

Rapid prototyping of social group dynamics in multiagent systems

Matthias Rehm · Birgit Endrass

Received: 15 January 2008 / Accepted: 19 December 2008 / Published online: 20 February 2009
© Springer-Verlag London Limited 2009

Abstract In this article we present an engineering approach for the integration of social group dynamics in the behavior modeling of multiagent systems. To this end, a toolbox was created that brings together several theories from the social sciences, each focusing on different aspects of group dynamics. Due to its modular approach, the toolbox can either be used as a central control component of an application or it can be employed temporarily to rapidly test the feasibility of the incorporated theories for a given application domain. This is exemplified by applying the toolbox to different applications.

1 Introduction

Virtual agents have been employed in many games and entertainment applications with the aim to engage users and enhance their experience. However, to achieve this goal it does not suffice to provide only for sophisticated animation and rendering techniques. Rather, other qualities have to come into focus as well, including the provision of conversational skills as well as the simulation of social competence that manifests itself in a number of different abilities. Important progress has been made in the area of embodied conversational agents focusing on dyadic interactions between a single user and a single agent (see, e.g., Cassell et al. 2000; Prendinger and Ishizuka 2004). Scaling

up to multiple users and/or multiple agents poses some new challenges.

In multiagent systems, interactions between agents are often based on rules and plans for the single agent assuming rational behavior. But when people interact, dynamic group processes take place depending on social rules but also on such irrational aspects like personality or emotion. For multiple agents, their individual behavior has to be accompanied by coherent group behavior, which will not simply emerge by itself if some agents are put together because a group is more than just a bunch of single characters that happen to be at the same location. Instead, a group is constituted of relations between the different group members that influence how they will behave and communicate among each other. Thus, endowing agents with social group dynamics will allow them to build relationships among each other ideally following theories from social psychology. This is important because of two reasons. On the one hand, Reeves and Nass (1996) have shown that people tend to socialize with technical artifacts, and virtual agents are an ideal vehicle for projecting assumptions about human–human interaction to the interaction with technical systems. Interacting with multiple agents results in the need for consistent and believable group behaviors for the agents. On the other hand, commercial games such as “The Sims” exemplify that the simulation of social skills can render interactions between virtual characters more believable and engaging.

Different approaches have been presented to handle aspects of social (group) behaviors. For instance, Prendinger and Ishizuka (2001) investigate the relationship between an agent’s social role and the associated constraints on emotion expression. They allow a human script writer to specify the social distance and social power relationships among the characters involved in an application, such as

M. Rehm (✉) · B. Endrass
Multimedia Concepts and Applications, Augsburg University,
Eichleitnerstr. 30, 86159 Augsburg, Germany
e-mail: rehm@informatik.uni-augsburg.de

B. Endrass
e-mail: endrass@informatik.uni-augsburg.de

a multiplayer game scenario. Another approach has been taken by Rist and Schmitt (2002) who aim at emulating dynamic group phenomena in human–human negotiation dialogues based on socio-psychological theories of cognitive consistency dynamics (Osgood and Tannenbaum 1955). To this end, they consider a character's attitudes towards other characters and model a character's social embedding in terms of liking relationships between the character and all other interaction partners. Prada and Paiva (2005) as well as Pynadath and Marsella (2005) developed a social simulation tool as a backend to interactive pedagogical drama applications. While the development of social relationships in the approach by Prada and Paiva (2005) is mainly determined by the type of social interactions between them, Pynadath and Marsella regard the beliefs of agents about other agents as a key factor of social interaction and rely on a theory of mind to explicitly represent the beliefs of agents about other agents.

In the systems described above, social behaviors are mainly reflected by the agents' communicative behaviors. In contrast, Guye-Vuillème and Thalmann (2001) concentrate on the simulation of social navigation behaviors in virtual 3D environments including the social avoidance of collisions, intelligent approach behaviors, and the calculation of suitable interaction distances and angles. Their work is based on an operationalization of empirically grounded theories of human group dynamics, such as Kendon's group formation system (Kendon 1991).

Whereas all of the above mentioned work concentrates on modeling one specific theory, the objective of our work is to supply a tool for prototyping several theories in isolation or combination for a given multiagent system to increase the transparency of the system and to investigate how the inclusion of a certain theory influences the behavior of the agents. Social psychological theories concentrate on different aspects of group processes. Sometimes knowledge about a given situation might influence one's behavior, for instance if politely asked for directions one will give an answer even if only to acknowledge that one does not know the directions. Sometimes the mere presence of others might have an effect on one's behavior, for instance resulting in stage fright or social loafing. To make these different processes available for multiagent applications we have developed a toolbox that integrates different types of group processes and allows for either using it as the central control component or for rapidly prototyping such processes to assess their feasibility in a given application domain. In this article, we will first describe the different group processes and the necessary representations of group relationships among agents. Then we will give an overview on how the processes were implemented.

Integrating the user in the social group dynamics of multiple agents is another interesting challenge. The above mentioned approaches have different answers to this problem. In Prendinger and Ishizuka, the interaction is scripted confining the user to a given path. Rist and Schmitt let the user interact with a single agent who then acts as the user's representative in a multiagent negotiation process. In the Perfect Circle game of Prada and Paiva, the user has the possibility of suggesting courses of action to facilitate the problem solving task focusing primarily on task relevant interactions. At the end of this article, we present three exemplary applications that make use of the toolbox as a control device allowing for a unified approach to integrate the user in the social group dynamics of the agents.

2 Categorizing theories of social group dynamics

Following Goethals (1987), we distinguish between three groups of theories that explain social group dynamics from different perspectives.

1. *Theories focusing on social knowledge* “stress an active approach to understanding the social world” (Goethals 1987). Instead of just (passively) reacting to the presence or actions of other people, this understanding of the relationships and interaction rituals allows for choosing appropriate actions in social encounters with others.
2. *Theories focusing on self validation* “emphasize an active approach to self validation, they also envision actively initiating as well as responding to social influence” (Goethals 1987). Self validation becomes possible by comparing one's individual action tendencies and values with other group members (or other groups) allowing for personal development and adaptation to group norms.
3. *Theories focusing on social influence* “emphasize passive responding to social influence rather than active initiating of social influence” (Goethals 1987). In contrast to the above mentioned groups, theories of social influence describe what influence the mere presence or the actions of others have on an individual. It is necessary to distinguish between the influence group members and members of other groups exert on the individual.

At least one theory from each group is implemented in the current version of the toolbox. Because we deal with social interactions between agents in a multiagent system, we have to explain how the agents and their relationships among each other are defined before we can describe the implemented theories and their impact on the agents and their relationships in detail.

2.1 Representing social relations

The profile of a single agent is characterized by its name, gender, age group, social status, and personality. An agent's personality is represented by a vector of values along a number of psychological traits, such as extraversion or agreeableness and is based on the five-factor model of personality (McCrae and John 1992). Furthermore, the model takes into account an explicit representation of the relationships between agents. Interpersonal relationships are described by the degree of liking, familiarity, trust and commitment following Guye-Vuillème (2004):

- Liking: this dimension denotes group members' emotional attraction to one another. Rubin's (1973) conceptualization of liking underlines affection and respect as its two major components, and most definitions include the idea of social "closeness".
- Familiarity: as group members interact, they get to know each other, which increases their ability to predict the others' behavior and allows for a better communication and synchronization of their actions. The concept does not overlap with the Liking dimension because it is very common for a relationship to develop on the familiarity dimension without a corresponding growth in emotional attraction (e.g., co-workers).
- Trust: the definitions of this dimension in psychology are often very broad, and are based on such concepts as dependability, i.e., the feeling that you can rely on your partners when it matters, or faith, i.e., the belief in a relationship continuing indefinitely.
- Commitment: the last dimension can be described as a social force acting for continuing the relationship in the future.

The values on these dimensions are either specified in advance or derived from known properties of the agent's profile. For instance, agents with a similar social status are considered to trust each other more than agents with a different social status.

2.2 Social knowledge

Interaction process analysis (IPA) (Bales 1951) is based on a classification of social interactions that take place in small groups and essentially distinguishes between social-emotional factors that refer to the social relationships within a group, such as positive or negative feedback to group members, and task-oriented factors that refer to group tasks, such as asking questions or summarizing and offering directions. Twelve different IPA types can be distinguished:

1. Social-emotional interactions
 - Positive: Agree (1), Show Solidarity (2), Show Tension Release (3)
 - Negative: Disagree (4), Show Antagonism (5), Show Tension (6).
2. Task-oriented interactions
 - Questions: Ask for Opinion (7), Ask for Suggestion (8), Ask for Orientation (9)
 - Answers: Give Opinion (10), Give Suggestion (11), Give Orientation (12).

According to Bales (1951), this classification is fully inclusive. Consequently, every human interaction can be placed into one of the twelve IPA categories. To generate a stream of interactions, two aspects have to be considered. The personal attributes of the agent like its social status or its personality influence if and what kind of interaction the agent is likely to initiate. On the other hand, there are certain "rules of interaction" that are generally followed like an increased probability that the reaction to a question is an answer. Thus, IPA defines an active approach to interactions in the social world, taking into account certain regularities that can be observed and that most people adhere to in interactions.

Another aspect of IPA theory, the development of groups, can be added to the behavior component. It describes the influence that the social configuration of a group has on future interactions. For instance, the more positive interactions two agents perform, the more the trust dimension of their personal relation will increase. In addition, their motivation for interacting with each other as well as the probability for another positive social-emotional interaction increases.

2.3 Self validation

Congruity theory by Osgood and Tannenbaum (1955) enhances Heider's balance theory (Heider 1946) and represents an attempt to explain the development of interpersonal relationships in a cognitive model. This theory is based on the hypothesis that interpersonal relationships are influenced by simple cognitive configurations that are either balanced or unbalanced. Based on the assumption that people tend to avoid unbalanced configurations or cognitive dissonances, the theory allows predicting how people change their attitudes in a certain social situation. In this way, the theory allows us to describe changes in social relationships as a side effect of interactions. This includes even interactions in which individuals are not involved themselves but which they merely observe. Such an observation may influence their cognitive configurations.

Suppose Resi likes Sepp as well as Heidi (symbolized by a “+”-sign in Fig. 1 left). So far, Resi assumes that Sepp also likes Heidi (symbolized by the dashed arrow in Fig. 1 left). However, when talking to Sepp, it turns out that Sepp has no interest at all in Heidi (symbolized by the solid line and “-”-sign in Fig. 1 middle). As a consequence, Resi’s cognitive configuration of the situation becomes unbalanced. Since the conversation caused an unwanted dissonance in Resi’s state, she will either change her attitude towards Heidi or her attitude towards Sepp (Fig. 1 right). Of course, such a change will take place gradually over a number of interactions. Figure 1 (right) just exemplifies the extreme cases.

2.4 Social influence

Two different theories of social influence have been taken into accounts that explain different ways in which the presence of others influences one’s behavior. Social impact theory (Jackson 1987) defines this influence in a close analogy to physical phenomena. Such as the amount of light visible on a table depends on the number of light sources, their distances to the table and their strength, the same holds true for the influence which others exert on a target person. The social impact on a target person is calculated taking into account the strength, immediacy, and number of source persons where strength comprises features such as status or power, and immediacy represents the physical distance between source and target. As any of the above mentioned factors increases, the impact on the target also increases. This impact can be lessened if the number of target persons increases. For instance, if a subject has to perform a song, stage fright increases with the number of people in the audience and their status (Jackson and Latané 1981). It decreases if the subject has not to perform alone.

A number of studies are concerned with changes in productivity due to social impact. Latané et al. (1979) exemplified the phenomenon of social loafing with a simple experiment. People were asked by the experimenter to clap their hands. Individual subjects showed more effort than group performers. Smith and Glass (1980) showed a

similar effect of the number of target persons by looking into research concerned with the relationship between class size and learning success. Their hypothesis was that with increasing size of the class, students should feel less and less social impact from the teacher who has a much higher status in this situation. It was shown that this is indeed the case up to a certain class size (afterwards one student more or less does not matter so much). Thus, productivity increases with a higher social impact.

Self attention theory (Mullen 1987) is a theory of self regulation that explains behavior modifications if one is the subject of one’s own attentional focus. In this case, violations of standard social norms will become more salient. Thus, one could argue if the attribution of self attention to the group of social influence theories is correct. It may also be an appropriate candidate for the self validation group. Because the effects of self attention become only apparent if others are present and because its definition is grounded in this presence of others, the current classification is appropriate.

Whereas social impact concentrates on the mere presence of others, self attention focuses on the effect of one’s peer group—and their presence or absence—on one’s behavior. The effects of self attention are described in relation to one’s own and another subgroup. If more peers are present, this results in decreased self awareness because single individuals will easier go unnoticed. Think about a soccer fan that is more ready to misbehave in a group of his peers on a Saturday afternoon on their way to the game than on his way to work on Monday morning.

The so called other-total ratio is used to describe this effect for the interaction between arbitrary groups. It represents the proportion of the total group that is comprised of people in the other subgroup. The higher the other-total ratio, the more the individual aspires to self regulation. If for instance, the target person’s own subgroup consists of three people including himself and the other subgroup of a single person, then the other-total ratio is $1(\text{other})/4(\text{total}) = 0.25$ (Fig. 2a). Now consider the case where the target person is alone and interacts with one other person (Fig. 2b). In this case, the other-total ratio is $0.5:1(\text{other})/2(\text{total})$. At last consider an agent that is alone and has to interact with a subgroup that is different from his own (Fig. 2c). The other-total ratio in this case is $4(\text{other})/5(\text{total}) = 0.8$. From left to right, self attention for one’s own actions increases. In a, a person is surrounded by his




Fig. 1 Example for balanced and unbalanced configurations




Fig. 2 Examples of configurations for the calculation of the other-total ratio (OTR): *S* self and peers, *O* others

peers and can easily go unnoticed whereas in c, the person is all by himself and can be individually identified increasing his sense of self awareness.

Correspondingly, the need for self regulation increases from a to c. Mullen (1987) reports on his analysis of transcripts from class discussions and attributes the student’s participation to the varying other-total ratios present in these discussions. With higher self awareness, which means smaller groups of students, participation increased significantly on a number of measures like number of turns, seconds spent talking, or number of words spoken.

3 Implementing theories of social group dynamics

In the previous section, we described which theories have been chosen for our toolbox. In this section, we present details on how these theories were implemented. Because IPA represents an active approach to structure interactions, it serves as the fundamental building block for the toolbox. The other theories influence interactions in more subtle ways, either by altering the cognitive representation of personal relations (congruity theory) or by increasing the probability for certain behaviors (social impact, self attention). But only IPA produces an interaction category that can be mapped to the agent’s behavior repertoire in a given application.

3.1 Social knowledge

IPA has already been successfully employed in other systems of social group dynamics (e.g., Guye-Vuillème 2004; Prada and Paiva 2005). Following Guye-Vuillème (2004), it is realized in the toolbox by two types of parameters that force the agents to interact. For each agent, the variables proactivity and reactivity are calculated for every interaction category. Together, they describe how likely it is for the agent to perform the given interaction category.

Proactivity depends on an agent’s personality and its personal relations (see Sect. 2.1). Thus, an agreeable agent with positive relations to another agent has a high motivation for positive social-emotional interactions. Table 1 displays the dependencies of proactivity on the two variables personality and personal relations.

Additionally, proactivity is subject to the social status of an agent. For instance, individuals with a higher social status are more confident that their suggestions are taken up by an interaction partner, which increases their tendency to proactively give advice. Reactivity describes the influence that group members exert on each other by their interactions. Thus, it describes the motivation to interact within a certain category as a reaction to a previous action, e.g., questions excite answers, positive actions tend to evoke

Table 1 Dependencies between interaction category and dimensions of personality and personal relation to calculate the proactivity variable

IPA category	Relevant personality dimension	Relevant dimension of personal relations
Agree	Intelligence Agreeableness	Liking
Show solidarity		Liking Familiarity
Show tension release	Agreeableness Intelligence	Liking
Disagree	Agreeableness Intelligence	Liking
Show antagonism		Liking Familiarity
Show tension	Agreeableness Intelligence	Liking
Ask for opinion	Extraversion Conscientiousness	Trust Familiarity
Ask for suggestion	Extraversion Conscientiousness	Trust Familiarity
Ask for orientation	Extraversion Stability	Trust
Give opinion	Extraversion Conscientiousness	Liking
Give suggestion	Extraversion Conscientiousness	
Give orientation	Extraversion Stability	

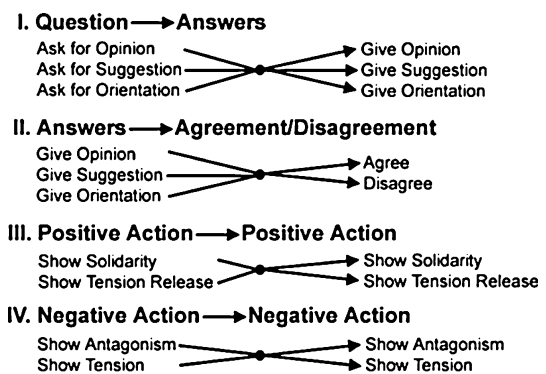


Fig. 3 Tendencies to fit social norms

other positive actions. Bales (1951) summarizes some common action sequences that can be utilized to this end (see Fig. 3).

For the process of action selection, an agent calculates the probability for each of the IPA categories, depending on the two parameters proactivity and reactivity. An interaction is triggered for the category with the highest

probability if it exceeds a given threshold. To finetune this model, a number of parameters can be set like the decay rate for the reactivity.

As can be seen from Table 1, proactivity depends partly on aspects of the social relationships between the two communicating agents. Such relationships on the other hand develop while two agents interact with one another, thus depending on the interaction categories used by the agents. Consequently, there is a mutual influence between the social configuration of the group and the interaction categories that will be chosen. For instance, the more positive interactions two agents perform, the more the trust dimension of their personal relation will increase. In addition, their motivation for interacting with each other as well as the probability for another positive social emotional interaction increases.

3.2 Self validation

Congruity theory has also been successfully implemented in other systems (e.g., Rist and Schmitt 2002). Another layer is added to each agent containing a cognitive net that holds subjective assumptions of the agent about his relations to other agents, his attitudes towards objects, and about presumed relations between other agents and attitudes of other agents towards objects. Relations are expressed as positive or negative and correspond to the liking aspect of the standard representation of personal relations (see Sect. 2.1). Thus, it augments this standard representation by incorporating a cognitive model about the whole group. To initialize an agent’s cognitive net, the liking aspect of the standard representation of the personal relations is exploited. The relation between self and any other agent is identical to the standard representation. The assumptions about the relations between other group members are calculated taking into account the agent’s own relation with the agents in question to achieve a balanced configuration:

$$\text{Rel}(\text{agent}_1, \text{agent}_2) = \text{Rel}(\text{self}, \text{agent}_1) \times \text{Rel}(\text{self}, \text{agent}_2)$$

Afterwards, changes in the cognitive net might happen if an interaction is triggered. An agent can either perform the action, which will not result in changes, or it can be the target or the observer of an action. As the target, the agent is directly involved in the interaction in the role of the addressee. To assume the role of an observer, two interacting agents have to be in the vicinity of the agent. To define such vicinity, Hall’s (1966) ideas on proxemics have been incorporated in the toolbox. Interactions can only be observed if they happen in an agent’s social area (up to 3.6 m). Changes in an agent’s cognitive net are calculated according to Osgood and Tannenbaum (1955):

1. Positive information:

$$\Delta S = \left(\frac{|E_O|}{|S| + |E_O|} \right) \times (E_O - S)$$

$$\Delta E_O = \left(\frac{|S|}{|S| + |E_O|} \right) \times (S - E_O)$$

2. Negative information:

$$\Delta S = \left(\frac{|E_O|}{|S| + |E_O|} \right) \times (-E_O - S)$$

$$\Delta E_O = \left(\frac{|S|}{|S| + |E_O|} \right) \times (-S - E_O)$$

where

- S: value of sender assessment
- ΔS: change in assessment of the sender
- E_O: value of topic assessment
- ΔE_O: change in assessment of the topic.

Interaction categories are ranked and assigned weights according to their influence on the agent’s cognitive model. A *Show Antagonism* action is a stronger negative expression than a *Show Tension* action and thus leads to stronger changes in an agent’s cognitive net. So far, changes in the cognitive net do not have an impact on the process of action selection. This is achieved by updating the liking aspect of the agent’s standard representation of personal relations.

3.3 Social influence

Social impact theory states that the impact felt by an individual is dependent on strength, immediacy, and number of source and target persons in the following way:

$$\text{Social Impact} = \frac{(S \times I \times N)_{\text{source}}}{(S \times I \times N)_{\text{target}}}$$

where

- S: strength
- I: immediacy
- N: number of persons.

To identify the crucial distance from which the impact increases, Hall’s (1966) observations about different areas of personal and public space were considered. If an agent B enters the social area of another agent A (<3.6 m), it will contribute to the impact felt by agent A. Agents within this area are divided into source and target agents. To decide for the appropriate category, we refer to the liking aspect of the respective social relation. Source agents (liking ≤0.5) increase social impact whereas target agents (liking >0.5) decrease it. To specify the strength of sources and targets, we compute a mean value of the social status of the agents

for each group. The same is done for the immediacy, which is calculated on the distance between the agent and the sources and targets.

Productivity increases with higher social impact. This phenomenon is modeled by varying the threshold to interact and by adapting the probability for task-oriented interactions. If the social impact felt by a given agent exceeds a pre-defined level, the threshold to interact is decreased and the probability to conduct a task-oriented interaction is increased in proportion to the amount of impact felt by the agent. If the social impact does not exceed the pre-defined level, the opposite holds true and individuals tend to hide in the crowd. Thus, the threshold to interact increases and the probability for a task-oriented interaction decreases. Self attention theory explains behavior modifications caused by self regulation, which can be described by the so called other-total ratio (see Sect. 2.4):

$$\text{OTR} = \frac{\text{other}}{(\text{other} + \text{self})}$$

where

OTR: other-total ratio

self: number of agents in own subgroup (liking >0.5)

other: number of agents in other subgroup (liking ≤ 0.5).

Like before, only agents within the social area are taken into account. To divide these agents into the self- and other-subgroup the liking value of the personal relation is used. Self attention theory states that the higher the other-total ratio is, the higher is the wish to adhere to social norms. Two versions of describing this effect have been implemented.

In a straightforward version, social norms are interpreted as being modest and polite. Thus, with an increasing other-total ratio the probability for positive social-emotional interactions increases. If the other-total ratio falls below a pre-defined threshold, self attention is marginal and social norms might be disregarded. In that special case, the probability for positive social-emotional interactions decreases. In a more complex version, social norms are adapted to follow well known tendencies in communicative behavior, like giving an answer if posed a question. Figure 3 gives an overview of the action tendencies that have been integrated and that follow Bales' (1951) description. Technically, this is realized by increasing the reactivity for the appropriate interaction category.

4 Using the behavior toolbox as a control component

In the last two sections we have presented the theories that were incorporated into the toolbox and gave some details on the implementation. We are now able to generate

behavior sequences for multiple agents making it suitable as a control component for a multiagent system. In this section, we tackle the challenge of integrating the user into the social group dynamics of the agents and afterwards present some sample applications that make use of the toolbox as a control device.

4.1 Integrating the user

The challenge of integrating the user in a multiagent system which simulates social group dynamics is due to the fact that the user has to develop social relations with the agents to get involved in their social system. This poses the following technical and conceptual problems:

- Navigating: the agents navigate freely in their environment and some of the implemented models rely on the spatial immediacy of interlocutors. Thus, the user should either be able to navigate in the environment, or models like social impact have to be unplugged.
- Communicating: the agents interact via interaction categories. The user has to be able to produce such categories or actions of the user have to be mapped to appropriate interaction categories.
- Developing relationships: the agents' subjective representation of their relationship towards others is calculated by the toolbox according to the selected models. How can this be realized for the user? Should the user interact based on his felt relationship towards his interlocutor or based on the system-calculated relationship parameters?

Exemplary solutions to these problems are detailed with the Virtual Beergarden, a virtual meeting place for agents and users. The Virtual Beergarden was developed to serve as a test bed for interactions between multiple agents and users. There is no special task to be solved apart from meeting other agents, communicating with them and by building up relations with them.

4.1.1 Navigating

To solve the navigation task, a pressure-sensitive dancing pad is employed in the Virtual Beergarden, which is used in many computer games. The user can navigate through the scenario in a first person view and join or leave other agents (see Fig. 4 right). Spatial behavior of the agents adheres to the F-formation system described by Kendon (1991) and takes Hall's (1966) ideas on proxemics into account.

Hall's analysis is primarily concerned with distances, distinguishing four different areas, which are related to behavior changes that occur if someone enters these areas. Kendon takes also the orientation of the interlocutors into account. Depending on their social relations, people will

Fig. 4 The Virtual Beergarden as a meeting place for agents and users. Users navigate by using a dancing pad



orient themselves differently when joining others in public places. Orientations are either closed (interlocutors would not be disturbed by others) or open (interlocutors allow others to enter the conversation). Like the other agents the user has to take these spatial organization patterns into account. If, for example, he gets too close to one of the agents, this results in a position change by the agent to re-establish the appropriate distance and formation (see Rehm et al. 2005 for a detailed account).

4.1.2 Communicating

In the beergarden, agents use natural language utterances to communicate instead of interaction categories. To generate the natural language utterances for an agent, the system makes use of a corpus-based statistical approach that relies on the agent's social relation towards the addressee and the interaction category that is produced by the toolbox (see Rehm et al. 2007). Because the user is situated on the dancing pad in front of a large projection of the 3D environment (see Fig. 4 right), he cannot use standard input devices like the keyboard to type in his utterances.

Because of the general domain of the utterances and the predicted low accuracy of a speech recognition component in this case, the integration of such an input component was not an option either. Instead, we created a semi-automatic way for the user to communicate with the agents. Instead of directly speaking to the agents or selecting pre-defined utterances like it can be seen in many computer games, the user chooses the interaction category he thinks is appropriate. Based on this choice, the system generates a natural

language utterance that is transmitted to the addressee. To switch from navigation to communication mode, the user presses the start button on the dancing pad. The corner buttons of the pad are mapped to the interaction categories (see Fig. 5). By pressing the button for a category the user can switch through the different possibilities available for each category. If the user wants to calm down a somewhat heated discussion, he has to press the button in the upper right corner. This activates the group of positive social-emotional interactions. The default interaction category for this group is "Agree". Pressing twice, the category changes from "Agree" over "Show Solidarity" to "Tension Release", which is the appropriate category for the user's intention. The interaction category is sent to the system by pressing the select button.

By making the statistical language generation component available to the user, he is freed from the burden of learning the right phrases to interact with the agents. Whereas in standard game applications the number of available phrases for a given situation is rather confined, in our approach a new utterance is generated each time the user chooses a category. This ensures a rich repertoire of different utterances and prevents repetitive or boring dialogues. The variability of utterances only depends on the training corpus.

4.1.3 Developing relationships

As we have seen in Sect. 3, interactions generated by the toolbox are based on the social relations between the interlocutors. To create a suitable utterance, the language

Fig. 5 The user communicates by using the dancing pad



engine does not only take the interaction category into account but also the relation between speaker and addressee, i.e., a semantic representation of the utterance is created as input to the language engine based on the interpersonal relations of the speaker towards the addressee and an interaction category (see Fig. 6). Integrating the user in this process is an interesting challenge because either the user has to provide his subjective impression of his relation towards the agent he is interacting with or the system has to calculate the user's relation parameters and use these for the generation task.

The first approach has the obvious disadvantage that it would be counterintuitive and very tedious if the user had to provide his subjective impression every time he wants to communicate. Treating the user as just another agent and integrating him in the calculations of social group behavior by the system, renders the simulation more flawless but poses the problem to the user that his “felt” impression of his relations to the other group members might not be in accordance with the calculated relation parameters.

To keep the interface intuitive, we realized the second option. Thus, the user is treated by the system as just another agent who has the only special feature that he supplies his own interaction category. But his relations towards the other group members are calculated and updated according to the chosen theories. Relations between group members are always subjective for a given individual. Thus, although agent A may have a high liking value towards agent B, this is not necessarily true for agent B in regard to its relation with agent A. As a consequence, the user has two possibilities of monitoring his current status in the group. He can turn on a relation monitor, which depicts the relation parameters liking, familiarity, trust, and commitment towards the current interaction partner. To support the user in establishing relational bonds towards the other agents he can also switch on a liking monitor, which depicts the current liking value of the interaction partner towards the user. Thus, if the agent doesn't like the user, the user might try to remedy this by engaging in additional positive interactions.

On the technical side, the toolbox had to be slightly modified because originally it only handled interactions between agents. Now the user, his interactions and personal

relations come into focus. Therefore a “user agent” is created. It has the same features like the virtual agents including personal relations. As the user was not involved in the group before, the four dimensions of social relations are presumed as neutral, which indicates that the user does not know the agents and vice versa.

4.2 The Perfect Circle

Whereas the Virtual Beergarden is a system specifically designed for testing interaction methods between a user and multiple agents, we also wanted to apply the toolbox also to an existing system for exemplifying the use as a tool for rapidly prototyping different theories of social group dynamics. This was done with the game “Perfect Circle” by Prada and Paiva (2005). In the game, the user plays the role of an alchemist that has joined a group of four other alchemists to undertake the quest for the rainbow pearl. The pearl is hidden in one of the elemental planes, which can only be reached through magic portals that are activated by the powers of gemstones. The group is progressively challenged with the task of opening a portal (see Fig. 7 left). They need to gather and manipulate the gemstones in order to get the required ones that will open the portal (goal in Fig. 7 left). Members of the group have different skills (skills in Fig. 7 left) and may engage in social-emotional interactions during the performance of the task. They can propose an action, manipulate or use the gemstones, express their opinions about the others' proposals by agreeing or disagreeing with them; and can encourage or discourage the others (a trace of the interactions is given in the bottom row of Fig. 7 left).

For this task-oriented application we tested the following combinations of theories. IPA based group development was appropriate to simulate the development of personal relations, for example if one agent encourages another several times this leads to a better personal relation. Furthermore, congruity theory was selected. Consequently, personal relations change if someone does something for the task. Changes can also be observed when interactions take place, which do not include an active participation of the agent. Self attention theory seemed to be a good choice as it deals with group constellations. But as tests with the toolbox showed, it would not be an advantage to apply it in this application. The other-total ratio is quite high, which leads to the result that the agents want to adhere to social norms and try to be friendly. In the context of the application, the game experience suffered from such a reaction, as the agents encouraged each other too often and did not try to solve the task anymore. Instead, to solve the task social impact theory was more appropriate, because it deals with productivity. The higher the social impact is that an individual feels, the more it wants to interact in a task-oriented

```
<SentenceFrame>
<Sentence>Let me get you another beer.</Sentence>
<MeaningElement attribute="category" value="maintain.se.pos.show_solidarity" />
<MeaningElement attribute="relation.liking" value="High" />
<MeaningElement attribute="relation.familiarity" value="High" />
...
<MeaningElement attribute="personality.extraversion" value="Med" />
<MeaningElement attribute="personality.agreeableness" value="High" />
...
<MeaningElement attribute="topic" value="beer" />
</SentenceFrame>
```

Fig. 6 Utterance and underlying semantic representation

Fig. 7 The Perfect Circle game (left) and Second Life (right)



way. As the alchemists in the game felt a quite high social impact (other alchemists with a high social status observed their work) they concentrated on task-oriented interactions and tried to perform well.

Thus, applying the toolbox to the existing system allowed us to test rapidly which theories are suitable in the application domain without having to implement them from scratch.

4.3 Second Life

The Virtual Beergarden as well as the Perfect Circle game are research prototypes that are used by only a limited number of users. Second Life (SL) represents the first massive 3D multiplayer platform that is not primarily concerned with gaming but aims at establishing a general virtual meeting place. Thus, every conceivable type of interaction is in principle possible, be it buying or selling virtual or real goods, or be it playing out as a real DJ in a virtual club. Figure 7 (right) gives an impression of the environment. Two agents—one an unknown user, the other controlled by our toolbox—have met in front of Augsburg's city hall and talk to each other. Central feature of SL is the use of virtual agents as interaction devices which can either represent a real user (avatar) or can be non-player characters (bots). Consequently, SL represents a multiagent system where users in the form of avatars and autonomous virtual agents can engage in social interactions. This offers for the first time the opportunity to test multiagent system techniques in unconstrained tests with an unlimited number of participants in what can count as a “natural” environment for the users. To exploit this possibility, we created a control architecture for autonomous agents in SL that integrates the Behavior toolbox for social group dynamics.¹ The control architecture had to integrate the following components:

- Low-level behavior control: for animating an agent, sending and receiving speech events, and for navigating through the environment, SL provides an open source client which was modified to handle the special needs of coordinated verbal and nonverbal behavior.
- High-level behavior control: to abstract from the tedious work of controlling every parameter for the agent in SL, an abstract control module was realized (BotControl). The BotControl represents the interface between SL and the third-party application. It provides the necessary control methods for agents in SL which can be incorporated in arbitrary applications handling the low-level behavior routines of the agents as well as the event handling for SL events occurring in social interactions.
- Chatterbot functionality (AIML): users communicate in SL via a chat system producing natural language utterances. To realize believable linguistic behavior for our agents, chatterbot functionality was integrated into our system. To this end, a widely used AIML based chatterbot program was extended to deal with interaction categories from the behavior toolbox as a pattern structuring mechanism.

So far, a pilot test has been run relying on IPA to ensure that users interact with our agent. The agent was placed into a sparsely populated area of SL for 7 days. In this time 39 users interacted with the agent. Average interaction time was 6:34 min, ranging from under 2 min to up to half an hour. Some users even wanted to add the agent to their list of friends. This pilot run convinced us that the idea of running a large scale evaluation study in Second Life is indeed feasible.

5 Conclusion

In this article we presented a toolbox that integrates several theories of social group dynamics, which focus on different aspects of group interactions and development. The goal was to provide a tool that can serve as a flexible control

¹ The software is available on request from rehm@informatik.uni-augsburg.de.

component for multiagent systems and that allows for rapidly testing different theories to assess their feasibility in a given application domain. To exemplify that this is possible, the toolbox was integrated into three different applications as a control module.

Acknowledgments The work described in this article is funded by the German Research Foundation (DFG) under research grant RE 2619/2-1.

References

- Bales RF (1951) Interaction process analysis. Chicago University Press, Chicago
- Cassell J, Bickmore T, Campbell L, Vilhjalmsson H, Yan H (2000) Designing embodied conversational agents. In: Cassell J, Sullivan J, Prevost S, Churchill E (eds) Embodied conversational agents. MIT Press, Cambridge, pp 29–63
- Goethals GR (1987) Theories of group behavior: commentary. In: Mullen B, Goethals GR (eds) Theories of group behavior. Springer, New York, pp 209–229
- Guye-Vuillème A, Thalmann D (2001) A high level architecture for believable social agents. *Virtual Real J* 5:95–106
- Guye-Vuillème A (2004) Simulation of nonverbal social interaction and small groups dynamics in virtual environments. PhD thesis, University of Lausanne
- Hall ET (1966) The hidden dimension. Doubleday, New York
- Heider F (1946) Attitudes and cognitive organization. *J Psychol* 21:107–112
- Jackson JM (1987) Social impact theory: a social forces model of influence. In: Mullen B, Goethals GR (eds) Theories of group behavior. Springer, New York, pp 111–124
- Jackson JM, Latané B (1981) All alone in front of those people: stagefright as a function of number and type of coperformers and audience. *J Pers Soc Psychol* 40:72–85
- Kendon A (1991) Conducting interaction: patterns of behavior in focused encounters. Cambridge University Press, Cambridge
- Latané B, Williams K, Harkins S (1979) Many hands make light the work: the causes and consequences of social loafing. *J Pers Soc Psychol* 37:822–832
- McCrae RR, John OP (1992) An introduction to the five factor model and its applications. *J Pers* 60:175–215
- Mullen B (1987) Self-attention theory: the effects of group composition on the individual. In: Mullen B, Goethals GR (eds) Theories of group behavior. Springer, New York, pp 125–146
- Osgood CE, Tannenbaum PH (1955) The principle of congruity in the prediction of attitude change. *Psychol Rev* 62:42–55
- Prada R, Paiva A (2005) Believable groups of synthetic characters. In: Proceedings of AAMAS 2005
- Prendinger H, Ishizuka M (2001) Social role awareness in animated agents. In: Proceedings of agents '01, Montreal, pp 270–277
- Prendinger H, Ishizuka M (eds) (2004) Life-like characters: tools, affective functions, and applications. Springer, Heidelberg
- Pynadath DV, Marsella SC (2005) PsychSim: modeling theory of mind with decision-theoretic agents. In: Proceedings of the international joint conference on artificial intelligence, pp 1181–1186
- Reeves B, Nass C (1996) The media equation—how people treat computers, television, and new media like real people and places. Cambridge University Press, Cambridge
- Rehm M, André E, Nischt M (2005) Let's come together—social navigation behaviors of virtual and real humans. In: Maybury M, Stock O, Wahlster W (eds) Intelligent technologies for interactive entertainment. Springer, Berlin, pp 124–133
- Rehm M, Endrass B, Wissner M (2007) Integrating the user in the social group dynamics of agents. *Soc Intell Des*
- Rist T, Schmitt M (2002) Avatar arena: an attempt to apply socio-physiological concepts of cognitive consistency in avatar-avatar negotiation scenarios. In: Proceedings of AISB '02 symposium on animated expressive characters for social interactions, pp 79–84
- Rubin Z (1973) Liking and loving: an invitation to social psychology. Holt, Rinehart, and Winston
- Smith ML, Glass GV (1980) Meta-analysis of research on class size and its relationship to attitudes and instruction. *Am Edu Res J* 17:419–433