

UNIVERSITY OF UTAH GRADUATE SCHOOL

SUPERVISORY COMMITTEE APPROVAL

of a dissertation submitted by

Henry Fuchs

I have read this dissertation and have found it to be of satisfactory quality for a doctoral degree.

June 3, 1975

Date

Robert P. Plummer

Robert P. Plummer
Chairman, Supervisory Committee

I have read this dissertation and have found it to be of satisfactory quality for a doctoral degree.

June 3, 1975

Date

Steven V. Coons

Steven V. Coons
Member, Supervisory Committee

I have read this dissertation and have found it to be of satisfactory quality for a doctoral degree.

June 3, 1975

Date

Elliott I. Organick

Elliott I. Organick
Member, Supervisory Committee

FINAL READING APPROVAL

To the Graduate Council of the University of Utah:

I have read the thesis of Henry Fuchs in its final form and have found that (1) its format, citations, and bibliographic style are consistent and acceptable; (2) its illustrative materials including figures, tables, and charts are in place; and (3) the final manuscript is satisfactory to the Supervisory Committee and is ready for submission to the Graduate School.

June 30, 1975
Date

Robert P. Plummer
Robert P. Plummer
Member, Supervisory Committee

Approved for the Major Department

Anthony C. Hearn
Anthony C. Hearn
Chairman/Dean

Approved for the Graduate Council

Sterling M. McMurrin
Sterling M. McMurrin
Dean of the Graduate School

ACKNOWLEDGMENTS

I am deeply grateful to my thesis supervisor, Robert Plummer, who has been a constant source of assistance and encouragement throughout the course of this project.

I want to express my appreciation to the other members of the committee, Elliott Organick and Steve Coons, for many hours of discussions, suggestions and insights -- as well as for responding so quickly and constructively to early drafts of this dissertation.

I would like to thank Ivan Sutherland, not only for being the initial inspiration for the project, but also for serving as a committee member while he was still at Utah.

I would like to express my appreciation to Rich Riesenfeld and Elaine Cohen for help and enthusiasm, especially at crucial periods in this project.

It should be noted that it is Mike Milochik to whom all praise is due for the excellent photography, in spite of some less-than-excellent sources from which to work.

Thanks are due to the many friends and fellow students, who willingly lent their bodies to be digitized in a none-too-comfortable environment.

Finally I would like to thank my parents for decades of faith, support, and encouragement -- without which none of this would have been possible.

TABLE OF CONTENTS

Acknowledgements	iv
Abstract	vii
Chapter 1 Introduction	1
1.1 Problem Statement	1
1.2 A Sampling of Previous Methods	2
1.2.1 Direct Manual Measurement	2
1.2.2 Mechanical Moving Devices	4
1.2.3 A Holographic Method	6
1.2.4 A Moire Method	8
1.2.5 Multiple 2-D Images	8
1.2.6 Controlled Illumination on Objects	14
Chapter 2 Design Philosophy	17
2.1 Hardware Design Considerations	19
2.2 Analysis System Considerations	20
Chapter 3 The Hardware Sensing System	21
Chapter 4 Data Acquisition, Analysis and Object Reconstruction	21
4.1 Basic Data Acquisition Programs	21
4.2 Analysis and Object Reconstruction Methodology	27
4.2.1 Sorting	28
4.2.2 Coherence	28

4.2.3 Applicability to the Present Situation	28
4.3 Description of Analysis and Reconstruction Algorithm	30
4.3.1 Analysis on Each Cutting Plane	33
4.3.2 Inter-Plane Reconstruction	37
Chapter 5 Conclusions and Further Development	47
5.1 Conclusions	47
5.2 Further Development	49
5.2.1 Hardware Improvements	49
5.2.2 Improved Analysis and Reconstruction Methods	51
Appendices	53
A. Descriptive List of Major Software Modules	53
B. Data File Formats	55
C. Sample Program Execution with Commentary	66
References	84

ABSTRACT

Described are the design and implementation of a new range-measuring sensing device and an associated software algorithm for constructing surface descriptions of arbitrary three-dimensional objects from single or multiple views.

The sensing device, which measures surface points from objects in its environment, is a computer-controlled, random-access, triangulating rangefinder with a mirror-deflected laser beam and revolving disc detectors.

The algorithm developed processes these surface points and generates, in a deterministic fashion, complete surface descriptions of all encountered objects. In its processing, the algorithm also detects parts of objects for which there is insufficient data, and can supply the sensing device with the control parameters needed to successfully measure the uncharted regions.

The resulting object descriptions are suitable for use in a number of areas, such as computer graphics, where the process of constructing object definitions has heretofore been very tedious. Together with the sensing device, this approach to object description can be utilized in a variety of scene analysis and pattern recognition applications which involve interaction with "real world", three-dimensional objects.

CHAPTER 1

INTRODUCTION

1.1 Problem Statement

Researchers in seemingly-diverse areas are often concerned with the acquisition of object descriptions. In artificial intelligence, for instance, a large part of most scene analysis systems is devoted to generating a description of objects in the system's working environment, whether this be a table-top scene of toy blocks, a rocky Martian surface, or a work-station on an auto assembly line.

In computer graphics, much time is spent attempting to create accurate pictorial images of real and imaginary objects. While the descriptions of imaginary objects are often created with the aid of an associated computer-aided design system, the descriptions of real objects usually has to be generated by laborious, largely-manual measurement techniques.

The interest in object descriptions is not limited to computer users. A prosthesis manufacturer may want to match the new artificial leg with the user's natural one, but they may not have the facilities to take more than a few basic measurements.

Researchers in artificial intelligence (specifically robotics) have been among the ones most intensely involved in the development of systems for the automatic acquisition of object descriptions. Most of their systems, however, have relied on a picture-oriented sensor, usually a TV camera. This report hopes to demonstrate that a significantly different kind of sensor, a computer-controlled rangefinder, may also prove useful for

some of these tasks. The design and implementation of such a rangefinding system is described. To demonstrate the feasibility of this approach, a simple scene-analysis algorithm is implemented, which can generate, solely from the range data, descriptions of objects in the sensor's field of view. It is hoped that this research will stimulate other attempts at sensing systems more readily adaptable to the computer than the human-oriented TV camera.

1.2 A Sampling of Previous Methods

1.2.1 Direct Manual Measurement

The most elementary method of digitizing objects is by direct, manual measurement. With the aid of yardsticks, plumb lines and calipers, a great many solid objects can be successfully measured, and the set of values later input to a computer system.

This idea of being able to specify an arbitrarily complex three-dimensional object with a set of simple measurements is hardly a recent development. The Renaissance artist Leon Battista Alberti, in his book *Della Statua*, published in 1440, describes a method for the accurate measurement of the human form (Figure 1-1). He claimed that by using his method, different parts of the same statue could be constructed at separate places and would still be able to fit together [10].

Today's approach, still basically the same, is often to mark all points of interest -- "key" points -- on the surface of the original object and then measure the distance of each of these points from a common reference position (Figure 1-2). The surface is then defined as a topological net over these key points. Of course, many tedious hours must be spent to carefully measure the position of each selected point on the original



Fig. 1-1: Alberti's "Definer" (from [10])



Fig. 1-2: Manual measurement today (from [18])

object. The results are, however, often surprisingly effective. Although this method is not practical for serious, large-scale digitizing, it should be noted that it has several advantages over the other more sophisticated methods. Obviously, it requires almost no equipment -- hence no cost, except of course for manual labor. The resulting descriptions also tend to be very compact, since the user naturally wants to minimize the number of points that he has to measure. Although the less compact descriptions resulting from the more sophisticated automatic methods can often be trimmed in size according to some algorithm, it turns out that the subjective criteria used by humans are usually more effective.

1.2.2 Mechanical Moving Devices

An obvious next step to the simple manual approach is to substitute some computer sensing device for the user's calipers and yardsticks. The human user still has to define the surface points and their interconnections, but now he can just move some pointer around the object and tell the computer when the "current" position is of interest.

One device of this type is the so-called three-dimensional "crane" (Figure 1-3). This is a mechanical arrangement of rods and gears which allows sliding movement in each of the three axial directions. Through the amount of turning of each of three gears, the computer can calculate the distance extended along each axis. The user simply positions the tip of the crane's arm to a point of interest and instructs the computer --through a foot switch in this case-- to note the current position. Although this method is much faster than the completely manual approach, it is still very time-consuming. More serious is the severe limitation to the range of object sizes which can be measured. Being a mechanical device, there are also considerations of its bulkiness and the inertia and slippage of its moving parts.




Fig. 1-3: Direct measurement with a 3-D "crane" (from [18])



Fig. 1-4a: Detail of fishing-line digitizing unit (from [20])



Fig. 1-4b: Fishing line-connected 3-D digitizing handle ("wand") resting on tripod (from [20])

A large variety of devices of this same general type have been constructed. One such device has been in use at the University of Utah for a number of years [20]. It uses three separate spring-loaded fishing-reel/fishing-line units (Figure 1-4a). These three assemblies are placed around the top corners of the working volume and the ends of the three fishing lines are all connected to the tip of a pointing device (Figure 1-4b). From the amount of rotation on the shaft of each reel, the length of fishing line rolled out can be calculated. Assuming that the lines are unobstructed, the three-dimensional position of the pointer tip can be calculated from the three separate lengths of the fishing lines.

1.2.3 A Holographic Method

Gara, Majkowski and Stapleton of General Motors Research Laboratories report the development of a novel new digitization technique [8]. Although their system does not have direct applicability to the interaction-oriented scene analysis applications, it may provide a solution to the off-line object-digitization problem.

Their method consists of first taking a controlled, high-quality holograph of the object of interest, then extracting surface measurements from the developed holograph with a special-purpose computer-controlled video viewing system. The surface measurements are calculated by moving the video detection system about the object's holographic real image. As seen in Figure 1-5, there is a large angular orientation between the face of the video detector and the object's surface image to allow both in and out of focus parts of the image to hit the video detector's surface. The intersection of the object's surface and the video detector's face is the locus of points on the detector face at which the image is in sharpest focus. Figure 1-6 shows a typical video image from the detector. (To aid in this focus-determination process, an optical




Fig. 1-5: Orientation of video detector to holographic real image




Fig. 1-6: Image from video detector showing in- and out-of-focus areas




Fig. 1-7: Signal from part of one scan line, showing a region of focus (all three figures from [8])

interference pattern was projected onto the object's surface when the holograph was taken.) Figure 1-7 shows the video signal from one scan line, the point of optimum focus being at the peak of the signal's envelope. After the optimum-focus locations are determined in a single video image, the system incrementally moves the video detector in an attempt to track the object's surface contours. In this way, all visible surfaces of the object can eventually be measured.

1.2.4 A Moire Method

Speight, Miles and Moledina report the application of a Moire method (suggested by H. Takasaki [19]) to the 3-D measurement of slaughtered animal carcasses [17]. Figure 1-8 shows an overhead view of the geometric arrangement of camera, flash-gun light sources, sliding grid and the carcass of interest. The resulting photographic image contains contour lines each of which is of equal depth from the grating plane (Figure 1-9). Although the actual digitization in the reported system was largely a manual operation, there do not seem to be any theoretical obstacles to the automatic processing of these contour maps. The main limitation to applying this method to the more general object-description problem may lie in the method's difficulty in accurately capturing complex, detailed, rapidly-varying surface structures.

1.2.5 Multiple 2-D Images

Acquisition of three-dimensional information from multiple two-dimensional photographic images is not just a widely used digitization method, but is the basis for an entire technical field, stereo-photogrammetry -- most likely inspired by the human stereo vision system.




Fig. 1-8: Geometry for generating Moiré patterns (from [17])




Fig. 1-9: Moiré contour maps of 4 views of a lamb carcass (from [17])

The geometry is beautifully simple. If a picture of a three-dimensional environment is considered to be an image drawn on a window-pane in front of the viewer's eye, then a point on that picture must correspond to a spot in the environment somewhere along the line defined by the viewer's eye and the point on the image.

Given another eye-image pair at a different orientation to the object, and assuming the point of interest is in view in both images, the point's three-dimensional position is simply at the intersection of the two lines of sight (Figure 1-10).

A variety of methods are based on this simple idea. A common technique consists of marking the points of interest on the object itself, then taking pictures from at least two different viewpoints (Figure 1-11). If the camera/eye positions and orientations are not known, they can be calculated from the correspondence between the picture positions and the known 3-D positions of a number (at least 6) of "reference" points in the object's environment [14].

When marking the subject is not practical, other methods can be employed. The common practice is to take two pictures from locations only a small distance from each other -- similar to the two human-eye views. An operator then looks at these images through a suitably adjusted stereo viewer and perceives the three-dimensionality of the object. By moving a pointer in each view until they "merge" in the virtual three-dimensional environment, he performs the correspondence which previously consisted of manually marking the object. From the X,Y distances of the pointers in each image, the three-dimensional position of the perceived point can be calculated [10]

An obvious advantage of this viewing approach is that an indefinitely large number of points can be digitized, since with the marking method, only the actual points marked can be measured. But with a stereo viewer, the accuracy of the measurements depends not




Fig. 1-10: Triangulation from two images




Fig. 1-11: Actual digitization using two images with a data tablet
(from [18])

only on the interocular distance, but also on the visual acuity -- and the patience ! -- of the operator.

Several attempts have been made to eliminate the need for the human operator to specify the correspondence for each point to be measured. Levine et. al. [12], expand on an earlier algorithm of Julesz[11] which is based on the observation that when two views of a scene are taken from nearby positions, relative to the objects, then the differences between corresponding parts of the two images is largely an X-axis offset, with the amount of offset related inversely to the distance of the object from the viewer (see Figure 1-12).

The technique then, is to digitize the two images and cross-correlate parts of corresponding scan lines. The X offset of the best fit can be used to calculate the three-dimensional position of the point defined by the center of the two matching scan-line segments.

Some initial success has been reported with this method. The obvious difficulty is that the viability of the offset-difference assumption (due to the depth-variation of the object surface) is often inversely related to the distance of the object from the viewing position; the assumption is reasonable for distant or flat regions where the view from both eyes is essentially the same, but it is often invalid for close-by objects, as with the face of a person, for whom one view may contain one side of the nose, and the other view may contain the other side. On the other extreme, if the image of the local surface is too similar in both images -- e.g. flat side of a building, a new sidewalk -- then there will also be difficulties due to a lack of characteristic features with which to achieve a high cross-correlation.

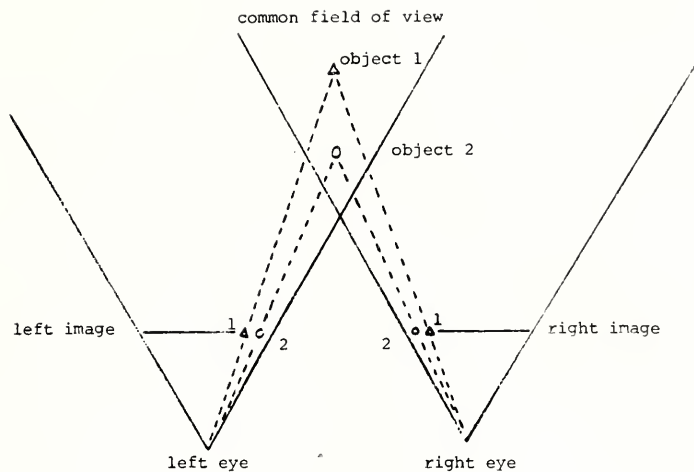




Fig. 1-12: Cross-correlation of stereo images for depth determination



(inputs)



(output)

Fig. 1-13: Object reconstruction from projected silhouettes (from [4])

As part of a more extensive project on computer vision, Baumgart[3,4] demonstrates a technique of reconstructing 3-D objects from their image silhouettes. The geometry is similar to the triangulation technique in Figure 1-10, except that in this case instead of line-of-sights being projected, the cone-shaped projections of a silhouette are mapped into the object space. The object by definition is constrained to lie entirely within each and every one of these projections. Baumgart describes it as being like "the old joke about carving a statue by cutting away everything that does not look like the subject." Figure 1-13 shows 3 silhouettes of a plastic horse and a view of the reconstructed object. It is of interest to note that the input views were all from the horse's left side, while the view of the reconstructed object is of the horse's right side. Due to the projective nature of this method, however, surfaces with full concavities cannot be adequately reconstructed.

1.2.6 Controlled Illumination on Objects

Methods for extracting three-dimensional information from multiple two-dimensional projections are not limited to considerations of photographs only. If the geometry of Figure 1-10 is reconsidered, it can be observed that a pencil-beam of light can replace one of the two lines-of-sights used in the triangulation process. In this way, one of the photographs could be eliminated; the beam of light would be seen -- if not obscured by some object -- as a bright reference point in the remaining photograph. This kind of a system yields itself naturally to automatic processing; the origin and orientation of the pencil-beam of light can be placed under computer control and the photograph can be input as a video picture. If the object under investigation can be examined at length, then an arbitrary number of points on its surface can be digitized (Figure 1-14).




Fig. 1-14: Digitization with one image and a pencil-beam of light

Fortunately, the geometry of such a system is over-solved, and thus can be simplified. Instead of two lines, a plane and a single line are sufficient to uniquely define a surface point (Figure 1-15). A number of investigators have used this kind of a system [1,2,5,13,16]; a plane of light, rather than a pencil-beam, is projected onto the object's surface. In this way, from a single video image, the system can extract not just one surface position, but rather a large number of points along the visible intersections between the plane of light and the object's surface.

The method developed for the present system is in some ways the inverse of the above "plane-of-light and single video image" design. The system is described in Chapter 3.




Fig. 1-15: Digitization with one image and a plane-beam of light

CHAPTER 2

DESIGN PHILOSOPHY

It is important in design considerations to review not only what the present limited system may be able to do, but also to delineate the scope of capabilities desired for the eventual "ideal" system.

is the basic goal of this research? It is to design a system which can easily acquire three-dimensional data and use it to construct surface descriptions of arbitrary three-dimensional objects. The general idea is to build a system in which the hardware sensor(s) and the software analysis algorithms interact to produce a more capable system than would be possible without this interaction.

The desired system would have a sensor whose orientation to the object(s) could be altered to allow input data from various views of the object. This could be accomplished in a number of different ways. There could be several sensors mounted at strategic locations around the system's environment. There could be a system-controlled device -- an arm, a turntable -- which could move the object. The sensor itself could be movable -- on a track, on a computer-controlled arm, or mounted on a moving platform. Of course, the particular application would influence the configuration design. For example, the "moving platform" model would be the one most likely to be used for a robotics application.

The scenario would go something like this. The object to be scanned is placed in the system's environment. The sensor starts scanning the environment according to some initial control parameters -- scanning the entire environment at a cursory level of

detail, or perhaps scanning only until a close-by object is encountered. The analysis system -- let's call it Analyzer -- processes this initial scan data and begins to construct its object descriptions. These not yet being complete, Analyzer calculates the control parameters the sensor needs to obtain the additional input data. The sensor again gathers some data, now according to the new specifications. Analyzer processes the new scan data and integrates it into its developing object-description structures. It again determines whether it needs additional input data. If it does, it again calculates the sensor control parameters. Again, the sensor is instructed to obtain more data, according to its new set of control specifications. This interaction between the sensing device and Analyzer continues until some "completeness" criterion in Analyzer is satisfied.

This approach has several advantages over a simpler method. First, only the data which is needed is actually acquired by the sensor. In this way, neither the sensor nor Analyzer is burdened with unnecessary data. The level of detail can now vary with the specific application. If the task is to give object descriptions to help navigate a robot through an obstacle-filled environment, then one or two requests to the sensor may be sufficient. If, on the other hand, the task is to digitize some arbitrary object for a computer graphics system, then Analyzer could successively direct the sensor to those regions of the object which have yet to be charted with adequate detail. The level of input detail could also be context-dependent. If the objective is to inspect an assemblage for missing bolts, then the level of detail could be modified as Analyzer directs the sensor towards the regions of interest -- after, of course, it has located these regions from earlier inputs.

