$\mathcal{P}3$: a three–dimensional digitizer prototype*

J.Serrat, X.Varona, A.López, X.Roca, J.Villanueva
Computer Vision Center and Departament d'Informàtica
Edifici O, Universitat Autònoma de Barcelona
08193 Cerdanyola, Spain.
`joans@cvc.uab.es`

Abstract

$\mathcal{P}3$ is a prototype for automatic optical three-dimensional digitizing of small objects. That is, given a free-form object, it is able to obtain the x, y, z coordinates of points uniformly spread all over its surface. When this cloud of points is dense enough, a mesh of triangles can be derived which approximates quite faithfully the three-dimensional shape of the object. The underlying digitization principle is the well known active triangulation through structured light, more specifically, a line light generated by a laser diode.

Keywords: active triangulation, calibration, mesh, hearing aid.

1 Introduction

This project has been developed under contract with a company devoted to hearing aid solutions. The aim was to build a three–dimensional digitizer prototype for ear channel moulds. They are necessary to build custom intra channel hearing aid devices, that is, to fit the particular anatomy of each subject. Presently, the manufacturing process of shells is fully manual (figure 1) and the three-dimensional digitization would allow in a near future to address the automatic fabrication of shells by means of rapid prototyping machines or 3D printers.

Three–dimensional digitization, in a broad sense, means to be able to build a 3D solid model of the scanned part, a sort of reverse engineering. But this is a two–steps process : first, coordinates of points on the objects surface must somehow be obtained ; second, these points are the input of a solid model building method, that fits a surface to them. We have just addressed the first step, because there are diverse techniques, implemented in commercial and public domain programs, to perform the second one.

---

* This work has partially been funded by grant PETRI 95–0407–OP. The CVC spin–off VyRA integrated most of the hardware components

Section 2 is about the design considerations and physical components. Next section addresses the perhaps main difficulty we came across, namely, the calibration of the camera and several moving devices. Finally, in section 4 we present some results.

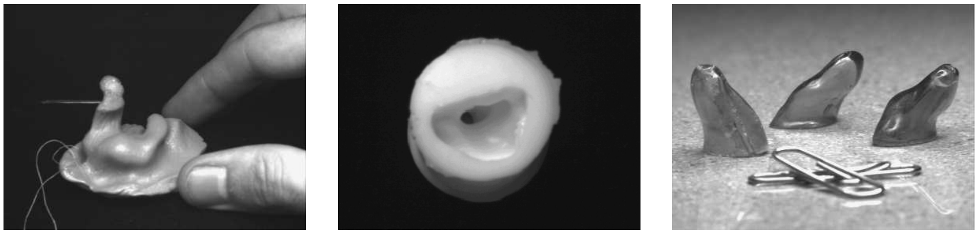

original mould negative hearing aid shells

Figure 1: Present manufacturing process

2 System design

The system design has mainly been driven by the customer requirements, namely : 1) accurate digitization of small moulds (roughly within a 6cm^3 cube), 2) uniform surface coverage, that is, no hidden parts are allowed, 3) fast scanning (under 10 minutes per part), 4) fully automatic digitization, 5) off–the–shelf or standard components, and 6) extremely simple user interface. In addition, it must be able to deal with deep concavities.

As mentioned above, the 3D digitization technique adopted is active triangulation. The reason is its simplicity, both conceptually and with regard to its implementation. The idea is that, once a camera is calibrated, we know, for each pixel $\mathbf{m}_i$ in the image, the equation of the line that passes through the lens optical center and hits that pixel. Thus, all points in the scene along that line would be imaged at that same pixel (figure 2). In order to avoid the ambiguity in the determination of the three–dimensional coordinates of such a scene point $\mathbf{M}_i$, we only need to add a further constraint, namely, that it belongs to a certain known plane. This can be guaranteed just by generating that plane with a laser diode and a cylindric lens. Now, points $\mathbf{m}_i$ are distinctly perceived in the image because laser light is quite thin and very intense even though the lens iris is almost closed. Note however, that this idea requires the line and plane equations to be known with precision (see section 3).

Yet, calibration does not assure the digitization of the whole object, but only the computation of the x, y, z coordinates of those points on the object's surface

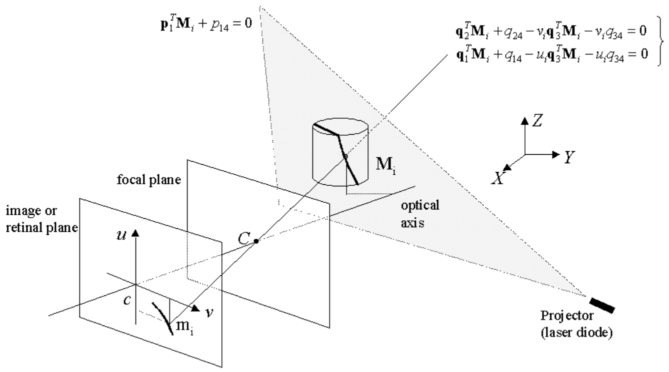

Figure 2: Active triangulation principle on a pin–hole camera

that can be viewed by the camera *when* the laser line hits them. Therefore, it is obvious that the object or the line light plane should move. The first idea in order to get the maximum coverage of the object's surface was to rotate it around a pan (vertical) and a tilt (horizontal axis), while keeping the laser fixed 4a. However, this proved insufficient and large gaps remained undigitized. Besides, concavities are very difficult to reach at the same time by both the camera and the laser line. The solution was to rotate also the laser line, so that for each pan and tilt position, a first surface mirror which reflects the laser line, rotates around a vertical axis, thus scanning the object face visible to the camera. In this way, almost no hidden regions are left to explore, even if its shape presents concavities or complex folds. There is a trade–off between coverage and both digitization time and unnecessary point density at planar or convex regions. This has lead us to the following settings : 4 pan angles, 3 tilt angles and about 30 mirror angles, for which the digitization takes 12 minutes, producing thousands of points.

Figure 3a shows the external aspect of the digitizing hardware, presented as a 'black–box' to the final user. He or she has only to fix the object on the pan–tilt unit within it, through a small gate on the cover. Once closed, the digitization proceeds without further intervention.

In figure 3b we see all the hardware components. Besides the most visible (camera, laser, mirror, pan–tilt unit) there are other 'hidden' components worth to mention. Each axis is rotated by a standard step motor. In order to have finer steps, the pan motor is coupled with a gear box which yields 2400 steps per cycle. The mirror axis needed much more steps because only a small range of angles is of interest. Therefore, instead of a gear box, a microstep driver was used, which splits electronically each actual step in tens or hundreds of 'virtual' steps. The mirror and the tilt motors need to be reset at the beginning of operation and between scans. To this end, an optical reset sensor for the mirror axis and a mechanical one for the tilt platform were included. The sensors, motors and step driver are commanded from the PC through the serial port. Simple commands like `pr-200` (anti–clockwise 200 steps of pan motor) or `mz` (mirror reset) are sent to a C–programmed micro controller which generates the corresponding trains of pulses for the motors and polls the reset sensors.

3 Calibration

The camera is calibrated according to the well known Tsai's method [2]. Even though it is surely not the best one, there is a public domain, well documented, complete and robust implementation in C[3]. As other methods, it only needs a list of corresponding image and scene point coordinates in order to compute the values of the extrinsic and intrinsic parameters [1]. To this end, we have built a pattern of black solid circles and a linear axis which moves it in front of the camera (figure 4b,c). As the distance among circle centers is known, we only have to take a sequence of images while the pattern approaches the camera, at 1cm or 0.5cm steps, and then compute with sub-pixel accuracy the location of the circles centers in the images.

Now, we are ready to calibrate the laser planes, that is, to obtain each plane equation according to the camera's reference coordinate system, for all planes spanned by the different mirror angles. This can easily be done by taking a new series of images, where the laser line hits a plane that is shifted by the linear axis unit along the z axis. For each known z value, an image for each mirror position is taken. Then, given the u and v image coordinates of a pixel that has received the laser light and the value of the corresponding z, we are able to compute its x and y spatial coordinates. The set of all (x, y, z) points lying on a laser plane are input to a minimum square error function to compute that plane's parameters (figure 5a).

The pan and tilt axes have to be calibrated, too. The reason is that when we are digitizing the points illuminated by the laser line, we need to transform their coordinates back to the coordinate system of the first pan and tilt position. That is, we must perform an inverse rotation around each axis. In order to compute the vector in space that is the rotation axis, we digitize three sets of points, each one on a plane rotated around that axis. Realize that we can do this because the camera and the laser planes have been previously calibrated. Three planes define two bisector planes, whose intersection line is the rotation axis (figure 5b).

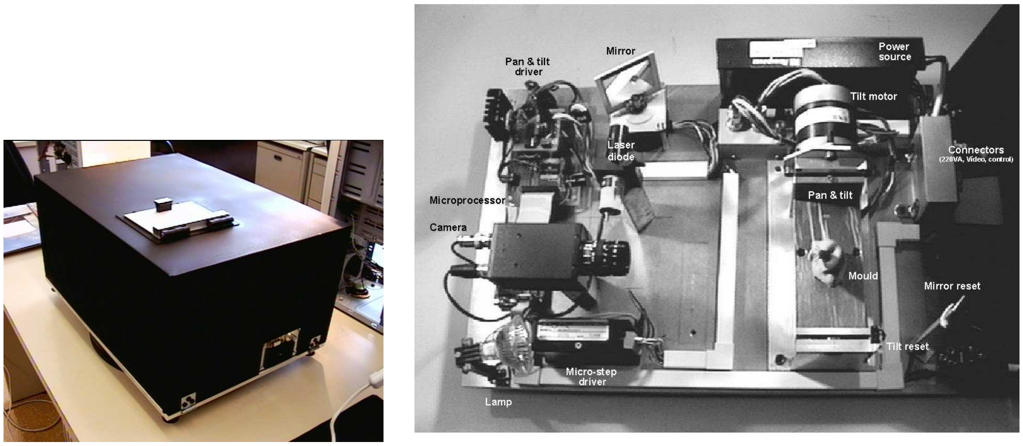

Figure 3: External and internal views of the digitizing unit

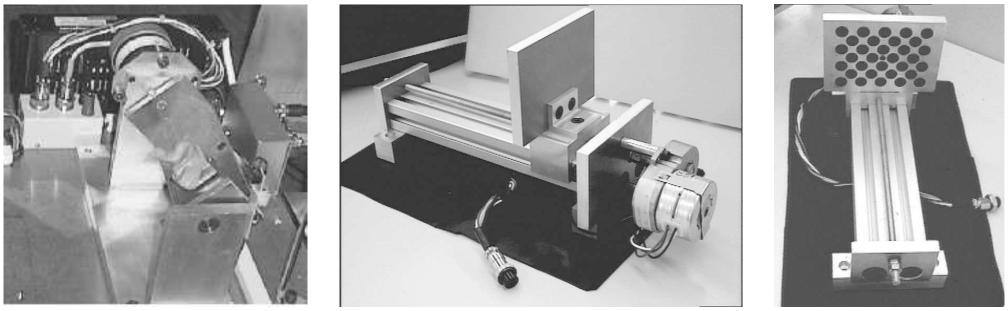

a b c

Figure 4: a) pan–tilt unit, b) linear axis unit for camera and laser–mirror calibration, c) pattern for camera calibration with Tsai's method

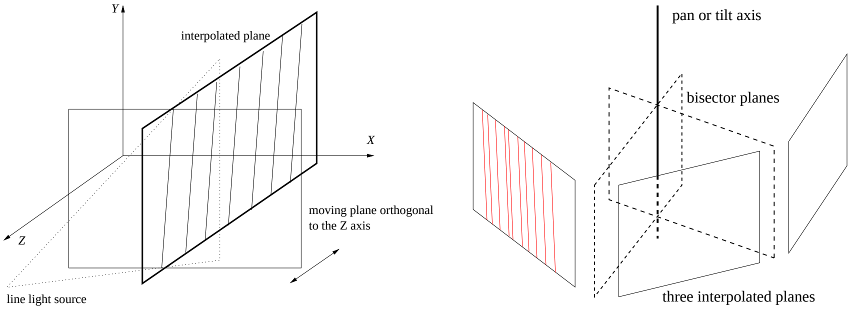

a b

Figure 5: (a) plane calibration of a single mirror angle by moving a plane surface along the z axis, (b) pan and tilt axes calibration by intersection of bisector planes.

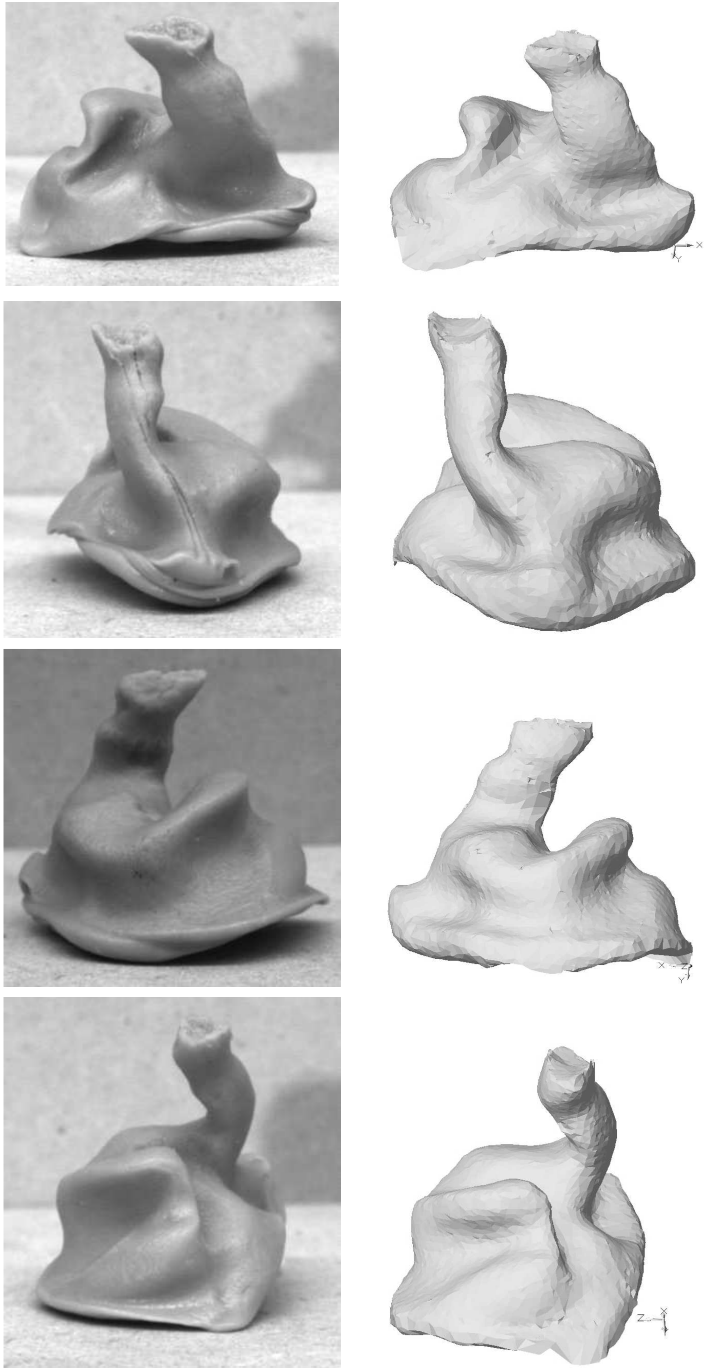

Figure 6: Views of an actual and digitized mould

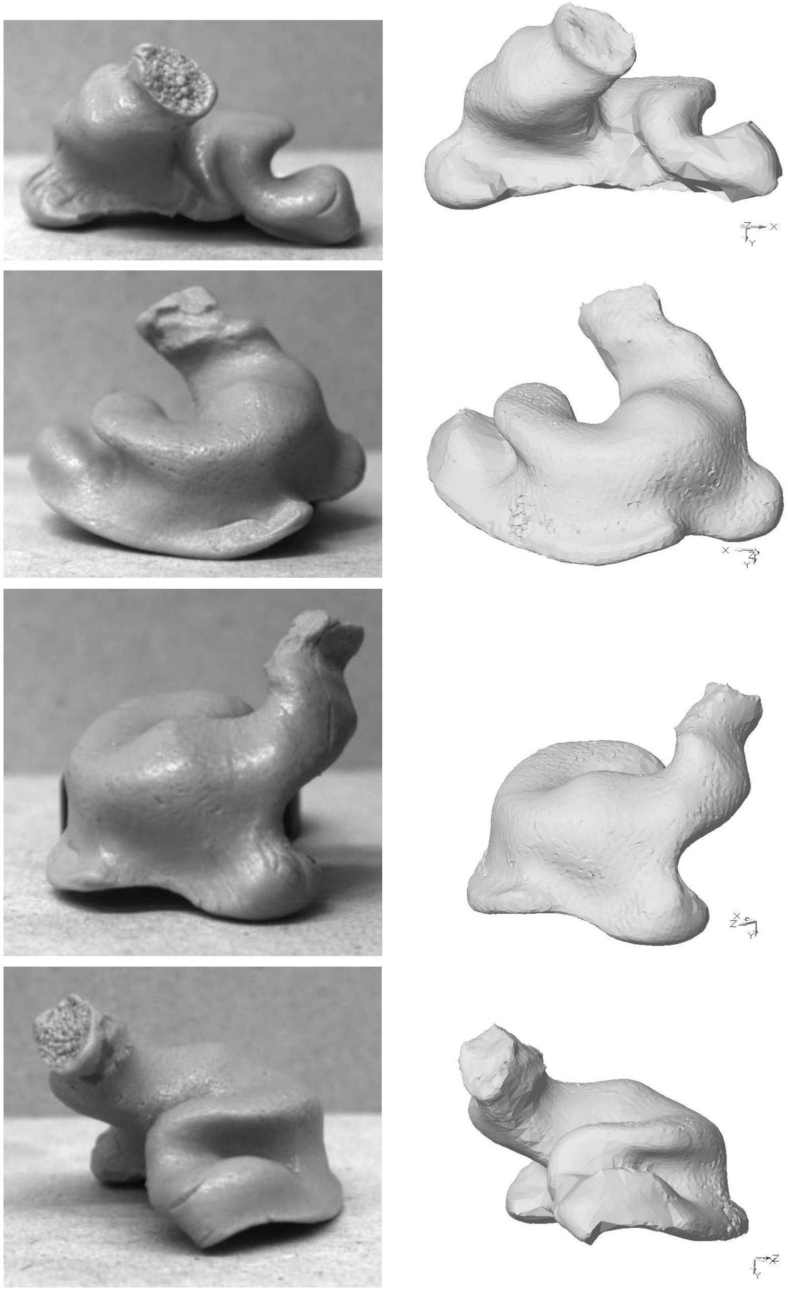

Figure 7: More results for a second mould

4 Results

The output of this system is just a text file containing the coordinates of a cloud of points in the three-dimensional space. Now, we must input this file to some program which builds a solid model. We have tested our system with Wrap from Geomagic (`www.geomagic.com`). It is able to generate a triangular mesh through the α–shapes algorithm, closely related to the Delaunay triangulation. Due to system error sources, these models have to be smoothed in order to achieve a more realistic appearance. The main source of errors are the step motors and the gear box. They suffer from small mechanical imprecissions which affect the quality of the digitization because there is a difference between nominal and actual pan and tilt rotation angles. All in all, the mean error estimated in the validation phase is 0.2mm, and the maximum (but much less frequent) is 0.5mm.

Figure 8 shows the model building steps with Wrap. Figures 6 and 7 compare several faces of an actual mould and renderings derived from its model. Note that we are able to obtain a very faithful representation, although micro-texture and very small details are lost (anyway, they were not relevant for the ear mould application).

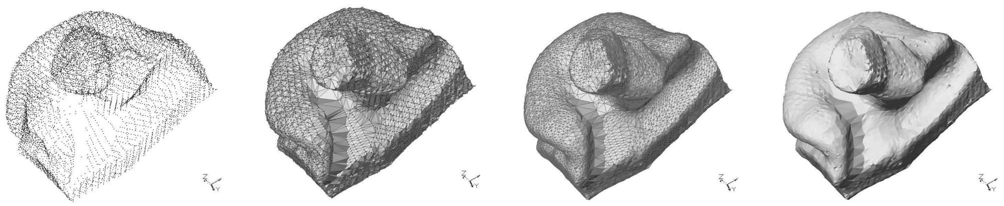
a b d c

Figure 8: (a) sampled (20%) point cloud, (b) triangulation with the α–shapes algorithm, (c) smoothing, (d) same than (c) without edges.

References

[1] O. Faugeras. *Three-dimensional computer vision : a geometric viewpoint.* Chapter 3. MIT Press, 1993.

[2] R.Y. Tsai. A versatile camera calibration technique for high accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE Journal of robotics and automation*, Vol 3 No. 4, 1987.

[3] Tsai's method web page. Reg Wilson.
`www.ius.cs.cmu.edu/afs/cs.cmu.edu/user/rgw/www/TsaiCode.html`