A Cognitive Model for Social Role Compliant Behavior of Virtual Agents

Jeroen de Man[1], Annerieke Heuvelink[2], and Karel van den Bosch[2,⋆]

[1] VU University Amsterdam
j.de.man@vu.nl
[2] TNO
{annerieke.heuvelink,karel.vandenbosch}@tno.nl

Abstract. 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.

1 Introduction

The norms and values shared within a social group shape for a significant amount the behavior of its members. For example, a western male adult behaves differently in a family setting than when in a football stadium, and yet again different when in church. 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. Sometimes, an individual may find itself in the context of two (or more) different social groups whose norms conflict. For example, if the father brings his 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 weigh and balance the norms and expectations of both groups. How the father will behave 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 someones behavior, it is necessary to know the current social group(s) that person belongs to, the role of the individual within the group,

⋆ With thanks to co-supervisors Tibor Bosse and Jurriaan van Diggelen.

Y. Nakano et al. (Eds.): IVA 2012, LNAI 7502, pp. 303–310, 2012.

and the norms that apply to the specific group and role. When one is unfamiliar with the applicable norms, it can be hard to understand why someone behaves in the way he or she does. Erroneous interpretations of someones 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 commanders and their teams to interpret their behavior accurately and timely (McFate, 2005). There is a growing awareness that training plays an important role in preparing the soldier for missions in unfamiliar settings (Muller et al., 2011). In this paper we present work on the development of a model that generates the behavior of socially compliant intelligent agents that can be used for such training.

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. 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 individuals to behave not according to their own personality, but according to some *prototype* of a particular (social) group. As people fulfill a particular role in a social group they ‘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, [..] The mental programs from these various levels are not necessarily in harmony.’ The model we propose is designed to handle this issue arising from combining multiple social groups.

Culture, a particular type of social group, has already been incorporated in various cognitive models. One line of research models culture by means of norms and obligations (Conte et al., 1999). Castelfranchi et al. (2000) proposed an architecture in which agents are able to communicate, adapt and violate norms. Unfortunately, there is a lack of research in ‘intelligent violation of norms’ (Castelfranchi et al., 2000). The Cultural Cognitive Architecture proposed by Taylor & Sims (2009) uses *schemas* and *appraisal theories of emotion* in their model. A downside of this architecture is that it currently only consists of a theoretical framework. FAtiMA is an architecture incorporating emotions and personality (Dias & Paiva, 2005) and has been expanded to incorporate culture using *symbols* and *rituals* (Aylett et al., 2009; Mascarenhas & Paiva, 2010), however the method does not generalize well to social groups.

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. Based on the current utilities of an agent, a Socio-Cultural Satisfaction (SCS) is calculated whereby intrinsic utilities are used as a weight for the respective state. The work of Bulitko et al. (2008) shows 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.

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 occupants are modeled using a set of prototypes, in this example related to a Muslim culture. The occupants can be either a family, a group of militant soldiers, or a combination thereof.[1]

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 social group (e.g. muslima, head of family). These roles produce the agent's behavior through so-called (role) prototypes. Fig. 1 shows an example, the various concepts are explained in the following paragraphs.

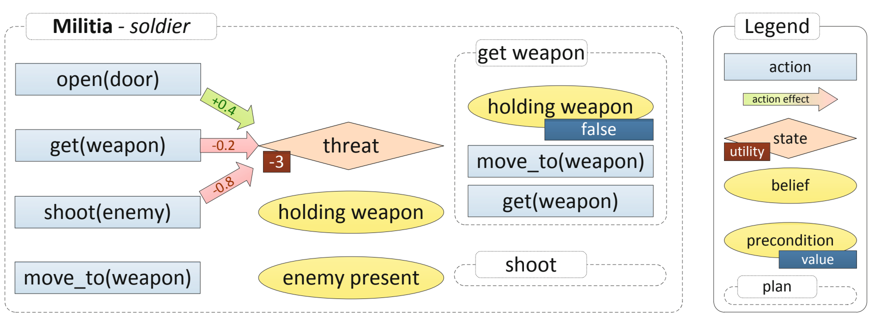

Fig. 1. Visualisation of the prototype of the role *soldier* of the social group **Militia**. The 'get weapon' plan is unfolded; the 'shoot' plan is folded.

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

Beliefs. Beliefs represent an individual's knowledge of the world (Rao & Georgeff, 1995). For example, a Muslim woman has a belief expressing whether she is veiled or not; a militant soldier has a belief whether or not he is carrying a weapon (see Fig. 1: a soldier has beliefs about `holding weapon` and `enemy present`). Which beliefs are required by each prototype are defined beforehand, however the actual value or content of beliefs can change at any time.

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). Like proposed by Bulitko et al. (2008) a state '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 Fig. 1. In our model, 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. Thus, 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 Fig. 1 consists of the actions `move_to(weapon)` followed by `get(weapon)`. Plans can have a precondition added, e.g. `holding weapon(false)` as getting a weapon is not needed when the agent already has a weapon.

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 (see the arrows denoting the action effects in Fig. 1). 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 value 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.

Observation Functions. Action effects model the *expected* effects of actions, but the *actual* effects may be different. For instance, whether the threat for a

[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.

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.

3.2 Agents as a Combination of Prototypes

The previous section described the components of a prototype denoting a role within a social group. 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. The following paragraphs explain these concepts in more detail.

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 or her social groups, an agent can also differentiate between each of its roles. This is modeled by defining a set of *static* modifiers that describe the relative importance of each prototype for the agent. A second type of modifiers are the *dynamic* modifiers. The motivation for adding these modifiers comes from the process of *depersonalization* (see Sec. 2 Background Research) that denotes that context affects the importance of roles. To model this, dynamic modifiers are calculated for each social group based on the number of people present belonging to that particular group. The static and dynamic modifiers regulate the current importance of the different prototypes for an agent, which affects the resulting agent behavior. The importance of a particular prototype relative to the other roles is determined by calculating a weighted average of the static and dynamic modifiers for each group.

Observation Processing. An agent processes input (e.g. observations) to respond to its environment. A prototype makes inferences on states denoted by the state transitions, and inferences on beliefs denoted by belief transitions. 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. 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. Only 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 according to expected results. The agent first simulates all action effects for each applicable plan, thereby calculating simulated values for the states taking the prototype weights into account. To evaluate the effects of a plan, a comparison is made between the initial state values and the expected values when executing the plan. For each prototype an agent belongs to, the predicted outcomes of plans are scored. Then, an over-all comparison is conducted, taking the weights of the prototypes into account. The plan having the best outcome is selected.

4 Proof of Concept

The model has been implemented using Jadex (Pokahr et al., 2005) and a small scale evaluation study was conducted as a proof of concept. 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. Within these scenarios, each agent should act according to its roles, taking the context (e.g. other agents) into account.

To instantiate the behavior model for this context, we considered the social groups family, militia and Islam with roles such as child, soldier or muslima. Important states in this context are for example threat with a high intrinsic utility for a soldier and decency with a high intrinsic utility for an Islamic woman. Afterwards, for each role relevant plans such as opening the door or getting a weapon were identified. For each of the actions in these plans, the expected effects on the various states were modeled as being low, medium or high. The following observations were made.[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

[3] For a more detailed description of the implementation and the results see Man (2011).

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 agent together with its family (second scenario) responds to the door knock by first retrieving a weapon and afterwards 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), it responds by retrieving a weapon and taking cover instead of opening the door. So, in a hostile context, the agent adapts its behavior accordingly.

5 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 virtual environments, 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.

The model was implemented in Jadex and tested using a house-search scenario. We found 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 research 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 set of intrinsic utilities and action effects; emotion could be incorporated by relating the current values of states to different emotions which in turn may affect the intrinsic utilities of particular states; or agents could learn from differences between expected and actual effects. 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.: But that was in another country: agents and intercultural empathy. In: Proc. of AAMAS 2009, Richland, SC, pp. 329–336 (2009)

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

Castelfranchi, C., Dignum, F., Jonker, C., Treur, J.: Deliberative Normative Agents: Principles and Architecture. In: Jennings, N.R. (ed.) ATAL 1999. LNCS, vol. 1757, pp. 364–378. Springer, Heidelberg (2000)

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

Dias, J., Paiva, A.C.R.: Feeling and Reasoning: A Computational Model for Emotional Characters. In: Bento, C., Cardoso, A., Dias, G. (eds.) EPIA 2005. LNCS (LNAI), vol. 3808, pp. 127–140. Springer, Heidelberg (2005)

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

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

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

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

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

Muller, T., van den Bosch, K., Kerbusch, P., Freulings, J.: LVC training in urban operation skills. In: Proc. of EURO SISO/SCS 2011, The Hague (2011)

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

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

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

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

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

Taylor, G., Sims, E.: Developing Believable Interactive Cultural Characters for Cross-Cultural Training. In: Ozok, A.A., Zaphiris, P. (eds.) OCSC 2009. LNCS, vol. 5621, pp. 282–291. Springer, Heidelberg (2009)

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