Evaluating a Framework for Representing Cultural Norms for Human Behavior Models

Steven Solomon
Matthew J. Hays
Grace Chen
Milton Rosenberg

Institute for Creative Technologies
University of Southern California
13274 Fiji Way
Marina del Rey CA 90292

Keywords: Human Behavior Model, Cultural Knowledge Representation, Evaluation, Cultural Norms, Cognitive Schemas, Learning, Worked Example

ABSTRACT: *Cultural awareness training is seen as a necessity in the military, in international business, and in diplomacy. The Culturally-Affected Behavior project has defined a framework for encoding cultural norms and values that facilitates the creation of human behavior models having cultural knowledge separate from domain knowledge. We evaluated a simulation based on the framework as a tool for learning cultural norms. Users were provided a worked example of a meeting with a first virtual character as a training session, and were subsequently able to distinguish appropriate socio-cultural actions from inappropriate actions in a meeting with a second character of the same culture, and in a judgment survey.*

1. Introduction

A promising application of immersive environments is cultural awareness training. Users interact with virtual characters from a foreign culture and learn through the interactions which interpersonal and social actions are considered appropriate and inappropriate in the culture. Culturally-Affected Behavior (CAB) (Solomon, van Lent, Core, Carpenter, & Rosenberg, 2008) is a framework for modeling cultural norms and values that facilitates the creation of human behavior models. CAB provides a language for ethnographers to encode cultural norms and values as well as biases and stereotypes. The resulting cultural models are used to affect the responses of human behavior models. In human behavior models that use CAB, socio-cultural sources of behavior are explicitly and modularly represented. Critically, these models' cultural knowledge is separated from their domain knowledge (i.e., a model's understanding of its societal role).

The focus of the present research effort is to describe and validate a simulation based on the CAB framework as a tool for teaching cultural norms and values. The simulation consists of two virtual characters with a shared culture model but different task models (i.e., they have different domain models). We seek to show that users interacting with this CAB-based simulation can learn which actions are appropriate and inappropriate in the characters' culture.

Section 2 describes an AI agent architecture that CAB extends (the Virtual Humans architecture) and social science theories on which CAB is based. The aspects of culture on which the CAB framework focuses are socio-cultural norms and culturally shared mental representations of others' beliefs. Our extension of the Virtual Humans architecture and the CAB simulation prototype are described in Section 3.

In our evaluation of learning, we exposed participants to two CAB human behavior models having the same culture. The interaction with the first character was a training session, in which the participants followed a worked-example script consisting of all socio-cultural actions. The worked-example script served as a step-by-step demonstration of how to conduct an entire meeting with the first character. Participants received responses from the human behavior model for each action. The training session was followed by the negotiation meeting with the second character, in which participants chose which actions to perform. Finally, the participants filled out a judgment survey on the actions, rating them as appropriate or inappropriate in the CAB characters' culture. A second (control) group of participants performed the training session and the negotiation meeting, and filled out the final judgment survey, but did not have the characters' responses to use as a basis for judging the actions. Section 4 details the evaluation of learning and its results.

2. Related Work

2.1 Virtual Humans

CAB is based on Virtual Humans (Gratch & Marsella, 2004), an AI architecture for human behavior modeling. Virtual Humans is a cognitive model; its major components—the task model, the emotional model, and the dialogue model—are based on theories of human cognitive constructs. The emotional model of the Virtual Humans has been assessed for its consistency with human emotional responses (I/O validity) and for consistency with human inference (process validity) (Gratch, Marsella, & Mao, 2006).

The task model of Virtual Humans is used to represent domain plans and the underlying world state from the perspective of the human behavior model. In the Stability and Support Operations application of the Virtual Humans, the user plays the role of a U.S. Army Captain who must persuade a local doctor (the Virtual Human) to move his clinic to a new location because of a nearby skirmish. The doctor character's domain goals include helping patients and remaining neutral in the conflict. In the negotiation, users can offer the doctor resources (e.g., transportation and medical supplies).

The task model formalism, based on a decision-theoretic framework, determines the human behavior model's choices. The human behavior model consists of multiple task models called plans. All plans are authored before runtime and achieve a goal state (e.g., "patients-are-helped"). The model includes a default plan and one or more alternative plans. These plans have utility values. The user's objective is to get the utility value of an alternative plan ("run-clinic-there") to be greater than the utility value of the default plan ("run-clinic-here") so that the human behavior model switches to the alternative plan the user desires.

Task models consist of tasks, states, and effects. *Tasks* represent simulation actions which can be performed by the user or the human behavior model. Tasks may be performed immediately in the simulation or in the future. *States* have intrinsic utility values that represent the relative importance that the human behavior model has for the state (weighed against other states). Tasks have *effects* on states. Effects have a sign (+/-) which represents whether a task reinforces or detracts from a state. Each task has a probability, which represents the likelihood that the entity responsible for the task (either user or human behavior model) will perform the task. One factor in the model's calculation of a task's probability is intention probability. *Intention probability* represents how confident the human behavior model is in the user to perform future, agreed-upon tasks. Intention probability is calculated from trust, which is determined outside of the task models.

A state's *current utility value* represents the immediate importance of that state to the human behavior model. The current utility value is based on the intrinsic utility value of the state, the tasks that affect the state, and effects of the tasks on the state. The utility value of a plan as a whole is calculated as the sum of the current utilities of all of that plan's states. The human behavior model begins with an initial, default plan. Users' actions affect the probabilities of various tasks, which in turn update the utility value of the states (and thus the default and alternative plans). The basic task model architecture of Virtual Humans is shown in Figure 1.

Cultural knowledge in Virtual Humans-based domains is embedded throughout the dialogue model and the task model, but not in a modular fashion. The rules,
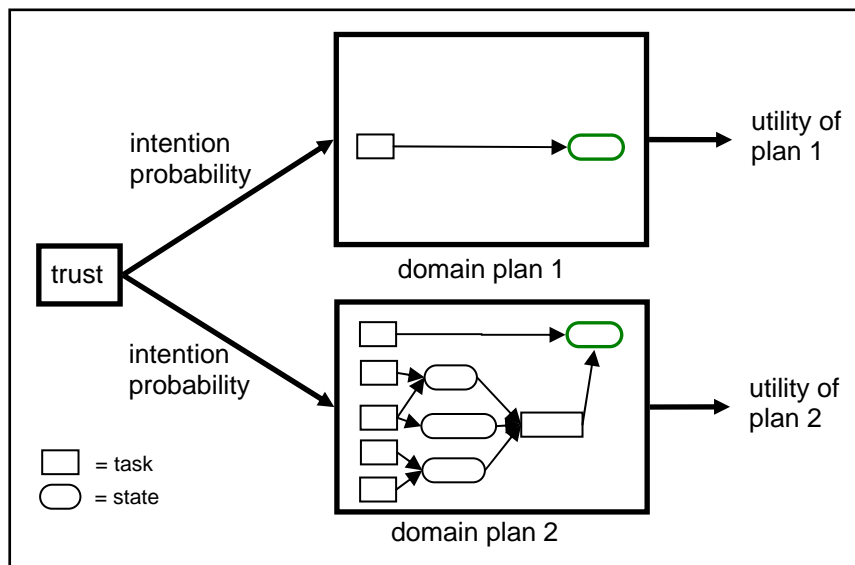

Figure 1: Virtual Humans Architecture

data, and other software constructs that implement the socio-cultural aspects of the human behavior model's knowledge are not segregated from the constructs that model the domain knowledge. Section 3 describes how socio-cultural knowledge is represented in CAB.

2.2 Social Science

The cultural knowledge representation of CAB is grounded in schema theory and the theory of mind. *Schemas* are codified forms of an abstract behavior or concept. In *schema theory*, culture is represented as a collection of schemas that is shared among members of a culture (D'Andrade, 1992). A *cognitive schema* is an organization of knowledge around a particular concept which is triggered by a word, symbol, or image. The *constitutive rules system*, a type of cognitive schema, represents shared conceptual meanings (D'Andrade, 1984). A constitutive rules system is defined as a set of if-then rules that defines some concept. The execution of certain behaviors (i.e., tasks) imputes a cultural state that relates back to the concept. For example, the constitutive rules system for Islam (the religion) contains rules such as "if X drinks alcohol, then X is not observant of Islam" and "if X gambles, then X is not observant of Islam". This type of cognitive schema is triggered when someone observes or performs any of the actions associated with the schema.

The *theory of mind* posits that individuals ascribe mental states (e.g., intentions, beliefs, and values) to themselves and to others (Nichols & Stich, 2003). In a social interaction, individuals generate two types of mental representations. The first representation consists of the individual's own knowledge, intentions, beliefs, and values. The second representation, which is contained in the first, is a belief model of the *other* individual's knowledge, intentions, beliefs, and values. Both representations are used in generating behaviors. In the absence of personal information, many of our initial beliefs about another's mental processes, beliefs, and intentions are based on stereotypes about the other person's culture. Thus, modeling theory of mind as applied to culture is an approach to modeling cross-cultural stereotypes and biases.

3. Modeling Socio-Cultural Norms and Generating Culturally Affected Behavior

CAB extends the Virtual Human task model formalism to model constitutive rules systems by creating cultural models ("socio-cultural networks") that consist of socio-cultural tasks, states, and effects, similar in construction to a task model. In CAB, the task model formalism becomes a language for expressing socio-cultural norms—a key aspect of cultural knowledge. The socio-cultural network overlays the default and alternative domain task models; together, a domain task model and the socio-cultural network form one culturally specific task model. By swapping the socio-cultural network with a new socio-cultural network, it is possible to achieve a degree of cultural knowledge modularity: maintaining the same domain human behavior model but with a different culture model. As in the Virtual Humans architecture, the human behavior model executes the default plan ("perform-normal-police-duties-plan") until the current utility value of the alternative plan ("help-participant-fix-market-problem-plan") exceeds the current utility value of the default plan.

The task model formalism is also used to model mental states as posited in the theory of mind. By default, the states and utilities of a task model associated with a human behavior model are from the point of view of the human behavior model. States[1] such as "agent's-view-that-community-is-helped" and "agent's-view-of-the-participant's-view-that-community-is-helped" model the human behavior model's own mental representation of a state of the world and his perception of the *user's* mental representation of a state of the world.

Each tuple consisting of (task, effect, state) in a socio-cultural network models a rule of the constitutive rules systems approach to representing cultural concepts. For example, the tuple ("give-alcohol", -0.95, "agent's-view-that-agent-is-observant-of-Islam") models the rule "drinking alcohol detracts from being observant of Islam" and the tuple ("show-picture-of-wife", -0.5, "agent's-view-that-agent-is-respectful-of-modesty") models a rule "looking at pictures of women in Western clothing detracts from being respectful of modesty" for the Iraqi-Sunni concept of modesty. In other words, looking at a picture of a woman in Western clothing causes the Iraqi-Sunni human behavior model to consider himself less respectful of modesty. The task model quantifies the constitutive rules systems approach by assigning probabilities to the effects and utility values to the states, which allows an author of a cultural model to reflect the fact that, within a culture, some effects and norms are more important than others. The tuple expression (task-name, (+/-) effect-probability, state-name) constitutes the grammar of a language for representing cultural norms and mental representations for a given culture.

---

[1] For state identifiers, we observed the naming convention that "agent" refers to the human behavior model and "participant" refers to the user. Thus, "agent's-view-that-agent-is-*X*" encodes the agent's view that he himself is *X*.
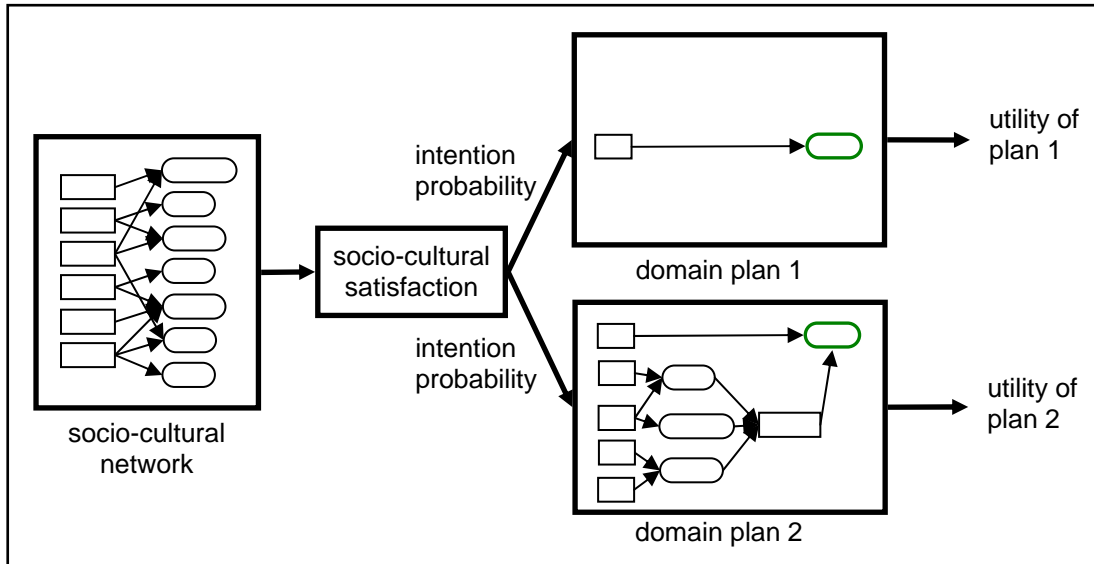
Figure 2: CAB Architecture

An important output of the socio-cultural network is the derived utility of the network, which we call the *socio-cultural satisfaction* (SCS) of the human behavior model. SCS represents the human behavior model's appraisal of the current situation against the set of its own socio-cultural norms, and is calculated in the same way as a current utility value in a domain plan. CAB uses SCS to affect the utility of domain plans, in that it is used in place of trust in the Virtual Humans to calculate intention probability, one of the factors in the calculation of task probability in the Virtual Humans. The lower the value of intention probability, the less likely these tasks are to happen, which in turn lowers the current utility values of states affected by the tasks and the overall current utility value of the associated domain plan. Thus, the overall effect of low SCS is to make the human behavior model less likely to agree to courses of action that the user wants. The architecture of CAB is shown in Figure 2.

The CAB prototype consists of the CAB language, the extension of the Virtual Humans behavior model, and culture models integrated with the graphics front-end and animation and voice capabilities of an existing training simulation environment. The extended Virtual Humans architecture, task models, and culture models are implemented in Java and Jess (a rule engine).

We have implemented socio-cultural networks for a number of cultures. For the prototype with character graphics and animations, we have implemented an Iraqi-Sunni culture model and a German culture model (based on non-expert data). For the evaluation of learning reported in the present paper, we have implemented a model of Appalachia Scots Irish culture

(as described in Section 4). CAB also includes default domain plans and alternative domain plans for two human behavior models: a police captain and a doctor.

In the prototype, the human behavior model implements a police captain character overlaid with either the Iraqi-Sunni or German socio-cultural network. The user can choose to meet with either the Iraqi police captain, named Farid, or the German police captain, named Fritz. Currently, the user interacts with only one character at a time. However, there is no limitation in the CAB architecture that would prevent the user from simultaneously interacting with multiple characters. When the human behavior model is loaded, a data file containing the police captain task model is loaded, followed by a data file containing the socio-cultural network. By loading a different socio-cultural data file, CAB is able to swap the culture of the human behavior model. Last, a file for the individual virtual character is loaded. It contains the character's dialogue utterances. Figure 3 shows screenshots of the Farid and Fritz police captain characters in the prototype.

To interact with the character, the user selects predetermined actions from a set of menus. The police captain—independent of culture—begins the meeting with the execution of a default plan ("perform-normal-police-duties-plan"). The user's goal is to persuade the police captain to provide police cooperation in solving a problem with a local market. As the user and police captain perform a mix of domain and socio-cultural actions (tasks), the current utility value (as influenced by the SCS of the socio-culture network) of the police captain's alternative plan, "help-participant-fix-market-problem" is updated (as is the value of the default plan). When the derived utility of the alternative plan

Figure 3: Farid and Fritz: virtual characters with an identical task model but different cultures

exceeds the utility of the default plan, the police captain agrees to cooperate with the user.

4. Learning Evaluation

We set out to validate that users can learn to distinguish positive and negative meeting actions from the point of view of a common culture shared by the simulation characters with whom the users interact. The ability to distinguish positive actions (those yielding a net utility value increase and eliciting favorable responses) from negative actions (those yielding a net utility value decrease and eliciting unfavorable responses) is a critical part of cultural awareness learning. Cross-cultural awareness is itself part of social- and meeting-skills learning.

The evaluation consisted of three parts: a training session (in which participants followed a worked example of a complete meeting with a first character); a negotiation meeting (in which participants chose actions to perform in a meeting with a second character of the same culture as the first character); and a judgment survey (in which participants rated all actions as "good" or "bad" from the point of view of the culture of the characters they encountered). Worked examples (i.e., step-by-step procedures for performing a task) are commonly found in intelligent tutoring systems. For example, Conati and Carenini (2001) describe a system that generates worked examples that are tailored to a student's competency level. The *worked-example effect* holds that learners that have studied worked examples perform better in problem solving than learners that have not studied prior worked examples (Cooper & Sweller, 1987).

The evaluation involved a fictitiously named culture (Shiri-Stawks) and two officials from this culture: Shel, a police captain, and Shir, a doctor. In actuality, the Shiri-Stawks culture was modeled in the CAB framework after the Appalachia Scots-Irish culture, using Webb (2004) as a reference. The decision to

fictionalize the name of the culture was based on the need to remove any predisposed bias or knowledge that participants may have had about the actual culture.

Although the CAB prototype integrates graphics and animation with culture models, the CAB interactions in the present study were text-based; participants selected text actions and received text responses. By withholding visual character cues in the form of dress, physical features, and accents, we removed the possibility that participants' prior biases would influence their judgments of the characters' culture.

Participants were divided into two groups. In the experimental group, participants read lists and made selections of meeting actions, interacted with the CAB characters, and then filled out a final survey. In the control group, participants made the same evaluations and selections of actions (and completed the survey), but did not interact with characters in the CAB simulation environment.

4.1 Hypothesis

We thought that users interacting with CAB characters who have different domains—but a common culture—could learn to distinguish positive from negative socio-cultural meeting actions (from the point of view of the culture). Without this interactivity, we thought that control participants would be unable to distinguish positive from negative socio-cultural actions.

4.2 Method

Participants

Participants were 37 employees of USC's Institute for Creative Technologies. Seventeen participants (four women, 13 men) were assigned to the experimental condition and 20 (five women, 15 men) were assigned to the control condition.

Materials and Procedure

The CAB prototype consists of the CAB language, the extension of the Virtual Humans behavior model, and

culture models (Solomon et al., 2008). The CAB simulation environment consists of virtual characters representing non-U.S. cultures. The United States Army has sent out a negotiator (the role played by the participant) to meet and negotiate with the characters to reach an agreement to solve a current problem.

The Shiri-Stawks culture model consists of 36 tasks, 35 states, and 68 effects. Some of the defining characteristics are: disdain for outside authority and elites, respect for the military and gun rights, competitiveness, and a strong sense of personal honor. An example task which may be performed by the human meeting participant is "give a bottle of expensive French wine". This task has a positive effect on the related states "agent's-view-that-agent-is-elitist" and "agent's-view-of-participant's-view-that-agent-is-elitist", both of which have negative intrinsic utility values (i.e., they are considered undesirable by the human behavior model). In other words, giving the virtual character a gift of expensive French wine causes the character to view himself as an elitist, and also causes him to believe that the participant thinks he is an elitist. This task also has a positive effect on the state "agent's-view-that-participant-is-elitist," meaning that the human behavior model views the participant (the French wine giver) as an elitist. Finally, the task has a small positive effect on the state "agent's-view-that-agent-is-familiar-with-participant", which has a positive intrinsic utility value. That is, the act of giving a gift is seen as a behavior that results in two people becoming more familiar with each other. Other example participant tasks in the Shiri-Stawks model include "talk about your family's long history of military service" (positive effect on "agent's-view-that-participant-is-connected-to-the-military") and "brag about your home-town sports team" (positive effect on "agent's-view-that-agent-is-competitive").

Training Session
The experimental procedure had three steps: a training session, a negotiation meeting, and the CAB Estimated Learning (CABEL) judgment survey. In the training session, experimental participants were told that they are negotiators who will need to speak with a Shiri-Stawks citizen. Because no one is familiar with the Shiri-Stawks culture, experimental participants were informed that they would practice interacting with Shel, a police captain from Shiri-Stawks, in a mock negotiation before going into the real negotiation with Shir, a doctor from Shiri-Stawks. The mock negotiation consisted of a complete, worked example of a negotiation session in which the participant was guided step-by-step through an interaction with Shel. Experimental participants were given a form that listed all socio-cultural and domain actions available in the CAB prototype. They were instructed to execute each

action in the indicated order and to read and evaluate Shel's responses to grasp the Shiri-Stawks culture. Experimental participants were told to notice which submitted actions created favorable responses from the character and which submitted actions created unfavorable responses. Since the character's responses were conveyed by text only, participants used their own judgment to determine which responses were favorable and which were unfavorable. The training session form contained all of the possible socio-cultural and domain actions for the police captain character, of which all participants were unaware.

Control participants received the same information and materials, but did not use the CAB prototype. Instead, they read and wrote a summary next to each action on the sheet. Their summaries indicated why they thought each action would be a good or bad idea to execute. They were instructed to base their assessments on common sense; they were told that they had never met a Shiri-Stawks person and had no way of knowing anything about the Shiri-Stawks culture. Because their experience was entirely paper-based, control participants never received feedback from the CAB system about the decisions they made.

Negotiation Meeting
After the experimental participants had gone through the entire script and submitted all of the actions in the training session (with Shel), they began the negotiation meeting (with Shir). Experimental participants were informed that doctors were being attacked and medical supplies were being stolen. They were again provided with a sheet of paper that included all possible actions, but the actions were now separated into seven different, ordered groups. For example, greeting actions were grouped and they preceded negotiation actions. Each group contained at least one socio-cultural action or domain action with a negative utility (as determined in the cultural model). From each group of actions, experimental participants were instructed to submit a subset of the actions to the CAB prototype to interact with Shir. Each group of actions was preceded by instructions that specified how many actions were to be submitted from that group. Experimental participants were allowed to choose the actions in any order within the group, as long as they finished submitting the actions in that group before moving on to the next group of actions. When experimental participants chose to submit an action, they were to make a mark next to that action on a sheet of paper that was given to them immediately before they submitted the action. Upon submitting each action, they were to attend to Shir's reaction. They based their action selection on our explanation that Shir and Shel—as Shiri-Stawks—share the same cultural norms and values; they were to choose the actions to which they thought Shir would
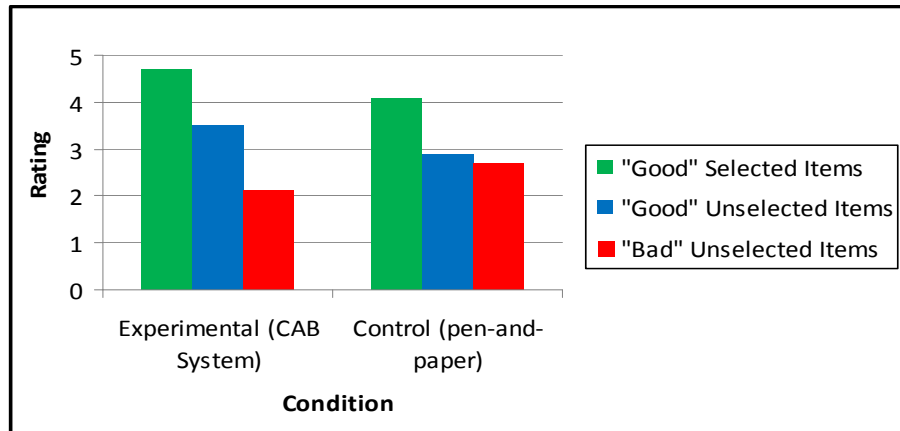
Figure 4: CABEL Judgment Survey Ratings (by Condition)

respond *most favorably*. When they had submitted the required number of actions from every group, the negotiation meeting ended.

Control participants were instructed to maintain the same mindset as they had during the training session. They were instructed to choose the actions to which they thought a Shiri-Stawks person would respond most favorably. Again, they received no feedback from the CAB system on the decisions they made.

CAB Estimated Learning Judgment Survey
After the negotiation meeting, participants completed the CAB Estimated Learning (CABEL) Judgment Survey. This survey listed all of the socio-cultural actions available to users of the program. In the survey, participants were asked rate each action from the point of view of the Shiri-Stawks culture when performed in a negotiation meeting with a Shiri-Stawks official. They provided their ratings on a scale from one to five (i.e., 1 = a bad action, 3 = a somewhat bad, but also somewhat good action, 5 = a good action). When they completed the survey, they were debriefed and thanked for their participation.

4.3 Results and Discussion
As one would expect, experimental participants rated the positive-utility actions that they selected (M = 4.701, SD = 0.241) as better than the positive-utility actions that they did not select (M = 3.521, SD = 0.957): $t(16)$ = 5.361, p < .001. Control participants also rated the positive-utility actions that they selected (M = 4.090, SD = 0.362) as better than the positive-utility actions that they did not select (M = 2.936, SD = 0.504): $t(19)$ = 12.329, p < .001. It stands to reason that, had they thought the unselected actions were better, they would have selected them, instead.

Participants chose to execute actions with negative utility values (as specified in the cultural model) too infrequently for meaningful statistical analyses to be performed on those data. As a result, the measure of interest becomes participants' ratings of unselected actions with either a positive or negative utility value on the CABEL Judgment Survey. We compared experimental participants' ratings to those provided by the control participants. Our goal was to determine whether interacting with the CAB prototype improved experimental participants' ability to discriminate between "good" and "bad" socio-cultural actions.

The experimental participants rated the (unselected) negative-utility actions (M = 2.108, SD = 0.356) as worse than the unselected positive-utility actions: $t(16)$ = 6.371, p < .001. On the contrary, control participants did *not* rate the (unselected) negative-utility actions (M = 2.733, SD = 0.386) as statistically significantly worse than the unselected positive-utility actions: $t(19)$ = 1.831, p = .08. This difference between control and experimental participants' ratings of the negative-utility actions versus the unselected positive-utility actions was statistically significant: $F(2, 70)$ = 34.147, MSE = 7.211, p < 0.001. Figure 4 illustrates these results.

These data support our claim that experience interacting with the CAB prototype was sufficient to enable experimental participants to distinguish negative socio-cultural actions from positive ones. This ability to discriminate extended to items that they did not choose to execute in the negotiation meeting. The control participants were exposed to all of the actions and encouraged to determine why some might be preferred to others. As a result, neither response familiarity, time on task, nor attention to content can explain their inability to make the same discrimination.

5. Conclusion and Ongoing Work
Human behavior models are an important element of effective simulations for cultural awareness training. Presently, we are integrating human behavior models into simulations that provide Army officers the opportunity to practice negotiating and interacting

appropriately with meeting partners from other cultures. Our human behavior models incorporate cultural models based on a framework for encoding cultural knowledge in the areas of norms, biases, and stereotypes. The CAB framework provides a common representation of cultural knowledge that is separate from a human behavior model's domain knowledge representation. We have described the ability to modify a human behavior model's culture by "swapping in" the socio-cultural elements of a model.

In our evaluation of learning, we have shown that users interacting with a CAB-framework simulation consisting of virtual characters (a police captain and doctor) with a shared culture model, can learn which actions are appropriate and inappropriate in the characters' common culture. Thus, the CAB framework is an effective tool in building simulations to be used for cross-cultural training. We have been evaluating how effective this tool is (e.g., comparing experience with CAB to other types of training).

In our ongoing efforts, we seek to determine whether the CAB framework is expressive and intuitive enough for ethnographers to encode cultural knowledge in the form of a cultural model that we have described. We have been working with ethnographers who have expertise in Mainland Chinese culture and Japanese culture to create models that will be integrated into human behavior models. From these experts, we also hope to discover expressive abilities that are lacking in the current CAB framework. One such area is cultural event schemas, which are script-like sequences of actions that are culturally-specific and apt to occur under certain conditions and in certain scenarios. Another area of exploration is culturally-specific dialogue models (i.e., dialogue knowledge in the cultural models that is used to generate dialogue and, in particular, culturally-specific speech acts). In the next phase of research, we plan to integrate these abilities into the CAB framework.

6. Acknowledgements

7. References

Conati, C., & Carenini, G. (2001). Generating Tailored Examples to Support Learning via Self-Explanation, *Proceedings of the 17th Joint International Conference on Artificial Intelligence* (IJCAI 2001), Seattle, WA.

Cooper, G., & Sweller, J. (1987). Effects of schema acquisition and rule automation on mathematical problem-solving transfer. *Journal of Educational Psychology*, 79, 347-362.

D'Andrade, R. (1984). Cultural meaning systems. In R. Schweder and R. Levine (Eds.), *Culture theory: Essays on mind, self, and emotion*. Cambridge: Cambridge University Press.

D'Andrade, R. (1992). Schemas and motivation. In R. D'Andrade and C. Strauss (Eds.), *Human motives and cultural models* (pp. 23-44). Cambridge: Cambridge University Press.

Gratch, J., & Marsella, S. (2004). A domain independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 5, 269-306.

Gratch, J., Marsella, S., & Mao, W. (2006) Towards a Validated Model of "Emotional Intelligence". *Twenty-First National Conference on Artificial Intelligence* (AAAI06). Boston, MA.

Nichols, S., & Stich, S. (2003). *Mindreading: An integrated account of pretence, self-awareness, and understanding of other minds*. Oxford: Oxford University Press.

Solomon, S., van Lent, M., Core, M., Carpenter, P., & Rosenberg, M. (2008). A Language for Modeling Cultural Norms, Biases, and Stereotypes for Human Behavior Models. *Paper presented at the 17th Behavior Representation in Modeling & Simulation (BRIMS) Conference*, Providence, RI.

Webb, J. (2004). *Born Fighting: How the Scots-Irish Shaped America.* New York: Broadway Books.

Author Biographies

STEVEN SOLOMON is a Research Programmer at the ICT.

MATTHEW J. HAYS was a Visiting Researcher at the ICT. He is currently a Ph.D. candidate in Cognitive Psychology at the University of California, Los Angeles.

GRACE CHEN was a Visiting Researcher at the ICT. She is currently a graduate student in Psychology Research at California State University, Long Beach.

MILTON ROSENBERG is a Special Project Manager at the ICT.