Virtual Reality for Social Skills Training

Marco Gillies, Xueni 'Sylvia' Pan
Goldsmiths, University of London

Abstract

Professional work involves a high degree of technical knowledge that must be acquired through years of study at school and university. However, it also involves other, more implicit, skills that cannot be straightforwardly learned from books or lectures. An important example is social skills, interacting with other people in potentially challenging professional situations. These social skills can only be learned through practice in realistic situations. However, real situations do not support the kind of deliberate practice required for improving skills. This paper proposes that Virtual Reality can provide a method of practising social skills that potentially allows deliberate practice in a safe setting. It surveys a range of studies by Mel Slater and colleagues that demonstrate the ability of Virtual Reality to reproduce the experience of social interaction. These show that people respond realistically to virtual humans, opening up the possibility of using them as a means of practicing social interactions. Two experiments in particular show this potential for medical training, involving general practitioners practising difficult consultations with patients. The paper ends with suggestions for what is required to enable deliberate practice of social skills in Virtual Reality.

Keywords: Virtual Reality, Virtual Humans, Social Skills, Learning, Deliberate Practice

Introduction

Pioneering Virtual Reality (VR) Researcher, Jeremy Bailenson, in the title of his recent book (Jeremy Bailenson, 2018) proposes that VR allows us to have "Experience on Demand". VR makes it possible to directly experience a situation, whether it is standing on a plank over a high drop, or walking through a refugee camp, in a way that is not possible in other media. While other media might allow us to watch such an experience at a distance, the immersive nature of VR allows us to feel as if we are experiencing it directly from our own first-person perspective. For Bailenson, a key benefit of this is educational. While many academic subjects can readily be learned from books and lectures, there are certain skills and competencies that can only be learned, or can be learned far better, from experience. Bailenson gives the example of a Quarterback's decision-making process in American Football. Quaterbacks must make important tactical decisions about what play to make in less than a second. Book learning can help in this context, but actual play is far better, though it is very expensive, and dangerous to practise in a full match. Bailenson and colleagues have shown that VR can greatly improve the effectiveness of quarterbacks' training by immersing them in a play scenario in which they have to make precisely the types of decision they would have to make on the field.

Virtual Reality is therefore a powerful medium for the type of experiential learning that is very difficult, if not impossible to learn from books, but can be very expensive or even risky for novices to practice for real. An important example of this type of skill is social interaction. While most of us learn effective social skills for interacting with friends and family throughout our lives, many professions involve complex and difficult social situations that we do not encounter in our daily social lives, whether it is a doctor telling a patient they have a terminal disease or a police officer interviewing a suspect. This position paper will argue that VR is potentially an excellent medium for practising these skills, and will make some theoretical arguments for why that is.

It will begin by discussing the challenges of social skills training. It will then explain unique features of VR and why it is so effective for experiential learning. The following sections will discuss virtual characters and the simulation of social interaction in VR, through a description of a series of experiments by Mel Slater and colleagues that show its effectiveness. The paper will then apply the ideas of these experiments to the problem of social skills training.

Social Skills Training

Donald Schön (Schön, 1983) noted that professional work involves two types of knowledge. The first is technical: engineers need to understand physics and materials and doctors need to understand human biology and pharmacology. This type of knowledge is relatively easily taught in a traditional way with books and lectures, and is the basis of professional training. However, there is another, equally important type of knowledge that is harder to teach. Engineers need to know how to design and doctors how to diagnose. This knowledge is not about following standard procedures, it involves creative thought. What is more, it is not explicit knowledge in the way that physics and biology are, the skills involved are largely tacit, in that practitioners are able to do them, but mostly cannot put into words how they do them. This makes it very hard to teach them explicitly, they can only be learned through experience. Social skills are a clear example of this type of knowledge.

What is more, as Ericsson (Ericsson, 2008) has shown, experience itself is not enough. Doctors who are poor diagnosticians can continue to practice for decades without improving. Ericsson's research indicates that what is needed is a particular form of practice called *deliberate practice* which has a number of characteristics (this list is a condensed version of the list in Ericsson & Pool, 2016):

- Deliberate practice is deliberate in the sense that it is focused on improving performance, not simply doing the task. This means having time dedicated to practice itself, rather than simply learning "on the job", otherwise the learner would have to focus on optimal performance on the task, rather than optimal learning.
- Deliberate practice takes a learner out of their comfort zone. It challenges learners to do tasks they cannot yet achieve (though they should also not be too far beyond the learner's abilities).
- Deliberate practice requires well defined goals and feedback on progress towards those goals. Learners are able to adapt their learning based on this feedback to ensure they improve their performance.

The above three characteristics define a form of practice called *Purposeful Practice* which is more effective than naïve practice, but not as effective as full deliberate practice. This requires additional characteristics, particularly the following:

- Deliberate practice develops skills that are already well understood and have standardized pedagogies and practice regimens that are supported by a skilled teacher.

According to Ericsson (Ericsson & Pool, 2016), deliberate practice aims to develop superior mental representations needed for particular tasks. The term “mental” clearly applies to abstract tasks like mathematics or writing, but Ericsson also applies the term to more physical tasks such as sports or musical performance. In this case it applies to brain circuits that represent physical actions and movements. In Ericsson’s description these are representations of physical movement, but the majority of physical skills are likely to be sensorimotor: the mental representations involved include mappings between sensory input and movement. This view of skills aligns with O’Regan and Noë’s (O’Regan & Noë, 2001) view that perception and action are based on sensorimotor contingencies, lawful relationships between our movements and our sensations. These contingencies are learned. Some in very early childhood, but some are associated with advanced skills are learned later in life, whether they are the sensorimotor relationship between turning a car steering wheel and what we see out of the windscreen or the relationship between pressure and haptic feedback involved in palpation: the medical diagnostic procedure of pressing on a body part. Deliberate practice can therefore be seen as learning sensorimotor mental representations.

How do these ideas apply to social skills? Social interaction involves many complex, multi-level mental representations. Some are relatively abstract, such as reasoning about other people’s motivations or beliefs. Others are much closer to sensorimotor skills. Reading facial expressions is a subtle perceptual skill. It may be purely perceptual, but in actual conversation it is important to interpret the relationship between our actions and our partner's responses. Social skills therefore involve complex sensorimotor representations that map the interplay of our verbal and non-verbal behaviour with that of our conversational partner.

One important factor, particularly in difficult social situations, is that they are strongly emotionally charged. A doctor breaking bad news to a patient, or interacting with an angry patient, must use all of their social skills effectively, while managing their own complex and potentially challenging emotions (possibly including factors such as work stress and cognitive load). This makes these scenarios very difficult.

Social skills and mental representations are largely tacit, we can read expression and pick up social cues without being able to explain how we do it. This places them in Schön's category of professional skills that cannot be learn in the same way as technical knowledge. While some skills, such as reasoning about other people's motivations and goals could be learned from paper scenarios or videos, the sensorimotor skills involved in reading and responding to body language are difficult to learn in any other way than real experience.

However, real social interactions are unlikely to support Ericsson's purposeful or deliberate practice. When interacting with a real patient a doctor must attend to the actual situation and cannot focus on their own learning as is required by deliberate practice. Since real social interactions occur naturally, it is not possible to control their difficulty. This means it is not possible to ensure that learners are outside their comfort zone, but not too far. There is unlikely to be meaningful feedback, people rarely express negative feelings fully and so professionals can leave an interaction unsure whether it was successful or not.

Deliberate practice relies on artificial situations that are geared towards improving skills, that have controlled difficulty and provide feedback relative to set goals. A common approach is a role play, in which pairs of learners work together to practise a social interaction. This can be effective to a degree, however, the situation is obviously artificial and far from a real situation. It is therefore unlikely to have the emotional charge that characterizes a real social situation. In some professions, such as medicine, learners practise with actors. This can feel more realistic than a role play with another student, but is expensive. Students are unlikely to be able to afford the many repetitions that characterize deliberate practice.

For professional social skills to significantly improve we need a way of performing deliberate practice in artificial, but experientially realistic, social situations. The rest of this paper will argue that Virtual Reality could provide this type of practice.

Virtual Reality

Virtual Reality is a medium characterized by immersion: the feeling of being transported into another environment, rather than merely watching it from the outside. This is generally achieved using a display that completely surrounds a participant, typically a head mounted display where the screens are directly in front of the participant's eyes. This display is updated based on the participant's head movements, so that turning their head will result in a new view point, just as it does in the real world.

Mel Slater (Slater, 2009) has proposed that VR consists of 3 major illusions that can be induced in participants:

- **Place Illusion** is the illusion that we are transported to a different place, whether real or entirely virtual. Slater's theory is that place illusion is caused by what O'Regan and Noë (O'Regan & Noë, 2001) call Sensorimotor Contingencies, the patterns of relationship between our movements and our sensory perceptions. This is primarily created through head tracking. The ability to turn or move our heads to look at something anchors us in an environment and makes us feel as if we are there, rather than simply watching from the outside.
- **Plausibility Illusion** is a higher-level illusion that the virtual world we enter is real. While place illusion is caused by low-level sensorimotor correlations, plausibility is caused by higher-level more complex and indirect correlations between our actions and the rest of the world. For example, leaves of a tree moving as we brush past them or another person nodding as we speak.
- **Embodiment Illusion** is the illusion that the body we have in a virtual world is our body. Like the other illusions it is caused by correlations, this time between the movements of our real body and those of our virtual body.

These three illusions are what make a virtual reality experience feel like a real experience. The correlations between our movements and perceptions reproduce those of the real world in a way that other media cannot reproduce. This means that the three illusions are key to why VR is effective for experiential learning. O'Regan and Noë (O'Regan & Noë, 2001) define our experience of the world in terms of sensorimotor contingencies. Slater proposes that VR reproduces these sensorimotor contingencies. Learning in VR is therefore experientially very similar to real world learning. We can therefore learn from a VR experiences as we would from a real experience, but we could not from a book or video.

More precisely, any learning is about developing new mental representations, as proposed by Ericsson (Ericsson, 2008). For many physical and practical skills, these representations are sensorimotor representations, in that they represent relations between our actions and our perceptions (for example, the "hand-eye coordination" associated with skills ranging from catching balls to carving wood). Since virtual reality can reproduce the same sensorimotor relationships as the real world, it can enable us to learn sensorimotor representations that carry over to the real world. This is true for a variety of basic sensorimotor skills, but what about social skills? Can VR simulate the experience of interacting with another person?

Virtual Humans

The power of VR to reproduce real world experiences does extend to social interactions, as shown in a series of experiments by Mel Slater and colleagues. They investigated interactions between people and virtual humans. These virtual humans are animated 3D characters that appear life-sized in front of participants in VR. The characters in these experiments are shown in figure 1.

The first experiment was a virtual audience (Pertaub, Slater, & Barker, 2002). Participants were asked to prepare a short speech and then perform it in immersive Virtual Reality in front of an audience of virtual characters, shown in figure 1, top left. The audience was not particularly graphically realistic, even by the standards of the day, but they were programmed to have realistic body language typically of the audience for a business presentation. There were two experimental conditions, each with different body language. The first was a positive audience, which looked at the speaker attentively and smiled and nodded during the presentation. On the other hand, the negative audience had body language that indicated boredom and even hostility. They slouched, looked around, yawned, muttered to each other and even fell asleep or walked out. The speakers with the negative audience performed significantly worse than those with the positive audience. They evaluated their speech significantly worse, even in categories like the quality of their preparation, which should have nothing to do with the audience. Anecdotally, experimenters noted that participants struggled to talk to the negative audience with some even directly criticizing the audience for their behaviour. This experiment, for the first time, showed that people respond to virtual people in immersive virtual reality in a similar way as they would to real people. The participants knew they were not talking to real people (no attempt was made to deceive them), the characters did not look particularly realistic, and though they had convincing body language this was largely pre-programmed, rather than responding to the participant (though there was an experimenter who triggered certain responses at appropriate movements). Nonetheless, participants were affected emotionally by their experience as if they had given a talk to a real audience.

This shows people's capacity to interpret and respond virtual characters as if they are real. This was supported by an experiment by Freeman *et al.* (Freeman et al., 2008), in which participants were placed in a VR recreation of a typical commuter train modelled on the London Underground. The researchers were investigating VR as a way of treating paranoid ideation: unrealistic negative thoughts about other people. They therefore created a number of virtual humans to share the underground train. These were deliberately programmed with neutral behaviour that did not respond to the participant (small shifts of posture, occasional looking around). The intention was that it was objectively provable that the characters were not making negative judgements about the participant or behaving aggressively. The experiment cited above was conducted with members of the public, not mental health patients. However, even in the general population a sizable proportion of people showed paranoid thoughts about the virtual characters. For example, “There was an aggressive person - his intention was to intimidate me and make me feel uneasy”. Others read positive attitude into the characters: “It was nice, much nicer than a real experience - people aren't so forthcoming with their feelings in a real situation. Thought they were pretty friendly”. This shows that people spontaneously interpret virtual humans as if they are real and have emotional responses to them.

Slater and colleagues (Slater et al., 2006) also created a VR replication of Stanley Milgram’s famous obedience experiments. These controversial experiments involved participants apparently giving electric shocks to other members of the public, even to the level of making the other person lose consciousness. The electric shocks were not real and the other person was, in fact, an actor. However, the experience was extremely traumatic for the participants. The experiments prompted the need for rigorous ethics in research, and it would not be possible to run them now. Slater *et al.* (Slater et al., 2006) however, ran the experiment in VR with a virtual character, without ever attempting to make the participant believe that the character they were giving "shocks" to was real. Despite the virtual nature of the experiment, the majority of experimental participants experienced distress at their actions and a minority did, in fact, feel so distressed that they refused to continue the experiment.

These three experiments show that people do respond realistically to virtual characters, even if they consciously know that they are not real. These responses are not simply play acting, but potentially very strong emotional responses. This shows that virtual reality can reproduce the experience of a social situation, including emotional aspects, and therefore has the potential for social skills training. In fact, scenarios similar to the virtual audience are currently being used to help people overcome fear of public speaking[1].

However, none of these scenarios is like a normal social interaction, since two-way interaction was absent or severely limited. The London Underground scenario involved no actual interaction. For the virtual audience, the interaction was mostly one way, the participant spoke to the audience. The Milgram recreation could be seen as a two-way interaction, but a highly artificial, formalized one that bears little resemblance to a normal conversation. To address this limitation, Xueni Pan, Slater and colleagues (X Pan, Gillies, & Slater, 2015; Xueni Pan, Gillies, Barker, Clark, & Slater, 2012) have developed experiments that involve realistic conversations with virtual characters.

Responsive Virtual Characters

Both of the experiments used a computational architecture, the PIAVCA model, by Gillies et al. (Gillies, 2008) shown in figure 2. It enables real time, responsive interaction that is both verbal (using language) and non-verbal (body

[1] For example, http://www.presentationsimulator.com/fear-public-speaking/

language). It would have been highly challenging to automatically generate real time speech for the kind of free-form conversation using the technology of the time (and it remains challenging with current technology). The PIAVCA model therefore used a Wizard of Oz approach, with an experimenter choosing appropriate utterances from a set of pre-recorded audio clips. However, using Wizard of Oz would not have been possible for non-verbal responses (body language). In part this is because they require very fast responses, faster than would've been possible for an experimenter to respond, but also because body language is largely subconscious, we do it without thinking about what we do, and so an experimenter would not necessarily know what actions would be appropriate in a given context. The PIAVCA model therefore took a semi-Wizard Of Oz approach inspired by the work of Vilhjálmsson and Cassell (Vilhjálmsson & Cassell, 1998). The verbal responses were chosen by an experimenter, but the non-verbal behaviour was generated automatically based on sensor data from the participant. The sensors were similar to those available on a modern head mounted display: 6 degrees of freedom head tracking and voice from a microphone. Though simple, this sensor enabled a range of behaviour such as maintaining a realistic social distance or detecting when the participant was speaking and responding with feedback such as nodding and increased gaze.

In the first experiment (Xueni Pan et al., 2012), male participants were approached by a female virtual character in an environment designed to look like a typical bar. The character was able to hold a conversation with participants using pre-recorded speech and realistic body language. All participants engaged the character in conversation. Though the conversation was designed to be pleasant, the scenario is also a potentially stressful one. All participants showed raised biometric markers of physiological arousal. This was particularly pronounced for participants who reported that they were not in a relationship (they were asked this by the virtual character). Again, this experiment shows that people respond realistically and emotionally to characters, and that this effect also applies to full social interactions with virtual humans.

A follow-on experiment investigated the effect of a virtual human's personality on the interaction, with a virtual character that displayed behaviour associated with either low or high social anxiety (Xueni Pan, Gillies, & Slater, 2015). Participants were asked to interview the character, asking questions of a personal nature that were potentially embarrassing. While the content of the character's speech was the same in both conditions, non-verbal behaviour, both body language and tone of voice, conveyed either social anxiety or confidence. At the end of the interview the character left the room to answer her telephone and did not return. Participants were then left to wait alone. They had been told that they could ring a bell to ask her to return. Participants who had interacted with the anxious character waited significantly longer before ringing the bell. This experiment shows that people obey social norms towards a virtual character, being hesitant to disturb her and also that characters' non-verbal behaviour can have a significant effect on people's responses to them.

Taken together, these experiments show that virtual reality is able to reproduce social experiences in a way that prompts people to respond realistically, both in terms of behaviour and emotion. This is largely due to people's ability to read realistic intentions into a virtual human's behaviour. In particular, people interpret virtual humans' behaviour as responding to them. In the case of the bar experiment, the character's behaviour was genuinely responsive, but in the case of the London Underground experiment, participants interpreted the behaviour as responding to them, even though it was not. In both cases the characters' behaviour is an example of plausibility illusion: a virtual environment that is interpreted as correlating with the participant's behaviour. These are likely to be examples of learned social sensorimotor contingencies, relationships between one person's actions and their perception of their conversational partner. A simple example is the relationship between a participant's speech and the virtual human's feedback behaviour such as nodding or smiling.

The success of these experiments opens up the possibility of using virtual reality for professional social skills training. The following section describes two experiments which look specifically at professional social skills for medical practitioners and which could be used for training.

VR for Medical Social Skills

Medical practitioners need to interact with other people in many difficult and highly stressful situations, including cooperating with other professionals on high pressure procedures, telling patients that they have fatal diseases, or dealing with demanding patients. These require complex social skills, but training is very difficult because the situations involved are highly safety critical or high pressure and it is difficult to allow trainees to have access to them in real life. Virtual Reality could therefore be an excellent way to practice this skill, as shown by the work of Benjamin Lok and colleagues (Kleinsmith, Rivera-Gutierrez, Finney, Cendan, & Lok, 2015; Robb et al., 2016). This section describes two experiments by Pan *et al.* which involve VR recreations of difficult social situations typical of medical practice.

The first experiment (Xueni Pan et al., 2016) involved general practitioners (family doctors). One of the most difficult situations for modern doctors is antibiotic prescription. Prescribing antibiotics in large quantities increases the possibility that a new bacteria will emerge that is resistant to these drugs, raising the spectre of many deadly and untreatable diseases. So antibiotic resistance is in many ways caused by doctors prescribing antibiotics when they should not. This tendency to prescribe is in large part due to pressure from patients, who perceive that antibiotics will help them even if they are not medically appropriate (for example for flu and colds). In the experiment, experienced doctors and trainees had a consultation with a virtual woman and her virtual mother who is presented as suffering from an illness that clearly should not be treated with antibiotics. The woman was very insistent that her mother should have antibiotics and did not accept any of the doctors' arguments that they were not appropriate.

The results were striking. Of the 9 trainees tested, 8 prescribed antibiotics. Of the 12 experienced doctors, 7 prescribed antibiotics. These results are remarkable. The doctors were all aware that the patients were not real, no attempt was made to deceive them. They were all also aware of the dangers of antibiotic resistance and the importance of not prescribing antibiotics. They were also aware of being watched by an experimenter who may have judged them as less professionally competent for prescribing antibiotics. Nevertheless, a large majority did prescribe. This indicates that the doctors felt subject to social pressure from the virtual patients in much the same way as they would from real patients.

This result supports the conclusions of the other studies described above: that people respond realistically to virtual people, and in particular respond with realistic emotional responses. In addition it shows that virtual humans can reproduce the kind of social pressure that we are subject to in many difficult social situations. In particular, it shows that it can reproduce the key social and emotional elements of professional social situations. At the moment professionals perform as badly in VR as they would in the real world. However, because it is sufficiently realistic experientially, it also has the potential for improving doctors' performance because they can practice in a virtual situation that reproduces many of the key features of the real situation.

The second experiment (Xueni Pan et al., 2018) concerned another important task for general practitioners: child safeguarding. In the UK, doctors are required to identify and report potential child abuse issues. This is a difficult task given that they only typically have 10 minutes with patients and their consultations will concern medical issues which may be unrelated to any abuse that is occurring.

The scenario designed involved a consultation with an adult male patient seeking advice on surgery for a serious medical condition. The patient brought along his son to the consultation. Though the consultation was about the father's health and no issues around safeguarding were explicitly raised, there were several cues to child abuse (there were two conditions, with stronger and weaker cues).

The doctors' ability to detect the cues depended on the level of cue (doctors were more likely to detect abuse with more obvious cues) and also their personality (neuroticism was associated with not detecting abuse, and extraversion with detecting it).

The task involved detecting subtle verbal and non-verbal cues in the interaction between father and son, while simultaneously dealing with the medical situation at hand. The first result (level of cue) suggests that detecting social cues had an important effect on detecting abuse. The second result potentially suggests that doctors with better social skills were better at detecting abuse. This, in turn, suggests that medical social skills training might improve detection of safeguarding issues.

Both scenarios show that VR can be used to reproduce medical social situations and that they can present important social challenges to doctors. They, however, test very different skills. The first is dealing with, and withstanding, social pressure in order to achieve an appropriate outcome. This scenario is about how a doctor interacts with patients and responds to their requests. The second, deals with detecting social cues to abuse. The skill is more about observing cues in the context of a complex medical scenario that is not related to those cues. This shows that VR can test a range of social skills.

More importantly, they both suggest educational benefits. While in the context of the experiment, the VR scenario could be seen as testing doctors' skills, they could also be used to practise and improve those skills. Both scenarios involved complex skills that were challenging for the participants, but that could potentially be improved with more practice. This type of VR scenario could allow them to practise those skills in a safe environment. We will discuss the value of VR for learning social skills in the next section.

Discussion

This survey of experiments shows that VR scenarios with virtual humans are able to recreate the experience of difficult social interaction, suggesting they can be used for learning social skills. In particular, they reproduce many of the implicit and emotional aspects that make social skills so hard to learn and practise in other ways.

Compared to real professional situations (e.g. real medical consultations), VR is cheap and low risk. This makes it very suitable for practising social skills. However, following Ericsson (Ericsson, 2008; Ericsson & Pool, 2016), experience and practice are not sufficient for developing true expertise. The practice must be *deliberate* in the sense that it must conform to a number of constraints. Let us examine each in turn to see how VR training supports it.

The practice must be deliberate in the sense that it must be for the explicit purpose of improving performance, not simply for performing as well as possible, and learners must attend explicitly to improvement. VR is well suited to this kind of practice because it is not a real situation. In a real consultation doctors must attend to performing as well as possible for the benefit of the patient, not to improving their own performance. The fact that the situation does not have the pressure of real performance can give learners more space to be aware of and reflect on their performance.

The practice must push learners outside their comfort zone, but not too far. The two scenarios described above were challenging for participants. However, repeated practice with the same scenario would quickly become easy. This raises the importance of having a wide range of scenarios and that these scenarios present different levels of difficulty so they remain challenging as the learner’s skills develop. This is very possible in VR, where the exact details of all interactions can be controlled much more carefully than in real life. Scenarios can be set up to vary in a number of dimensions, including different aspects of difficulty. This requires a lot more content creation than the experiments describe here, but is very possible given sufficient budget.

Deliberate practice requires clear goals and feedback on those goals. None of the experiments above involve an element of feedback, so this is a key issue for educational development of VR. Goals must be set by educators and will vary according to the skills required. Feedback must be provided based on these goals. VR has the benefit that it is possible to record large amounts of data about interactions. It may be possible to extract automated feedback metrics from these data, but the complexity and subtlety of social skills means that the usefulness of automated metrics is likely to be limited. Feedback is likely to rely on teachers’ judgements and the ability to review a performance in VR (which is relatively straightforward in VR).

The final requirement of deliberate practice is that it is done in the context of a well understood pedagogy that enables a formalized practice regimen. This kind of formal pedagogy is likely to emerge slowly (if at all). On this factor a fully deliberate practice regimen is likely to take many years to emerge. Until then social skills learning is likely to be closer to Ericsson’s weaker (though still effective) form of practice called Purposeful Practice.

Conclusion

This paper has presented a number of studies that have demonstrated that VR can reproduce important experiential aspects of social interactions and can therefore form the basis of a form of deliberate practice regimen for professional social skills.

There are important developments that are still required. Some are technical, like the development of a wide range of scenarios that are variable in terms of difficulty and also the development of methods for giving feedback. These require resources, but can be achieved relatively quickly. The harder challenges are pedagogical: developing a set of goals for practice and a pedagogical regimen equivalent to standardized training in music or chess. This is likely to take many years to develop, even if it is possible at all in such a complex domain as social skills.

References

Ericsson, K. A. (2008). Deliberate practice and acquisition of expert performance: A general overview. In *Academic Emergency Medicine* (Vol. 15, pp. 988–994). http://doi.org/10.1111/j.1553-2712.2008.00227.x

Ericsson, K. A., & Pool, R. (2016). *Peak: Secrets from the New Science of Expertise*. Eamon Dolan/Houghton Mifflin Harcourt.

Freeman, D., Pugh, K., Antley, A., Slater, M., Bebbington, P., Gittins, M., … Garety, P. (2008). Virtual reality study of paranoid thinking in the general population. *British Journal of Psychiatry*, *192*(4), 258–263. http://doi.org/10.1192/bjp.bp.107.044677

Gillies, M. (2008). Piavca: A Framework for Heterogeneous Interactions with Virtual Characters. In *IEEE Virtual Reality Conference 2008 (VR 2008), 8-12 March 2008, Reno, Nevada, USA, Proceedings* (pp. 255–256). IEEE.

Jeremy Bailenson. (2018). *Experience on Demand: What Virtual Reality Is, How It Works, and What It Can Do*. W. W. Norton Company.

Kleinsmith, A., Rivera-Gutierrez, D., Finney, G., Cendan, J., & Lok, B. (2015). Understanding empathy training with virtual patients. *Computers in Human Behavior*, *52*, 151–158. http://doi.org/10.1016/j.chb.2015.05.033

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *The Behavioral and Brain Sciences*, *24*(5), 939-973; discussion 973-1031.

Pan, X., Collingwoode-Williams, T., Antley, A., Brenton, H., Congdon, B., Drewett, O., … Delacroix, S. (2018). A Study of Professional Awareness Using Immersive Virtual Reality: The Responses of General Practitioners to Child Safeguarding Concerns. *Frontiers in Robotics and AI*, *5*, 80. http://doi.org/10.3389/frobt.2018.00080

Pan, X., Gillies, M., Barker, C., Clark, D. M., & Slater, M. (2012). Socially anxious and confident men interact with a forward virtual woman: an experimental study. *PloS One*, *7*(4), e32931. http://doi.org/10.1371/journal.pone.0032931

Pan, X., Gillies, M., & Slater, M. (2015). Virtual Character Personality Influences Participant Attitudes and Behavior - An Interview with a Virtual Human Character about Her Social Anxiety. *Frontiers in Robotics and AI*. http://doi.org/10.3389/frobt.2015.00001

Pan, X., Gillies, M., & Slater, M. (2015, February 10). Virtual Character Personality Influences Participant Attitudes and Behavior - An Interview with a Virtual Human Character about Her Social Anxiety. *Frontiers in Robotics and AI*. Retrieved from http://discovery.ucl.ac.uk/1469972/1/frobt-02-00001.pdf

Pan, X., Slater, M., Beacco, A., Navarro, X., Swapp, D., Hale, J., … Delacroix, S. (2016). The Responses of Medical General Practitioners to Unreasonable

Patient Demand for Antibiotics - A study of medical ethics using immersive virtual reality. *PLoS ONE*.

Pertaub, D.-P., 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*. http://doi.org/10.1162/105474602317343668

Robb, A., Kleinsmith, A., Cordar, A., White, C., Wendling, A., Lampotang, S., & Lok, B. (2016). Training Together: How Another Human Trainee's Presence Affects Behavior during Virtual Human-Based Team Training. *Frontiers in ICT*, *3*, 17. http://doi.org/10.3389/fict.2016.00017

Schön, D. A. (1983). *The Reflective Practitioner: How Professionals Think in Action*. Basic Books. Retrieved from http://books.google.co.uk/books?id=ceJIWay4-jgC

Slater, M. (2009). Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Philos Trans R Soc Lond B Biol Sci*, *364*(1535), 3549–3557. http://doi.org/10.1098/rstb.2009.0138

Slater, M., Antley, A., Davison, A., Swapp, D., Guger, C., Barker, C., … Sanchez-Vives, M. V. (2006). A virtual reprise of the Stanley Milgram obedience experiments. *PLoS ONE*. http://doi.org/10.1371/journal.pone.0000039

Vilhjálmsson, H. H., & Cassell, J. (1998). BodyChat: Autonomous Communicative Behaviors in Avatars. In *second ACM international conference on autonomous agents*.

Figure 1: Slater's experiments on Virtual Characters. (left to right and top to bottom): the virtual audience(Pertaub et al., 2002), the London underground (Freeman et al., 2008), the Milgram recreation(Slater et al., 2006), the virtual bar(Xueni Pan et al., 2012), the socially anxious virtual character(X Pan et al., 2015) and the medical consultation(Xueni Pan et al., 2016).

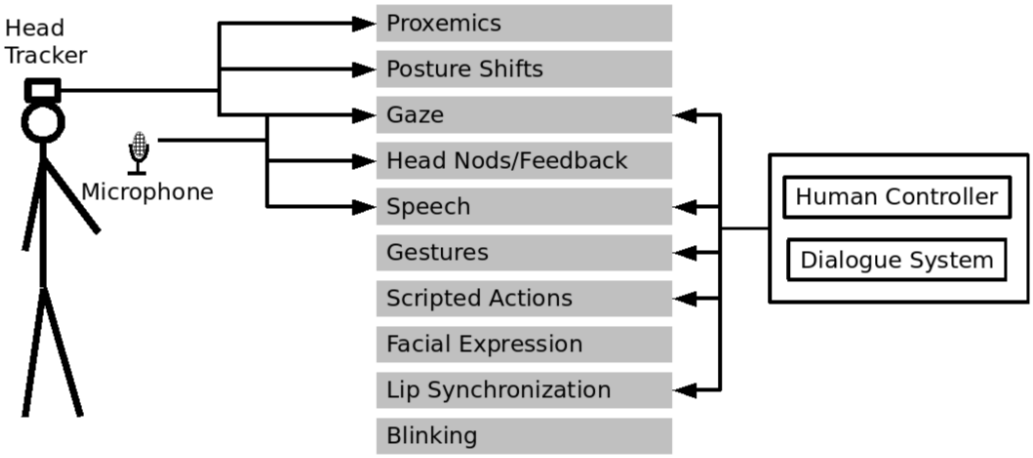

Figure 2: The PIAVCA model of virtual character behavior used in most of the experiments discussed in this paper.