Individuality and Contextual Variation of Character Behaviour for Interactive Narrative

M. Gillies*
*Department of Computer Science,
University College London
Malet Place, London WC1E 6BT, UK
m.gillies@cs.ucl.ac.uk

I. B. Crabtree†
†BT plc
Adastral Park, Ipswich IP5 3RE, UK
barry.crabtree@bt.com

D. Ballin‡
‡BT plc
Adastral Park, Ipswich IP5 3RE, UK
daniel.ballin@bt.com

Abstract

This paper presents a system for generating non-verbal communication behaviour suitable for characters in interactive narrative. It is possible to customise the behaviour of individual character using a system of character profiles. This allows characters to have a strong individuality and personality. These same profiles also allow the characters' behaviour to be altered in different contexts, allowing for suitably changing behaviour as the story unfolds.

1 Introduction

Characters are vital to narrative and their behaviour is central to expressing the unfolding story. Their actions play a major part of creating a narrative but other, more subtle, behaviour is also important. Non-verbal communication (also known as "body language") plays a large part in defining characters and displaying the mood of a scene. In this paper we present a system for generating non-verbal communication behaviour suitable for interactive narrative.

Diversity of characters is vital to creating stories, the interplay and conflicts between contrasting personalities is one of the most important elements composing narrative. The differences between characters must be clearly visible in their behaviour, and non-verbal behaviour is one of the most important expressions of personality. As such, it must be possible to customise characters, giving each their own specific mannerisms and behaviour. However, characters should not always act in the same way. A key element of narrative is that as the story unfolds the characters' situation changes and in particular the emotional tone of the story alters. This means that the behaviour of the characters should be able to change to express their new situations. This can happen in many ways. For example, characters should behave differently depending on who they are interacting with, and the relationship between them. Characters can also take on different goals, and different behaviour is appropriate in different places and situtations.

We present a system of character profiles that both allows end users to customise characters, and also allows the characters' beahviour to change in different narrative situations. The profiles system is build on the demeanour framework. Demeanour is a system for generating non-verbal communication. It contains a behaviour language that is used to author autonomous behavioural controllers. These controllers contain a number of parameters that can alter the behaviour of a character.

2 Related Work

Our work builds on a body of work on autonomous characters for virtual environments, for example, Blumberg and Galyean (1995); Badler et al. (1993); Tu and Terzopoulos (1994); Perlin and Goldberg (1996), and Rickel and Johnson (1999). There has been extensive research on autonomously producing expressive behaviour of a number of types including facial expression (Pelachaud and Poggi (2002)), eye gaze (Cassell et al. (1999), Rickel and Johnson (1999) and Gillies and Dodgson (2002)), gesture (Cassell et al. (1999)), style of motion (Chi et al. (2000)) and, like our current implementation, posture (Cassell et al. (2001a), Bécheiraz and Thalmann (1996)).

Maya et al. (2004) have investigated how to create variation between animated characters. They use XML based profiles which are merged with an XML based specification of the affective content of a particular piece of speech, using an XSLT based system, to produce a final piece of behaviour. However, they do not provide any user friendly system for customising characters, nor does their system work in real time. The use of profiles and context dependence has also been used in other types of agent technology, for example, Soltysiak and Crabtree (1998).

3 Non-verbal behaviour

This section describes a behaviour network for non-verbal communication that we have developed. It models the way people relate to each other or their attitude to each other and is based on the work of Argyle (1975). In our model the attitude of one person to another is expressed through posture and, to a more limited degree, gesture. It is discussed in more detail in Gillies and Ballin (2003).

Though there is an enormous variety in the way that people can relate to each other Argyle identifies two fundamental dimensions that can account for a majority of non-verbal behaviour, affiliation and status. Affiliation can be broadly characterised as liking or wanting a close relationship. It is associated with high levels of eye gaze and close postures, either physically close such as leaning forward or other close interaction such as a direct orientation. Low affiliation or dislike is shown by reduced eye gaze and more distant postures, including postures that present some sort of barrier to interaction, such as crossed arms. Status is the social superiority (dominance) or inferiority (submission) of one person relative to another, we will not discuss it directly in our examples.

Figure 2 shows in diagrammatic form a fragment of the attitude behavioural controller that deals with affiliation (status is calculated in a similar way) and posture (eye gaze is discussed in section 3.2). At the top of the diagram the actual value for affiliation is calculated as a weighted sum of a number of factors (for the sake of clarity not all the factors used are actually shown). This is done in two stages, firstly the factors depending on the character itself are calculated. These factors are represented as parameters (here 'liking of other' and 'friendliness' are shown). Then factors depending on the other character's behaviour ('close' and 'distant') are added in, these are taken directly from the controller of the other character. As the behaviours associated with positive and negative affiliation are very different it is split into

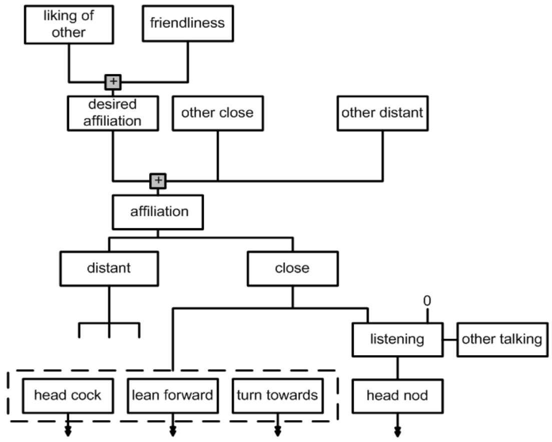

Figure 2: A section of a behavioural controller.

two terms, 'close' which is equal to the affiliation and 'distant' which is its negation. 'Close' is then mapped into actual behaviour (as is 'distant' but it is not shown in the diagram). At semi-regular intervals a new combination of the various behaviours ('head cock', 'lean forward' and 'turn towards') is produced, this combination is always proportional to the value of 'close'. These behaviour types are passed as parameters to the underlying animation system. Another affiliative behaviour is head-nodding, but this is only shown when the other person is talking. This behaviour is controlled by a switch node ('listening'), based on whether the other character is talking. If 'other talking' is true then 'head nod' is proportional to 'close' otherwise it is zero. Figure 3 shows examples of body language generated by the Demeanour framework.

3.1 Posture and Gesture

Human bodies are highly expressive; a casual observation of a group of people will reveal a large variety of postures. Some people stand straight, while others are slumped or hunched over; some people have very asymmetric postures; heads can be held at many different angles, and arms can adopt a huge variety of postures each with a different meaning: hands on hips or in pockets; arms crossed; scratching the head or neck, or fiddling with clothing. Computer animated characters often lack this variety of expression and can seem stiff and robotic; however, posture has been relatively little studied in the field of expressive virtual characters. It is a useful cue as it is very clearly visible and can be displayed well on even fairly graphically simple characters.

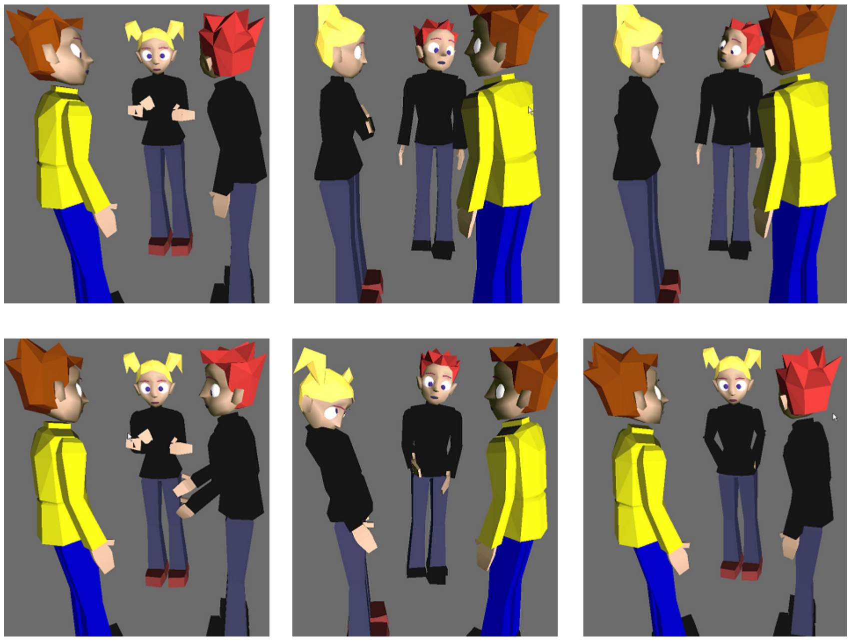

Figure 1: Three characters displaying expressive behaviour. The characters display gaze, posture and gestures behaviour and different attitudes to each other. The female character displays a negative attitude (low affiliation) whereas the male character dressed in yellow displays more positive attitude and the male character dressed in black has a mixed attitude.

Research on posture generation has been limited relative to other modalities. Cassell et al. (2001a) have investigated shifts of postures and their relationship to speech, but not the meaning of the postures themselves. As such their work is complimentary to ours. Coulson (2002) uses an OCC model of emotion to generate postures. Bécheiraz and Thalmann (1996) use a one-dimensional model of attitude, analogous to our affiliation, to animate the postures of characters. Their model differs from ours in that it involves choosing one of a set of discrete postures rather than continuously blending postures. This means that it is less able to display varying degrees of attitude or combinations of different attitudes.

The generation of gestures has been studied by a number of researchers. For example, Cassell et al. (1999) have produced a character capable of extensive non-verbal behaviour including sophisticated gestures. Chi et al. (2000) present a way of generating expressive movements, similar to gestures, using Laban notation. Gestures are closely related to speech and should be tightly synchronised with it. Cassell et al. (2001b) present a system that parses text and suggests appropriate gestures to accompany it. Gestures are less closely related to attitude than posture, though some connection can be made, for example head nodding while listening is a generally affiliative gesture.

As described in the previous section the attitude model generates a high level description of the behaviour of the character in terms of a value of each of a number of behaviour types. The behaviour modules themselves must translate this description into concrete behaviour. Each behaviour type can be expressed as a posture in a number of different ways, for example space filling can involve raising to full height or putting hands on hips while closeness can be expressed as leaning forward or making a more direct orientation (or some combination thereof). Actual postures are calculated as weighted sums over a set of basic postures each of which depends on a behaviour type.

The basic postures were designed based on the description in Argyle (1975) and Mehrabian (1972), combined with informal observations of people in social situations. The weights of each basic posture are the product of the value of its behaviour type and its own weight relative to the behaviour type. The weights of the basic postures are varied every so often so that the character changes its posture without changing its meaning, thus producing a realistic variation of posture over time. Each basic posture is rep-

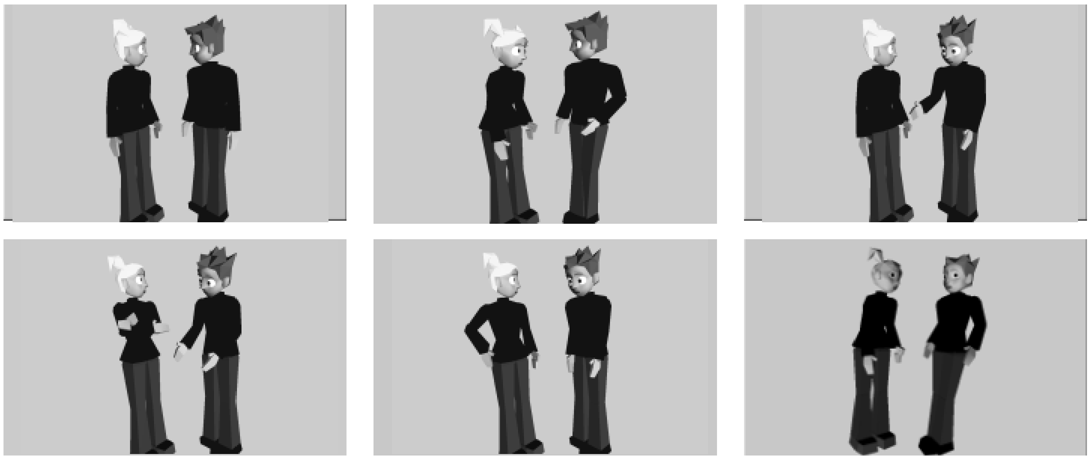

Figure 3: Examples of body language generated by the Demeanour framework to reflect different attitudes between the characters. Clockwise from top left: mutual gaze; close and relaxed postures; the male character is gesturing while talking; the female character has a distant, hostile posture; the female character has a high status, space filling posture the male character has a low status, submissive posture; the male character is relaxed (a high status posture) and the female character has a close posture.

resented as an orientation for each joint of the character and final posture is calculated as a weighted sum of these orientations.

Gesture is generated using the same body animation system as postures, the main difference being that gestures are multi-frame animations and so weighted sums must be performed over a number of frames. They are also no longer merely static poses that can be held for a period of time; they must be repeated at appropriate intervals. More importantly gestures are more closely integrated with the flow of conversation and so must be synchronised with conversation. Of course as the conversation is textual the synchronisation does not have to be as exact as it would be with spoken language. We also do not attempt to parse text so gestures are not strongly connected to the meaning of the text as in Cassell et al. (2001b). Our gesture model serves only to indicate when someone is talking and to express a degree of attitude. Figure 3 shows examples of postures and gestures.

3.2 Eye gaze

Natural eye gaze is critical to the realism and believability of an animated character. This is because eye gaze is fundamental in showing interest levels between characters and as means of anticipating events. Typically a person will look to another before exhibiting any behaviour, such as moving towards them or speaking to them. In conversation, a listener will typically spend a large proportion of their time looking at the speaker. A complete lack of gaze towards the speaker is a clear message of the lack of interest of the audience towards the speaker and will be picked up very quickly. Conversely, mutual gaze, in which two people are looking into each others' eyes is a powerful mechanism that induces arousal in the individuals, so typically mutual gaze is short (of the order of a second).

Argyle and Cook (1976) have done extensive studies with pairs of individuals to understand levels of eye gaze, and mutual gaze, and has detailed results covering (among other things) conversations and the level to which individuals will look at the other while speaking (35%) and listening (75%) etc. We have used these results to influence our model of gaze and mutual gaze in-group settings. Eye gaze is also related to attitude. Higher affiliation results in higher levels of eye gaze. Argyle and Cook have demonstrated compensatory behaviour for eye gaze. People react to higher levels of eye gaze by reacting with more distant postures, and conversely people will look at each other less if they are placed close together.

Existing simulations of eye gaze fall into two broad categories. Chopra-Khullar and Badler (1999) and Gillies and Dodgson (2002) simulate the eye gaze of characters navigating and performing actions in an environment but do not handle social factors of gaze between people. Our work is closer to the other type of simulation that deals primarily with social gaze.

Garau et al. (2001) and Colburn et al. (2000) simulate the patterns of eye gaze between pairs of characters based on frequencies of mutual gaze. Vilhjálmsson and Cassell (1998) use eye gaze to help regulate the flow of conversation by indicating when a speaker is about to finish talking, when someone wants to start or end a conversation and other similar information. Rickel and Johnson (1999), in their character based virtual reality tutoring system, use gaze primarily as a method of indicating to the user an area of interest in the environment. Thórisson (1998) simulates eye gaze in the context of more general work on multi-modal communicative behaviour during conversation.

Each character has a set of foci of interest, which are objects that it will look at. The level of interest is specified as the proportion of time spent looking at that object. So for example if the character is in conversation with another character, while talking the level of gaze will be set to (say) 35%, and whilst listening to about 75% to approximate the natural gaze levels in conversation between two people.

However, this base value is also affected by the affiliation attitude between the character that is looking and the one that is being looked at. A close attitude increases proportion of gaze (up to a maximum of 100%) and distant behaviour reduces it (to a minimum of 0%). The exact formula used to determine the actual eye gaze is:

$$g = g_{cond} - g_{cond}\frac{distant}{d_{max}} + (1 - g_{cond})\frac{close}{c_{max}}$$

where g is the proportion of time spent gazing at the target on average. g_{cond} is the gaze proportion due to the condition (talking, listening or neither). $distant$ and $close$ are the values for the close and distant attitudes and d_{max} and c_{max} are the values at which the gaze proportion is either 0 or 1.

In conversation between people a person will look at another then look away, usually by averting their gaze rather than moving their head, but they are not looking specifically at any other object, just averting their gaze. In our model we achieve this by having a number of 'halo' points around the head of a character that can be selected to look at if we need to look away, and have no other object that demands our attention.

4 Profiles

Demeanour provides a system of character profiles for off-line customization by end-users or world designers. By world designers we mean expert content creators with some programming skills or at least the ability to handle technologies such as XML used when defining character behaviour and adjectives for profile creation (see below). Interfaces for end users are aimed at typical computer game players, not experts but familiar with instant messaging and 3D navigation. Player character's are controlled mostly through a text chat interface, through which players can enter text to be spoke and emoticons which control the character's behaviour (as well as choosing profiles). The user interface is shown in figure 4.

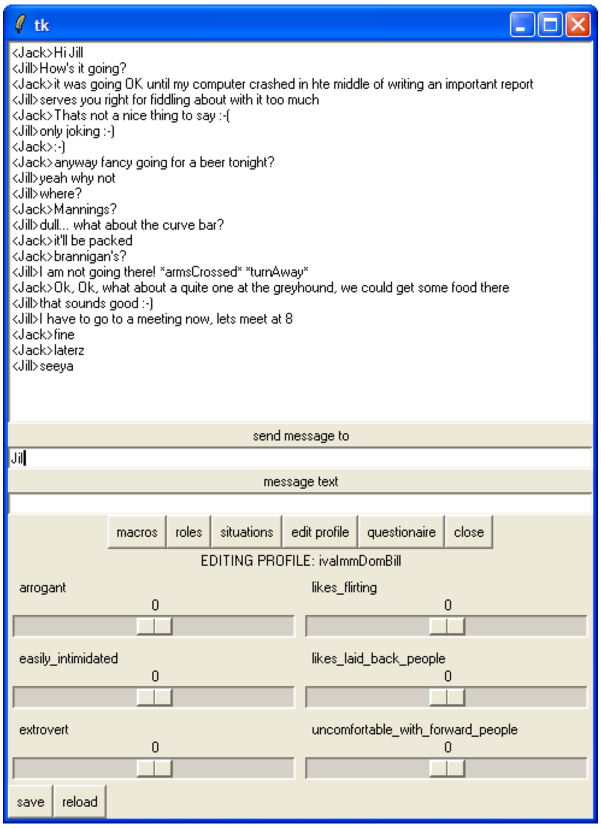

Figure 4: The user interface used for real time control fo chara

A profile is a set of data that determines the unique behaviour of a character, i.e. how it differs from other characters. In Demeanour a character's behaviour is generated by a parameterised behavioural controller (the structure of the controllers is discussed in more detail in Gillies and Ballin (2004)). This controller can be the same for each character but changing the values of the parameters allows for different behaviour.. Customization is possible by altering the values of the parameters e.g the weighting for how the closeness behaviour of other character affects a char-

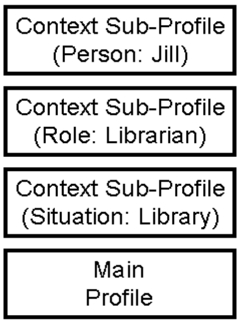

Figure 5: The profiles stack containing a number of loaded contextual profiles

acter's affiliation. A profile can set this weighting to a positive value to achieve reciprocating behaviour, negative for compensation and a low or zero value for indifference to the other's status.

Thus a profile consists of a number of values for parameters of the behavioural controller. These values are stored in an XML-based format separate from the controller definition. When a profile is loaded into a behavioural controller the values in the profile are used to set the parameters of the controller (profile values are matched to parameters by name). Profiles are used as a means of customising a character, and a means of providing contextual variation. This means there will be a number of profiles loaded in a controller at any given time. They are stored in a stack as shown in figure 5. The base of the stack is always the main profile that contains the context independent customisations of a character. Above this, a number of context dependent profiles are loaded as described in section 4.2. When a new context profile is loaded it is added above all the previously loaded profiles in the stack but below the temporary and conversation profiles. Profiles higher up the stack will override profiles lower in the stack, so recently loaded profiles override older ones and user input overrides other profiles. However, this process can be controlled by giving priorities to values within a profile. Values can have two priorities, *required* and *optional*. Required values always override values lower down the stack but optional values only override other optional values, and so are only loaded if no profile has a required value for that parameter.

4.1 Profiles for customisation

The primary function of character profiles is the customisation of characters. End users should be able to customise the behaviour of their character and designers of virtual worlds should be able to provide variety in the autonomous agents in their world. Each character has a main profile, at the base of the profile stack, containing values for the parameters of the behavioural controller that determines the unique behaviour of that character. This is the main customisation system for a given character.

To be an effective customisation method, easy to use tools must be provided for designing profiles. The most direct method is for the user to choose values to parameters whether by hand editing files or via a user interface. However, parameters are often closely linked to the internal workings of the behavioural controller and not necessarily intuitive to end-users, so this method should generally be confined to world designers and advanced users.

We propose the use of "adjectives". These are names in natural language that describe a particular character trait or group of traits that is understandable to end-users. These adjectives are mapped onto actual settings of the internal parameters, each adjective affecting a number of parameters. For example, 'extrovert' might combine dominance with high affiliation while 'easily intimidated' might indicate compensation behaviour to dominance (i.e. responding submissively to dominant behaviour). Each adjective is a fixed set of parameter values and therefore is itself a self contained profile. The adjectives themselves are chosen by world designers. They can be designed at the same time as the behaviour network, through direct profile authoring tools as above. An end-user designs their profile as a combination of the adjectives. They are presented with series of sliders each labelled with an adjective name, the values of the sliders represent the proportions of the various adjectives. The values contained in the adjectives are multiplied by the slider values and summed to obtain the final profile. This provides a customization tool that is easy to use, abstracts from the internal workings of the controller, and is itself easily customizable by world designers. Figure 4 shows an example of the user interface for choosing adjective weights.

4.2 Profiles and context

As described in the introduction the variability of human behaviour is not solely between individuals but within individuals. People behave very differently in different contexts and it is important to also model this sort of variability. The importance of this type of adaptation is brought out in work by MacNamee et al. (2002) and Maya et al. (2004). Goffman (1972) provides a fascinating description of how people's behaviour varies in different contexts. This is particularly true for characters in narrative where the character's behaviour must reflect the unfolding story.

In order to handle this sort of variability Demeanour uses a system of sub-profiles for specific context. A sub-profile is a small set of parameter values that are loaded in a given context to alter the behaviour of the main profile. These are loaded above the main profile in the stack as shown in figure 5, with more recently loaded sub-profiles overriding older ones.

The variation of a person's behaviour in different contexts can depend on a number of different factors and so these contexts themselves can have different meanings, for example, a relationship with a colleague may define a context for interaction with that colleague but the context would also depend on whether they are at work or in a social context. We divide contextual sub-profiles into three types depending on when they are loaded and to some degree who designs them. The system could be augmented to add a number of other types.

Person sub-profiles define a relationship to a particular individual and thus represent the attitude to that person. They are loaded when an interaction starts with that individual. The parameter 'Liking of other' in figures 2 is an example of a parameter that represents an attitude to a particular person and is suited to inclusion in a person sub-profile. Person profiles are generally designed by end users.

Role sub-profiles represent the behaviour when the character is performing a particular role. These roles are often related to work, for example, a waiter behaves very differently when actually serving customers than when not working or when interacting with other staff in the restaurant kitchen (Goffman (1972)). A practical example, might be eliminating all low affiliation behaviour when a waiter is at work to ensure politeness. Role sub-profiles can be loaded by the end-user or automatically loaded in a given situation. Role profiles can be created by world designers to give a user character specific behaviour in a role that the user might not have forseen or they can be created by an end user.

Situation sub-profiles are specific to a particular environment or narrative context, for example, flirting behaviour might be disabled in an office environment but re-enabled in an office party context. They are loaded automatically when the character enters a situation and applied to all characters in that situation.

5 Conclusion

We have described the Demeanour framework's system of profiles which is applied in a number of ways, end user customisation, context dependence and real time control. This system shows promise in producing some of the variability and adaptability of real human behaviour.

Acknowledgements

We would like to thank BT plc. for sponsoring this research. We would also like to thank the UCL Department of Computer Science Virtual Environments and Computer Graphics group for their help and support, and Amanda Oldroyd for the use of her character models.

References

M. Argyle. *Bodily Communication*. Routledge, 1975.

M. Argyle and M. Cook. *Gaze and Mutual Gaze*. Cambridge University Press, 1976.

N. Badler, C. Philips, and B. Webber, editors. *Simulating Humans: Computer Graphics, Animation and Control*. Oxford University Press, 1993.

P. Bécheiraz and D. Thalmann. A model of nonverbal communication and interpersonal relationship between virtual actors. In *Proceedings of the Computer Animation '96*, pages 58–67. IEEE Computer Society Press, June 1996.

B. Blumberg and T. Galyean. Multi-level direction of autonomous creatures for real-time virtual environments. In *ACM SIGGRAPH*, pages 47–54, 1995.

J. Cassell, T. Bickmore, L. Campbell, K. Chang, H. Vilhjálmsson, and H. Yan. Embodiment in conversational interfaces: Rea. In *ACM SIGCHI*, pages 520–527. ACM Press, 1999.

J. Cassell, Y. Nakano, T. Bickmore, C. Sidner, and C. Rich. Non-verbal cues for discourse structure. In *41st Annual Meeting of the Association of Computational Linguistics*, pages 106–115, Toulouse, France, 2001a.

J. Cassell, H. H. Vilhjálmsson, and T. Bickmore. BEAT: the behavior expression animation toolkit. In *ACM SIGGRAPH*, pages 477–486, 2001b.

D. Chi, M. Costa, L. Zhao, and N. Badler. The emote model for effort and shape. In *ACM SIGGRAPH*, pages 173–182. ACM Press/Addison-Wesley Publishing Co., 2000.

S. Chopra-Khullar and N. Badler. Where to look? automating visual attending behaviors of virtual human characters. In *Autonomous Agents Conference*, 1999.

A. Colburn, M. Cohen, and S. Drucker. The role of eye gaze in avatar mediated conversational interfaces. Technical report, Microsoft Research, 2000.

M. Coulson. Expressing emotion through body movement: a component based approach. In Ruth Aylett and Lola Cañamero, editors, *AISB workshop on Animating Expressive Characters for Social Interactions*, Imperial College London, April 2002.

M. Garau, M. Slater, S. Bee, and M. A. Sasse. The impact of eye gaze on communication using humaniod avatars. In *ACM SIGCHI*, pages 309–316, 2001.

M. Gillies and D. Ballin. A model of interpersonal attitude and posture generation. In Thomas Rist, Ruth Aylett, Daniel Ballin, and Jeff Rickel, editors, *Fourth Workshop on Intelligent Virtual Agents*, Kloster Irsee, Germany, September 2003.

M. Gillies and N. Dodgson. Eye movements and attention for behavioural animation. *Journal of Visualization and Computer Animation*, 13:287–300, 2002.

Marco Gillies and Daniel Ballin. Integrating autonomous behavior and user control for believable agents. In *Third international joint conference on Autonomous Agents and Multi-Agent Systems*, Columbia University, New York City, July 2004.

Erving Goffman. *The presentation of self in everyday life*. Pelican, 1972.

Brian MacNamee, Simon Dobbyn, Padraig Cunningham, and Carol O'Sullivan. Men behaving appropriately: Integrating the role passing technique into the aloha system. In Ruth Aylett and Lola Cañamero, editors, *AISB workshop on Animating Expressive Characters for Social Interactions*, Imperial College London, April 2002.

Vincent Maya, Myriam Lamolle, and Catherine Pelachaud. Influences on embodied conversational agent's expressivity. In Patrick Olivier Ruth Aylett and Marc Cavazza, editors, *AISB workshop on Language, Speech and Gesture for Expressive Characters*, pages 75–85, Univesity of Leeds, March 2004.

A. Mehrabian, editor. *Nonverbal Communication*. Aldine-Atherton, 1972.

C. Pelachaud and I. Poggi. Subtleties of facial expressions in embodied agents. *Journal of Visualization and Computer Animation.*, 13:287–300, 2002.

K. Perlin and A. Goldberg. Improv: A system for scripting interactive actors in virtual worlds. In *Proceedings of SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, pages 205–216, New Orleans, Louisiana, August 1996. ACM SIGGRAPH / Addison Wesley.

J. Rickel and W. L. Johnson. Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Applied Artificial Intelligence*, 13:343–382, 1999.

S. J. Soltysiak and I. B. Crabtree. Automatic learning of user profiles – towards personalisation of agent services. *BT Technology Journal*, 16(3):110–117, 1998.

K. Thórisson. Real-time decision making in multimodal face-to-face communication. In *second ACM international conference on autonomous agents*, pages 16–23, 1998.

X. Tu and D. Terzopoulos. Artificial fishes: Physics, locomotion, perception, behavior. In *ACM SIGGRAPH*, pages 43–49, 1994.

H. H. Vilhjálmsson and J. Cassell. Bodychat: Autonomous communicative behaviors in avatars. In *second ACM international conference on autonomous agents*, 1998.