Running Head: FROM MICROCOGNITION TO MACROCOGNITION

From Microcognition to Macrocognition:

Architectural Support for Adversarial Behavior

Christian Lebiere

Psychology Department

Carnegie Mellon University

Bradley J. Best

Adaptive Cognitive Systems

Please address all correspondence to:

Christian Lebiere

Psychology Department - Baker Hall 345A

Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213

Email: cl@cmu.edu

Phone: 412-268-6028

Abstract

Asymmetric adversarial behavior is a complex naturalistic domain that involves multiple macrocognitive processes. Traditional techniques that have been applied to that domain from game theory and artificial intelligence have generalized poorly from simplified paradigms to real world conditions. We describe a competing approach rooted in cognitive architectures and distinguished by multiple levels of processes that involve complex interactions of microcognitive constructs. The naturalistic requirements of the task impose upon the cognitive architecture additional constraints beyond those involved in modeling laboratory experiments. We describe a number of improvements to the cognitive architecture designed to boost its robustness in uncertain, unpredictable and adaptive environments characteristic of adversarial behavior. What emerges is a symbiotic relation between macrocognitive processes that drive improvements in microcognitive constructs that in turn provide a computational account of the realization of those processes in human cognition.

Keywords: Adversarial behavior, decision-making, cognitive architectures, macrocognition

Introduction

Adversarial behavior is a pervasive aspect of human activity. It happens every day in naturalistic settings as diverse as driving, work, and interpersonal relationships. In more abstract settings, it is featured prominently in paradigms ranging in complexity from simple games (e.g. board games) to first-person shooter video games (e.g., Doom, Quake, etc.) to command-and-control simulations in military settings (e.g., Command and Conquer) to massive multi-player online environments (e.g., EverQuest). Adversarial behavior has been a central application area of Artificial Intelligence from its early days (e.g., Samuels' checker player) and a key form of performance benchmark (e.g., the Friedkin Prize for chess, Robocup). In economics, a formal theory, game theory, has been developed to account for the unique aspects of behavior that it brings forth.

In stark contrast, adversarial behavior has received comparatively little attention in cognitive psychology, for a number of reasons both practical (e.g., the increased difficulties in running experiments involving multiple participants and the increasingly divergent paths their behavior takes with time) and theoretical (e.g., the lack of control and resulting difficulties in deriving statistically valid results). However, we have shown that, despite that lack of interest, cognitive models of adversarial behavior can provide high-levels of functional behavior (e.g., Anderson & Lebiere, 2003) as well as close correspondence to human performance (e.g., West & Lebiere, 2001; Lebiere et al., 2003; Sanner et al., 2000). Moreover, those models have exposed the flaws in game-theoretic accounts of adversarial behavior (e.g., West, Lebiere & Bothell, 2006; Lebiere, Wallach & West, 2000) as well as the functional advantages of cognitive architectures over machine learning approaches (Lebiere et al., 2003).

As one moves from simple, controlled and bounded domains to more naturalistic ones (e.g., Military Operations in Urban Terrain), complexities arise. Central characteristics of adversarial behavior, including unpredictability and adaptivity, stress some underlying assumptions of cognitive architectures. Asymmetric adversarial behavior further exacerbates those difficulties by putting one of the parties at a fundamental disadvantage, thus compelling it to resort to unusual or innovative tactics to improve its chances of prevailing. Key aspects of the practice if not the theory of cognitive modeling such as the ability to anticipate and reflect in the design of the cognitive model the structure of the problem-solving process and the representation of information received from the environment that are mainstays of experimental psychology designs are under stress in adversarial environments. The change in the nature of task from static, tractable experimental design to constantly changing, dynamic adversaries raises the bar for flexibility and adaptivity in cognitive models.

Klein et al. (2003) include naturalistic decision-making, sensemaking and situation assessment, planning, adaptation and replanning, and problem detection as the key macrocognitive functions that they have identified in domains they have examined. They further stress that, as Gobet and Simon (1996) and many others have pointed out, recognition processes have a primary role in decision-making. Potentially more important to the domain of adversarial reasoning, however, is the insight that human decision makers with any level of experience in a domain typically generate a plausible approach to a problem on their first try (i.e., their approach will certainly not be random) before further refining it. The general challenge for behavioral modelers is to capture these capabilities in a robust computational form that allows for prediction and testing of hypotheses. Lebiere, Gonzalez & Warwick (this issue) argue that the macrocognitive functions and processes listed above and the microcognitive architecture mechanisms, while at different levels of scale and abstraction, are inherently complementary.

First, however, we examine alternative approaches to adversarial behavior and their limitations.

Alternative Approaches

Game Theory

Since the inception of the field of game theory, researchers have been interested in applying the theory to real-world situations with the goal of improving the quality of decision making in public policy, especially in the area of adversarial reasoning as it relates to defense. Early applications to situations between adversaries that were described as “zero sum”, a situation where what benefits one of the adversaries is to the detriment of the other, seemed promising. However, many real-world situations simply do not fit this description -- outcomes that are harmful for one opponent may also be harmful for the other, or compromises that benefit both might be possible -- and basing actions on prescriptions generated from zero-sum assumptions may be problematic.

Game theorists gained leverage on these more complex situations through the application of Nash's Theorem, which states that it is always possible for an adversary to select a best course of action if their opponents also select their best course of action. These situations can be analyzed in terms of their equilibrium states, which are essentially courses of action that the adversaries cannot improve on. In many real-world situations, unfortunately, there is no single equilibrium, and the optimal course of action depends on the course of action selected by the adversary, creating a dynamic system characterized by oscillations.

For example, the game commonly known as “chicken”, where two drivers approach each other at high speed with the goal of intimidating the other driver into swerving, can only be “won” if a driver's opponent believes they will not swerve and therefore themselves deviates.

Thus, each driver must convince the other of its own commitment to irrational action. This dependence on opponents signaling their own rationality and knowing the intentions and motivations of others crosses over into human behavior, and is unfortunately not easily accounted for from the perspective of an analysis of normative (optimal) behavior. Generally, game theory's assumptions of optimal behavior based on complete information about all courses of action and their consequences is simply not supported in many empirical situations – while game theory can identify optimal behavior it struggles to account for human behavior.

Machine Learning

The field of machine learning, including the application of probabilistic nets (e.g., Bayes nets), has had many successes in modeling human behavior by simply being agnostic about the exact structure of the underlying cognition, its rationality or optimality, and instead has approached cognition by mathematically mirroring its effects through extracting regularities from existing data. In the domain of adversarial reasoning, however, this approach suffers from a critical weakness – sufficient data often do not exist from which to calculate all probabilities. The structure and content of situations about which we might have appropriate data are only a remote approximation to situations about which we wish to make predictions, and there is little hope or no possibility for obtaining the constraining data. Thus, even if related data do exist, the degree of extrapolation necessary may invalidate the validity of machine learning techniques.

What is required, instead, is an approach that allows reasonable extrapolation in the absence of data. We assert that this is only possible by extracting the cognitive processes and mechanisms responsible for adversarial behavior, in the form of a cognitive architecture, and applying those mechanisms to novel situations.

We illustrate those challenges using a task involving teams of adversarial agents in a synthetic environment. Key aspects of needed functionality for cognitive agents in that environment are the needs to flexibly and robustly represent the current context, access declarative information such as plans and strategies, manage conflicting goals, and reconcile goal-directedness and reactivity in mapping context to strategies. We describe approaches at providing that functionality within the context of the ACT-R (Adaptive Control of Thought – Rational) cognitive architecture (Anderson & Lebiere, 1998) and attempts at validating it using traditional modeling methods.

The Cognitive Approach

Rather than extend traditional Artificial Intelligence (AI) approaches to address the shortcomings described previously, our approach leverages existing theories of human cognition to develop a theory of human adversarial behavior. Our hypothesis is that adversarial behavior that more closely reproduces human behavior maximizes transfer between tasks (by leveraging common cognitive mechanisms and strategies) and learning within tasks (rather than assuming optimality or tuning to existing data). Cognitive computational methods, using an approach that depends on recent progress in cognitive psychology, neuroscience, and cognitive modeling at many levels of description, can be used to predict and model human adversarial behavior in a more realistically human way than more traditional AI approaches would either allow or encourage. This is because these cognitive methods have a fundamental commitment to emulating (at least at some level of abstraction) the same processes and mechanisms that research has shown to be characteristic of human information processing. In particular, substantial research has been conducted in generating architectures that combine symbolic and statistical properties for fast, heuristic analysis, and employ partial matching techniques to produce robust pattern recognition, both of which are hallmarks of human cognition. These approaches can be applied to modeling both the selection of behavior and the identification of relevant knowledge. Further, they incorporate stochastic processes as an architectural feature, in a way that makes the adversarial behavior difficult to predict in exactly the same way human behavior is often difficult to predict. In our approach, asymmetric tactics and strategies are driven by learning and adaptation supported by the underlying cognitive architecture, so that agents and their behavior evolve in response to the environment, where the evolved behavior includes both fine-grained responses to the opponent and the broader tactical means at their disposal. The key challenge lies in bridging the gap between the macrocognitive tasks involved in adversarial behavior and the microcognitive architectures of cognitive psychology, such as the ACT-R cognitive architecture employed in our research.

For high-fidelity modeling of human decision-making, the ACT-R cognitive architecture (Anderson & Lebiere, 1998; Anderson et al., 2004; Anderson, 2007) provides a comprehensive modeling framework validated by hundreds of models over a great variety of cognitive phenomena[1]. ACT-R has been applied to model data ranging from simple cognitive psychology phenomena to interaction with complex dynamic systems such as driving a car or interacting with an anti-air warfare control station. The architecture (see Figure 1) is composed of a set of modules, including vision, motor and declarative memory modules, coordinated by the procedural module through limited-capacity buffers. Each processing step within a module is synchronous and massively parallel while communication between modules is serial and asynchronous. This fine-grained decomposition of processing (with each individual module operation lasting from tens to hundreds of milliseconds) is ensured by the small-scale of representation in the various modules, such as "atomic" chunks holding no more than half-a-dozen constituent values in declarative memory, and production rules that match and operate on

[1] See the ACT-R web site at http://act-r.psy.cmu.edu for a repository of models and papers.

the set of buffers in parallel without complex backtracking. Activity in the modules has been correlated with functional Magnetic Resonance Imaging (fMRI) BOLD response in specific brain regions, bringing neuroscience constraints to bear on the architectural organization (Anderson, 2007).[2]

INSERT FIGURE 1 HERE

While ACT-R is based upon a methodology relying on the decomposition of complex tasks into basic cognitive operations at the sub-second scale, its basic cognitive mechanisms are constrained by the rational analysis of cognition (Anderson, 1990) that provides a statistical basis to every aspect of its performance and learning. It is a hybrid system that combines a tractable symbolic level enabling the easy specification of complex cognitive functions, with a subsymbolic level that tunes itself to the statistical structure of the environment, which provides the graded characteristics of cognition such as adaptivity, robustness and stochasticity. Those qualities endow ACT-R models with basic capacities of inference, planning, reasoning, learning and decision-making that are both powerful and general without the computational complexity and specialization of standard AI techniques. As such, ACT-R provides the foundation for our approach to adversarial behavior. However, as we will see, ACT-R needs to be changed substantially to scale up to the challenges of robustly implementing the macrocognitive functions involved in adversarial behavior in naturalistic settings.

---

[2] While a concern for neuro-anatomic accuracy might seem too far removed from our focus, especially with relating microcognitive mechanisms to macrocognitive tasks and processes, one cannot pick and choose which constraints on human behavior to account for, and only an integrated framework across all levels of description and functionality can provide a satisfactory theory of human behavior (e.g. Jilk et al, 2008).

Levels of Behavior

Figure 2 presents an overall schematic view of our framework of adversarial behavior, and specifically the interplay between the macrocognitive (naturalistic) domain constraints and the microcognitive (architectural) processes. This complete framework is essential to capturing all of the cognitive aspects of adversarial behavior. Blue boxes and arrows correspond to cognitive architectures and processes, respectively. Red ovals correspond to domain knowledge, and the red arrow corresponds to manual knowledge extraction processes. The green box represents the task environment and the green arrow the interaction of the system with that simulation.

INSERT FIGURE 2 HERE

There are three distinct functional levels to providing non-traditional, asymmetric tactical behavior. Those levels are (1) the selection of existing tactics, (2) the evolution of new tactics from existing ones, and (3) the discovery of new tactics from the constraints of the task and the means available. For instance, consider the task of facility guarding as a domain of asymmetric adversarial behavior. The goal of the defenders is to protect the facility from attack while minimizing the impact of the facility's operations, while the goal of their opponents is to hamper the facility's operations either by direct attack or by necessitating intrusive defense tactics. The first level consists in choosing between an existing set of tactics, e.g., firing from a nearby building, lobbing mortar shells from a distance or driving a car bomb to the building entrance. An example of tactical evolution would be the combination of two of the tactics above (the weapon firing and car driving) into a new tactic, a drive-by shooting of the facility. The third level, tactical discovery, would require the innovation of an entirely new tactic unrelated to the existing ones, such as poisoning the building's water supply.

Each level provides a separate functionality and uses qualitatively different mechanisms. Therefore, it is essential to make the distinction between those functions and the mechanisms that are well suited to perform them to avoid overlooking some of these functions or trying to perform them with inadequate mechanisms. We will describe each function and the mechanisms that apply to them in some detail in the following sections.

Strategy Selection

The first level (a.k.a. Level 1) is the selection of existing tactics.[3] The source of those tactics is not important here: they could have been encoded from empirical data or cognitive task analysis, or they could have originated from within the model itself such as through the use of default heuristics. The key at this level is to choose from a set of existing options. At first this may seem like a fairly trivial function. Existing strategies have an established record of efficacy and we can simply deploy the best one(s). However, this obvious approach misses the fundamental point of the cat-and-mouse game played between our technically and numerically superior forces and the asymmetric adversaries whose chief advantage is the element of surprise. Despite having superior forces, our warfighters cannot prepare for everything and the adversaries' most effective tactics are those that are least expected. Context is therefore key: the most effective asymmetric adversary is the one who can choose the options that are least expected at a particular moment and in particular circumstances. A prominent example of this pattern in Iraq is the "Whack A Mole" chase to tamp down the insurgency in places where it is

[3] In this section and following ones, we use the terms "tactics" and "strategies" interchangeably. This is not to imply that their meaning is identical ("strategy" refers to the global plan whereas "tactics" refer to the more local means to achieve it) but rather that the same analysis applies to both levels of behavior.

currently flaring up, followed by a movement of insurgents to previously calm places with low levels of troops.

Even identifying and encoding existing tactics, of course, is a challenging issue. Knowledge extraction and engineering are notoriously difficult processes, requiring considerable resources and often providing disappointing results. The key aspect of our framework that makes this method tractable is that the process is not expected to result in a perfect, complete, and final set of tactics and strategies but instead is merely a starting point, a representational skeleton upon which the cognitive architecture and models will elaborate. Next, given a set of applicable tactics, the ACT-R model selects the most suitable tactic or strategy (that is, the one most likely to benefit from the element of surprise) for execution using elementary behaviors. Two points merit emphasis: one is the statistical nature of the selection process that provides a robust and adaptive mechanism for adversarial behavior, and the other is the integration benefits of having a general cognitive architecture to implement both tactic and strategy selection and behavioral implementation. Finally, the selected behaviors interact with the task environment, leading to an adaptive feedback loop to the tactic and strategy repository. This key final step records the pattern of successes and failures of individual tactics, which in turn drives future selection processes.

Using the ACT-R architecture, we have developed a first-principles methodology for performing Level 1 adversarial action selection and have applied it successfully to a broad range of adversarial situations (West & Lebiere, 2001; West, Lebiere & Bothell, 2006; Lebiere et al., 2003). This framework has proven general enough to handle both symmetrical and asymmetric situations. The only differences are the actions available to each agent, and the knowledge about them and the general structure of the adversarial environment. Those adversarial settings all had the same characteristics: simultaneous decisions that attempt to exploit an opponent's expectations. Those cognitive agents were validated both in terms of matching human agents performance and decision-making profile and in terms of absolute performance against other human and cognitive agents (West et al., 2006). The key aspect of this kind of decision-making is adaptivity. Accordingly, our agents perform this function by leveraging the statistical primitives of the cognitive architecture rather than attempting a symbolic reasoning approach that is bound to be circular (e.g., typical "he knows that I know that he knows…" reasoning). That architectural approach has a number of fundamental benefits. First of all, it is very general and does not require any domain-specific data to parameterize the decision-making because it relies on architectural mechanisms that are by definition invariant across tasks. Second, because it is statistical rather than symbolic in nature, it straightforwardly includes stochasticity in the decision-making process in the form of a softmax (Boltzmann) decision rule that adds noise to the quantity associated with each tactic (production rule utility for procedural memory or chunk activation for declarative memory) before selecting the highest-valued option. Finally, because of both its architectural and statistical nature, those mechanisms are easily parameterizable to provide opponents of different strength and predictability levels. This adjustment of opponent strength is invaluable in training application and can be accomplished independently of domain and without requiring extensive re-engineering typical of symbolic approaches (Rehling et al., 2004).

Strategy Evolution

The second level (a.k.a. Level 2) is the evolution or combination of new tactics from existing ones. The genetic analogy is particularly relevant here: existing tactics either mutate into new ones by adding a new twist or parameter variation, or a new tactic emerges that combines aspects of two or more existing tactics. Examples of the former, i.e., tactic mutation, include the recent use of non-conventional weapons (e.g., chlorine bombs) in cars and truck suicide bombings, or the use of new places to hide Improvised Explosive Devices (IEDs), such as animal carcasses or under freshly paved roads. Examples of the latter, i.e., tactic recombination, include 9/11 itself and the recent combination of criminal and terrorist tactics, such as extortion or arson to fund the insurgency or ensure support of the local population. The innovative aspect of the 9/11 attacks resulted primarily from their original combination of two classic terrorist tactics, the suicide bombing and the aircraft hijacking, with the latter serving the functional needs of the former.

Evolution is driven by the ACT-R cognitive architecture through a combination of parameterizing existing tactics and mutating existing tactics into new combinations. ACT-R is well suited to accomplish those tasks through its combination of symbolic and statistical processes. Sanner et al. (2002) have developed a model of the Level 2 type of tactical evolution. That model works by combining symbolic reasoning (by decomposing an action in its functional constituents) with statistical adaptation (by evaluating each component using its previous experience) using the architectural integration of those techniques provided by the ACT-R cognitive architecture. Symbolic methods are essential in achieving the key recombination aspect. However, by themselves, they would be lost in a vast combinatorial space of possibilities. Conversely, statistical techniques such as neural networks or genetic algorithms that lack the symbolic structure for framing the problem space have been shown to be two to three orders of magnitude slower in converging to human-like performance (Sanner et al., 2000). Finally, this hybrid approach has the same advantages as those described above for selection, including generality, stochasticity and parameterizability. The primary difference is that instead of viewing each action independently, i.e. reactively, the impact of each individual action is evaluated in conjunction with other ones to determine their effectiveness. That process of evaluation usually has a limited look-ahead short of a full exploration of the search space. More coherent strategies are then starting to emerge from the collective consideration of individual low-level tactical moves.

Strategy Discovery

The third, and probably hardest, level (a.k.a. Level 3) is the discovery of new tactics. It is unclear how common this aspect is, because the introduction of entirely new tactics in asymmetric warfare is an unusual event. Even tactics that seem strikingly new, such as the car bombing of the marine barracks in Beirut in 1983, dates back to the 1920s when an anarchist detonated a bomb hidden inside a horse-drawn wagon in front of the New York Stock Exchange. The popularity of such tactics waxes and wanes, (e.g., car bombings flared up again after WWII in Palestine), demonstrating the importance of the selection of tactics, including those that have fallen into disuse. As uncommon as truly new tactics and strategies might be, the discovery process is not limited to them. For instance, requiring an expert to provide a comprehensive list of all possible known tactics might prove prohibitively difficult. Therefore, discovery can serve an essential role not only in discovering entirely new tactics for applications such as wargaming and Course of Action analysis but also to elaborate upon the initial set of tactics generated by human experts and knowledge engineers. Hence it is essential to address this level as well.

To perform the harder, more purely symbolic and logical task of discovering new strategies, we need to integrate planning capabilities into the ACT-R architecture to perform the inferences from the environmental conditions to innovative tactics and strategies applicable in that context. Planning is essential to the discovery of new strategies because those do not arise by accident but instead result from a means-end process that requires both structured and adaptive planning capabilities. Previous attempts to integrate planning into ACT-R in the past have been very limited (e.g., Best & Lebiere, 2006). Our approach is to aim for a general planning capability built on top of a cognitive architecture to avoid the computational intractability of many traditional approaches or the brittleness of narrow domain-specific approaches. Human planners only consider a small part of the search space but are often capable of providing efficient and effective (though not necessarily optimal) solutions to problems defying traditional algorithmic approaches. Their solutions are often substantially different from those found by algorithms (Freed et al., 2004), usually resulting in additional robustness and flexibility.

Architectural Support

While we have argued that applying cognitive architecture techniques leads to more accurate and powerful models of adversarial behavior than obtained by traditional AI techniques, applying the ACT-R architecture to this challenging new domain still requires a number of architectural enhancements. Most importantly, these new architectural developments enhance the robustness of the architecture beyond the level required to model typical laboratory experiments. While those models are typically hand-coded for static tasks, adversarial behavior stresses the flexibility, scalability and robustness of the architecture, because not every condition can be foreseen in adversarial situations. This forces the architecture to take on some of the responsibility for knowledge management and representation that is currently assumed by the modeler.

INSERT FIGURE 3 HERE

Those new, interrelated architectural developments affect each key module of the architecture as illustrated in Figure 3 and include:

- Explicit option evaluation: A more effective way of evaluating a sequence of alternatives in the declarative module. Currently the dynamics of the subsymbolic activation calculus prevent the full and effective consideration of a set of options and often lead to a premature focus on just a few possibilities (Young, 2003).
- Production selection and generalization: A more powerful way of combining context and experience in the procedural module, to enable the implementation of complex plans of actions while remaining reactive to environmental changes. This process fundamentally depends upon the new representation of context (below), which must unify internal and external sources of information.
- Robust context maintenance: A more powerful way of representing in the imaginal module the current processing context to better support decision-making at each step of the planning process. This involves implementing an effective theory of working/episodic memory in the architecture, a long-acknowledged need (e.g., Lovett et al., 1999).
- Goal management: An improved mechanism for managing goals in the goal module, in order to better support the structure of the planning process. Earlier versions of ACT-R had a rigid goal stack that was abandoned as cognitively implausible but an adequate replacement in the form of an intentional module implementing a more limited and flexible memory for goals never materialized.

We will discuss each these architectural efforts in some detail in the following sections.

Explicit Option Evaluation

A key aspect of both action selection and planning involves sequentially examining a number of possible options. While one expects to significantly prune the search space, especially through learned statistical heuristics, the problem remains of iterating through a subset of possible options. Young (2003) has pointed out that architectures such as ACT-R, which include mechanisms to reinforce activation upon access but do not include explicit control structures such as refraction (a mechanism that temporarily prevents the reapplication of a just-used action), might encounter out-of-control looping problems where one piece of knowledge becomes so active that it is constantly retrieved at the expense of others. This problem has also occurred in the development of higher-level languages that include implicit or explicit retrieval loops for checking logical conditions (Jones et al., 2006; 2007).

The current version of the ACT-R architecture (Anderson, 2007) prevents a chunk from being marked as having been retrieved, the traditional method for preventing out-of-control looping (e.g., Anderson et al., 1998). The current architectural solution to this problem is to apply the finst ("fingers of instantiation") theory (Pylyshyn, 1989), as it was implemented in the ACT-R visual module to declarative memory, essentially preventing a chunk from being retrieved again for a fixed, usually brief, amount of time or cycles. That solution is impractical for the planning process because a significant amount of time passes between iterations due to the exploration of the subspace associated with each option. We have defined a variant of the activation process, inspired by neural mechanisms such as long-term depotentiation, to provide a robust and effective solution to that problem in a cognitively and neurally plausible fashion by adding a fast-decaying inhibition term to the activation equation to balance the effects of the long-term reinforcement process. Our exploration of that solution has also led to more varied and less predictable behavior for action selection as well as for planning. This architectural change corresponds to a change in the subsymbolic level of the declarative memory module.

Production Selection and Generalization

As Best and Lebiere (2006) discovered in their model of Military Operations in Urban Terrain (MOUT) operations, effective operation in a complex dynamic environment involves a mix of following (and developing) a plan of action and reacting to changes in the environment, especially those resulting from decisions and actions by adversaries, which might be in conflict with the original plan. The problem is that the production matching process is currently quite rigid, with production conditions having to be defined in advance to anticipate all possible combinations. What is needed is a way to combine the learned production utilities that reflect the benefits of following a plan with a pattern-matching sensitivity to environmental events that allows more directly relevant productions to respond. Best and Lebiere (2006) have defined a mechanism that extends the successful similarity-based partial matching mechanism from declarative knowledge retrieval to production rule matching and selection. It combines utility and degree of match to provide an adaptive, flexible and powerful mechanism for action selection. From a neural perspective, it extends the architecture's ability to capture the similarity-based generalization resulting from distributed representations, currently limited to declarative memory located in posterior cortical regions, to the procedural memory system located principally in the basal ganglia and associated frontal cortical regions (e.g., Barnes et al., 2005). Of particularly importance for planning is the fact that the conditions of application of a production rule implementing a given tactic are not pre-specified but instead expand and contract around a canonical case with the success of that tactic. Therefore, this capability allows the discovery of new tactics by allowing each action to be considered in a novel context. The combination of such quantitative innovations leads to qualitatively new tactics. This architectural change corresponds to a change in the subsymbolic level of the procedural memory module.

Robust Context Maintenance

Planning must represent the current state of the system. Often in ACT-R the context gets reduced to a single state slot in the goal, or when the demands are higher, to contextual information kept in a special dedicated buffer (e.g., the self buffer used by Best & Lebiere, 2006, for modeling behavior in spatial environments). However, both options often prove to be inflexible in access and overly limited in content. Even more worrisome, the explicit context management in the form of the definition of task-specific structures and step-by-step manipulation of information in and out of those structures has led to a complex modeling process fraught with degrees of freedom and concerns about over-fitting (e.g., Roberts & Pashler, 2000). We are exploring a more flexible representation of short-term context that is managed automatically by the architecture on similar principles as long-term memory, resulting in both simpler and more robust models.

Planning is a particularly demanding activity regarding context representation and maintenance because it involves navigating across a set of varying context states. What is needed to support planning and other context-intensive processes is a working memory functionality to automatically and flexibly aggregates the local context in which to ground decisions. The key new component is a contextual module added to the architecture to accumulate short-term information and bind it into episodic assemblies. The structure of such information is expected to be more dynamic than the current fixed chunk structures used for storing semantic information. This context representation also plays a key role in multitasking by allowing the development of a merged context as opposed to the common and unsatisfactory treatment in terms of alternating contexts (e.g., Salvucci, 2005). This architectural change corresponds to the introduction of a new working memory module that provides a more general and flexible way of representing the processing context than the current imaginal buffer.

Goal Management

Planning usually involves manipulating goals as one explores the search space. Earlier versions of ACT-R (Anderson & Lebiere, 1998) had a goal stack, which was subsequently removed because its assumption of perfect memory can be easily refuted by empirical evidence. However, since the demise of the goal stack, ACT-R has lacked a mechanism to support goal operations despite the presence in the architectural diagram of an unspecified "intentional module" supporting the goal buffer (Anderson et al., 2004). This often results in attempts to recreate the goal stack, such as direct links from a subgoal back to its parent goal, that have many of the drawbacks of the goal stack such as inflexibility without any of its usability benefits. Our experience with high-level languages that compile into cognitive architectures such as ACT-R (Jones et al., 2006; 2007) indicates that complex behavior in dynamic environments relies heavily on goal management. A special module to support limited (in terms of capacity) but flexible (in terms of retrieval ability) goal operations would be of great benefit for planning. There is neuroscience evidence (e.g., Ridderinkhof et al., 2004) for an increase in pre-frontal cortex activation (typically associated with context maintenance tasks such as goal management) that can be seen as correlated with storage and retrieval of past goals. This architectural change corresponds to the introduction of a new intentional module that provides a more general and flexible way of representing and retrieving past processing contexts than past techniques such as a goal stack (Anderson & Lebiere, 1998).

Benefits and Validation

The architectural changes described above contribute in a number of ways to the three levels of our adversarial modeling framework. At Level 1, changes to the activation equation allow a more flexible consideration and the various options instead of having to hardwire their order of evaluation in the model structure. The generalized production utility mechanism allows their evaluation to be more sensitive to the specific conditions under which the decision is being made. At Level 2, the adaptivity resulting from the generalization of production rule matching enables individual tactics to be considered outside of their normal range of applicability, leading to the emergence of new tactical assemblies. Also, the enhanced representation of context allows for more sensitive and accurate evaluation of tactical options, leading to more effective and challenging adversarial performance. All the new architectural developments play a key role in the implementation of the most demanding Level 3. The new activation equation allows for a more robust consideration of the options available at a given state. The new production matching mechanism enables a more flexible consideration of the tactical options, resulting in innovative new tactical combinations. The new goal management mechanism enables a robust and flexible planning process that is not constrained to rigid, arbitrary tree exploration but can display the opportunism of human planning. Finally and most importantly, the enhanced context representation enables more powerful insights and a more efficient planning process by uncovering new contextual cues.

What is emerging from this approach is a mutually beneficial relationship between microcognition and macrocognition: microcognitive techniques and infrastructure provide a better computational basis than AI and other traditional techniques for the implementation of macrocognitive processes, while the implementation of macrocognitive tasks in demanding naturalistic setting provides constraints that force microcognitive theories and tools to confront the complexity beyond the boundaries of the laboratory. Validation of the architectural enhancements described here, which is currently underway, involves a complex interaction between the development and validation of these mechanisms using controlled laboratory data. Also, it involves the integration and application of these architectural variations to complex tasks that will prompt methodological advances in their own right beyond traditional modeling patterns.

Applications

Attempts to apply game theory and other traditional techniques to macrocognitive processes have run into significant roadblocks since these environments are often naturalistic (i.e., not laboratory-based), and strategies and behaviors exhibited are often one-off (i.e., they have never happened before and are unlikely to ever repeat). This does not, however, diminish the need to simulate and predict human behavior in these domains. If anything, these are exactly the domains where we need the ability to predict human behavior the most.

Our approach has been to codify the processes needed to generate this behavior in a computational form, based on past behavior, and then to extrapolate using those encoded computational processes. We believe that this potential for grounded extrapolation is a key potential benefit of this work. Thus, while the architectural extensions described above are grounded in behavioral data collected from human performance, we are also working to apply those extensions within specific domains to demonstrate their ability to generate novel behavior.

One domain in which we are working is a simulator involving vehicles and terrain called dTank (Ritter et al., 2007). Although this simulator was originally aimed at simulating tank battles, it is both light enough computationally and flexible enough for other applications, and thus we have extended it to support randomly generated maps and scenarios involving friendly, hostile, and neutral forces, both on foot and within vehicles. We are currently using this domain as a proof of concept to demonstrate that the architectural mechanisms developed do, in fact, produce robust and novel human behavior in adversarial domains.

Conclusions

Adversarial human behavior, especially that behavior exhibited in asymmetric environments where opponents differ in resources, capabilities, and willingness to engage in different strategies, is an extremely important domain to understand, and for which simulations would be exceedingly useful. However, the combination of the lack of data available for modeling human behavior in these environments and the tendency of actual human behavior to follow non-normative patterns has stymied formal approaches based on both game theory and statistics. The alternative we advance here to address the lack of data and capture the non-optimal behavior of human reasoners is to pursue the modeling of the processes implicated in adversarial behavior, rather than the behavior itself, extrapolating the processes to novel situations to generate predictions. In this endeavor, asymmetric adversarial behavior provides an excellent domain for bringing microcognitive methods to bear in understanding macrocognitive processes.

Acknowledgements

This research is supported by the Office of Naval Research (award N00014-07-C-0912). We would like to thank David Jilk for his work on the development of the dTank environment.

References

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

Anderson, J. R., & Lebiere, C. (1998). *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum Associates.

Anderson, J. R., Bothell, D., Lebiere, C. & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language, 38*, pp. 341-380.

Anderson, J. R. & Lebiere, C. L. (2003). The Newell test for a theory of cognition. *Behavioral & Brain Sciences 26*, 587-637.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y . (2004). An integrated theory of the mind. *Psychological Review 111, (4)*. 1036-1060.

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: Oxford University Press.

Barnes, T. Kubota, Y., Hu, D., Jin, D.Z., and Graybiel, A.M. (2005). Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature, 437*, 1158-1161.

Best, B. J. & Lebiere, C. (2006). Cognitive agents interacting in real and virtual worlds. In Sun, R. (Ed) *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. NY, NY: Cambridge University Press.

Gobet, F., & Simon, H.A. (1996). The roles of recognition processes and look-ahead search in time-constrained expert problem solving: Evidence from grand-master-level chess. *Psychological Science, 7(1)*, 52-55.

Jilk, D. J., Lebiere, C., O'Reilly, R. C. and Anderson, J. R. (2008). SAL: an explicitly pluralistic cognitive architecture. *Journal of Experimental & Theoretical Artificial Intelligence,* *20:3*, 197-218.

Jones, R.M., Crossman, J. A., Lebiere, C., & Best, B. J. (2006). An abstract language for cognitive modeling. In *Proceedings of the 7th International Conference on Cognitive Modeling*. Trieste, Italy

Jones, R., Lebiere, C. & Crossman, J. A. (2007). Comparing Modeling Idioms in ACT-R and Soar. In *Proceedings of the 8th International Conference on Cognitive Modeling*. Ann Arbor, MI.

Klein, G., Ross, K.G., Moon, B.M., Klein, D.E., Hoffman, R.R., & Hollnagel, E. (2003). Macrocognition. *Intelligent Systems, IEEE Volume 18, Issue 3*, pp. 81 – 85.

Lebiere, C., Wallach, D., & West, R. L. (2000). A memory-based account of the prisoner's dilemma and other 2x2 games. In *Proceedings of International Conference on Cognitive Modeling 2000*, pp. 185-193. NL: Universal Press.

Lebiere, C., Gray, R., Salvucci, D. & West R. (2003) Choice and Learning under Uncertainty: A Case Study in Baseball Batting. In *Proceedings of the 25th Annual Meeting of the Cognitive Science Society*, 704-709.

Lebiere, C., Gonzalez, C., & Warwick, W. (submitted). Emergent complexity in a dynamic control task: A qualitative model comparison. Special issue of *Journal of Cognitive Engineering and Decision-Making on Microcognition and Macrocognition*.

Lovett, M. C., Reder, L. M., & Lebiere, C. (1999). Modeling working memory in a unified architecture: An ACT-R perspective. In Miyake, A. & Shah, P. (Eds.) *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control*. New York: Cambridge University Press.

Pylyshyn, Z.W. (1989). The role of location indexes in spatial perception: A sketch of the

FINST spatial-index model. *Cognition, 32*, 65-97.

Rehling, J., Lovett, M., Lebiere, C., Reder, L. M., & Demiral, B. (2004) Modeling complex tasks: An individual difference approach. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 1137-1142) . August 4-7, Chicago, USA

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004). The Role of the Medial Frontal Cortex in Cognitive Control. *Science, Vol. 306. no. 5695*, pp. 443 – 447.

Ritter, F., Case, S., Bhandarkar, D., Lewis, B., Cohen, M. (2007). dTank Updated: Exploring Moderated Behavior in a Light-weight Synthetic Environment. *Proceedings of the 16th Conference on Behavior Representation in Modeling and Simulation*, 51-60. Orlando, FL: U. of Central Florida.

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review 107(2)*, pp. 358-367.

Salvucci, D. D. (2005). A multitasking general executive for compound continuous tasks. *Cognitive Science 29(3)*, pp. 457-492.

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

West, R. L., & Lebiere, C. (2001). Simple games as dynamic, coupled systems: Randomness and other emergent properties. *Journal of Cognitive Systems Research, 1(4)*, 221-239.

West, R. L., Lebiere, C. & Bothell, D. J. (2006). Cognitive architectures, game playing and human evolution. In Sun, R. (Ed) *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. NY, NY: Cambridge University Press.

Young, R. (2003). Should ACT-R include production refraction. In *Proceedings of the 10th ACT-R Workshop*. Pittsburgh, PA.

Figure Captions

Figure 1. ACT-R architectural diagram including main architectural modules and buffers, and communication processes. Two new modules (red) are being developed as part of this research.

Figure 2. Overall view of adversarial behavior system, including cognitive architecture processes (blue), domain knowledge (red) and task environment (green).

Figure 3. Architectural modules supporting adversarial behavior and their functions, processes and knowledge components.