Immersive Environments: Real Problems, Virtual Solutions

Professor Paul Chapman
Head of the School of Simulation and Visualisation
Glasgow School of Art

September 2018

Abstract
In this article I will review the evolution of computer graphics and the long anticipated 'coming of age' of virtual reality. I will highlight the power of multidisciplinary teams and describe a number of diverse projects undertaken at the School of Simulation and Visualisation (SimVis) at The Glasgow School of Art that each utilise immersive environments. Specifically, medical visualisation, pharmaceutical visualisation, dangerous sports and hazardous work environments. Finally, I will describe how scientific tools can generate beautiful 3D immersive point cloud images.

Background
Before joining The Glasgow School of Art, I had always worked very much within a single discipline. I graduated in computer science and in my early career, I would always find myself surrounded by fellow computer scientists or electronic engineers. Whilst this seemed to be effective for the work I was doing (working in the offshore industry), I did find that everyone's thought patterns seemed to be very similar resulting in a rather linear approach to group problem solving.

Joining The Glasgow School of Art and SimVis was a breath of fresh air with its intense learning and research environment exploiting the interface between science, technology, industry and the arts. What makes SimVis unique is the exciting melting pot of different disciplines. Computer scientists sit next to 3D modellers who in turn sit next to psychologists, product designers, artists and photographers. Everybody has their own way of thinking and tackling a problem and each individual brings something fresh to the group enabling truly innovative work to be achieved within the team.

The Evolution of Computer Graphics
As a child growing up in the 70s and 80s I was very much part of the home computer revolution. For the first time, we (the general public) were able to purchase (at a very reasonable cost) our own programmable computers for home. Devices like the Sinclair ZX81, 48K and the Commodore 64k found their way into millions of people's homes enabling computer nerds like me to start programming and experimenting with the limited graphics and memory available.

Since then, computers and graphics cards have improved significantly. One of my favourite demonstrations of this advancement is the rapid evolution of computer graphics rendering since the 1980s. Figure 1 (left) shows a typical monster rendered from the early 1980s using a Sinclair ZX81. Colours were limited to black and white and the resolution was an enviable 64x48 in graphics mode. In Figure 1 (right) we see the 2014 modern day 'monster' equivalent ('Alien Isolation' ©2018 SEGA). The difference in quality is clear and significant and it must be recognised that it has been the huge demand from the games industry that has forced the rapid evolution of high powered, low cost graphics cards that now handle over a billion colours, real-time raytracing and the processing power to render photorealistic design concepts interactively. Thanks to the games industry, we can now take advantage of these advances in computer graphics for non-gaming, serious applications. This then begs the question, what will computers be capable of rendering in the next 35 years?

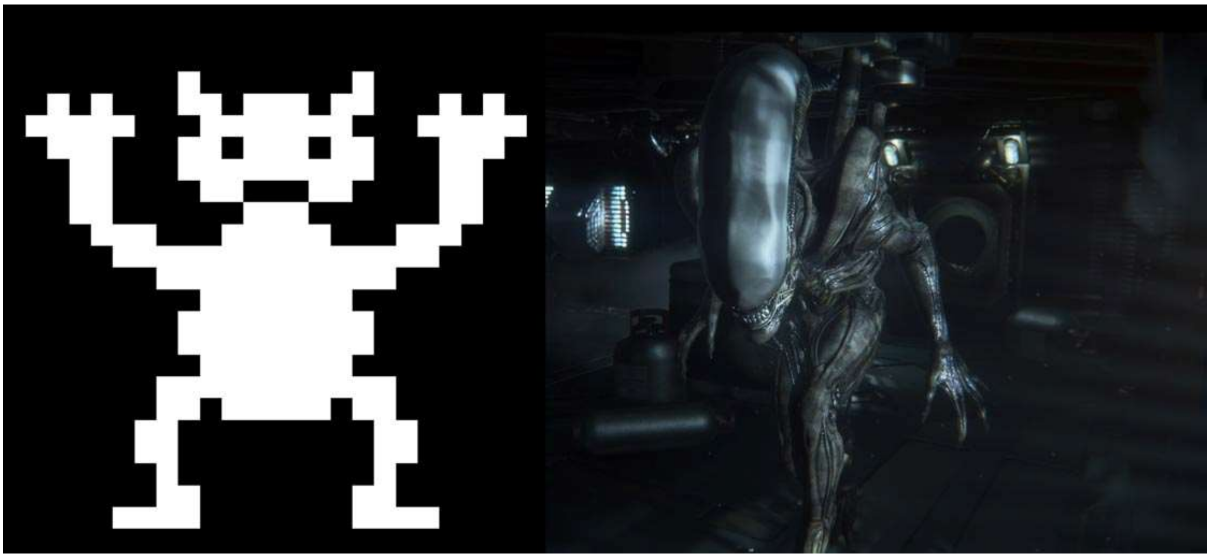

Figure 1 – Computer Graphics Monster Evolution 1981-2014
Alien Isolation © 2014 SEGA used with permission

Evolution of Virtual Reality

Virtual Reality has had many false starts, most notably during the mid 1990s. Whilst working as a computer scientist at the University of Hull, I distinctly remember films like the Lawnmower Man that got the media and general public excited and hyped about VR. However, when the press descended on university computer science departments around the world they were extremely disappointed with the bulky, low resolution, high latency helmet mounted displays that showed no relation to the exciting pre-rendered graphics of the film. Indeed, VR products such as Nintendo's Virtual Boy (released in 1995) were a complete failure.

Gartner Hype Cycle

Looking back at the first wave of failed virtual reality products, one is reminded of the work of the American researcher and futurist Roy Amara and his famous adage relating to forecasting the effects of technology:

"We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run" [1]

Amara's law is beautifully illustrated by the Gartner Hype Cycle (Figure 2). The respected consultancy firm Gartner provides this graphical representation annually to track the gradual adoption of a technology or product [2]**.** Their Hype Cycle is divided into five phases:

The technology trigger. A new product/technology 'breaks through' via prototypes. There may be proof of concept demonstrations which can trigger significant media interest and publicity. At this stage it is very rare for a usable product to exist.

Peak of Inflated Expectations. The technology will be implemented by early adopters and there will be a lot of publicity relating to its successful (and unsuccessful) implementation.

Trough of Disillusionment. Interest dissipates as implementations of the technology fail to deliver. Investors continue to support the technology only if the problems can be addressed and the technology improved.

Slope of Enlightenment
Second and third generation products emerge from companies and the technology starts to see more investment. More examples of how the technology can provide real returns on investment start to become understood.

Plateau of Productivity
The technology is extensively implemented and its application is well-understood resulting in mainstream adoption. Standards start to arise for evaluating technology providers.

Figure 2 shows Gartner's Hype Cycle for AR and VR technologies for 2017/18. AR for 2017 (yellow circle) and 2018 (green circle) is positioned within the Trough of Disillusionment. Even with the significant financial investments of recent years, the technology is still seen as being disappointing. Conversely, in 2017 VR was positioned on the Slope of Enlightenment and in 2018, VR is not present. Why? Because Gartner considers the technology to have matured to the point where it is no longer an emerging technology [2].

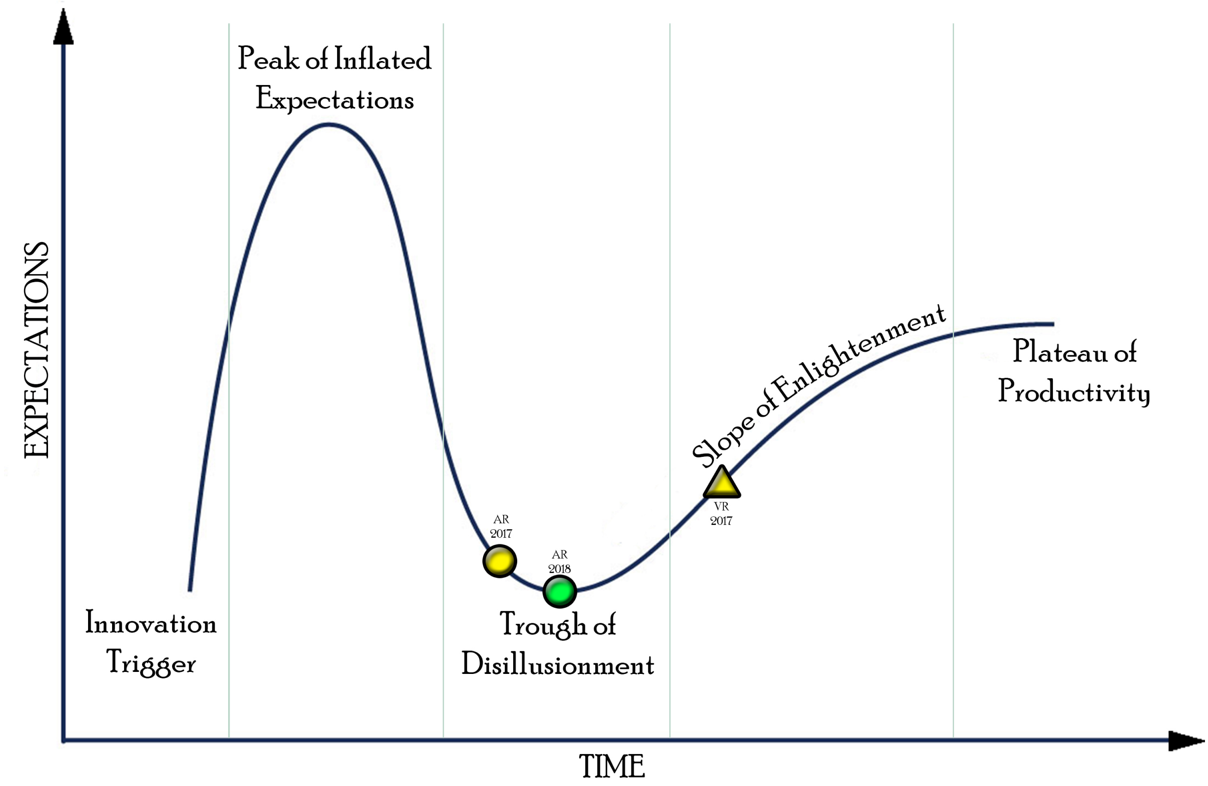

Figure 2 – Gartner Hype Cycle 2017/2018 – Augmented Reality (AR) and Virtual Reality (VR).
In 2018, VR has matured sufficiently to move it off the emerging technology class of innovation profiles.

The rapid maturity of VR is further evidenced by a greater understanding by the general public of its applications. A recent example of this in Scotland is the adoption of VR technology by local government through the installation of headsets in primary and secondary schools leading to a greatly enhanced student learning experience. After VR's numerous false starts, it is satisfying to see the technology finally breaking through and of course it is only going to improve.

The next sections of this article will describe a number of projects at SimVis that have successfully utilised VR for serious applications outside of the gaming sector.

Medical Visualisation

For over a decade, in close partnership with the world's anatomical and medical experts, SimVis has been developing new visualisation systems for improving healthcare. Our sharp focus has been on developing accurate, interactive 3D models of human anatomy for learning and training purposes. As a natural adjunct to this resource, we have also created immersive, interactive training simulations of challenging scenarios that pharmacists or dentists are likely to face throughout their career.

Anatomical Visualisation

Traditional methods of learning complex human anatomy are still centred on 2D images and illustrations in medical textbooks. Over the last ten years, SimVis has been developing an extremely detailed clinical-normal anatomical model of the human body. This model has been painstakingly created from first principles by a multi-disciplinary team including medical experts, computer scientists and artists beginning with innovative data capture including directed cadaveric dissection, live surgery, CT and MRI scans. The resulting model, the Definitive Human, has been verified by trusted experts in the field including senior world renowned anatomists from several countries.

Figure 3 shows the completed Definitive Human at various levels of construction. Using standard 'games ready' computers, all learners from medical students to practicing surgeons are now able to systematically and regionally strip the body apart for a completely novel and inspiring anatomical learning experience.

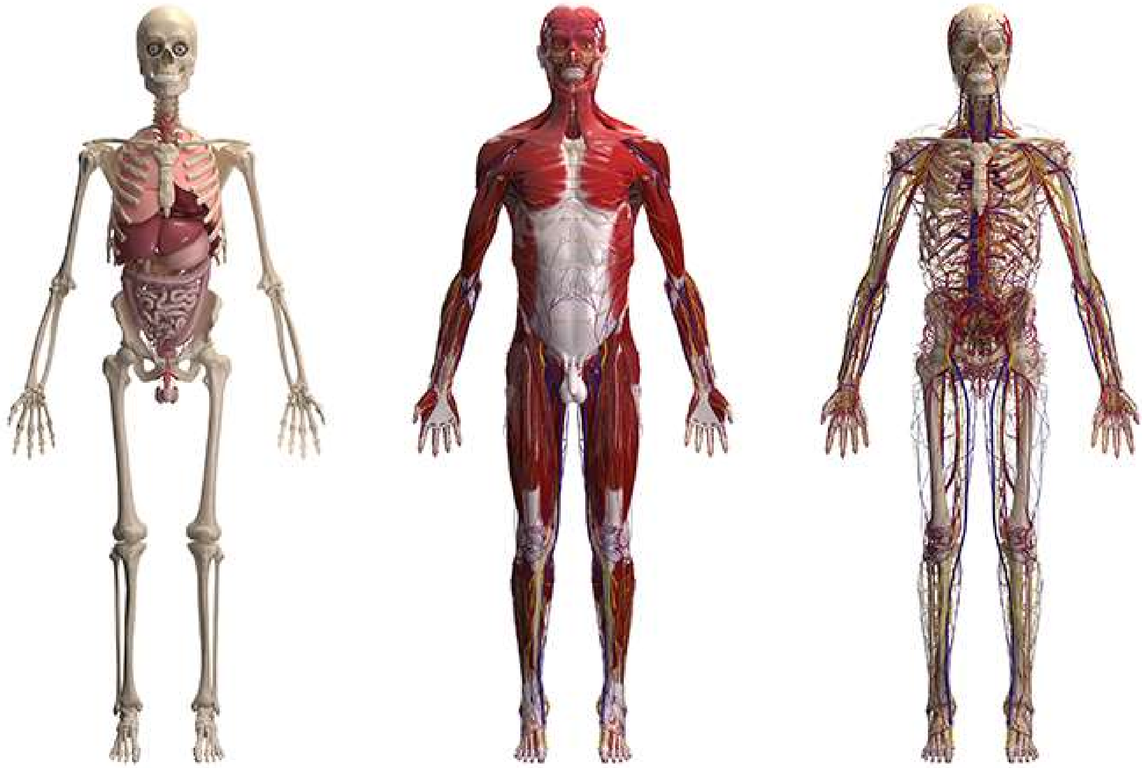

Figure 3 – The Definitive Human Dataset

In recent years, we have combined our completed Definitive Human model with a VR interface. This has radically improved the potential for meaningful interaction and understanding of this extremely complex dataset. Users can now walk up to a full size human male and using their hands, pull the Definitive Human apart for personal exploration. Every anatomical component is clearly labelled and once selected and examined can be returned to its original position at the touch of a button. To enhance the experience, the user can grab a virtual overhead light and hold it up to items of interest, such as the skull, in order for the light and shadows to detail the complex morphology of the bones. As a new trusted resource, the Definitive Human is proving to be the preferred method of learning for anatomy students and is also being used successfully within surgical training. Our head and neck anatomy model is a key learning resource across the NHS (National Health Service) and is supporting dentists throughout the UK. Figure 4 shows a medical student interacting with the Definitive Human dataset using the new HP Z mobile VR backpack enabling them to move freely, untethered by cables, within a realistic virtual environment.

Figure 4 – Virtual Reality Interaction of the Definitive Human Dataset

Avatar Training in Pharmacy and Dentistry

Facing and handling difficult conversations is a daily challenge within the medical profession with dentists, doctors, nurses and pharmacists having to successfully communicate with their patients. Such conversations may be related to trying to change and manage behaviours around healthy eating, drug abuse, alcohol abuse and even child protection. The conversation may be related to having to support a patient with their diet choices, denying a patient a requested drug or supporting a patient who is suffering from depression. Traditionally healthcare training in this area would include role-play scenarios but actors can be expensive and the training may differ from one day to another.

As an answer to this question, in 2012, SimVis in collaboration with our colleagues at NHS Education Scotland, created from tried and tested traditional approaches a new set of virtual scenario based training programmes (The Virtual Patient) to support education for care providers in the area of substance misuse (mainly drugs and alcohol) [3]. The key outcome was to produce immersive training scenarios that had been written by experts in the field. The student or healthcare professional can now experience the training as part of a repeatable tutorial or lecture-based training session. Alternatively, if preferred, they can download the application and experience the training from the comfort of their own home.

The storylines employed were carefully developed by NHS Education for Scotland (NES) and permit the user to take the role of the lead pharmacist or dentist. They are then able to have conversations with the patient and see the repercussions of decisions that they make. The user is subsequently able to return to specific decision points in order to investigate how a different response may have played out.

The Virtual Patient scenarios were developed using a motion capture system to acquire realistic human motion including complex human movement such as those associated with injecting drugs. As the software was developed in Unity, users may experience the training in an immersive VR environment as though they were the medical professional sitting at the table.

Figure 5 – Substance Misuse Virtual Training

Using The Virtual Patient program as a means of training postgraduate trainees to deliver pharmaceutical care to substance misuse patients has proven to be extremely successful both as a means of increasing confidence and knowledge. During the testing phase, over 80% of students preferred the avatar training system to traditional classroom based training.

Immersive Environments in Pharmaceutical Engineering

There are significant challenges around digitalisation in science, engineering and manufacturing. This is in part due to the complex nature of the data generated and the challenges in creating useful data sets with the scale required to allow big data approaches to identify patterns, trends and useful knowledge. Companies such as Facebook, Amazon and Google are extremely effective at incorporating the power of predictive data analytics into their operations. Unfortunately, much of the pharmaceutical manufacturing research and development community is lagging behind and struggle with modest, poorly interconnected datasets, which ultimately tend to have short useful lifespans.

A result of these poor, under-utilised datasets, is that it is largely impossible to avoid starting at the beginning of the process for every new drug that needs to be manufactured. Of course this is extremely costly with new medicines currently doubling in cost every nine years ($1 billion US Dollars currently funds only 50% of a new drug) [4]. It is clear that addressing this issue is key for the sustainability of the industry and future medicines supply.

ARTICULAR (ARtificial inTelligence for Integrated ICT-enabled pharmaceUticaL mAnufactuRing) is an exciting new collaborative project between SimVis and The University of Strathclyde's Institute of Pharmacy and Biomedical Sciences. Within ARTICULAR, we seek to develop novel machine learning approaches, a branch of artificial intelligence research, to learn from past and present manufacturing data and create new knowledge that aids in crucial manufacturing decisions. Machine learning approaches have been successfully applied to inform aspects of drug discovery, upstream of pharmaceutical manufacturing, where large genomic and molecule screening datasets provide rich information sources for analysis and training artificial intelligences. They have also shown promise in classifying and predicting outcomes from individual unit operations used in medicines manufacturing, such as crystallisation. For the first time, there is an opportunity to use AI approaches to learn from the data and models from across multiple previous development and manufacturing efforts and then address the most commonly encountered problems when manufacturing new pharmaceutical products. Specifically, the processes and operations to employ, the sensors and measurements to deploy to optimally deliver the product and the potential process upsets and their future impact on the quality of the medicine manufactured.

All of these data and the AI 'learning' are made available via specially created VR interfaces incorporating gesture and voice inputs alongside more traditional approaches such as dashboards. These immersive interfaces are already helping us understand the pharmaceutical manufacturing process design and complex data being captured in real-time. Detailed, interactive 3D visualisations of drug forms, products, equipment and manufacturing processes and their associated data are being created that provide intuitive access across the transformation from the drug molecule to the final drug product. This provides a unique tool, allowing the user to see their work and engage with their data in the context of upstream and downstream processes and performance data. Virtual and augmented reality technologies in the lab/plant environment visualise live data streams for process equipment that significantly improves user experiences providing data rich, customisable interactive visualisations to aid researchers in their work, allowing them to focus on the meaning of results and freeing them from menial manual data curation steps.

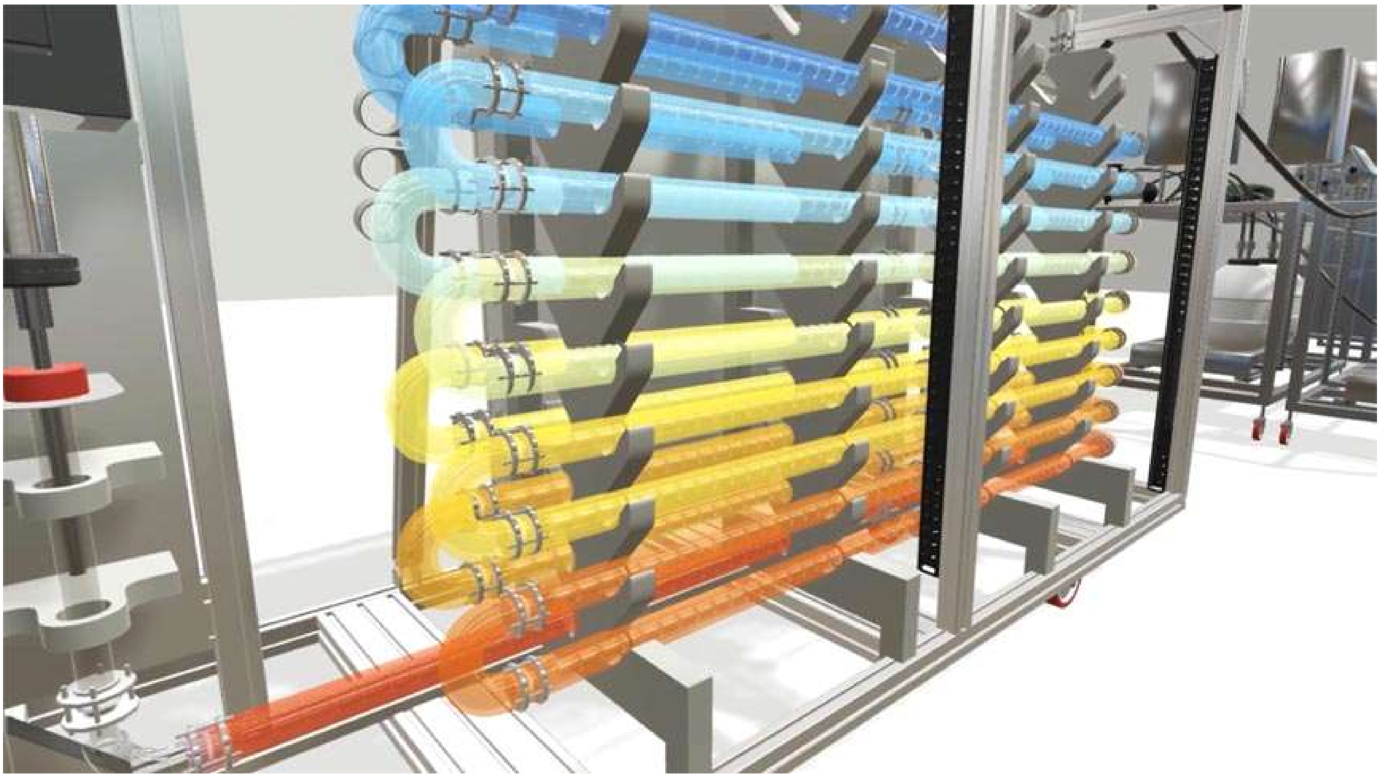

Figure 6 – Crystalliser digital twin showing temperature

Figure 6 shows an accurate, virtual reality 'digital twin' of a crystalliser that forms part of the pharmaceutical manufacturing pipeline. An accurate 3D digital model of the crystalliser and other lab equipment were created using laser scanning and photogrammetry techniques. These photogrammetric and point cloud datasets were then modelled up in 3D Studio Max and imported into the real-time games engine Unity.

Data relating to temperature and mass of the formed crystals can then be streamed in real-time from the crystalliser to a network and into the VR system which then applies appropriate colour palettes directly to the digital twin visualising either temperature or mass of forming crystals. Users can then walk into an accurate real-time VR model of their experiment, view their experiment's data and interrogate hardware in a more intuitive environment. In addition, they can be physically separated from their experiment, perhaps in another country, yet are still able to 'walk around' their laboratory and interrogate their results in real-time.

VR Training

This VR digital twin also provides an excellent training opportunity where we can educate pharmaceutical manufacturing students how to operate complex laboratory equipment without 'clogging' up the labs with training events. Remotely, for example, from the comfort of their own home, students can walk up to the equipment in VR and experience and interact with the hardware (Figure 7) as part of an immersive lab-based training programme. This has cost saving advantages as the bottleneck of training requirements at the start of term is a significant challenge to the university with hundreds of students needing training on complex equipment within their first term of study.

Figure 7 – VR interaction and training with pharmaceutical laboratory hardware

VR Dangerous Environments

Another area of interest for SimVis has been facilitating training within dangerous environments. For example, dangerous sports, offshore and hazardous decommissioning activities.

When the body and mind are under pressure, the ability to make good decisions can break down. Simple tasks that we may have repeated many times are forgotten in the blink of an eye when under duress. A good example is that of a trainee skydiver. During their training they may be shown several images and asked what to do under certain situations. For example, they may be shown a sequence of photographs of a parachute that has only partially deployed and they will be asked for their response. The problem is that viewing photographs or videos is not an immersive experience for the student and although the student may be able to answer the question correctly 10/10 on the ground. If they are in the unfortunate position of experiencing a real parachute failure then the added variables of falling to earth at 120mph, spinning uncontrollably and hearing the rushing of the wind in their ears results in total sensory overload. At this point they might not remember what they were taught on the ground and this is why a number of accidents occur with trainee skydiver and paraglider pilots.

One solution is to provide the trainee with more realistic immersive training before taking to the air, similar to a flying simulator for professional airline pilots. In 2008, the British Hang-gliding and Paragliding Association (BHPA) recruited SimVis and Hull University staff to develop the world's first virtual reality paragliding simulator.

An important aspect to developing a simulator requires the developer to have a very good first person understanding about what is being simulated. Consequently, as a development team, we participated in skydiving and paragliding training gaining all the necessary qualifications so that we were able to fly gliders and be sure that the simulation we were creating was as accurate as possible. Professional paragliding pilots and the British Hang Gliding and Paragliding Association (BHPA) also provided significant technical input into the project.

Figure 8 – World's first VR paragliding simulator developed by Chapman, Ward and Currer

The simulator allows paragliding pilots to experience what it is like to be at the controls hundreds of feet in the air but without leaving the ground. The pilot sits in a real paragliding harness, puts on a helmet (which also includes a VR HMD) and sees the landscape panning out in front of them through the VR headset. The simulator uses real paraglider controls so it actually feels that the user is flying. As well as steering, the pilot can 'weight shift' by moving their body in the harness. The simulation has modelled thermals so that a student can learn what it feels like to be caught in thermal currents or suddenly be confronted by the presence of another glider. Decisions around potentially life threatening situations can be safely resolved on the ground using the simulator.

It is imperative when creating simulations for dangerous activities that the model is as accurate as possible. For example, if an incorrect stall point on the VR paraglider simulator was built, then the trainee pilot will build an equivalent incorrect user model and muscle memory of the stall point. The more time they spend in the simulator, the more hard-wired and stronger their internal incorrect model will be. Flying a real paraglider may well result in the pilot stalling the wing due to them being pre-programmed with the incorrect setup.

Virtual Reality has enormous potential for supporting training in the areas of dangerous environments and health and safety training. Other SimVis VR projects include preparing engineers to walk around a decommissioned nuclear power station (Figure 9) that has no electrics and therefore no lights other than the light on the engineer's hardhat. After VR training, engineers have a basic idea of what the conditions will be like as they enter the pitch black power station.

Another example includes training for offshore workers in Trinidad in how to perform the extremely dangerous and complex procedure of filling a floating tanker with liquefied natural gas (LNG). SimVis developed a simulator to train offshore workers on how to control the arms that connect to the ship's manifold. The VR simulator enables complex procedures to be simulated, practised and monitored in the correct sequence. Future versions of the system will allow us to modify parameters related to weather and sea state.

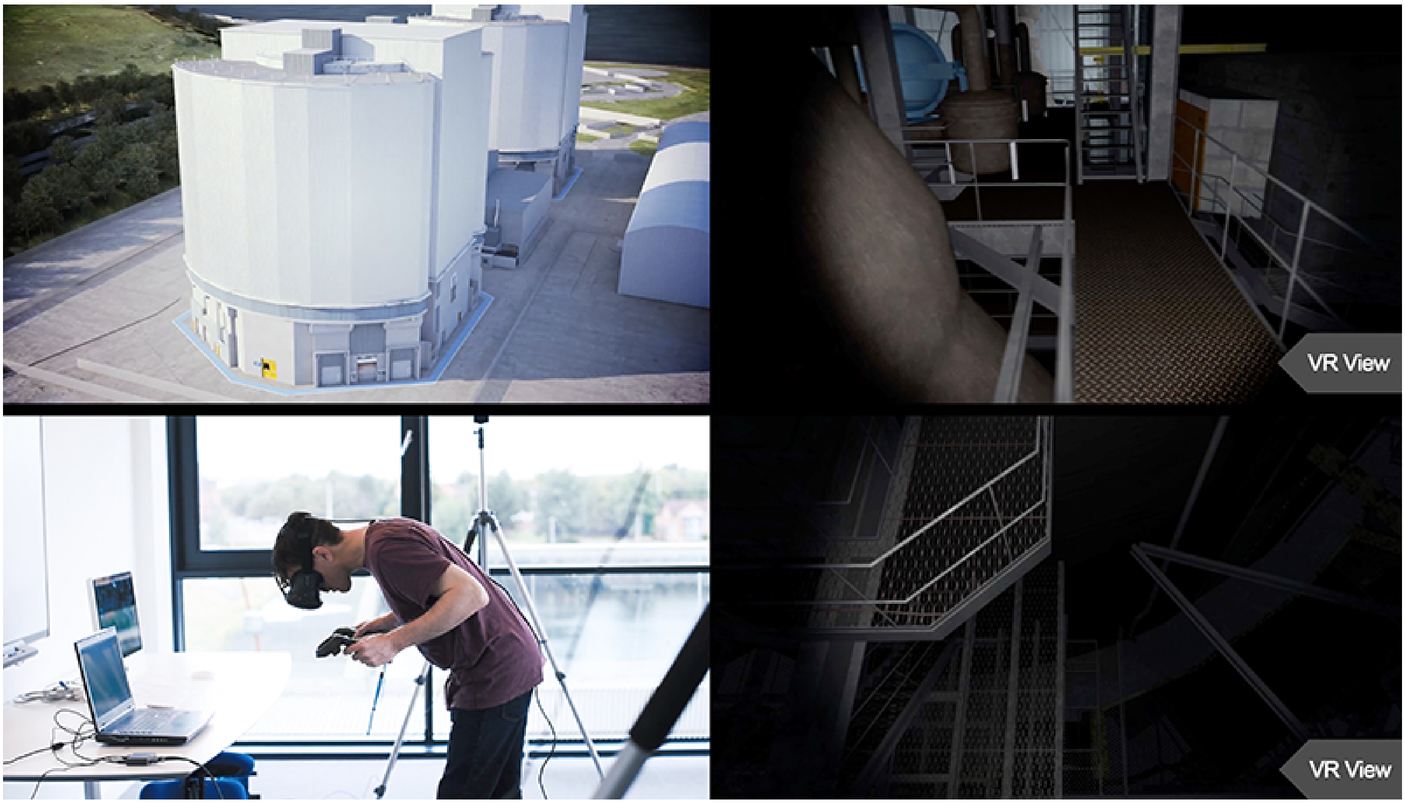

Figure 9 – VR Training for walking around a decommissioned nuclear power station

Art of the Point Cloud

The case studies described in this article have all benefited from the applied use of laser scanners and photogrammetry in order to generate the complex datasets that have been modelled (our human cadaver, lab equipment and nuclear power stations). There is however, a beauty to the raw point cloud data before surface reconstruction and the 'makeup' of photorealistic rendering is applied. Point cloud images blend the disciplines of the sciences and the arts, precision and measurement with aesthetics and creativity. Point clouds are a catalyst for wonder and enquiry, providing new perspectives across traditional views. Point clouds force us to wonder what we are looking at. Where and how was this image created? What is the story behind the data?

A point cloud is a set of data points usually displayed in an X, Y, Z Cartesian coordinate system. Each data point usually represents the surface of an object so the generation of point clouds is often associated with survey work where a detailed image of a surface is required. We use 3D scanning devices to generate accurate point clouds of, for example, a small archaeological artefact, the large ramparts of an ancient castle or perhaps a detailed surface of a patient's broken leg. A 3D scan will give us a permanent snapshot in time, a spatially accurate 3D photograph. Police can use 3D scans to revisit a complex crime scene to take important measurements previously missed. 3D scans of heritage sites can be used to create accurate digital models that can be shared worldwide in VR enabling people to virtually visit the sites, while medical 3D scans of patients can be used by doctors to plan complex brain surgery.

All these examples require different 3D scanning technologies to capture the point clouds. With advances in hardware and software, we are now able to generate extremely accurate and large point cloud datasets quickly and cost effectively.

How are point clouds generated?

There are a number of different scanning technologies available for creating point clouds. The most popular systems and techniques include laser scanning, photogrammetry, CT scanning and sonar.

Laser scanning

Laser scanning is one of the most popular methods for generating point clouds and is used extensively for surveying. 3D laser scanning (also known as high definition surveying or LiDAR) is a system that uses lasers to capture and measure real-world environments in 3D. Laser scanners emit a laser pulse that reflects back from the surface of interest. During the receiving stage, the scanner will either use time of flight (TOF) or a phase-based system which compares the phases of the output signal and the return signal to calculate the distance the beam has travelled. This process is performed up to a million times a second resulting in an extremely dense point cloud representing the surfaces of the environment. The beauty of raw laser scanned data can be seen in Figure 10.

Figure 10 – Laser Scan Point Cloud: Strange Fruit Scott Page, Oakland, California 2018

Photogrammetry

Structure from motion photogrammetry is a technique for estimating 3D structures from 2D image sequences. By processing a large number of high resolution photographs of the object of interest (taken from various locations), 'rays' can be calculated from each camera to points on the object. These rays are then mathematically intersected to produce the X, Y, Z coordinates making up the point cloud. Recent advances in computer processor power and digital cameras have resulted in an increase in the quality and popularity of photogrammetry to generate 3D models.

Sonar

In a similar manner to laser scanning (timing the return signal of light), we can time the return signal of sound waves to measure the depth of the seabed. A sonar system will measure the length of time it takes for sound to travel from a boat-mounted transducer to the sea floor and back. We are then able to use time measurements of the return echo signal to understand the depth of the water and generate a number of depth points (our point cloud) representing the seabed or objects on the seabed. We often use the term ‘bathymetry’ to refer to the depth of water. In the same way that topographic maps represent the 3D features of overland terrain, bathymetric maps precisely visualise the topography of the seabed and objects that lie underwater. An example of a bathymetric dataset generated from sonar is the beautiful point cloud image of the sunken SS Richard Montgomery and its surrounding seabed (Figure 11).

The image was created from 4.8 million sonar soundings of the wreck of the SS Richard Montgomery, which ran aground in the Thames Estuary, near Sheerness, in 1944. It was carrying 1,400 tonnes of explosives at the time and, because of the unexploded ordinance in its hold, the vessel remains a hazard to shipping to this day. The colours indicate depth and highlight the scouring of the seabed by the tide around the wreck. The main shipping channel can be seen at the bottom right, surprisingly close considering the ship’s cargo

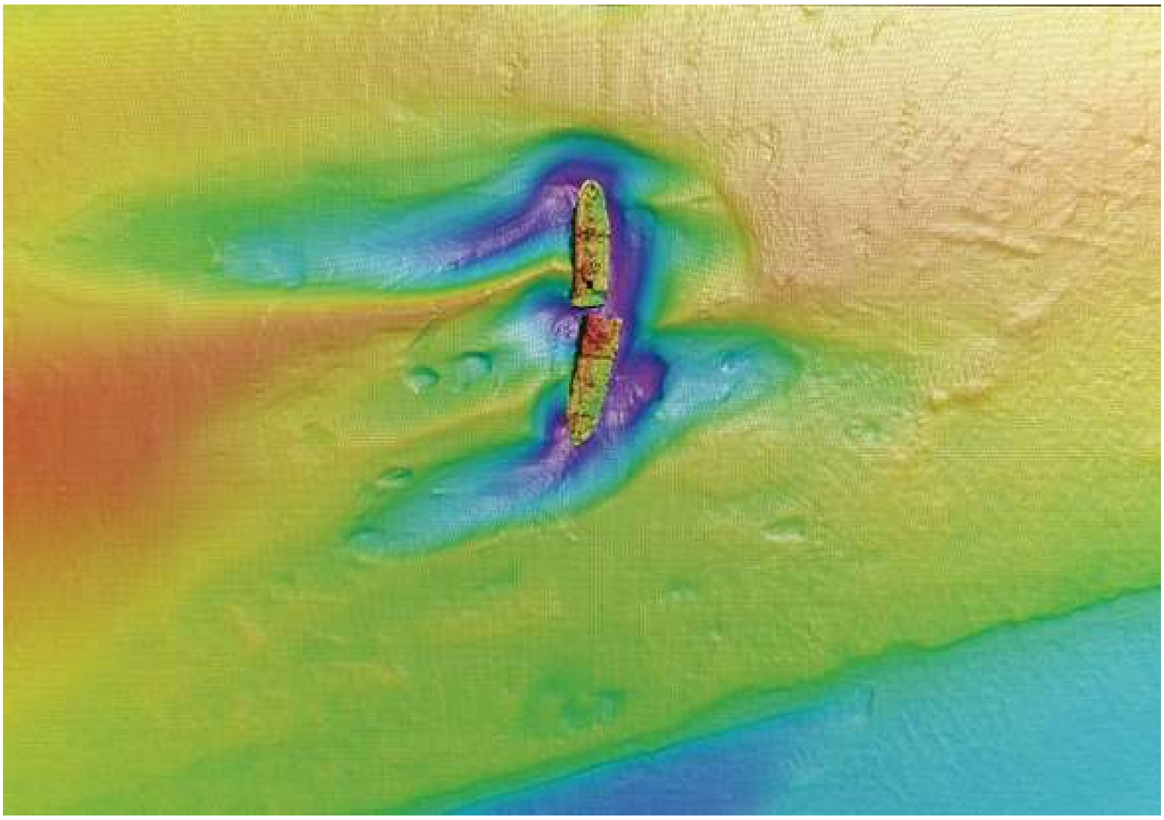

Figure 11 – Sonar Point Cloud: The SS Richard Montgomery

CT Scanning

3D scanning technology is used extensively in the medical domain. CT (or CAT scans) produce cross-sectional images of the body using x-rays. CT has become a mainstay for imaging the human body and is used extensively to diagnose conditions such as damage to bones and organs. MRI (magnetic resonance imaging) scanners are also used extensively in the medical domain and produce internal 3D images but use a powerful magnet and radio waves rather than the CT’s x-rays. MRI scanners give higher detail in soft tissues whilst CT scanners tend to be better for imaging bone structures.

These scanning technologies all provide us with high-resolution data but the surprising bi-product is the extraordinarily beautiful images created as part of the process. Our recent book The Art of the Point Cloud [5] is a celebration of some of the most beautiful and striking 3D scans from around the world and uses all the scanning technologies described above.

When these datasets are further coupled with immersive technologies such as VR, we have yet another methodology for interaction and creating new pieces of art. Users can now get 'inside' datasets and their artworks. In addition, with software such as Google's Tilt Brush, Artists can now move away from the traditional 2D planar canvas and experience painting in three dimensions. Literally painting around themselves.

Summary

Virtual Reality has finally come of age. We are finally able to move away from the gimmick installations and 'dragon demos' that have become ubiquitous at technology trade shows. There are now real applications of the technology available that can create tangible benefits to end users. VR's absence from Gartner's Hype Cycle for 2018 demonstrates a maturity of VR that is evidenced by industry's rapid and increasing adoption and application of the technology to real world problems. SimVis has successfully implemented VR for medical training, pharmaceutical manufacturing, dangerous sports and hazardous work environments. In addition, we have looked at how scanning technology can generate stunning images that can be viewed both traditionally and in immersive environments. The next five years will see an explosion of higher quality immersive helmet mounted displays that will provide increases in resolution and field of view as well as improved ergonomics. This is only the beginning for virtual reality.

Acknowledgements

Thanks to all the staff at the School of Simulation and Visualisation at The Glasgow School of Art. Dr Blair Johnston and Professor Alastair Florence at Strathclyde University. I would also like to thank NES, RCSEd and the Scottish Funding Council.

REFERENCES

[1] Ratcliffe S, ed. Roy Amara 1925–2007, American futurologist. Oxford Essential Quotations (4th ed.). Oxford University Press. 2016.

[2] Cearley D, Burke B, Searle S, Walker M. Top 10 Strategic Technology Trends for 2018: A Gartner Trend Insight Report, Gartner Publishing, March 2018.

[3] Zlotos L, Power A, Hill D, Chapman P. A scenario based Virtual Patient program to support substance misuse education: a pilot in pre-registration pharmacist training, American Journal of Pharmaceutical Education. 2016, 80:3.

[4] DiMasi JA, Grabowski HG, Hansen RA. Innovation in the pharmaceutical industry: new estimates of R&D costs. Journal of Health Economics 2016, 47:20-33.

[5] Chapman P, Mitchell D, McGregor C, Wilson L, Rawlinson A. Art of the Point Cloud, Wild Harbour Books, 2018.