Using Environmental Annotations and Affordances to Model Culture

Ryan McAlinden, Michael van Lent, William Clevenger, Wen Tien

USC Institute for Creative Technologies
13274 Fiji Way
Marina del Rey, CA 90292
{mcalinden, vanlent, clevenger, tien}@ict.usc.edu

Abstract

This paper details the demonstration of an annotation and affordance-based software model intended to introduce cultural and social influences into a non-player character's (NPC) decision-making process. We describe how recent research has supported the need to begin incorporating the effects of culture into the interactive digital domain. The technical approach is presented that describes the software techniques for embedding and utilizing culturally-specific information inside of a virtual environment, as well as the design and implementation of a deterministic Markov Decision Process (MDP) to model the affects of culture on the AI.

Motivation

Commercial games focus on providing the user with an engaging and immersive virtual experience, often with somewhat simple AI mechanisms such as finite state machines and linear scripts. These static techniques typically rely on primitive environmental information to aid decision making such as collision volumes, node networks and basic NPC information (position, speed, team affiliation). However, as game complexity has increased non-player character (NPC) behavior inside of these environments are being forced away from these static techniques, such as seen in Black & White 2's experience model (Pottinger 2000). Collision volumes and pathnode networks may be sufficient for path-planning and creating basic associations between elements in the environment, but it is the patterns, landmark references and cultural influences that help shape real human actions (movement, gestures), perceptions and emotions, all of which are absent from games today (Hofstede and Hofstede 2005). One exception is *The Sims* in which the surroundings convey to the AI more goal-directed information (i.e. affordances) to assist it in meeting certain objectives (Cass 2002). The relationship that exists between this embedded environmental information and how it is used to satisfy a set of goals is a derivative of affordance theory (Gibson 1979). The advantages and disadvantages of this approach are addressed later but the premise, derived from Gibson, is that environmental surroundings convey to humans certain information that directly map to their decision-making process (Gibson 1979). Environment in this context refers to the objects (static or dynamic) in a game that traditionally make up the level: buildings, ground-plane, props, features, vehicles, etc. The approach we present here is similar and relies on explicitly embedding cultural information inside of the environment rather than in the agent's knowledge base or memory.

It may be difficult to acknowledge but our biases, stereotypes and inability to accept races, religions, or ethnicities aside from our own dramatically affect how we sense and act on a daily basis: "what is different is dangerous" (Hofstede and Hofstede 2005). Many of our strongest beliefs and ideals are rooted in cultural and social influences (Hofstede and Hofstede 2005) and allowing players to interact with NPCs that include the cultural nuances of human behavior can significantly enhance the immersion and entertainment of games. For example Beverly Hills and South Central Los Angeles, although not geographically far apart, have very different cultural contexts. Most people (no matter their race or socio-economic status) would act and interact with the environment very differently in these two different contexts. This introduces the notion of culturally-induced environmental effects on human behavior, which we model as cultural annotations; the socio-economic, political, ethnic, historical, episodic, demographic, religious… information associated with an agent's environment that are important in determining how/why they behave the way they do. Our approach is to embed cultural information about the environment in the virtual environment rather than representing this information in the agent's knowledge or memory.

This approach is inspired by *The Sims* which embeds information about how a character can interact with an object and how that interaction affects the character in the model of the object. There are four fundamental differences between our approach and *The Sims:* 1) we add cultural aspects to the information embedded in the environment, 2) we introduce a model of how cultural information affects behavior, 3) we embed information not only in individual objects but in ground regions, and 4) we add the ability to modify cultural parameters to increase the variability of NPC actions.

Another unique feature of this approach is the support for dynamic annotations and affordances that can be modified during runtime. This adds to the environment (and in-turn the agent) the notion of events or history that may alter the AI's behavior. For example, if a faction of AI agents from one political affiliation overtakes an area of town that was previously controlled by another political affiliation the regional annotation for these areas can dynamically be recalculated including storing the previous annotation. This allows episodic or historical information to be embedded in the environment that can then be incorporated into the AI's decision-making. Similar support for singular annotations is also possible (e.g. a bank that has just been robbed emits a HasCrime() affordance that is broadcast to police officers).

There are several implications, both positive and negative, for embedding annotations and affordances in the environment versus directly in the AI's memory. 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 3D environment cultural annotations and affordances on objects and regions of the ground. Lastly, dynamic annotations and affordances 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 (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. 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.

Technical Approach

The first step in embedding culture within the AI is to markup the environment. This is accomplished by an artist or modeler "tagging" the level (i.e. buildings, props, ground-plane) with cultural annotations. For example, regions may be drawn and buildings marked to delineate higher vs lower socio-economic neighborhoods. In addition to these cultural attributes, environmental objects have associated with them traditional annotations and affordances as is seen in *The Sims* such as what type a building is (bank, post office) and what they afford an agent (hasMoney(), mailLetter()).

This markup process is performed during the construction of the environment in a 3D authoring tool such as Maya. An artist or modeler launches the "Tagger" tool and selects existing groups of geometry (e.g. buildings) to classify with individual annotations or affordances. Once each of the objects has been annotated, the vertices of the ground-plane are "painted" in Maya, with each color representing a regional cultural attribute (e.g. a neighborhood with a particular religious affiliation). Once tagged, the environment is exported to the virtual environment where the information is then exposed to the underlying game engine (in our case Unreal Engine v2.5). All of this embedded information is explicitly represented in the terrain models (i.e. static meshes) and are sensed/accessed by the agents through subscription channels. This approach allows any AI system with an I/O mechanism to integrate these object and regional annotations with relative ease. We have prototyped an initial AI system using an MDP that incorporates both the agent's and environment's cultural attributes to alter its decision-making.

The response to this system has been primarily voiced by content creators of the terrain/level (artists) and AI developers. The Tagger tool was well-received and does not present an undue burden on the artist creating the level. The AI developer has found that once the culture model was represented in code, there was extensive variability possible in entity behaviors simply by throttling the probabilities of the transition table and/or dynamically modifying the embedded annotations and affordances of the environment. An area of future work includes engineering a POMDP solution that incorporates observational uncertainty and is more representative of the way humans actually perceive their environment.

References

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.

Cass, S. 2002. Mind Games. *IEEE Spectrum, Vol 39 (12),* 40-44. http://ieeexplore.ieee.org/iel5/6/23652/01088444.pdf?arnumber=1088444

Pottinger, D. The future of game AI. Aug 2000. *Game Develop. Mag. Vol 7 (8)*, 36–39.