Virtual Reality Simulator for Medical Auscultation Training

Luis Andrés Puértolas Bálint[1], Luis Humberto Perez Macías[2], and Kaspar Althoefer[1]

[1]Queen Mary University of London, London, UK
[2]Benemérita Universidad Autónoma de Puebla, Puebla, Mexico

l.a.puertolasbalint@qmul.ac.uk

Abstract. According to the Oxford English dictionary, auscultation is "the action of listening to sounds from the heart, lungs, or other organs, typically with a stethoscope, as a part of medical diagnosis." In this work, we describe a medical simulator that includes audio, visual, pseudo-haptic, and spatial elements for training medical students in auscultation. In our training simulator, the user is fully immersed in a virtual reality (VR) environment. A typical hospital bedside scenario was recreated, and the users can see their own body and the patient to increase immersion. External tracking devices are used to acquire the user's movements and map them into the VR environment. The main idea behind this work is for the user to associate the heart and lung sounds, as heard through the stethoscope with the corresponding health-related problems. Several sound parameters including the volume, give information about the type and severity of the disease. Our simulator can reproduce sounds belonging to the heart and lungs. Through the proposed VR-based training, the medical student ideally will learn to relate sounds to illnesses in a realistic setting, accelerating the learning process.

Keywords: Auscultation, Pseudo-Haptic, Virtual Reality, Simulator, Medical, Heart, Lungs, Pulmonary.

Introduction

We are currently experiencing a virtual reality (VR) boom in the medical industry. VR is still not capable of addressing all aspects of medical training. However, it has been shown to work very well in some cases. This is the case in problems such as phobia treatment [1] to surgical training [2]. In this work, a different approach based on VR for the training of medical students in auscultation is proposed. Keeping a view on affordability, so that the proposed medical training can be applied widely. VR development software is becoming more sophisticated [3]. Lower hardware prices and greater software functionality is leading to a point where VR training makes both practical and economic sense. This rapidly changing landscape has encouraged the development of training simulators that take advantage of the latest technologies. Programming VR has recently become accessible. While in not such a distant past, graphics engines were restricted to game developing. Graphics engines are now being used in serious simulations such as industry 4.0 [4] and robotics [5]. The available VR development tools have reached the maturity of allowing a single programmer or a small team of programmers to create a medical simulator in a short period. External medical input is also needed.

Since we aim to have a high degree of realism of the experienced motions and interactions in a simulator, it is essential to model the simulator's physics as accurately as possible. Since we want to obtain results quickly (but also accurately), we use an existing graphics engine that allows the rapid creation of realistic simulation environments. Modern graphics engines have the option to simulate a wide range of physical properties that happen. This allows VR developers to generate VR environments that ensure realistic motion of objects, (including visualization and tracking of human limbs) and the appropriate physical interaction between these objects (including collisions). Today, two graphics engines, Unreal and Unity, arguably dominate the market and are used by most developers. In this work, Unity is used. Through its asset store, Unity also allows us to incorporate portions of code and 3D models made by the community of VR software developers, into the application that is being programmed. This speeds up the development process considerably. We incorporated some valuable prefabricated Unity assets, including a hospital room [6], an accurate sinatomic model to be used as a patient [7], and a sophisticated medic avatar [8]. The acquired Unity models were slightly modified using Autodesk's Fusion 360. These changes include the clothing of the medic's and the patient's avatars inside Unity.

To maximize realism and minimize the cost of the simulator, we emphasize on the accurate modeling of natural human motion. Realistic collision and the visualization of the user's whole body in VR was one of the top priorities in the design of the simulator. It is crucial for us to appropriately model the user's hands and strive for realistic interaction with other objects in the environment. This includes avoiding the penetration of one object with another. For this auscultation simulator, much effort was put into acquiring our sound base, ensuring the highest possible level of realism. Some current auscultation simulators employ mannequins [9, 10], while others achieve VR auscultation using a head-mounted display [11]. These systems provide useful training tools for medical students. After careful consideration, the conclusion that mannequins are not necessary was reached. They are too expensive for widespread application, and virtual objects can do this in a better way.

To enhance immersion, aspects from the area of pseudo-haptics were incorporated. The use of the physics capabilities of our graphics engine, to detect collisions between the user's simulated hands and solid objects was also implemented. Upon collision, the realistic computation of expectable behavior of the hand is experienced. This way, it is tricking the mind into believing that there is actual resistance to the user's hands. Enhancing immersion in VR is very important, with other researchers coming to similar conclusions [12–14].

This paper is organized as follows. Section 1 provides an overview of how auscultation learning is currently studied. In Section 2 (Methods and Materials), we present our VR-based simulator for the training of medical students. In Section 3 (Pilot study), we investigate the capabilities of the simulator. Section 5 concludes our paper.

1   Background in auscultation training

Medical students usually learn auscultation through audio clips, graphs, and textbooks with images. In some training sessions, the medical professors knock on a table twice. The sound intensity and speed of the knocks are varied to emulate heartbeats caused by a health problem. Currently, medical students learn about auscultation in their study program. It is known that health problems that happen in specific organs affect the way the body sounds. Several health issues can be detected by hearing these sounds. Experienced medics can identify the differences between healthy sounds and those that are altered because of a health-related problem. They also view images showing the correct placement of the stethoscope on the patient. Learning in this way is protracted. Students find it difficult to fuse the different elements of auscultation into one whole understanding of the process. For example, it is difficult to understand which sound relates to which position of the stethoscope on the patient's body. Here, we create a more natural way of linking the placement of the stethoscope to the sounds and the associated illness. All this happens in a realistic environment, also incorporating the student's limbs' movements, visual impressions as well as the connected sounds. Because medics are separated in different knowledge areas, they will often only study a small part of the full spectrum of auscultation. This will make them likely to miss out on learning several sounds whose interpretation would help them diagnose a broader range of diseases.

We could only find two basic VR auscultation simulators currently being available [9, 11, 18]. Other auscultation simulators make use of mannequins and stethoscopes; employing mannequins increases the costs of such systems.

On the back, the six top auscultation points are P1-P6 where "P" stands for "pulmonary" and the two lower points are B1-B2, where "B" stands for "base" (see Figure 1 (right)).
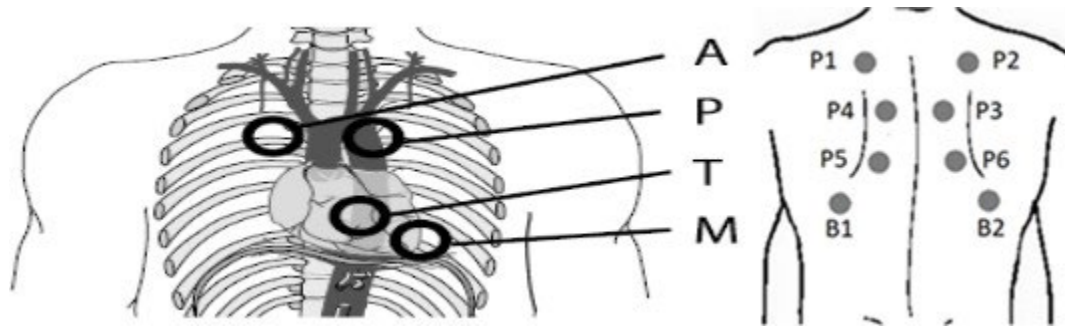
Fig. 1. (Left) Auscultation points on the chest. We can see the aortic (A), pulmonary (P), tricuspid (T), and mitral (M) valves auscultation points [15]. **(**Right) Auscultation points on the back of the patient. Adapted from Altan et al. [16]. If the stethoscope is not placed on these points, no relevant sounds will be heard, neither on the real patient nor our simulator.
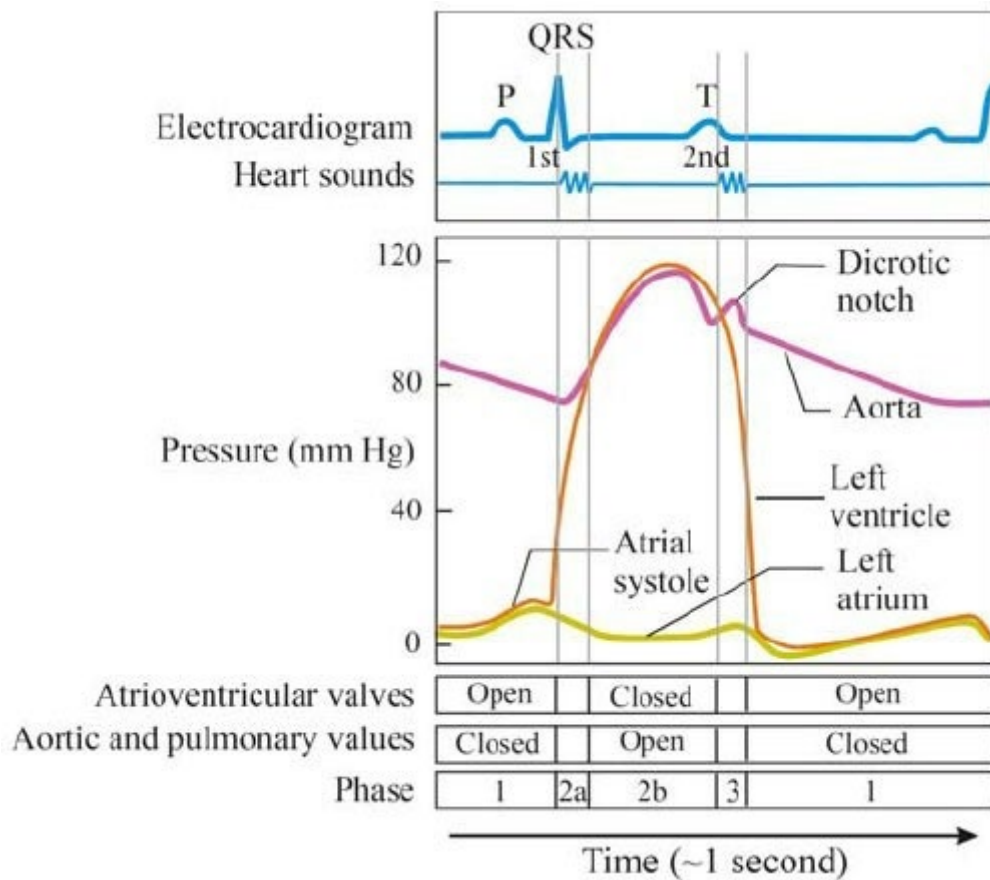

Fig. 2. During the learning of auscultation, students make use of tables and graphs like the one shown here [17]. These graphs are used to give a visual representation of the sounds that one is expected to hear during auscultation. These graphs mostly represent the first heartbeat, the second heartbeat, the sound intensity, and the timing between these two. Additional information in the form of textual descriptions (including sound amplitude, frequency, and type of disease) is usually provided.

2 Proposed Methods and Materials

The proposed simulator was created using the Unity graphics engine, making use of its physics simulation capabilities. Employing the assets from the Unity store, we download a realistic 3D environment resembling a standard hospital room. This is a place where auscultation is likely to happen. Our VR simulator requires users to have a commercially-available head-mounted display (HMD), compatible with Unity. From all the choices available at the time, the Oculus Rift was chosen in the end. This is because it is arguably the most comfortable, and less heavy HMD currently available.

To enhance immersion, the whole user's body movement was incorporated into the simulator. To capture the user's whole-body movement, a Leap Motion tracking system is attached to the HMD. The Leap Motion tracking system uses a camera system to observe its environment. It can extract the hands' and forearms' positions, and motions from the acquired images. This approach creates an impression in the user that their hands are moving in the created virtual environment, sharply increasing immersion. Alternative approaches such as incorporating a Kinect V2, or a newer variant called the Orbbec Astra to track the whole of the body were considered [19]. Adding a Kinect or an Orbbec Astra, greatly improves the tracking of legs and arms. The Leap Motion alone can estimate the whole-body position, with the help of the Oculus Rift's positional system. Incorporating the Kinect or the Orbbec, increases the cost and the space needed for the simulator.

To achieve a realistic interaction, it is essential to have realistic virtual collisions. This includes the whole virtual scene, actors, and objects. A collider (which is an integrative part of the modeled VR components) informs the graphics engine when an object is colliding with another object (for example the user's hand and the patient's body). Some of the mesh colliders we use in the simulation are shown in Fig. 3. In our experience, most simulators put too much emphasis on the graphical element of a scene's representation. This approach takes most of the computational resources by reducing the quality of the colliders, significantly affecting realism [20]. We believe that current simulators (as well as video games) use the wrong approach to achieve realism, which is to use most computational resources to have the best possible graphics. A priority to maximize the use of the physics libraries was given to the simulator, instead of spending most of the computational resources in the graphical quality.
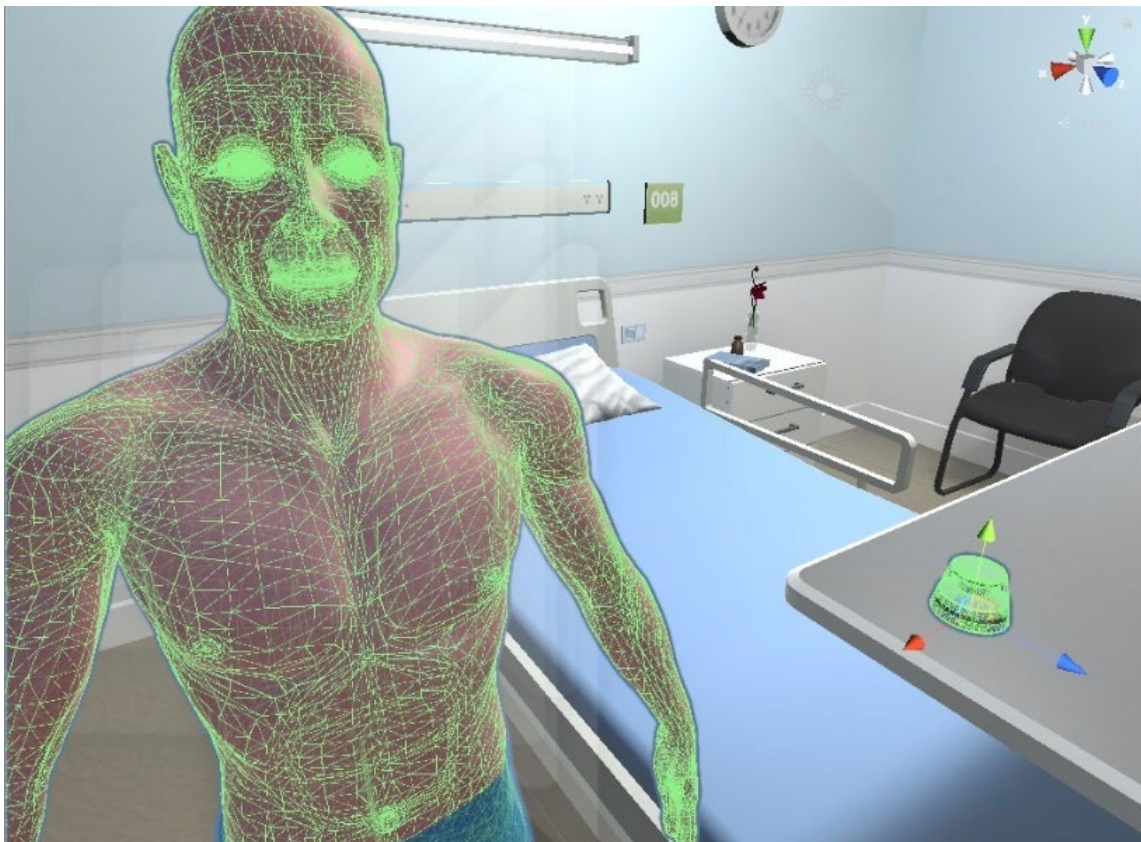

Fig. 3. Image of the proposed simulator, with a patient inside a standard hospital room. On the desk, we can see a wireless stethoscope. We added the most sophisticated colliders available, which are the mesh colliders (shown in green colour) to increase the realism of physical interaction in the simulator. This is how we create the illusion of touch or pseudo-haptics [21].

Using the Leap Motion tracking system and the described mesh colliders, the physical interaction of the user with objects in the VR environment was achieved. This means that the users can

use their hands to grab the stethoscope. The described features have been made available by the "Humanoid Control Plus" asset from the Unity Store. Since the simulator was developed using a potent laptop, we changed the colliders from cylindrical to mesh colliders [20]. Work is being done to improve the grasping of objects by using particles with Nvidia's FleX plugin [22], which is a state-of-the-art physics library that can be obtained as a Unity asset. This asset seems to work by detecting whether a hand is closed, and an object is near enough this hand to be grasped. With the FleX plugin which is part of Nvidia's PhysX [22], we can add particles, friction, stiction, stiffness, clothes, fluids, and other characteristics to make the simulator more realistic.
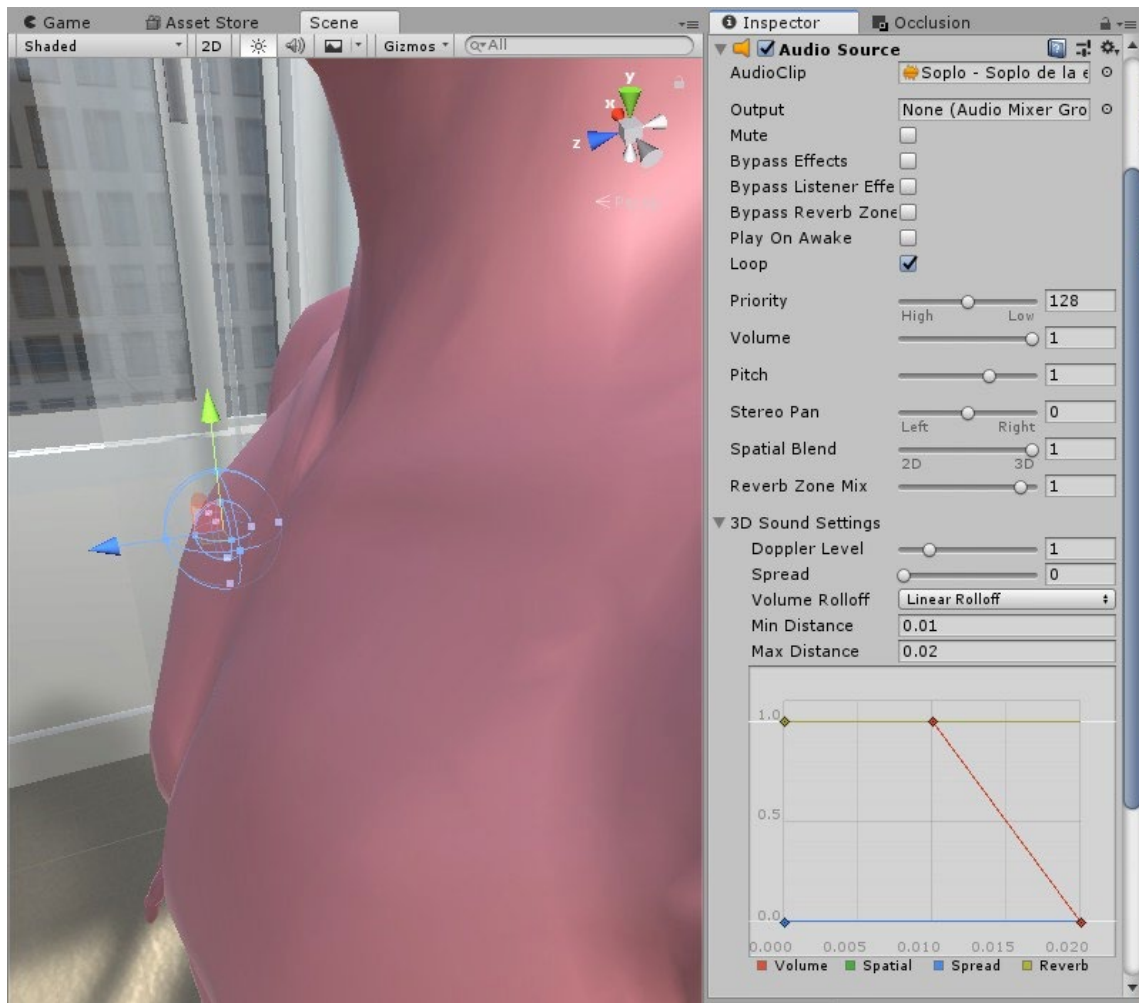

Fig. 4. After setting the scene, we add an audio listener and several audio sources [23][23]. The audio listener is usually positioned in the head of the character, centrally between the character's ears. We can define the audio listener as the ears of the simulation. As in the real world, the simulator provides sounds that are louder the closer we are to the sound source. However, if we find ourselves far enough, we will not hear that audio source. The scene contains only one audio listener, which is the stethoscope. Each of our 12 auscultation areas has audio sources, whose main parts are two spheres and with the respective audio clips that can happen in those places. We can see in this figure that there are several options to configure an audio source. The most important part is the volumetric sound consisting of two spheres. As we can see in the graph, the minimum distance is formed by a sphere of 0.01 distance units. A sphere of distance units forms the maximum distance on which the sound will propagate. A unit in Unity corresponds to one meter in the real world. This means that if the audio source (stethoscope) goes through the volume formed between these two spheres, the audio will play. These volumes were tuned until we could replicate a realistic experience in auscultation.

A table with health-related problems with their own sound (see Table 1). Our has an unknown number of possible states.

Auscultation point	Diagnostic							Sound %
Mitral	Normal	Mitral stenosis	Mitral insufficiency	Acute mitral insufficiency	Chronic mitral insufficiency	Mitral valve prolapses	Austin flint murmur	100
Tricuspid	Normal	Tricuspid stenosis	Tricuspid insufficiency	-	-	-	-	100
Pulmonary valve	Normal	Pulmonary stenosis	Pulmonary insufficiency	-	-	-	-	100
Aortic	Normal	Aortic stenosis	Aortic insufficiency	Acute aortic insufficiency	Severe acute aortic insufficiency	-	-	100
P1	Normal	-	-	Pleural rub	-	Sibilance	-	30
P2	Normal	-	-	Pleural rub	-	sibilance	-	30
P3	Normal	-	-	Pleural rub	Crepitations	sibilance	-	30
P4	Normal	-	-	Pleural rub	Crepitations	sibilance	-	30
P5	Normal	Rhonchus	-	Pleural rub	Crepitations	sibilance	-	30
P6	Normal	Rhonchus	-	Pleural rub	Crepitations	sibilance	-	70
B1	Normal	Rhonchus	-	Pleural rub	Crepitations	sibilance	-	70
B2	Normal	Rhonchus	-	Pleural rub	Crepitations	sibilance	-	100

Table 1. Overview relating auscultation points to diagnosis. It is noted that the sound amplitude is an essential element when distinguishing diseases. In the feasibility study, presented here, we explore a limited range of sounds. Future work will aim at considering the impact of all auscultation points for all the related heart diseases.

With the use of object-oriented programming, we coded 13 scripts. The first twelve scripts (used as objects) are placed in the auscultation points, and the main script is placed inside the patient. The twelve scripts linked to the auscultation points, make use of one method for each possible sound that we considered for so far (shown in Table 1). Every time the simulator runs, a feasible random health problem is created. The main script (inside the patient) oversees the task of generating a random health-related issue in the patient for the user to experience during the VR auscultation. Another option is to have a completely healthy patient can learn to identify them. Future work will explore cases of combined diseases, and the full use of our 23-sound database. Although this will be more challenging for the user, it will add more realism and has the potential to prepare medical students who use the simulator for real-life situations in a better way. Further study will also extend to other diseases in the thoracic area, beyond the set of illnesses discussed here. We will also explore how many combinations can happen; it is possible that they are more than previously thought. To further improve the simulator, more diverse models are needed, simulating different genders, differently sized people and people from different ages. Other aspects could also be considered, such as incorporating sounds from valve prostheses, pacemakers, stents and other types of implants routinely used on patients nowadays, and, hence, necessary to be included in medical training. Another critical aspect of future work will be to carry studies involving medical professionals and medical students.
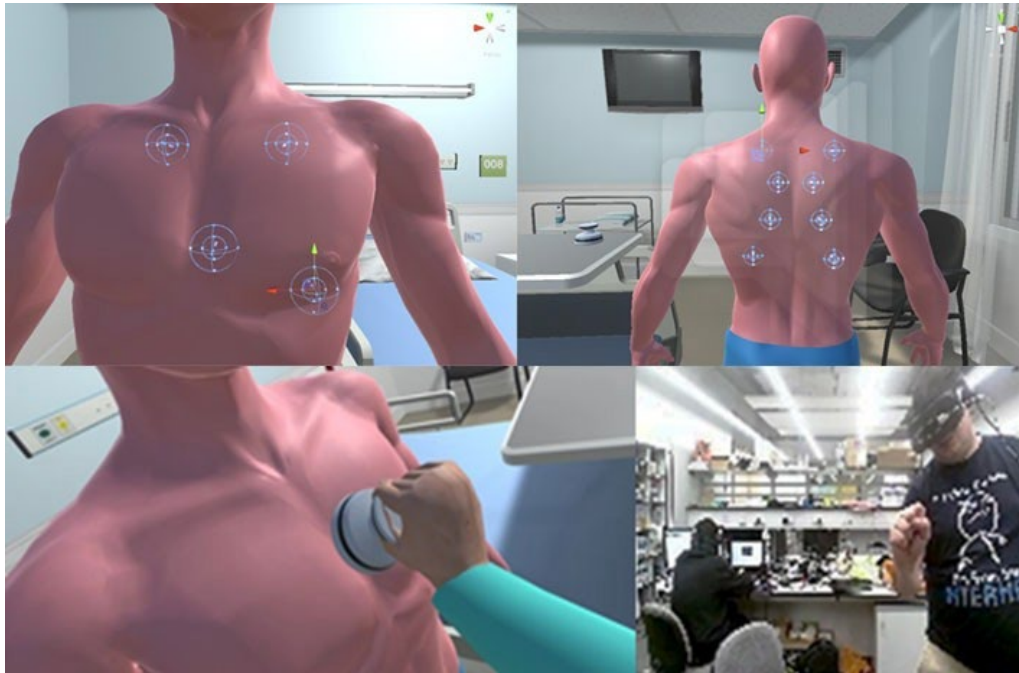
Fig. 5. (Top) Our VR patient with 12 auscultation points. Each auscultation point has two spheres. The inner sphere is where the sound experienced is the loudest. Moving from away from the inner sphere, the sound intensity linearly decreases coming to zero at the outer border of the bigger sphere. The three-dimensional spatial sound was also incorporated. (Bottom left) HMD image showing the user-held stethoscope on the chest of the patient. Because of the colliders integrated with the body and the stethoscope, the stethoscope will not penetrate through the body and create the illusion of resistance – this effect is called pseudo-haptics. (Bottom right) We see the user closing his hand to grab the stethoscope and moving it around to auscultate the patient.
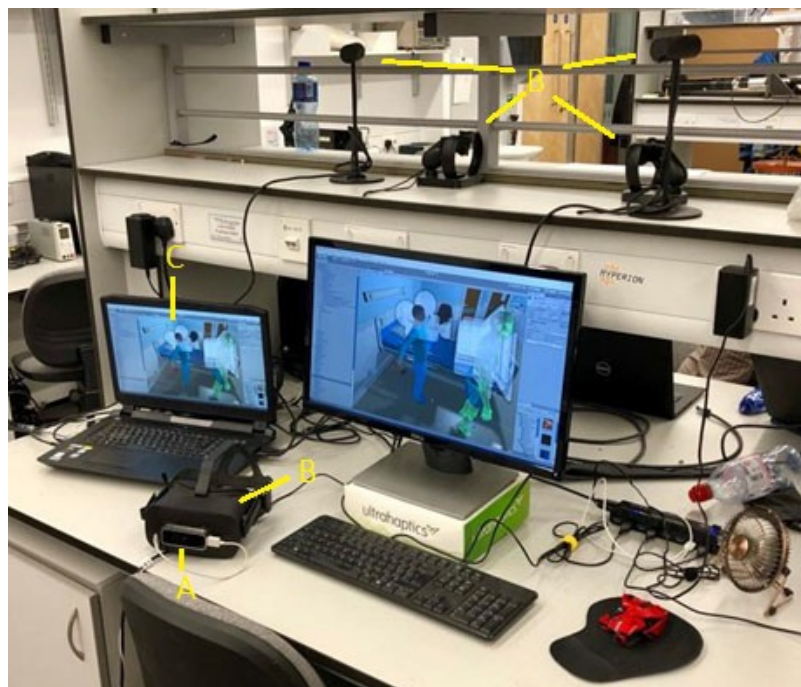

Fig. 6. Hardware configuration to run the auscultation training simulator. A) The Leap Motion. B) The Head Mounted Display with its controllers and tracking stations. C) A VR ready laptop.

It is essential for a training simulator to be as realistic as possible to maximize immersion and training outcomes. At the same time, our research aims to create a low-cost solution of a trainer so that medical students across the world can be trained without incurring high costs. To increase immersion, we make use of two important components: (1) we visualize the user's hands in the virtual reality environment in order for the user to get the feeling of being part of the pre-

sented environment, and (2) through the use of colliders on visualized objects we ensure that objects including the user's hands properly collide with each other and cannot penetrate each other. The visualization of the user's hands is achieved using a Leap Motion sensor mounted at the front of the HMD. Using the output from the Leap Motion sensor, the user's hands can be registered with the environment and the virtual objects within.

The colliders that are integrated with the VR objects ensure that the user's hands cannot move through objects. For the user, this creates the feeling that objects actually have density. This is an aspect of pseudo-haptics whereby the human mind is tricked to haptically experience a physical interaction even though there is no actual interaction between the user and the environment. In our experience, pseudo-haptics has shown to improve immersion significantly. Pseudo-haptics have, for example, shown that users can be tricked to believe in resistive forces when pressing down the piston of a syringe solely through appropriately manipulated visual feedback. Using pseudo-haptics, as done here, noticeably increases the feeling in the user to be part of the pre-sented reality and is likely to enhance the training performance.

3   Pilot Study

There is an unknown amount of realistic possible combinations. To reduce this combination, we ran a pseudo-random test. We first took 228 random samples (with python, as it proved too difficult to do it in C#). From this sample and to the best of our knowledge, table 3 was created. We debugged the code by reviewing 40 pseudo-random samples (C#) repeatedly until the resulting combinations were 100% accurate.

Auscultation point	Diagnostic						Sound %
Mitral	Normal	Mitral stenosis	Mitral insufficiency	Acute mitral insufficiency	Chronic mitral insufficiency	Mitral valve prolapses	Austin flint murmur	100
Tricuspid	Normal	Tricuspid stenosis	Tricuspid insufficiency	-	-	-	-	100
Pulmonary valve	Normal	Pulmonary stenosis	Pulmonary insufficiency	-	-	-	-	100
Aortic	Normal	Aortic stenosis	Aortic insufficiency	Acute aortic insufficiency	Severe acute aortic insufficiency	-	-	100
P1	Normal	-	-	Pleural rub	-	Sibilance	-	30
P2	Normal	-	-	Pleural rub	-	sibilance	-	30
P3	Normal	-	-	Pleural rub	Crepitations	sibilance	-	30
P4	Normal	-	-	Pleural rub	Crepitations	sibilance	-	30
P5	Normal	Rhonchus	-	Pleural rub	Crepitations	sibilance	-	30
P6	Normal	Rhonchus	-	Pleural rub	Crepitations	sibilance	-	70
B1	Normal	Rhonchus	-	Pleural rub	Crepitations	sibilance	-	70
B2	Normal	Rhonchus	-	Pleural rub	Crepitations	sibilance	-	100

Table 2. Our findings from a sample size of 228. This table remains the current standard we consider for measuring the combinations. From up to down, the algorithm has to choose one sound from each row. If the first sound picked is gray, it must go with other gray colours. The same applies to the black colour. The white can always happen in almost any possible combination.
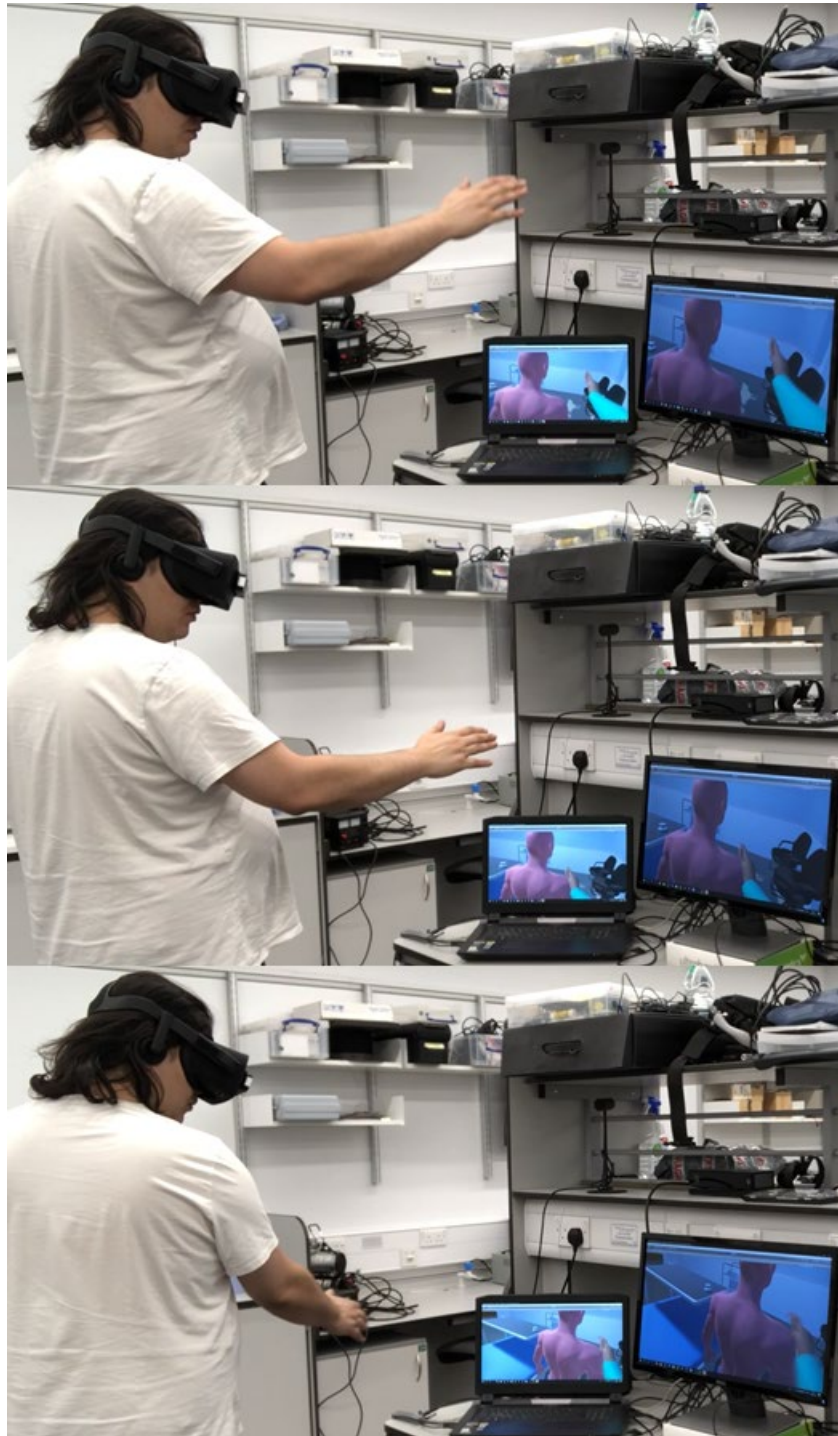
Fig. 7. From top to bottom. The first image shows no contact with the patient. The second image shows contact with the patient, where the real world and the virtual one overlap in both spaces, and there is contact with the patient's collider. The third image shows the pseudo-haptic effect, where the overlap is no longer existent. However, the third image tricks the mind of the user (medic) into thinking his hand is in the position of the virtual space, thanks to the complex colliders.

4   Conclusions

A virtual reality simulator for auscultation training of medical students was successfully programmed. Our work suggests there is potential for learning auscultation with an immersive VR environment. Accurately linking the sounds that the user can hear when using the simulator, to the diseases that produce these sounds in patients. This concept can reproduce the sounds of a range of diseases in the thorax and use an estimate to modulate their volume depending on the placement of the stethoscope. Through the inclusion of colliders and pseudo-haptics principles,

we were able to achieve a good level of immersion and feel of physical interaction with the simulated environment, while keeping costs and hardware complexity low. Four medics from different fields and levels of education were consulted for this research, one of them a co-author of this document.

Acknowledgment

References

1. Cruz, M., Amann, E.: Virtual reality and tactile augmentation in the treatment of spider phobia: a case report. Cap. Without Borders Challenges to Dev. 35, 39–52 (2010).
2. Aggarwal, R., Tully, A., Grantcharov, T., Larsen, C.R., Miskry, T., Farthing, A., Darzi, A.: Virtual reality simulation training can improve technical skills during laparoscopic salpingectomy for ectopic pregnancy. 1382–1387 (2006).
3. Wang, M., Reid, D.: Virtual reality in pediatric neurorehabilitation: Attention deficit hyperactivity disorder, autism and cerebral palsy. Neuroepidemiology. 36, 2–18 (2011).
4. PolliniMatteo, R.P.F.: The Use of Augmented Reality Glasses for the Application in Industry 4.0. Augment. Reality, Virtual Reality, Comput. Graph. (2017).
5. Tom Erez, Y.T. and E.T.: Simulation Tools for Model-Based Robotics: Comparison of Bullet, Havok, MuJoCo, ODE and PhysX. In: In Robotics and Automation (ICRA), 2015 IEEE International Conference on. pp. 4397–4404 (2015).
6. Hospital Room - Asset Store, https://assetstore.unity.com/packages/3d/props/interior/hospital-room-57399.
7. Anatomic Pack - Asset Store, https://assetstore.unity.com/packages/3d/props/anatomic- pack-79612.
8. Humanoid Control Plus - Asset Store, https://assetstore.unity.com/packages/tools/animation/humanoid-control-plus-104965.
9. Takashina, T., Clinic, T.: A New Cardiac Auscultation Simulator. 872, 869–872 (1990).
10. Kagaya, Y., Tabata, M., Arata, Y.: Variation in effectiveness of a cardiac auscultation training class with a cardiology patient simulator among heart sounds and murmurs. J. Cardiol. 70, 192–198 (2017).
11. Vargas-Orjuela, M., Uribe-Quevedo, A., Rojas, D., Kapralos, B., Perez-Gutierrez, B.: A mobile immersive virtual reality cardiac auscultation app. 2017 IEEE 6th Glob. Conf. Consum. Electron. GCCE 2017. 2017–Janua, 1–2 (2017).
12. Gruchalla, K.: Immersive well-path editing: Investigating the added value of immersion. Proc. - Virtual Real. Annu. Int. Symp. 157–164 (2004).
13. Stengel, M., Grogorick, S., Eisemann, M., Eisemann, E., Magnor, M.A.: An Affordable Solution for Binocular Eye Tracking and Calibration in Head-mounted Displays. Proc. 23rd ACM Int. Conf. Multimed. - MM '15. 15–24 (2015).
14. Bowman, D.A., Mcmahan, R.P., Tech, V.: Virtual reality - how much immersion is enough. (2007).
15. Bhoi, A.K., Sherpa, K.S., Khandelwal, B.: Multidimensional analytical study of heart sounds: A review. Int. J. Bioautomation. 19, 351–376 (2015).
16. Altan, G., Kutlu, Y., Pekmezci, A.Ö., Nural, S.: Deep learning with 3D-second order difference plot on respiratory sounds. Biomed. Signal Process. Control. 45, 58–69 (2018).
17. Israel, S.A., Irvine, J.M.: Heartbeat biometrics: a sensing system perspective. Int. J. Cogn. Biometrics. 1, 39 (2012).
18. Abbas, A., Fahim, A.: An automated computerized auscultation and diagnostic system for pulmonary diseases. J. Med. Syst. 34, 1149–1155 (2010).
19. Călin, A.D., Coroiu, A.: Interchangeability of kinect and orbbec sensors for gesture recognition. Proc. - 2018 IEEE 14th Int. Conf. Intell. Comput. Commun. Process. ICCP 2018. 309–315 (2018).
20. Matsas, E., Vosniakos, G.C.: Design of a virtual reality training system for human–robot collaboration in manufacturing tasks. Int. J. Interact. Des. Manuf. 11, 139–153 (2017).
21. Andrés, L., Bálint, P., Althoefer, K.: Medical Virtual Reality Palpation Training using Ultrasound Based Haptics and Image Processing. CRAS. (2018).
22. Macklin, M., Müller, M., Chentanez, N., Kim, T.-Y.: Unified particle physics for real- time applications. ACM Trans. Graph. 33, 1–12 (2014).
23. Unity - Scripting API: AudioSource, https://docs.unity3d.com/ScriptReference/Audio Source.html.