural and declarative components. From this one distinction, many others follow, as elaborated in subsequent subsections. This distinction has clear cognitive and neural validity. People can possess abstract declarative knowledge of how to do something but be procedurally incapable of doing so (e.g., new drivers or golf players), and vice-versa (e.g., touch typists often cannot recall where the keys are located). Neurally, the basal ganglia are critical for initiating procedural actions, whereas the cortex and hippocampus support declarative knowledge.

In the Leabra framework, different types of processing are supported by the neural specializations present in the basal ganglia, compared with those present in the hippocampus and cortex. The basal ganglia system is strongly modulated by dopamine, which signals reward and

punishment information. Positive reward reinforces associated procedural actions, while negative feedback reduces the likelihood of producing associated actions. A similar, more abstract form of reinforcement learning is present in the ACT-R procedural system.

On the other hand, the neural properties of the hippocampus have been shown in the Leabra framework to be critical for the rapid learning of new arbitrary information without interfering with existing knowledge. Specifically, having a relatively few neurons active at one time (“sparse representations”) causes neural representations to separate from each other, minimizing interference. This rapidly acquired knowledge can, over time, be integrated into more overlapping, distributed representations in cortical areas, supporting the ability to draw sophisticated inferences and generalize to novel situations. The declarative system in ACT-R integrates both of these properties: new chunks of knowledge, encoded as combinations of existing chunks, can be rapidly formed; chunks that are used more frequently over time gain higher levels of activation and correspond to more expert knowledge; similarities can be defined between symbolic chunks to drive semantic generalization to related situations.

Although dissociable, the procedural and declarative systems interact intimately in any complete cognitive process. In ACT-R, the firing of productions is driven by the active contents of the declarative and other memory buffers, and the result of production firing is the updating of these buffers. In Leabra, the basal ganglia procedural system is tightly linked with the prefrontal cortex, which maintains task-relevant information in an active state over time. One of the primary functions of the basal ganglia in the brain is to drive the updating of these prefrontal active memory states. These prefrontal areas then influence activation states throughout the rest of the cortex via strong top-down excitatory projections. Each area of posterior cortex has an associated prefrontal area, with which it has strong bidirectional excitatory connectivity. Thus, we associate the buffers of ACT-R with these prefrontal representations of corresponding posterior cortical areas.

3.2 Reinforcement Learning for Procedural Processing

Both ACT-R and Leabra include reinforcement learning mechanisms to shape the procedural processing system. This form of learning uses success and failure information to shape the probability of selecting a given action in the future, and is dissociable from the form of learning that shapes cortical and hippocampal declarative representations. Although the detailed equations differ, there is considerable similarity between the two architectures in the computational principles underlying this learning, and both agree that the basal ganglia are its central neural locus.

3.3 Declarative Learning and Processing Mechanisms

Both Leabra and ACT-R make use of Hebbian-style learning mechanisms to modulate the strength of representations in declarative memory. Such learning mechanisms are based on the history of activation of the information stored in declarative memory; but critically, not on the success or failure of a particular action taken using that memory. This fact clearly dissociates these mechanisms from procedural reinforcement learning, and numerous cognitive experiments have validated this property of declarative memory.

In terms of processing information already stored in declarative memory, the concept of spreading activation is critical to both architectures. In ACT-R, activation spreads among declarative chunks in proportion to their associative strength. In Leabra, a similar activation spreading dynamic occurs, in that coarse-coded distributed representations in posterior cortical areas cause associated representations to overlap and share activation states.

3.4 Visual Mechanisms

With respect to vision we can distinguish between visual perception and visual attention. On the visual perception front Leabra offers a detailed theory and will be important for parsing the raster format that is anticipated for BICA Phase II. Efforts to incorporate direct perception into ACT-R have been limited to date, although there have been some proposals for extensions with systems like Robert St. Amant's "Segman." We will explore possible synergies in these approaches in the future.

ACT-R has a moderately functional overall theory of top-down attention that has been applied to vision, and which we will explore in connection with Leabra attentional processes. For example, Mike Byrne has been developing a "rational analysis of attention" on the premise that the system attempts to maximize the information uptake. This allows background biases such as a preference for rare objects to be combined with an immediate and explicit desire to find a red object. We will research how Byrne's equations for the salience of an object (a lot like activation equations in declarative memory) map to Leabra. Separately, Dario Salvucci has developed a rather sophisticated theory of eye movements for ACT-R. Salvucci's equations relate information uptake and probability of a movement to foveal distance, and also deal with the timing of saccadic programming. As Leabra visual processing is strongly dependent on successful foveation, a mapping of this theory onto Leabra control mechanisms could be quite powerful.

4. Initial Concrete Implementation of SAL

The SAL team has built a demonstration model representing a preliminary synthesis of the two architectures. Our goal was to anticipate the challenges we will face in implementing a truly integrated and embodied architecture for Phase II. This demonstration performs a simple version of the "Easter Egg Hunt" challenge suggested by discussions of the Phase II test problems; specifically, the SAL agent searches for a target object within a familiar environment. For this demonstration, we adopted a "first order" form of architectural integration, whereby one of the cognitive modules in an ACT-R model is replaced with a Leabra network. In this case, the Leabra model is capable of processing raw bitmap images in a way that the ACT-R model was not capable of doing; similarly, extant Leabra models are not capable of organizing problem solving behavior over a period of several minutes, as required to search for the target object in a complex environment. Thus, this SAL model represents a new level of functionality that goes beyond the capabilities of its constituent architectures. Given that this is the simplest form of integration, we are optimistic that much more interesting and powerful forms of cognition can be captured as our integration efforts develop further.

It is also worth noting that very little new work was required to make this model operational. We had already established a mechanism for ACT-R and Leabra interactions in preparation for demonstrations at the August technical meetings, including an attentional blink model and a model of the Haimson radar search task. In the attentional blink model, we combined the top-down control capabilities of ACT-R with the graded visual representations of Leabra, allowing us to account for aspects of psychological phenomena that neither architecture demonstrates individually. In the Haimson radar search task, ACT-R and Leabra shared a symbolic representation (a name) for objects of interest and interacted dynamically. Using an existing ACT-R model for searching environments and the Leabra model of visual object recognition, we simply adapted the ACT-R task instructions and trained the Leabra model on relevant visual stimuli.

In the demonstration (Figure 3), the SAL agent is embodied within an Unreal Tournament simulation environment consisting of three rooms containing three categories of objects. It is familiar with the environment in that it has access to navigation points and object location points in symbolic form. Further, it has been trained to perceptually identify the three object categories from a variety of viewing angles and distances. An operator instructs SAL to find the desired target via a typed command (“find armor”). SAL must then navigate the rooms (Figure 4); view and perceptually identify each object; and when it recognizes the desired target, navigate to it, and picks it up.

As noted, this combined model is implemented using the Leabra system as a perceptual front-end for ACT-R, by effectively replacing the ACT-R “Vision What” module (Figure 5); viewed conversely, ACT-R serves as a top-down control for the Leabra vision model. ACT-R must decide which navigation point to visit next and which objects to view from that navigation point. Upon selection of the object to view, ACT-R provides the digital image that is a snapshot of the view of that object from SAL’s current location. Leabra attempts to identify the object in the center of the image, and responds with its conclusion in symbolic form. If the object matches the specified target, ACT-R navigates to it, picks it up, and navigates back to the starting point; otherwise it selects a new object to view or a new location to which to navigate.




Figure 3: The SAL Demonstration Model Architecture. Navigation through the known environment is performed by the ACT-R component, and visual object identification is handled by a Leabra neural network component. The neural network replaces ACT-R's "vision: what" module; the ACT-R production system replaces Leabra's prefrontal cortex / basal ganglia element.

The ACT-R component of the model deals with the challenges of performing the search in an efficient manner. It plans where to begin searching, and uses its episodic memory and inhibition of return capabilities to remember where it has already searched. The Leabra component of the model was trained to recognize the object categories by repeated presentations of each object from a variety of perspectives and distances and in different room backgrounds. Its non-embodied performance on novel examples is 96%, when using a variety of backgrounds and object angles; simplifying the object angles and ensuring perfect foveation increases its performance to 100%.

This model provides a number of benefits:

- It illustrates a simple connection of the two architectures operating together, and demonstrates that it is possible to bridge the gap between their different levels of description.
- It serves as a first "embodiment" of a combined model operating in a simulation environment.
- It demonstrates that the system can adapt to new environments and be tasked with new instructions without reprogramming.




Figure 4: A sample view of a room with single object in the virtual environment.




Figure 5: The state of the Leabra network as it finishes correctly identifying the armor in the scene above. The network mimics the properties of the visual system, such as: a hierarchy of visual areas, with shortcut connections between them; representations of the image at several scales; and learning from visual experience.

By combining the two architectures in a straightforward way in the context of this task, the immediate future research efforts are obvious and compelling. There are a number of clear directions for improvement:

1. The combined model must use its “Visual What” system for navigation as well as object identification. It must learn to recognize doorways, room corners, obstacles, and other important navigational cues, and the control system must learn what to do with these inputs.
2. Rather than directly specifying a next location by its coordinates, the system must navigate using simpler locomotive operations, such as “move forward” or “turn right.” These operations could then be further elaborated to interact with more realistic effectors by a motor module based on previous Leabra work on learning motor movements. This was done in the ACT-R MOUT system, and that model can be adapted for SAL.
3. Rather than automatically orienting directly toward an object or navigational cue, the combined model must perform visual search in a way that is guided by both perceptual inputs and task demands. This will require a deeper and more sophisticated integration of ACT-R and Leabra, combining body movements and saccades with dual goals of learning the navigation options and attempting to find and identify objects in the environment. This deeper integration is likely to involve integrating ACT-R models of visual salience and eye movements with Leabra calculations graded activation information.
4. The system must be able to purposefully learn about new perceptual inputs. The Leabra perceptual system will need to recognize novelty in objects, actions, or outcomes, and the cognitive system will respond by making a goal of learning more about the novel element and taking active steps to do so. This approach is inspired by the theory that human infants act as scientists, making hypotheses and performing simple experiments to test them.
5. Since navigation points and object locations will no longer be available in symbolic form, the model’s internal representations of “where it has been” and “where it has looked” must instead be connected to perceptual representations of elements or items identified in the environment. This will also drive a deeper integration of ACT-R and Leabra, by requiring either a neurally-based episodic memory, symbolic chunks that have non-symbolic perceptual components, or both.
6. The control model must become more robust to perceptual and motor errors. Since the visual system will sometimes identify or search incorrectly, and the motor system may bump into obstacles or walls, the system must recover from these errors and choose a viable strategy for recovering and continuing. The ACT-R community has had some experience with achieving this robustness in the MOUT system and we believe this experience can be applied to SAL models.
7. The particular model needs to be capable of more flexibly learning from experience. In particular we hope to take advantage of the work in ACT-R on combined learning from instruction and observation. With respect to language processing we will need a system that can more robustly respond to instruction and this may require taking advantage of the more continuous and approximate representations that Leabra allows.

This demonstration project has been useful in revealing some challenges that arise in constructing a truly embodied and flexible cognitive architecture. These challenges highlight gaps in existing psychological theory, and suggest that building embodied architectures may be crucial to progress in the field. The project has also established the surprising power and flexibility available with even a simple combination of ACT-R and Leabra architectures. The two architectures are quite compatible despite two different levels of focus, suggesting that these two approaches have converged on a correct overall theory of cognition. The combined SAL architecture is capable of accomplishing tasks that are fundamentally beyond the reach of either ACT-R or Leabra operating in isolation. This success suggests that substantial progress has already been made in understanding the human mind in enough detail to replicate it.

5. Future Evolution of the SAL Architecture

In addition to the more task-specific challenges facing the SAL model outlined above, there are a number of more general and far-reaching issues that will shape the future evolution of the architecture. Overall, we hope to evolve this architecture from a relationship of mutual codependency between components of two separate systems to a more synthetic combination of the two systems. Below we consider the consequences of such a synthesis for some of the modules.

5.1 Procedural

The ACT-R production system represents a high functionality system that provides the needed control in SAL. Leabra's basal ganglia model represents a much more detailed system that is closer to the neural realities. In the synthesis we will constrain the production system to reflect that neural reality and in the process actually increase its functionality.

- (a) **Action Selection.** In the current ACT-R system, a production will fire only if it matches exactly. Exact matching is not tenable in light of either the biology or the needed functionality. Among the productions that do match in ACT-R, selection of the one to fire is made on the basis of learned reinforcements. This constraint that selection only begins to apply after matching is complete is again not tenable in light of either the biology or the needed functionality. The Leabra system would suggest that a dynamic threshold for matching emerges as a function of a competition among candidate productions – a less than perfect match can fire if it is the best thing currently available. This may be critical for allowing the system to learn new productions by reshaping old ones in new ways, as partial matches based on existing knowledge are co-opted and modified for new tasks.
- (b) **Production Learning.** One of the functionally powerful mechanisms in ACT-R is production compilation, by which new productions are created. In a typical example, one production will request a critical piece of information be retrieved from declarative memory, the information will be retrieved, and a second production will act on it. Production compilation replaces this with a single step in which the action is directly taken without retrieval. This is critical to the process by which instructions come to directly control behavior. From the Leabra perspective, this kind of learning involves the development of new representations in both prefrontal cortex and basal ganglia, and with

sufficient levels of repetition, may become independent of these systems and be encoded directly between the parietal and motor frontal areas. Thus, there are likely to be important shifts in the focus of neural activity over the course of production learning, which may produce important functional benefits in terms of reducing central capacity bottlenecks for highly practiced procedures. In the context of a navigation and target-search task, the ability to avoid obstacles and perform local navigation may become highly automated and free up more resources for visual search and higher-level route planning.

- (c) **Pattern Matching.** One of the critical questions is exactly how complex a pattern can be recognized in a single production cycle or a single pass through the basal ganglia. The simpler production rules in ACT-R can be realized in Leabra-like processes. However, Anderson has identified a more powerful kind of rule involving what he calls dynamic pattern matching which seems critical for human intelligence. In particular, they are critical for learning from instruction and demonstration – the typical means of social communication of knowledge. This can be supported through Leabra’s dynamic gating system. The current mechanism in ACT-R is only able to learn dynamic pattern-matching productions from other dynamic-pattern matching productions – it is not able to generalize explicitly matched rules to dynamically matching ones. Considering such mechanisms, in conjunction with Leabra capabilities such as dynamic gating, provides a good area for exploration.

5.2 Declarative

When ACT-R retrieves a chunk, it selects the most active one. The activation of a chunk reflects its past frequency of occurrence, its strength of association to the current context, and how well it matches the retrieval probe. All of these factors are combined to yield a quantity that reflects the likelihood that the chunk is the desired memory. A series of blending models have been developed in ACT-R for merging the contribution of multiple chunks into a single retrieved memory. In Leabra, there are actually two underlying systems supporting declarative retrieval: the hippocampus and posterior cortex. These systems have different characteristics. The hippocampus behaves more like ACT-R single-chunk retrieval, in that a single coherent chunk is typically retrieved, and it is highly sensitive to context and probe match. However, the posterior cortex can support overlapping distributed representations of multiple chunks at the same time, with each making a graded contribution to the overall memory retrieval process. This is more like the ACT-R blending models. We plan to integrate the Leabra and ACT-R perspectives into a more effective declarative memory.

Part of the effectiveness of declarative memory is the ability to incrementally absorb facts and adjust its generalization threshold to reflect the increasing knowledge base. In Leabra, that property arises from the gradual increase in the size of connection weights as a function of practice. In ACT-R however, while the absolute activation level of chunks increase with practice, their discriminability does not. We have experimented with modifying the ACT-R activation and partial matching equations to more closely reflect the computations in Leabra.

This work is an instance of a different sort of integration between ACT-R and Leabra where properties of one are absorbed into the other at a different level of abstraction. This approach is complementary to the integration strategy described earlier and indeed facilitates it.

5.3 Motor

The current motor actions are issued as discrete requests that are not guided by changing sensory information. A more Leabra-like implementation would have the specific parameters of these actions emerge as a result of a strong constraint-satisfaction process that takes into account many variables (precise location of things, speed of motion, slope of the floor, etc) to produce the desired goal.

5.4 Visual

The current visual system in ACT-R can only use top-down constraints to select objects to attend to. When these top-down constraints fail to find adequate guidance it is left to select among the objects randomly. There is new work within the ACT-R group on visual salience and how that can provide bottom-up influence. Merging bottom-up and top-down constraints will be critical in the anticipated BICA environments. Leabra provides guidance about how to coordinate the bidirectional top-down and bottom-up effects.

In addition to these module-specific considerations, there are numerous broad-based issues that are common across many different modules. For example, the way that learning is shaped by emotion and motivational states, which in turn are strongly influenced by social interactions, is a critical aspect of human cognition that SAL will need to address more directly. Some of this is captured in the existing reinforcement learning models in each architecture, but these issues really go beyond the confines of the procedural learning system, and shape representations and processing throughout the system.

6. Conclusion

We are just at the very beginning of what will hopefully be a long and fruitful process of breaking down longstanding barriers between different architectural “camps” in the field, and developing a truly synthetic and powerful understanding of the human cognitive architecture. The joining of forces represented by the SAL team already represents an unprecedented accomplishment of the BICA program, and we look forward to many more. We are confident that the BICA goal of developing a dynamically taskable, adaptive cognitive agent that can be deployed in a wide range of novel environments and task conditions is achievable with the synthesis of ideas represented by our team.

7. Supplementary Reports on Research Activities

7.1 Report 1: Models of Algebra Learning

The ACT-R group has been working on learning in the domain of algebra. They have taken this as a miniature for exploring the taskability issues that will arise in BICA Phase II. An

environment has been created for presenting problems and instructions to ACT-R or students much like a common environment is imagined for the BICA Phase II. Empirical explorations have included study of standard algebra taught to children and an isomorph appropriate for instruction to adults. Different studies have looked at learning from typical textbook instruction, learning from examples, and learning from exploration. These reflect the modes of learning that will be required of the BICA agent. Successful ACT-R models have been developed of learning from instruction and learning from examples that are able to predict the learning trajectories of actual students. However, the ACT-R group is still working to characterize the rather remarkable success that students have in learning from discovery. Particularly important for our ongoing efforts to model learning from exploration is understanding how the existing knowledge of students guides their exploration and enables them to interpret the outcomes of this exploration. Again it will be critical in BICA Phase II to be able to characterize the role of prior knowledge in learning about the environment.

One of the important outcomes from this effort was the realization that the current pattern matching in the ACT-R system was not powerful enough to enable processing the abstract relationships in instruction and example. Initial explorations revealed that we could capture the kind of learning students were doing with the more powerful SOAR pattern matcher but that the SOAR pattern matcher was so powerful as to be completely unrealistic biologically – being able to do exponential search in a single match. This led to a restricted version of pattern matching called dynamic pattern matching. One of the early results of interactions between the Leabra and ACT-R research group was the realization that the Leabra dynamic gating system provided a neural model for dynamic pattern matching in ACT-R. One of our future goals is to use this Leabra work to provide a more careful analysis of how dynamic pattern matching should be implemented in ACT-R, what its limitations are, and how to characterize its time costs relative to regular ACT-R pattern matching.

Another aspect of this research has been to look at the learning of algebra in an fMRI scanner. There is a mapping of the ACT-R modules onto specific brain regions and we have been able to use activation in these regions to inform our models of these tasks. As such it represents the potential role of fMRI in BICA Phase II.

The imaging experiment looked at the brain signature obtained while participants performed certain algebraic transformations. It manipulated two factors. One was whether the transformations were relatively simple algebraically (e.g., $3*4X=24 \rightarrow 12X=24$) or relatively more complex (e.g., $3*(4+x)=24 \rightarrow 12+3x=24$). The second was whether this was early or late in the learning. Figure 6 illustrates the results obtained from four cortical regions. The dotted lines connect the actual data and the solid lines are the predictions of the theory.

- (a) Figure 6a shows the response in the region of the motor cortex that controls hand movement and corresponds to the manual modules. Since it required more hand motions to execute complex transformations there is greater activation in this region for complex transformations. However, the number of hand movements did not vary early to late. While participants took longer earlier and so the response is stretched over a greater time span, the total area under the curves is the same.

- (b) Figure 6b shows the response in the region of the fusiform gyrus, which we have found to tap the high-level activity of the visual module which is all ACT-R represents. This shows both greater activation for complex equations and greater activation early. Like the motor region the greater activation for more complex equations is predicted because more encoding is required to enable a complex transformation. The ACT-R model does not really predict the learning effect. The solid lines in Figure 6b were only produced by the ad hoc reduction of encoding time for the late curves. This points to a place where a SAL model with Leabra-based visual learning would do better.
- (c) Figure 6c represents activity in a prefrontal region that many researchers have found to reflect retrieval from declarative memory, presumably because it holds controlling retrieval cues. The pattern we see here is one that has been observed in almost every study we have done manipulating complexity and practice – which is large effects of both. An effect of complexity is predicted because more retrievals are required for more complex equations and an effect of practice is predicted because the major dimension of learning in the ACT-R model is the drop out of some declarative retrievals and a reduction in the time to perform others.
- (d) Figure 6d shows the activation pattern in the anterior cingulate cortex, which is believed to reflect control operations. Like the other regions it shows greater activation for the more complex condition but the effect of practice is complicated. In the case of simple equations there is no effect on area under the curve but in the case of complex equations there is actually an increase in area under the curve with practice late in the performance of the transformation. This reflects the fact that mastering later transformations makes students sensitive to decisions about signs that they ignored earlier. These sign decisions come late in entering the transformation.

These data illustrate how imaging data can confirm the theory in ways that range from expected to surprising (e.g., Figure 6a to Figure 6c to Figure 6d) and at the same time indicate places where the analysis needs to be modified (i.e., Figure 6b).




Figure 6: BOLD response in four cortical regions for simple and complex equations (early and late in practice). Dotted lines connect actual data and solid lines are predictions from ACT-R modules.

7.2 Report 2: Attentional Blink Model

To demonstrate that the SAL architecture can model behavioral phenomena that are difficult or impossible to model with either ACT-R or Leabra alone, we developed a SAL model of the Attentional Blink phenomenon.

In the experiment associated with Attentional Blink, participants are presented with rapid streams of 20 characters (at a presentation rate of 100 ms/character), most of which are digits (distracters), and some (0, 1 or 2) are letters (targets). The goal for the participant is to identify and report the targets (the letters). In streams with two targets, response accuracy differs depending on how many distracters are presented between the two targets. This distance is called the lag, where lag 1 means no distracters in between the targets, lag 2 one distracter, etc. If the lag is five or more, accuracy on both targets is the same, around 80%. When the lag is less, accuracy on the second target is worse than on the first target with one exception: when the lag is one, meaning the targets are immediately sequential; accuracy on both targets is again identical. However, in that case participants make a different error: they sometimes report the two targets in the wrong order.

To model the task we used a SAL prototype that is architecturally similar to Figure 3, where the “Vision what” module of ACT-R was replaced by a modified Leabra vision model. This modified Leabra model does not reset its activations between stimuli, and its output is a set of graded activation values at each time step, rather than a final symbolic determination. Having been previously trained on the character set used in the human experiment, the Leabra network was presented with the stream of characters that the participants also perceived. Due to the speed of the input, the network was not always able to reach peak activity for a particular classification, and would sometimes be in a transition in between two classifications when queried by the ACT-R part of the model. Figure 7 shows a sample graph of output activation of the vision module.




Figure 7: Example of activation rise and fall for 7 output cells (for characters 2, 3, 5, 7, 8, 9 and P). Vertical lines indicate where a new character was presented to the network. Note that it takes around 10 cycles (20 cycles maps onto 100 ms) before a new character produces a rise in activation in the output layer. If the network is sampled at a moment when multiple characters are active (e.g., 2 and 9 at cycle 115) it will pass all candidates on to ACT-R.

The Leabra component not only provides a realistic account of how a letter is processed, but also explains why the order of targets is sometimes mixed up when the two targets are right after each other: if the network perceives two characters at the same time it has no way to deduce the order, so has to make a random guess (participants report that it seems like the two letters are superimposed on each other).

Once Leabra has recognized characters, ACT-R has to classify and possibly memorize them. The ACT-R model assumes that targets are stored in the imaginal buffer. In order to determine whether something is a target, its category (letter or digit) has to be retrieved from declarative memory first. The rapid presentation rate puts a heavy strain on the capacity of the architecture to keep track of all the input. Once a first target has been detected, the additional task of storing it creates additional load on the system, creating a ripple effect (similar to a traffic jam) on processing further down the line. It turns out that the peak of this “cognitive traffic jam” is 200 to 300 ms after the presentation of the first target, which is exactly where the attentional blink effect peaks. Figure 8 shows the comparison between model and data.




Figure 8: Data and model results of the attentional blink experiment. Left: proportion of correct T2 responses, showing the attentional blink effect. Right: proportion of switch errors, where both targets are named correctly but in reverse order.

The model of attentional blink combines the strengths of the two architectures: the fine-grained perceptual capabilities of the Leabra architecture can explain the reversals at lag 1 (Figure 8, right), and ACT-R's serial processing constraints on individual modules can explain the blink effect (Figure 8, left).

7.3 Report 3: Additional ACT-R Research Thrusts

While for phase I the ACT-R part of this effort has been concentrated at CMU, it also included a number of consultants who have been longstanding members of the ACT-R community and are envisioned to play a much more substantial role in Phase II to help us tackle the difficulty and complexity of the tasks ahead. Those consultants involve 12 people distributed over 8 institutions:

Kevin Gluck and Glenn Gunzelmann – Air Force Research Laboratory Mesa

Greg Trafton – Naval Research Laboratory Washington DC

Wayne Gray and Michael Schoelles – Rensselaer Polytechnic Institute

Dario Salvucci and Frank Lee – Drexel University

Michael Byrne – Rice University

Frank Ritter – Pennsylvania State University

Peter Pirolli and Stuart Card – Xerox Palo Alto Research Center

Richard Lewis – University of Michigan

In addition, Michael Matessa at Alion Science & Technology has also been involved as a subcontractor.

During phase I, those consultants have investigated architectural issues of central interest to the goal of BICA and in particular to the challenge tasks and environments that have been discussed. To avoid scattering our efforts and to bring significant resources to bear on those problems, they

have organized into larger teams focused on a small number of research questions that are both of fundamental scientific interest and directly relevant to the BICA challenge tasks. The teams considered constraints ranging from functional performance to neuroscience evidence in evaluating architectural designs for the hybrid SAL architecture. Those research questions, together with the teams focused on them, are:

7.3.1 Spatial modules and navigation (AFRL, NRL, Rice)

This research thread focuses on developing a new spatial module for SAL. The current ACT-R visual module only provides a relatively flat and static view of the world and significant enhancements are required to operate effectively in the 3-dimensional dynamic world of the challenge problems as well as to reflect a more complete picture of the neuropsychological evidence regarding human spatial abilities. This new module would work in concert with the improved and integrated versions of the visual and motor modules and would provide competencies including support for frames of reference, mental imagery, magnitude estimations, spatial transformations and navigation.

7.3.2 Situation Awareness/Multimodal Integration/Episodic Memory (Drexel, RPI, AFRL)

This research thread focuses on developing memory competencies that integrate external experiences across time and across sensory modalities. The current ACT-R declarative memory represents information of different points in time or sensory modalities as independent chunks without any links to a common integrated picture of the environment associated with episodic memory capacities. This research effort has explored ways of maintaining robust continuous situation awareness in a dynamic environment by developing memory representations and processes that integrate experience across time and sensory modalities. These do not take the form of a new architectural module but instead consist of an elaboration of the existing memory representation and retrieval processes.

7.3.3 Motor modules/Robotic Embodiments (Rice, PSU, NRL)

This research thread focuses on developing motor modules consistent with the simulated embodiment of the BICA agent, including simplified lower-body, articulated upper body and complex manipulators. The current ACT-R motor module is limited to two-handed keyboard actions and needs considerable generalization to support the proposed embodiment. This research thread also investigates and develops architectural assumptions such as direct perceptual-motor module links and continuous control of motor actions. Because of the importance on learning and development and the difficulty in programming controllers for complex activators, this thread also focuses on learning mechanisms for the new motor module in concert with cognitive learning mechanisms such as production compilation. This coordinated learning across architectural modules constitutes a significant challenge and a major innovation.

7.3.4 Language/Ontologies (Michigan, Xerox PARC, Alion)

This research thread focuses on developing language capabilities to support the dialogue with teacher(s) and fellow BICA agents in the proposed task environment. These include support for understanding instructions, language acquisition similar to human development, and support for the role of language in cognition. This thread also focuses on representational issues, such as the adoption of common ontologies for integrating models.

7.3.5 *Social Interactions/Theory of Mind (Alion, Drexel)*

This research thread focuses on the social capabilities involved in interacting with teacher(s) and BICA agents in the task environment, including inferring and understanding the beliefs, intentions and actions of others, sharing mental pictures, and working in cooperation and competition. This thread also focuses on understanding the role of mirror neurons in learning from demonstration and imitation. This will require integration both with the visual/spatial modules to provide a representation comparable to that provided by mirror neurons and with the motor modules to allow the closing of the loop from visual input to motor action and back. This could take the form of a separate intentional module, of direct connections between perceptual and motor modules, of cognitive skills modulating those interactions, or any combination of the above

All of these research directions involve interaction between with the teams (e.g. the second topic will need to take as input the output of the spatial modules developed under the first topic) as well as with the CMU team, requiring continuous attention to integrating these efforts conceptually (as well as in software) into a coherent framework.

7.4 Report 4: Proposed Challenge Task – Assembly Required

Our team developed and submitted the following description of a potential challenge task targeted to the BICA objectives:

7.4.1 *Challenge Task*

Develop a cognitive robot that can take written instructions consisting of text and/or graphics and assemble from constituent parts any assembly, device or machine allowed by its effectors. The work will require the use of tools. The robot is allowed to supplement the written instructions by asking any question to an independent supervisor. The work can be performed independently or by a group of robots working cooperatively.

7.4.2 *Ecological Validity*

Using tools and building devices is an essential human activity. The exponentially increasing trajectory of human evolution can be traced directly to those activities. The same applies to children. Children manipulate blocks and other small objects from an early age and their ability to construct increasingly complex assemblies (stacking blocks, building legos, assembling magnetix, meccano and erector sets, etc) seems to parallel their general cognitive development.

7.4.3 *Cognitive Requirements*

This challenge task covers a broad range of possible instantiations (basically walk into any store and pick up any toy or device that requires assembly) and requires a broad range of integrated cognitive abilities such that a pure engineering approach is unlikely to be effective. Those cognitive capabilities include but are not limited to:

- Language: reading and understanding complex instructions (an understated skill that children spend much school time acquiring and perfecting); asking questions of and interacting with supervisor.

- **Spatial Abilities:** object recognition to identify pieces and parts of assembly, perspective-taking to adopt relative frames of reference, spatial planning to perform manipulation steps, navigation to retrieve parts or tools scattered across the environment.
- **Reasoning:** instruction steps can be made arbitrarily abstract and complex to require significant problem-solving, including planning, scheduling and mental simulation.
- **Social Interaction:** communicate and work with other robots to perform actions that the individual robot cannot perform by itself, e.g. planning complex actions and manipulation that requires more effectors or strength than available to a single robot.
- **Motor Skills:** learn new actions required by a new assembly domain. Learn to use new tools. Domain-specific effectors could be added to the robot to allow it to manipulate objects that are not tractable with its default effectors (e.g. magnetic “hands” to manipulating small magnetic objects). While fine motor skills might cause interesting problems, they should be limited to avoid emphasizing low-level robotic issues.

7.4.4 Robotic Platform

Finding the right robotic platform for this task is a major challenge. Requirements include a reliable and widely available (probably commercial) platform, a compatible simulation environment for development purposes (e.g. PlayerStage, Webots, Microsoft Robotics Studio) and a competent set of manipulators. Mobility is desirable but the best arms and hands tend to be stationary for reasons of weight, power and/or stability.

Potential high-end platforms might have included Sony’s QRIO humanoid robot (no longer in production) and Honda’s ASIMO (expensive and limited availability), but some humanoid research robots might still be a possibility. At the low end would be toys like the RoboSapien v2, a two-foot tall robot retailing for only \$200, which would make large numbers affordable to research groups for both parallel development and team-based challenges. It is remote controlled by default but has been adapted with camera and PDA controllers to play robot soccer by research teams in Germany. In the middle of the range in terms of price and sophistication are non-humanoid robots that have some mobility and manipulation capacities, such as the Neuronics RCS5000 or iRobot Packbot.

7.4.5 Test Instantiation

An actual challenge test would require both a specific robotic platform and a specific assembly task to accomplish, consisting of a set of pieces, written and/or graphic instructions and target assembly. The same robotic platform could be used for a number of different assemblies. The set of pieces and the target configuration will be physically limited by the robot manipulator and the set of tools available. In other words, it should be possible for a human controlling the robot(s) remotely to complete the assembly.

Performance on the task can be measured in a number of ways. The coarser measure would be proportions of correct assemblies performed and time to complete them. Finer behavioral measures would include techniques and strategies used to perform the assemblies, such as use of

tools and order of operations. Eye-tracking data could conceivably be collected and perhaps even fMRI if human subjects can control the robot(s) remotely from inside a magnet.

Learning could be emphasized during the test in a number of ways. The most straightforward would be to repeat the test multiple times and judge the improvement in the performance measures above across iterations. More subtle measures of learning transfer could also be collected, including transfer across different assemblies and positive (speedup) and negative (errors) transfer across similar subassemblies differing in some respect (e.g. reverse symmetries).

7.4.6 Civilian Applications

Assembly plays a prominent role in both high-technology industries (e.g. computer, automobile) as well as in low-cost industries in developing countries (e.g. Chinese workshops). Despite some high-profile successes of industrial robotics (e.g. robots painting cars or putting on windshields), the impact has been relatively low, partly due to cost and political considerations, but also partly due to the lack of flexibility of robotic platforms. An inexpensive, fully taskable robot could achieve in the industrial world the kind of impact that cheap domestic robots (e.g. Roomba) are starting to make in the commercial world.

7.4.7 Military Applications


Robotic devices have started to appear in a number of tasks from UAVs to IED mitigation. However, their roles tend to be limited to dangerous situations and require considerable human supervision. While public attention in military deployments is focused on combat roles, the number of troops that must be deployed to support combat units is often considerably larger (almost 10-to-1 in Iraq). Robotic devices that could perform even some of those support roles could serve a key role in reducing manpower requirements in dangerous deployments.

7.4.8 Potential Platforms

PlayerStage: <http://playerstage.sourceforge.net/>
Webots: <http://www.cyberbotics.com/products/webots/>
Microsoft Robotics Studio: <http://msdn.microsoft.com/robotics/>
QRIO: <http://www.sony.net/SonyInfo/QRIO/>
ASIMO: <http://world.honda.com/ASIMO/>
RoboSapien V2: <http://www.robosapienv2online.com/>
RoboSapienSoccer: <http://www.nimbro.net/rs/index.html>
Neuronics RCS5000: http://www.neuronics.ch/cms_en/web/index.php?id=201
iRobot Packbot: <http://www.irobot.com/sp.cfm?pageid=138>
iRobot Roomba: <http://www.irobot.com/sp.cfm?pageid=95>

7.5 Report 5: SAL Integration Protocol

The integration of ACT-R and Leabra in Version 0.8 (such as used for the Unreal Tournament integration demo described in the main body of the report) consists of two different file transport/semaphore protocols connected via a glue script. Both protocols rely on the fact that file rename and file remove operations are atomic in the operating-system. The diagram below shows the overall organization:



Note also that both protocols can be used in more general form than described here. This document merely describes the protocols as they are used in the current implementation of the Glue Script.

All three modules (ACT-R, Glue Script, and Leabra) run at the same time and have wait modes where they alternate sleeping and looking for input. Each module exits upon receiving the appropriate completion signal as described below. In theory the modules can be loaded in any order, although in our experience this is not robust.

7.5.1 *ACT-R Protocol*

The ACT-R protocol uses three files for requests and three analogous files for responses. The files are:

	Command	Response
Data	Stream_cmd.dat.txt	Stream_res.dat.txt
Temp	Stream_cmd.tmp.txt	Stream_res.tmp.txt
End	Stream_cmd.eof.txt	Stream_res.eof.txt

“Stream” would be replaced with the name of the data stream or interface, e.g., “Haimson.” The Data file is the container for the actual data for the request or response; the End file is an indicator that there are no more requests or responses, and the Temp file is a temporary container for the next data until the software receiving the request or response is ready to receive it.

More specifically, when requesting a visual identification from Leabra, the ACT-R model places content for a request in the Command Temp file and waits until the Command Data file does not exist. When the Command Data file no longer exists, ACT-R renames the Command Temp file to be the new Command Data file. ACT-R waits until the Response Data file exists, and when it does, it processes this response and then deletes the Response Data file.

When the final request has been generated, ACT-R creates the Command End file. Its existence signals that no more Command Data files will be created; however, there may or may not be an unprocessed Command Data file in existence at this point. Symmetrically, when ACT-R sees a Response End file, it knows there will be no more Response Data files created, beyond the one that may or may not currently exist. Note that the content of both End files is irrelevant.

The format of the Command Data file for object identification tasks is:

```
filename llx lly urx ury
```

where filename is a path specification to a JPG image, and the other values represent a rectangular region of the image using a 0.0 to 100.0 Euclidian coordinate system with independent x and y scales. The file should have only one line and should end in a newline.

The format of the Response Data file is:

```
size answer cycles
```

where size is the number of fields in the response (in this case always 2), answer is the name of the object identified, and cycles is the number of Leabra settling cycles required for the identification. This file also should have only one line and end in a newline.

7.5.2 Leabra Protocol

The Leabra protocol relies on a single Image Batch file for requests and produces its responses by appending sequentially to a Trial Log file. The Image Batch file is named Stream.bat.oia, where “Stream” is the name of the data stream. Leabra waits until an Image Batch file exists, and when it does it reads it, deletes it, and proceeds to process the indicated input. Conceptually Leabra updates the log file upon completion of processing the input, but see below for more detail on the response sequence.

The format of the Image Batch file is:

```
ground-truth filename Object llx lly urx ury
```

where ground-truth can indicate the correct answer or a ? if the answer is not available, filename is the path to the image file to be processed, “Object” is simply that text, and the rest of the values are a rectangular region of the image using a 0.0 to 100.0 Euclidian coordinate system with independent x and y scales (i.e., the same as in the ACT-R protocol).

If the Image Batch file exists but is empty, this indicates that there will be no further inputs.

The Trial Log file is a tab-delimited file containing one result on each row. The number of Leabra cycles used in processing is in column 7, and the result of the identification is in column 8. Leabra does not actually update the Trial Log file until just prior to processing the next input request. Thus, to perform synchronous identifications it is necessary to run an additional input through Leabra. The simplest way to accomplish this is to run each input twice.

7.5.3 Glue Script

The glue script is a relatively brief Bash script that coordinates between these two protocols. At a high level, it:

1. Cleans up after any previous runs
2. Waits for a Command Data file from ACT-R
3. Upon receiving one, translates it into an Image Batch file for Leabra
4. Gives Leabra the same Image Batch file twice

5. Watches for changes to the Trial Log file and when it is updated, produces a Result Data file for ACT-R
6. Upon receiving a Command End file, it produces an empty Image Batch file for Leabra and a Response End file for ACT-R

It performs these operations while obeying the semaphore rules for each protocol; e.g., it does not produce a Response Data file or an Image Batch file unless the file does not already exist, and it deletes the files as needed and appropriate.

7.6 Report 6: Directions for Further SAL Integration

7.6.1 Motivation

Explorations in Phase I have focused on a modular integration of ACT-R and Leabra. There were a number of good reasons for it: a natural complementarity between the strengths and weaknesses of each architecture at the module level, the ability of getting running demonstrations off the ground quickly, and a good fit between the architectural constraints of the low-bandwidth inter-module synchronization and the computational requirements of inter-process communication. However, there are also good reasons for at least exploring a deeper integration of the two architectures. The most obvious one is that it coincides with DARPA's vision of the kind of ground-breaking achievements. The deeper one is that the architectures have complementary strengths and limitations that are only partially addressed by a module-level integration.

Leabra's primary strength is its considerable learning flexibility. Its primary limitations are its computational complexity (in terms of operations per cycle) and its learning complexity (in terms of cycles per learning episode). Leabra's computations might scale better in this respect than other algorithms such as back-propagation, and certain components of the architecture (e.g. the hippocampus model) are specifically designed for that purpose, but it still remains that even on small tasks the computational requirements are significant, and will likely scale poorly as tasks get more complex, possibly preventing real-time computation even though some tasks (e.g. robotic embodiment) might require it, and/or requiring extremely long developmental runs that would limit architectural explorations.

ACT-R's strength, in addition to its generality and behavioral accuracy, is its computational tractability. ACT-R models run many times faster than real time even on desktop or laptop hardware and its algorithms (e.g. declarative retrieval and procedural matching) parallelize naturally if the need for special-purpose hardware arose. Its primary limitation is the fact that while it is very flexible at learning new skills and facts, it cannot learn new representations and therefore depends on the modeler to specify the features of the skills and memories that it learns. In relatively small models where the external features of the stimulus are few and simple enough to serve as direct encoding of the experiences (e.g. many instance-based models developed in ACT-R) or when a good set of features is known and can be encoded in the model (even if large as long as combinatorial (de)composition works, e.g. the ACT-R backgammon model), that is not a key shortcoming. But in complex environments where deep internal statistical regularities have to be discovered (e.g. Granger-Hawkins-type processes of learning layers of representations), then something else is needed.

One can see from these strengths and limitations the basis for a tighter, more complementary integration: use Leabra for its ability to learn internal representations and use ACT-R for its tractability, generality and behavioral accuracy. The key of course is to find a way for the two paradigms to communicate with each other such that they can both perform their key functions without being limited by the shortcomings of the other or the interface between them, a problem that has historically dogged attempts at developing hybrid systems (e.g., Wermter & Sun, 2000).

A couple of key clarifications:

- This view of a deeper, tighter integration of the two architectures is not a search for an intermediate level of representation but instead a way of combining two layers of abstraction that have each been extremely successful individually. Thus this proposes to build on the successes of ACT-R and Leabra rather than replace them with something entirely new and untested, an approach that at best would require years to develop based on experience with past architectures. Moreover, this can be viewed as a relatively straightforward intra-module extension of the already existing inter-module integration.
- For sake of specificity, this view will be described specifically for the long-term declarative memory module, but its principle should also apply with some variations to the procedural, perceptual/motor and other modules. Indeed, we will suggest that it can also serve as a basis for learning representations across modules.

7.6.2 Principle

The basic principle beyond this integration of ACT-R and Leabra is to extend the modularity approach currently used between existing modules of the architecture(s) to apply within modules as well. As mentioned previously, this will be discussed here primarily in the context of long-term declarative memory, but should be applicable to other modules as well as to inter-module communication.

There are a number of arguments for internal modularity from each relevant viewpoint. From the environmental (task) standpoint, (de)composition of knowledge and skills is a well-established fact. We are able to learn new knowledge on top of previous knowledge, but newly acquired knowledge seldom if ever interfere with previously unrelated knowledge, and never catastrophically. For ACT-R, chunk types provide a natural decomposition of long-term knowledge into exclusive hierarchical subsets that provide both efficiency of access and insulation against interference. In Leabra, while fully distributed representations suggest a completely connected network, in practice either sub-modules can be adopted by distinct representational subspaces, or orthogonal, non-overlapping representations for unrelated concepts will develop that will accomplish substantially the same effect. In the brain, hyper-columns provide that kind of sub-system modularity within the flat cortical organization. From a functional point of view, modularity is required to be able to achieve model composition. Otherwise, one could not assemble a model of a complex task from independently developed models of component skills, but instead one would have to develop increasingly complex models tabula rasa each time, an increasingly unfeasible task.

If modularity of knowledge and skill is a given, then one can apply the same specialization argument within each module as has been applied between modules. However, in light of the complementary strengths of weaknesses of each architecture described in the previous section, instead of permanently dedicating a piece of functionality to one particular architectural level, either paradigm could be applied at any particular point in time depending upon the functionality required. In particular, Leabra could be applied when performance is insufficient and new internal representations must be learned to capture the regularities of that type of knowledge. Conversely, ACT-R could be used when an adequate representation has been either specified or learned to provide efficiency and behavioral accuracy. We will discuss in the next section the technical mechanisms that would make this complementary integration of the two architectures work.

7.6.3 Mechanisms

7.6.3.1 Learning Similarities

At the symbolic/procedural level, generalization in ACT-R is accomplished through production variabilization, which has been shown to correspond to directing the flow of information between brain areas. At the subsymbolic/declarative level, generalization is accomplished through similarities between chunks akin to dot-product between distributed representations in neural networks. However, ACT-R lacks a mechanism for learning those similarities. They have typically been set to reflect external domain structure (e.g. numbers, physical stimuli, etc) or according to algorithms such as LSA reflecting commonalities between more abstract concepts (e.g. words, etc). An essential objective of ACT-R/Leabra integration would be to use Leabra to learn ACT-R representations (i.e. similarities), providing a direct representational bridge between the two frameworks.

The most tractable way to achieve that representational link is incrementally on the type modularity basis discussed in the previous section. Figure 9 illustrates the connection between the two representations. The ACT-R representation below the dotted line describes the representation of one chunk type. It consists of a number of chunks of that type, each composed of a number of slots. Assuming the chunks used as component values in those slots already have known distributed equivalents in Leabra (either by a previous application of this mapping or simply as part of a bootstrap set specified by the user), the goal is to learn the representational similarities between the chunks built from those component values. The chunk slot values in the ACT-R type representation can then be converted into distributed representation in Leabra (above the dotted line), then fed into a hidden layer charged with representing the binding of those values together and thus storing those chunks in memory. Each chunk of the given type can be presented that way, basically feeding in the slot values using the existing ACT-R/Leabra representation mapping, and then learning the distributed representation of chunks of those types that achieves the best memory generalization. The similarities between the new ACT-R chunks can then be directly inferred between the dot-product of the representation of those chunks in the hidden Leabra layer. That mapping between ACT-R chunks and their Leabra distributed representations can then be added to the previous mappings, ensuring the bootstrapping to the next step when these chunks in turn are used as components of other chunks.




Figure 9: Bridging ACT-R and Leabra representations.

If one allows the representation of the slot component values in the input layer to reflect the learning of the new chunks in the output layer instead of assuming them fixed, then an updated version of those representations could be added to the ACT-R/Leabra mapping. This is similar to the LSA concept of having the representation of a word reflect its context, in this case the chunk that it is a part of. It would also have the advantage of relaxing the bootstrapping requirements by allowing members of a new chunk type to be used as slot values in that type itself rather than requiring a strict type hierarchy.

As an aside, it might be interesting to compare Figure 9 above to Figure 4 in Norman & O'Reilly, 2003. There seems to be a rough upside-down correspondence. The Leabra hidden layer would correspond to EC, with the slot representation input layer corresponding to other cortical association areas. ACT-R takes the role of hippocampal representation, with the slot representation corresponding roughly to area CA3 and the chunk representation to area CA1. The role of the dentate gyrus in achieving the transformation in representation between the two systems is similar to that of the mapping between Leabra distributed representations and ACT-R chunk similarities. The input-output relationships between the two systems are reversed however, and it is not clear what can be gained from pursuing this thread.

7.6.3.2 Learning Chunk Subtypes

ACT-R has a single-inheritance type hierarchy. Thus, one could view a type as a subtype of a

common type chunk, with every layer down the type hierarchy holding a narrower set of chunks. However, there is no mechanism for learning chunk types. The learning of representations described in the previous section provides an opportunity for expanding the type hierarchy as well. Chunks are mapped onto a continuous representation space in Leabra. If they form natural representational subdivisions of that space, a clustering/partitioning algorithm can detect the separation and create multiple subtypes for the original chunk type. An instance is storing arithmetic facts then detecting that they form two separate representation clusters of addition and multiplication facts, which in turn leads to the creation of subtypes of arithmetic type.

7.6.3.3 Learning Intermediate Types

Since an ACT-R chunk type corresponds roughly to a single layer of connections in Figure 9, it is possible that a complex type could not be learnable efficiently and/or with proper generalization without extra layers of units. The common solution to improve the learning power of the network is to add additional hidden layers (see Figure 10). Those additional layers in turn become new types. The first layer becomes one type (or more if it partitions such as the technique above applies) re-encoding the stimulus in terms of higher-order features, and the second layer is another type encoding the entire stimulus in terms of the higher-order features.




Figure 10: Switch from ACT-R type to Leabra network and back.

A good example for the need of these intermediate types to represent complex information is the ACT-R backgammon player (Sanner et al., 2000). Attempting to represent the entire board plus the evaluation of a given move in a single chunk would stretch cognitive plausibility as well as preclude adequate generalization. Instead, a move is represented in terms of intermediate high-level features such as capture and block, together with some key parameters of those features such as position. The features in our model were intuitive and simple enough to result from explicit deliberation or instruction, but often those features are implicit and result from learning of statistical regularities. Leabra would provide the means to learn those features that enhance the problem representation in ACT-R without requiring explicit knowledge engineering in the model.

While this discussion focuses primarily on declarative memory representation, this system could extend to cover multiple modules. In particular, the first network could correspond to perceptual sub-modules while the second network corresponds to higher-level cortical memory module. In effect, the encoding of the raw stimulus into higher-order features is accomplished perceptually while memory of the entire stimulus is not encoded in terms of its basic constituents but instead in terms of the higher-order chunks. This would provide a way to develop new representations across modules, in particular perceptual and procedural/memory modules in a way consistent with many findings on the representation of expertise (e.g. Chase & Simon, 1973).

7.6.3.4 Integration

One can visualize (see Figure 11) long-term memory as composed of a number of ACT-R and Leabra sub-modules interacting with each other irrespective of their own composition.




Figure11: Intra-module Organization of SAL Architecture.

A number of technical challenges are raised by that view of ACT-R/Leabra architectural integration.

The first question is when to decide which architecture to use and when to switch between them. One could see ACT-R being the default representation (for memory – again different tradeoffs might apply to different modules) based on its tractability and behavioral accuracy. However, that depends strongly on the right representation in terms of slot values for a given chunk type. If a representation proves insufficient then the ACT-R sub-module could be replaced by a Leabra module. Insufficient performance could be measured in a number of ways, from an inability to predict future events (e.g. in a prediction-based framework similar to Hawkins) to poor performance in controlling systems based on memory retrievals (e.g. instance-based models).

The Leabra network would take the same input-outputs and attempt to generate an internal representation that captures the regularities of the domain. Once the performance of the network reaches a plateau in its learning and performance, it could be again replaced by more efficient ACT-R module(s) that now have the representation needed.

The next question is how to switch between ACT-R and Leabra representations. Switching from ACT-R chunk type to Leabra network requires specifying a distributed representation equivalent to the ACT-R symbolic representation. That representational mapping of inputs and outputs to the network can be bootstrapped from inter-module representation formats and is part of the infrastructure of SAL (more on that in the next paragraph concerning the SAL infrastructure) and of course the primary objective of the network is to discover the internal representation. Once the representation translation has been determined, the knowledge content of the ACT-R sub-module can be used to provide the input-output patterns to train the Leabra network. More difficult is switching back from Leabra to ACT-R. As mentioned above, at least two ACT-R chunk types will be necessary to represent a multi-layer network with the first chunk type mapping inputs to hidden representation and the second mapping hidden representation to outputs (more than two might be required for a network with more than one hidden layer). The key problem is to generate the symbolic representation approximating the learned hidden unit representation (symbolic representation for inputs and outputs can be retrieved from the opposite transition). Statistical techniques such as principal component analysis or clustering could be used to find the main components of the hidden unit patterns. Symbolic chunks could then be created for each of these main components. Actual hidden patterns could then be encoded as new chunks with similarities set to these main component chunks that represent the overlap between the actual hidden unit representation and the representation of the main components. By exploiting the fact that ACT-R is not a purely symbolic framework but one capable of similarity-based partial-matching, loss of information in the translation from Leabra network to symbolic ACT-R chunks can be kept to a minimum.

The final question is the nature of the infrastructure providing the glue between the ACT-R chunk types and the Leabra networks. The first purpose of that infrastructure is to preserve the mapping between equivalent concepts in terms of ACT-R chunks and Leabra distributed representations. That mapping is initialized with the mapping for inter-module representations (e.g. as has already been done for communications between a Leabra visual module and ACT-R, leaving aside for now the question of learning those representations as well). As more Leabra networks are created, the mapping between the new hidden representation and the new chunks created to span that representation space is added to the overall mapping between Leabra and ACT-R. The second purpose is to manage the information flow between sub-modules within the module. That flow will depend substantially upon the specific module. One could see a number of entities managing that aspect, with different interpretations. If Leabra is managing the system, then the ACT-R submodules basically become a level of abstraction above point neurons, roughly at the area level. If ACT-R is managing the system, then the Leabra submodules become akin to the ACT-RN (Lebiere & Anderson, 1993) type memory modules with a somewhat different neural implementation. If SAL becomes the system manager, then the intra-module integration just becomes an extension of the inter-module communication. The latter option is the more likely as SAL becomes the communication framework between ACT-R and Leabra modules and sub-modules.

8. Related Approaches

Several parts of this scheme have been tried in various combinations.

ACT-RN (Lebiere & Anderson, 1993) was an early attempt to implement ACT-R using standard neural network constructs such as Hopfield networks for declarative type memories. Each ACT-R chunk type was given its own network, to keep the representation tractable (storage in Hopfield nets increases linearly with the number of units, but computation in terms of number of connections grows with the square of the number of units) and to allow each network to discover the regularities of each memory type without interference from other memory types. ACT-RN never proved to be a practical system compared to its symbolic cousin. This was largely due to poor choices of implementation, but it also reinforced the notion of moving up to higher (e.g. symbolic) levels of abstraction whenever possible.

One important feature of this architecture is that it decomposes learning at the neural (Leabra) level into a set of relatively small, tractable problems rather than leaving it as a prohibitively large (billions of neurons and trillions of synapses in the human cortex) computational problem if left whole. A number of algorithms have been proposed along those lines (e.g. cascade-correlation (Fahlman & Lebiere, 1990)), and while they do encounter the tradeoff between generalization and compact representation on the one hand and learning tractability on the other, they do provide a practical way of making the basic neural learning algorithms tractable for large, incremental problem sets.

A number of hybrid architectures have proposed to integrate symbolic and connectionist systems. A common problem has been that they either statically divide the problem, leaving each part half-covered, or they duplicate the full functionality into both symbolic and connectionist systems (e.g. CLARION (Sun, 2002)), in essence requiring both paradigms to do everything and obtaining the worst of both worlds. A number of attempts have also been made at extracting rules from neural networks (e.g. Andrews, Diederich & Tickle, 1995), but this extraction often targeted a purely symbolic rule-based system that lacked ACT-R's similarity-based generalization characteristics, and lacked the interactiveness of this approach that switches back and forth between systems as required by the task performance.

But whatever their shortcomings (or precisely because of them), there is much to be learned from previous approaches at integrated hybrid intelligent systems.

9. References

Andrews, R., Diederich, J., & Tickle, A. B. (1995). Survey and critique of techniques for extracting rules from trained artificial neural networks. *Knowledge-Based Systems*, 8(6), 373-389.

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1), 55-81.

Fahlman, S. E., & Lebiere, C. (1990). The cascade-correlation learning architecture. In D.S. Touretzky (Ed.). *Advances in Neural Information Processing Systems 2*. Los Altos, CA: Morgan Kaufmann.

Lebiere, C., & Anderson, J. R. (1993). A connectionist implementation of the ACT-R production system. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* (pp. 635-640). Boulder, CO: Institute of Cognitive Science.

Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary learning systems approach. *Psychological Review*, 110(4), 611-646.

Sanner, S., Anderson, J. R., Lebiere, C., & Lovett, M. C. (2000). Achieving efficient and cognitively plausible learning in backgammon. In *Proceedings of the Seventeenth International Conference on Machine Learning* (pp. 823-830). San Francisco: Morgan Kaufmann.

Sun, R. (2002). *Duality of the mind: A bottom-up approach toward cognition*. Mahwah, NJ: Lawrence Erlbaum Associates.

Wermter, S., & Sun, R. (2000). *Hybrid neural systems*. Berlin: Springer-Verlag.