A Culturally-enhanced Environmental Framework for Virtual Environments

Ryan McAlinden
William Clevenger
University of Southern California
Institute for Creative Technologies
13274 Fiji Way
Marina del Rey, CA 90292
310-574-7817
mcalinden@ict.usc.edu;clevenger@ict.usc.edu

Keywords:
Artificial Intelligence, Culture, Affordance Theory, Games, Terrain

ABSTRACT: *This paper details the design and implementation of an embedded environmental framework that introduces cultural and social influences into a simulation agent's decision-making process. We describe the current limitations associated with accurately representing culture in virtual environments and military simulations, and how recent research in other academic fields have enabled computational techniques to begin incorporating the effects of culture into AI and behavior subsystems. The technical approach is presented that describes the design and implementation of a hierarchical data model, as well as the software techniques for embedding culturally-specific information inside of a virtual environment. Finally, future work is discussed for developing a more comprehensive and standardized approach for embedding this culturally-specific information inside of the virtual domain.*

1. Introduction

"To an American Soldier, culture can be likened to a minefield- dangerous ground that, if not breached, must be navigated with caution, understanding, and respect" (LTC Wunderle, 2005).

Until recently, traditional theatre operations conducted by the U.S. Military have involved large-scale, well-structured interactions with a symmetric opposing force (OPFOR). However, as the threat has shifted to become more asymmetric and fragmented, the Military has been forced to adopt new techniques and methodologies to better understand and manage it. Such threats are not only combative but cultural in nature, and in both cases have a tendency to embed themselves within the general population making it very difficult to identify and discern who the perceived threats are, as well as the strategies useful in managing them. Iraq and Afghanistan are exemplary examples of this for they illustrate an insurgent threat that is not only dispersed across various elements of the terrain, but embedded culturally within the general population, making the situation inherently nebulous to control. Concurrently, there is a general failure by both instructors and automated systems (e.g. simulators) within the training and modeling and simulation (M&S) communities to adequately represent and model aspects of these cultural threats which are important to both standard military operations as well as operations "other than war" (Panagos, 2004). It has been documented that various layers of culture (national, ethnic, social, and religious) can significantly affect how individuals and groups of individuals organize and behave, particularly when their primary motivations are not solely derived from a particular agenda but instead from a diversity of in-groups (Triandis, 1989). Therefore, a cultural representation in training systems and simulation environments is imperative to understanding how and why these asymmetric threats behave the way they do, and what strategies should be developed to best manage and mitigate them. This paper details one approach for this representation that explicitly embeds cultural information within a simulated environment to allow an AI or behavior system to culturally interpret its surroundings. Also discussed is one possible AI implementation for utilizing this embedded information to produce observably goal-directed, culturally-influenced entities within the virtual domain.

1.1 Motivation

It may be difficult for us as humans to acknowledge but our biases, stereotypes and inability to understand/accept races, religions, or ethnicities aside from our own dramatically affect how we sense and act on a daily basis. Many of our strongest beliefs and ideals are rooted in cultural and social constructs (Hofstede and Hofstede, 2005), though research in applying these types of influences to AI or behavior subsystems has been minimal. This is possibly due to culture often being perceived as a higher-level abstract concept that can be difficult to quantify and represent in a logic-constrained computational environment due to the magnitude and combinatorial complexity of various elements that must be formulated into a cohesive and quantifiable model. However, several strides have been made in other research fields to understand the effects of culture on human behavior, including business/management, healthcare, security, and of course psychology and sociology, which is a promising start to representing these effects in simulations.

Triandis (1988) and Hofstede/Hofstede (2005) are two examples that evaluate and apply culture to specific domains (Psychology and Business, respectively) to better understand and classify human behavior. Though aspects of their approach are not unique, they both attempt to quantifiably characterize human behavior from certain cultural parameters (individual vs collective, masculine vs feminine,…), which enables a computational domain to then take as input for altering agent behavior. Derived from the types of models presented by Triandis and Hofstede, cultural inputs to an intelligent agent can be grouped into first and second-order aspects of culture. First-order aspects of culture include the descriptors most people would use to define their cultural identify (race, ethnicity, nationality, politics, religion, economic status, age, gender…). Second-order aspects of culture attempt to characterize more directly how the first-order aspects influence beliefs, attitudes and actions (individuality, egalitarianism, risk acceptance/avoidance, short/long term oriented, task/relationship focused…). Thus, first-order descriptors, such as nationality and religion, lead to second-order descriptors, such as individual or collective, which influences behaviors relating to power distance, occupation, the family, school, workplace and the State (Hofstede and Hofstede, 2005). Though these characteristics are not true of a population in its entirety, they have been suggested by Triandis (1988) as one of the most promising delineations in understanding the way culture relates to social psychological behavior, which can assist us in developing a more comprehensive understanding of cultural affects on behavior.

1.2 Embedding Culture in the Environment

In addition to there being a lack of cultural models present within current simulation systems there is also an overwhelming reliance on scripted and static techniques for modeling AI behavior. The problem is particularly prevalent in the game industry, where scripted AI represents the most common technique for determining NPC actions (Baker, 2002). One reason for the overuse of scripting is the lack of information present in the environment useful to the AI's decision making. Current M&S environments typically rely on primitive elements of the terrain for an agent's decisions and often at a very low-level such as used for path-planning and navigation, and nowhere near the level of fidelity required for representing complex and variable agent behavior such as culture. Geometry, collision surfaces, ground type, pathnodes and their networks are well-suited for basic mobility and projectile calculations but fail to accurately convey higher-level pertinent information that may be useful to agent's set of goals. For example, if an agent has a goal to `secureNeighborhood()` there is currently little information contained within the environment that would indicate how that goal may be satisfied. Certainly aspects like terrain navigability and pathing are useful for such a task but there is no concept of the neighborhood, or the components that make up a neighborhood (predominant religious affiliations, socio-economic classes, political alliances). As a result, bland and scripted AI sequences are created that fail to exhibit adequate intelligence beyond basic movement and tactical doctrine. As described earlier, the focus of military operations have shifted where the representation and understanding of non-combative elements (i.e. culture) is essential, which highlights a requirement for much more complex human behavior in M&S environments. Ultimately, it is the patterns, landmark references and cultural influences, not the geometry and their facades, that shape both an agent's low-level actions (movement, gestures) and high-level perceptions and emotions. Our approach is to embed this contextual information directly in the virtual environment (i.e. terrain) and have the AI (scripted or otherwise) use this information in its decision-making. By embedding this type of data it allows agents to apply context to the objects around them (such as usability and satisfiability) and as a result provide a more immersive and realistic simulation experience.

There are several implications, both positive and negative, for embedding annotations and affordances in the environment versus directly in the AI's knowledge base. One advantage is that the knowledge is represented and stored in a format that is independent of any single AI system or agent. This allows use of the information by many systems and does not require each agent to have a separate copy of the cultural context of the environment in memory. A second advantage is authoring. It is easier, and requires less programming skill, to add to a simulated environment cultural annotations and affordances on objects and regions of the terrain. Lastly, dynamic annotations and affordances (i.e. environmental attributes that change over time) support an environment that is episodic and historical, and agents can use that information to maintain a sense of history without being forced to keep all of that information in local memory. However, there are a couple of disadvantages to embedding culture within the environment. One is that it increases the sensing "cost" by agents, which directly results in higher required throughput between the game and AI. A second disadvantage is that agents are still required to know how to react to the cultural information around them (i.e. once the cultural descriptors are sent to the AI they must do something with them) which still requires some (perhaps a great deal) of culturally-related information to be stored in the agent. Additionally, identifying what information is relevant to an agent is vital to this approach and an area we have only cursorily examined. We have prototyped an example of such an AI system in the context of a Markov Decision Process (MDP) but understanding how the cultural annotations affect behavior could certainly be a very large challenge.

The research presented here identifies one potential hierarchical framework and methodology for embedding culture within a virtual environment for use by an AI or behavior subsystem. This framework serves as an example for a more generic standard that may eventually be developed that incorporates social and cultural constructs in the minds of simulation entities. Similar to how the SEDRIS initiative addresses a standardized representation and interchange of environmental data (SEDRIS, 2006), this eventual standard would layout a cohesive methodology and process for incorporating culture into the M&S domain.

1.3 Definitions

This work draws on several areas of research and to understand the breadth and scope of embedding culture in virtual simulation environments, several definitions must be clarified.

Culture

According to Hofstede and Hoftstede (2005), culture is the collective programming of the mind that separates one group of people from another. Examples of culture include language, technology, economic, political and educational systems, religious and aesthetic patterns, and social structures (Triandis, 1989). The mental programming model presented by Hofstede (Figure 1.1) illustrates the relationship of culture with personality and human nature.

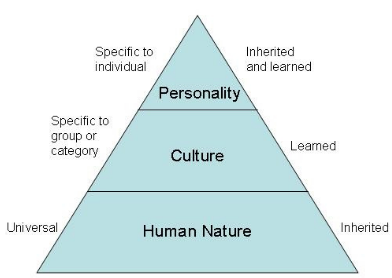

Figure 1.1: Three Levels of Uniqueness in Mental Programming (Hofstede and Hofstede, 2005)

As a core concept culture is learned and not innate, though the borders that delineate what we learn versus what we inherit is a matter of debate among social scientists (Hofstede and Hofstede, 2005). It also encompasses an entire group or category of people rather than an individual, though the group may contain a very small sampling of people (e.g. a militia).

The Agent

When the term agent is used by AI researchers, it typically (though not always) refers to a single entity, be it a squad commander, a fixed-wing aircraft, or a fire-team. According to Jacques Ferber (1999) and for the duration of this paper, an agent can be a physical or virtual entity that can act, perceive its environment (fully or partially), communicate with others, is autonomous and has skills to achieve its goals and tendencies. The key here is that the self is an active agent that promotes differential sampling, processing and evaluation of information from the environment, which in turn leads to differences in social behavior (Triandis, 1989). The self is capable of operating autonomously, or as part of a larger in-group. It may or may not share similar characteristics with other agents in the same environment.

The Group

The group is a much more complex concept and is defined as any number of entities (i.e. agents) considered as a single unit (WordNet, 2005). The key term here is *considered*, which we refer to as entities that share similar first-order cultural characteristics, such as the same political party or socio-economic status. These groups can further be divided into appropriate second-order categories that define their overall behavior such as individual or collective.

As described earlier, to make the effects of culture on groups possible in a real-time, computationally-constrained environment, generalizations must be characterized that may not be true across an entire population but are across a large cross-section. For example, not every Westerner places a high degree of importance on education, though the individualistic nature of that society may support that broad characterization to be made. The minutia of individualistic and collective societies is outside the scope of this paper; it is presented to help identify the similarities certain individuals within a common group share, and how they may be generalized across a large sampling for representation inside of a virtual environment.

Affordance Theory

Affordance theory is based on James J. Gibson's notion of perceived possibilities for an agent's actions (Gibson, 1979). Affordances, according to Gibson, are those environmental perceptions that allow an agent to meet an internal set of goals. This is most often exemplified with structural attributes of physical objects that humans innately recognize as being able to perform a certain function, such as sitting in a chair. Our approach is a derivative of Gibson's original premise such that physical objects not only have attributes which are innately recognizable but socially recognizable, such as going to school (to learn) or going to the post office (to mail a package). Additionally, when implementing a multi-layered cultural model where agents are not relegated to strictly strategic actions (i.e. symmetric simulations) a functional understanding of the objects in the environment is essential (Cornwell et. al, 2003). It should be noted that the context of affordance theory is often presented at a very low level, such as opening a door or picking up a coffee mug; however, modeling such low-level actions often results in poor performance and scalability when implemented in a virtual simulation system. Therefore, we have decided to represent the affordances and internal goals of agents at a much higher level, and leave out the lower-level details of action execution (i.e. opening a door to get into the house), which can sufficiently be executed using a traditional finite state machine (FSM) or script-based approach. The affordances specifically identified in our framework are those properties of an environmental feature (e.g. building) that can satisfy an agent's higher-level goals. For example, a bank `HasMoney()` and an agent that has a goal to `getMoney()` would match that affordance with satisfying the goal. However, an important distinction between this and the traditional affordance approach is the incorporation of culture into the overall goal-evaluation process, described in detail in Section 2.

1.4 Related Work

Research in the field of applying culture to M&S systems and virtual environments is minimal. The Political Geographical Religious Economic and Demographic Simulation (PGREDS) Modeling System is an example of applying these influences to the traditional simulation domain, specifically the OneSAF TestBed (Panagos, 2004). The focus here was on the compilation of disparate models that were then analyzed by a Resolver to produce a set of rules that dictate entity actions in the simulation. Though our technical approach is considerably different, we were able to draw on several areas of the PGREDS architecture to understand how cultural influences should be organized within an overall framework.

Analogous work in the virtual simulation (i.e. game) community is seen with *The Sims*, a real-time strategy game which allows players to create and control virtual characters that interact with one another. Using affordance theory, each agent perceives and acts within its environment based on emitted information from objects in the world, such as structures, vehicles, props and other agents (Cass, 2002). The affordance-based model used in *The Sims* not only enhances the variability of the AI but also provides a unique authoring environment whereby artists, programmers and even users define what objects are capable of emitting for use by the agents. This was a key feature in leading us to use an affordance-based approach for our system.

2. Technical Approach

There are two fundamental steps to embedding culture within the environment. The first is defining the framework of descriptors that exposes the appropriate cultural and social information to the behavior system in an organized manner. The decision to embed salient information inside the environment (versus the agent) was based on the well-established use of affordance theory in PMFServ (Cornwell et. al, 2003) and *The Sims* (Cass, 2002) and the need to create as much interdependence as possible between the AI agents and their environments. Additionally, by embedding classifying data in the objects of the terrain and dynamically mapping their associative uses to an agent, we are abstracting away the scripted details of each agent's execution. Drawing upon the SEDRIS Environmental Data Coding Specification (SEDRIS, 2006), a hierarchical data model has been designed to support the incorporation of first-order culturally-relevant data within the physical environment. This currently includes, but is not limited to: religious denominations, political affiliations, and socio-economic classes. Because culturally-specific data tends to be more qualitative and abstract than other environmental information (geometry, collision cylinders…) EDCS (SEDRIS, 2006) is well-suited for the encoding and communication of such information. Using its framework for classifications, attributes and enumerants, many of these cultural descriptors can be included as part of the physical environment's data representation. Additionally, because it represents a standard for environmental information, it allows behavior and AI subsystems to use a well-established programming specification for accessing and using the data.

The second step is to adequately annotate the simulated environment with the information defined in the framework above. Objects within the virtual environment are defined as pertinent features such as buildings, vehicles or props (trees, telephone poles). These objects have associated with them two sets of attributes: annotations and affordances. The annotations identify an object's first-order cultural descriptors while the affordances are indicators emitted from an object that can satisfy an agent's goals. This is a similar approach taken in PMFServ (Cornwell et al., 2003) and The Sims (McLean-Foreman, 2001) and adds variability and flexibility to the AI by moving some aspects of knowledge into the environment. As mentioned earlier, our innovation to this well established idea is to embed cultural annotations and higher-level affordances in the environment. This turns out to be a very natural approach to representing the cultural context of an agent's surroundings as people regularly associate cultural concepts with objects (a Sunni Mosque) and regions (a lower-class neighborhood).

The process of embedding annotations and affordances inside of a virtual environment begins with an authoring tool such as Maya. The modeler simply selects individual objects in the environment (buildings, vehicles, props…) and "tags" them with the desired annotations and affordances, analogous to *The Sims*.

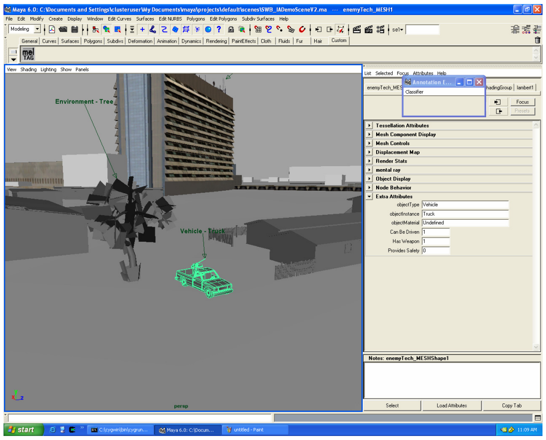

Figure 2.1: The "Tagger" for Maya

Once the objects have all been "tagged" the vertices of the ground-plane are "painted" with different descriptor values with a separate layer for each first-order descriptor (i.e. a socio-economic layer, a religious layer, a political layer…). Because the objects are placed on the ground (directly connected to the ground-plane), each of the objects inherit a set of default cultural annotations based on the underlying regional annotations (e.g., a middle-class Christian church). Once the level/terrain has been fully annotated, it is exported to the virtual environment where the information is then exposed to the underlying game engine. All of this embedded information is explicitly represented in the terrain models (i.e. static meshes) and are sensed by the agents through lookup tables (further described in Section 2.1).

At this point, agents placed in the virtual environment can poll nearby objects or the underlying ground-plane to access the cultural context of their surroundings. This mechanism is also modifiable during runtime by selecting individual or groups of objects and reclassifying their annotations, affordances or regional attributes as necessary (a lower-class community grocery store that is converted into an upscale restaurant), which is a useful feature when an anomalous structure or prop is placed in an area in which it typically wouldn't reside (e.g., a U.S.-controlled barracks in a predominately Sunni neighborhood of central Baghdad) and in large, heterogeneous urban areas where the diversity of in-groups are especially prevalent.

2.1 Implementation

Using the commercial Unreal Engine (v2.5) we have created a virtual environment that simulates a small urban scene populated with observably goal-directed, culturally-influenced human agents. The Unreal engine was selected because of its flexibility as a renderer and available toolset that includes a robust GUI toolkit and scripting language for creating and controlling the simulation. The engine also includes a large reusable code base including a robust and flexible set of classes to represent and control non-player characters (NPCs). The NPCs (i.e. agents) developed for this effort move about this urban environment using the embedded annotations and affordances to accomplish a set of predetermined goals. The relationship between the agents and the environmental annotations/affordances they use in their decision making is illustrated in Figure 2.2.

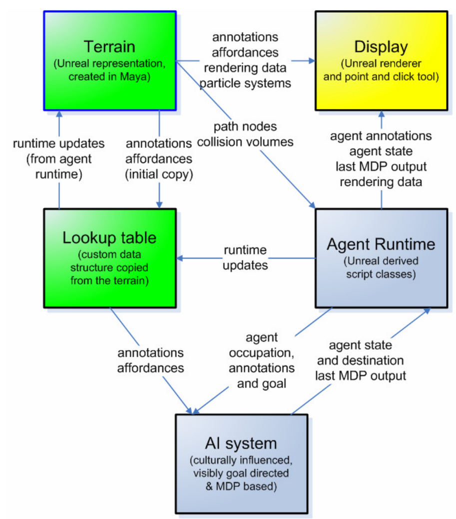

Figure 2.2: Agent-Environment I/O

Each of the agents has a limited knowledge base that consists solely of this predetermined goal set and their own physical/cultural attributes. They are initialized with no knowledge of the world and operate entirely on the annotations and affordances embedded within it. Individual agent goals include higher-level objectives such as `Work()`, `Eat()` and `GetMoney()` and are represented within a goal stack. Agent attributes include culturally relevant information such as age, gender, occupation, socio-economic status and political affiliation. Incorporating Triandis' research (1989), each agent belongs to one of several possible in-groups that make up the community in the simulation. Here an in-group is an aggregate of agents that share a similar political affiliation and/or socio-economic status. This representation has associated with it broad generalizations that influence either an agent's evaluation of a goal or its execution. Based on the in-group characterization, agents belonging to similar in-groups will use their cultural attributes to alter their behavior during runtime. For example:

- Agents belonging to more individualistic in-groups move more quickly, especially when they are executing a goal (Hofstede and Hofstede, 2005)
- Each in-group has a weight assigned to each of the possible goals. This weight represents the importance of accomplishing a goal for that particular in-group. The end result is that agents will be more or less likely to try to accomplish a goal based on their cultural background
- When not actively executing a goal an agent will return to a "friendly" area that matches the agent's cultural in-group completely or partially
- If possible, agents will seek to accomplish goals in areas that match their own cultural affordances
- Certain groups will not gather publicly with members of opposite gender

The actions in the simulation are finite state machines (FSM) that represent an agent's interaction with the environment: idle, execute, wander and gather. The idle action is used to carry out a scan of the environment for a specific annotation or affordance that satisfies one of the agent's internal goals. The execute action is called by an agent to move to a feature within the environment and satisfy a particular goal (`execute(Work)`). The wander and gather actions are used when an agent decides not to execute a goal due to cultural considerations: wander moves the agent to a specified location in a culturally friendly region and gather will move the agent to a specific point with agents that share the same or similar cultural affiliations. Though not entirely representative of actual human behavior, wander and gather allow us to introduce the notion of a community by congregating culturally-similar NPCs in the same region.

All of this information has been represented in the context of a Markov Decision Process (MDP) that determines each agent's behavior in the simulation. The policy graph for this MDP is illustrated in Figure 2.3. An important distinction between this and traditional MDP policies is the use of a randomization function that allows the traversal of the graph in several ways to add variety to agent behavior. The common notation for representing an action/state sequence in an MDP is: $T(s', a, s) = \Pr(s' | a, s)$ which is the probability of ending up in state *s'* given an action *a* in state *s*. The states in our model are represented as a tuple: `{goalValue, politicalVal, socioEconomicVal}`. The `goalVal` is the weighted value of a goal (`Work()`, `Eat()`) taking into consideration the agent's distance to the goal, a boolean whether the goal is located in a "friendly" region, and the cultural priority of the goal, which indicates how important the goal is given the agent's political and/or socio-economic affiliation. For example, attending a particular political event may be reserved solely for those individuals who belong to a certain political party.

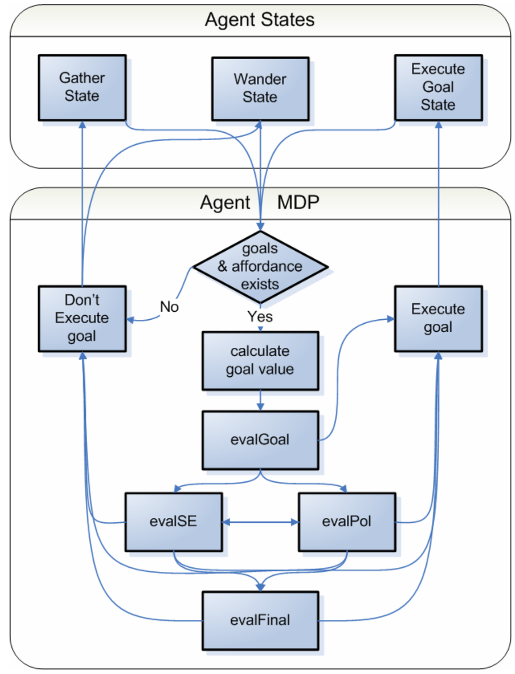

Figure 2.3: MDP Policy Graph

The political and socio-economic values (`politicalVal`, `socioEconomicVal`) are enumerations that contain what region the agent is currently in: 0—unknown region, 1—friendly, 2—hostile. The actions are the nodes in the policy graph (executeGoal, evalGoal, etc) that produce state values, transition probabilities are dictated by the policy to determine what edge is traversed given a particular state, and the reward function is the cost associated with traversing various edges of the graph.

When traversing the policy the agent will consider one and only one of its goals from the goal stack. If the agent has no goal it will wander or gather. Otherwise it will perform a scan of the world looking for features that can satisfy its current goal. If no appropriate feature exists the agent will enter the wander or gather action to move or remain in a culturally-safe region. If a feature does exist, the one with the highest goal value will be selected and all others will be ignored. The goal value is then used to put the goal into one of three enumerations: desirable (2), neutral (1) or undesirable (0). This enumeration is used by the agent's MDP to decide upon an action. Once an action has been decided the agents carry it out through FSMs that move them to their destinations using Unreal's built-in pathing system. Upon completion of a goal (which in this implementation is just a set time interval after arriving at the goal's location) the goal is popped from the goal stack and the next one is selected. The MDP is then initialized again to determine the appropriate course of action. In addition to polling for static annotations and affordances within the environment agents also have the ability to manipulate the affordances of buildings at runtime. This introduces into the environment the notion of history, which can be immensely useful for an agent's decision-making. For example, an agent with occupation "criminal" can commit a crime in a building that then adds to the building's affordance list `HasCrime()`. Police officer agents can listen for this affordance and respond accordingly and upon resolution of the crime (arrest, file police report,…) the affordance will be removed.

An additional feature that has been added to the system is a "point-and-click" GUI tool that allows a user to view cultural information about the environment. This includes a series of keyboard controls for displaying and hiding the regional annotations. When a layer is toggled on (see Figure 2.4) the ground-plane is illuminated with a color that corresponds to a particular cultural influence and affiliation. We have developed the system such that the agents behave according to the display of the colors. For example, if no regional influence is turned on the agents will not take that cultural parameter into account when deciding upon a behavior. However, if it becomes enabled they then must weight their MDP evaluation to account for this new cultural influence.

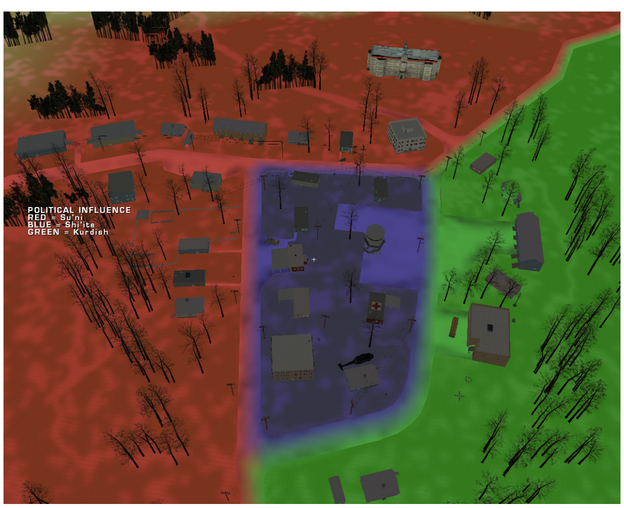

Figure 2.4: Culturally-annotated Regions

2.2 Results

The environment on which this system was demonstrated is a ~ 2km² area with ~30 structures. The regions are divided between 3 socio-economic groups: upper-class, middle class and lower class, and three political groups: Su'ni, Shi'ite and Kurdish. The city is divided up along the roads so that there is some variety in the combination of political and social groups.

An initial performance analysis shows where the current limits of the systems are and where improvements can be made. We are particular interested in the number of agents that can exist in the world while maintaining an acceptable frame rate. In early versions of the system the frame rate would be noticeably bogged down after only 20-40 agents were added to the world. Decreasing the size of the textures from 512x512kB to 128x128kB reduced the memory footprint by a factor of 16 and greatly improved performance. With the new textures we are able to include between 100 and 200 agents while keeping the frame rate above 50 fps. This is about number of agents necessary to give the relatively small town the look and feel of an active community. Figures 2.5 and 2.6 below illustrate the relationship between the number of agents running simultaneously in the virtual environment and the Framerate/CPU utilization.

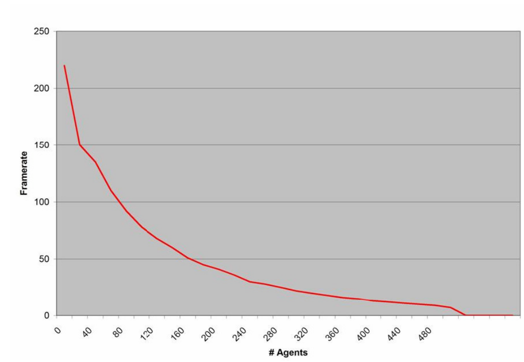

Figure 2.5: # Agents vs Framerate

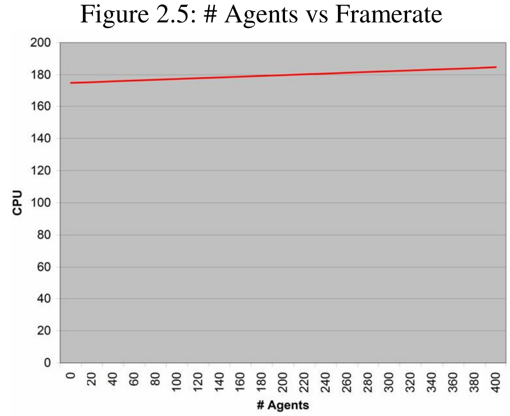

Figure 2.6: # Agents vs CPU Utilization

The results highlight two issues: 1) that the number of agents is nearly inversely proportional to the frame rate, and 2) that additional agents add almost no extra load to the CPU. These results, combined with our experience with the texture memory, seem to indicate that the current bottleneck is with the rendering of the agents. In a test to determine the absolute maximum number of culturally-influenced agents that could be included simultaneously in the scene everything was stripped from the agent but the MDP and the agent's I/O mechanism with the environment's annotations/affordances. This includes the removal of all rendering, collisions, animations and path finding. From here we were able add 500 agents to the simulation and still operate at over 50 fps. One possible future improvement will be to move functionality out of UnrealScript and into a faster native C++ implementation. This will speed up the AI cycle of the simulation, which may be important in supporting tens of thousands of culturally-influenced NPCs.

3. Conclusion & Future Work

Embedding cultural annotations and affordances within a simulated environment is only one approach to introducing culture into the AI. One of the most significant challenges with this approach is presenting the AI with relevant, contextual information that it can use without overburdening the system. For example, there are countless cultural meanings and representations for an object like a sword and attempting to model each is nearly impossible. Though recent research has enabled broad characterizations to be made that relate certain environmental attributes with certain types of human behavior (masculine/feminine, collective/individual), there is still a significant amount of research to be conducted to determine precisely how useful this approach is.

One of the near-term efforts of this research will be to test the flexibility of embedding cultural information in the environment by reusing the same MDP implemented here in an entirely different annotated location and evaluate its results without changing any of the existing AI codebase. Additionally, we plan on creating a more comprehensive set of first and second-order cultural descriptors and their relationships that are more representative of not only individual but group behavior. For example, the notion of the family carries with it varying connotations across a variety of different cultures, and it is important that the group and their associated annotations/affordances are sufficiently represented.

Finally, to promote the adoption and use of this framework across the M&S community, effort will be directed towards standardizing the representation of culture in an established data model. The EDCS provides the templates and extendibility for this and writing a comprehensive EDCS that is specific to embedding culture within the physical environment is a logical next step. The first and second-order descriptors described in this paper will mapped to other areas of research in culturally-derived behavior to form a broader classification scheme that may be used by a variety of AI systems, virtual environments and military simulations.

4. References

Baker, T. (2002). Game Intelligence: AI Plays Along. *Smart Computing, Vol 2(1),* 56-60.

Cass, S. (2002). Mind Games. *IEEE Spectrum, Vol 39(12),* 40-44.

Cornwell, J., O'Brien, K., Silverman, B., & Toth, J. (2003). Affordance Theory for Improving the Rapid Generation, Composability, and Reusability of Synthetic Agents and Objects. *12th Conference on Behavior Representation in Modeling & Simulation.*

Ferber, Jacques. (1999). *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*. Harlow: Addison Wesley Longman.

Gibson, J.J. (1979). *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Hofstede, G., & Hofstede, G.J. (2005). *Cultures and Organizations: Software of the Mind*. McGraw-Hill.

McLean-Foreman, J. (2001). *An Interview with Will Wright*. http://www.gamasutra.com/features/20010501/wright_01.htm

Panagos, J., Franceschini, R., Zhang, J., Comer, B., & Grosse, J. (2004). Political Geographical Religious Economic and Demographic Simulation Modeling. *Proceedings of the Fall 2004 Simulation Interoperability Workshop*.

Roth, M. (2005). Markov Models for Multi-Agent Coordination. http://www.cs.cmu.edu/~mberna/mrrg/notes/roth_mdps.ppt.

SEDRIS. (2006) Retrieved from the web February 14, 2006. http://www.sedris.org/ab_1trpl.htm. http://www.sedris.org/edcs.htm.

Triandis, H. (1989). The Self and Social Behavior in Differing Cultural Contexts. *Psychological Review, Vol 96 (3),*506-520.

Triandis, H., Bontempo, R., Villareal, M., Asai, M., & Lucca, N. (1988). Individualism and Collectivism: Cross-Cultural Perspectives on Self-Ingroup Relationships. *Journal of Personality and Social Psychology, Vol 54 (2),*323-338.

WordNet. (2005). WordNet Search- 2.1. Retrieved from the web January 24, 2006. http://wordnet.princeton.edu/.

Wunderle, W. (2005). Through the Lens of Cultural Awareness: Planning Requirements in Wielding the Instruments of National Power. *June 2005*

Presentation to USC's Institute for Creative Technologies.

4. Acknowledgements

The project or effort described here has been sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

Author Biographies

RYAN MCALINDEN is a Lead Software Engineer at USC's Institute for Creative Technologies. His interests and experience are in several areas of software development and research, including commercial game technology, AI, terrain and military simulations.

WILLIAM CLEVENGER is a Research Programmer at USC's Institute for Creative Technologies. He recently completed his degree in computational mathematics at the University of California, San Diego.