Chapter 6
Interactions with Virtual People:
Do Avatars Dream of Digital Sheep?
***

Mel Slater   Maria V. Sanchez-Vives

Department of Computer Science   *Instituto de Neurociencias*
University College London, Gower Street,   *Universidad Miguel Hernandez de Alicante*
London WC1E 6BT, UK   *Alicante, Spain*
www.cs.ucl.ac.uk/staff/m.slater   *http://in.umh.es/?page=personalsg&key=50*

1. Introduction

In his celebrated book 'Do Androids Dream of Electric Sheep' the author Philip K. Dick explores the relationship between humans and humanoid androids that seem to be human, even superhuman, in every way. The theme is taken up in the film Bladerunner, based on this original story, where the policeman Dekker is required to terminate renegade 'replicants', but falls in love with one himself. In the more recent movie 'A.I.' the question is raised whether you can love a robot (in this case a child) who exhibits every sign of loving you – one slogan of the movie was 'His love is real but he is not'. The book and each of these movies, and many others, explore the moral implications of relationships between humans and machines, machines albeit that are constructed to behave as if human, often also with super-human powers. Their behaviour is so varied, realistic and compelling that the observer is forced to assume that these machines have achieved consciousness, that they know themselves, have emotions and feelings, and know that they have these feelings. Indeed in the case of Dekker there is a tantalizing hint that perhaps he himself is unwittingly a replicant.

These movies and many others paint a popular conception of a world in the not too distant future populated by physically embodied replications of people – robots (entirely non-organic materials) or androids (mixtures of non-organic and organic materials, such as feature in the Terminator movies). The story lines then revolve around relationships between these beings and real humans, whether relationships of exploitation (the replicants in Bladerunner are essentially slaves), proxy love (in A.I. the main artificial character is to be a new son in the family) or war (as in The Terminator). In any case the assumption is that the natural development of artificial intelligence and robotics research will eventually lead to such entities becoming essentially mass produced consumer products available to fulfil a variety of human needs and roles.

This paper explores another form of artificial entity, ones without physical embodiment. We refer to 'virtual characters' as the name for a type of interactive object that have become familiar in computer games and within virtual reality applications. We refer to these as avatars: three-dimensional graphical objects that are in more-or-less human form which can interact with humans. Sometimes such avatars will be representations of real-humans who are interacting together within a shared networked virtual environment, other times the representations will be

of entirely computer generated characters. Unlike other authors, who reserve the term 'agent' for entirely computer generated characters and avatars for virtual embodiments of real people; the same term here is used for both. This is because 'avatars' and 'agents' are on a continuum. The question is where does their 'behaviour' originate? At the extremes the behaviour is either completely computer generated or comes only from tracking of a real person. However, not every aspect of a real person can be tracked – every eyebrow move, every blink, every breath – rather real tracking data would be supplemented by inferred behaviours which are programmed based on the available information as to what the real human is doing and her/his underlying emotional and psychological state. Hence there is always some programmed behaviour – it is only a matter of how much. In any case the same underlying problem remains – how can the human character be portrayed in such a manner that its actions are believable and have an impact on the real people with whom it interacts?

This paper has three main parts. In the first part we will review some evidence that suggests that humans react with appropriate affect in their interactions with virtual human characters, or with other humans who are represented as avatars. This is so in spite of the fact that the representational fidelity is relatively low. Our evidence will be from the realm of psychotherapy, where virtual social situations are created that do test whether people react appropriately within these situations. We will also consider some experiments on face-to-face virtual communications between people in the same shared virtual environments. The second part will try to give some clues about why this might happen, taking into account modern theories of perception from neuroscience. The third part will include some speculations about the future developments of the relationship between people and virtual people. We will suggest that a more likely scenario than the world becoming populated by physically embodied virtual people (robots, androids) is that in the relatively near future we will interact more and more in our everyday lives with virtual people – bank managers, shop assistants, instructors, and so on. What is happening in the movies with computer graphic generated individuals and entire crowds may move into the space of everyday life.

2. Virtual Environment, Immersion and Presence

In most of what follows we will be describing experiments and results that take place within virtual environments (VE) (or 'virtual reality' VR). By a virtual environment we mean a computer generated 'place' in which it is possible for people to interact. There may be events occurring in this place, and there are various forms of interaction. At one extreme the place is static (nothing changes in it) but the human participant can move around within it taking arbitrary positions and orientations. At the other extreme, there may be many events taking place, and the participant is able not only to look at (hear and feel) what is happening but also to intervene and change the course of events.

Virtual environments may be more or less immersive. Immersion breaks down into a number of factors (Slater and Wilbur, 1997):

- Inclusiveness – is the extent to which all sensory data is generated only from within the virtual environment.

- Extensive – is the number of sensory modalities that are accommodated. A system which has vision and sound is more immersive than one that has vision alone.

- Surrounding – is the extent to which the virtual sensory data can be generated from any position and orientation.

- Vividness – the degree of fidelity to every day reality – for example, a system that is able to display shadows in real time is more immersive than one that cannot display shadows.

- Egocentric – information is displayed to the participant to the sense organs in the normal sense of everyday reality. In other words, they see through their own eyes from a first person point of view as if they were there, rather than looking at scenario from the outside (an exocentric point of view).

- Proprioceptive matching – there is a correlation between what they feel as they are moving and what they see, feel and hear as a response. For example, when they feel they are moving their body the sensory response should be appropriate to this – when they turn their head the visual and auditory sense data should match the head turn in exact correlation.

All of the above are ideals, and note that they describe the objective features of a system. Two systems can be at least partially ordered with respect to their degree of immersiveness. Presence, however, is a phenomenon that may arise on the basis of immersion. This is the extent to which individuals respond as if they are in a real world. This response is at multiple levels – low level physiological responses, unconscious behavioural responses, volitional behavioural responses, feelings and emotions, patterns of attention, and so on through to high level cognitive responses – all associated with the feeling of acting in a real place.

In the context of relationships between people and virtual people we can also consider the question of 'presence'. In real life there are typical responses that occur when people interact. For example, eye contact is a particularly important form of interaction which can evoke strong responses – especially if it is held too long. If person A gets too close to person B in a situation which is culturally deemed to be inappropriate (e.g., during a business conversation) it is likely that B will attempt to back away, and feel strong emotions. For people with particular syndromes some types of social interaction can provoke powerful responses. Someone with a fear of public speaking will show strong anxiety responses when forced to be in front of an audience and speak. People with generalized social phobia or shyness will react with strong anxiety to many different types of social situation – such as eating in public, simply attending a party, interactions with members of the opposite sex and so on. People with paranoid delusions will invent entire stories about what is happening around them based on the smallest evidence – a random glance, a coincidental turning away of someone else, two other people who happen to be looking at them while talking amongst themselves, and so on. Confronted with such social situations within a VE, where the other characters are virtual characters, the extent to which these responses are also generated is a sign of presence. In the next section we examine some evidence for this in the context of applications inspired by psychotherapy.

3. Anxiety in Social Situations as a Surrogate for Presence

In this section we consider some examples taken from the realm of psychotherapy. We consider two conditions – social phobia and paranoid ideation. We consider whether people with these conditions will experience similar responses as they would in every day reality. The particular type of social phobia we consider is 'fear of public speaking'. Will people with this condition have the same anxiety speaking to an entirely virtual audience as they would speaking to a real audience? If yes, this is a sign of presence. With paranoia – will people who tend to towards paranoid thoughts (that other people are against them) in everyday life also exhibit such symptoms when the 'other people' are virtual? Again this would be a sign of presence. Of course remember throughout that everyone knows that in fact there are no real people there – so what we are considering are people's automatic responses, not their perhaps higher level thought that 'I know this is not really happening but …' The 'but' is of crucial importance – 'I know this is not really happening but I still feel anxious when those people look at me'.

In this paper we do not at all address the formal scientific results – these are available in other publications which are referenced in the appropriate sections. Here we concentrate only on qualitative aspects of the results.

3.1 Fear of Public Speaking

At UCL in 2000-2001 we carried out an experiment in which more than 40 people were exposed to virtual audiences which had three different types of behaviour – they were static (did not move at all), dynamic and showing very positive responses towards the audience, or dynamic and showing very negative responses towards the audience (Pertaub et al., 2001, 2002). Each person experienced only one of these conditions. Some examples of the positive and negative audiences are shown in Figure 1, each consisting of the same eight male virtual characters who changed posture and facial expression, and also made verbal comments during the progress of the 5 minute talk.

Each person experienced the audience using a head-tracked stereo head-mounted display, as shown in Figure 2. A remote operator unseen by the subjects signalled the sequence of responses of the virtual audience, only choosing the timing of each next response, to ensure that each person saw the same thing.

The statistical results indicated that for those who saw the positive or static audience their reported anxiety provoked as a result of the talk correlated with their usual anxiety in everyday life with respect to fear of public speaking. However, irrespective of everyday life anxiety in relation to public speaking the general trend for those who experienced the negative audience was a very strong anxiety reaction. The experimenters noted anecdotally that such subjects had changes in body posture and skin colour, and overall demeanour after experiencing the negative audience indicating a strong negative reaction. (For ethical reasons each person who was assigned to the negative audience group experienced also the positive group before they went away).

Positive Audience

Negative Audience

Figure 1. Virtual Audience for the Fear of Public Speaking Experiment

Figure 2. Head-Tracked Stereo Head-mounted display used in Fear of Public Speaking Experiment

Here are some of the comments made by the subjects who experienced the positive audience:

'It was clear that the audience was really positive and interested in what I was saying and it made you feel like telling them what you know.'

'I felt great. Finally nobody was interrupting me. Being a woman, people keep interrupting you in talks much more… But here I felt people were there to listen to me.'

‘They were staring at me. They loved you unconditionally, you could say anything, you didn't have to work'.

Here are some responses to the negative audience:

'It felt really bad. I couldn't just ignore them. I had to talk to them and tell them to sit up and pay attention. Especially the man on the left who put his head in his hands; I had to ask him to sit up and listen…. I entered a negative feedback loop where I would receive bad responses from the audience and my performance would get even worse…. I was performing really badly and that doesn't normally happen.'

'I was upset, really thrown. I totally lost my train of thought. They weren't looking at me and I didn't know what to do. Should I start again? I was very frustrated. I felt I had no connection to them. They weren't looking at me. I just forgot what I was talking about.'

These comments might seem to be unsurprising – especially the reactions to the negative audience – for after all, their behaviour was extremely hostile. However, it is important to remember that *there was no audience there!* The situation was entirely virtual. What happened expressed the power of the virtual reality to evoke a response that was similar to that of reality – a 'presence' response.

In a later study we compared people's responses to a virtual reality public speaking scenario where the audience was dynamically behaving but neither negative nor positive but neutral in its response. Half of the subjects of this study had anxiety in relation to public speaking and the other half were confident public speakers. Moreover, there were two scenarios – an empty virtual seminar room or the seminar room with the neutrally behaving audience. Half of each group were assigned to the empty room and half to the populated room. Our prediction was that those with public speaking anxiety would respond with greater anxiety to the populated room than to the empty room, but for the confident speakers it would not make much difference whether the room was populated or not. These results were observed across a range of subjective measures of anxiety response and also by heart rate patterns (Slater et al., 2004).

3.2 Paranoid Ideation

Paranoid ideation is the typical pattern of thinking displayed in cases of paranoia. It is characterised by suspiciousness and beliefs that one is being followed, plotted against, persecuted, and so on. There are degrees of paranoid tendencies in a population – ranging from none at all, all the way through to psychotic illness. In 2003 together with colleagues at the Institute of Psychiatry, Kings College London, an experiment was carried out that tested whether the range of paranoid thoughts typically present in a normal population (but excluding people with psychosis) could be reproduced within a virtual reality (Freeman et al., 2003). This experiment was carried out in an immersive projection system (sometimes called a 'Cave') illustrated in Figure 3.

Figure 3(a) shows someone standing in a white box. In fact the floor and three walls are projection screens onto which a stereo image is back-projected. The glasses worn by the person have left and right lenses switching on and off in synch with the images displayed on the screens, thus creating a stereo illusion.

Images are displayed in Figure 3(b) resulting in the illusion of a kitchen. The person perceives a 3D stereo scene, and since his head position and orientation is tracked the images on the four screen-walls knit together to form one overall 3D scenario with high immersion. More information about such systems can be found in (Cruz-Neira et al., 1993).

(a) The projectors are turned off in the ReaCTor

(b) The scene shows a kitchen

Figure 3. Trimension ReaCTor – A Cave-like system

The subjects in the paranoia experiment had a simple task to move through a virtual library. The characters in the library would look at them and make some facial expressions, and objectively they maintained a neutral attitude towards the subjects. Illustrations are given in Figure 4.

The statistical results supported the hypothesis that paranoid thoughts were triggered in the virtual reality in correlation with subjects' propensities to experience these in everyday reality. Here we quote some of the remarks made by the subjects in post-experimental interviews.

> "The two people to the left, I didn't like them very much – well, I don't know, maybe because when I entered the room I felt I was being watched and then they started talking about me. The other people were more neutral and more inviting except the guy with the beard."

Figure 4. The Scenario for the Paranoia Experiment

"It was probably more real to me than I expected it to be. At some point, I was trying to navigate around a table and almost found myself saying sorry to the person sitting there. I felt that they were getting annoyed with me for doing that…"

"It was really weird, because they were all definitely in on something and they were all trying to make me nervous. It was clear that they were trying to mock me, they kept on looking at me and when I looked back, they were uuhh… The guy with the suit was really weird because he kept smiling at me and it was quite sinister."

It should be noted that there were no sounds in these environments – so what the subjects heard were entirely made up from within their own minds. The results were quite remarkable – people reacted strongly to the virtual characters, even though objectively everyone knew that there were no people there at all. These studies have been followed up and repeated again, with publications pending.

4. Face-to-Face Meetings Between Real People in a Virtual Reality

Above we considered the results of encounters between real people and virtual people. In this section we briefly consider what happens when real people, who may be physically separated by thousands of kilometres, meet face-to-face in a networked virtual environment. By face-to-face we mean that each person sees a virtual character representing the other one. This virtual character speaks in real-time using the real voice of the remote partner, and at least some of the movements (in particular head movements) of the remote person are reflected in the movements of the virtual character.

An example is shown in Figure 5 – where a person is in the immersive projection system interacting with another who is physically remote. They can talk to one another and have the impression of being in the same extended 3D space which is experienced in stereo. For example, they can make 'eye contact' and (virtually) stand very close to one another if they so wished. Of course the remote person may not be in a projection system – they may be using a head-mounted display or even sitting in front of a conventional monitor, keyboard and mouse system, seeing the virtual scenario on the screen in front of them.

Figure 5. A meeting in virtual reality

Many experiments have been carried out examining two and three-party interactions in such environments (Slater et al., 2000; Schroeder et al., 2001). There are several findings. First greater immersion confers greater social power – typically in a mixed encounter where one person is in a more immersive system than the others (e.g., one is in a head-mounted display and the others on desktop display systems) the one in the more immersive display will have a tendency to become the leader (other things being equal).

The avatar representation must have the capability for the people to express themselves – at the very least make eye contact, make gestures. If these basic features of everyday communication are not possible, then social order tends to break down, with some people even believing that others are deliberately trying to be disruptive and thwart the desires of the others. Similarly problems of network delay so that people for example are unable to immediately hear answers to their questions may become major impediments to smooth social interaction.

The avatar representations must match its corresponding behavioural capabilities. For example, in one experiment one avatar looked very sophisticated (photorealistic) but could do nothing – not move its head nor limbs. The other people represented by simple block-like extremely cartoonish avatars became suspicious of the intentions of this more sophisticated looking person. In another experiment (Garau et al., 2003) people experienced either more sophisticated looking or less sophisticated looking partner avatars, and behavioural capability with respect to eye movement was either more or less in conformance with every day expectations during the flow of a conversation. It was found that consistency between appearance and behaviour was very important to maintain presence between the participants.

In spite of technical problems the possibilities are nevertheless very interesting. Here we have people who are physically in remote places who nevertheless can coexist and interact within a shared 3D space, talk and accomplish physical tasks together. One such experiment concluded that when two remote people carried out a task similar to solving a Rubik cube type puzzle together their performance within two 'Cave' systems was almost as good as the performance of people who did the same thing in reality (Steed et al., 2003).

5. Presence and Perception

Above we have informally reviewed some of the evidence demonstrating that people do respond to virtual characters as if they were real. How can this happen? Why do we react to characters that are crude human representations in a way that is similar to how we react to humans? Do emotions overrule our cognitive functions? Understanding these phenomena would take us deep into the workings of the brain and in this section we offer some speculations that may point to where to look for an explanation. First of all, we should consider how we perceive the external world. Perception is constructed from information from different sensory modalities (visual, auditory, tactile, etc.) that arrives to the brain in the form of electrical impulses. For instance, all the information of the visual scene travels codified in a binary code (impulse/no impulse) in a 3 mm wire which is the optic nerve comprised of some one million axons (for each eye). This information is relayed to the thalamic nucleus and from there it will reach the primary visual cortex, to project then to higher or association areas. Different aspects of the visual scene (contours, colours, movement, etc.) will then have to be reconstructed such that they create our internal visual representation of the environment. This processing, from the exterior world all the way up to associative cortices is called bottom-up processing and is driven by so-called feed-forward projections. However, there is evidence to support the view that our perceptual processing does not solely work as a video camera that films and re-plays the environment. One piece of evidence is that the neurons located in the primary visual cortex, and even in the thalamus receive more connections from other neurons in the cerebral cortex (or feed-back connections) than from the sensory organs (Bullier *et al.*, 2001). What this means in functional terms is that the visual reconstruction that these neurons are achieving is deeply influenced by information that already exists in the cerebral cortex, which is called the top-down processing (e.g. Li et al., 2004). What top-down connections convey are the internal factors that affect our sensory perception, including our previous experiences, expectations, motivation, attention, etc. The evidence for the functional correlate of this top-down influence is easily experienced when we look at images like the Kanizsa triangle and similar illusions e.g. (Kojo et al., 1993), confirming the fact that sensory perception is not a passive but an active process where perceiving and interpreting are carried out simultaneously. This implies that our perception is finally determined by the confluence between two streams of information, the bottom-up and the top-down, the external world and the "internal" one.

How does this relate to our perception of avatars? In studies of presence in virtual environments it has been observed that in order to achieve a high level of presence it is not crucial to have a highly realistic visual representation of the environment (Sanchez-Vives and Slater, 2005). A coarse representation may induce a high level of presence; participants in a VE just need to be given some *minimal cues*. The reason why minimal cues can be enough to induce presence is because our perception is an active process and the cortex fills in some of the missing information (Ramachandran and Gregory, 1991; De Weerd et al., 1998). Since our perception is tightly linked to interpretation, the non-sense information is often eliminated and the missing information is filled-in based on previously experienced schemes. One classical example of fill-in processes is the one that takes place in the *blind spot* of our retina. The retina has a discontinuity in the back of the eye at the location were the optic nerve emerges, the optic disc. This area of 1.5 mm diameter lacks photoreceptor cells and should therefore appear as a hole in our perceived visual scene. This does not happen however because that empty space is appropriately filled-in by the brain from surrounding information. In a similar way, probably a

few cues from the avatars such as face elements, eye following, body movements and so on are enticing enough for the brain to perceive this and to react to it as if it were a human.

There is an additional element that makes interaction with avatars special compared to with other elements in the VE: the emotional content. Faces are considered as emotional stimuli. Emotional visual stimuli, are detected faster, they evoke enhanced responses in the visual cortex, and they capture attention more readily than other significant objects in the scene (Lang et al., 1998; Bradley et al., 2003). Face processing occurs even in non-conscious, pre-attentive states, and immediately recalls attention and induces a response. An area in the brain, the fusiform gyrus (inferotemporal cortex)(for a review see Haxby et al., 2000), is specialized in responding to face stimuli. This explains why localized lesions in the cerebral cortex can affect specifically the ability to recognized faces, as it was nicely illustrated in (Sacks, 1998) "The man who mistook his wife for a hat". This area seems to work as well comparing visual inputs with internal representations or recalled images (Frith and Dolan, 1997), or, as we explained above, one of the areas where bottom up and top down information meet each other.

What is so special about faces? Facial expression is one of the elements forming emotion, along with autonomic response (changes in heart rate, respiration, blood pressure, etc.) and with the subjective feelings characteristic with emotion (sadness, fear…). It is also a fast form of non-verbal communication that may have had an important survival role in evolution. For this reason, face recognition is deeply engrained in the brain wiring and even newborn humans express the capability of recognizing schematic representations of faces (Turati et al., 2002). Due to this hard wiring, we willingly identify faces and respond to them, even if they belong to highly schematic virtual characters.

To summarise: top down processing implies that if sufficient minimal cues are provided within a virtual environment, the brain processing fills in missing information. The second element, critical to avatar interaction is that the processing of the human form, especially faces, carries emotional aspects, which may still further minimise the degree of fidelity that needs to be depicted within a VE in order to evoke 'as if it were real' responses.

6. Speculations on the Role of Avatars in the Future

The discussion above was based on a review of scientific work – experiments that were designed the probe the relationships between people and virtual people, and between remotely located real people communicating via avatars, and a brief review of the relevant neuroscience literature about how this process may work. This section is entirely speculative, looking at implications, painting a portrait of future possible proliferation of virtual characters in everyday social life.

We are already familiar with entirely virtual characters in film. For example, in the film Gladiator entire crowd scenes were constructed with computer graphics. In the film Final Fantasy every individual was a computer generated character with almost believable bodies, movements and facial expressions. In films such as Shrek, the characters have excellent postures, gestures, motions and facial expressions, but are mostly humanoid rather than human. In spite of their non-humanness we still laugh at their antics and even maybe identify and sympathise with their plight. The fundamental difference in technology between such movie based virtual characters and those who populate virtual environments is that the movie characters cannot be

generated in real-time (or anything approaching real-time). The rendering and animation requirements are so complex that typically several computers have to work together to produce a frame of animation over several minutes or hours. For a movie every aspect of the rendering and animation has to be as perfect as is possible. Speed is sacrificed to quality. In a virtual reality exactly the opposite must occur – quality is sacrificed for speed, because it is completely impossible to interact with a character which has its display updated once every few seconds or minutes instead of 20 or 30 times a second at the minimum.

However, it is only a question of time. As processing power in general, and in particular the speed and capabilities of graphics processors essentially double once a year, it will soon be possible to populate virtual realities with highly realistic looking and behaving virtual characters – years rather than decades.

Moreover projection and display technology is advancing – albeit at a much slower pace. The idea of Mixed Reality is to bring virtual elements into the realm of everyday life. A 'Cave' is a highly specific example of this – we project onto the walls of a particular built-for-the-purpose room. But the goal is to be able to project anywhere and everywhere, and to blend the virtual into the real in a seamless fashion. 'Presence' still has the same meaning as before – we respond to the virtual as if it were real. However, now the virtual and real form a new totality – neither real nor virtual but something more than the sum of its separate real and virtual parts.

Imagine you walk by a restaurant deciding whether or not to go in. What is important to you is the atmosphere inside – does it seem friendly are there people in there, are they enjoying themselves? You look into one place – and it seems empty – you walk on. You look into another place, and it is full of people. You step inside – and soon come to notice that most of the 'people' in there are in fact virtual characters, programmed to behave as if they were having a great time. Amongst the virtual people are a few real people. Now you can distinguish the difference. You know that in a few years time even this difference will not be noticeable.

Actually the situation is more complex. You notice that at one table there are real people and other virtual people. They all seem to be sitting around the same table and they are conversing. In fact what is happening is that there are real people in this restaurant and real people in another restaurant which might be in another continent, and they are all maintaining the Mixed Reality illusion of sitting down to have dinner together.

Later you need to take a train to get home. You go up to the ticket office inquiry desk because you want to know the most efficient way to get to your home station. The person behind the other side of the glass offering you advice is in fact virtual. The program is sophisticated enough to understand your question and respond with appropriate information. You wonder why anyone has gone to all the trouble to present the information to you in this very sophisticated way, rather than just, for example, on a text display. However, interaction between humans can provide a certain level of reassurance, comfort. After giving you the information the virtual inquiry clerk smiles warmly at you, and wishes you well on your journey. Somehow you cannot help responding in kind, smiling back and wishing the clerk a pleasant evening – even though you know that your reaction is absurd at some level. Absurd but human – it seems that we are programmed to behave in such a way when all the outward signs give the impression that

someone else is there and will respond to you. (The experiments described in the previous section lend support to this supposition).

As you look at the crowds filling up the station concourse, you realize that in fact there are very few real people actually there. There are many reassuring looking people together all having a good time, giving an overall pleasant ambience to this station scene. You're not sure who is real and who is virtual – except that you see a group of children (presumably real) hand in hand with full-sized solid looking characters out of a Disney cartoon.

7. Conclusions

The combination of computer graphics, virtual reality, mixed reality, projection and display technology, artificial intelligence make it more likely that our future will be one where humans interact in their daily lives with virtual characters, rather than the long-predicted time of the coming of the robots. The robots will come but they will be displayed, not physical entities that hurt when they bump into you. As we have seen, being virtual does not reduce the probability that people will react to them with appropriate affective responses.

For people of a certain age today this is already true in a specific sense – people who spend hours playing with computer games already are interacting daily and significantly with virtual characters. People who engage in on-line societies and on-line multi-participant games already partly live in a parallel universe with different normative values, different culture, different objectives.

In time to come such virtual characters individuals and crowds will permeate our society, taking on individual roles as information providers, entertainers, advertisers, advisors, counsellors – and maybe friends, partners, and perhaps – the idea of A.I. – proxy pets and children. At some time in the future we will be in a real scene and no longer know without substantial investigation who is real and who is virtual.

Do avatars dream of digital sheep? Today they don't but we act towards them as if they do. Tomorrow – maybe the answer will be that in fact they do.

Acknowledgments

References

Bradley MM, Sabatinelli D, Lang PJ, Fitzsimmons JR, King W, Desai P (2003) Activation of the visual cortex in motivated attention. Behav Neurosci 117:369-380.

Bullier J, Hupe JM, James AC, Girard P (2001) The role of feedback connections in shaping the responses of visual cortical neurons. Prog Brain Res 134:193-204.

Cruz-Neira C, Sandin DJ, DeFanti TA (1993) Surround-screen projection-based virtual reality: the design and implementation of the CAVE. In: Proceedings of the 20th annual conference on Computer graphics and interactive techniques, pp 135-142: ACM Press.

De Weerd P, Desimone R, Ungerleider LG (1998) Perceptual filling-in: a parametric study. Vision Res 38:2721-2734.

Freeman D, Slater M, Bebbington PE, Garety PA, Kuipers E, Fowler D, Met A, Read CM, Jordan J, Vinayagamoorthy V (2003) Can virtual reality be used to investigate persecutory ideation? J Nerv Ment Dis 191:509-514.

Frith C, Dolan RJ (1997) Brain mechanisms associated with top-down processes in perception. Philos Trans R Soc Lond B Biol Sci 352:1221-1230.

Garau M, Slater M, Vinayagamoorthy V, Brogni A, Steed A, Sasse MA (2003) The impact of avatar realism and eye gaze control on perceived quality of communication in a shared immersive virtual environment. In: CHI '03: Proceedings of the conference on Human factors in computing systems, pp 529--536: ACM Press.

Haxby JV, Hoffman EA, Gobbini MI (2000) The distributed human neural system for face perception. Trends Cogn Sci 4:223-233.

Kojo I, Liinasuo M, Rovamo J (1993) Spatial and Temporal Properties of Illusory Figures. Vision Research 33:897-901.

Lang PJ, Bradley MM, Fitzsimmons JR, Cuthbert BN, Scott JD, Moulder B, Nangia V (1998) Emotional arousal and activation of the visual cortex: an fMRI analysis. Psychophysiology 35:199-210.

Li W, Piech V, Gilbert CD (2004) Perceptual learning and top-down influences in primary visual cortex. Nat Neurosci 7:651-657.

Pertaub DP, Slater M, Barker C (2001) An experiment on fear of public speaking in virtual reality. Stud Health Technol Inform 81:372-378.

Pertaub DP, Slater M, Barker C (2002) An experiment on public speaking anxiety in response to three different types of virtual audience. Presence-Teleoperators and Virtual Environments 11:68-78.

Ramachandran VS, Gregory RL (1991) Perceptual filling in of artificially induced scotomas in human vision. Nature 350:699-702.

Sacks O (1998) The Man Who Mistook His Wife for a Hat: And Other Clinical Tales: Touchstone.

Sanchez-Vives MV, Slater M (2005) From Presence to Consciousness Through Virtual Reality. Nature Reviews Neuroscience 6:8-16.

Schroeder R, Steed A, Axelsson AS, Heldal I, Abelin A, Widestrom J, Nilsson A, Slater M (2001) Collaborating in networked immersive spaces: as good as being there together? Computers & Graphics-Uk 25:781-788.

Slater M, Wilbur S (1997) A framework for immersive virtual environments (FIVE): Speculations on the role of presence in virtual environments. Presence-Teleoperators and Virtual Environments 6:603-616.

Slater M, Sadagic A, Usoh M, Schroeder R (2000) Small-group behavior in a virtual and real environment: A comparative study. Presence-Teleoperators and Virtual Environments 9:37-51.

Slater M, Pertaub D-P, Barker C, Clark D (2004) An Experimental Study on Fear of Public Speaking in a Virtual Environment. In: IWVR 2004. Lausanne, Switzerland.

Steed A, Spante M, Heldal I, Axelsson A-S, Schroeder R (2003) Strangers and friends in caves: an exploratory study of collaboration in networked IPT systems for extended periods of time. In: Proceedings of the 2003 symposium on Interactive 3D graphics, pp 51-54. Monterey, California: ACM Press.

Turati C, Simion F, Milani I, Umilta C (2002) Newborns' preference for faces: what is crucial? Dev Psychol 38:875-882.