Teamwork, Communication, and Planning in ACT-R Agents Engaging in Urban Combat in Virtual Environments

Bradley J Best and Christian Lebiere
Human Computer Interaction Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA
bjbest@cmu.edu, cl@cmu.edu

Abstract

Agents must be able to effectively employ teamwork, coordination, navigation, and planning to succeed in multi-agents virtual environments. These high-level abilities depend on reliable navigation and perception of the environment. We present cognitively plausible agents used for training in virtual simulations of Military Operations on Urban Terrain (MOUT). These cognitive agents are developed using the ACT-R cognitive architecture, an architecture used to model human performance in a wide variety of psychological experiments. These agents use their real-time perceptions to shape their interactions with other agents and human players on both the friendly and opposing sides of the conflict. Planning is accomplished on-line by combining schematic plans with the current context, resulting in flexible and appropriate action. Roles within plans are negotiated between agents on an as-needed basis while plans are selected based on the availability of other agents.

1 Introduction

This project is part of an effort to develop leap-ahead human-immersive technology for naval training as part of the VIRTE project. VIRTE aims to provide training tools that will prepare warriors for increasing complexity and chaos by supplementing and complementing live simulations with virtual and wargaming simulations. [Lyons et al., 2002]

Towards this end, we are working to incorporate current understanding of human behavior and learning theories into systems, leverage commercially available advanced technology, incorporate this current understanding into products, and then transition these products to the naval forces.

Our goal is to improve the cognitive validity of synthetic soldier entities in dismounted infantry simulations. This effort involves the ACT-R cognitive architecture for modeling and the Unreal Tournament virtual reality game platform for a simulation environment. To this end, we have refined and developed a framework for robust spatial reasoning, communication, planning, and teamwork between cognitive agents.

2 ACT-R

ACT-R is a unified architecture of cognition developed over the last 30 years at Carnegie Mellon University. At a fine-grained scale it has accounted for hundreds of phenomena from the cognitive psychology and human factors literature. The most recent version, ACT-R 5.0, is a modular architecture composed of interacting modules for declarative memory, perceptual systems such as vision and audition modules, and motor systems such as manual and speech modules, all synchronized through a central production system (see Figure 1). This modular view of cognition is a reflection both of functional

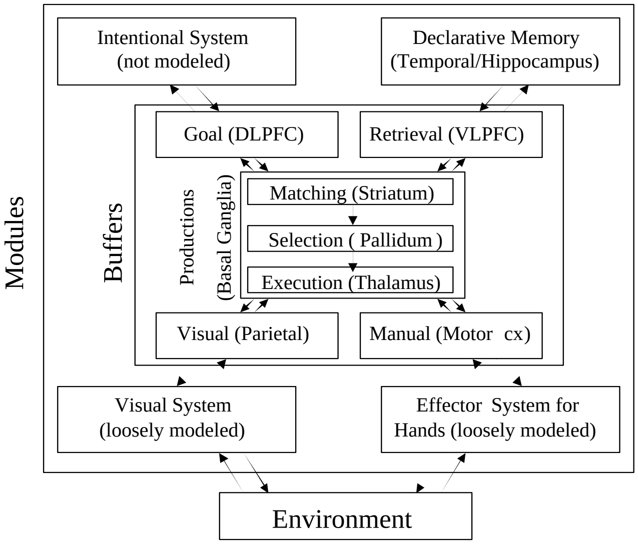

Figure 1: ACT-R Architecture

constraints and of recent advances in neuroscience concerning the localization of brain functions. ACT-R is also a hybrid system that combines a tractable symbolic level that enables the easy specification of complex cognitive functions, with a subsymbolic level that tunes itself to the statistical structure of the environment to provide the graded characteristics of cognition such as adaptivity, robustness and stochasticity.

The central part of the architecture is the production module. A production can match the contents of any combination of buffers, including the goal, which holds the current context and intentions, the retrieval buffer which holds the most recent chunk retrieved from declarative memory, visual and auditory buffers that hold the current sensory information, the manual buffer which holds the current state of the motor module (e.g. walking, firing, etc). The highest-rated matching production is selected to effect a change in one or more buffers, which in turn trigger an action in the corresponding module(s). This can be an external action (movement, firing, etc) or an internal action such as requesting information from memory. Retrieval from memory functions in a similar manner. A pattern specified in a production is sent for matching in declarative memory. Each chunk competes for retrieval, with the most active chunk selected and returned in the retrieval buffer. The activation of a chunk is a function of its past frequency and recency of use, the degree to which it matches the requested pattern, plus stochastic noise. Those factors confer memory retrievals, and behavior in general, desirable "soft" properties such as adaptivity to changing circumstances, generalization to similar situations, and variability. [Anderson and Lebiere, 1998]

The current goal is a central concept in ACT-R, which naturally provides strong support for goal-directed behavior. However, the most recent version of the architecture (ACT-R 5.0) is less goal-focused than its predecessors by allowing productions to match to any source of information, including the current goal, information retrieved from declarative memory, objects in the focus of attention of the perceptual modules and the state of the action modules. This emphasis on asynchronous pattern matching of a wide variety of information sources better enables ACT-R to operate and react efficiently in a dynamic fast-changing world. Thus ACT-R is capable of flexible goal-directed behavior which gives equal weight to internal and external sources of information.

There are three main distinctions in the ACT-R architecture. First, there is the procedural-declarative distinction that specifies two types of knowledge structures – chunks for representing declarative knowledge and productions for representing procedural knowledge. Second, there is the symbolic level, which contains the declarative and procedural knowledge, and the sub-symbolic level of neural activation processes that determine the speed and success of access to chunks and productions. Finally, there is a distinction between the performance processes by which the symbolic and sub-symbolic layers map onto behavior and the learning processes by which these layers change with experience.

Human cognition in a domain like MOUT has two principal components. The first is the knowledge and procedures codified in military doctrine and instilled in military forces through extensive training. The second is the natural cognitive abilities of soldiers that manifest themselves in tasks as diverse as memory, reasoning, planning and learning. The fundamental advantage of an integrated architecture like ACT-R is that it provides a framework for modeling basic human cognition and integrating it with specific domain knowledge.

The advantage of a symbolic system like ACT-R's production system is that, unlike connectionist systems for example, it can readily represent and apply symbolic knowledge of the type specified by military doctrine. Moreover, the specific type of knowledge encoded in that doctrine are rules specifying what to do in a given condition, a type of knowledge particularly well-suited for representation as production rules. Symbolic knowledge however does not define performance. If it did, humans would behave more like robots, acting rigidly, deterministically, with neither errors nor inspiration. In ACT-R, performance is defined by the parameters at the sub-symbolic level that determine the availability and applicability of symbolic knowledge. Those parameters underlie ACT-R's theory of memory, providing effects such as decay, priming and strengthening. But they also play a broader and more general role: they make cognition adaptive, stochastic and approximate, capable of generalization to new situations and robustness in the face of uncertainty. Those qualities provide ACT-R models with 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. Finally, because of the continuous nature of the sub-symbolic layer, architectural parameters can be varied to incorporate the effect of behavior moderators to obtain a gradual degradation of behavior in the face of fatigue, stress, injury and other conditions.

3 Simulation Environment

The Simulation environment used in this project is Unreal Tournament, a real-time first-person virtual reality game. This environment has the benefits of being inexpensive, running on modest hardware, having a well-developed programming API, directly supporting network and multi-player interaction, and possessing a wide user base.

3.1 Unreal Tournament

UT is a (mostly) open source game with a thriving online community that emphasizes networked play between synthetic and real participants. We have developed a programming interface from ACT-R to UT based on the GameBots interface [Kaminka et. al, 2002] and have obtained very good responsiveness between human players and synthetic agents. A single dedicated server can handle at least 16 players or agents, making it more than suitable for the interaction between a MOUT squad and a small group of opponents. UT also provides high-level abstractions for actions, allowing us to focus on the cognitive aspects of the project. For perception, UT provides fast regular updates of the environment (at configurable time intervals that default to 100 milliseconds) as well as asynchronous updates of special events. The simulation provides updates of the global coordinates of the various objects or artifacts in the current field of view (FOV), which are transformed into a quantitative representation.

3.2 An Autonomous Mapping Agent

For an agent to navigate and act within a space it must have a representation of the environment that supports these actions. Although simply reacting to other nearby agents can simulate rudimentary behavior, complex planning and teamwork, which are required to provide a worthwhile training experience, require a more complete representation of space. This representation can be constructed from basic elements available within the particular virtual environment, in this case UT.

The representation we have used is generated from a process that can be divided into two parts: a low-level implementation-dependent feature extraction process and a method for translating this to a cognitive-level representation usable by the agent. Note that the abstract representation is implementation-independent. Implementations on other platforms would focus on extracting low-level primitives available in that environment and mapping them onto the cognitive-level representation.

The low-level spatial primitives available in UT are fairly sparse. An agent in UT can request information about whether a particular point is reachable from its current location, meaning that there is a clear navigable straight path between the agent's location and the target location. This data is analogous to range-sensor data, and can be used to build higher-level representations. Because of the primitive nature of the spatial primitives and the relatively slow rate at which they are provided by the environment, we needed to build offline a representation of the space for use by the synthetic agents. For that purpose, we developed an autonomous mapping agent that navigates the environment (typically, the inside of a single building) and gradually builds a representation of the space (see Figure 2).

Using the range-sensing data as the only spatial primitive, the exploration agent builds up a usable cognitive-level representation of the space. Using a Hough transform to detect straight lines within the data, a cognitive-level description that consists of walls, corners, rooms and openings is constructed. From this static representation, the dynamic perceptual presence of architectural primitives relative to the agents' current location can be determined in real-time. [Best et al., 2002]

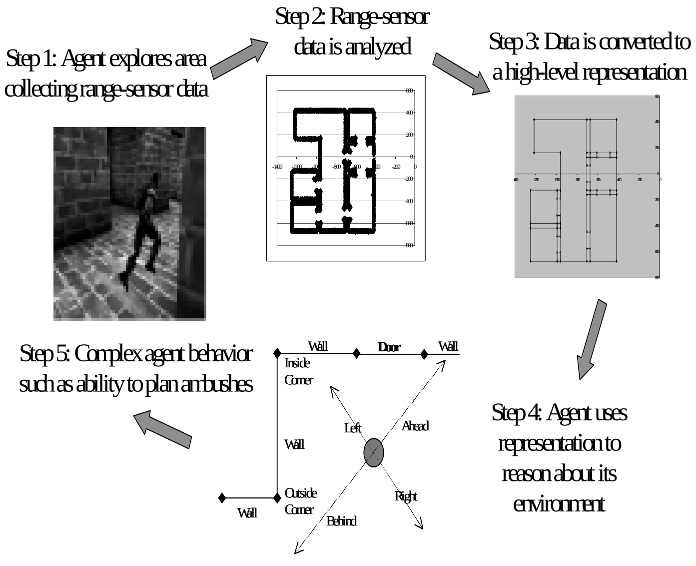

Figure 2: Map Construction by Autonomous Agent

4 Perception and Spatial Representation

The ACT-R/MOUT framework currently provides for visual, auditory, and tactile senses. A sampling of the primitives available through these senses and their attributes is listed in Figure 3 below.

The ACT-R/MOUT framework uses architectural features to describe the visual surroundings of an agent. These visual primitives include walls, corners (intersections of walls), doorways, stairwells, and entrances. In addition there is an extensible set of object primitives including other agents and items such as weapons, ammunition, and furniture. Auditory primitives include verbal messages (encoded as text), and sounds such as weapon fire, footsteps, etc., which are localized in space. Tactile inputs are limited to sensing contact with architecture or objects in the environment, and receiving injury (potentially from those objects). This allows the agent to, for example, feel its way along a wall, or to stop backing up when it makes contact with another player.

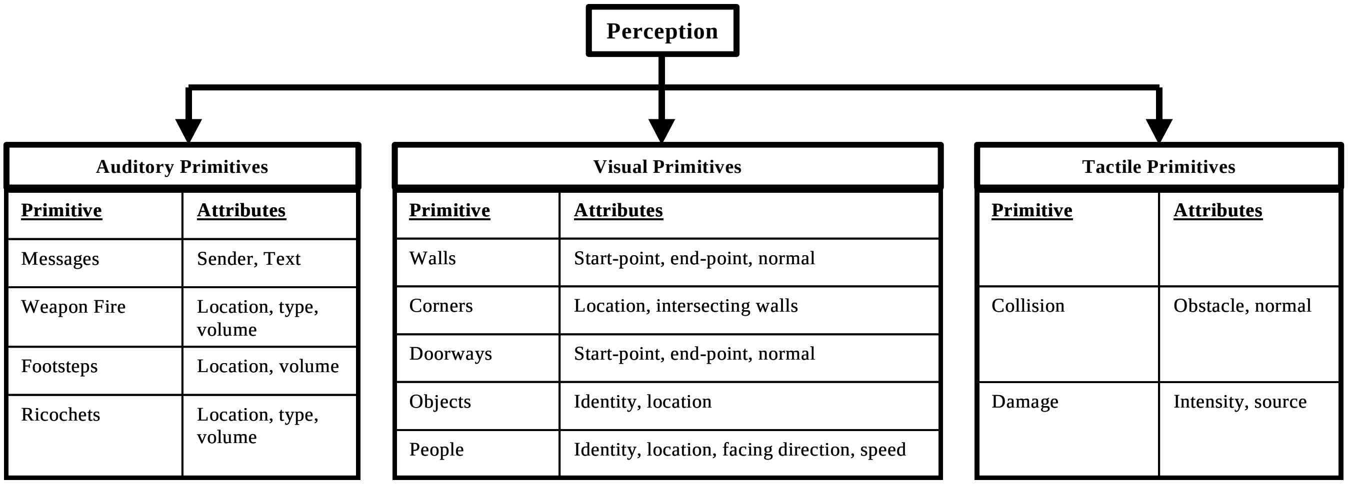

Auditory Primitives

Primitive	Attributes
Messages	Sender, Text
Weapon Fire	Location, type, volume
Footsteps	Location, volume
Ricochets	Location, type, volume

Visual Primitives

Primitive	Attributes
Walls	Start-point, end-point, normal
Corners	Location, intersecting walls
Doorways	Start-point, end-point, normal
Objects	Identity, location
People	Identity, location, facing direction, speed

Tactile Primitives

Primitive	Attributes
Collision	Obstacle, normal
Damage	Intensity, source

Figure 3: Perceptual Primitives

4.1 Spatial Representation

In order to perceive, react, navigate, and plan, it is necessary for the agents to have a robust spatial representation. The agents use a representation based closely on findings from studies of human spatial cognition. Agents can represent things in two fundamental ways: where something is relative to the agent's location, or egocentrically (e.g., something is 10 feet to my left); or where something is in absolute terms relative to a world coordinate system, or allocentrically (e.g., something is at a particular latitude/longitude). In addition, the agents' design is influenced by work on autonomous robots (which also borrows heavily from studies of human spatial cognition). [Frank, 2000]

The egocentric representation of an item includes both the distance to the item and its relative bearing. Distance and bearing are both represented quantitatively and qualitatively.

Distance is represented both as: 1) absolute distance to the target, and 2) descriptive distance to the target. A descriptive distance is how distant something is relative to the current visual horizon, and ranges across "here" (within 1/8 the distance to the horizon), "near" (between 1/8 and 1/4 the distance to the horizon), "far" (between 1/4 and 1/2 the distance to the horizon), and "very far" (beyond 1/2 the distance to the horizon). This encoding is similar to the three spatial modules recently proposed as an extension to ACT-R. [Schunn and Harrison, 2001]

Bearing is represented as both: 1) absolute compass bearing to target relative to current orientation (e.g., 30 degrees to the left, 5 degrees up), and 2) descriptive bearing to the target. A descriptive bearing may be "right", "left", "ahead", "behind", or any of the four intermediate bearings "ahead right", "ahead left", "behind right", or "behind left" (see Figure 4 for a diagram of this representation). In keeping with studies of human spatial cognition indicating that humans maintain egocentric representations with much more accuracy than allocentric representations, the agents rely extensively on egocentric representations. [Klatzky, 1998]

The allocentric representation of an item includes the location of an item in the world coordinate system (in this case, x, y, and z) and its orientation relative to that coordinate system (pitch, yaw, and roll – the angles relative to the axes). An allocentric representation is particularly important in reference to maps (which are typically defined relative to some world coordinate system), and correspondingly to planning and navigation tasks.

Allocentric and egocentric representations are complementary. The egocentric representation of an object always has a corresponding allocentric representation. Many studies of human spatial cognition have been directed at how people translate from one representation to the other (e.g., map following). [Klatzky, 1998]

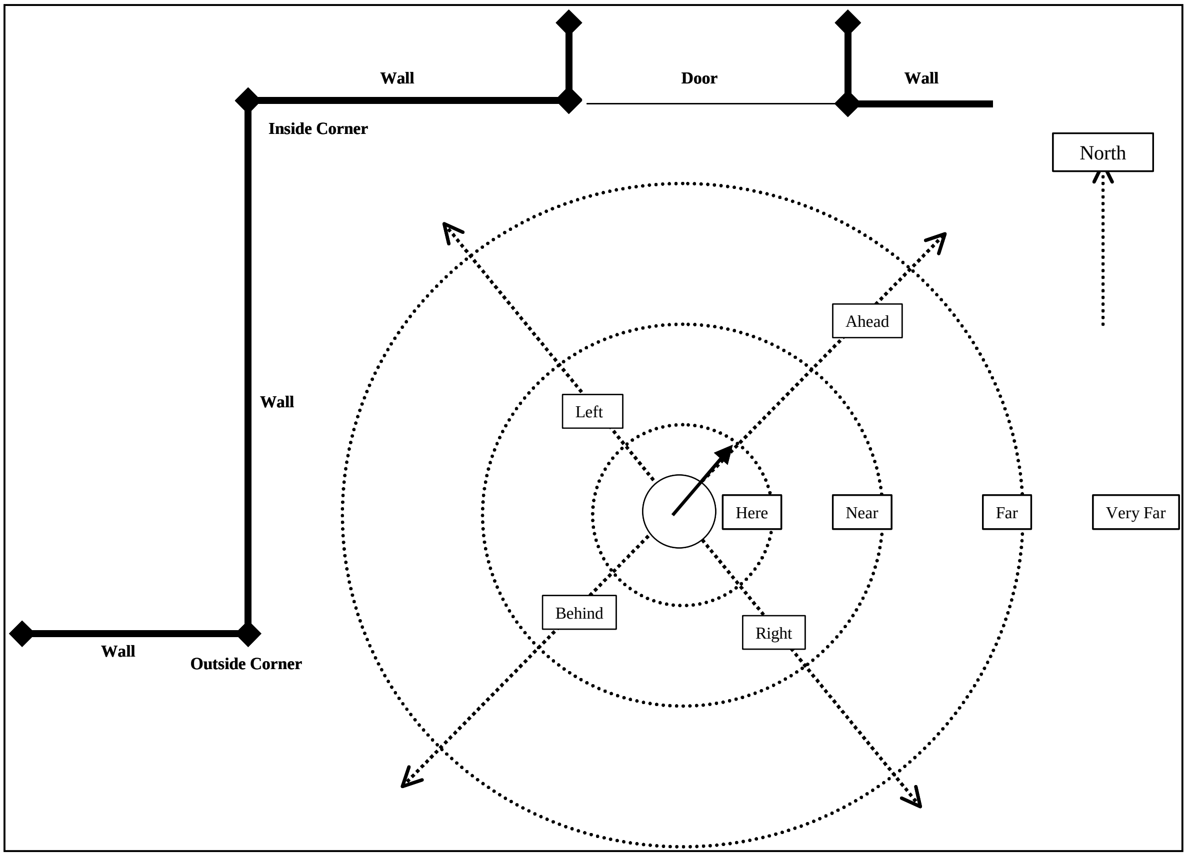

Figure 4: Spatial Representation Showing Egocentric and Allocentric Reference Frames

5 Navigation

The most fundamental skills needed to cope with the demands of the virtual MOUT environment are basic navigational abilities such as locomotion, determination of location and orientation, and recognition of locations and use of paths to travel between locations.

Locomotion, the most fundamental behavior, is simply moving from one location to another. It does not require attention once initiated, and thus may occur in parallel with other activities such as communicating and firing a weapon. As implemented, locomotion involves starting to move to a destination, continuing movement while not at the destination, abandoning that movement if an obstacle is encountered, and halting the movement upon arrival at the destination. Besides destination, velocity can also be specified and agents can move by walking, running, or slowly stalking. They can also direct their attention elsewhere and proceed to the destination sidestepping while facing another direction. Receiving an order or reacting to a new threat can also interrupt locomotion, or it can be handled (e.g. by firing a weapon while running) in parallel with it.

Higher order navigational behavior involves an interaction of the cognitive map of the environment (the allocentric reference frame) with the current visual scene (egocentric cues) and memory for goals and past events (paths followed and destinations). As such, it represents a significant theoretical challenge in both cognitive psychology [Klatzky, 1998] and robotics [Beetz, 2002].

Agents in this simulation use a node-link-node representation for rooms and pathways between them. Attacking agents may be given a pre-defined path representation (mission instructions). Alternatively, they may build up a representation of rooms visited, as well as the episodic trace of items and other agents seen there. When moving from the current room through a doorway to a new room, the agent creates a chunk in declarative memory corresponding to that path. Defending agents, who are assumed to have intimate knowledge of the area to be defended, are given a complete representation of the rooms and pathways connecting them within a building. This allows them to fluidly and quickly choose paths for attack and escape.

The attackers may be forced to rely on knowledge gained during interacting with the environment, instead of on predefined knowledge. In addition to remembering the path followed, attackers may also encode individual moves at particular situations. This is similar to the heuristic applied by some people who “retrace their footsteps” when trying to find their way. These previous moves can be actions relative to landmarks (e.g., turn left at the L-shaped hall), actions relative to an allocentric frame (e.g., proceed at a compass bearing of 90 degrees), or actions relative to an egocentric frame (e.g., turn left 45 degrees). These representations are complementary, and are commonly used by people as the context allows. Landmarks are often preferred, but in a situation where landmarks are impoverished, people quickly adopt the other strategies. In a combat context, an allocentric frame such as that encoded in a map is often used. This is particularly useful in military situations for exchanging information about threats, destinations, and movements, and in planning, since allocentric coordinates such as GPS coordinates are unambiguous, while egocentric coordinates depend on knowing the egocentric orientation of the perceiver and are therefore often less useful.

6 Communication, Teamwork, and Planning

To succeed in a MOUT environment, effective communication, teamwork, and planning is essential. The agents described in this paper engage in elaborate planning and teamwork consistent with MOUT doctrinal approaches.

6.1 Communication

Communication is a prerequisite for teamwork and planning. Within the UT environment, the simplest form of communication is through passing text messages. We have created a grammar that defines a protocol for the agents' use in signaling, acknowledging, sending and receiving orders, communicating intent, and specifying the type and location of a contact (e.g., friendly fire, from location (x,y,z)).

The most fundamental of these, simple communication, is potentially non-verbal and simply involves the passing of signals and the acknowledgment of their receipt. For example, saying "On the count of three, go" requires the receipt of the signal "three" while ignoring other signals. In the UT environment, this is implemented by passing text messages between the agents. These messages take the form of English phrases for the sake of readability and extensibility to interactions with human players.

The passing of orders typically involves an instruction to execute a schematic plan, very much like running a play from a playbook. These plans include actions such as clearing an L-shaped hallway, supplying covering fire, moving to a particular location, standing guard, providing assistance in storming a particular room, providing covering fire, or retreating from overwhelming fire. Each of these plans is defined in military doctrine to a level of detail that includes the actual phrases to be used in communications. This enables a minimal verbal instruction to carry a very substantial content. Many MOUT squads also evolve their own extensions to standard doctrine communications, further improving their efficiency and reducing an enemy's ability to understand their signals.

In addition to orders, agents can also share information. The most common information shared is a spot report of enemy activity. A spot report includes a brief description of the enemy forces spotted including their numbers and armament if known, their location, and their movement (if any). Other agents may use this information to provide coordinated ambushes and attacks.

6.2 Teamwork

Teamwork requires a team, the members of which are known. Rather than establish this in advance, within our simulations agents negotiate roles with each other. The fundamental organization within urban combat is a team of two combatants. The next level of aggregation is of a squad consisting of two teams of two combatants. In the military, these roles are often predefined. Within this simulation, these roles are determined on the fly to allow more flexibility in working with scenarios.

Upon introduction to the UT environment, agents first pair up. They simply seek a nearby teammate available for pairing. Once a pair enters into a partnership, they then negotiate for leadership of the team using a simple horse-race strategy: The first to claim leadership by saying so is the leader. The same strategy is used to determine the lead team of the squad. This self-organization, and more generally the agents' reactivity to their environment, provides a good deal of robustness to the agents' behavior.

The status of an agent's team and squad membership is critical to planning in MOUT. Human combatants are trained on appropriate techniques for performing actions based on the presence of supporting forces. Similarly, the agents use their membership status to constrain the plans they form to accomplish goals. For example, a team leader, when rounding an L-shaped hallway, may choose to take the position along the back wall. A team subordinate will then take the other position along the front wall approaching the corner.

6.3 Schematic Plans

A large part of the teamwork exhibited by these agents hinges on sharing common knowledge about how to approach certain tasks. The details on how to do this come directly from military doctrinal manuals [e.g., MCWP] and are routinely taught to trainees as part of their fundamental training. Each agent knows, as a trained human combatant does, what actions to perform when playing any of the roles in different scripts. This knowledge is stored as a set of chunks in declarative memory of the agent. These chunks, analogous to a schema, are a somewhat general description of what to do in a certain situation. The details are then filled in by productions that interpret the declarative script given the currently perceived environmental context. Table 1 gives an example of a production that selects the next step in an action plan as well as three steps of the plan.

```
(p get-next-action              (action-L1
   =goal>                         isa action
      isa action                  plan take-L-corner
      plan =plan                  index 1
      index =index                type move
      type nil                    argument inside-corner)
      argument nil              (action-L2
   =action>                       isa action
      isa action                  plan  take-L-corner
      plan =plan                  index 2
      index =index                type wait
      type =type                  argument go)
      argument =argument       (action-L3
➔                                 isa action
   =goal>                         plan  take-L-corner
      index (1+ =index)           index 3
      type =type                  type move
      argument =argument)         argument around-corner)
```

Table 1: Schematic Plan for Clearing an L-Corner

Plans in which an abstract script is filled with details later are sometimes referred to as “skeletal plans”, or sometimes simply as “scripts”. We have chosen to use “schematic plans”, since the plans we are dealing with here have a spatial component, and are most easily visualized using a schematic diagram. [Stefik, 1995]

For example, when clearing an L-shaped hallway, the procedure for clearing the hallway is well-defined (see Figure 5). A pair of attackers will split up and take position along the front wall (agent L in the diagram) and back wall (agent F) respectively. Agent L then moves forward close to the corner while agent F waits for a signal. Once in position, agent L signals agent F to move. Agent F then advances to a position almost directly across the hall from agent L. At this point, agent L waits for F to signal the next move. Upon agent F’s signal, L and F simultaneously move into the hallway, L staying close to the corner and dropping to a crouch while F sidesteps along the back wall. This brings both of their weapons to bear on the hallway simultaneously, while allowing both of them an unobstructed field of fire including the whole hallway.

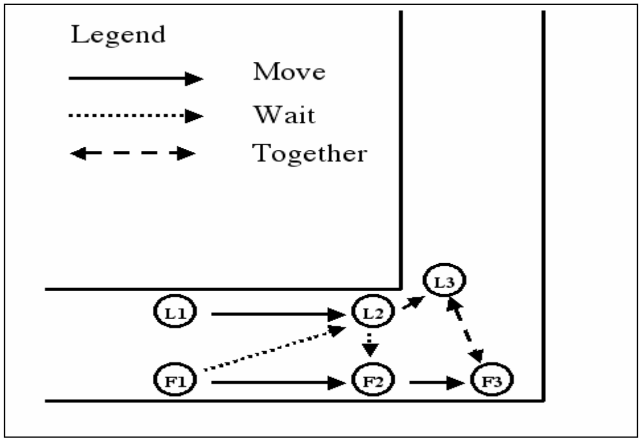

Figure 5: Two UT Agents Clearing an L-corner

These schematic plans, then, are scripts with a spatial component that describe how multiple agents are expected to work together in a particular situation. For both human combatants and agents, this bypasses the potentially deadly inefficiency of trying to decide what to do at each step. Each agent knows what to do, and what to expect from a partner. Signals are predefined and the potential for confusion is relatively low. The MOUT environment provides a clear example of a domain where teamwork is explicitly taught at a fine level of detail. [MCWP]. A visual snapshot of two agents performing this script in UT is presented in Figure 6.

Figure 6: UT Agents Clearing L-shaped Corner

6.4 Planning

Due to the dynamic nature of the task environment, it is not possible to fully develop a plan prior to performing the first actions. Instead, a rough plan consisting of abstract steps is developed. The abstract steps themselves can be implemented by the schematic plans described above. In turn, individual actions to accomplish the schematic plans are combined with elements from declarative memory and perception to form an action plan on an as-needed basis (see Figure 7 for a diagram of the hierarchical planning framework). This provides flexibility and robustness in the actual actions taken since they are planned with the immediately perceived context in mind.

Figure 8 shows a plan during execution. In this case, the agent has just performed the actions to clear an L-shaped hallway and has noticed the entrance to a room. The planning has proceeded to execute a schematic plan to clear the room. The planning will continue by determining the sequence of actions necessary to clear the room, and grounding those actions in the currently experienced context. Once that room is cleared, if there are any remaining tasks involved with clearing the first floor, they will be inserted into the plan. Otherwise, the agent will move on to clearing the second floor, with the first step being the schematic plan of gaining entrance to that floor.

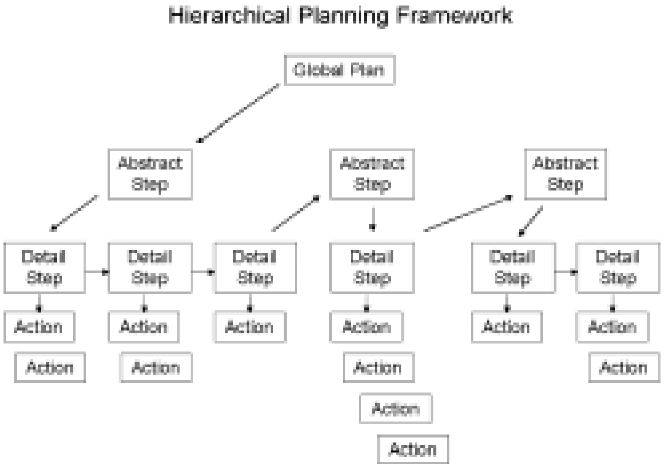

Figure 7: Planning Framework

This method of planning has roots in means-ends analysis and has much in common with skeletal planning and hierarchical match algorithms. Since the plan can be modified at several abstract levels, it may be better described as hierarchical planning. However, the individual action steps themselves are highly constrained while the planning at the more abstract levels is less constrained. This significantly reduces planning complexity since the sequence of action nodes is most often predefined by a schematic plan. The interesting implication is that human combatants have developed schematic plans to deal with exactly those situations that present many options. In any case, this type of hierarchical planning, modified by on-the-fly circumstances, provides planned, goal directed behavior that is sensitive to context. The abstract plan of clearing the two floors will not change under most circumstances, but the details of carrying out these steps often cannot be known in advance. [Schank and Abelson, 1977; Stefik, 1995] This provides an efficient compromise between the need for flexibility in robustly adapting one's behavior to unforeseen (and unforeseeable) circumstances with the need for efficiency in executing any actions in dealing with immediate threats. This tradeoff is representative of many everyday though less dramatic human environments, e.g. driving.

7 Discussion

The agents discussed in this paper leverage current knowledge about spatial perception and reasoning, and integrate it with the ACT-R cognitive architecture to provide robust, cognitively plausible behavior. These agents engage in communication, teamwork, and planning to achieve their goals in the MOUT environment.

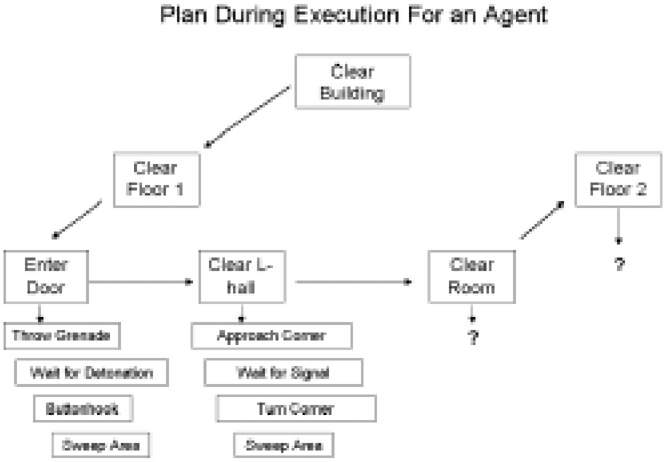

Figure 8: Abstract Plan During Interpretation

Each agent within this framework is autonomous. It has no knowledge of whether another player is a synthetic agent or human player, and the only knowledge exchanged between agents is passed in text messages. In fact, the only connection between the agents is through the simulation environment, so the agents themselves may run on physically distinct machines or, at the least, in separate computational processes. Despite the clear dependence on individual cognition and lack of collective control or management, collective behavior emerges. Agents take on roles as needed. They negotiate leadership of teams, give orders to subordinates, share plans and goals, and coordinate their behavior.

The agents use schematic plans, which are spatially scripted plans that describe teamwork in terms of locations relative to architecture, signals to be sent and received, and actions to take. These skeletal plans are combined with the products of real-time perception to create concrete action plans grounded in the current context. This allows for a blend of reactivity and planning that would be difficult to achieve with a more rigidly structured plan.

The use of schematic plans, borrowed directly from MOUT doctrine, greatly simplifies determination of action plans. In this case, the domain is well-understood and documented, allowing the process of human planning to be folded into the cognitive models. It is interesting to note that these frameworks for action simplify planning at an intermediate level of detail, and do not appear at more abstract levels. This allows very complex behavior to arise by using these building blocks made up of incompletely specified sequences of actions.

The current task environment, simulated military operations in urban environments, is highly constrained in the range of expected behavior in response to situations. These constraints are explicitly expressed in the form of military doctrine. Improvisation is rarely desirable in a combat situation. Warriors depend on each other to perform according to well-scripted plans, similar to plays run by sports teams. Teamwork is accomplished through agreeing on which play to run in advance instead of through attempting to design plays on the fly. This allows combatants to minimize communication and confusion, control exposure to enemy threats, concentrate firepower, and maximize chances of survival and success. However, this efficient behavior is a function of the knowledge level and is not the only kind that can be displayed by the architecture. Different levels of training and knowledge would lead to cognitive agents that display irrationality, poor performance and other self-defeating behaviors.

Due to the hostility inherent in the MOUT environment, learning by trial and error is infeasible. Instead, warriors are trained in doctrinal approaches that have been developed through the experience of others. This allows contemporary warriors to profit from the lessons learned by former combatants. An added benefit of this is that this knowledge has been codified in a declarative form. The result is a domain where teamwork and planning is well understood, making simulating that planning and teamwork with cognitive modeling feasible. An interesting possibility of high-fidelity cognitive agents is that they provide a scalable opportunity for wargaming and inventing new tactics and strategies, perhaps in response to a fast–changing technological environment. Unlike humans, agents can learn from failure and death and gradually evolve better strategies by incorporating their experiences through the architecture's learning mechanisms.

Although the schematic plans detailed here are described in MOUT doctrine in detail, the plans are still presented in an abstract form. A major focus of this research has been to develop a framework that supports those abstract plans and allows their combination with current context to provide a concrete action plan. This has required the development of a general spatial representation for this task environment so the ACT-R architecture and the military doctrine authors are speaking the same language.

Synthetic opponents found in gaming environments usually have superhuman capabilities or crippling shortcomings (or sometimes both). The combination of basic cognitive abilities with a representation of military doctrine has provided agents that have both human limitations and human flexibility and expertise.

Acknowledgments

This work was funded by grant N00014-02-1-0020 from the Office of Naval Research.

References

[Anderson and Lebiere, 1998] John R. Anderson and Christian Lebiere. The Atomic Components of Thought, Erlbaum, Mahwah, NJ 1998

[Beetz, 2002] Michael Beetz. Plan-Based Control of Robotic Agents, Lecture Notes in Artificial Intelligence 2554, 2002.

[Best et al., 2002] Bradley J Best, K. Christine Scarpinatto, and Christian Lebiere. Modeling Synthetic Opponents in MOUT Training Simulations Using the ACT-R Cognitive Architecture. Proceedings of the 11th Conference on Computer Generated Forces and Behavior Representation. University of Central Florida, 2001

[Frank, 2000] Andrew Frank. Spatial Communication with Maps: Defining the Correctness of Maps Using a Multi-Agent Simulation. Spatial Cognition II, 80-99, 2000.

[Gillis, 2000] Patrick D. Gillis. Cognitive Behaviors for Computer Generated Command Entities. U.S. Army Research Institute Technical Report, 2000.

[Kaminka et al., 2002] Gal A. Kaminka, Manuela Veloso, S. Schaffer, C. Sollitto, R. Adobbati, Andrew N. Marshal, Andrew S. Scholer, and S. Tejada. GameBots: the ever-challenging multi-agent research test-bed, In Communications of the ACM, January 2002.

[Klatzky, 1998] Roberta L. Klatzky. Allocentric and Egocentric Spatial Representations: Definitions, Distinctions, and Interconnection, Spatial Cognition 1-18, 1998.

[Lyons et al., 2002] Denise Lyons, Joseph Cohn, and Dylan Schmorrow. Virtual Technologies and Environments for Expeditionary Warfare Training. Proceedings of the IEEE Virtual Reality 2002.

[MCWP] Marine Corps Warfighting Publication (MCWP) 3-35.3, Military Operations on Urbanized Terrain (MOUT)

[Newell and Simon, 1972] Allen Newell and Herbert A. Simon. Human Problem Solving. Englewood Cliffs, NJ: Prentice Hall, 1972.

[Schank and Abelson, 1977] Roger Schank and R. P. Abelson. Scripts, plans, goals and understanding. Hillsdale: Erlbaum, 1977.

[Schunn and Harrison, 2001] Christian Schunn and Anthony Harrison. ACT-RS: A neuropsychologically inspired module for spatial reasoning. Proceedings of the Fourth International Conference on Cognitive Modeling, 267-268. Mahwah, NJ: Erlbaum, 2001.

[Stefik, 1995] Mark Stefik. (1995). An Introduction to Knowledge Systems. San Francisco: Morgan Kaufmann.