Verification Procedures in a Medical Imaging Application

Neil Muller, Evan de Kock, Ruby van Rooyen

iThemba Labs, P. O. Box 722, Somerset West, 7129, South Africa

Abstract

At iThemba Labs, we are constructing a patient positioning system for proton therapy. This system will use stereo vision techniques to accurately position the patient relative to the beam line. To do this, we need to be sure that the vision system is functioning correctly. In addition to various procedural checks, we implement a variety of software tests to try and detect errors in the vision system.

1. Introduction

Proton radiotherapy is a useful treatment method for a number of lesions. The dose distribution properties of the proton beam allow for high doses to be delivered to the target volume while keeping the dose to the surrounding tissue to a minimum. Due to the high cost associated with this treatment, it is often reserved for lesions that are difficult to treat with conventional radiotherapy techniques, especially lesions close to critical structures. For more information see for example [1].

iThemba LABS has been involved with proton therapy for over ten years. Due to cost restrictions, iThemba LABS uses a fixed beam-line to deliver the proton dose, and uses a robotic manipulator to position the patient. The position of the patient during setup and treatment is monitored by a number of cameras and stereo techniques are used to calculate the patient's position at any time. A critical issue is the high positioning accuracy required. For further discussion on the system and some of the previous work on the vision aspects see [2], [3] and [4].

As we need accurate positioning of the patient, we need stringent controls on the quality of the vision system. In addition, the system will be used by people who are not imaging experts, thus we need to have simple quality assurance procedures, and, where possible, automate as many quality checks as possible.

2. System Description

Figure 1: Treatment Environment.

The treatment environment currently under construction (illustrated in figure 1), shows the treatment layout. We have a total of 9 cameras observing the patient during treatment. During a treatment session, at least 3 cameras will be used to calculate the position of the markers on the mask (see the mask shown in figure 2). As the relationship between these markers and the target volume is known (see [3] for details), we can use the observed marker coordinates to calculate the position of the target volume in the room.

Since we need to monitor the patient position in real time and react to patient motion, we restrict ourselves to a subset of the cameras during a treatment session to minimise the computational load. We also try to ensure that the marker detection problem is as simple as possible, by using retro-reflective markers, which are illuminated by green light-emitting diodes (LEDs), and narrow bandpass filters on the cameras, to ensure that the images are as noiseless as possible. For more details, see [4] and [2].

3. Procedural Checks

3.1. Calibration

Obviously the camera's need to be well calibrated. Thus we need a simple and easily repeatable calibration procedure. We use a fairly standard calibration cube, and ensure that the cube position can be repeatably placed in a surveyed position by using a suitably designed jig.

We use distinct textures for each side of the cube (as

Figure 2: Treatment Mask.

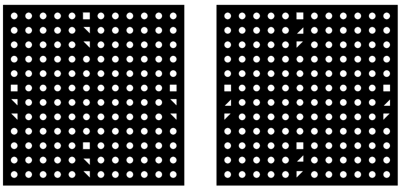

Figure 3: Calibration Textures

shown in figure 3). This allows us to detect cases where (due to maintainance work) the camera has been connected to the incorrect frame-grabber. These errors can be detected in software and brought to the attention of the operator.

3.2. Radiation Damage to the Cameras

The cameras will be placed in the treatment vault. As we have a high energy proton beam i the vault, we can expect camera damage to occur due to beam scattering effects. While the cameras will be shielded, this will only serve to slow the rate of radiation damage to affordable levels. We need to detect damage so that we can both compensate for small amounts of camera damage and detect when the camera needs to be replaced. See [2] for more details.

This is done by taking a dark-frame for each camera as part of the maintenance cycle. Radiation damage to the CCD shows up as stuck pixels, which are easily detectable in the dark frame. The camera is declared overly damaged once either the number of stuck pixels exceeds some maximum number or we get a group of stuck pixels in a large enough cluster to adversely affect the accuracy of our marker detection routines. See [2] for further details.

A section of a dark frame, showing some damaged areas, is shown in figure 4.

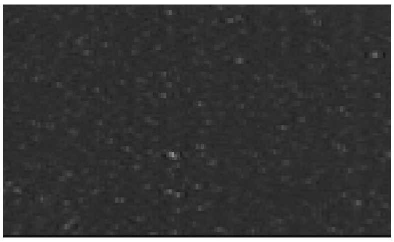

Figure 4: Dark Frame Section.

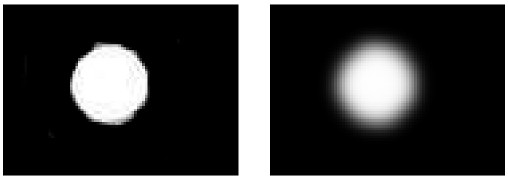

Figure 5: Markers.

4. Software Verification

4.1. Focus

Accurate marker detection requires that the images be well focused. We rely on a simple procedure to ensure that the images are in focus.

We note that, if the image is in focus, edges will be sharp. Since the retro-reflective markers are designed to be easily visible, we naturally look to use the edges of these markers to test the focus of the images. Furthermore, our markers have a simple geometric structure (circles), so the edges are well defined and can easily be identified by taking the gradient across the image.

For each edge, we obviously get a peak in the gradient. Thus we need a measure of the width of this peak to determine if the image is in focus. We use the traditional full-width-half-maximum (width of the peak at half the maximum value of the gradient peak) to measure the edge sharpness. Provided this is below a suitable threshold, we declare the target to be in focus.

Two markers, one in focus and one out of focus are shown in figure 5. The gradients across the middle of the markers are shown in figure 6. As can be clearly seen, the peaks for the out of focus marker are much wider than the focused case.

In principal, we can generate test cases in any arbitrary direction, but for simplicity, we test only the horizontal and vertical directions. This gives four peaks for each marker (2 horizontal edges and 2 vertical edges). We accept a marker as focused if all of these are below the threshold. If too many markers are out of focus, we abort with an error requiring that the focus be corrected.

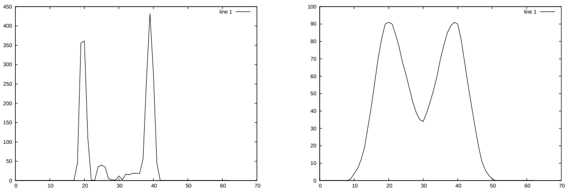

Figure 6: Gradient across marker.

The focus is checked for all cameras during calibration and rechecked on a opportunistic basis during the course of a treatment session.

4.2. Calibration Verification

There are a variety of events which occur between the calibration of the cameras as part of the maintenance routine, and the use of the camera system. Thus we need to have some method to confirm that the camera calibration is still valid. The ideal solution would be to have some set of fixed points, observable by the cameras, whose stereo reconstructed position could be checked after calibration and before each treatment session. Unfortunately, due to the tight working environment, and the need to be able to manoeuvre the patient freely, we do not have space for this.

We do however, have multiple views of the markers on the patient mask. Thus we can at least confirm that the camera calibration is still consistent between the different cameras. As several markers will be observed by multiple camera pairs, we can confirm that the reconstructed positions of these markers all agree to within acceptable limits.

4.3. Rigidity Constraints

Although our marker detection code uses several tests to try and eliminate spurious markers due to lighting reflections and similar effects, the computational constraints limit the amount of effort that can be spent testing the markers. Thus we still have the possibility that spurious markers may be identified at some point in the process.

If such a spurious marker (due to reflection for instance is observed by multiple cameras, the system will try and reconstruct the 3D position of this point and treat it as a marker on the mask. Since this point does not correspond to a marker, spurious markers can thus impact on the accuracy of the system. Alternatively, the stereo code can potentially match a marker on one camera view to a spurious marker in another camera view, although the geometric constraints on the stereo matching process will ensure that there are only a very limited number of positions in which this can occur.

Both problems can be addressed by noting two important properties. If the patient is at rest, a spurious detected marker should not have a significant impact on the system. We are monitoring patient motion, and thus an additional marker that does not move poses no great problem.

When the patient is moving the spurious marker should be easily detectable as it should violate the constraint that the markers are fixed to a rigid body, and thus the relationship between this marker and the true markers will not be consistent with the rigid body assumption. Consequently, by monitoring the relative distances between the markers after the stereo reconstruction phase, we can easily eliminate markers that violate the rigidity constraints.

5. Conclusions

In this paper, we discuss various techniques we use, both in procedural specifications and in system design and implementation, to ensure that we can verify the correctness of our system. This verifiability is an essential component of any system designed for medical applications.

6. References

[1] S. Webb, *The physics of three-dimensional radiotherapy: Confromal radiotherapy, radiosurgery and treatment planning,* Insitute of Physics Publishing, Bristol and Philadelphia, 1993.

[2] R. van Rooyen, “Fast, robust detection of circular retro-reflective targets,” in *Proceeding of the Fourteenth Annual Symposium of the Pattern Recognition Association of South Africa*, Nov. 2003, pp. 21–26.

[3] Evan de Kock, Brian O’Kennedy and Neil Muller, “Calibrating a stereo rig and ct scanner with a single calibration object,” in *Vision, Modeling and Visualization 2002*, G. Greiner, H. Niemann, T. Ertl, B. Girod and H.-P. Seidel, Eds., 2002, ISBN 3-89838-034-3.

[4] Evan de Kock, Neil Muller, Denys Maartens, Jan van der Merwe, Deon Muller, Ruby van Rooyen, Andre van der Merwe, Jan Eksteen, Neil von Hoesslin, Dirk Wagener and Jan Hough, “Integrating an industrial robot and multi-camera computer vision systems into a patient positioning system for high-precision radiotherapy,” in *35th International Symposium on Robotics*, Mar. 2004.