

Social Perception and Steering for Online Avatars

Claudio Pedica^{1,2} and Hannes Vilhjalmsson¹

¹ Center for Analysis and Design of Intelligent Agents, Reykjavik University, Iceland

² Camerino University, Italy

Abstract. This paper presents work on a new platform for producing realistic group conversation dynamics in shared virtual environments. An avatar, representing users, should perceive the surrounding social environment just as humans would, and use the perceptual information for driving low level reactive behaviors. Unconscious reactions serve as evidence of life, and can also signal social availability and spatial awareness to others. These behaviors get lost when avatar locomotion requires explicit user control. For automating such behaviors we propose a steering layer in the avatars that manages a set of prioritized behaviors executed at different frequencies, which can be activated or deactivated and combined together. This approach gives us enough flexibility to model the group dynamics of social interactions as a set of social norms that activate relevant steering behaviors. A basic set of behaviors is described for conversations, some of which generate a social force field that makes the formation of conversation groups fluidly adapt to external and internal noise, through avatar repositioning and reorientations. The resulting social group behavior appears relatively robust, but perhaps more importantly, it starts to bring a new sense of relevance and continuity to the virtual bodies that often get separated from the ongoing conversation in the chat window.

1 INTRODUCTION

Massively Multiplayer Online Role Playing Games (MMORPGs) are a rapidly growing form of mass entertainment delivered over the Internet in the form of live game worlds that persist and evolve over time. Players connect to these worlds using client software that renders the world from their perspective as they move about and meet fellow players. The community is the cornerstone of these games. Therefore, any effort spent on supporting communication and social interaction between players has to be considered valuable for the application. When games wish to use avatars to represent players in environments where they can virtually meet face-to-face, all the animated behaviors that normally support and exhibit social interaction become important. Since players cannot be tasked with micro-management of behavior, the avatars themselves have to exhibit a certain level of social intelligence [1] [2]. The purpose of this avatar AI is in fact twofold: to give players helpful cues about the social situation and to ensure that the whole scene appears believable and consistent with the level of game world realism.

This paper presents ongoing work, which is a collaborative effort between a major MMORPG developer and a research center in the field of embodied conversational

agents that specializes in multi-modal behavior generation based on social and cognitive models. The work presented here is one piece of the project which deals with modeling the relatively low-level motion dynamics in conversational interaction, which despite several ongoing efforts, is largely an unsolved issue.

Many approaches propose interesting solutions for generating the stream of actions that an agent, or in our case, an automated avatar needs to perform in order to believably simulate the positional and orientational movements of a human engaged in a conversation. Each action usually triggers some motor function directly at locomotion level in order to animate the agent. The sequence of actions usually needs to pass through an intermediate layer in order to achieve the desired fluidity of movements and reactions. This extra level between action generation and locomotion, is responsible for smoothing the agent's overall behavior by applying steering forces directly to the underlying physical model. Therefore a steering layer provides a suitable solution for filling the gap between two consecutive actions, generating a net continuous fluid behavior. This approach is particular well suited for modeling unconscious reactions and motion dynamics in conversational interactions.

A conversation defines a positional and orientational relationship of its participants. This arrangement has been described by Kendon [3] as an instance of an F-formation system. Moreover, since a conversation is not a fully rigid formation, external events from the environment or individual behaviour of single participants may produce fluctuations inside this system that lead to compensational rearrangement that avatars can automate without requiring input from their human users. In fact, this is the kind of reactive behaviour which is ill suited for explicit control [1].

2 RELATED WORK

2.1 Automating Avatar Control

In most commercial avatar-based systems, the nonverbal expression of communicative intent and social behavior relies on explicit user input. For example, in both Second Life and World of Warcraft, user can make their avatars emote by entering special emote commands into the chat window. This approach is fine for deliberate acts, but as was argued in [4], requiring the users to think about how to coordinate their virtual body every time they communicate or enter a social situation places on them the burden of too much micromanagement. When people walk through a room full of people, they are not used to thinking explicitly about their leg movements, body orientation, gaze direction, posture or gesture, because these are things that typically happen spontaneously without much conscious effort [3]. Some of these behaviors are continuous and would require very frequent input from the user to maintain, which may be difficult, especially when the user is engaged in other input activities such as typing a chat message. In the same way that avatars automatically animate walk cycles so that users won't have to worry about where to place their virtual feet, avatars should also provide the basic nonverbal foundation for socialization.

Automating the generation of communicative nonverbal behaviors in avatars was first proposed in BodyChat where avatars were not just waiting for their own users to issue behavior commands, but were also reacting to events in the online world according to preprogrammed rules based on a model of human face-to-face behavior [4]. The focus was on gaze cues associated with establishing, maintaining and ending conversations. A study showed that the automation did not make the users feel any less in control over their social interactions, compared to using menu driven avatars. In fact, they reported they felt even more in control, suggesting that the automated avatars were providing some level of support [1]. The Spark system took this approach further by incorporating the BEAT engine [5] to automate a range of discourse related co-verbal cues in addition to cues for multi-party interaction management, and was able to demonstrate significant benefits over standard chat interaction in online group collaboration [2]. The Demeanor system [6] blends user control at several different levels of specification with autonomous reactive behavior to generate avatar posture based on affinity between conversation partners.

However, both Spark and Demeanor assume that the user will bring their avatar to the right location and orient correctly for engaging other users in conversation. Interestingly, even though users of online social environments like Second Life appear sensitive to proximity by choosing certain initial distances from each other, they rarely move when approached or interacted with, but rely instead on the chat channel for engagement [7]. Since locomotion and positioning is not being naturally integrated into the interaction when relying on explicit control, it is worth exploring its automation as well. Some have suggested that once an avatar engages another in a conversation, a fixed circular formation should be assumed [8]. This is a simplification, as the circle we often see is merely an emergent property of a complex space negotiation process, and therefore the reliance on a fixed structure could prevent the avatars from arranging themselves in more organic and natural ways. Therefore we decided on an approach that combines a higher level organizational structure (that of conversations) with low level steering behaviors, affected by a social force field inspired by work on dynamic group behavior in embodied agents such as [9]. We focus on continuous social avatar positioning in this current work, while keeping in mind that other layers of behavior control introduced in previous work will need to be added for fully supporting the social interaction process.

2.2 Dynamic Group Behavior

Numerous works have been published in the area of dynamic group behavior. Most of them concern simulating crowds of people or formations of animals, such as flocks of birds or schools of fish. These kinds of global collective phenomena have been modeled with different approaches but the most interesting and successful of them define the group dynamic as an emergent behavior. In this direction, there are basically two main approaches:

- A particle-based system approach, where particles are animated in real time by the applications of forces.

- An agent-based systems approach, in which each agent is managed in real time by rules of behavior.

The main difference between these two approaches is the order of magnitude of the simulated group behavior. Particle-based systems are well suited for modeling global collective phenomena (such as group displacement and collective events) where the number of individuals is huge. They are all quasi-similar objects and the overall focus is on the global collective behavior. Couzin et al. [10] use an interesting approach to model schools of fish. They define three proximity zones for each particle, where each zone can exert a force on the particle's constant velocity. These forces have different priorities and each of them models a particular behavior. Heigas et al. [11] use a physically-based particle model to simulate emergent human crowd behavior. They use a model of fluid dynamics that incorporates two elementary repulsive forces to simulate jamming and flowing behavior. Treuille et al. [12] present a model for crowd dynamic continuously driven by potential fields. The integration of global navigation planning and local collision avoidance into one framework produces very good video results. Furthermore, their model seems to integrate well with several agent-based models promoting interesting future integrations. Particle-based systems present good performance and scalability but they put more emphasis on the large-scale collective movement and do not model the micro-dynamic of interpersonal interactions.

On the other side, agent-based systems are more suitable for modeling small scale group dynamics. Thalmann et al [13] use complex finite automata to determine the behavior of actors. These automata represent intelligent autonomous behaviors defined by a set of clever rules. The purpose of the model is still to simulate human crowds but this time the introduction of structured behavior of groups and individuals is remarkable. A hierarchical model describes the behavior of each part, but still the rules of social interaction such as conversation are not taken into account. The approach of Reynolds [14] to modeling emergent collective phenomena with an agent-based system, consist of adding a steering motor force to the agent-particle. Despite its simplicity, the idea of an agent-particle with a force-driven behavior leads to the modeling of many complex collective behaviors such as crowd path following, leader following, queuing at a doorway and flocking. Although there isn't any specific reference to conversation dynamics, the elegance and the simplicity of the solution was of great inspiration for the current work. Very comprehensive is also the work of Shao et al. [15], which present fully autonomous pedestrian interacting in a virtually reconstructed Pennsylvania Station. Interesting for the purposes of our work is how they use perceptual information to drive low level reactive behaviors in a social environments. Still motion dynamics for conversations and social interactions are taken into account.

Moving to a more fine-grained social behavior, Rehm et al. [16] recognize the value of social proxemics and formation theories. They use them to inform their models of dynamic distance and orientation between pairs of humanoid agents based on their interpersonal relationship. The focus of their work is modeling how the relationship changes over time based on simulated personality and a theory driven classification of small group interactions. Their underlying model provides an interesting representation of the interpersonal relationships in focused social interactions and how they influence

distances and orientations in small groups of two participants. While interpersonal relationships are necessary for simulating full simulation of social group dynamics, they are not sufficient, as is evident from Kendon's work [3].

The state-of-the-art work on dynamic movement and positioning of social agents, which inspired the current research, is the work of Jan et al. [9] who propose to drive behavior by a social force field. At any point in time, an agent is motivated to move toward a certain position and this motivation is modeled as a force which depends on the position of the other participants. Each social force is activated by certain rules and then they are used to compute a destination point for the agent. While the approach looks promising, the main problem is that motivational forces applied to agent orientation are not taken into consideration.




Fig. 1. The first picture is a top-down view of a conversation; notice the participants social spaces and conversation domain with its o-, c- and b-space. The second picture shows an avatar focused and attentive to another; notice how the head is oriented toward its target and how the beam gets narrow to represent an avatar focused gaze.

3 APPROACH

3.1 Underlying Social Theory

To define a formal model that approximates the physical group dynamics of a real conversation, we start from the basic concepts of social theory. Like some of the more interesting previous work, our primary inspiration has been Kendon's [3] research that defines a conversation as an instance of an F-formation system. Individuals in a conversation tend to arrange in a way that gives all of them equal, direct and exclusive access to a common space. This positional and orientational arrangement is called an F-formation and the set of the behavioural relationships among the participants defines a behavioural system called the F-formation system. From the definition of the

F-formation comes an explanation of the usual circular arrangement of people in conversation with more than two participants: it is simply the best way to give everybody equal access to a common focused space. At the moment, we are not dealing with interpersonal relationships, emotions, affiliation and other important social aspects of group dynamics because they are extensively addressed in many other works. In our approach we start from the fundamental concepts of the F-formation theory in order to cover those aspects of social group dynamics still open and not taken into account by the state-of-the-art avatar environments. From the simple fact that participants in a focused social interaction share a common space with equal access rights to it, comes a series of compensatory movements that have to be simulated before one can hope to completely model social group dynamics.

The common focused space, called the o-space (innermost white circle in Fig. 1, left), is necessary for the social interaction. All the participants perform a series of behaviours that aim to create and defend the o-space. The set of these behaviours defines a system of positional and orientational reactionary behaviours for maintaining an equilibrium in the system. Typically the o-space is sustained through the appropriate orientation of the lower body.

Any external or internal perturbation to the system leads to a series of compensational movements to again reach a stable formation. From this point of view, conversation formations are more similar to bubbles than rigid circles. However, the shape of the formation strongly depends on the number of people involved, the interpersonal relationship among them and from environmental constraints (e.g. the position of furniture inside a room). In the present work, we start by concentrating on simulating the dynamics of conversations in an open space setting, assuming they are strangers with equal social status. Thus we are first dealing with something that tends to approximate a circular formation, but other settings will follow.

Kendon explains a connection between the F-formation system and Goffman's concept of a frame [17]. A frame is a set of rules and social norms that all the participants in an interaction silently accept. The frame comes from the experience of the individual and states what behaviours are meaningful and what conduct is expected in that particular face-to-face interaction. The process of frame-attunement is tightly linked to the F-formation system. By actively maintaining a formation, participants inform each other that they share the same frame for the situation. This further reinforces the definition of an F-formation as a system and moreover describes the system as a unit of behaviour at the interactional level of organization, not at the individual level. Thus, an F-formation system, and consequently a conversation, is a shared social unit.

Since an F-formation system defines a positional relationship among the participants, it is quite natural to take into account the work of Hall and his Proxemics Theory [18]. Basically, in his work Hall claims the existence of four important distances between people as they interact. From shorter to longer, these are: intimate, personal, social and public distances (colored lines in Fig. 1). Proxemics can be seen as an interaction parameter in the sense that for a particular interaction a distance within a specific range is respected. Usually, normal conversations between acquaintances take place in the social

area of the participants, while the personal area is reserved for interaction between close friends. In the present work, we are first looking at acquaintances carrying on normal, daily conversations, but plan to look at other relationships as well.

As part of their autonomous behaviour, avatars must exhibit proper reactions to the positional and orientational fluctuation inside conversation groups and larger changes within their various proximity areas. This is true, even if the avatars are otherwise under user control. In fact, the avatars cannot trust their users to always exhibit proper distances between them in order to preserve the believability of the interaction. In our automated approach, this complex social dynamic is defined as emergent avatar behaviour where a set of rules based on proximity perceptual data, defines a social force field that guides the active steering behaviours of each individual. Furthermore, the idea of conversation as a frame, and thus a social unit, leads to an architecture where we need to treat each social interaction as an important entity in the world.

3.2 The Social Simulation and Visualization Platform

For conducting this research, we felt we needed a special virtual environment that supported and highlighted the type of information and behavior we wanted to develop. We created a tool called CADIA Populus that combines full online multi-player capability with clear visual annotation of the social environment in terms of the theoretical models being used to inform avatar behavior (See Fig.1). The avatars themselves are human 'clay' figures with articulated necks, drawn to scale in the environment for accurately reflecting distances. Users can drive their avatars around and strike up conversations with each other using the built-in text chat. A single user can generate any number of avatars and switch between them at will, which is perfect for manipulating the social situation. A flexible and powerful framework for avatar automation is provided with the tool, part of which is inspired by the OpenSteer³ model which allows steering behaviors to be activated and deactivated. Access to the social environment, through what we call a social perception interface, can be used to trigger behaviors. Examples of social perception include how many individuals are within a certain proximity range (intimate, personal, social and public) or whether the avatar has stepped within the domain of an active conversation according to Kendons F-Formation model. Notice that the perception of social and public space has a blind cone behind the avatar, of respectively 90 and 150 degrees. Both perceptual visualization and the set of active behaviors can be shown inside the environment. The tool is written in Python and uses the Panda 3D⁴ game development library from CMU and Disney.

3.3 Social Situations and Activation of Steering Behaviors

In order to let an avatar generate the proper reactive behavior in every social situation, we activate a set of steering behaviors each of which represents a rule underlying the

³ <http://opensteer.sourceforge.net/>

⁴ <http://www.panda3d.org>

social frame that the avatar is participating in. For example, as soon as an avatar joins a conversation (i.e. enters a new social frame) it will start to keep personal distance, social equality, group cohesion, common attention and domain awareness. Each of them are behaviors in the steering layer and they can generate motivational forces for repositioning or reorienting the avatar. The activation of this set of behaviors produces the reactive dynamics expected from a group of people that has accepted the same social frame, i.e. share the same social norms. In our platform, a steering behavior usually implements activation/deactivation rules, generates motivational force or calls other steering behaviors. Instead of computing steering forces, like in Reynolds [14], our behaviors compute motivational forces such as a desired velocities or a desired direction for a particular avatar's body part. Each behavior has a priority, a weight and a frequency. Priority and weight are used during combination of forces. In particular, we linearly combine (using weights) the desired velocities generated by the active highest priority behaviors and choose the desired direction generated by the highest priority behavior. Some behaviors define in their body part activation rules that depend on perceptual data. For example, the behavior for keeping personal distance will generate a repulsion force only if an individual steps into the avatar's personal space.

In order to provide an example of how this behavior activation works, we are going to succinctly describe how an avatar joins a conversation. As soon as an individual gets close enough to an ongoing conversation, it steps inside the conversation domain. If it keeps moving closer to the conversation, the individual receives an *associated* status. An associated person will be allowed to join a conversation if certain requirements are met. Since our conversations are only modeled as open at the moment, the requirements are very relaxed. In fact is sufficient to have the body oriented toward the o-space of the conversation (intersection of transactional segments) and not move too much (i.e. not be running somewhere). Once an avatar is allowed to join, it is considered inside the conversation social frame, and therefore it is necessary to activate the proper set of rules listed above in order to adapt the agent's behavior to the ongoing social situation and smoothly blend in.

3.4 Conversation Force Field

As mentioned earlier, as soon as an avatar enters a new conversation, a particular set of steering behaviors is activated and run. Some of them generate the low level reactive behaviors of the avatar, which produce repositions and reorientations for keeping the formation's equilibrium. Our purpose is to simulate the compensational movements of the formation when the system is subjected to some fluctuations and for doing so we define a field of forces that drives the avatar movements. This force field is produced and updated by the following steering behaviors: `KeepPersonalDistance`, `KeepConversationEquality` and `KeepConversationCohesion`. Each of these behaviors has its own priority and generates motivational forces at a given frequency using perceptual information. `KeepPersonalDistance` prevents avatars from stepping inside the personal space of someone else generating a repulsion force. The activation rule generates a repulsion only when someone steps into the avatar's personal space. `KeepConversationEquality`

forces an avatar to keep a shared group space of constant size, generating an attraction or a repulsion force toward a specific point. An orientational force toward the group is also generated. This behavior produces motivational forces only if someone is inside the avatar's social space. KeepConversationCohesion prevents an avatar from getting isolated from a group and keeps participants close enough to each other, by generating an attraction force when someone is inside the avatar's public space. Furthermore, an orientational force is generated towards those participants, that for some reason, are going far away from the conversation. The cohesion forces have the highest priority over all the other forces, while repulsion and equality forces are blended together. What follows is a more detailed discussion on how to compute these forces for the social force field of a conversation.

To maintain a minimal distance between individuals we calculate a *repulsion force* (Fig. 2) as follows.




Fig. 2. Repulsion force diagram.

Let N_p be the number of individuals inside the personal area of the avatar, $r \in \mathfrak{R}^3$ the position of the avatar and $r_i \in \mathfrak{R}^3$ the position individual i inside the personal area and Δ_p the personal distance:

$$F_{repulsion} = -(\Delta_p - d_{min})^2 \frac{R}{\|R\|} \quad (1)$$

where $R = \sum_i^{N_p} (r_i - r)$ and d_{min} is the distance of the closest individual.

To sustain the o-space with the individual inside the social area, we calculate an *equality force* (Fig. 3) as follow. This equality force acts like a repulsion or attraction force, depending on the difference between the distance of the avatar from the center of the group in its social area and the mean of every group member's distance from the same center.




Fig. 3. Equality force diagram.

Let N_s be the number of individuals in the avatar's social area, and m the mean distance of the members from the group center. We can calculate the centroid, mean distance and equality force and equality orientation as follow:

$$c = \frac{1}{N_s+1} \left(r + \sum_i^{N_s} r_i \right) \quad (2)$$

$$m = \frac{\|r-c\|}{N_s+1} \sum_i^{N_s} \|r_i - c\| \quad (3)$$

$$F_{equality} = \left(1 - \frac{m}{|c-r|} \right) (c - r) \quad (4)$$

$$D_{equality} = \sum_i^{N_s} (r_i - r) \quad (5)$$

To avoid being isolated, individuals are attracted toward the o-space by means of a *cohesion force* (Fig. 4).

Let N_a be the number of individuals in the avatar's public area, $o \in \mathbb{R}^3$ the center of the conversation and s the radius of the o-space. Then we can calculate the cohesion force and cohesion orientation as follow:

$$\alpha = \frac{N_a}{(N_s + 1)} \quad (6)$$

$$F_{cohesion} = \alpha \left(1 - \frac{s}{\|o - r\|} \right) (o - r) \quad (7)$$

$$D_{cohesion} = \sum_i^{N_a} (r_i - r) \quad (8)$$

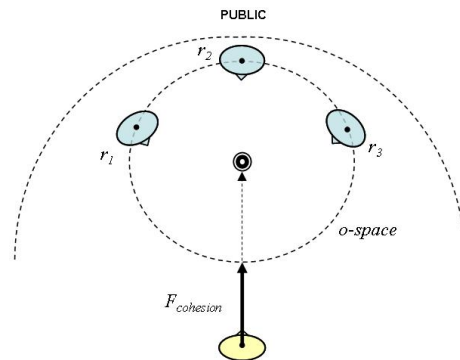


Fig. 4. Cohesion force diagram.

the scaling factor α for the cohesion force is used to reduce the magnitude of this force if the avatar is surrounded by individuals in its social area. The cohesion force is intended to be stronger for those participants far away and isolated from the conversation.

At each simulation time step, using these forces we have a force field that drives the avatar's repositioning and reorienting behavior. Notice that the conversation force field is activated only when the conversation is not in a steady formation and a rearrangement is necessary. Therefore, we are going to perform forces calculation only when strictly necessary. In fact, `KeepConversationEquality` and `KeepConversationCohesion` have a further activation/deactivation rule that runs the behaviors only if strong movements in the conversation is perceived and the formation is not well formed anymore.

3.5 Conversation Unit

In our social force field model, some forces are calculated using parameters (e.g. the center of the o-space or its radius) provided by a *conversation unit*, a small separate module shared by all the participants that represent the conversation as a social unit. The rationale for a high-level social interaction unit comes from the necessity of dealing with group behavior, where an individual is aligned with the same behavioral frame of the other participants but still maintain its own individuality. This is a first step towards simulating the aforementioned frame attunement process. The next step is to keep a conversation frame inside each avatar's mind in addition to the shared conversation unit in the world. The former will maintain all of the avatar's own beliefs about the conversation and its participants and, moreover, will activate context-specific reactive behavior. The latter will uniquely represent the conversation as a world object, with its own position, dimensions and generally visible information about its members. This way it will be possible to let an agent perceive a conversation through its perception system and then, recognizing that conversation as a specific social situation (eg. a frame in its mind), activate the proper reactive behaviors in order to produce a continuous dynamic. Another advantage of this approach is that the conversation unit will be a

suitable entry point in case will be necessary to impose scripted behavior on a group of agents for artistic purposes.



Fig. 5. An open setting with several ongoing conversations.

4 CONCLUSIONS AND FUTURE WORK

While formal evaluation of the effectiveness of the approach for multi-player games has not been performed yet, the visual results (available in video form⁵) and informal interaction tests support the direction of the chosen approach. One of the most immediate impacts on the user is the feeling that conversation groups are robust elements in the environment, that organically maintain sensible formations as players bring their avatars in and out of contact with other players. Additionally, social perception and reactive maneuvering seems to give the avatars a heightened level of realism and a greater relevance to the ongoing social interaction, potentially avoiding the dreaded disconnect between the visual environment and the chat box.

An obvious current limitation that we plan to address is that the locomotion model we currently apply is a simple point mass that, while good in its simplicity, is really far from legged locomotion. We will investigate better techniques for combining steering forces, avoiding jitter and achieving real-time performance with truly large crowds of agents. We will also continue to refine the force field model but, most of all we want to further investigate the benefits of explicitly representing conversation frames. We

⁵ <http://www.ru.is/faculty/hannes/movies/CADIAPopulus/iva2008.html>

believe we have the right underlying model here to simulate a range of different social situations within the same complex environment.

Finally, as we move up from the lower levels of purely reactive behaviors we continue to add higher levels of intelligent control, incorporating some of the behaviors that we have ready from previous work, while also adding new ones that address social scenarios that the specific MMORPG setting calls for. The user interface will also need to grow and adapt to handle the increased complexity without risking the seamless user experience. This is something we believe automation makes easier, but that is an empirical question we look forward to answering.

ACKNOWLEDGMENTS

We are very grateful to the other project team members, in particular Marta Larusdottir, Dr. Kristinn R. Thorisson and Eng. Eric Nivel, for numerous discussions and suggestions. We would also like to thank Torfi Olafsson and the rest of the team at CCP Games for their productive collaboration. Special thanks go to Dr. Adam Kendon for his precious material on face-to-face interaction and valuable personal communication. This work is supported by the Humanoid Agents in Social Game Environments Grant of Excellence from The Icelandic Research Fund.

References

1. Cassell, J., Vilhjalmsón, H.: Fully embodied conversational avatars: Making communicative behaviors autonomous. *Autonomous Agents and Multi-Agent Systems* **2**(1) (1999) 45–64
2. Vilhjalmsón, H.: Animating conversation in online games. *Lecture Notes in Computer Science* **3166**(International Conference on Entertainment Computing) (2004) 139–150
3. Kendon, A.: *Conducting Interaction: Patterns of behavior in focused encounters*. Cambridge University Press, New York (1990) Main Area (nonverbal behavior).
4. Vilhjalmsón, H., Cassell, J.: Bodychat: Autonomous communicative behaviors in avatars. In: *Autonomous Agents*, ACM Press (1998) 269–276
5. Cassell, J., Vilhjalmsón, H., Bickmore, T.: Beat: the behavior expression animation toolkit. In: *SIGGRAPH01*, New York, NY, ACM Press (August 12-17 2001) 477–486
6. Gillies, M., Ballin, D.: Integrating autonomous behavior and user control for believable agents. In: *Autonomous Agents and Multi-Agent Systems*, ACM Press (July 19-23 2004) 336–343
7. Friedman, D., Steed, A., Slater, M.: Spatial social behavior in second life. In: *7th International Conference on Intelligent Virtual Agents*. Volume 4722., Berlin, Springer-Verlag (September 17-19 2007) 252–263
8. Salem, B., Earle, N.: Designing a non-verbal language for expressive avatars. In: *Collaborative Virtual Environments*, ACM (2000) 93–101
9. Jan, D., Traum, D.: Dynamic movement and positioning of embodied agents in multiparty conversation. In: *Proc. of the ACL Workshop on Embodied Language Processing*. (June 2007 2007) 59–66
10. Couzin, I., Krause, J., James, R., Ruzton, G., Franks, N.: Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology* (218) (2002) 1–11

11. Heigeas, L., Luciani, A., Thollot, J., Castagne, N.: A physically-based particle model of emergent crowd behaviors. In: Proc. of GraphiCon. (September 5-10 2003)
12. Treuille, A., Cooper, S., Popovic, Z.: Continuum crowds. In: SIGGRAPH 2006 Papers, New York, NY, USA, ACM (2006) 1160–1168
13. Musse, S.R., Thalmann, D.: Hierarchical model for real time simulation of virtual human crowds. *IEEE Transactions on Visualization and Computer Graphics* 7(2) (2001) 152–164
14. Reynolds, C.W.: Steering behaviors for autonomous characters. In: Proc. of the Game Developers Conference, San Francisco, CA, Miller Freeman Game Group (1999) 763–782
15. Shao, W., Terzopoulos, D.: Autonomous pedestrians. In: Proc. of the ACM SIGGRAPH Symposium on Computer Animation, ACM Publishing (July 29-31 2005)
16. Rehm, M., Andre, E., Nisch, M.: Let's come together - social navigation behaviors of virtual and real humans. In: *Intelligent Technologies for Interactive Entertainment*. Volume 3814., Berlin, Springer-Verlag (November 30 - December 2 2005) 124–133
17. Goffman, E.: *Frame Analyses: An Essay on the Organization of Experience*. Harvard University Press, Cambridge, MA (1974)
18. Hall, E.T.: *The Hidden Dimension*. Doubleday, New York, NY (1966)
19. Brown, B., Bell, M.: Cscw at play: 'there' as a collaborative virtual environment. Volume 6., ACM (November 6-10 2004) 350–359
20. Gillies, M., Dodgson, N.A.: Behaviourally rich actions for user-controlled characters. *Computers & Graphics*, **28**(6) (12 2004) 945–954
21. Shi, J., Smith, T.J., Graniere, J., Badler, N.: Smart avatars in jackmoo. In: *Virtual Reality*, IEEE (March 1999 1999)