<u>**Dissertation Proposal**</u>

A Systems-Engineering Study of the Registration Errors in a Virtual-Environment System for Cranio-Facial Surgery Planning

Richard L. Holloway
Department of Computer Science CB 3175
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-3175

1. Introduction

Head-mounted displays (HMDs) can be divided into two categories: *opaque* (the real world view is totally or partially blocked by the display apparatus) and *see-through* (the computer image is superimposed onto the real world). Each type has its own uses and advantages, depending on the application. See-through head-mounted displays (STHMDs) can be further subdivided into *optical* STHMDs and *video-merge* STHMDs. Optical STHMDs use some form of beam splitter to merge the real-world view with the computer-generated view, while video-merge STHMDs use small cameras and a video mixer to do the same thing. This research focuses on optical STHMDs.

Optical STHMDs have existed since at least 1968, when Ivan Sutherland built the first stereo head-mounted display at Harvard [Sutherland 68], but their usefulness has not been widely explored except in certain military applications. Other see-through systems are cited in the bibliography; these systems typically do not address the issue of registration in any detail.

The main advantages of see-through head-mounted displays are:

1. They allow the computer-generated objects to be superimposed onto the real world, thus enhancing reality by augmenting it with useful data;

2. They do not obscure the real-world visual cues that humans use for navigation and balance.

While the promise of such displays is great, there are a number of technical problems that must be overcome in order for them to be useful. The main problem is that of *registration*. According to <u>The American Heritage Dictionary</u>, one of the definitions for *register* is “A state of proper alignment or positioning.” The STHMD system’s main purpose is to superimpose computer generated virtual objects onto a real-world scene. The virtual objects must be *registered* with the real environment so that they appear to be in the correct position and orientation at all times. If this is not done correctly, the objects will either appear to wander when they should remain still, or stay still, but in the wrong position or orientation (distortion error is discussed separately below).

If the system's function is to superimpose a virtual object on the real environment in a particular place, then the system is failing at that task to the extent that the virtual object and the real environment are misregistered.

I propose to perform a systems-engineering study to identify and quantify the registration errors in a real STHMD system by deriving a registration error model for it and then testing the model by measuring the physical system's performance with respect to those parameters. In specific, I plan to measure each component's contribution to the overall registration error; from the original dataset to the final images seen by the user (a detailed list is given in Section 7). In order to focus this research, I will study a system (hereafter referred to as "the x-ray vision system") designed to be used for a real-world problem: cranio-facial surgery planning. Studying a system targeted for a real application has (at least) two advantages: It defines when the system is "good enough", since no system will ever be perfect; and it reduces the number of parameters to a level where the system can actually be built and studied. While this study is targeted for the cranio-facial surgery application, the results should be easily generalizable to other applications with significantly less effort than starting from scratch.

2. Issues Not Treated

To build an effective system for cranio-facial surgery is a considerable task. My collaborator, Dr. Jefferson Davis (UNC Department of Plastic Surgery), feels that for the system to be useful, it should superimpose an accurate 3D model of the patient's bony tissue within a millimeter of where the real bone is. Even if this data is superimposed accurately, there is no guarantee that the human visual system will perceive the depth relationships correctly, especially since some occlusion cues will be in conflict with each other. Finally, even if the data is perceptually clear, there is no guarantee that such a system is more effective than current practice.

In order to address the issues of perception and effectiveness, the system must first succeed at registering the virtual anatomy with the real anatomy; otherwise, any perceptual issues relating to superimposition will be muddled with registration problems, and questions of effectiveness will be muddled with any failings in solving the first two problems.

Our experience has shown that currently available technology is not yet accurate nor fast enough for the registration requirements of this application. Since the registration problem must be resolved first, my research addresses an understanding of this problem. I therefore will not consider perceptual issues: My analysis will stop at the user's eyes[1] and will not concern itself with how the real and virtual images are processed by the visual system beyond that point. Moreover, I will not address the question of whether the x-ray vision system is or will be more effective than current practice.

The only issue addressed by this research proposal is the derivation and validation (with a real-world system) of a registration error model for a STHMD system for cranio-facial surgery planning. It is my hope that such a model will give future system designers a better idea of the sources and sensitivities of error in a STHMD system so that we will eventually be able to build a usable system for this application. This is a formidable task in itself, and I hope it will lead to further research on perceptual issues and questions of effectiveness.

---

[1] To be precise, at the user's first *nodal points* (defined below).

3. System Overview

The purpose of the system is to superimpose a three-dimensional image of the virtual object(s) into the real environment such that the virtual object appears to be fixed in space as the viewpoint is changed. In order to do this, the system must have a STHMD on which to display the images, a tracker for monitoring head movement, an image-generation system for creating and updating the images of the virtual object, and a host computer to coordinate the activities of these components.

The UNC system will have the following components:

See-Through Head-Mounted Display: The STHMD for this study will be "seethru2", a simple, low-resolution, LCD-based, 30˚ field-of-view STHMD which was built for this research.

Tracker: The task of the tracking subsystem is to report the position and orientation of the user's head and hand for each display cycle. There are many different technologies used for this task; we currently have Polhemus and Ascension magnetic trackers, mechanical trackers (the Argonne Remote Manipulator) and the UNC ceiling tracker. I will experiment with different trackers in order to find one that works the best with the STHMD system.

Graphics: The graphics engine for the system will be Pixel-Planes 5, a highly parallel raster graphics engine developed at UNC. The host computer sets up the transformations for the viewing calculations and then passes this information to Pixel-Planes 5 for display. The images are rendered into two NTSC frame buffers, whose signals are passed through NTSC encoders and then on to the displays in the STHMD. Pixel-Planes 5 can currently render over two million Phong-shaded triangles per second, which makes it well suited for displaying complex medical data sets. The system will initially use polygonal models for the anatomical data, since these models can be derived using existing software and displayed using existing graphics libraries. If this is inadequate in any way, volume rendering techniques can be substituted at some programming cost.

Host Computer: Currently a Sun 4. The host coordinates the activities of the tracker and the image-generation system. It receives position and orientation information about the user's head and hand from the tracker system. The host then processes this data and sends it on to the graphics system, which updates the computer graphics model accordingly.

4. The Application: Cranio-Facial Surgery Planning

4.1 Overview

The cranio-facial surgeon’s task is to move or reconstruct pieces of the patient’s bone in order to make the soft tissue end up in the right place. This is both a technical and artistic task, requiring the surgeon to form a mental model of the bone underlying the soft tissue using CT scans of the soft and bony tissues. Given the difficulty of this task, much of the hard work goes on during the planning phase, so that by the time the surgeon begins the operation, he has a clear plan of what he must do. One of the hardest parts of the planning task is understanding the complex relationship between the bone and the soft tissue at each point. Due to this difficulty, surgeons occasionally find “surprises” in surgery– anomalies of the bone or soft tissue that they had not perceived during the planning phase. While these are rare, a system that improved the surgeon’s understanding of the hidden anatomy would reduce the possibility of such incidents.

While current practice currently yields acceptable results, a system that enabled the surgeon to do the planning phase better (faster, more accurately, or with greater understanding) would be quite useful.

4.2 Current Technology

Currently, cranio-facial surgeons have the following tools at their disposal:

- Clinical examination
- Diagnostic acumen
- Dental impressions
- Moulages (facial impressions)
- Photos of patient
- Routine x-rays (cephalometrics)
- Computed Tomography (CT) scans (best for bone images)
- Magnetic Resonance (MR) images (good for soft tissue)
- Physical measurements of patient
- Standardized measurements from x-rays
- 3D reconstructions of CT or MR images viewed on conventional CRT
- 2D images made from 3D reconstructions

While the imaging tools are quite useful and important, they have limitations. The primary limitation is that they are typically examined in a 2D fashion separately from the patient. The CT scans and MR images alone are 2D, and the surgeon has to reconstruct the 3D information in his head from them. The 3D constructions from these images solve this problem, but are typically displayed on a 2D screen away from the patient. Even the addition of stereo viewing of 3D constructions (which is done, but is not standard practice) does not address the inability to interact with the patient while seeing the images.

4.3 The Planning Process

The typical surgical plan for cranio-facial reconstruction is developed using a standard medical procedure known as "SOAP", which stands for *Subjective, Objective, Assessment* and *Planning*. These four stages are described below.

Subjective: The patient does a self-evaluation and determines subjectively that something is wrong.

Objective: The surgeon examines the patient and performs an objective evaluation of what is wrong.

Assessment: The surgeon makes a diagnosis on what is wrong, based on evaluations of the tests and scans.

Planning: The surgeon forms a plan to correct each problem.

While these steps are a useful conceptual model, the distinctions between them are often blurred. For example, planning may well begin at the moment the surgeon first sees the patient, and may continue through the assessment phase to the planning phase.

The typical scenario for patient treatment with the SOAP method is as follows:

Physical Examination

The surgeon examines the patient to get a gut feeling for what is wrong. During this process, the surgeon does the following:

- makes judgements about what is abnormal about this patient, using normative data
- takes measurements, both of bony and soft tissue
- judges the patient's overall health

Imaging

After examining the patient in person, scans and/or casts may be ordered to get more information on the patient. This may include any or all of the following:

- routine x-rays (cephalometrics)
- CAT scans (for bone; axial and 3D)
- MR scans (for soft tissue)
- pre-operative photos of patient
- dental casts
- facial casts (moulage)

Re-examination

After the images and casts are evaluated, the patient may come back for an additional examination to compare the scans with the patient in person. Sometimes the surgeon visually superimposes the scans onto the patient in order to form a better mental image of the underlying structure. He may also superimpose one scan onto another to achieve the same effect. Additional scans may be ordered at this point as well. The entire procedure up to this point may be repeated until the surgeon feels that he has enough data for a surgical plan.

Planning

The surgeon studies the scans, the casts, and his notes in order to plan where the bone should be cut, where the pieces should be placed, and how best to place them in order to make the soft tissue end up in the right place.

4.4 Surgical Planning with the X-Ray Vision System

The x-ray vision system would aid in process described above by allowing the surgeon to see the 3D CT reconstructions superimposed on the patient. The surgeon could have the patient scanned early in the diagnostic process and could then use the x-ray vision system at any point during the assessment and planning phases. After registering the reconstructed data with the real patient, the surgeon could use the system to judge soft-tissue depths, make measurements, and study the relationships between the bone and soft tissue in the patient. This would essentially overlap the assessment and planning phases, since the surgeon could view the scans and the patient at the same time.

This augmented system should ideally have the following inherent advantages over current practice:

- The surgeon does not have to look away from the patient in order to see the reconstructions.
- The surgeon can see the reconstructions superimposed in the correct locations on the real patient.
- The view the surgeon gets of the patient's soft tissue is the best possible, since he is seeing the surface tissue itself.
- The surgeon can palpate the soft tissue at the same time that he views the bone in order to get a better understanding of the relationship between the two.
- Making measurements of the soft and bony tissues is easier since the registration of the patient's landmarks is unambiguous in real life, as compared to a CRT screen.
- Viewpoint-changing is more natural, since the surgeon simply moves his head to get a different view, as opposed to some other artificial interface (such as a joystick).
- The session with the patient is a necessary part of any planning which could never be replaced entirely by a session with a 3D reconstruction; thus, this enhances that phase rather than trying to replace or compete with it.
- Superimposing the soft and bony tissues with film doesn't work very well- the superposition afforded by the x-ray vision system is clearer[2] and more natural.
- The views of the virtual anatomy are 3D, vs. 2D for conventional CT scans.

The hope of this system is that it would allow the surgeon to have shorter and fewer sessions with the patient, yet result in more accurate measurements and greater situation awareness for the surgeon.

5. Error Model

5.1 Overview

The patient is scanned initially to acquire a 3D dataset of the bone, which is then processed and displayed in (ideally) the same location as the soft tissue. The real and virtual datasets are both three-dimensional surfaces. We have the choice of either trying to compare the two surfaces for misregistration, or using some lower-dimensional data (such as edges or points) to characterize each dataset. Since current practice in cranio-facial surgery uses landmarks for analyzing the patient's anatomy, it is natural to use these as a measure of how well the real and virtual datasets are registered. In addition, there are a host of existing techniques for characterizing the "distance" between two sets of landmarks. A thorough treatment of these techniques is given in [Bookstein 91]. The distance between real and virtual landmark pairs will give a measure of how far the system is off at any time and for any given position. Thus, the goal is to find an error vector ε_i between an arbitrary pair of real and virtual landmark points $\mathbf{p}$ and $\mathbf{p}'$ as a function of space and time.

---

[2] Not clearer in the sense of resolution (film has much higher resolution), but in the natural way which it presents the information on the patient's anatomy.

In order to do this, each *operation* (e.g., scanning, rendering, tracking) in the system must be modeled to determine how it generates and propagates error as a function of the system's parameters. Some of the system parameters (such as head position as measured by the tracker) will be modeled as random variables, since their values are not stable over time and space. I will derive a model for each operation and test it by measuring the registration error in our existing x-ray vision system. For simplicity, I will begin by making the following assumptions, which will be tested by measurement:

1. A first-order model will be sufficient;
2. There will be no interactions between parameters;
3. Parameters modeled as random variables will be normally distributed.

By building the system incrementally and testing the model at each stage, I should be able to tell if and when any of these assumptions is incorrect and to make the necessary changes in the model.

5.2 Types of Registration Error

I propose that the registration error for a given point can be divided into two categories:

1. Alignment error- Misregistration due to relative rotation and/or translation of the real and virtual datasets. This error can be corrected for if we can find the proper transformation between the two objects' coordinate systems.

2. Distortion error[3]- even if the alignment of the real and virtual objects were as good as possible, there could still be a point-by-point misregistration due to scaling errors or distortion of the real and/or virtual objects that cannot be corrected by translation and rotation.

The first category of error is easier to correct for than the second, since there is a unique correct solution and it does not involve warping or scaling the virtual object in order to make it line up with the real one. The second type of error is much more difficult, since there may be multiple solutions of arbitrary complexity.

5.3 Error Descriptions

I will discuss each error source, starting from the original CT scan and working toward the images that the user's eyes will see while using the system. The items listed generally correspond to operations in the system, which both generate error through imperfections in the process and propagate error in their inputs. Each error will be labeled as an alignment or distortion error (or both).

1. CT scanning errors (distortion): CT scans do an imperfect job of representing the patient's anatomy. Examples of errors present in any scan are: blur from patient movement during scan, aliasing due to finite slice thickness and detector/beam size, and beam-hardening artifacts. These artifacts may affect where the surgeon digitizes the landmarks for the virtual object.

---

[3] In this document, I will distinguish this broad category of "distortion error" as defined here from optical distortion, which is a type of monochromatic aberration.

2. Volume visualization (distortion): Both volume and surface rendering methods introduce errors in the final image- volume methods introduce some amount of blur and surface methods are subject to artifacts. As with (1), these may affect the initial placement of the virtual-object landmarks.

3. Dataset registration (alignment): Since the landmarks for the real and virtual datasets must be digitized by hand, there will be error in their locations. This error will be propagated by the registration algorithm, which will also introduce some amount of computation error.

4. Static tracker error (alignment): Any error in the tracker’s head position and orientation data will cause corresponding errors in the superimposed image position. Static error is how much the tracker is off when the user holds perfectly still for long enough for the tracker's output to stabilize.

5. Dynamic tracker error (alignment): Depending on the type of tracker used, the error may be different when the tracker is in motion. In particular, many magnetic trackers pulse their measurement signals sequentially in time for each axis, which causes more error for higher-frequency movements.

6. Tracker delay (alignment): Even if the tracker were perfectly accurate at all times, there is always some amount of delay between the time the tracker makes the measurement and the time that measurement is available to the system for use in display calculations. This delay causes the virtual image to appear to jump away from the real objects for sudden movements, and then snap back once the user's movement has stopped.

7. Image-generation delay (alignment): Once the user's eyepoint data is known, there is an additional delay before the updated images appear on the screen due to rendering calculations and the raster refresh cycle. LCDs may introduce yet more delay.

8. Rendering error (alignment, distortion): The final image is rendered into a discrete frame buffer, which means that the final image will have some amount of either aliasing or blur. There is also some computation error in the transformation from world space to eye space. Also, errors in the estimation of the field of view and center of projection in the computer graphics model of the system will introduce errors.

9. Display (distortion): All displays represent images imperfectly to varying degrees: CRT images are refracted at the glass faceplate and have beam-deflection distortion, and finite electron-beam spot size, to name a few. Currently available LCDs have even larger sized pixels and may have other problems.

10. Optics (distortion, alignment): The lenses typically introduce a number of monochromatic and chromatic aberrations, including optical distortion. These aberrations move the virtual landmark points in a number of different ways. The beam splitter may introduce refraction alignment error and distortion error since it has finite thickness.

11. STHMD construction (alignment, distortion): The STHMD itself will introduce some alignment errors and distortion errors, since it may be built with the screens improperly positioned relative to the optics. For example, a small error in the placement of the screen relative to the optics may map to a large change in the

distance to the screen's image; tilting the screen relative to the optics will introduce keystone distortion.

12. STHMD-Head calibration (alignment, possibly distortion): Even if the STHMD is built perfectly, it will still be placed and worn on a human head. If the STHMD is not aligned properly when put on, there may be alignment or distortion error. User head motion may cause slippage of the HMD on the head which will cause more misalignment. Beyond, that, the user's eye movement will move the pupil of the eye (which is usually the exit pupil for the system), which will move the first nodal point of the eye from the assumed center of projection, which introduces errors.

13. Stereoscopic error (alignment, distortion): The errors listed so far can be considered to be strictly monocular; ie., they must be considered even if the STHMD has only one viewpoint. This error is the failure of a point's two projections on each virtual screen to correspond to its proper location in 3D. The projectors for the two projected points may or may not cross, leading to either an erroneous 3D location or an inconsistent stereo image. This is a *derived* error in that it is completely a function of the previously listed errors and thus is not an error source per se.

5.4 Analysis and Research on Further Improvements

Once the model has been verified for the real system, I will perform a sensitivity analysis to show how much a given improvement in any part of the model will impact the final registration error distribution.

Next, I will research new technologies and techniques for reducing the system error, using the results from the sensitivity analysis. In particular, I will determine an error budget for each component of the system (based on the overall error) and see whether any existing or predicted technology can meet the budgeted amount for each component.

6. Testing the Model

6.1 Overview

To test the model described above, I will perform a component-by-component analysis of our existing, 30˚-FOV STHMD system. Basically, I plan to dismantle the existing system and measure each component's properties with respect to the model, then to re-build the system piece by piece and verify the model for each incremental subsystem.

The measurement apparatus will be a "virtual user"- a pair of tiny video cameras mounted inside a mannequin head. The cameras will "see" what a user wearing the STHMD would see and will be used to measure the misregistration of the virtual and real objects.

My plan is to use a top-down, model-driven approach. I will first derive the error model for the parameters expected to be the biggest error sources. I will then perform a sensitivity analysis for these parameters, using typical values for parameters such as scene complexity, user head velocity, tracker error variance, etc. Next, I will build the necessary measurement setup and test the model for the included parameters. The model will be revised and expanded based on discrepancies between predicted and observed errors. New parameters will be added to the model and then tested with the lab setup. A complete description of the research plan is given in section 7.

6.2 Special Equipment

The following equipment (at least) will be needed:

- *Optical bench.* I will use the optical bench that was used for calibrating the photodiodes for the ceiling tracker. It has a plotter platform mounted vertically at one end for precisely positioning LEDs (currently infrared, but we can swap in some visible LEDs).

- *Two CCD cameras w/ mount.* The cameras should be small enough to fit inside a mannequin head, probably be monochrome to simplify measurement, and have at least NTSC resolution. NTSC resolution (O(500x500)) should allow us to resolve two points separated by a distance of less than 1 mm at a distance of 0.5m with a field of view of 20°[4]. We can go to higher resolution cameras if we are willing to spend more money on special frame-grabbers. The cameras should have an electronic shutter for use in dynamic registration experiments. Mount should have adjustments for fine-tuning alignment of cameras relative to each other. The mount itself will probably have a fixed connection to the optical bench. For rotations, a turntable with a rigid attachment to the optical bench can be made with an appropriate coupling for the camera mount.

- *Frame grabber(s).* I'll need two frame-grabbers or a single stereo frame-grabber capable of taking in about a second's worth of video and making it immediately available for processing. Mike Bajura's frame grabber (640x512, double buffered) may be appropriate for this task. Ideally, we should be able to find the centroids of the projected points in real time so that we don't have to store too many frames. If we have to, we can ship the data to the D-2 recorder, but this introduces more complications (e.g., it may have lower resolution than the frame grabber and its digital output is in an obscure format).

- *Dynamic measurement apparatus.* There are a few ways to measure the system's dynamic performance. The key criterion is that we have to be able to predict or measure both the position, velocity and acceleration of the virtual head during each trial so that we can verify their effects on the error. Vern Chi suggested that it may be possible to get this data from the Sarcos arm; we could then use the arm to move the virtual head and use its data to derive the velocity and acceleration for describing the motion. Another possibility is to use Ron Azuma's rate-gyro/accelerometer setup to measure arbitrary motions.

7. Plan

(Note that the actual execution order of the plan items will be determined by the results of the sensitivity analysis in step 1 and a PERT chart analysis of the entire project. The order used in the plan below is a straightforward, bottom-up creation of a STHMD system.)

---

[4] Having a smaller FOV reduces distortions in the camera lenses and increases the resolution in the area of measurement. 20° at 0.5 corresponds to an object width of about 18cm, which is roughly the width of a human head.

1. Choose parameters in system that I expect to be major error sources and derive error model for them. Do sensitivity analysis and determine ranking of parameters.

2. Procure two CCD cameras and calibrate them so that we can derive the real-world position of a point given its projection in the two camera images. This can be done with the optical bench and plotter setup used for calibrating the ceiling tracker photodiodes and existing camera-calibration techniques. The algorithm in [Shih 93] may be appropriate for the calibration since it accounts for optical distortion.

The cameras' outputs will be connected to frame grabbers, which will store the composite real and virtual object images for analysis.

3. Derive error term for STHMD lenses. Measure these separately using standard optical techniques (may have to send lenses to an optical shop to do this). Then use a test pattern and the cameras to verify distortion figures.

Parameters:
- optical aberrations
- virtual image size, transformation

4. Replace test pattern with LCD screens and calibrate them for overscan, tilt, linearity, resolution, etc. Next, test parameters of computer graphics model: center of projection, field of view, mismatches between eyes for stereo, frame-buffer resolution, (anti-)aliasing. Errors in three-space can be derived from the errors in the left and right screen spaces.

Parameters:
- screen: size, resolution, linearity, transformation relative to optics
- graphics model parameters
- eye movement (movement of nodal points)

5. Incorporate principal-component analysis (as described in [Toennies 90]) into existing x-ray vision application for registering the real and virtual datasets. This will give the best fit of the virtual landmark set to the real landmark set. Bending energy [Bookstein 91] can be used to measure the degree of remaining misregistration.

Parameters: error in real and virtual landmark positions, algorithmic error

6. Introduce beam splitters and verify error term for refraction and misalignments. Test registration algorithm for some simple, easily measurable objects. Measure their landmark positions in world space and attempt to register the virtual and real datasets. Camera's location/orientation in world space should be known from the calibration procedure already done.

Parameters:
- beam-splitter refraction, transformation relative to screen/eyes
- registration algorithm

7. Mount the cameras inside a mannequin head and place the 30˚ FOV STHMD on it. Evaluate required calibration procedure for determining transformation between STHMD and eyes (this may be non-trivial). If feasible, derive and implement calibration procedure. Explore effects of moving STHMD on head and incorporate these terms into the error model.

Parameters: STHMD-eye transformation: calibration procedure and variation in xform

8. Measure latency of image-generation system. Plot latency as a function of polygon count. To reduce the parameter space, use typical anatomical models viewed at typical distances (including some minimum-complexity dataset). The measurements can be done both with a photodiode on a screen as an input to a scope, as well as via frame-counting with the camera's output.

Parameter: image-generation delay as a function of polygon count

9. Find error distribution and latency for the best available magnetic 6DOF tracker via static and dynamic tests.

Parameters: static and dynamic tracker error, latency

10. Build setup for measuring dynamic performance of system and model. We will need some system whose motion can be precisely characterized in order to specify the motion input to the STHMD system. See the discussion in "Special Equipment" below for details. We will not be able to measure the system's performance for all possible inputs, but should be able to take measurements on some typical[5] motions to verify the model for the dynamic case. In particular, rotation about the vertical axis and translation from left to right are two motions that should be critical for this application.

If there is not too much interference due to metal, connect the magnetic tracker sensor to STHMD and verify its error contribution.

Parameters: image-generation and display delays, STHMD movement/non-rigidity

11. Do sensitivity analysis for virtual object derived from CT scan using typical error bounds for scanners and volume visualization techniques. If the resulting error is not insignificant relative to the other terms studied so far: Scan a simple phantom object with a CT scanner, polygonalize resulting dataset, and superimpose virtual version onto the real object. Verify that the error distribution corresponds to that predicted by the model.

12. Perform full-system sensitivity analysis, develop error budget, research new technologies.

---

[5] To find out which motions are typical, I will do a study with Dr. Davis during which he examines a virtual patient. His head will be tracked and the tracker data logged; the session will also be videotaped. I will then analyze this data to determine what motions are typical and what sort of energy spectrum is associated with this task.

8. Bibliography

Altobelli, D., R. Kikinis, J. Mulliken, W. Lorensen, H. Cline, F. Jolesz. 1991. Intraoperative Navigation in Craniomaxillofacial Surgery. *J. Oral and Maxillofacial Surgery*. 49:8:57.

Altobelli, D., R. Kikinis, J. Mulliken, W. Lorensen, H. Cline, F. Jolesz. 1991. Computer Assisted Three-Dimensional Planning for Cranio-Maxillofacial Surgery: Osteotomy Simulation. *J. Oral and Maxillofacial Surgery*. 49:8:89.

Bookstein, Fred. 1991. Morphometric tools for landmark data : geometry and biology. Cambridge University Press.

Chesnais, Pascal. 1988. A graphic/photographic arthroscopy simulator. Masters thesis, Massachusetts Institute of Technology.

Deering, Michael. 1992. High resolution virtual reality. Computer graphics. 26:2:195.

Drascic, David and Paul Milgram. 1991. Positioning accuracy of a virtual stereographic pointer in a real stereoscopic video world. *Proc. SPIE 1457 Stereoscopic Displays and Applications II.* 1-12.

Glines, C. V. 1986. Brain Buckets. *AIR FORCE Magazine.* 87-90.

Hirose, Michitaka. 1991. Visualization Tool Applications of Artificial Reality. (From author; no publication information).

Schmandt, Christopher. 1982. Interactive three-dimensional computer space. SPIE *Processing and Display of Three-Dimensional Data.* 367:155-159.

Shih, Sheng-Wen, Y Hung, W Lin. 1993. Accurate linear technique for camera calibration considering lens distortion by solving an eigenvalue problem. Optical engineering. Jan. 32:1:138.

Sutherland, Ivan. 1968. A head-mounted three-dimensional display. *Proceedings of the 1968 Fall Joint Computer Conference.* Thompson Books. 757–764.

Toennies, Klaus, J. Udupa, G. Herman, I. Wornom, S. Buchman. 1990. Registration of 3D objects and surfaces. IEEE CG&A. May, p 52.