An Affordable Virtual Reality Training Aid for Anaesthetists.

PETER BENTLEY
UNIVERSITY OF HUDDERSFIELD

The use of virtual reality for medical training applications is increasing rapidly. Until now, the systems have been mostly to aid surgeons, but there exists another group of medical staff that would benefit from virtual reality also - the anaesthetists. There is a great need to increase the experience of trainee anaesthetists without endangering patients. An affordable virtual reality training aid for anaesthetists would accomplish this. In the context of some currently available products, the requirements for such a system are examined, including hardware, software, cost and design.

Why Virtual Reality?

Virtual reality (VR) systems are being developed world-wide for a huge variety of medical applications. Systems for training surgeons for laparoscopic and endoscopic surgery techniques, for the teaching of anatomy, for the analysis of medical scanners and even to help monitor the progress of recoveries, are in development [1-6,8]. However, the use of VR for the training of anaesthetists - an equally important area - does not appear to have been examined yet.

This need for a training aid for novice anaesthetists arose from discussions with anaesthetists during the development of an expert system[1]. It seems that, as with surgeons, there is a major problem when training new anaesthetists: dangerous and unusual complications can only be read about or described to the trainee, they cannot actually be observed by the trainee (for obvious reasons). This means that when the inexperienced anaesthetist finds himself in an emergency situation where the patient is reacting very badly to the anaesthetic or operation, he has only the knowledge gained from books to guide his actions. If the anaesthetist had actually *experienced* a similar situation during training, it seems highly probable that his actions would be very much better in real life.

Since the construction of a complete training operating theatre with human 'dummy' to simulate the patient and training versions of the monitoring equipment would cost a fortune to both develop and construct, the best solution seems to be a virtual reality training simulator. This would allow a realistic simulation of the theatre, patient and monitoring equipment to be achieved at a more reasonable cost, and would take up very little space [6].

---

[1] The software was to monitor the outputs from anaesthetists' equipment and check to see whether they became unstable. This could arise both from faulty equipment or the deteriorating condition of the patient. If the latter was the case, it would suggest in real time possible causes of the problem.

This 'virtual' operating theatre would then be able to simulate various operations taking place, the trainee playing the part of the anaesthetist. The tutor would specify an appropriate 'situation' (i.e. unusual complication), and the virtual patient and monitoring equipment would behave in the appropriate manner. The novice anaesthetist would then administer the appropriate remedy (or not) and the virtual environment would be updated once more. (Note that in some cases, the best action might be to just observe the patient more closely - overreacting to spurious data can be worse than not acting at all.) In this way a trainee would gain valuable experience of such situations without putting any real patients at risk [7,8].

VR Hardware

Since this application is specifically for medical purposes, cost is an important feature. A top of the range VR system would be very effective for training, but its cost would be very high. Hospitals and medical schools, as with many establishments today, have very limited budgets, so a £50,000 (or more) system would not have many buyers. For this reason, it is essential that the cost of the completed system be kept as low as possible.

This requirement immediately effects the choice of hardware to be used. The most expensive part of a VR system tends to be the headset with head-tracking devices [1] p.95. Now, for the system described above, clear, detailed graphics are necessary to allow the user to see the simulated monitoring equipment display, and to observe the virtual patient's condition. Most affordable headsets are based on colour LCD technology, all of which have very low resolutions - typically little more than around 185 by 139 pixels per eye (often quoted as being higher). This is plainly inadequate - the user would have to 'virtually' place his head a few centimetres from the simulated monitor display in order to make out any of the details, and by that time the patient's condition could be terminal. The main alternative is a headset based on cathode-ray tube technology. Amongst these types are miniature CRTs[2], as favoured by the military and fibre coupled displays[3]. Both give images of excellent resolution - often more than 1000 by 1000 pixels, but typically the images are a very small size, meaning the user makes no use of his peripheral vision. These kinds of systems are certainly the best for displaying detailed simulations, but tend to cost tens of thousands of pounds.

Another requirement of the system is that the user must be able to remain 'immersed' in the virtual theatre for large periods of time. (Note that, at least initially, the simulation will be of the maintenance stage of anaesthesia and a session could last for as long as three hours.) Now, all but the very best headsets currently suffer from a major problem - time (or phase) lag. This is the extremely disorientating effect when the headset wearer moves his head, and the images he sees are slow to catch up. The restriction of the level of comfortable movement the user can make causes a feeling of being underwater, and is not desirable for long periods of time.

---

[2] Twin minature CRTs are mounted in front of the eyes to be viewed via the optics. The health issues from having CRTs so close to the eyes of the user are still somewhat unresolved.

[3] Images from remote CRTs are transmitted through coherent fibre optic bundles to the eyes using a collimating lens system [9].

Another common problem with the cheaper headsets is the quality of the optics used to allow the user to focus onto the twin displays. The quality can be so poor that stereo fusion becomes a real strain, resulting in nausea, dizziness and headaches after short periods of time. So again, the only suitable headsets and head-trackers appear to be the most expensive ones [9,10].

For these reasons, and many more, it is clear that until the current display and head-tracking technology is improved, with a large reduction in cost, VR headsets are just not a feasible option for a realistic training aid [5]. Instead, traditional monitor technology will almost certainly have to be used.

Input devices also pose problems. The trainee would hardly want to spend a couple of weeks becoming used to a specialised device such as a pressure ball, nor would he want to spend thirty minutes of every training session calibrating a glove. The remaining suitable low cost devices include a variety of cheap '3D' mice such as the Logitech 3-D Mouse and Multipoint's Z-Mouse, the ordinary 2D mouse and standard analogue joysticks. Any of these simpler and less advanced devices allow full, and fairly accurate movement around a virtual world, and take very little time to learn to use. So again, unfortunately, because of the current inadequate and high cost input devices, the choice must be for hardware a few steps back from the virtual reality ideal of full immersion in the virtual world.

The nature of the problem of simulating a virtual operating theatre requires quick and accurate movement occasionally. For instance, if the patient's condition suddenly deteriorates, the trainee must be able to change his virtual view to examine the patient and monitoring equipment as quickly as in real life - something not achievable by waving one of the current obscure input devices about. This means that some pre-set movement sequences are necessary, e.g. a 'look at patient's eyes' or 'close-up of E.C.G.' command - yet another step away from virtual reality. The problem of *fast* and accurate movement in a virtual world has not yet been addressed in virtual reality research - currently today's technology still struggles just with slow accurate movement.

An important aspect of simulating anaesthetists' equipment, is sound. There are various continuous noises such as the artificial respirator that help the anaesthetist to monitor the patient's condition by ear as well as by eye. Audible alarms are vital in the simulation also, to allow the quick recognition of problems. Ideally these sounds would be generated in stereo, and 'convolved' to give the illusion of sound coming from the right place in the virtual world, but this is not a vital requirement and the extra cost for the hardware is high. A basic mono-output capable of playing simple digital samples is sufficient, and many computers today already have this capability.

Finally, the last and perhaps most important part of the VR hardware must be selected - the computer itself. Today there exist a large number of high speed processors, with even faster ones coming on the market every few months. A simple calculation can be made to determine the minimum performance needed by the processor: assume that for the virtual operating theatre a minimum of 1000 polygons would be required. Ideally, for smooth animation at least 10 frames per second are required. The processor (and display hardware) therefore must be able to cope with 10,000 polygons per second. Note that there would almost certainly be a combination of flat-shaded, Gouraud-shaded, and texture-mapped polygons, effectively meaning the graphics workstation would need to be able to render between 10,000 and 15,000 basic flat-shaded polygons per second[4]. Obviously if stereoscopic images were required, this figure could double.

Performance ratings given by the various workstation manufacturers inevitably vary both in value and accuracy, but it seems likely that the lower-end machines have to be ruled out for this processor-intensive task. A Pentium PC with advanced graphics card would be able to cope with a low-resolution, monoscopic output, a Sun SPARCstation would manage the full stereoscopic display, and an Onyx Silicon Graphics Workstation could handle the job at a very impressive number of frames per second. With prices dropping continuously, it will not be long before the higher performance machines become affordable for a task such as this, but currently it seems a high-end desktop or average workstation (such as a Pentium PC or SPARCstation) gives the best performance for the money.

VR Software

When developing a new VR system, the choice of software[5] can be just as important as the hardware. Indeed, if a poor choice of software is made, it can actively prevent the use of certain hardware. (It is not uncommon for the choice of the software to be made first, hardware second.) The requirement must therefore be for software that is affordable, compatible with the required hardware, and produces fast and realistic virtual worlds. Other, less important, but still desirable features are for the software to be portable across many hardware platforms, and for there to be additional 'world designing' tools to allow the easy design of realistic items in the virtual world.

At the present time, there are few virtual reality software development tools capable of fulfilling all of the above requirements. Many only support specialised hardware or are far from affordable. The more affordable solutions range from the more costly high-level development packages to the cheaper lower-level development libraries (usually for a language such as C or C++). The former solutions allow quick prototyping of the virtual worlds, sometimes at the cost of speed, whilst the latter ones usually only allow slow prototyping, producing quick virtual worlds. Obviously the balance of development time against cost and performance for the proposed system would determine the best compromise.

---

[4] The more advanced hardware can perform specialised graphics functions such as texture mapping in the hardware itself, thus reducing the speed difference between rendering basic flat-shaded polygons.

[5] Note that only software development tools and libraries are being considered here, it is assumed that few people would want to write an entire Virtual Reality Engine from scratch.

Examples of such products available today include 'Superscape VRT 3' from Superscape Ltd, 'WorldToolKit' from Sense8 Corporation and 'OpenGL' from Silicon Graphics. Superscape provides excellent high-level development and world-designing software, but the price is still somewhat high. Performance and realism are not the best, because of the need to always have Superscape in memory when running virtual worlds and because only PCs are supported. WorldToolKit is a more affordable software solution, if somewhat basic. The software consists of a set of C libraries to allow virtual reality rendering at high speed on a large number of different platforms (including PCs, Suns, and the SGi Reality Engine). There is no form of world-designing software included, but the libraries allow the input of standard CAD files, so a normal CAD package can be used. It also supports a large range of input and output devices. OpenGL is the cheapest and lowest-level solution, providing the basic graphical libraries required, but no support for standard data formats (the programmer must define his own). Again, a range of workstations are supported. Additional software to help the design of the worlds is available from various companies.

Designing a Virtual World

Realism in a virtual operating theatre is essential. The trainee must be able to see the difference between a simulated ill patient and a healthy one; a patient with a square robot-like face with even coloration would not be very beneficial to this. Now, data for the virtual world would normally be stored as collections of polygons[6]. The use of flat-shading [7] when rendering items exhibiting properties of diffuse reflection such as floors and walls is quite acceptable, but for the patient causes problems. Ideally, the parts of the patient visible (for the purposes of the simulation, normally just the head and hands) should be smooth and flesh-coloured, and lit by the moveable overhead lighting of the operating theatre. Flat-shading a face made up of polygons would only highlight the artificial-looking corners and flatness in the simulation, so some alternative is required.

There are two main alternatives - change the model of the patient, or change the way the model is rendered. The first method simply involves increasing the number of polygons used to represent the patient. The closer the size of individual polygons get to the pixels on the display device when rendered, the smoother the appearance is[8]. Unfortunately, the more polygons used, the more processing power is required as discussed earlier, so this is not the best solution. The other method involves the use of an alternative shading algorithm when rendering objects that need to appear smooth. Such algorithms do exist (e.g. Phong, Gouraud) and are supported by the available software. The use of such algorithms must be limited to only those objects that require a very realistic appearance though, since a much larger degree of processing time is needed [11].

---

[6] The perhaps better representation of Bézier or parametric surfaces is not yet widely supported by the available software.

[7] For example constant shading, or interpolated shading, where the whole of every polygon is a single colour.

[8] The benefits of this method are much reduced by the Mach band effect, where the eye is drawn to the boundary between two shades. Decreasing the polygon sizes helps little to remove the effect.

So the problem of displaying in real time a realistically shaded face is not a great one, but how to actually obtain the data for the face in the first place? There are two possibilities - data could be created by hand, with the help of an artist to keep the features realistic, or a real face could be digitised somehow. The first option is acceptable if an artist with a great deal of spare time is available, but the second seems the best for a very accurate and realistic simulation. The digitising of a face in 2D is elementary, but in 3D the problem is not trivial. There are many techniques today to achieve this, involving everything from lasers to 'pin sculptures', but the simplest seems to be Moiré fringe techniques as used by researchers at the National Engineering Laboratory. Light is projected through a grating onto the face and then photographed through a second grating. The result is a contour map of the face which can then be used to form the full three dimensional model. The accuracy is high, since the data is captured in seconds avoiding errors due to movement. The same method could be used to create models of hands also. (There is another, more costly solution to the problem of obtaining the data - buy it from one of the increasing number of sources.)

Sharp details in the simulation could be created by texture mapping. For instance, the graphical display of the monitoring equipment could be produced by a sequence of bitmaps mapped onto the 3D virtual monitor. Details such as watery eyes could be just digitised images of real eyes mapped onto the eyes of the patient. Again, texture mapping (especially when perspective-corrected and shaded) can involve much processing power, and should be kept to a minimum. The sensible use of texture mapping can provide extraordinary realism to a simulator, however.

There are also some simple ways to improve the speed of a simulation. It is unnecessary to create full models of every object in the operating theatre - for instance a cupboard that cannot be opened hardly needs to be filled with detailed virtual contents (in fact, it does not even need a back to it). Those parts of the patient covered by cloth do not even need to exist, as long as the cloth gives the appearance that the patient is beneath it. Using these standard methods the polygon-count of a virtual world can be reduced drastically, thus improving the speed of the entire simulation.

Controlling VR

The use of virtual reality for an anaesthetist's training aid is limited to the front-end of the software only, i.e. the part the user sees and interacts with. What goes on behind the scenes to actually control the actions and responses of the virtual operating theatre is just as important. A full simulation of monitoring equipment with their own graphical displays and sound, the patient's external and internal reactions to the operation and anaesthetist's drugs all must be achieved in real time. However, this does not mean that the simulation cycle (the updating of every part of the simulation) need correspond to the frame-rate of the VR display. A lower or even unsynchronised simulation cycle would be acceptable, and quite possibly essential to keep the speed of the system to a sensible level[9].

The diagnosing expert system developed in previous work relied heavily upon the use of Prolog (a specialised, but slow logic-based programming language). It seems likely that a similar system will need to exist as part of the simulation controller. However, since speed is vital for the sake of smooth movement, it seems sensible to implement the simulation-controlling expert system in the same language as the VR front-end (almost certainly the object oriented language C++).

Now, virtual reality is inherently object oriented. Not only are there separate objects visible and often movable in a virtual world, but the underlying representations are often coded as separate objects. It would perhaps be an attractive solution therefore, to continue this form of representation down to the level of the virtual world controller, and have the simulation take place in terms of messages between objects and inheritance. So the patient-object would send messages to the monitor-object describing its condition, the monitor object then displaying the information. The surgeon-object would send messages describing its actions to the patient-object, making the latter react in some manner. It should be apparent that these computer representations match very closely the things in real life that they are simulating, meaning the simulation would be that much more accurate.

To create a simulation controller using a combination of object oriented methods and an anaesthetic knowledge based system (expert system) would almost certainly involve the use of a number of specialised artificial intelligence techniques [12]. Once designed though, the result should be a highly modular, expandable and realistic training aid.

Conclusions

Virtual reality may not be a new idea, but today the technology is plainly either overpriced or inadequate (and often both). The only way to create a usable and affordable training aid for anaesthetists is to make compromises. A sensible solution today would require a fast computer with basic sound, monitor and mouse. (For a training aid such as this, the level of immersion as created by VR headsets is not essential, only desirable.) With the right development software, a quite acceptable desktop VR system could be produced.

Serious consideration must be taken in the design of the virtual world - how much detail is needed and how much speed? The two requirements are often not compatible. There is also more to a simulator than just the graphics. A rigorously designed controller is essential for a simulator that acts as well as it looks.

---

[9] In other words, unlike the virtual world, the graphical displays of the virtual monitoring equipment (for example) need only be updated two or three times a second.

There is, however, a real need for an affordable virtual reality training aid for anaesthetists. Once completed, it would form part of a more comprehensive training programme for anaesthetists, and would be of real benefit. It would also be a novel and commercially viable use for virtual reality today.

References

[1] Pimentel, K. & Teixeira, K. *Virtual Reality Through the New Looking Glass.* Intel/Windcrest/McGraw-Hill, 1992.

[2] Rheingold, H. *Virtual Reality.* New York, NY:Summit Books, 1991.

[3] Helsel, S.K. & Roth, J.P. *Virtual Reality - Theory, Practice and Promise.* Meckler, 1991.

[4] Avis, N. J. *Medical Applications of Virtual Environments.* Proceedings from the 1st UK VRSIG Conference, Nottingham University, 1994.

[5] Stone, R.J. *Virtual Reality: Interfaces for the 21st Century.* Proceedings from Advanced Information Systems: the New Technologies in Today's Business Environment, Incorporating the Sixth International Expert Systems Conference, Oxford, 1991, p.99-110.

[6] Theasby, P.J. *The Virtues of Virtual Reality.* GEC Review v7:3, UK, 1992, p.131-45.

[7] McLellan, H. *Virtual Environments and Situated Learning.* Multimedia Review v2:3, USA, 1991, p.30-7.

[8] Middleton, T. *Applications of Virtual Reality to Learning.* Interactive Learning International v8:4, UK, 1992, p.253-7.

[9] Kalawsky, R.S. *The Realities of using Visually Coupled Systems for Training Applications.* Proceedings of the SPIE - The International Society for Optical Engineering v1695, USA, 1992, p.72-82.

[10] Rowley, T.W. *Problems of Head Mounted Display.* IEE Colloquium on 'Real World Visualisation - Virtual World - Virtual reality (Digest No.197), London, 1991, p.12/1-4.

[11] Foley, J. D. et al. *Computer Graphics Principles and Practice, 2nd Ed.* Addison Wesley, 1991.

[12] Williams, P. *Applying AI to Virtual Environments.* AI Expert v7:8, USA, 1992, p.26-9.