Head Tracking Based Avatar Control for Virtual Environment Teamwork Training

Stefan Marks*, John Windsor†, Burkhard Wünsche‡

*School of Computing and Mathematical Sciences, Faculty of Design and Creative Technologies
Auckland University of Technology
Private Bag 92006, Auckland 1142, New Zealand
email: stefan.marks.ac@gmail.com
www: http://cs.auckland.ac.nz/~stefan

†Department of Surgery, Faculty of Medicine and Health Sciences
The University of Auckland
22 Princes Street, Auckland 1010, New Zealand
email: j.windsor@auckland.ac.nz

‡Department of Computer Science, Faculty of Science
The University of Auckland
22 Princes Street, Auckland 1010, New Zealand
email: b.wuensche@auckland.ac.nz
www: http://cs.auckland.ac.nz/~burkhard

Abstract

Virtual environments (VE) are gaining in popularity and are increasingly used for teamwork training purposes, e.g., for medical teams. One shortcoming of modern VEs is that nonverbal communication channels, essential for teamwork, are not supported well. We address this issue by using an inexpensive webcam to track the user's head. This tracking information is used to control the head movement of the user's avatar, thereby conveying head gestures and adding a nonverbal communication channel. We conducted a user study investigating the influence of head tracking based avatar control on the perceived realism of the VE and on the performance of a surgical teamwork training scenario. Our results show that head tracking positively influences the perceived realism of the VE and the communication, but has no major influence on the training outcome.

Keywords: virtual environment, teamwork training, head tracking, non-verbal communication, head-coupled perspective

First presented at the International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (GRAPP) 2011, extended and revised for JVRB

1 Introduction

In recent years, virtual environments (VEs) have become increasingly popular due to technological advances in graphics and user interfaces [MSL+09]. One of the many valuable uses of VEs is teamwork training. The members of a team can be located wherever it is most convenient for them (e.g., at home) and solve a simulated task in the VE collaboratively, without physically having to travel to a common simulation facility. Medical education and training facilities have realised this advantage and, for example, created numerous medical simulations within Second Life or similar VEs [DPH+09, TSHW08, DYHK07, DWB09].

When looking at teamwork, communication is a vital aspect. An ideal VE would therefore facilitate all communication channels that exist in reality – verbal as well as non-verbal. Due to technical limitations, this is not possible, and therefore, existing communication in VEs is currently mostly limited to voice. Other channels like text chat, avatar body gestures, facial expressions have to be controlled manually and thus do not reflect the real-time communicative behaviour of the user.

Analysis of communication in medical teamwork has shown that nonverbal communication cues like gesture, touch, body position, and gaze are equally important to verbal communication in the analysis of the team interactions [CMBS07]. VEs that do not consider those nonverbal channels are likely to render the communication among the team members less efficient than it would be in reality.

We propose an inexpensive extension of a VE by camera-based head tracking to increase the 'communication bandwidth'.

Head tracking measures the position and the orientation of the user's head relative to the camera and the screen. The *rotational* tracking information can be used to control the head rotation of the user's avatar, as shown in Figure 1. That way, other users in the VE can see rotational head movement identical to the movement actually performed physically by the user, like nodding, shaking, or rolling of the head.

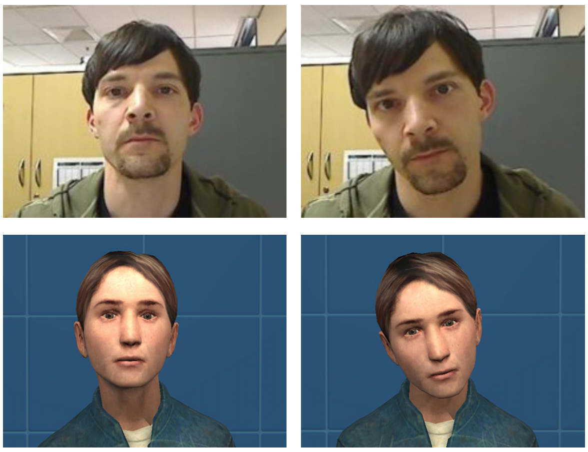

Figure 1: Example screenshots for avatar head movement being controlled by physical rotational movement of the user's head.

The *translational* tracking information can be used to control the view 'into' the VE, as demonstrated in Figure 2. This so called Head Coupled Perspective (HCP) enables intuitive control, like peeking around corners by moving sideways, or zooming in by simply moving closer to the monitor. The use of head tracking information has therefore the potential to simplify the usage of a VE by replacing non-intuitive manual view control by intuitive motion-based view control. Especially in medical applications, this has the potential to free the hands of the user, enabling the use of other simulated instruments or tools, e.g., an endoscope.

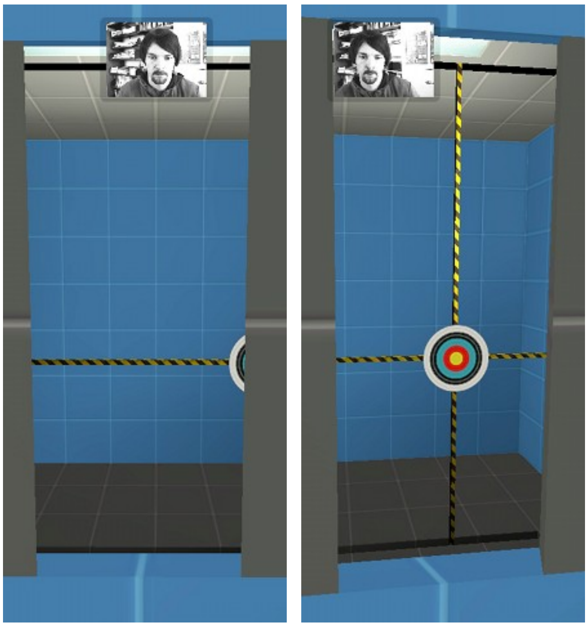

Figure 2: Example screenshots for HCP. The user can control the viewport by translational movement of the head.

This paper presents the results of an experimental study designed to measure any influence that the introduction of head tracking has on teamwork communication, teamwork efficiency, and perceived realism of the VE.

2 Related Work

Medical professionals require a large number of skills ranging from technical skills to perceptual, analytical, communication, and process skills. In order to determine requirements and the optimal role of a VE-based teamwork simulation, we reviewed the fields of medical education, medical simulation, and teamwork training. The full evaluation can be found in [Mar11].

Training of individual technical skills is often performed using simple rubber and plastic models, which provide a realistic look and feel for a reasonably low price. Some critics of medical simulators argue that, embedded within a proper curriculum, this type of training is sufficient to achieve the technical skills training objective [Win07]. Modern virtual reality simulators with haptic feedback (for an overview, see, e.g., [DAAM07]) can offer more complex training scenarios and have the advantage of automated and objective computerised assessment, but suffer from a disproportional higher price tag.

After mastering the technical skills, the medical trainee has to leran about procedural skills, which require a different set of simulations. Complex computer-controlled manikins that simulate various vital signs and respond to medication as well as a multitude of other interventions can be used, e.g., in critical care units, to train emergency response teams in near-lifelike situations (see, e.g., [KAK+10]). Besides the technical and procedural skills, the trainees also have to demonstrate good communication and decision making. A breakdown in these factors is as critical for the well-being of the patient as technical errors [LEW+04]. Therefore, recording, assessment, and feedback on these factors is vital for effective training.

A severe disadvantage of manikins and real-life settings is the high price, making it infeasible for smaller hospitals or rural facilities – specifically in developing countries. Simulation centres that provide a multitude of simulators and training facilities are increasingly set up, but still, the trainees have to be transported to and accommodated near these centres, temporarily removing them from their original workplace.

This is where medical simulation in VEs offers alternatives. Training can happen from the home computer or from a dedicated workstation at the workplace, at convenient times and without the need to travel to a central simulation facility. Simulation scenarios can range from one-to-one conversations to large scenarios like terrorism attacks, natural disasters, or mass accidents (e.g., [TSHW08, DYHK07]).

Due to interaction limitations, training in VEs has to focus on procedural aspects, teamwork, and communication. Since communication includes more aspects than just the verbal channel, it is desirable to reproduce as many non-verbal channels in the VE as possible [Man04].

A great body of research exists about non-verbal communication within VE, indicating, e.g., the importance of eye gaze [VvdVV00, VSvdVN01, GSBS01, GSV+03], head or body pose [GVCP+99], or facial expressions [SHS+00]. In general, an increase in communication channels, especially non-verbal, leads to an increase in effectiveness and/or perceived realism of communication within VEs. However, obtaining the information for these additional channels can be problematic.

In the early stages of research on this topic, data was captured by using magnetic tracking sensors, facial markers, data gloves, stereo cameras, etc. (see, e.g., [OKT+93]). Some of this technology is either too expensive for general use, or too sensitive to changes in the setup, or too obtrusive to the user. With the development of better hardware and advanced tracking algorithms, it became possible to use less obtrusive and less expensive technology.

As an example, the authors of [CPP+02] track the head of the user with a single monocular camera to extract the rotational information. This can easily and efficiently be transmitted during a video call to simulate the head movement of the user's avatar on the receiver's side. However, the focus of this paper is more on information reduction than on virtual environments.

Wang, Xiong, Xu, Wang, Zhang, Dai, and Zhang use the 2-dimensional position of the face within the camera image to control the 2D-movement of a game character [WXX+06]. This idea is extended into the 3rd dimension by the authors of [YQG08], where a head-mounted LED line is used to track the position and rotation of the user's head. This information is used again to control a game instead of an avatar.

Using only a single camera, Sko and Gardner present a range of interaction techniques based on 3-dimensional translation and rotation tracking data [SG09]. A predefined set of head gestures is recognised and associated with certain actions in a game. Slightly tilting the head sideways is used for peering around a corner. Leaning forwards is interpreted as zooming. Head rotation is used for a slight change in the view direction, whereas head translation is used for HCP. These techniques focus on a single user, but the authors have not extended their research on the possibilities for multi-user scenarios.

In a study conducted previously to the one described in this paper, we found that users can clearly perceive head motions, facial expressions, and mouth movement of avatars, and that camera-based HCP is more natural and intuitive to use than manual view control [MWW10].

3 Questions and Hypotheses

In this paper, we investigate the following hypotheses:

H1 When head tracking is enabled, the participants demonstrate more overall head movement than when head tracking is disabled.

H2 When head tracking is enabled, the amount of head movement is greater when participants talk to each other compared to when head tracking is disabled.

H3 When head tracking is enabled, the participants perceive the other avatars as more natural than when head tracking is disabled.

H4 When head tracking is enabled, the participants perceive the communication with the other participants as more natural than when head tracking is disabled.

H5 When head tracking is enabled, the participants look more at each other's avatars than when head tracking is disabled.

H6 When head tracking is enabled, the participants solve the teamwork task better than when head tracking is disabled.

4 Design

The following factors were involved in the design of the study:

- The number of participants was limited to three due to the amount of available game engine licenses.
- It should be impossible or at least infeasible for only one or two participants to solve the scenario efficiently and instead require the effective teamwork of all three users involved.
- To facilitate comparison of the results, the scenario should be of equal length throughout all repetitions.
- To avoid too much of a learning curve effect, the scenario should vary at least in parts.

4.1 Overview

We chose a simplified surgical procedure for the simulation scenario. It involves three participants: a surgeon, an anaesthesiologist, and a nurse. Each of these three roles has very clear and defined tasks:

- The surgeon performs the necessary steps of the surgical procedure on the patient.
- The anaesthesiologist monitors the patient's vital signs and stabilises the patient when critical events occur.
- The nurse is responsible for supplying the surgeon with the correct instrument and the anaesthesiologist with the appropriate medication.

This clear definition of the roles and tasks together with the necessary interactions is depicted in Figure 3. Every role has to communicate with every other role throughout the scenario. Ineffective communication would inevitably lead to ineffective teamwork.

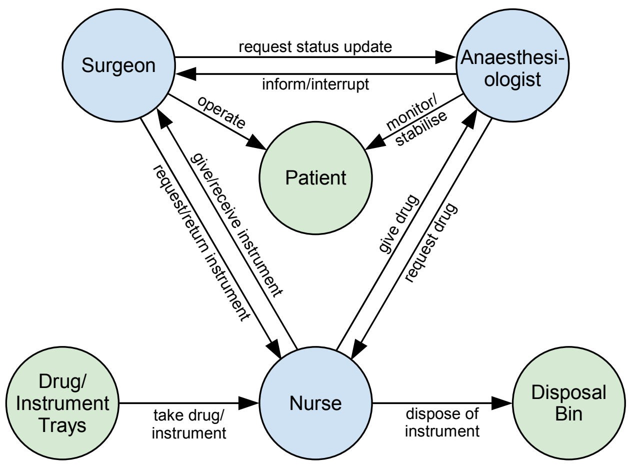

Figure 3: Diagram of the relations and interactions between the three roles and the patient.

For measuring the effectiveness of the teamwork of a group of participants, we used the following two metrics:

- The total time from the beginning of the first step of the surgical procedure until the completion of the last step.
- The relative amount of time that the patient's vital signs are critical (e.g., blood pressure too low, heart rate too high).

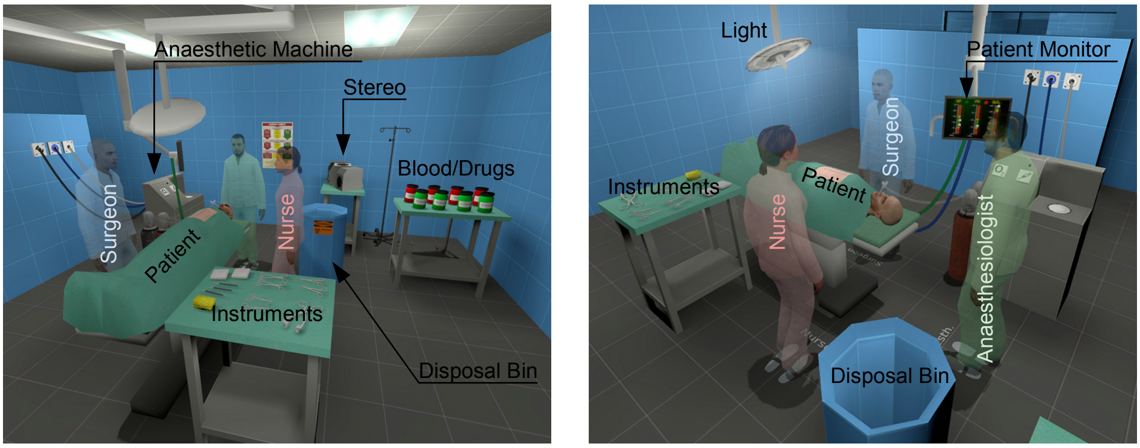

Figure 4: Two views of the operating theatre implemented in our simulation. The characters were made partially transparent in order to allow a better perception of the environment.

For each group of three participants, we conducted six experiments, rotating the roles and disabling/enabling the head tracking so that no participant would have the same role in two successive experiments, and every participant would experience each role with head tracking disabled and enabled.

4.2 Surgical Procedure

The selection of the nature of the simulated surgical procedure was mainly based on the abilities of the game engine we used for creating the VE, a modified version of the Source Engine [Cor07, MWW07].

In general, larger objects can be handled more easily in this VE than small objects. Therefore, we chose the relatively large torso of the patient as the operating field, as opposed to the leg or other extremities. This allowed us to design relatively large instruments that would cause less problems for the participants to move and position.

For the same reason, we chose to design a rather large incision along the middle of the upper body, revealing organs of the digestive system, such as stomach, liver, pancreas, small intestines, and colon.

Located in the centre of this incision are the stomach and the pancreas. After consultation with staff of the surgical department, we decided on the removal of dead and infected tissue from the pancreas of the patient as the surgical procedure to simulate. In medical terms, this procedure is called *pancreatic necrosectomy*.

4.3 Role Description

The positions of the three roles involved in the surgical procedure and the layout of the room can be seen in Figure 4.

4.3.1 Surgeon

The surgeon operates on the right side of the patient. The task of the surgeon is to apply the correct instruments in the right order to the patient to complete the procedure step by step.

The steps of the simulated surgical procedure are visualised in Figure 5. To proceed from one step to the next, the indicated instrument has to be used. The surgeon uses the instrument on the patient by touching the operating field with it. Caused by the operating room light, the instrument will cast a shadow, which assists in positioning the instrument. The instruments are displayed in Figure 6.

When the instrument is the correct one for the task, it starts to blink blue and a progress bar becomes visible in the lower right corner of the screen. Each step takes a certain amount of time, also listed in Figure 5. When the instrument has been applied for the required time, the progress bar disappears, the instrument blinks green once and can now be disposed of by the nurse. At this point, most instruments also change their appearance, e.g., the scalpel, the scissors, the clamp, and the needle are not clean any more but are partially covered in blood.

When the surgeon applies a wrong instrument, it blinks red and has no effect.

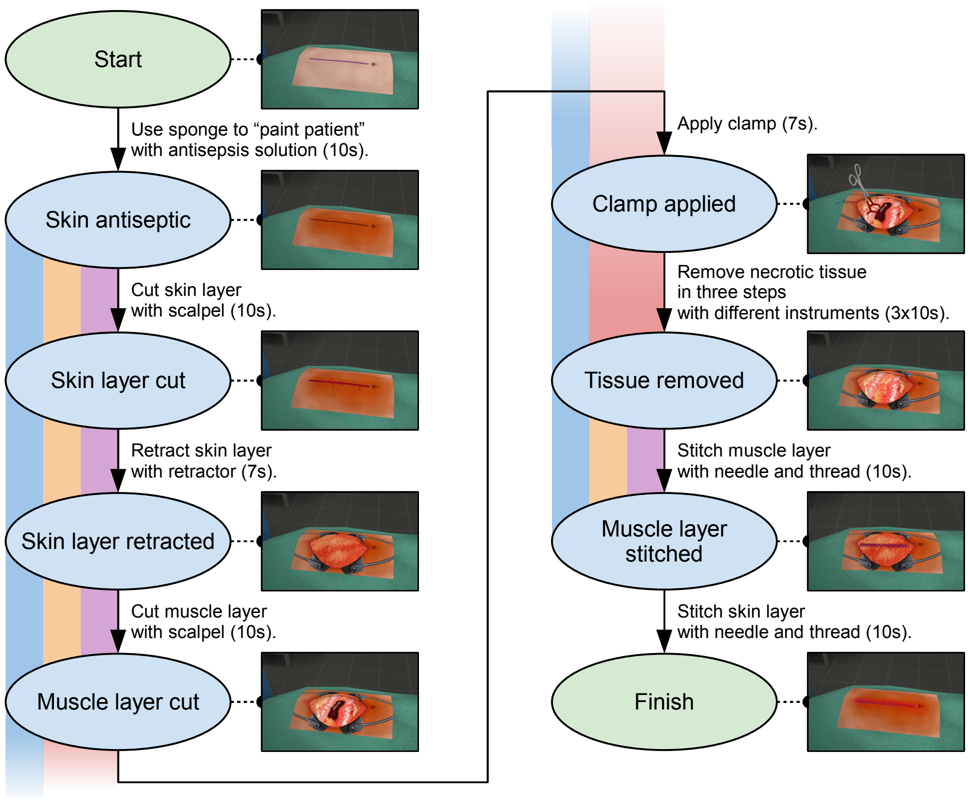

Figure 5: Overview of the steps of the simulated surgical procedure and the instruments necessary to reach the next step. The coloured vertical areas symbolise the sections of the procedure during which the events in Figure 7 can occur.

4.3.2 Anaesthesiologist

The anaesthesiologist is located at the head of the patient, close to the airways. The role of the anaesthesiologist is to observe the patient's vital signs on the patient monitor and to counteract four different types of events that will occur at random times during the operation:

1. The patient's blood oxygen saturation level falls lower than 80%. This will cause a *hypoxaemia* alarm. In this case, the anaesthesiologist has to press the 'O_2' button of the anaesthetic machine to supply additional oxygen and cause the blood oxygen saturation to rise back to normal again.

2. The blood pressure drops due to bleeding during the procedure, resulting in *hypotension*. In this case the anaesthesiologist has to start a blood infusion to stabilise the blood pressure again. The blood pressure will continue to drop as long as the bleeding is not stopped by the surgeon who has to apply pressure by using pads.

3. An increasing heartbeat rate is an indicator of *tachycardia* and has to be treated by administering β-Blockers. The procedure does not have to be stopped during this condition, but as it is the nurse who has to hand the drug to the anaesthesiologist, the surgeon might have to wait for the next instrument.

4. A decreasing heartbeat rate is a sign of *bradycardia* and has to be treated by administering Atropine to the patient. Similar to tachycardia, the procedure does not have to be stopped, but might be delayed by the nurse having to pass the drug to the anaesthesiologist.

4.3.3 Nurse

The nurse is located opposite of the surgeon, on the left side of the patient. The role of the nurse is to:

- switch on the lighting of the operating field,
- hand the requested instrument to the surgeon,
- take back the used instrument from the surgeon and dispose of it, and

- hand the requested drugs to the anaesthesiologist.

4.3.4 Role Enforcement

Physical and logical barriers prevent any of the users from taking over the role of any other user.

- The anaesthesiologist cannot reach the table with the drug and blood bottles, because the stereo (see Section 4.7), a small tray, a stand for intravenous fluids, and the disposal bin are blocking the way. These objects are not physically simulated and can therefore not be moved. To get to the table, it would require walking around the whole room, which in turn would increase the total time needed for the whole scenario. For the same reason, it is infeasible for the nurse and the surgeon to try to reach the position of the anaesthesiologist.

- The surgeon and the anaesthesiologist cannot easily reach the position of the other, because they are blocked by the gas lines to the anaesthetic machine.

- The surgeon can see the instruments and guide the nurse by looking at them, but it is impossible to reach them from the other side of the patient. To take an instrument, it would require to walk around the patient, which would also increase the total time for the whole procedure.

- The simulation script prevents any other role than the surgeon from using instruments on the patient. When other players try to apply an instrument on the patient, it will not have any effect.

- Running and jumping is disabled to prevent participants with a higher experience in first person perspective games to jump over the patient or objects to reach other users' locations.

4.4 Events

During certain phases of the procedure, the events described in Section 4.3.2 will occur. In Figure 5 and Figure 7, the times at which the events can happen are visualised by different colours.

During a certain phase of the operation, especially while the clamp is applied or the necrotic tissue is being cut, unexpected and severe bleeding might occur. This has to be stopped by applying pressure with a pad. The operation can only be continued when the bleeding has stopped. Bleeding is visualised by a

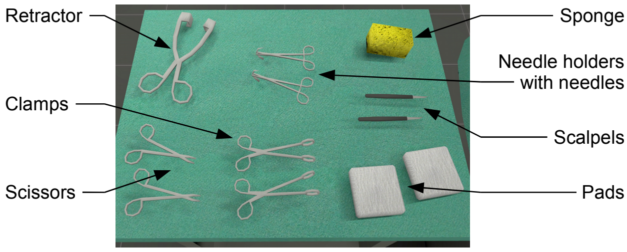

Figure 6: This screenshot shows the instruments for the surgical procedure. It is the task of the nurse to hand over the required instrument to the surgeon who is standing opposite on the other side of the patient.

red particle system simulating a medium intense blood flow inside the operating field. At the same time, the anaesthesiologist will notice a loss of blood pressure. Therefore, bleeding is an event that is easy to recognise by all three roles.

At random times before and after this part, tachycardia and bradycardia will occur. Those events are most apparent to the anaesthesiologist who, in the event of tachycardia, will see the pointers on the patient monitor rise, or, in the event of bradycardia, fall. The surgeon and the nurse will only have the audio cue of the beeping heartbeat monitor becoming faster or slower.

Finally, at another random point in time, the oxygen saturation will drop significantly. The anaesthesiologist has to keep an eye on this value to prevent hypoxaemia. Again, this event is most apparent to the anaesthesiologist, and, from a distance, visible for the nurse.

All four events are spaced temporally so that they do not occur at the same time. Only when the participants do not react to an event and instead continue with the operation, the effects of multiple events will eventually build up. However, this did not occur during the study. All teams reacted promptly to the events.

4.5 Face Validity

The simulated procedure is a very simplified version of a real surgery. We had to find a compromise between face validity, the medical term for realism, and simplicity for several reasons:

- The game engine has limitations in terms of physical simulation and animation.

- The interaction of the user with the VE has to be kept simple.

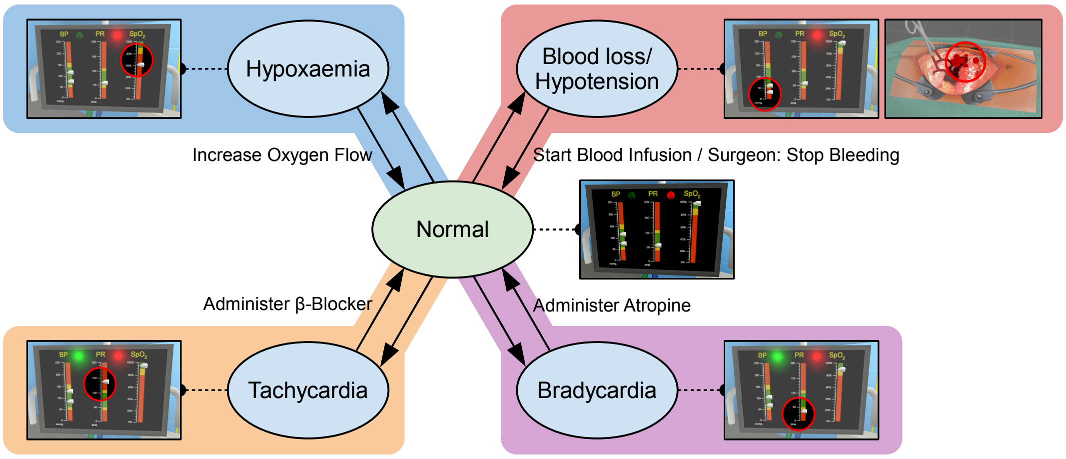

Figure 7: Overview of the four events that the team has to react to during the simulated surgical procedure. The events with a coloured background occur at a random point of time during the steps of the procedure in Figure 5 with the same coloured background. To return to the 'Normal' state, the action next to the arrow has to be taken.

- Most of the participants were non-medical students who did not have a solid background in medicine.

For both participant groups, medical and non-medical, an explanation was given beforehand that this scenario is a very simplified simulation and does not always realistically reflect procedures, devices, and effects.

In summary, the following aspects of the simulation are not realistic:

- The avatars don't wear face masks which, in reality, are mandatory for infection control [LE02].
- A medical team is usually comprised of more than only tree members. In our case, the roles of, e.g., the running nurse (responsible for fetching instruments and gear from outside of the operating theatre) and anaesthesiology nurse (mainly responsible for assistive tasks required by the anaesthesiologist) are not covered.
- The anaesthetic machine and the patient monitor have a very simplified functionality and a different look compared to real machines.
- In case of blood loss between 15% and 30% of the total blood volume, defined as *Class II hemorrhage* by the American College of Surgeons, saline fluid is usually sufficient to stabilise the blood pressure [OTC08, p. 564]. A blood transfusion is only necessary when the patient has lost larger amounts of blood. In addition, due to its viscosity, blood would not flow fast enough into the circulatory system of the patient.
- In case of blood loss, the blood pressure does not immediately decrease. Instead, the heart starts to beat faster to maintain the blood pressure despite the loss of blood volume [OTC08, p. 564]. During this phase, only the diastolic blood pressure starts to drop, caused by the tachycardia and contracting blood vessels, while the systolic pressure stays more or less the same.
- The incision is in general too large for an abdominal procedure. In addition, removal of necrotic pancreas is usually not done by an open operation. A multi-centre study [vSBB$^+$10] has shown that endoscopic procedures result in less complications and a lower mortality rate among patients compared to open procedures.
- In the simulation, the pancreas is immediately visible. In reality, a larger part of the pancreas is hidden behind the stomach.
- The treatment for the critical events in the simulation is extremely simplified.
- Instruments that have been dropped on the floor can be used again on the patient without the risk of infection (see further commentary to this issue in Section 8.4).

- A team coming together for a real world medical simulation already has a reasonable background of medical knowledge and experience. During the simulation, this knowledge and understanding of their roles enables them to completely focus on the teamwork and the task. In contrast, several of the participants in our VE scenario had no medical knowledge and had to learn their roles *as well as* having to work together as a team. For some of them, having to learn the use of the VE was another additional task.

4.6 Room Design

The overall design of the room was based on photos of real operating theatres. 3D modelling of the instruments and other objects in the operating theatre was done in the 3D open-source software Blender [Fou11]. As for the room, we used photos and schematic diagrams for textures and measurements, but enlarged the majority of the instruments by 20 % to 50 % to enable easier handling in the VE.

4.7 Sound

To disrupt voice communication and encourage an increased use of non-verbal communication channels, we added two sound sources to the simulation. Random announcements can be heard in intervals of 20 s to 40 s through a speaker mounted in the ceiling. Those announcements are meant to disrupt the communication between the participants on a regular, but short basis.

A second sound source is a CD player with speakers in the operating theatre, which is common practice in operating theatres [UFS$^+$08]. It randomly plays one of six music tracks that are all about 6 min to 8 min long and differ in style: Rock, Funk, Folk, Jazz, Trip Hop, and Ambient. The anaesthesiologist can change the track, but cannot switch it off or alter the volume.

The music is intended to be a continuous source of noise, making it slightly harder for the participants to fully rely on pure voice communication. Therefore, we purposely violate a general recommendation that, during a surgical procedure, the music should be turned to a low volume in general and completely off in cases of emergency [HT90].

5 Setup

For the VE, we use the Source Engine, running a heavily modified version of a multi-player deathmatch game. The head tracking is done by an external program that uses the commercial tracking engine faceAPI [Mac10]. This program runs in the background and communicates with the VE, continuously transmitting updates of the user's head position and rotation.

The delay between actual head movement and the change of the view or the avatar's head direction is approximately 150 ms. None of the participants reported that this delay interefered with their tasks. However, research suggests that a large latency significantly slows down any tasks performed within an VE [MAEH04]. For our study, the delay was identical on all machines, and minimal bewteen server and clients due to a local network connection. Therefore, we did not look at any effects of latency in our study. In case of a global simulation scenario however, additional delays between the clients and the server would have to be taken into account.

We used three DELL Optiplex 745 desktop computers[1] as clients and the development machine[2] as the server.

The simulation was run with a 60 Hz framerate on the server as well as on the clients.

The client desktops were equipped with a Logitech QuickCam Pro 9000 webcam. This camera provided a stable framerate of 30 fps with only the available room lighting. Autofocus ensured that even with movement away and towards to the monitor, the heads of the participants always stayed in focus.

It was not possible to organise three separate rooms for the participants, so we had to use other methods to ensure that they could not hear or see each other directly. One solution was the use of Logitech G35 headsets which covered the ears of the participants completely and provided sufficient isolation from outside sounds. That way, the participants had to rely solely on communication channels provided by the game engine.

To prevent the participants from seeing each other

[1] Intel Core™2 Duo CPU, 2.13 GHz, 3 G memory, Windows Vista Enterprise 32, Radeon X1300 graphics card with 256 M graphics memory, DELL 2007FP 20.1 inch flat panel monitor with 1600 × 1200 pixel resolution

[2] Intel Core™2 Quad CPU, 2.4 GHz, 4 G memory, Windows XP Professional, NVIDIA Quadro FX 570 graphics card with 256 M graphics memory

or each other's screens, we set up office walls that occluded direct line of sight between the three client workstations and in addition provided further acoustic separation of the participants. A photo of the complete setup is shown in Figure 8.

The control of the user's avatars within the VE was limited to controlling the view direction with the mouse, controlling walking movement of the avatar with the keyboard keys W, A, S, D, and clicking the mouse to press buttons or to grab and release objects. In addition, when head tracking was active, translational head movement of the user would change the view of the VE, whereas rotational head movement would rotate the head of the avatar, but without changing the user's view direction on the screen.

6 Methodology

We recruited participants by printed advertisements in the buildings of the departments of Computer Science and the Medical School, and by emails sent to classes and postgraduate students of the department of Computer Science.

More than 30 people answered to the invitation emails and advertisements. However, due to financial and time constraints, we had to limit the amount of participants to 30. The age of the participants, 7 of them being female, covered a range from 20 to around 44 years (see Figure 9 for more details). The gender distribution within the teams formed by the participants ranged from purely male to up to two female participants per team.

At the time of the study, all participants had used a computer often or more. 8 of them had not played computer games at all within the three months before the study. 25 of the participants have had experience in using a webcam, e.g., by using it for video conferencing.

On average, one user study took around 90 min to complete. Due to this extended duration, we reimbursed the participants by handing out grocery vouchers valued 10 NZD. We assumed that this amount would encourage participants sufficiently without introducing the risk that they would 'just do it for the money'.

We explained the overall goal of the study to the participants, mentioning the head tracking and the possibility to control the view and the head movement of the avatar. Afterwards, we gave a brief overview of the procedure and the roles and their responsibilities. The participants were encouraged to interrupt and ask questions at any stage. The user study was approved by The University of Auckland Ethics Committee on 9.9.2010 for 3 years, Reference Number 2010/422.

Before the actual study took place, we introduced the participants to the VE by running an introduction simulation without recording any data. The participants were able to get accustomed to the environment and to experiment with objects, instruments, and devices. Each participant had a set of three information sheets in front of them that described their role, the necessary instruments or devices, and the mouse and keyboard controls.

After the introduction and any further questions from the participants, we enabled data logging of the simulation and started the actual series of six experiments. The participants were *not* informed whether head tracking was disabled or enabled.

After each experiment, the participants filled in a short questionnaire about their opinion on teamwork and communication aspects of the simulation. During that time, we changed the information sheets according to the rotating roles of the participants.

At the end of the six experiments, the participants filled out an additional questionnaire page with general questions about the VE. The complete questionnaire can be found in [Mar11].

7 Results

7.1 Hypotheses

Hypothesis H1: *When head tracking is enabled, the participants demonstrate more overall head movement than when head tracking is disabled.*

For the verification of this hypothesis, we calculated the total amount of translational and rotational head movement in 1 s intervals to even out tracking jitter. The values from experiments with head tracking being disabled were compared to the values from experiments with head tracking being enabled.

The resulting data is not normally distributed, but instead strongly skewed towards lower values and tailing off towards higher values. For that reason, it is not possible to apply the t-test to this data. Instead, we chose the Wilcoxon rank sum test [MG08].

The results show a statistically significant but only minute difference in translational movement (median values 28.2 mm / 26.7 mm, Wilcoxon rank sum test with continuity correction, $p < 0.001$, 95 % CI

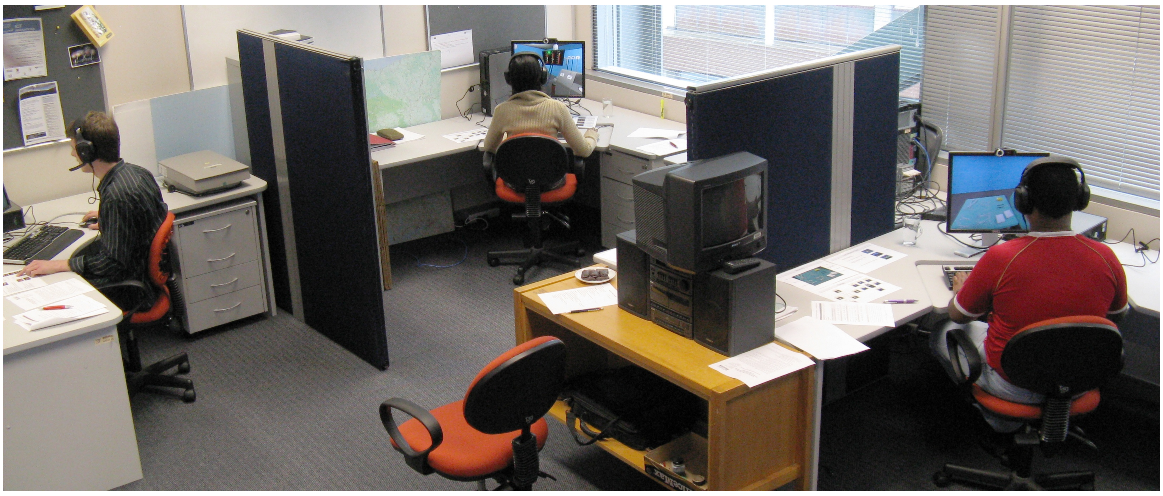

Figure 8: Photo of the experimental setup with three participants. The participants are isolated visually by office walls and acoustically by headsets that completely cover the ears.

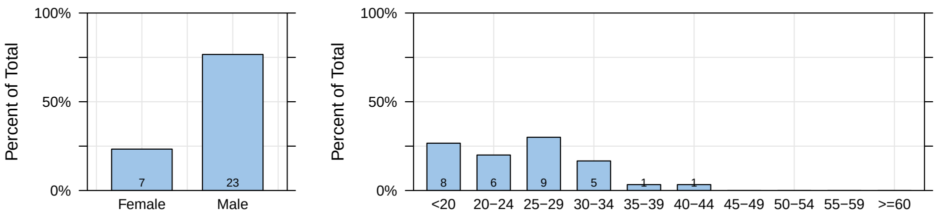

(a) Gender (b) Age Group

Figure 9: Demographic data of the participants of the multi user study.

0.71 mm to 1.51 mm) and a minute and statistically insignificant difference in rotational movement (median values 11.2° / 11.1°, Wilcoxon rank sum test with continuity correction, $p = 0.500$, 95 % CI −0.07° to 0.16°). Based on this outcome, Hypothesis H1 has to be falsified. The amount of head movement is *the same* irrespectively of head tracking being disabled or enabled.

Hypothesis H2: *When head tracking is enabled, the amount of head movement is greater while participants talk to each other compared to when head tracking is disabled.*

For the analysis of this hypothesis, we prepared the data in the same way as described for the hypothesis before, but selected only the subset of the measurements when the participants were talking. Again, we had to apply the Wilcoxon ranked sum test due to the skewed distribution of the data.

Like for the previous hypothesis, the results of the test for head movement of the participants when they were talking do either indicate statistically significant but minor differences, or statistically insignificant differences. Translational head movement differs only by −1.0 mm (median values 35.9 mm / 34.9 mm, Wilcoxon rank sum test with continuity correction, $p = 0.054$, 95 % CI 0.00 mm to 1.53 mm), and rotational movement by −0.2° (median values 14.0° / 13.8°, Wilcoxon rank sum test with continuity correction, $p = 0.183$, 95 % CI −0.06° to 0.42°).

Similar to the previous one, Hypothesis H2 has to be falsified. While talking, the amount of movement is *the same*, no matter whether head tracking is disabled or enabled.

Hypothesis H3: *When head tracking is enabled, the participants perceive the other avatars as more natural than when head tracking is disabled.*

For the verification of this hypothesis, we evaluated the questionnaires that the participants had to answer after each of the six experiments (see Figure 10).

We converted the answers to the 7-point Likert scales on the questionnaires to a numerical value using a translation scale of −3 for ‘Strongly Disagree’ to +3

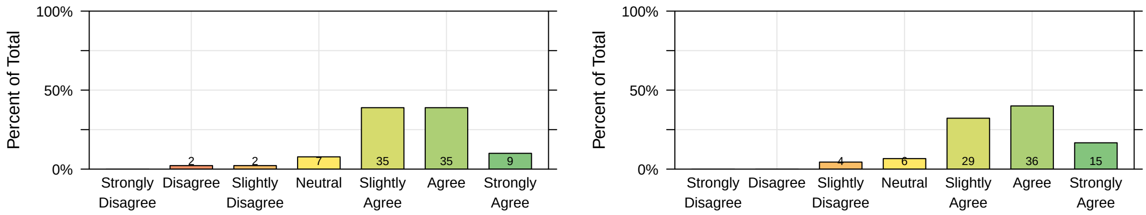

(a) Head Tracking disabled (b) Head Tracking enabled

Figure 10: Results of the questionnaires about how natural the participants perceive the teammember's avatars.

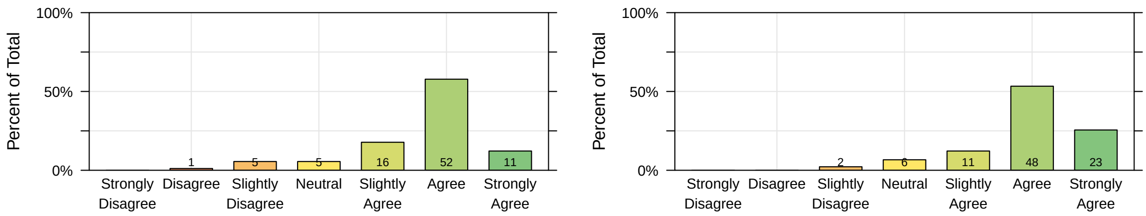

(a) Head Tracking disabled (b) Head Tracking enabled

Figure 11: Results of the questionnaires about how natural the participants perceive the communication with their team members.

for 'Strongly agree'. We then compared the answers of experiments 1, 3, and 5 (head tracking disabled) to the answers of experiments 2, 4, and 6 (head tracking enabled) using a t-test.

The comparison reveals a statistically non-significant difference of 0.2 (mean values 1.4 / 1.6, Welch Two Sample t-test, $p = 0.282$, 95 % CI -0.45 to 0.13), thereby weakly supporting the validity of the hypothesis. The participants perceive the other avatars slightly more natural when head tracking is enabled compared to when head tracking is disabled.

Hypothesis H4: *When head tracking is enabled, the participants perceive the communication with the other participants as more natural than when head tracking is disabled.*

We evaluated the answers to the question about how natural the communication is perceived (see Figure 11) in a similar way to the one described for the previous hypothesis. Again, we compared the results of experiments 1, 3, and 5 (head tracking disabled) to the answers of experiments 2, 4, and 6 (head tracking enabled) using a t-Test.

The comparison suggests that the participants perceive the communication slightly more natural when head tracking is enabled than when it is disabled (mean values 1.7 / 1.9, Welch Two Sample t-test, $p = 0.059$, 95 % CI -0.55 to 0.01), thereby supporting the hypothesis.

Hypothesis H5: *When head tracking is enabled, the participants look more at each other's avatars than when head tracking is disabled.*

The logged data from the simulation also recorded which user has been looking at which object or avatar and for how long. We used this information to calculate the relative amount of time each participant had been looking at one of the team members. The results, grouped by role, are visualised in Figure 12.

Though there are differences in the comparisons, those are again only minor and statistically insignificant. When head tracking is enabled, the surgeon spends $0.4\,\%$ more time looking at the colleagues (mean values $25.0\,\%$ / $25.4\,\%$, Welch Two Sample t-test, $p = 0.849$, 95 % CI $-5.15\,\%$ to $4.25\,\%$), the anaesthesiologist $1.6\,\%$ (mean values $11.7\,\%$ / $13.3\,\%$, Welch Two Sample t-test, $p = 0.430$, 95 % CI $-5.71\,\%$ to $2.47\,\%$), and the nurse $1.4\,\%$ (mean values $23.9\,\%$ / $25.3\,\%$, Welch Two Sample t-test, $p = 0.576$, 95 % CI $-6.05\,\%$ to $3.40\,\%$) compared to when head tracking is disabled.

Overall, these results add only very weak support for the hypothesis. Nevertheless, it is interesting to see how the values reflect the different observation patterns of the three roles. It is apparent that the anaesthesiologist has a different pattern than surgeon and nurse. This difference is most probably caused by the primary task of the role of the anaesthesiologist: to watch the patient monitor for changes in the vital signs.

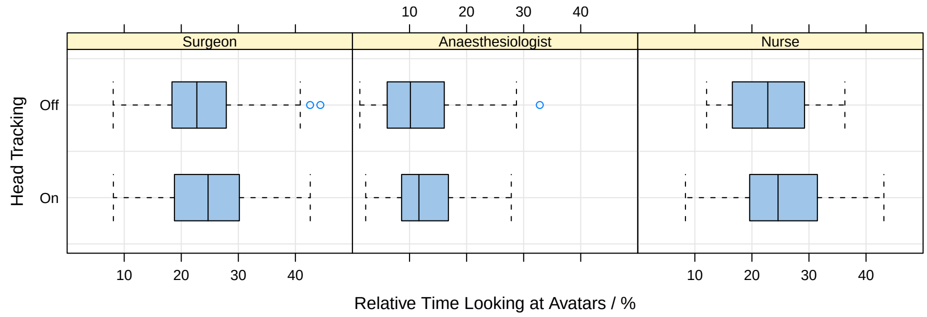

Figure 12: Plots of the relative amount of time the participants have been looking at the avatars of their team members.

Hypothesis H6: *When head tracking is enabled, the participants solve the teamwork task better than when head tracking is disabled.*

As described in Section 4.1, we calculated the efficiency of the teamwork based on two values that are extracted from the data:

- The total time for completion of each experiment.
- The relative amount of time that the patient is in a critical state.

For both values, we expected the influence of the head tracking to reflect on the results as shown in Figure 13. The participants were repeating a similar task six times, so we anticipated a learning curve that would demonstrate a tendency towards faster completion with each repetition, with a gradient that would ease out towards the end. The positive influence of head tracking would be visible by average values that are *below* the fitting curve for experiments 2, 4, and 6, whereas experiments 1, 3, and 5 without head tracking would produce results *above* the fitting line.

However, the measured results look very different. The learning curve is visible in experiments 2 to 6, but experiment 1 has a significantly lower completion time than experiments 2 to 3. The reason for this outlier is most probably the practice run that the participants had to complete before the first experiment was conducted.

For the practice run, they performed the teamwork task without their results being recorded and with help and minor intervention from the simulation administrator. After the practice run, the first experiment started with the participants having the same roles than in the practice run. They had gotten used to their roles and were able to finish the teamwork task relatively fast.

For experiment 2 and 3, the roles were rotated so that every participants had to learn the requirements and tasks of the new roles. The influence of this new learning process reflects in the similar results for experiments 2 and 3.

Finally, for experiments 4 to 6, the participants were repeating former role combinations. Because by now, they had experienced every role, they were better able to focus on the task.

Due to the unusual values of the boxplots results, it is not possible to model an average curve for the comparison of the results with and without head tracking. Instead, we grouped the values by head tracking ‘off’ (experiments 1, 3, and 5) and ‘on’ (experiments 2, 4, and 6) and compared the mean values using a t-test. The values do not differ significantly (mean values 285.2 s / 285.9 s, Welch Two Sample t-test, $p = 0.974$, 95 % CI -45.21 s to 43.74 s). Instead, the high p-value suggests that head tracking makes *no difference at all* on the completion times of the teamwork task.

Likewise, the comparison of the relative amount of time that the master alarm was on results in a statistically insignificant difference (mean values 16.7 % / 17.3 %, Welch Two Sample t-test, $p = 0.821$, 95 % CI -6.68 % to 5.33 %), suggesting additional support for the hypothesis that head tracking makes *no difference at all*.

8 Discussion

Overall, the participants enjoyed to work in teams within the simulation and had no major problems with the head tracking. In fact, the tracking was so unobtrusive, that several participants later commented that they were not at all aware of the tracking and the head movements.

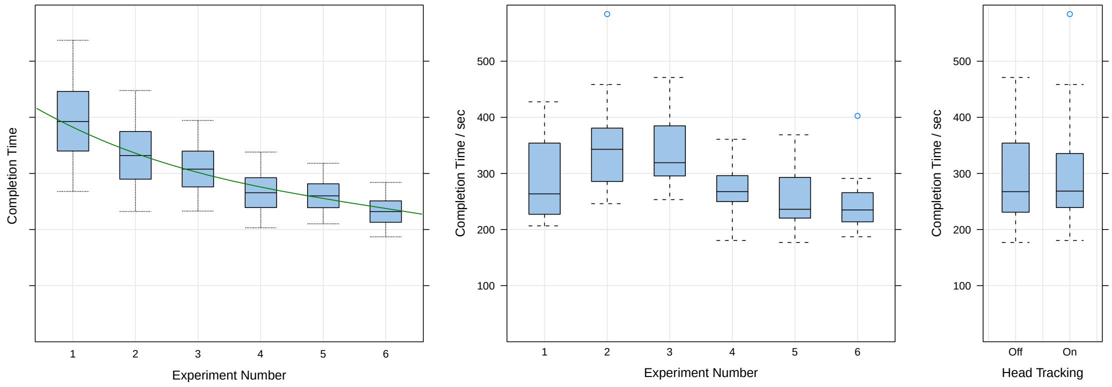

(a) Expected results

(b) Measured results

(c) Grouped by head tracking off/on

Figure 13: Box plots of the times the participants needed to complete the task. The left graph visualises a rough estimate of the expected results, demonstrating a typical learning curve (green) with completion times being slightly faster when head tracking is enabled and slightly slower when head tracking is disabled. The centre graph visualises the actual results, demonstrating a major difference to the expected results.

8.1 Head Tracking

There were slight problems resulting from a combination of head tracking jitter and handling of small objects, such as the scalpel. In these cases, some participants had problems passing on the scalpel from one to another. The tracking jitter of one simulation client would cause the object to move randomly, making it more difficult for the opposite user to grab it. In their comments, several participants agreed on the fact that small objects were most difficult to handle, especially the scalpel.

On the other hand, head tracking was mentioned as being very useful in precise placement of instruments or for zooming in on the patient monitor.

In general, the rotation of the avatar's heads by head tracking was perceived, although most participants stated that they were too preoccupied by the task to pay attention to the head movements. Occasions where the tracking would stand out were when the participants would physically look sideways to their role description sheets and their avatar would also perform that sideways look. Other occasions were, for example, head bobbing to the rhythm of the music.

8.2 Involvement

The participants demonstrated a high level of involvement, characterised by utterances such as 'Oh No!', when the master alarm would go off, or 'Don't die, please!' when the bleeding could not be stopped immediately. Sometimes, participants would even engage in role play and joke around or tease their team colleagues.

On some occasions, when teamwork was not optimal, and the group was especially competitive, the tone turned slightly aggressive. In one group, for example, the nurse ignored the request of the surgeon for the pad and instead decided to give a bottle of blood to the anaesthesiologist first. In terms of efficiency, this decision was not optimal, because as long as the bleeding was not stopped, the patient continued to lose blood. In this case, the first priority of the team should have been to stop the bleeding (i.e., give the surgeon a pad) and then to stabilise the blood pressure. The surgeon, being aware of this ineffectiveness, commented in a slightly annoyed tone: 'We need another nurse!', The nurse defiantly countered: 'I think we need another surgeon!'.

8.3 Communication

Every group quickly established a certain communication style that was kept throughout the six experiments. Those styles ranged from being short and concise, i.e.,

- Surgeon: 'Scalpel!'
- Nurse: 'There!'
- Surgeon: 'Got it!'

to being extended and descriptive, i.e.,

- Surgeon: 'Ok, next, I need a scalpel. Nurse, could you please give me the scalpel?'

- Nurse: 'Scalpel. Ok. Here it is.'
- Surgeon: 'Ok. Wait a second. Ok. Got it!'

Because of the difficulties in passing on especially small objects, verbal feedback for having taken an object (e.g., 'Ok. Got it!') from another avatar was always present.

Some groups also introduced an additional feedback layer when certain commands were given. By repeating the command or the name of the required object, they avoided misunderstandings as early as possible:

- Surgeon: 'I need the scalpel next.'
- Nurse: 'Ok. I'll get the scissors.'
- Surgeon: 'No. The scalpel. I need a scalpel.'
- Nurse: 'Sorry. I'll get the scalpel.'
- Surgeon: 'Yes, that's it.'

Groups that did not introduce such a communication protocol, seemed to be more prone to performing unnecessary actions or passing on the wrong object.

In addition to verbal feedback, some participants also used object interaction for the feedback:

- Surgeon: 'Can you please give me the retractor.'
- Nurse: Picks up the scissors. 'Is this the retractor?'
- Surgeon: 'No. That's the scissors. The retractor is the big one in the middle.'
- Nurse: Drops the scissors and picks up the retractor. 'This one?'
- Surgeon: 'Yes.'

8.4 Unexpected Behaviours

A disadvantage of the rotation of roles was that on several occasions, participants would take over responsibilities of other roles. Several nurses who had been the anaesthesiologist in the previous run would keep an eye on the patient monitor and take over parts or nearly all of the anaesthesiologist's responsibilities. Actions ranged from fetching the necessary drug in advance, to warning or even commanding the anaesthesiologist.

One participant commented:

> "[. . .] I am able at this point to predict what I need to do and watch the monitor so I can somewhat circumvent the communication process to a degree by already having in hand what I would otherwise be asked to get [. . .] . This is particularly the case with respect to the anaesthesiologist."

Another unexpected behaviour, observed especially with participants who indicated a high amount of computer game playing time, was the exploration of similarities and differences between the simulation and reality. When participants had, for example, dropped the scalpel on the floor, there was an initial hesitancy to use that scalpel again on the patient, caused by basic knowledge of the risk of infection in reality. However, when the participants found out that, in the simulation, this re-use did not have any negative consequences, some of them continued to use dropped instruments on the patient in favour to demanding a new and 'clean' instrument from the nurse and thereby losing time.

8.5 Miscellaneous

When being asked about the stressfulness of the roles, the majority of the participants agreed that nurse was the most stressful role, in contrast to anaesthesiologist, who was considered the most relaxed role.

One of the participant groups consisted of medical students, compared to mainly computer science students forming the other groups. The medical students did not appear to have more or less difficulties than other groups. However, in their feedback they picked out details that other participants did not notice, e.g., that the avatars were not wearing facial masks. This single fact was most distracting for one of the members of the group. They also provided more feedback concerning the face validity of the medical side of the setup as described in Section 4.5.

Music was perceived as relaxing and soothing, but at the same time loud enough to slightly impact the verbal communication, which was exactly what we had intended during the design.

Voice communication via the headsets worked very well. The headsets caused no discomfort for the participants and had no negative influence on face tracking. On a few occasions, participants would accidentally touch the mute button on the headset while they were adjusting the volume, but this mishap was usually detected very early into the simulation.

8.6 Related research

The results of this study can be positioned alongside other virtual environment training scenarios such as 3DiTeams, a military hospital simulation [TSHW08], or SimTech, a medical virtual teamwork training environment [DYHK07]. Similar to the virtual environment presented in this paper, communication in these two environments is possible using microphones and headsets. Text based chat is a second possible channel for communication. However, the avatars have no specific head movements and facial expressions that are controlled in realtime by the user. At most, the user can manually trigger a set of pre-animated avatar gestures like waving or head nodding and a limited set of facial expressions.

Since the two studies have an educational focus and look closely at the training outcome compared to traditional training methods like manikins, the shortcomings in non-verbal communication support are not mentioned or looked at in more detail. In contrast, our study analyses the impact of head-tracking based avatar control on the training outcome and the user experience itself. Besides improving the usability and the perceived realism of the simulation for the user, the additional data obtained by head tracking and event logging in the simulation could potentially improve teamwork performance evaluation by providing additional information to a human assessor or a computer-based assessment tool.

9 Conclusion

In summary, the evaluation of experimentally obtained data is showing

- that avatars with tracking-based head movement are perceived as more natural than avatars without tracking-based head movement, and
- that avatars with tracking-based head movement facilitate a communication that is perceived as more natural than with avatars without tracking-based head movement,

However, these results are either only weakly supported by the data or the difference in the participants' perception is only marginal.

Unfortunately, there is only very weak support for the hypothesis that participants look at each other's avatars more regularly when head tracking is enabled.

Finally, contrary to the original hypotheses, the analysis of the data also shows

- that the participants do *not* change the amount of head movement when head tracking is enabled compared to when head tracking is disabled,
- that, while talking, the participants do *not* change the amount of head movement when head tracking is enabled compared to when head tracking is disabled, and
- that teamwork performance does *not* improve when head tracking is enabled compared to when head tracking is disabled.

Those last three statements indicate an important fact: Although a new technology has been introduced with several possibilities for failure and measurement errors, the participants were largely unaware of that new technology. They perceived an improvement in the naturalness of the VE and an improvement of the quality of communication, but did not experience any major negative limitations imposed upon them by this new technology.

We assume that the main reason for the lack of influence of head tracking was the fact that the simulated scenario required more focus on the task and on handling the instruments than on the communication.

In addition, not all possibilities of this new technology have been explored yet. For our experiment, the head rotation was simply mapped onto the avatars without any further semantic analysis. One future extension will be to calculate the actual target of the user's head direction on the screen and to use that information to transform the tracking data so that the user's avatar looks at the same target in the VE. An additional level of realism would be introduced by eye tracking and a similar transformation of the avatar's gaze.

Automatic recognition of facial expressions is another future extension. This feature would not increase the complexity of the physical setup and would also not require any major changes in our simulation engine. However, we would expected that the inclusion of this communication channel can further increase at least the perceived naturalness and ease of use – if not more.

References

[CMBS07] John Cartmill, Alison Moore, David Butt, and Lyn Squire, *Surgical Teamwork: Systemic Functional Linguistics and the Analysis of Verbal and Nonverbal Meaning in Surgery*, ANZ Journal of Surgery **77** (2007), no. Suppl 1, A79–A79, ISSN 1445-2197.

[Cor07] Valve Corporation, *Source engine*, 2007, source.valvesoftware.com.

[CPP+02] Marius D. Cordea, Dorina C. Petriu, Emil M. Petriu, Nicolas D. Georganas, and Thomas E. Whalen, *3-D Head Pose Recovery for Interactive Virtual Reality Avatars*, IEEE Transactions on Instrumentation and Measurement **51** (2002), no. 4, 640–644, ISSN 0018-9456.

[DAAM07] B. Dunkin, G. L. Adrales, K. Apelgren, and J. D. Mellinger, *Surgical simulation: a current review*, Surgical Endoscopy **21** (2007), no. 3, 357–366, ISSN 0930-2794.

[DPH+09] Douglas Danforth, Mike Procter, Robert Heller, Chen Richard, and Mary Johnson, *Development of Virtual Patient Simulations for Medical Education*, Journal of Virtual Worlds Research **2** (2009), no. 2, 3–11, ISSN 1941-8477.

[DWB09] Scott Diener, John Windsor, and David Bodily, *Design and Development of Medical Simulations in Second Life and OpenSim*, EDUCAUSE Australasia Conference (Perth, Australia), 2009.

[DYHK07] Parvati Dev, Patricia Youngblood, W. LeRoy Heinrichs, and Laura Kusumoto, *Virtual Worlds and Team Training*, Anesthesiology Clinics **25** (2007), no. 2, 321–336, ISSN 1932-2275.

[Fou11] Blender Foundation, *Blender*, 2011, www.blender.org.

[GSBS01] Maia Garau, Mel Slater, Simon Bee, and Martina Angela Sasse, *The impact of eye gaze on communication using humanoid avatars*, CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems, 2001, pp. 309–316, ISBN 1-58113-327-8.

[GSV+03] Maia Garau, Mel Slater, Vinoba Vinayagamoorthy, Andrea Brogni, Anthony Steed, and M. Angela Sasse, *The impact of avatar realism and eye gaze control on perceived quality of communication in a shared immersive virtual environment*, Proceedings of the SIGCHI conference on Human factors in computing systems, 2003, pp. 529–536, ISBN 1-58113-630-7.

[GVCP+99] A. Guye-Vuilléme, T. K. Capin, S. Pandzic, N. Magnenat Thalmann, and D. Thalmann, *Nonverbal communication interface for collaborative virtual environments*, Virtual Reality **4** (1999), no. 1, 49–59, ISSN 1359-4338.

[HT90] B. Hodge and J. F. Thompson, *Noise pollution in the operating theatre*, The Lancet **335** (1990), no. 8694, 891–894, ISSN 0140-6736.

[KAK+10] Roger Kneebone, Sonal Arora, Dominic King, Fernando Bello, Nick Sevdalis, Eva Kassab, Raj Aggarwal, Ara Darzi, and Debra Nestel, *Distributed simulation - Accessible immersive training*, Medical Teacher **32** (2010), no. 1, 65–70, ISSN 0142-159X.

[LE02] Allyson Lipp and Peggy Edwards, *Disposable surgical face masks for preventing surgical wound infection in clean surgery*, Cochrane Database of Systematic Reviews **1** (2002), no. CD002929, ISSN 1469-493X.

[LEW+04] Lorelei Lingard, Sherry Espin, S. Whyte, Glenn Regehr, G. R. Baker, Richard Reznick, J. Bohnen, B. Orser, D. Doran, and E. Grober, *Communication failures in the operating room: An observational classification of recurrent types and effects*, Quality & Safety in Health

Care **13** (2004), no. 5, 330–334, ISSN 2044-5423.

[Mac10] Seeing Machines, *faceapi*, 2010, http://www.seeingmachines.com/product/faceapi.

[MAEH04] Katerina Mania, Bernard D. Adelstein, Stephen R. Ellis, and Michael I. Hill, *Perceptual sensitivity to head tracking latency in virtual environments with varying degrees of scene complexity*, Proceedings of the 1st Symposium on Applied Perception in Graphics and Visualization, APGV '04, 2004, pp. 39–47, ISBN 1-58113-914-4.

[Man04] Tony Manninen, *Rich Interaction Model for Game and Virtual Environment Design*, Ph.D. thesis, University of Oulu, Finland, 2004.

[Mar11] Stefan Marks, *A Virtual Environment for Medical Teamwork Training with Support for Non-Verbal Communication using Consumer-Level Hardware and Software*, Ph.D. thesis, The University of Auckland, 2011.

[MG08] Evie McCrum-Gardner, *Which is the correct statistical test to use?*, British Journal of Oral and Maxillofacial Surgery **46** (2008), no. 1, 38–41, ISSN 0266-4356.

[MSL+09] Paul R. Messinger, Eleni Stroulia, Kelly Lyons, Michael Bone, Run H. Niu, Kristen Smirnov, and Stephen Perelgut, *Virtual Worlds - Past, Present, and Future: New Directions in Social Computing*, Decision Support Systems **47** (2009), no. 3, 204–228, ISSN 0167-9236.

[MWW07] Stefan Marks, John Windsor, and Burkhard Wünsche, *Evaluation of Game Engines for Simulated Surgical Training*, Graphite '07: Proceedings of the 5th international conference on Computer graphics and interactive techniques in Australia and Southeast Asia, 2007, pp. 273–280, ISBN 978-1-59593-912-8.

[MWW10] Stefan Marks, John Windsor, and Burkhard Wünsche, *Evaluation of the Effectiveness of Head Tracking for View and Avatar Control in Virtual Environments*, 25th International Conference Image and Vision Computing New Zealand (IVCNZ) 2010 (2010).

[OKT+93] J. Ohya, Y. Kitamura, H. Takemura, F. Kishino, and N. Terashima, *Real-time reproduction of 3D human images in virtual space teleconferencing*, Virtual Reality Annual International Symposium, 1993, pp. 408–414, ISBN 0-7803-1363-1.

[OTC08] J. Patrick O'Leary, Arnold Tabuenca, and Lea Rhea Capote (eds.), *The physiologic basis of surgery*, 4 ed., Wolters Kluwer Health/Lippincott Williams & Wilkins, 2008, ISBN 9780781771382.

[SG09] Torben Sko and Henry J. Gardner, *Head Tracking in First-Person Games: Interaction Using a Web-Camera*, Human-computer interaction - INTERACT 2009, 2009, Lecture Notes in Computer Science, pp. 342–355, ISBN 978-3-642-03654-5.

[SHS+00] M. Slater, J. Howell, A. Steed, D.-P. Pertaub, and M. Gaurau, *Acting in Virtual Reality*, Proceedings of the Third International Conference on Collaborative Virtual Environments, 2000, pp. 103–110, ISBN 1-58113-303-0.

[TSHW08] Jeffrey Taekman, Noa Segall, Eugene Hobbs, and Melanie Wright, *3DiTeams - Healthcare Team Training in a Virtual Environment*, Simulation in Healthcare: The Journal of the Society for Simulation in Healthcare **3** (2008), no. 5, 112, ISSN 1559-2332.

[UFS+08] Yehuda Ullmann, Lucian Fodor, Irena Schwarzberg, Nurit Carmi, Amos Ullmann, and Yitzchak Ramon, *The sounds of music in the operating room*, Injury **39** (2008), no. 5, 592–597, ISSN 0020-1383.

[vSBB+10] H. C. van Santvoort, M. G. Besselink, O. J. Bakker, H. S. Hofker, M. A. Boermeester, C. H. Dejong, H. van Goor, A. F. Schaapherder, C. H. van Eijck, T. L. Bollen, B. van Ramshorst, V. B. Nieuwenhuijs, R. Timmer, J. S. Laméris, P. M. Kruyt, E. R. Manusama, E. van der Harst, G. P. van der Schelling, T. Karsten, E. J. Hesselink, C. J.van Laarhoven, C. Rosman, K. Bosscha, R. J. de Wit, A. P. Houdijk, M. S. van Leeuwen, E. Buskens, H. G. Gooszen, and Dutch Pancreatitis Study Group, *A Step-up Approach or Open Necrosectomy for Necrotizing Pancreatitis*, The New England Journal of Medicine **362** (2010), no. 16, 1491–1502, ISSN 0028-4793.

[VSvdVN01] Roel Vertegaal, Robert Slagter, Gerrit van der Veer, and Anton Nijholt, *Eye gaze patterns in conversations: there is more to conversational agents than meets the eyes*, CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems (New York, NY, USA), 2001, pp. 301–308, ISBN 1-58113-327-8.

[VvdVV00] Roel Vertegaal, Gerrit van der Veer, and Harro Vons, *Effects of Gaze on Multiparty Mediated Communication*, Graphics Interface, 2000, pp. 95–102, ISBN 0-9695338-9-6.

[Win07] John A. Windsor, *Whither, not wither, for Skills Training*, ANZ Journal of Surgery **77** (2007), no. 10, 811–811, ISSN 1445-2197.

[WXX+06] Shuo Wang, Xiaocao Xiong, Yan Xu, Chao Wang, Weiwei Zhang, Xiaofeng Dai, and Dongmei Zhang, *Face-tracking as an augmented input in video games: enhancing presence, role-playing and control*, CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems, 2006, pp. 1097–1106, ISBN 1-59593-372-7.

[YQG08] Jeffrey Yim, Eric Qiu, and T. C. Nicholas Graham, *Experience in the design and development of a game based on head-trackinginput*, Future Play '08: Proceedings of the 2008 Conference on Future Play, 2008, pp. 236–239, ISBN 978-1-60558-218-4.

Citation
Stefan Marks, John Windsor, Burkhard Wünsche
Head Tracking Based Avatar Control for Virtual Environment Teamwork Training, Journal of Virtual Reality and Broadcasting, 9(2012), no. 9, December 2012, urn:nbn:de:0009-6-35607, ISSN 1860-2037.