A cognitive model for social role compliant behavior of virtual agents

Soldiers in foreign missions need to be able to interact with local people while taking their cultural and social norms into account. Serious games offer an opportunity to acquire the knowledge and skills to act appropriately, provided that the virtual characters behave according to the norms to be learned. This paper presents research on how to model the characteristics of social groups into the constituent members of that group. A (virtual) person can belong to different social groups simultaneously (e.g. family, religious community; war tribe, etc). Each group has their own characteristics, such as common goals or a set of norms, which (partly) determine the behavior of the individuals. We developed a method to generate behavior of virtual characters as a function of the social groups they belong to. This is achieved through calculating plan utilities by taking into account the social groups, personal preferences, and the situational context. The method is tested using a military house-search scenario, revealing that our virtual characters acted in accordance with their social groups, even in the face of conflict between groups, by expressing behavior relevant to one or more of their social roles.

Part of this chapter has been published as:
de Man, J., Heuvelink, A., & van den Bosch, K. (2012). A cognitive model for social role compliant behavior of virtual agents. In *Intelligent Virtual Agents* (pp. 303-310). Springer Berlin Heidelberg.

9.1 Introduction

Belonging to a social group is -among many others- an important factor affecting an individual's behavior. There are many different social groups, like a family, a sports club, a religious community, a company, et cetera. Social groups within a particular culture (e.g. western culture, middle-eastern culture) share many norms and rules. However, there are also norms and rules that specifically apply to a particular social group within that culture. For example, a western male adult often behaves differently in a family setting than when in a football stadium, and yet again different when in church. Thus, social groups have their own set of norms and rules that partly prescribe how its members behave. These norms can be explicitly defined, but are more often than not implicit in nature. Furthermore, the role that a person holds within a social group also affects how someone behaves, or is expected to behave. For example, in a football club the captain behaves differently than a substitute player. Again, each role has its own set of behavior rules. Sometimes, an individual finds himself in the context of two (or more) different social groups whose norms do not entirely match, or that may even be in conflict. For example, if the father of the family brings his young children along to the stadium, his mates are likely to experience other behavior of him than they normally do. The situation compels the father to balance the norms and expectations of both groups. The resulting behavior is determined by many factors, like the importance that the social group attaches to particular behavior, the value that the individual assigns personally to displaying the particular behavior et cetera. Although it is difficult to predict the resulting behavior precisely, it is always a function of weighing profits and losses.

In order to understand someone's behavior, it is necessary to know the current social group(s) that person belongs to, the role of the individual within the group, and the norms that apply to the specific group and role. In contrast, when one is not familiar with the applicable norms, then it can be hard to understand why someone behaves in the way he or she does. Erroneous interpretations of someone's intentions may easily arise. And in some circumstances, this may have serious consequences. An example is the military. Current military missions are often staged in faraway countries with non-western cultures and unfamiliar social groups. Yet, when confronted with individuals, or a group of people, it is imperative for the commander and his team to interpret their behavior accurately and timely. Not recognizing a threat may endanger the own group. On the other hand, acting with violence against people or groups that have no evil intentions harms the respect and goodwill of the local population.

So, assessing accurately the nature and intentions of individuals and groups is of central importance for soldiers on foreign grounds (McFate, 2005). Knowledge of how people's actions are determined by the social groups of the local culture they belong to, is essential for behavior interpretation. There is a growing awareness that education and training can play an important role in preparing the soldier for missions in unfamiliar settings (Michael, 2006; Muller et al., 2011). Serious games offer the potential to develop contextually rich and flexible environments for training people how to evaluate behavior that is affected by social groups, and also to practice the skills needed to communicate in accordance with the norms of these groups (Core et al., 2006). In order to realize the potential of serious games to train people in recognizing social role compliant behavior, there is a need for adequate models that generate the behavior of the virtual players in the games. These virtual players need to be able to act according to the norms of their social groups. In this paper we present work on the development and evaluation of a model that generates the behavior of socially compliant intelligent agents. In particular, the question is addressed how to create agents that belong to different social groups simultaneously, and yet produce believable behavior given the context. As argued above, this can be difficult as the norms and goals of the different groups may raise friction. Pokahr et al. (2005a) said about this problem: 'the cumbersome task of ensuring that the agent will never process any conflicting goals at the same time is left to the agent developer'. Our research aim is to develop agents that can autonomously decide between conflicting goals forthcoming from their assigned social groups and roles.

9.2 Background research

Tajfel (1972) introduced the concept of social identity as 'the individual's knowledge that he belongs to certain social groups together with some emotional and value significance to him of this group membership', which served as the beginning of the social identity theory. Subsequent work has extended this theory to link the self-concept formed with social categories to the cognitive processes that underlie group behavior. Self-categorization theory (Turner et al., 1987) explains group behavior by stating that people can categorize themselves at different levels of abstraction. Most important here is the level that defines social identity; the ingroup-outgroup level. Categorization at this level evokes the so-called process of *depersonalization*, which 'brings self-perception and behavior in line with the contextually relevant ingroup prototype' (Hogg & Terry, 2000).

It is thus possible for an individual to behave not according to his own personality, but according to some *prototype* of a particular (social) group. Hogg & Terry (2000) define a prototype as the 'cognitive representation of features that describe and prescribe attributes of the group', thereby referring to features that 'maximize similarities within and differences between groups,

[..] including beliefs, attitudes, feelings and behaviors'. Such prototypes of social groups serve as the basis for our model. People fulfill a particular role in a social group and 'rapidly internalize social norms about what their roles entail' (Sunstein, 1996). Hofstede & Hofstede (2004) point to a problem forthcoming from this view: 'As almost everyone belongs to a number of different groups and categories at the same time, we unavoidably carry several layers of mental programming within ourselves, corresponding to different levels of culture. [..] The mental programs from these various levels are not necessarily in harmony.' This problem has to be taken into account when multiple social groups are combined.

Culture has been incorporated in various models proposed in related research. One possible approach is to model culture by means of norms and obligations. The notion of norms has been taken from social psychology to model coordination and cooperation between multiple agents realistically. Early work in this field built norms directly into the agents, often as explicit constraints (Conte et al., 1999). Castelfranchi et al. (2000) proposed an architecture in which agents are able to communicate, adapt and violate norms. Especially violation of norms is relevant for our research, because of social groups with conflicting norms. Unfortunately, there is a lack of research in 'intelligent violation of norms' (Castelfranchi et al., 2000) and it is difficult to find concrete work done in this area.

Much research focusing on training cultural awareness take into account cultural models. As culture is a particular type of social group the models may be applicable to our research. One line of research models culture by extending FAtiMA, an architecture incorporating emotions and personality (Dias & Paiva, 2005). FAtiMA is built up in two layers: the reactive layer processes events while the deliberative layer plans the actions the agent performs. Both Aylett et al. (2009) and Mascarenhas & Paiva (2010) expanded this model to create agents for an educational program. They added *symbols* and *rituals* to capture a culture. Symbols express how certain actions are interpreted by that culture, while rituals contain a sequence of actions specific to that culture. Although these models did create a way to model cultures independent from the agent itself, the method does not generalize well to social groups.

Taylor & Sims (2009) follow a different approach. Their architecture, named the Cultural Cognitive Architecture, uses *schemas* and *appraisal theories of emotion*. Schemas are a hierarchical representation of interactions and appraisal methods are used to evaluate observed interactions on different emotions. Based on these emotions, subsequent actions are chosen to fit within the current schema. A downside of this architecture is that it currently only consists of a theoretical framework. Not many processes have been formalized and there is no preconception on how to implement this model.

Solomon et al. (2008) describe a model of Culturally Affected Behavior

(CAB). Here, culture is modeled as a network of actions affecting mental states. A mental state has a current utility (to what extent it is valid in the given situation) and an intrinsic utility that reflects the importance of the mental state to the agent. E.g. a mental state *is-observant-of-Islam* has a high intrinsic utility to an agent representing a pious Muslim, but a low intrinsic utility to an agent representing a western atheist. The action to offer alcohol to a Muslim agent would decrease this agent's view that the person offering him alcohol is observant of Islam. Similarly, any action could be defined to change the current utilities of the agent's states. Based on the current utilities, a Socio-Cultural Satisfaction (SCS) is calculated whereby intrinsic utilities are used as a weight for the respective state. Bulitko et al. (2008) use this SCS as a value to assess whether or not an agent can be trusted to keep his promises. If this trust is low, the corresponding effects of his promises are not considered in deciding future actions. The work of Bulitko and colleagues also show that a variety of factors, such as emotion or personality, can be incorporated in determining an agent's actions using this method. Below, we propose a model where this method is used to create social compliant behavior.

9.3 Social compliant behavior model

In this paper our model of social compliant behavior is illustrated in a military context where a soldier needs to learn how to enter a house and to address the occupants. In this scenario soldiers closely interact with people of different cultures and customs. The local residents are modeled using a set of prototypes, in this example related to a Muslim culture. The residents can be either a family, a group of militant soldiers, or a combination thereof.[1]

9.3.1 Prototypes for roles in social groups

Social group and *role* are the main concepts of our model. An agent may belong to one of more social groups simultaneously (e.g. Islam, family) by fulfilling a role within each group (e.g. man, father). These roles produce the agent's behavior through so-called (role) prototypes. Figure 9.1 shows an example.

Beliefs. Beliefs represent an individual's knowledge of the world (Rao & Georgeff, 1995). For example, a Muslim woman has a belief expressing if she is veiled or not; a militant soldier has a belief whether or not he is carrying a weapon (see Figure 9.1: a soldier has beliefs about `weapon` and `enemy present`). Which beliefs are required by each prototype are defined beforehand, however the actual value of beliefs can change at any time.

[1] For a more detailed description of the model's functioning we refer to Man (2011).

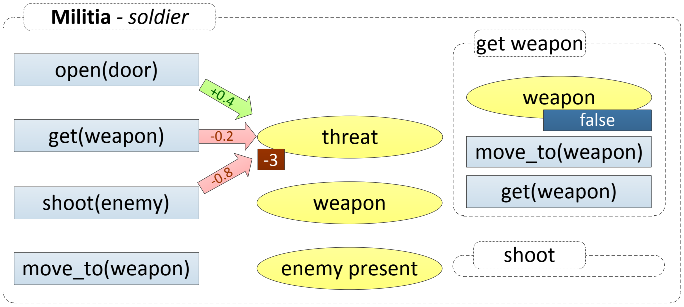

Figure 9.1: Visualization of the prototype of the role *soldier* of the social group *Militia*. The 'get weapon' plan is unfolded; the 'shoot' plan is folded.

States. States are the driving force behind plan selection and are therefore very important in the model. The term *state* is first used in the *Culturally-Affected Behavior project* where it is described to 'decode states of the world [..] and have intrinsic utility values that represent the relative importance that the human behavior model has for the state weighed against other states' (Solomon et al., 2008). In Bulitko et al. (2008) such a state is described as 'an attribute of the problem-solving state that has some intrinsic utility / concern-value to the agent'. Thus, states carry information about the world and their presence (or absence) have a value to the agent.[2] The difference with beliefs - that directly link to information about the world - is that states constitute more abstract mental concepts, e.g. compare the belief `enemy present` with the state `threat` in Figure 9.1. In our model, these states are not defined for individual agents, but for roles of social groups. For example, decency is very important for all Muslim women living in Islamic countries. Not wearing a veil in the presence of non-family members evokes a feeling of indecency that should be avoided. To model this, the prototype of the role woman for the social group Islam has a state `decency` with a high intrinsic utility.

Plans. Every prototype contains its own set of plans. A plan consists of a sequence of actions and (optional) preconditions. The `get weapon` plan of a soldier shown in Figure 9.1 consists of the actions `move_to(weapon)` followed by `get(weapon)`. Plans can have a precondition added, e.g. `weapon(false)` as getting a weapon is not needed when the agent already has a weapon.

[2] Note that a state does not refer to a particular configuration of the agent, unlike other approaches where a state describes a particular configuration of information.

Action effects. Plans are composed of multiple actions for which certain effects are expected. One type of expected effect is a change in one or multiple states and is modeled as an action effect. For example, the plan to open the door contains the action `open(door)`. By performing this action, we expect the current utility of the agent's state curiosity to decrease as it will find out what is on the other side. Similarly, some actions are expected to increase the current utility of a state; a Muslim woman veiling herself increases her decency. The effects that actions may have on the current utility of states can vary in degree. For example, for a militia man grabbing a weapon may decrease his feeling of threat somewhat, but shooting the enemy is likely to decrease his feeling of threat even more (see the arrows denoting the action effects in Figure 9.1).

Observation functions. Action effects model the *expected* effects of actions, but the *actual* effects may be different. For instance, whether the threat for a militia man indeed decreases depends on whether he has eliminated the enemy altogether. An agent establishes the actual effects by means of its observation functions. Observation functions are more detailed and context dependent than action effects. Observation functions are specific for a particular role. Spotting an enemy soldier is very threatening to persons belonging to a militia group, but not (so much) for civilians. Observation functions thus affect states. Moreover, they also affect beliefs. For example, if the observation is made that the door is being opened, the belief that the door is closed needs to be adjusted.

9.3.2 Agents as a combination of prototypes

The previous section described the components of a prototype denoting a role within a social group. In this section, these prototypes are combined into an agent. An agent is a collection of any number of roles. In addition, an agent contains modifiers to regulate the importance of the different roles as explained below. Furthermore, agents are preconfigured by an initial set of beliefs and states, whose content or value can change over time. During the run of a scenario, the observation functions of the different prototypes adjust these values in real-time. See Figure 9.2 for an overview of the agent processes.

Prototype importances. An agent can be related to various social groups via different roles. However, just as any person does not feel equally connected to all of his social groups, an agent can also differentiate between each of his roles. This is modeled by defining a set of static modifiers M_S that describe the relative importance of each prototype. A second type of modifiers are the dynamic modifiers M_D. The motivation for adding these modifiers comes from the process of *depersonalization* as explained in Background Research

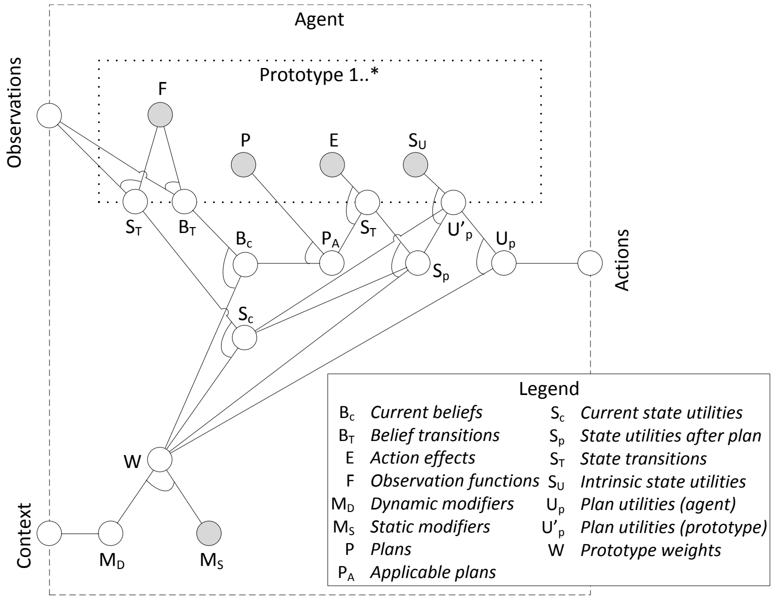

Figure 9.2: An overview of the processes within an agent

(Section 9.2). Obviously, context affects the importance of the different roles. This is modeled by calculating dynamic modifiers for each social group based on the number of people present belonging to that particular group.

The static and dynamic modifiers are used to regulate the importance of the different prototypes, which in turn has an effect on the resulting agent behavior. This is done by calculating a weighted average of these modifiers for each group. This weight represents the importance of a particular prototype, relative to the other roles and is based on both the context and the relative importance the agent assigns to each role.

Observation processing. An agent needs to be reactive to his environment and therefore needs to be able to process input. In the present model, input comes from observations. A prototype can make two types of inferences from observations; inferences on states denoted by the state transitions S_T and inferences on beliefs, denoted by belief transitions B_T. These transitions express how a belief or state should change for a particular prototype. For example, the belief `door_open` becomes true; or `curiosity` decreases by 0.8. For each prototype such a set of transitions follow from the observation functions. This could result in different, or even conflicting transitions when two or more prototypes are relevant. The actual transitions applied to the current beliefs and states are established using the prototype weights. If conflicting beliefs are derived, the belief transition of the group with the highest weight is processed. For state transitions, a weighted average of the effects is added to the current state value.

Plan deliberation. Deliberation starts with the formation of a list of plans that can be executed at that particular point in time (P_A). Each plan may have one or more associated preconditions. Every time the agent starts its plan deliberation, each precondition is checked against the current beliefs and only those plans for which every precondition is met are considered in the deliberation process.

Plans contain actions for which action effects on states may be defined in prototypes. Remember that action effects refer to expected outcomes and do not necessarily constitute actual effects. Plans are prioritized on the basis of which plan is expected to bring about the best results. To determine this, all action effects are simulated for each of the applicable plans. For each state which is affected by an action within a particular plan, a simulated value is calculated by using the prototype weights to add a weighted average to the current value.

To evaluate the effects of a plan, a comparison is made between the initial state values and the values that are expected to be achieved when executing the plan. To compare this with other prototypes, relative change rather than absolute change needs to be used (due to prototypes having different scales). Plan utilities can be calculated for each prototype with respect to the intrinsic utility of the states within that social group. Afterwards, an overall plan utility can be calculated using the prototype weights. This process is done for each of the applicable plans. The plan with the highest utility is selected for execution.

9.4 Model implementation

To implement our behavior model Jadex is used; 'a software framework for the creation of goal-oriented agents following the belief-desire-intention (BDI) model' (Pokahr et al., 2005b). Beliefs, states and action effects of prototypes are defined using XML; Java is used for observation processing, plan deliberation and the execution of plans. For each agent, a number of prototypes are combined and an initial configuration is manually set. A separate agent is used to interact with a graphical interface of choice.

To instantiate the behavior model, social groups and roles have to be defined. To implement the military house-search scenario four different social groups and a multitude of role prototypes are used. In addition, a number of states are defined that are relevant to the scenario and reflect underlying motivations to act. For example, a person is motivated to open a door because of the state `curiosity` of which the value is lowered by that action. In this section we elaborate on a few roles for illustrative purposes. Table 9.1 shows an overview of various prototypes (columns) and states (rows). Within this table, for each role a `+` or `-` expresses whether or not the state is desired, and the amount of signs its level of importance. For example, curiosity is an undesired state of medium importance for both an adult and a child, while decency is a desired state of high importance for a Muslim woman.

For each prototype, plans are defined that an agent can perform. For each of these plans, necessary conditions are defined as well as the actions that compose the plan. Table 9.1 gives an overview of some of the plans implemented for the house-searching scenario.

To enable agents to select appropriate plans, expected action effects are defined for each role (see Table 9.2). For example, for an adult opening a door is expected to decrease curiosity, denoted by `---`. Together with the fact that the state curiosity has a negative intrinsic utility for an adult (see Table 9.1), this action effect implements that there is a large chance that an adult will open the door after having observed a door knock (increasing its curiosity). However, for a militant soldier (belonging to social groups Person and Militia) curiosity is not the only factor. Opening the door is also expected to increase the level of threat, a state that also carries a negative intrinsic utility. Here, a conflict in goals is found between two prototypes. As an adult person you would like to open the door, while a militant soldier also considers the threat this might bring along.

For each action effect, a corresponding observation inference is defined (table omitted due to space limitations). In addition there are observation inferences that specify how observations lead to beliefs. For example, seeing a soldier will create a belief that an enemy is present. Furthermore, actions performed by the trainee result in events observed by the agent. These observations can bring about various inferences, depending on the role and social group of the agent. To illustrate this point if a trainee knocks on the door this brings about curiosity, but it also evokes a feeling of threat in militants.

	Person *adult*	**Person** *child*	**Family** *man*	**Family** *woman*	**Islam** *woman*	**Militia** *soldier*
threat	---	---				---
threat_family			--	---		
curiosity	--	--				
decency					+++	

Table 9.1: Intrinsic utilities of several states for various role prototypes

Plan	Social group	Condition	Actions
open_door	Person - adult	door_open: `false`	move_to(door) open(door)
investigate	Person - child	unknown: `not empty`	move_to(`unknown`) investigate(`unknown`)
protect_wife	Family - man	wife_present: `true`	move_to(`wife`) protect(`wife`)
protect_child	Family - woman	child_present: `true`	move_to(`child`) protect(`child`)
veil	Islam - woman	veiled: `false`	veil(true)
get_weapon	Militia - soldier	weapon: `false`	move_to(weapon) get(weapon)
shoot	Militia - soldier	enemy_present: `true`	shoot(`enemy`)

Table 9.2: Plans and their conditions and actions

	Person *adult*	**Person** *child*	**Family** *man*
open(door)	curiosity: `---`		
investigate(unknown)		curiosity: `--`	
move_to(unknown)		threat: `+`	
protect(*)			threat_family: `--`
shoot(*)			threat_family: `+++`
	Family *woman*	**Islam** *woman*	**Militia** *soldier*
open(door)		decency: `--`	threat: `++`
veil(true)		decency: `++`	
protect(*)	threat_family: `--`		
get(weapon)			threat: `-`
shoot(enemy)			threat: `---`

Table 9.3: Actions and their action effects

9.5 Exploring the model's validity

A small scale evaluation study was conducted to evaluate the proposed model. Three different scenarios were developed to investigate whether prototypes can be used to create social role compliant agent behavior. The context is house-searching by western soldiers in a culturally unfamiliar setting (Islamic), with different compositions of people in the house. The player is a western soldier, the agents represent members of different groups having varying roles. In the first scenario, the group in the house is composed of family members. In the second scenario, a militant soldier and his family occupy the house. In the third scenario, only militant soldiers are present in the house. Each agent should act according to its role, taking the context (e.g. other agents) into account. The following results were obtained.[3]

First, the model was able to produce behavior consistent with the agent's social groups and roles, even when agents belonged to multiple groups and served multiple roles. This can be illustrated, for example, by the behavior of the woman Muslim agent. When she hears the knock on the door, she wants to open the door to satisfy her curiosity. However, being an Islamic woman, opening the door would be indecent for her to do. This creates a conflict. The model acknowledges the conflict and uses the relative importance of her roles and states, leading her to decide not to open the door. However, the Muslim man is not restrained by this cultural norm and opens the door.

Second, agents' behavior is affected by multiple roles. For example, the militant soldier of the second scenario responds to the door knock by retrieving a weapon and opening the door. Retrieving a weapon comes forth from his role in the militia group, while opening the door is the default human response. This demonstrates that the model generates behavior affected by multiple roles.

Third, results show that agents' behavior is affected by context. When the militia soldier is accompanied by other militia members (third scenario), he responds differently to a door knock: he retrieves a weapon and takes cover. So, by designing a more hostile context, the agent adapts his behavior accordingly.

9.6 Conclusion

Soldiers in foreign missions need to be able to interact with local people while taking their (cultural) norms into account. These communication skills can be trained in simulations or games, provided that the virtual characters behave according to the norms of their social groups. This paper presents research on how to model the characteristics of social groups into the constituent members of that group. We developed a method to generate behavior of virtual characters as a function of the roles they fulfill in various groups. This is achieved by defining role prototypes reflecting what is important for a particular social group, what plans are available to change the current situation, and what effects are expected from particular actions. An agent, being a combination of role prototypes, uses this information for calculating plan utilities, taking into account the social groups, personal preferences, and the situational context.

[3] For a more detailed description of the results we refer to Man (2011).

The model was implemented in JADEX. This allowed us to test whether the agents would produce behavior consistent with their social group(s) and role(s). For testing, we used a house-search scenario, revealing that our virtual characters acted in accordance with their social groups, even in the face of conflict between groups, and were able to express behavior relevant to one or more of their social roles. It is not claimed that the model presented here has sociological or psychological validity. Although concepts and processes have been based on research in those areas, the current implementation has been constructed only at face value. In order to develop valid models, it is needed to utilize more data from relevant studies and theories, validate the created models in human subjects studies and experiment with larger implementations in a variety of domains.

The proposed model opens up various paths for future research: personality could be implemented as a separate prototype; emotion could be incorporated by relating the current utilities of states to different emotions which in turn may affect the intrinsic utilities of particular states; or agents could be made aware of differences between expected and actual effects to promote learning. The model is obviously not complete, but the architecture allows developing models that produce social compliant behavior in virtual agents. Even with just a few social groups, a large variety of agents can be created by making different combinations. This, in itself, is an important step in developing a wide range of scenarios for interactive training in socio-culturally appropriate behavior.

References

Aylett, R., Vannini, N., Andre, E., Paiva, A., Enz, S., & Hall, L. (2009). But that was in another country: agents and intercultural empathy. In *Proc. of AAMAS 2009*, (pp. 329–336). Richland, SC.

Bulitko, V., Solomon, S., Gratch, J., & Lent, M. v. (2008). Modeling culturally and emotianally affected behavior. In *Proc. of AIIDE 2008*, (pp. 10–15). Stanford, California, US.

Castelfranchi, C., Dignum, F., Jonker, C., & Treur, J. (2000). Deliberative normative agents: Principles and architecture. In N. Jennings, & Y. Lesprance (Eds.) *Intelligent Agents VI. Agent Theories, Architectures, and Languages*, LNCS 1757, (pp. 364–378). Springer Berlin/Heidelberg.

Conte, R., Falcone, R., & Sartor, G. (1999). Introduction: Agents and norms: How to fill the gap? *Artificial Intelligence and Law*, *7*, 1–15.

Core, M., Traum, D., Lane, H., Swartout, W., Gratch, J., Van Lent, M., & Marsella, S. (2006). Teaching negotiation skills through practice and reflection with virtual humans. *Simulation*, *82*(11), 685–701.

Dias, J., & Paiva, A. (2005). Feeling and reasoning: A computational model for emotional characters. In C. Bento, A. Cardoso, & G. Dias (Eds.) *Progress in Artificial Intelligence*, LNCS 3808, (pp. 127–140). Springer Berlin/Heidelberg.

Hofstede, G., & Hofstede, G. J. (2004). *Cultures and Organizations: Software of the Mind*. New York: McGraw-Hill, 2 ed.

Hogg, M., & Terry, D. (2000). Social identity and self-categorization processes in organizational contexts. *Academy of Management Review*, *25*(1), 121–140.

Man, J. d. (2011). *Composing agents from role prototypes of social groups*. Master's thesis, VU University, Amsterdam, Netherlands. Available from: www.few.vu.nl/~jmn300.

Mascarenhas, S., & Paiva, A. (2010). Creating virtual synthetic cultures for intercultural training. In *CATS 2010*.

McFate, M. (2005). The military utility of understanding adversary culture. *Joint Forces Quarterly*, *38*, 32–48.

Michael, D. (2006). Serious games: Games that educate, train, and inform. *Boston, MA.: Thomson Course Technology*.

Muller, T., Bosch, K. v. d., Kerbusch, P., & Freulings, J. (2011). LVC training in urban operation skills. In *Proc. of EURO SISO/SCS 2011*. The Hague, Netherlands.

Pokahr, A., Braubach, L., & Lamersdorf, W. (2005a). A goal deliberation strategy for BDI agent systems. In *Proc. of MATES 2005*, (pp. 82–94). Springer-Verlag.

Pokahr, A., Braubach, L., & Lamersdorf, W. (2005b). Jadex: A BDI reasoning engine. In *Multi-Agent Programming: Languages, Platforms and Applications.*, vol. 15, (pp. 149–174). Springer US.

Rao, A., & Georgeff, P. (1995). BDI agents: From theory to practice. In *Proc. of ICMAS 1995*, (pp. 312–319).

Solomon, S., Lent, M. v., Core, M., Carpenter, M., & Rosenberg, M. (2008). A language for modeling cultural norms, biases and stereotypes for human behavior models. In *Proc. of BRIMS 2008*.

Sunstein, C. R. (1996). Social norms and social roles. *Columbia Law Review*, *96*(4), 903–968.

Tajfel, H. (1972). Social categorization. In S. Moscovici (Ed.) *Introduction a la Psychologie Sociale*, vol. 1, (pp. 272–302). Paris: Larousse.

Taylor, G., & Sims, E. (2009). Developing believable interactive cultural characters for cross-cultural training. In A. Ozok, & P. Zaphiris (Eds.) *Online Communities and Social Computing*, LNCS 5621, (pp. 282–291). Springer Berlin/Heidelberg.

Turner, J., Hogg, M., Oakes, P., Reicher, S., & Wetherell, M. (1987). *Rediscovering the social group: A self-categorization theory*. Oxford: Blackwell.