The Use of Augmented Reality in the Operating Room: a Review

John Bartlett*
Department of Electrical and Computer Engineering,
University of British Columbia

ABSTRACT

This paper presents a survey of the use of augmented reality (AR) in medical applications, with special focus on the operating room. After introducing the motivation for this research, we present an overview of the current status of AR research to provide a foundation with which to understand the remainder of the paper. The following section discusses the application of AR to medicine, and also elaborates on the use of colour in medical imaging. The next section deals exclusively with minimally invasive surgery, presenting some of the achievements in AR systems as well as their limitations. Next we consider volume visualization, which is relevant here because so much of the medical data that would be useful to display is volumetric. The paper concludes that AR has a vast array of applications to different areas of medicine, and while many of them still need refinement, the outlook is good for clinical integration in the future.

Keywords: augmented reality, surgery, medical imaging, information visualization, volume visualization, minimally invasive surgery, laparoscopic surgery

1 INTRODUCTION

Although medical imaging technology has existed for over a century [36], its use is still mostly limited to diagnostic purposes. Technological advancements in the past few decades have led to images with higher contrast, better spatial resolution and less noise, but these are often pushed aside when it is time to perform surgery.

This is not to say that surgeons don't make use of medical imaging. On the contrary, preoperative planning allows the surgeon to formulate a strategy for the procedure and anticipate potential problems. However, it would be advantageous if this information could be presented not simply as a reference, but mapped onto the surgeon's field of view to provide an augmented reality interface (see Figure 1). This would enable the surgeon to rapidly locate important structures during time-critical stages of the operation while avoiding sensitive nerves and vessels.

This is where information visualization work is needed. Scientists and engineers have developed the necessary tools to acquire and process images, but we need to answer the question of how to best present this information in the operating room. A solution to this problem will require extensive cooperation between the medical imaging and infovis communities. The goal of this survey is to consolidate the knowledge relevant to this field and bridge the gap between problem-driven and technique-driven research. This report attempts to provide a complete perspective on how augmented reality is used in surgery today and what challenges it faces in the future.

Figure 1: The problem that we are trying to solve with augmented reality in the operating room: good medical information is available, but not in the most convenient format [35].

*E-mail: johnb@ece.ubc.ca

2 AUGMENTED REALITY

Augmented Reality (AR) enables users to perceive virtual objects. It differs from Virtual Reality (VR) in that users are not completely immersed in a virtual world, thus they perceive a hybrid of the virtual environment and reality (see Figure 2). A true AR system should also register virtual and real elements in 3-D and be interactive in real time. AR is advantageous because the virtual elements may enable the user to perceive information that is not directly visible in the real world. These techniques have been applied in the domains of maintenance and repair, annotation, robot path planning, entertainment, military aircraft navigation and medical visualization [1].

Figure 2: Virtual lamp and chairs are integrated in a real environment featuring a real table and phone in a simple AR example [1].

2.1 Augmented Reality Beyond Vision

Although AR is generally seen as a visualization tool, the strategy of enhancing our perceived reality can be applied to all senses. For example, headphones and microphones can be used to integrate sounds from the room with virtual sounds. AR may also be used to enhance our sense of touch. Haptic AR systems keep track of virtual objects and can provide haptic feedback through tactile effectors when the user encounters a virtual interface [1].

2.2 Projection Displays

AR systems with projection displays are designed to project the virtual information directly onto the real environment. This task is simplest when the real environment consists of planar surfaces that face the projector, but more complex surfaces can be effectively augmented using multiple calibrated projectors. Projection displays are an example of Spatial Augmented Reality (SAR), and have the advantages of scalable resolution, better ergonomics and a theoretically unlimited field of view. They can also reduce eye strain as the virtual objects are rendered at their real depths [7].

Head-mounted projective displays project images from the headgear onto retro-reflective surfaces in the environment, which reflect light back to the user [2, 7] (see Figure 3). These systems provide a large field of view and avoid parallax distortion errors that can be caused by a mismatch in inter-pupil distance. However, the head-mounted apparatus is cumbersome, resolution is limited and a certain amount of control over the AR environment is required [7].

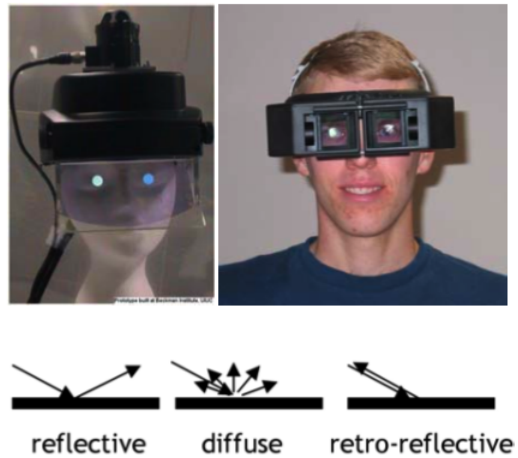

Figure 3: Examples of head-mounted projective displays (top) and diagram showing reflective properties of retro-reflective surfaces (bottom) [7].

2.3 Head-Mounted, Monitor and Handheld Displays

The dominant method for enhancing the user's visual perception of reality is to use a head-mounted display (HMD). This involves placing a visor directly in the user's field of view. Unfortunately, the proximity of the screen to the eyes limits the number of pixels and hence the resolution [7]. Modern devices are designed to be as light and unobstructing as sunglasses [2] (see Figure 5(c)). Monitor-based systems fall under the category of SAR, as the display is fixed in space and separated from the user [1]. Since interaction with the augmented real world is difficult and graphics can only be superimposed on a limited area of the user's field of view, the degree of immersion is limited [7]. Handheld displays can be seen as intermediate between HMDs and monitor systems (see Figure 4). These displays have the advantage of extending the effective field of view as they are moved around, but the user also has one less hand free to complete their task [2, 7].

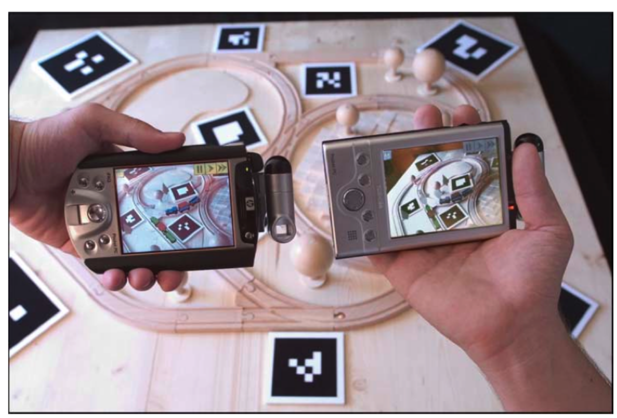

Figure 4: "The Invisible Train Game" allows multiple users to augment a scene using handheld displays [33].

These displays come in three main varieties: optical, transparent and video.

2.3.1 Optical Displays

An optical display has a see-through screen composed of optical combiners. The combiners are partially reflective, so that virtual images can be projected onto the user's field of view without completely obscuring the real world beyond the screen (see Figure 5). One downside of this technology is that the combiners necessarily reduce the amount of light that reaches the user from the real world. This effect can be reduced by combiners that selectively reflect the wavelengths of light from the virtual display and let other light through [1].

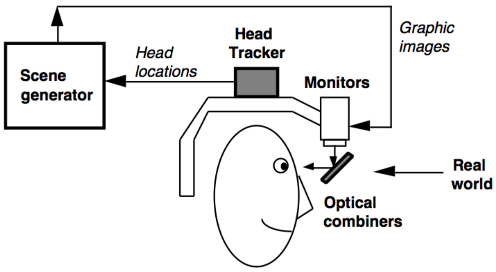

(a) diagram of a typical apparatus

(b) example of an optical HMD

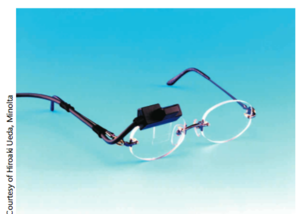

(c) modern HMDs are designed to be light and unobtrusive

Figure 5: Optical see-through head-mounted displays [1].

2.3.2 Transparent Displays

Transparent displays behave similarly to optical displays but they have no optical combiners. These systems can either use an LCD screen with no opaque backing or a simple transparent screen with a projector. Transparent displays generally have the same tradeoffs as optical displays [7].

2.3.3 Video Displays

The video approach to HMDs involves a closed-view video display that incorporates the real world by blending in images from head-mounted video cameras (see Figure 6). Likewise, AR video monitor displays blend a real-world video stream with virtual graphics on a monitor. In some cases, the user can use stereo glasses to allow depth perception [1].

The blending may occur by simple chroma-keying, which means replacing a monochrome virtual background with the real-world video, or by real-world placement of virtual objects using 3-D coordinates. The latter is preferable because it allows real objects to occlude virtual ones, but it requires extensive depth information [1].

2.3.4 Comparing Optical and Video Displays

Advantages of Optical Displays

- Optical displays are advantageous over video displays in their simplicity. They only have to deal with one video stream, as the real world is perceived directly. In contrast, video displays must digitize the real images, correct for distortion in the camera and then combine the real images with the virtual images computationally. This is more expensive both in terms of time and equipment cost.

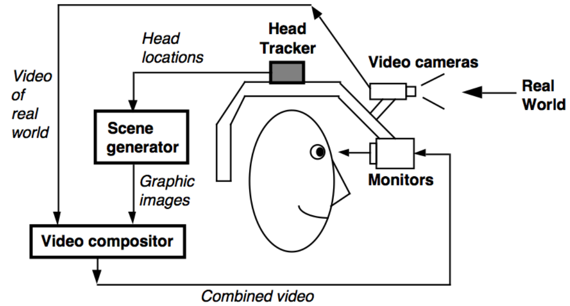

(a) diagram of a typical apparatus

(b) example of a video HMD

Figure 6: Video head-mounted displays [1].

- Real-world resolution is better in optical displays because it does not have to be represented with pixels as in the video case.

- Optical displays also avoid the video display issue of eye offset, which results from camera streams taken from a different position than the user's eyes.

- Finally, these displays reduce safety concerns by allowing the user to see even when the AR system is powered down [1].

Advantages of Video Displays

- One of the greatest advantages of video displays is that the combination of real and virtual video is entirely up to the designer. In optical systems, virtual objects cannot completely obscure real objects because the screen is semi-transparent. This may reduce the level of immersion in optical systems [1]. It should be noted, however, that some modern optical systems have introduced an LCD panel on the outside of the optical combiner, whose pixels can be selectively opacified to allow occlusion of real objects [2, 7].

- It is much easier to implement a wide field of view in video systems. Distortion increases as the user looks away from the centre of the view, but this can be corrected digitally in video systems. In optical systems, distortion of virtual objects cause registration issues with the undistorted view of the real world.

- Video displays also offer the luxury of delay matching, allowing movement in the real and virtual views to be synchronized.

- In optical displays, contrast can be an issue, as the eye can perceive a much greater range of intensities than we are able to create on a screen. Video systems can correct for this by controlling the range of intensities from the real-world stream.

- Lastly, the digital video information from the real-world can be used to improve registration or perform other types of analysis not possible with optical systems [1].

2.4 Retinal Displays

More recent research has led to the emergence of virtual retinal displays, which draw images directly on the retina with low-powered lasers (see Figure 7). These displays allow brighter, higher-resolution images with a wider field of view [7].

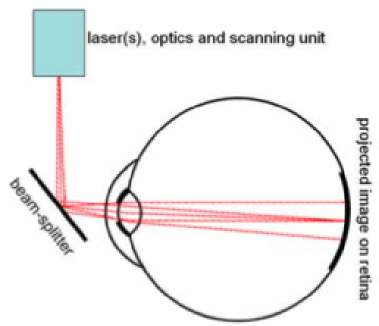

Figure 7: Simplified diagram of a retinal display [7].

2.5 Registration

Image registration or alignment is critical to preserve the illusion of AR. Accurate registration is even more important in medical applications because the registration may directly influence what tissue is cut or sampled. Errors in registration may occur due to optical distortion, tracking errors, discrepencies in system mechanics, time delays between the perceived motion of virtual and real objects, or calibration of the viewing system. In some cases, computer vision techniques have been used with video-based matching or fiducial markers to correct for registration error [1]. Since this paper aims to focus on the visualization aspect of AR, we will omit an in-depth discussion of the registration problem.

2.6 Sensors

Except for rare systems that deal only with a static real environment, AR implementations require sensors to receive a variety of inputs from the environment. Trackers are particularly important, especially when accurate registration is desired. Other sensors may be used to allow users to perceive information that cannot be detected without the aid of a machine [1].

Common varieties of trackers include mechanical, magnetic and optical models. Often, more than one type of tracker is used in an AR system to improve performance [2].

2.7 Tangible User Interfaces

One of the great advantages of augmented reality systems is the possibility of elaborate tangible user interfaces (TUIs). A TUI maps the manipulation of a physical input device to an intuitive interaction with the virtual data. These interfaces can be either space- or time-multiplexed. In space-multiplexed interfaces, a unique input device occupying its own space is tied to each function. In contrast, input devices in time-multiplexed systems can serve different purposes at different points in time, as is the case for the computer mouse. Space-multiplexed systems are generally faster to use because the user does not have to keep track of the function of each input device as it changes with time [6].

2.8 Information Visualization in AR

Once all the AR hardware is in place, the format in which virtual information is displayed must be determined. How should virtual depth be represented on a 2D display? How much virtual information can be displayed without adversely affecting what can be perceived from the real world? Is photorealism of virtual objects desirable? The answers to these questions will depend on the system being studied [2].

2.8.1 Depth Perception

A means of accurately conveying depth information of virtual objects to the user is essential for an AR system to provide a useful and believable depiction of an environment. However, this task is made difficult by the fact that display screens in non video-based AR systems are often at a fixed depth in front of a 3-D world. It has also been shown that users tend to consistently underestimate depths when using HMDs, although there is no consensus on the explanation for this phenomenon. Even more difficulties are encountered when trying to determine how to portray “x-ray vision”, where normally occluded objects are made visible by AR with purposely conflicting depth cues [16].

The ten generally accepted depth cues are binocular disparity, binocular convergence, accommodative focus, atmospheric haze, motion parallax, linear perspective and foreshortening, occlusion, height in the visual field, shading, and texture gradient. The relative importance of each cue varies with respect to scene structure and lighting, as well as the distance from the user. AR depth judgement experiments have shown that depth perception of real objects is better than that of virtual objects, and that accuracy decreases at greater distances. Experiments with an occluding real surface, often referred to as the “x-ray vision” condition, have resulted in underestimated depth judgements that improved when the occluding surface was removed [16] (see Figure 8).

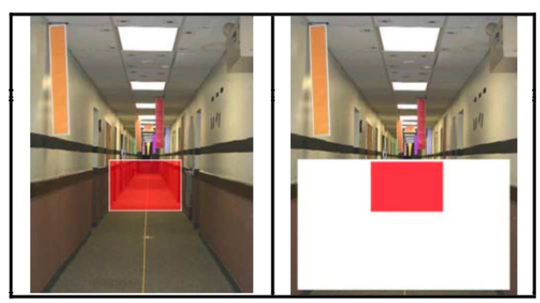

Figure 8: In a study by Jones *et al.*, participants were asked to manipulate the depth of a virtual rectangle to match a real reference object of the same colour, with and without the presence of an occluder. Depths were consistently underestimated, but accuracy improved when the occluding surface was removed [16].

2.8.2 Information Filtering

In many AR applications, more information is available than the user can perceive at once. In these cases, information filtering can be used to reduce clutter. Filtering generally requires the virtual data to be annotated to allow such interactions as object selection, semantic zooming and cutaway visualization. The AR MagicLens is a handheld device designed to make information filtering intuitive [17, 6] . In some systems, multiple lenses are used to perform several augmentations separately or in combination. This technique is well suited for segmented and categorized data [21]. An example is shown in Figure 9.

Figure 9: Example of information filtering using AR MagicLenses, where arteries and veins in the liver are coloured red and blue, respectively, while the parenchyma is made semi-transparent [21].

2.8.3 Photorealism

The requirement for rapid rendering of AR scenes means that simulation of believable virtual objects is difficult. The techniques used to model lighting and shading of virtual objects are often simplified and may not take into account the actual properties of the room. Studies have shown that the photorealism of virtual objects can be improved by taking into account the lighting of the scene and adding shadows where appropriate. In other applications, photorealism may be discarded by applying stylization to virtual and real objects alike in an effort to increase immersion. Stylization involves edge detection and colour segmentation to blur the lines between virtual objects and their environment (see Figure 10). This may be useful in AR applications where immersion takes precedent over realism [13].

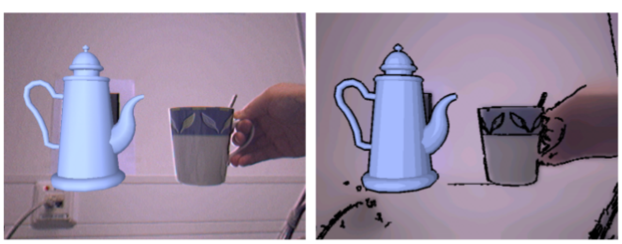

(a) standard AR implementation (b) with stylization

Figure 10: Example demonstrating the use of stylized AR to improve immersion of a virtual teapot in its environment [13].

3 Augmented Reality in Medicine

AR is clearly a powerful tool, and there are many scientific domains that could benefit. Medicine is particularly interesting as there are already highly-refined techniques for imaging the human body, but no convenient way to display this information where it would be most useful: in the physician's view of the patient.

An early medical AR technique, described in a paper by Lorensen *et al.* [18], involves registering 3D imaging data with a live video feed of the patient for surgical planning. This study was successful in that it allowed the surgeon to map the location of a brain tumour onto the patient's head before beginning the procedure (see Figure 11). However, this information was presented on a screen and not directly in the surgeon's field of view, so the surgeon had to look back and forth to make comparisons between the augmented reality interface and the patient.

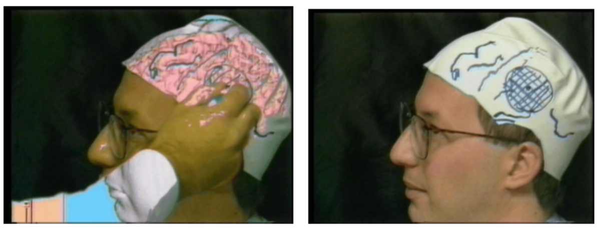

Figure 11: Early experiment using AR with luminance keying to superimpose a computer-generated model of the brain on a volunteer to facilitate pre-operative planning [18].

Schwald *et al.* [25] describe an augmented reality implementation that they named the AR Window, a semi-transparent display that projects medical imaging data directly into the surgeon's field of view (see Figure 12). Where similar devices use a half-silvered mirror to ensure the correct perspective, the AR Window uses eye-tracking technology. The paper serves mainly as a proof of concept, so more research is required to determine whether the display's usefulness outweighs its obstructive interface.

Figure 12: The AR Window is a semi-transparent display that uses eye-tracking technology to superimpose virtual image data on the surgeon's field of view [25].

Fuchs *et al.* [14] developed a tracked head-mounted system with 3D visualization that was able to enhance the surgeon's natural point of view and preserve motion parallax (see Figure 13). Currently, depth acquisition is too slow for surgical integration of this device to be feasible. The paper is also lacking a discussion of the practicality of the head-mounted display with respect to limitations on peripheral vision and restriction of motion due to bulkiness.

Figure 13: Tracked HMD used in study by Fuchs *et al.* [14].

3.1 Colour in Medical Imaging

Medical images were generally colourless until the appearance of real-time colour Doppler ultrasound imaging in the 1980s. The role of colour is now becoming more prominent with the increasing scale of medical image data and the emergence of more complex visualization techniques [23].

3.1.1 Early Use of Colour

One of the first uses of colour in diagnostic medical imaging occurred in the early days of x-ray technology, when radiographs taken with different x-ray energies were visualized in different colours on a single image (see Figure 14). Despite the extra information that was provided, this application never achieved widespread use, partly due to the fact that it required a higher dose of radiation than a standard radiograph [23].

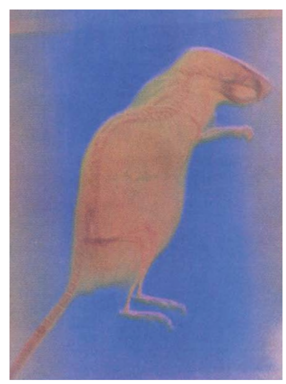

Figure 14: Colour image of a mouse using information taken from radiographs made at 40, 60 and 80 kV [23].

Colour Doppler ultrasound was introduced in the early 1980s. This technique uses the Doppler frequency shift of sound waves echoing off of moving blood to visualize blood flow. Red is used to represent flow towards the probe, blue for flow away from the probe, and other colours can be used to indicate the turbulence of the flow (see Figure 15). Saturation values are mapped to the magnitude of the phase shift, which corresponds to the speed of the moving blood. This innovative visualization strategy opened the door for the use of colour in other imaging domains [23].

3.1.2 Proliferation of Colour in Modern Imaging Techniques

The next important application of colour came with the introduction of functional magnetic resonance imaging (fMRI) and MR diffusion tensor imaging in the 1990s. These procedures use MRI

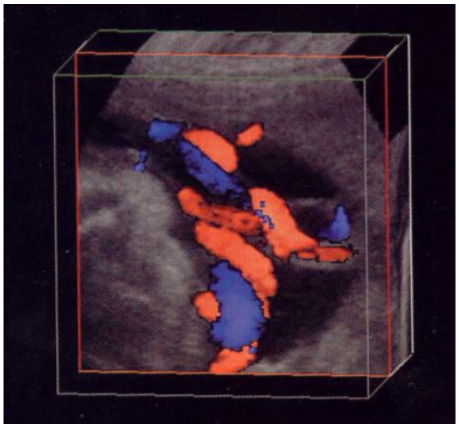

Figure 15: Colour Doppler ultrasound image showing blood flowing towards the ultrasound probe in red and blood flowing away from the probe in blue [23].

technology to visualize brain activity and brain water diffusion, respectively. The resulting images employed colour to highlight areas of interest within the larger image (See Figure 16). Colour-coded 3D renderings of anatomical data also became feasible around this time, but limited computational power made this a time-consuming process at first [23].

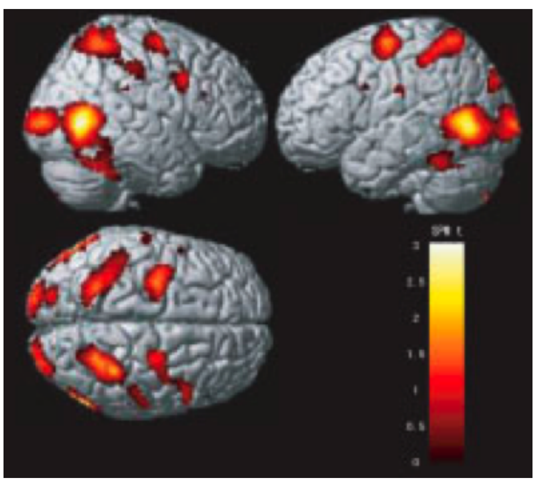

Figure 16: fMRI images use a colour scale to display brain activity, based on variations in blood oxygen level [22].

3.1.3 Combining Multiple Modalities

It may be in the physician's interest to combine multiple images from different modalities in order to achieve a more information-rich visualization. For example, CT images provide excellent bone contrast, while MRI is ideal for visualizing soft tissue [36]. Tracking systems or image registration algorithms can be used to align images taken using different modalities so they can be displayed together using some colour mapping. Another example is PET/CT, where CT images are enhanced with data from positron emission tomography. In this case, the CT data presents the underlying anatomy while PET indicates regions of high glucose uptake, indicating highly metabolic regions that may contain cancer cells [23] (see Figure 17).

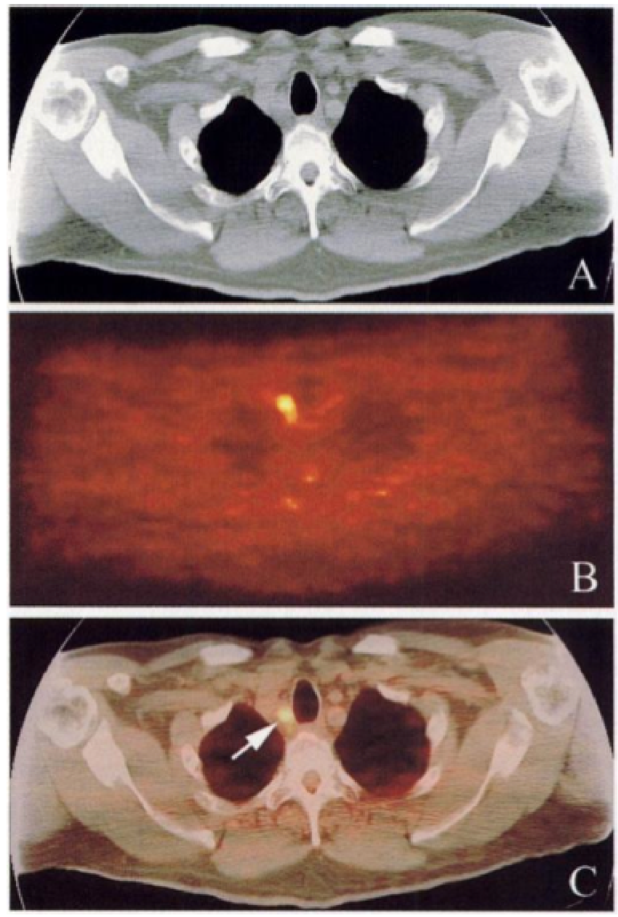

Figure 17: A PET/CT image allows rapid tumour localization. In (A), the CT image provides good resolution of the anatomy but the tumour is difficult to locate. (B) clearly indicates a region of high metabolism, but it is hard to place this information in an anatomical context. A fused PET/CT image is shown in (C), making the location of the offending lymph node immediately apparent [3].

3.1.4 Colour vs. Monochrome Images

Research in human perception has shown that continuous colour shifts improve a viewer's ability to distinguish between subtle shifts in intensity [23]. There is also some evidence to suggest that at low contrast, luminance and chromatic contrast may be processed independently by the brain. This suggests that redundantly coding the image intensity with both luminance and colour could improve a viewer's understanding of the data, perhaps highlighting low spatial frequency information with colour while the luminance contrast allows effective perception of data with high spatial frequency. However, it should be noted that colour coding has also been shown to distort the apparent shape of objects in some cases [10]. It is also unclear if the benefits reported in human perception experiments are as significant in practice. Hwang *et al.* [15] showed that B-colour ultrasound images provided similar accuracy to traditional greyscale images when evaluating left ventricular systolic function in coronary artery disease (see Figure 18).

Whatever the reason, colour enhancement of monochrome images was never truly accepted by the medical community. It appears that standard monochrome images are generally preferred by radiologists except where colour is being used to encode a separate variable [23].

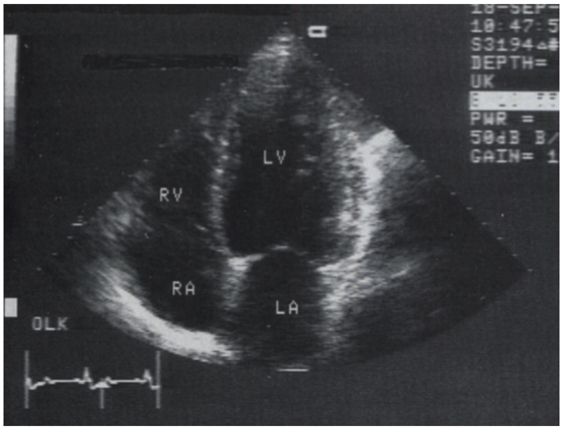

(a) greyscale image

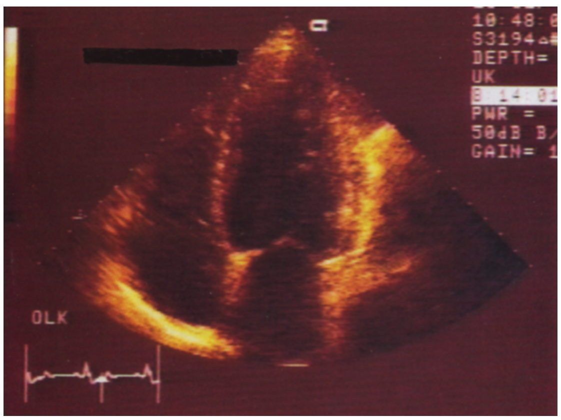

(b) "temperature" colour map

Figure 18: Visualizing an ultrasound image showing a four-chamber view of the heart with and without the use of colour [15].

3.2 AR for Needle Insertion

When needle insertion is required, for example in cases of biopsies or thermal ablation, image guidance is often used to avoid injury to sensitive structures such as nerves. Computerized tomography (CT) has been integrated into several commercially available guidance systems for this purpose [11]. Magnetic resonance imaging (MRI) has especially good soft tissue contrast, but the machines are usually too bulky to allow needle insertion during imaging. In some cases, the patients are required to enter and exit the imaging system repeatedly in order to ensure correct needle placement [32].

3.2.1 MRI-based AR-guided Needle Insertion

Wacker *et al.* [32] developed a system that made use of an augmented reality HMD to overlay the MRI data and an augmented needle on the physician's field of view during needle insertion outside of the MRI machine (see Figure 19(a)). The display consists of a stereoscopic view of the enhanced needle and patient with the target areas highlighted. The enhanced needle has a 7-cm virtual cylinder extending beyond its tip to allow alignment with the target before it enters the skin (see Figure 19(b)). An optical tracking system was used to correctly position the patient data and the virtual needle on the field of view. This technique was shown to be useful in phantom and animal experiments, but the authors noted that the requirement for patient immobilization presented some limits and challenges. For example, breathing motion correction may need to be applied when operating on some parts of the body.

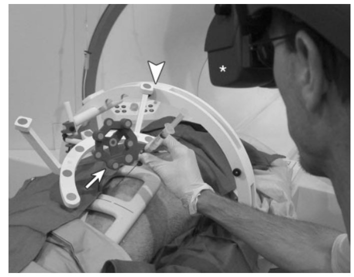

(a) An HMD is used to visualize the augmented scene.

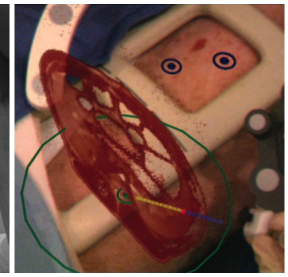

(b) The tip of the needle is virtually highlighted and extended, and MRI data is overlaid.

Figure 19: This AR-enhanced needle system made use of MRI data after the patient was removed from the machine [32].

3.2.2 CT-based AR-guided Needle Insertion

Members of this same group also implemented a CT-based AR approach to needle insertion [11]. Although CT-guided interventions are common, they have the drawback of high radiation exposure to the patient and medical staff, and the dose increases with the length of the procedure. The proposed technique used an HMD with cameras and optical trackers to overlay recently recorded CT data on the physician's field of view along with an enhanced needle. This system allowed rapid needle placement with error of just a few millimetres, but given the restrictions on patient movement and the increased complexity of the apparatus, it will likely be reserved for time-consuming or difficult procedures on parts of the body that can remain stationary for the duration of the operation.

3.3 Fluorescence-guided Procedures

Near-infrared (NIR) light is an attractive option for encoding information in the OR because it is invisible, safe, deeply penetrating and target-specific. Currently, radioactive tracers are used for target-specific mapping, but at the cost of a longer procedure, exposure to radiation, and a steep learning curve.

Tanaka *et al.* [28] implemented an intra-operative NIR fluorescence imaging system using a dichroic mirror that separates the NIR wavelengths from the visible light and directs it to an NIR-sensitive camera. A colour representation of the received NIR light can thus be superimposed on the video feed from the visible-light camera. Ten high-powered light-emitting diodes provide the NIR excitation light (see Figure 20).

This system allowed real-time visualization of the tracer injection in a porcine model, followed by identification of the sentinel lymph nodes (SLNs) within 15 seconds. The image guidance allowed dissection of non-nodal tissue to isolate the lymph nodes. The nodes appear to be larger and blurred on the tissue surface due to scatter but their location can still be determined (see Figure 21) [28].

This technique is now past the pre-clinical development stage, but further testing will be required to ensure it meets all regulations before it is implemented in a hospital setting. The authors also noted that while the approved, organic tracers proved effective, some inorganic and hybrid tracers may be even better if they pass toxicological screening [28].

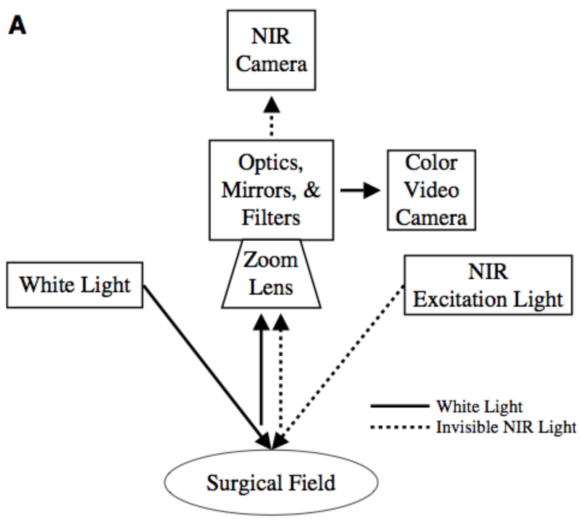

Figure 20: Diagram of the NIR fluorescence imaging system used by Tanaka *et al.* [28]. See text for details.

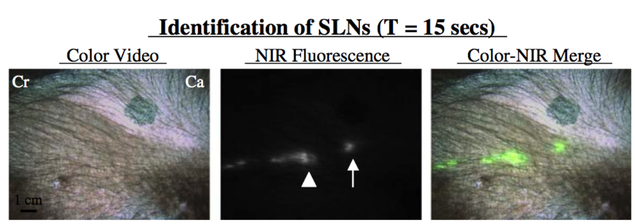

Figure 21: NIR fluorescence was able to identify 2 SLNs after 15s. The live video is shown on the left, the NIR fluorescence image in the centre, and the pseudo-coloured combination on the right [28].

4 Minimally Invasive Surgery

Minimally invasive surgery (MIS), also known as endoscopic surgery [12] or keyhole surgery [5], refers to procedures where the surgical instruments are inserted into the patient through small ports and the surgeon is guided by an endoscopic camera. Laparoscopy is a term specific to the abdominal region.

4.1 Advantages of Minimally Invasive Surgery

MIS is advantageous in that it leads to a reduction in surgical complications, operating times, and patient recovery times. However, this strategy limits what the surgeon is able to see and increases the difficulty of the procedure. What is needed is a way to increase the amount of information that is visible to the surgeon, using data from medical imaging modalities such as CT, MRI and ultrasound [34].

In addition, robotic minimally invasive surgical systems have been developed to aid surgeons with the transition to more difficult surgical procedures. These systems allow greater freedom of motion, better precision and improved hand-eye coordination relative to manual laparoscopy. Furthermore, many medical robots, such as the da Vinci™ surgical system, track the motion of the camera and surgical instruments in their own coordinate system and provide an Application Programming Interface (API) to integrate this information with other systems in the OR [34] (see Figure 22).

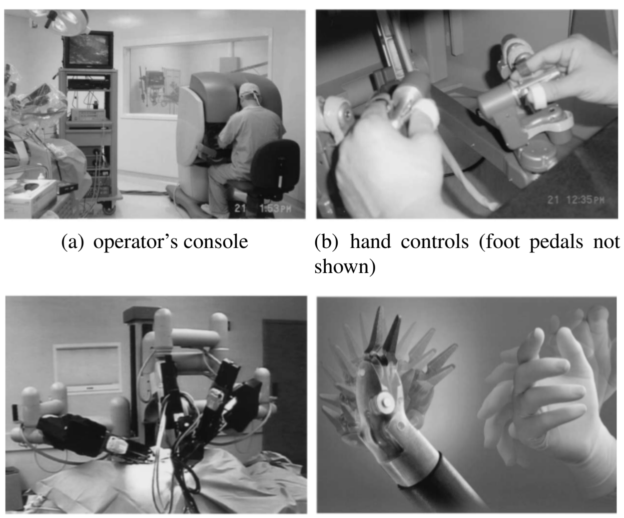

(a) operator's console (b) hand controls (foot pedals not shown) (c) robotic arms (d) range of motion

Figure 22: Assorted views of the da Vinci™ surgical robot [26].

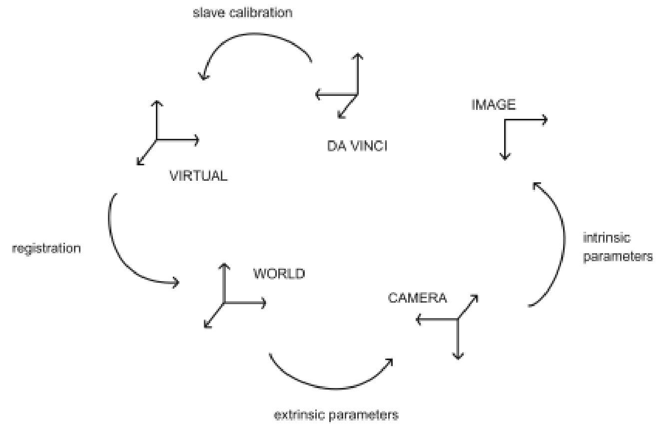

Figure 23: Diagram illustrating a typical registration problem when designing an AR system that works with a da Vinci™ surgical robot to overlay medical imaging data on the surgeon's field of view. Curved arrows are used to indicate which coordinate systems are aligned [34].

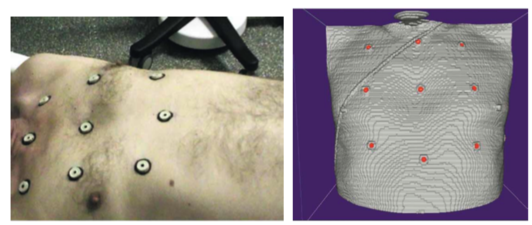

(a) fiducial markers anchored to the patient (b) fiducials are easily identified in acquired data

Figure 24: Fiducial markers may be placed on the patient prior to the preoperative CT scan to provide easily identified reference points to use during registration [12].

4.2 Registering Virtual Data to Video Stream

A paper by Wang *et al.* [34] examined the possibility of registering pre-operative data from MRI or CT scans to the endoscopic video feed from a da Vinci™ surgical robot. Initial results were not promising, as simple point selection exercises produced errors as large as 11 mm. The registration problem is common to a lot of medical AR applications (see Figures 23 and 24).

4.3 Port Placement

Port placement is usually considered during the pre-operative planning stage. A patient model is constructed from a CT scan, allowing the surgeon to determine the optimal port location for visibility and access to the structures of interest. Alternatively, placement may be determined computationally once the surgeon inputs the target points and their relative importances. In the case of robotic surgery, any restrictions on the placement of the robot's arms in the room will be taken into account [12] (see Figure 25).

On the day of the surgery, the virtual surgical plan will have to be registered to the coordinate frame of the OR. This can be done by relating the location of fiducial markers on the patient's body to their location in the virtual patient model [12] (see Figure 24).

4.4 Augmenting the Endoscopic View with Pre-operative Data

Since most of the visual information is conveyed to the surgeon with a screen or stereoscopic display, no new display apparatus is required to implement AR. Unfortunately, for a useful registration of the virtual model of the patient's interior, external fiducial markers will not provide adequate precision. Even if they did, organ shift and patient deformation on the operating bed would render the registration unusable. Thus, the surgeon must perform a short correction procedure intra-operatively by identifying some clear features in the video feed and the virtual data. Blood vessel bifurcations are often used as the point of bifurcation can be identified unambiguously. Once the system is fully calibrated, the rendering of the virtual data is constantly updated to match the camera view. Falk *et al.* [12] performed the first AR-navigated endoscopic cardiac surgery, but the registration error of 9 to 19 mm indicates that there is room for improvement.

More recently, an AR-assisted laparoscopic adrenalectomy was

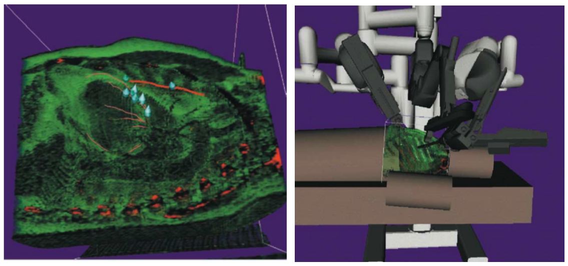

(a) visualization of target sites (b) virtual planning model

Figure 25: Computer interface designed to aid in port placement for robotic surgery [12].

performed by Marescaux *et al.* [20] with a human operator updating the image alignment in real time. The AR information helped to locate the tumour as well as nearby organs and blood vessels (see Figure 26). It was also useful in determining the appropriate dissection plane, and especially for identifying the main adrenal vein, which could then be safely isolated. The authors concluded that there were many benefits to the AR technique and that reliable, real-time automated registration algorithms would simplify the procedure.

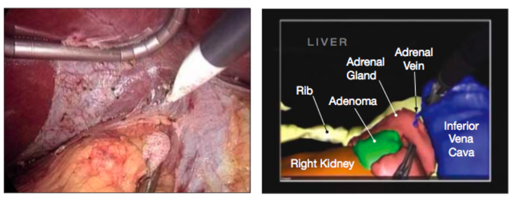

(a) typical laparoscopic view (b) 3D model generated from CT

(c) augmented laparoscopic view

Figure 26: Use of AR in laparoscopic adrenalectomy [20].

4.5 Surgical Training Simulations

The complexity of laparoscopic surgical procedures has led to an increase in the time spent training and evaluating surgeons before they can work on real patients. Since the surgeons performing these procedures mostly interact with a video display, VR and AR have been appealing options for training simulations. VR and AR systems also allow an objective assessment of the trainee's performance by tracking the position of the surgical tools [8].

Studies have shown that laparoscopic training is generally more effective when the simulations use realistic force feedback. AR training systems can provide accurate tactile feedback when the trainee performs the simulated procedure on real-world tissue phantoms that are augmented in their field of view. In contrast, realistic force feedback is quite difficult to implement in VR training systems. This has been cited as one of the reasons why AR simulations, and even inanimate box trainers, are generally favoured over VR simulations for laparoscopic surgery training [8] (see Figure 27). The use of AR haptic implementations in medical systems is discussed in Section 4.6.

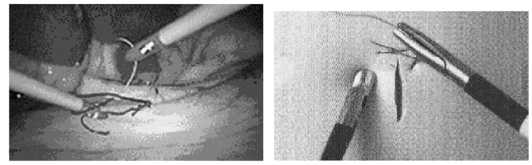

(a) virtual reality (b) augmented reality

Figure 27: Stitch operation in VR and AR training simulations.

4.6 Implementing Haptic Feedback in Medical AR

Although some proof-of-concept systems that implement haptic feedback in medical AR have been developed, there is much work to be done before these systems will be useful in the OR. Issues such as lag and tracking error need to become a focus if clinical integration is desired [4].

4.6.1 Lag

Lag can be a serious issue when implementing haptic feedback in AR systems because the computational demands are already so high. A system designed by Bianchi *et al.* [4], shown in Figure 28, used a distributed framework with separate physics and graphics servers so the necessary computations for the haptic interface could be performed independently of the visual AR computations. This also required a communication module and synchronization of the two servers to ensure synchrony of virtual events [4].

Figure 28: AR haptic feedback apparatus used by Bianchi *et al.* The optical tracker is shown in the inset [4].

4.6.2 Tracking Error

Another challenge in implementing haptic feedback is accurately tracking the instruments in space in order to precisely register haptic events to the surgical environment. Early studies reached an unacceptable precision of 15 mm, but this was improved to just a few millimetres by using a calibration grid in training systems. Bianchi *et al.* [4] were able to abandon the need to follow a grid by tracking the tip of a calibrated tool by following a marker at the other end (see Figure 29). The resulting tracking error was close to 1 mm but no evaluation of the system's medical utility was provided. The use of haptics in medical training simulations is discussed in Section 4.5.

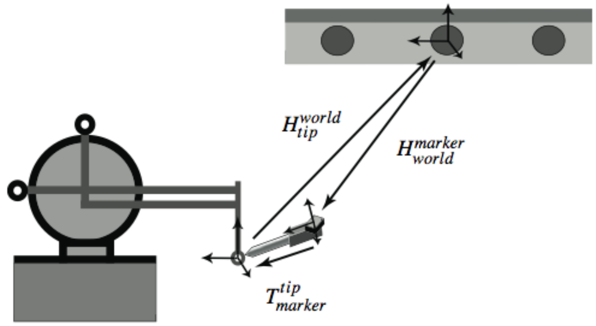

Figure 29: Diagram illustrating the tip-marker calibration in the system designed by Bianchi *et al.* The tip of the haptic device was fixed in space while the marker was rotated around in a sphere. The location of the marker was recorded by the tracking system at a series of time values, allowing localization of the centre of the sphere and hence relating the marker to the position of the tool tip. Another process was then used to calibrate the haptic system with the world [4].

5 Volume Visualization

Volume visualization presents a problem where there is much more information available than can be displayed at one time. Various strategies exist for filtering volume data, including volume reslicing, surface rendering and direct volume rendering (DVR) methods.

Volume reslicing is a simple method of visualizing volumetric data in arbitrary 2D slices. The user selects the position and orientation of the slice, and the reslicing algorithm interpolates between voxels (3D pixels) to determine the appropriate value to assign to each pixel in the resulting 2D image. In the case of medical data, physicians generally prefer the traditional slice views that they are familiar with from human anatomy texts [36] (see Figure 30).

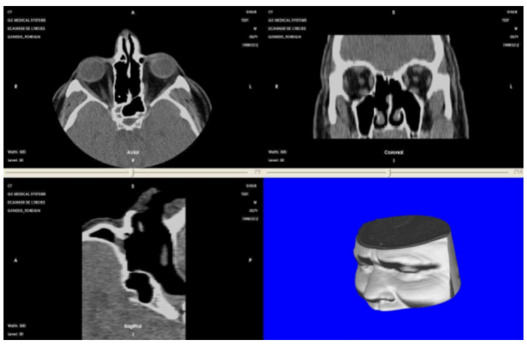

Figure 30: Reslices of a volume in the axial, coronal and sagittal planes, along with a 3D rendering [36].

Surface rendering is the rendering method of choice in the graphics industry, because no information about the inside of rendered objects is required. Rendered surfaces are composed of polygons, and they are displayed using an orthographic or perspective projection. One drawback of surface rendering is its absolute dependence on effective segmentation [36].

This paper will mostly focus on DVR as it represents data in 3D and it lends itself more easily to real-time visualization than surface rendering techniques [29].

5.1 Direct Volume Rendering

Direct volume rendering attempts to present volumetric data in its 3D context without requiring segmentation. Without some form of information filtering, only the outermost voxels would be visible. Methods used to form useful images in DVR include opacity reduction, volume clipping, selective ghosting and projection methods (see Figure 31). Most visualization tools provide some parameters that can be varied by the user [9].

5.1.1 Opacity Reduction

When reducing the opacity of volume data, some important information may be lost as objects are occluded by other semi-transparent objects. If opacities are reduced further, it becomes difficult to perceive object shapes. Many algorithms try to reduce these drawbacks by implementing complex transfer functions that use a more intelligent strategy for assigning opacity values. These functions may consider voxel values, gradients, curvature information, estimated material boundaries and distance from the viewer or focal point [9]. The transfer function design process is often referred to as classification because some pattern recognition is usually required to effectively assign optical properties to 3-D scalar data [24]. Other analogues to opacity reduction are screen-door transparency, where the occluding region is represented as a wire mesh with holes, and volume thinning, where the exterior of the occluding region is represented as a series of isosurfaces [31] (see Figure 32).

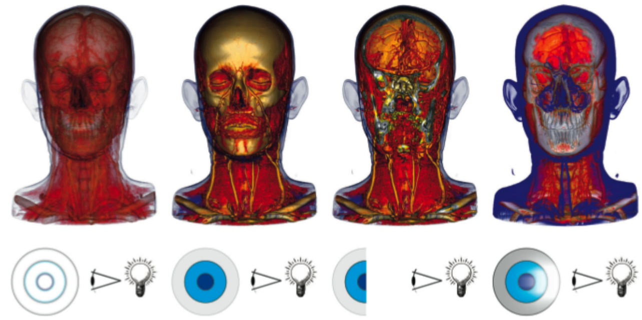

Figure 31: Direct volume rendering strategies, from left to right: simple opacity reduction, thresholding, clipping and a "context-preserving" strategy [9].

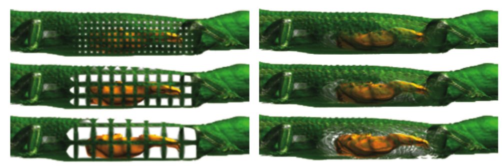

(a) screen-door transparency (b) volume thinning

Figure 32: Alternative forms of opacity reduction [31].

5.1.2 Volume Clipping

Volume clipping involves cutting away voxels (3-D pixels) based on their position. This technique allows the volume to be cut away up to a certain depth, exposing hidden information. Clipping algorithms do not usually consider the structure of the data being filtered, so it is up to the user to ensure that proper context is preserved. Some more sophisticated clipping interfaces present the user with a deformable mesh that specifies the clipping region or allow interactive carving of the volume [9, 31].

5.1.3 Selective Ghosting

Ghosting is an artistic technique used to show the underlying structure of objects by making unchanging, low-information density regions transparent while preserving the structures that provide the most information (see Figure 33). This approach is useful because salient occluding objects are preserved in order to maintain context. Some recent research in volume visualization has made an effort to adapt the principles of ghosting to volume rendering. The ghosting effect can be achieved by increasing the transparency of regions that are farther away from detected edges in the image [9].

5.1.4 Importance-based Filtering

Some rendering techniques have directly encoded "object importance" in the volume data, so pre-identified salient features are sure to be visible [31]. Importance-based methods may be useful when there is ample time to design renderings, but the requirement for segmentation and manual assignment of importance makes this strategy impractical for the frequently updated datasets encountered in AR. However, automatic segmentation and recognition algorithms can make importance- and classification-based methods more feasible. For example, in some medical data such as MRI,

Figure 33: Ghosting is a well known artistic technique for preserving context of a subject while presenting its underlying details [9].

different tissues may be represented by the same scalar values. In this scenario, a basic transfer function alone would not be sufficient, and some form of segmentation is desirable [24].

5.1.5 Projection Methods

Projection or ray-casting methods produce images by projecting a virtual ray from the viewer through the volume and recording the values of the voxels that it passes through. The output varies depending on what strategies are used for sampling, interpolation and merging [36]. Maximum Intensity Projection (MIP) is one of the simplest methods of this type. In this technique, the virtual rays keep track of the highest intensity voxel encountered along their trajectory. The value at this voxel is then presented at the appropriate pixel in the rendered image. Maximum Importance Projection works in a similar way, considering the importance value at each voxel in importance-based rendering methods [31] (see Figure 34). Other projection methods may use a threshold to determine, for each ray, the shallowest voxel that has a certain value [36].

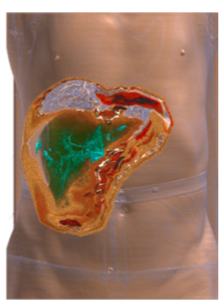

Figure 34: Maximum Importance projection, a derivative of maximum intensity projection, displays the voxel with the highest importance encountered along each ray's trajectory [31].

5.2 Medical Volume Visualization

It is often the case in medical visualization scenarios that the structures of interest are significantly smaller than the relevant context. However, the contextual information must be preserved as the final diagnosis may depend on the spatial relation between multiple features. Thus, many medical visualization scenarios can be approached as a focus and context task [31]. Visualization strategies can vary greatly between illustrative or exploratory applications and time-critical AR applications.

5.2.1 Illustrative Volume Visualization

It is not surprising that many medical volume visualization strategies borrow from anatomical illustration techniques. The strategy of applying non-photorealistic rendering (NPR) techniques to volume visualization has been dubbed "volume illustration" [29]. Some algorithms have adapted the notion of ghosting [9], while others are designed to model the entire illustration process [27]. Some work by Tietjen *et al.* [29] was inspired by medical atlases and research in perception to use lines to facilitate object recognition. Segmented focus and near-focus objects were emphasized with the use of lines and surface renderings while the context was presented with DVR (see Figure 35). Surgeons preferred images with less context information as long as the basic information for all context objects was always available.

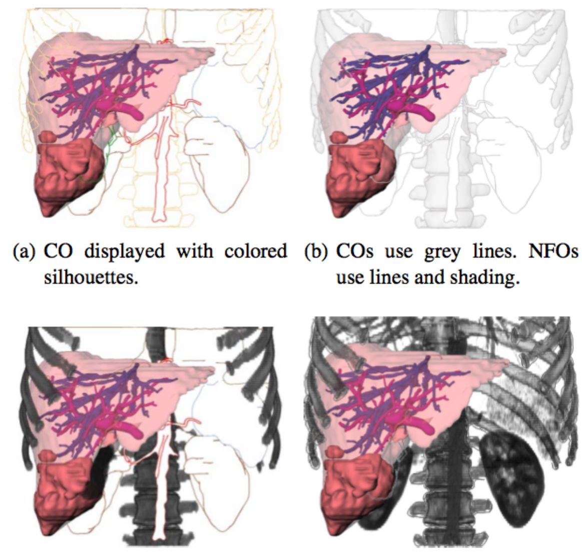

(a) CO displayed with colored silhouettes.

(b) COs use grey lines. NFOs use lines and shading.

(c) NFOs use DVR while COs are shown using colored lines on a white background.

(d) COs and NFOs are shown using DVR. NFOs also have lines added.

Figure 35: A technique by Tietjen *et al.* [29] renders the segmented focal object as a surface while using lines and shading to represent the near-focal objects (NFOs) and context objects (COs). A survey of physicians and medical laypersons showed that both groups preferred images (a) and (b), where only basic information was given for COs.

5.2.2 Exploratory Volume Visualization

Many volume visualization tools serve a significantly different purpose than static anatomical illustrations, so such strict adaptations of illustration techniques are not always appropriate. Systems designed to allow interactive volume exploration often provide a graphical editor to modify the transfer function for the volume (see Figure 36). This interface makes parameter selection slow and tedious. Improvements to this strategy impose constraints on the visualization parameters as selections are made in order to reduce the search space, or provide a histogram of the data values to guide the user towards making meaningful selections [30].

An alternate interface design turns the parameter selection into a visual one by generating images with a variety of parameter values and allowing the user to narrow their choices. A recent method by Tory *et al.* [30] placed more emphasis on the visualization of the parameter search space itself, by presenting a history of rendered images in a table and plotting the rendering parameters as parallel coordinates (see Figure 37). This made it easy to compare parameters between renderings and helped accelerate the user's convergence on a desired visualization. An expert evaluation concluded that this system outperformed traditional and table interfaces for data exploration and search tasks. This strategy marked a transition from image-based visualization to parameter-based visualization, and as such the parameters occupy most of the display space. The authors noted that some improvements could be made to allow full

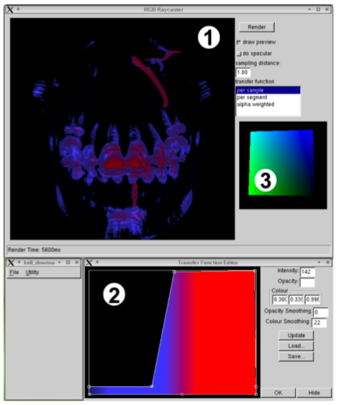

Figure 36: Standard volume visualization interface, with rendering window (1), transfer function editor (2) and zoom/rotation widget (3) [30].

transfer function interactivity and intuitive representations of multi-dimensional parameter values, and that constraints on display size may limit the scalability of the application.

5.2.3 Real-time Volume Visualization

Volume visualization in medical applications will commonly be under strict time constraints. Salama and Kolb describe the novel technique of opacity peeling for generic on-the-fly volume segmentation and rendering [24]. This method uses ray tracing and attenuation through the volume to determine how much can be seen from the outside, but resets the rays to full strength when they become insignificant or reach a strong gradient. Their algorithm was implemented with a GPU to allow rapid rendering of MRI datasets as they were acquired. "Opacity peeling" refers to the ability to remove layers corresponding to the depths where new rays were cast in the volumes. This technique proved particularly useful for looking beneath the skull and fat in brain MRI images (see Figure 38). The effectiveness of this method under time-critical conditions indicate that it is a candidate for integration in the OR. Although it was mentioned that the segmentations achieved by this method are generally poor compared to offline techniques, implementing this algorithm with more complex transfer functions may help to bridge the gap if it were to be run offline.

5.2.4 Visualizing Uncertainty

Awareness of measurements uncertainty is crucial in medical imaging because the consequences of a poorly-informed decision can be severe (see Figure 39). For example, the aortic width could easily be determined inaccurately when pre-operatively assessing an abdominal aortic aneurysm with unsuitable visualization parameters. This could prompt the life-threatening insertion of a stent of the wrong size. Experienced radiologists know not to blindly trust a DVR made with a pre-defined transfer function or "preset." They instead manually adjust the transfer function to sample the range of feasible values. Unfortunately, this process is time-consuming and there is no guarantee that all relevant values will be considered [19].

Past efforts in medical DVR used transparency or separate objects to represent regions of uncertainty. Lundström *et al.* [19] proposed a method of viewing error by rendering the output of a probabilistic transfer function using animation. The probability of a given sample falling under a certain classification is encoded by the fraction of frames in the animation in which it takes this form

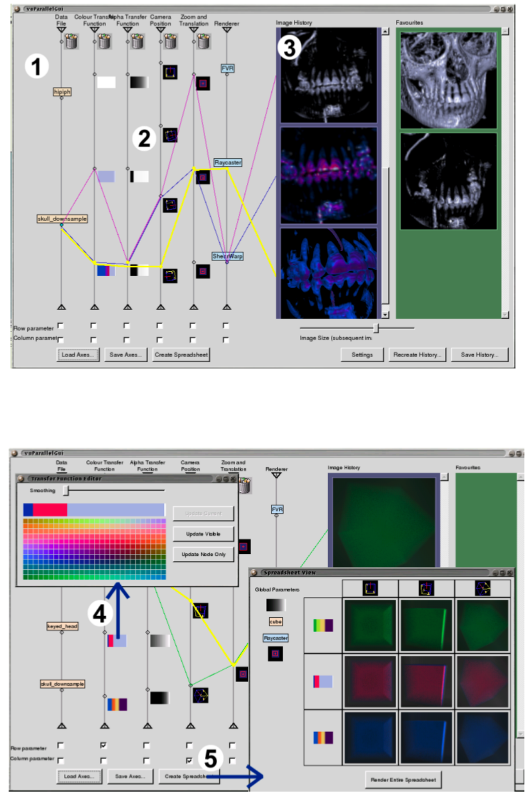

Figure 37: Parallel coordinates exploratory volume visualization interface, with: one axis per parameter (1), lines connecting parameters at appropriate values to their resultant rendering (2), and a history bar showing previous settings (3). (4) lets the user edit existing parameter nodes to make new ones, and (5) allows selection of parameters to compare in a small multiples view [30].

(see Figure 40). In order to produce these animations, they first had to develop a probabilistic transfer function to perform an automatic fuzzy classification of the tissues. Users were initially overwhelmed by the complexity of the large animated area, so a "sensitivity lens" was added. This tool allows the user to select the region where the animation should be applied, thus simplifying the rest of the visualization. This selective visualization technique is discussed in Section 2.8.2. The probabilistic animation method was shown to be more effective than a traditional static rendering when it simulated a sweep through the parameters of the transfer function. An added benefit of this method is that since the uncertainty is encoded separately from the data itself, no spatial precision is lost [19].

6 Discussion/Conclusion

It is evident that the applications of augmented reality are far-reaching. AR technology is permeating our daily lives and literally changing the way we see the world. However, the purpose of this survey was to examine its applications to medical technology and particularly to the operating room. This research was motivated by the question of whether AR can save lives.

Although most medical disciplines may be united by a similar goal, Medicine is also a very broad field. Where AR is a technology in search of a problem, Medicine presents endless problems that slowly find answers with the advancement of science. With such a

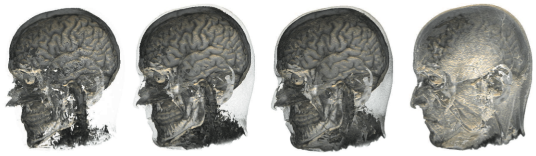

Figure 38: Renderings of the head and brain using the opacity peeling method. Note that the rightmost image is providing a rare view of the muscle layer directly below the skin [24].

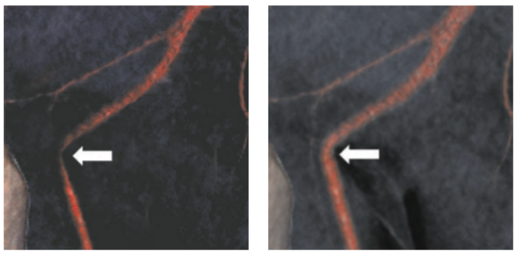

Figure 39: Commercial DVR software in clinical use presents the leftmost image by default, resulting in a surgery to remove the apparent stenosis. After exploring other settings, the image on the right shows that there is no stenosis and thus the procedure was unnecessary [19].

vast array of conditions and procedures, it is difficult to say which could benefit the most from AR. The topics covered in this paper show that different tools are suited for different problems, making it tough to compare effectiveness across disciplines.

Minimally invasive surgery was certainly a good area to consider first. AR is best used to enable humans to perceive what must normally be recorded by machine, and MIS is a field where critical operations are performed with very little opportunity to see what is happening. Medical imaging has made it possible to visualize immense amounts of data, but this is unfortunately pushed aside when surgery takes place. If AR could be made precise enough to rely on it for surgical decisions, surgeons could make use of this image data to effectively see right through their patients and localize critical structures immediately. The unpredictable nature of surgical procedures suggests that there is much to be gained.

At this time, however, there is still much work to be done. Many systems still suffer from tracking or registration errors that limit the reliability of the data. The results of the AR-guided adrenalectomy performed by Marescaux *et al.* [20] were certainly promising, but the system's dependence on a human operator for image registration indicates that some progress on the computational side is required if this technique is to see widespread use.

AR has also been successful in other medical domains. Near-infrared fluorescence imaging [28] and AR-guided needle insertions [32, 11] are close to being ready for clinical integration. As these and other systems are refined, the use of AR in hospitals will slowly become commonplace.

While AR is likely to generate many exciting new applications, the physicians will have the final word on what technologies get adopted in hospitals. Researchers who are developing new medical AR techniques need to keep the physicians needs in mind during the entire course of development. The developments thus far has shown a lot of promise, so we should expect patients to reap the benefits in the not-too-distant future.

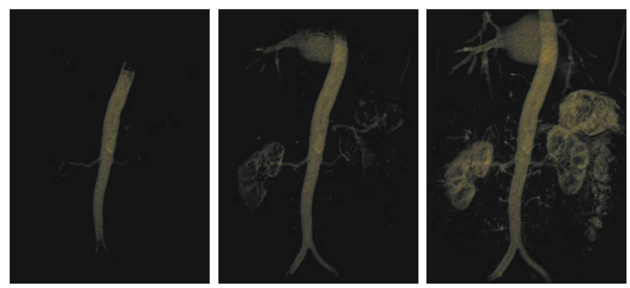

Figure 40: Sample frames showing an animation through the salient parameter values for a renal angiography image [19].

Disclaimer

The figures used in this document have been reprinted without the expressed permission of the authors. Please contact the referenced authors before using these images in a public setting.

References

[1] Ronald T. Azuma. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments*, 6(4):355–385, 1997.

[2] Ronald T. Azuma, Yohan Baillot, Reinhold Behringer, Steven Feiner, Simon Julier, and Blair MacIntyre. Recent advances in augmented reality. *IEEE Computer Graphics and Applications*, 21(6):34–47, 2001.

[3] T. Beyer, D.W. Townsend, T. Brun, P.E. Kinahan, M. Charron, R. Roddy, J. Jerin, J. Young, L. Byars, and R. Nutt. A combined PET/CT scanner for clinical oncology. *Journal of Nuclear Medicine*, 41(8):1369, 2000.

[4] Gerald Bianchi, Benjamin Knoerlein, Gabor Szekely, and Matthias Harders. High precision augmented reality haptics. In *EuroHaptics 2006*, pages 169–177, July 2006.

[5] Christoph Bichlmeier, Felix Wimmer, Sandro Michael Heining, and Nassir Navab. Contextual anatomic mimesis hybrid in-situ visualization method for improving multi-sensory depth perception in medical augmented reality. In *ISMAR '07: Proceedings of the 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality*, pages 1–10, Washington, DC, USA, 2007. IEEE Computer Society.

[6] Mark Billinghurst, Raphael Grasset, and Julian Looser. Designing augmented reality interfaces. *SIGGRAPH Comput. Graph.*, 39(1):17–22, 2005.

[7] Oliver Bimber and Ramesh Raskar. Modern approaches to augmented reality. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Courses*, page 1, New York, NY, USA, 2006. ACM.

[8] Sanne M.B.I. Botden, Sonja N. Buzink, Marlies P. Schijven, and Jack J. Jakimowicz. Augmented versus virtual reality laparoscopic simulation: What is the difference? *World Journal of Surgery*, 31(4):764–772, 2007.

[9] S. Bruckner, S. Grimm, A. Kanitsar, and M.E. Groller. Illustrative context-preserving exploration of volume data. *Visualization and Computer Graphics, IEEE Transactions on*, 12(6):1559–1569, Nov.-Dec. 2006.

[10] MS Chesters. Human visual perception and ROC methodology in medical imaging. *Physics in Medicine and Biology*, 37:1433–1476, 1992.

[11] Marco Das, Frank Sauer, U. Joseph Schoepf, Ali Khamene, Sebastian K. Vogt, Stefan Schaller, Ron Kikinis, Eric van Sonnenberg, and Stuart G. Silverman. Augmented reality visualization for CT-guided interventions: System description, feasibility, and initial evaluation in an abdominal phantom. *Radiology*, 240(1):230–235, 2006.

[12] Volkmar Falk, Fabien Mourgues, Louaï Adhami, Stefan Jacobs, Holger Thiele, Stefan Nitzsche, Friedrich W. Mohr, and Ève Coste-Manière. Cardio navigation: Planning, simulation, and augmented reality in robotic assisted endoscopic bypass grafting. *The Annals of Thoracic Surgery*, 79(6):2040–2047, 2005.

[13] Jan Fischer and Dirk Bartz. Stylized augmented reality for improved immersion. In *VR '05: Proceedings of the 2005 IEEE Conference 2005 on Virtual Reality*, pages 195–202, 325, Washington, DC, USA, 2005. IEEE Computer Society.

[14] Henry Fuchs, Mark A. Livingston, Ramesh Raskar, D'nardo Colucci, Kurtis Keller, Andrei State, Jessica R. Crawford, Paul Rademacher, Samuel H. Drake, and Anthony A. Meyer. Augmented reality visualization for laparoscopic surgery. In *MICCAI '98: Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 934–943, London, UK, 1998. Springer-Verlag.

[15] Zhen-hua Huang, Wei-yin Long, Gong-yuan Xie, Oi Ling Kwan, and Anthony N. DeMaria. Comparison of gray-scale and B-color ultrasound images in evaluating left ventricular systolic function in coronary artery disease. *The American heart journal*, 123(2):395–402, 1992.

[16] Adam Jones, Eric Kolstad, and Harvey S. Smallman. Egocentric depth judgments in optical, see-through augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 13(3):429–442, 2007. Member-Swan II, J. Edward and Member-Livingston, Mark A.

[17] Julian Looser, Mark Billinghurst, and Andy Cockburn. Through the looking glass: the use of lenses as an interface tool for augmented reality interfaces. In *Proceedings of the 2nd international conference on Computer graphics and interactive techniques in Australasia and South East Asia*, pages 15–18, 2004.

[18] W. Lorensen, H. Cline, C. Nafis, R. Kikinis, D. Altobelli, and L. Gleason. Enhancing reality in the operating room. In *Visualization, 1993. Visualization '93, Proceedings., IEEE Conference on*, pages 410–415, Oct 1993.

[19] C. Lundstrom, P. Ljung, A. Persson, and A. Ynnerman. Uncertainty visualization in medical volume rendering using probabilistic animation. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1648–1655, Nov.-Dec. 2007.

[20] Jacques Marescaux, Francesco Rubino, Mara Arenas, Didier Mutter, and Luc Soler. Augmented-reality-assisted laparoscopic adrenalectomy. *JAMA*, 292(18):2214–2215, 2004.

[21] Erick Mendez, Denis Kalkofen, and Dieter Schmalstieg. Interactive context-driven visualization tools for augmented reality. In *ISMAR '06: Proceedings of the 5th IEEE and ACM International Symposium on Mixed and Augmented Reality*, pages 209–218, Washington, DC, USA, 2006. IEEE Computer Society.

[22] Y. Moriguchi, T. Ohnishi, J. Decety, M. Hirakata, M. Maeda, H. Matsuda, and G. Komaki. The human mirror neuron system in a population with deficient self-awareness: An fMRI study in alexithymia. *Hum Brain Mapp*, 2008.

[23] K.J. Parker, M. Zhang, and NY Rochester. Color in medical imaging. In *Proc. IS&T/SIDs12th Color Imaging Conference*, pages 4–8, 2004.

[24] Christof Rezk-Salama and Andreas Kolb. Opacity peeling for direct volume rendering. *Computer Graphics Forum*, 25(3):597–606, 2006.

[25] Bernd Schwald, Helmut Seibert, and Tanja Weller. A flexible tracking concept applied to medical scenarios using an AR window. In *ISMAR '02: Proceedings of the 1st International Symposium on Mixed and Augmented Reality*, page 261, Washington, DC, USA, 2002. IEEE Computer Society.

[26] Gyung Tak Sung and Inderbir S Gill. Robotic laparoscopic surgery: a comparison of the da vinci and zeus systems. *Urology*, 58(6):893 – 898, 2001.

[27] N. Svakhine, D.S. Ebert, and D. Stredney. Illustration motifs for effective medical volume illustration. *Computer Graphics and Applications, IEEE*, 25(3):31–39, May-June 2005.

[28] Eiichi Tanaka, Hak Soo Choi, Hirofumi Fujii, Moungi G. Bawendi, and John V. Frangioni. Image-guided oncologic surgery using invisible light: Completed pre-clinical development for sentinel lymph node mapping. *Annals of Surgical Oncology*, 13(12):1671–1681, 2006.

[29] Christian Tietjen, Tobias Isenberg, and Bernhard Preim. Combining silhouettes, surface, and volume rendering for surgery education and planning. In *Planning, in: IEEE/Eurographics Symposium on Visualization*, pages 303–310. Springer, 2005.

[30] Melanie Tory, Simeon Potts, and Torsten Moller. A parallel coordinates style interface for exploratory volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 11(1):71–80, 2005.

[31] Ivan Viola, Armin Kanitsar, and Meister Eduard Grller. Importance-driven volume rendering. *Visualization Conference, IEEE*, 0:139–146, 2004.

[32] Frank K. Wacker, Sebastian Vogt, Ali Khamene, John A. Jesberger, Sherif G. Nour, Daniel R. Elgort, Frank Sauer, Jeffrey L. Duerk, and Jonathan S. Lewin. An augmented reality system for mr imageguided needle biopsy: Initial results in a swine model. *Radiology*, 238(2):497–504, 2006.

[33] D. Wagner, T. Pintaric, F. Ledermann, and D. Schmalstieg. Towards massively multi-user augmented reality on handheld devices. In *Third International Conference on Pervasive Computing*, pages 208–219. Springer, 2005.

[34] D. A. Wang, F. Bello, and A. Darzi. Augmented reality provision in robotically assisted minimally invasive surgery. *International Congress Series*, 1268:527–532, 2004. CARS 2004 - Computer Assisted Radiology and Surgery. Proceedings of the 18th International

Congress and Exhibition.

[35] Peter A. Woerdeman, Peter W. A. Willems, Herke Jan Noordmans, and Jan Willem Berkelbach van der Sprenkel. Auditory feedback during frameless image-guided surgery in a phantom model and initial clinical experience. *Journal of Neurosurgery*, 110(2):257–262, 2009.

[36] Z. Yaniv and K. Cleary. Image-guided procedures: A review. *Computer Aided Interventions and Medical Robotics*, 2006.