Graduate School Form 9
(Revised 8/89)

PURDUE UNIVERSITY
GRADUATE SCHOOL
Thesis Acceptance

This is to certify that the thesis prepared

By Charles Scott Foshee

Entitled

Goal Driven Three Dimensional Object Inspection from Limited View Backprojection Reconstruction

Complies with University regulations and meets the standards of the Graduate School for originality and quality

For the degree of Doctor of Philosophy

Signed by the final examining committee:

______________________, chair

______________________ co-chair

____

____

Approved by:

______________________ 10/1/90
Department Head Date

This thesis ☐ is ☒ is not to be regarded as confidential

____
Major Professor

GOAL DRIVEN THREE DIMENSIONAL OBJECT INSPECTION FROM LIMITED VIEW BACKPROJECTION RECONSTRUCTION

A Thesis
Submitted to the Faculty

of

Purdue University

by

Charles Scott Foshee

In Partial Fulfillment of the
Requirements for the Degree

of

Doctor of Philosophy

December 1990

this is dedicated to my family

ACKNOWLEDGEMENTS

I would like to acknowledge all the support I received from all my friends over the long months during which there seemed to be no end to this effort. I would especially like to thank my family for their never ending support and encouragement. It was their dedication to me which gave me the freedom to be committed to the completion of my research.

TABLE OF CONTENTS

Page

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Foshee, Charles Scott. Ph.D., Purdue University, December 1990. Goal Driven Three Dimensional Object Inspection From Limited View Backprojection Reconstruction. Major Professors: Edward J. Delp and Jan P. Allebach.

The automatic inspection of three dimensional objects is an important part of manufacturing. This process will allow detection of defects relating to the size and shape of parts during manufacturing and assembly. Tomography techniques which are currently in use in medical applications are too expensive and computationally intense for use in manufacturing. These techniques also provide more information about the object than is typically required for inspection applications. For example, a machined part may only need to have certain exterior shape characteristics for it to be properly manufactured.

The inspection problem will be approached using backprojection reconstruction techniques. This basic technique will be modified into a goal driven algorithm to inspect certain locations on the surface of the object which we shall call landmarks. These landmarks could be derived from the CAD phase of object design. The use of these landmarks will allow the inspection system to use a limited number of views in its analysis. The actual reconstruction process will make use of a modified volume projection technique which will be employed to make maximal use of the information incorporated into the views of the object.

A related problem to the inspection analysis is the interpretation of the results of the inspection. All information derived during inspection analysis is three dimensional in nature. A new technique for the interpretation of the inspection results is proposed which is based on the

goal driven nature of our system. The limitations of inspection accuracy as they relate to the system is also examined.

Before the inspection process can occur, it is necessary for the inspection views of the object to be selected. This selection process is not obvious because it is not apparent how different views interact with each other in their supplying information to the inspection process. A goal driven algorithm is presented which removes the redundancy within a set of needed views of the object. This goal driven view selection is consistent with the goal driven nature of the inspection process.

This entire process assumes that the object is capable of being inspected. An object characterization, called *Weak Convexity*, is presented for objects to be reconstructed using backprojection reconstruction. This characterization is sufficient but not necessary for inspection to be possible. Inspection is possible if the object has this property locally, at each landmark of interest.

CHAPTER 1

INTRODUCTION

1.1 Three Dimensional Computer Vision Background

Computer vision is the process of extracting scene information from sensory input. The form of the extracted scene information depends upon the goal of the computer vision system. Three dimensional computer vision is the process of extracting three dimensional scene or object information from an integration of multiple sensory inputs. This can range from a simple scalar result summarizing the scene to a complex three dimensional scene description which preserves much of the original sensor detail.

One analysis area which lends itself to three dimensional computer vision is industrial manufacturing. A common problem would be the manipulation of a machine part. A three dimensional computer vision system could be used to identify the part, plan the manipulation of a part by a robot, or inspect the part for proper shape. Each of these problems requires the sensing of features which describe the object. These features are then combined or examined, possibly with related scene data, to determine information about the object.

Many other types of problems can be cast as three dimensional vision problems. Some of these problems include object recognition; object motion planning; object shape extraction from stereo, shading, and motion; and shape extraction from two dimensional projections. Many different aspects of these and related problems have been examined. Ballard and Brown, Horn, Marr, and Grimson provide a thorough introduction to the history and recent state of computer vision [7, 39, 51, 82]. Siebers provides an introduction to the three dimensional aspects of computer graphics [114]. Wilson discusses three dimensional geometry as a modeling abstraction for Computer Integrated Manufacturing environments [139].

A variety of sensing techniques have been employed to acquire information about objects and three dimensional scenes. Most investigations have been done using cameras to sense the images used for analysis [100, 123]. Others have used ultrasound [122], X-rays [121], optical serial sections [79], active stripe coding [133], or CAD information [11, 103] to obtain data. In these later cases the technique is usually selected because of some unique properties of the object under investigation. For instance, X-rays and ultrasound may be used for medical applications where the object is contained within the human body.

One major area of interest in three dimensional computer vision is the area of object recognition. The goal of object recognition is to classify the object by type or identity. Besl provides an overview of the three dimensional object recognition problem [10]. He presents a description of the recognition problem and concepts associated with it. Three dimensional object recognition has been approached from a number of perspectives. Some of these include using a single view for feature extraction [49, 78], surface patch descriptions for matching [33], graph matching on surface models [61, 140], image sequences [68, 91], range information [102], evidence approach based upon notable features [60], features derived from multiple views [65], cross-section matching [145], multiple silhouette views of an object [17, 24, 25], model based Bayesian [18], dense range images [4], multiple sensors of a mixed type [44]. It has been pointed out that recognition algorithms should be scale invariant and dependent upon shape alone [90].

Another related area of three dimensional computer vision is the planning of robotic motion within a three dimensional scene. A survey of recent developments in motion planning is presented by Schwartz [108]. Motion planning is the selection of a series of coordinates in three space for the motion of an object. Many different methods have been employed for motion planning. These include a dynamic programming approach

[120], monte-carlo methods [8], search strategies [31], tube path restriction [119], and backprojections [32].

One analysis method, stereo, has been the basis of shape extraction techniques due to its unique analogy to the human vision system. Stereo analysis techniques have been used to extract three dimensional measurements of a scene [142], surface reconstruction [46], detailed scene information [45], and object reconstruction [40]. An argument for the use of structural relationships in the interpretation of three dimensional scenes has also been presented [89]. Stereo analysis using only a single image has been explored [118]. This technique uses a perspective transformation to synthetically generate the second view. An integration of stereo image pair analysis and sequential image analysis has been examined [134]. A key area for the success of stereo analysis is the proper determination of the camera calibration and correspondence between stereo images [3, 77].

Another technique for extraction of object shape is to make use of relative motion between the sensor and object to extract information about the object. It is normal in this analysis to obtain shape and orientation information about the object. Some of the techniques used to extract motion and structure include the use of orthographic views [52, 146] and Kalman filters [74, 141]. Motion in the presence of noise has been examined [136]. Also, a necessary and sufficient condition for the unique determination of object motion under certain conditions has been determined [143].

In the case of object recognition, object motion planning, and object shape extraction; the object was treated as a single item to be identified, moved, or described. In object reconstruction the goal is to recreate the detail of the object in our representation of the object. Many different techniques have been proposed for object reconstruction. Magee presents a review of multiview techniques [80]. Kanatani compares Euclidean and Non-Euclidean methods of three dimensional reconstruction [63].

Tomography [13] is an object reconstruction technique which reconstructs a mathematical descriptions of object surfaces. It can be generalized as the reconstruction of multidimensional functions from their line integral projections. Tomography is an analysis technique which often used non-visible electromagnetic radiation to probe the interior regions of an object for structure.

Other reconstruction techniques may be split into two different categories: orthogonal and perspective reconstruction. In either case, a number of views of the object are analyzed until sufficient accuracy in the reconstruction is achieved.

Orthogonal backprojection reconstruction typically uses data from some form of orthogonal view where there are no unknowns or noise in the system. Engineering drawings present one problem which is fully described by this type of data. Interpretation of engineering drawings by vertice generation has been examined to reconstruct three dimensional objects [11, 22, 41, 103]. Wire frame models have also been reconstructed from this data [117]. Finally, Hand input orthogonal views have also been used for data [36].

The second category relies upon imaged data. The use of imaged data for reconstruction has led to a variety of techniques. These techniques include using optical serial sections [75, 79], feature sets from image sequences [43], successive refinement of stereo depth information [113], sparse data [99], parametric estimation from projections [13], and an integrated method using passive sensing to rough in the object, followed by active sensing to refine the reconstruction [133]. Backprojection reconstruction is an object reconstruction technique which makes use of object silhouette projections to infer outer shape characteristics about an object. It has been widely used and will be discussed in more detail later.

Two aspects common to all reconstruction techniques are registration and conflict resolution. Registration is the process of aligning the different elements of a multisensor vision system. Registration is specifically examined for serial cross-section alignment [86]. Conflict resolution is the process of deciding between two conflicting pieces of information derived from two different sensing sources. A method of conflict resolution in the reconstruction of an object which is overdescribed by a set of lines and vertices has been examined [62].

One specific application of reconstruction is heart ventricle reconstruction. Heart ventricle study is useful in the detection of heart disease. Heart ventricle reconstruction has been examined from surface curvature data [14] and cross-sections from x-ray backprojections [121]. Another application area is industrial inspection [123]. Industrial inspection is the process of analyzing an object for its correct manufacture. This application of reconstruction will be examined later. Our focus, reconstruction through the backprojection of silhouettes, will then be applied to this area.

1.1.1 Areas Related to Three Dimensional Computer Vision

In order for three dimensional computer vision analysis to be performed, there are a number of different intermediate or related steps of the analysis which must be completed without dealing specifically with object analysis. Examples of these include data storage techniques, visibility of one part of the object space from another, and computer decision making.

Three dimensional computer vision requires a more compact internal representation scheme for data than a complete description of the scene. Many different methods have been proposed for three dimensional data representation. Some of these include decomposition into cross-sections [94], prism trees [95], rectangular parallelpipeds [66], boundary models [35], oct-trees and quad-trees [6, 135], and other hierarchical data structures [106, 107].

Beyond developing storage representations for three dimensional objects, there has been a variety of efforts in the manipulation of these data structures for specific applications. Walker presents a discussion of the reasoning which must go into the transformation between two dimensional and three dimensional representation schemes [132]. Chen discusses representation, display, and manipulation of objects for the purpose of medical applications [21]. Mao discusses oct-tree/quad-tree representations of two dimensional images / three dimensional objects and the conversion between the two with application to medical imaging [81]. Finally, geometric solid models have been proposed for the transfer of data between CAD systems [69]. In each of these representations, the representation scheme is chosen to facilitate the data analysis of interest.

Oct-trees provide a unique and compact method of storing three dimensional object information. An oct-tree for volume representation is a hierarchical data storage whose successive levels of branches represent a smaller refinement of the volume being represented. Jackins [59] examines the use and manipulation of oct-trees for three dimensional representation without discussing how they are generated. Region growing in quad-trees has been investigated [111]. Quad-trees construction from an image has been investigated [110]. An introduction to oct-trees and their construction has been presented [20]. An effort to give visible meaning to oct-trees has been given for ray tracing [104] and for line drawings [131]. [105] investigates neighbor finding in oct-trees.

Object analysis must take place in an environment where the object can be viewed by the computer vision system in a manner sufficient for analysis. Very little work has been done in the area of determining characteristics of objects which can be analyzed. Basic work has been performed by mathematicians in geometric analysis. Lee and Avis [5, 72] have investigated the determination of the internal visibility of polygons from edges of the polygon. Lee, Shin, and Fisk [34, 73, 112] have investigated algorithms to determine the kernal of a polygon. Culberson and Tor [29, 125] present algorithms for the generation of convex covering of polygons.

Because of the complexity of the analysis of three dimensional objects it is quite common for three dimensional computer vision systems to be confronted with some form of uncertainty in the results or analysis. In some instances it is possible for some artificial intelligence decision making techniques to be incorporated into the analysis. Smith [115] examines the control of backward inference in the decision making process. Martins [85] examines belief system revision. Lee [71] compares Bayesian and Dempster-Shafer reasoning. Shafer-Dempster reasoning has been applied to identification [12]. Shafer [109] discusses the implementation of Dempster's rule. Hummel [53] presents a Bayesian viewpoint of applications for Dempster/Shafer evidence theory. Ihara [56] presents a conditional probability extension on hypothesis based uncertainty testing. Search strategies have been well investigated [19, 57, 67, 92].

1.1.2 Industrial Inspection

Three dimensional industrial inspection is one example of an application area where it may be necessary to have a complete description of the object available for examination. A computer vision system for the purpose of inspection must be capable of examining objects and building a description of sufficient detail to perform an accurate inspection. In general, any part of the object might be significant to the object Thus, all parts of the object must be accessible for examination by the vision system. Different computer vision techniques are better adapted to examine different parts of object.

For instance, silhouettes can readily provide outer contour information of convex objects but can not provide details about concave portions of an object. Stereoscopic analysis can provide depth maps of concave portions of the object but are limited to areas which are in view of both image viewpoints. Limitations and strengths of each analysis technique available must dictate the selection of an appropriate technique.

Inspection analysis requires a comparison between the complete descriptions of a known desired object and a reconstructed unknown object to be performed. The comparison must make use of a distance measure to compare corresponding elements of the models to yield a meaningful description of differences. A meaningful description of the differences is reliant upon a priori knowledge of the significant parts of the object that are to be compared.

The computer vision analysis employed must be able to supply the detail required for the aforementioned inspection. There are two parts to achieving this goal. First, the system must combine sensory information in a manner which maintains the detail and accuracy of the sensory data after combination. This can be accomplished by an appropriate backprojection reconstruction system. Second, the system needs to examine the object in a manner which insures that all significant parts of the object are sufficiently described. This can be accomplished by an appropriate view selection for a backprojection reconstruction process.

1.2 Reconstruction Through Backprojection

Reconstruction through backprojection is a computer vision analysis technique whereby a number of views of an object are combined to provide a three dimensional description of an object. Unfortunately, any pairing of views can result in conflicting information about the object. Also, it is not apparent how many views are necessary for reconstruction or where the viewpoint should be located in order to obtain the "best" resulting reconstruction. Even the definition of "best reconstruction" is open to question.

Reconstruction backprojection is well qualified for the job of industrial inspection because it requires less registration effort than other multisensor techniques. This effort in the industrial inspection using backprojection reconstruction is motivated by the efforts of Tan [123]. Tan discusses backprojection through reconstruction, an approach to conflict resolution, the importance of camera calibration to accuracy and suggests the need for view planning. He implies the need for objects not to have concave regions for reconstruction to be possible, but does not develop a characterization for suitable objects. Also, although view planning is addressed, no attempt is made at proposing a view planning technique.

All backprojection reconstruction algorithms can be classified as using one of two types of data. They may be broken down into those which use orthographic projection in the backprojection routine and those which use perspective projections in the backprojection. Orthogonal backprojection reconstruction has been used to reconstruct objects using silhouettes [2] and for medical examination of the heart [122]. Perspective backprojections have been used to describe a robots workplace [48] and to reconstruct objects [83].

Closely related to backprojection reconstruction is the internal representation of the reconstruction space. Unlike other three dimensional vision analysis techniques, Backprojection Reconstruction seeks to obtain a complete description of the object. Because of the amount of information stored in this complete description, most reconstruction algorithms make use of some form of data compression. Potmesil [97] examines oct-tree models of entire scenes of objects which are obtained from silhouette perspective backprojection reconstruction. Lavakusha and Chien [26, 70] examines oct-tree models of objects which are obtained from silhouette orthogonal backprojection reconstruction. Veenstra [130] presents an algorithm for the oct-tree generation of a three dimensional cube whose orientation is known. Cyganski [30] discusses the implementation aspects and performance of reconstruction using silhouettes. Raviv [100] presents a method of reconstruction from the shadows of an object where the light source is near the camera. Noborio [93] presents an algorithm for the construction of an octree from multiple views.

The backprojection reconstruction technique relies upon the knowledge that a specific point imaged onto a two dimensional surface view point must lie somewhere along a line defined by the backprojection of the two dimensional surface point back through the optical system. This is shown in Figure 1.1. This figure shows how a single point can be backprojected and tells us that at least one point of the object must lie along this line. When this backprojection is performed we only know that there is at least one point on the backprojection which is part of the object. We do not know if there is more than one or where they are located.

The backprojection of any point in the view which is not part of the object is known not to be part of the object. In those instances, we know an entire line passing through the reconstruction volume which is not part of the object. In actuality, the backprojection is accomplished by backprojecting the lines which miss the object. The object is then found to be those points in the reconstruction volume which are not determined to be non-object. The backprojection of a view is the union of the backprojections of the points of that view.

Backprojection reconstruction can be thought of as iteratively intersecting additional backprojections of that object within a reconstruction volume. This is equivalent to the complement of the union

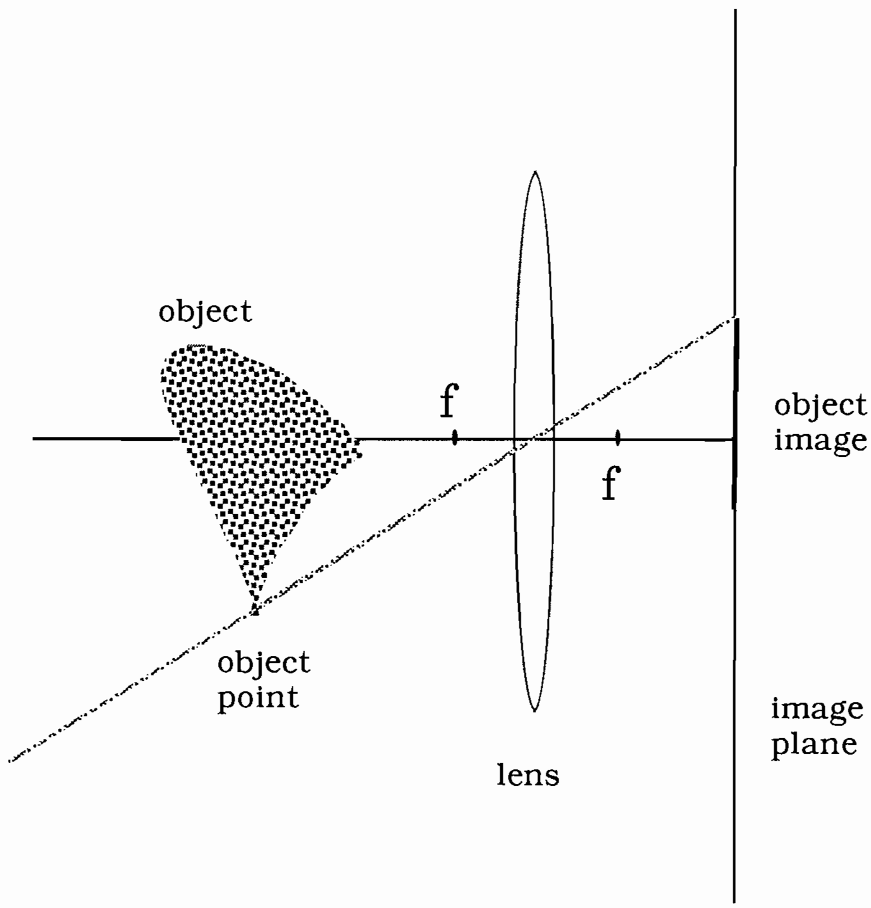

Figure 1.1 Backprojection of a Single Point from a Viewpoint through the Object.

of background backprojection lines. In both cases, this forms a solid which must contain the object which was originally imaged and is the subject of our reconstruction efforts. The reconstruction system could assume a perspective or orthogonal imaging system. Reconstruction takes place when more than one object view backprojection is combined with the other backprojected views.

The process of combining the view intersections has the effect of iteratively removing pieces of reconstruction space from possible object space. This process, shown in Figure 1.2, will result in an outer bound on the object undergoing reconstruction. Note that in the figure only one perspective of the reconstruction process is shown. Views are not restricted to planar views. They may be taken from any angle with respect to the object which the imaging system allows.

1.2.1 Areas Related to Backprojection Reconstruction

The best possible reconstruction is still limited to a boundary determined by this removal process. Some objects are not well described by this process. Figure 1.3 shows a planer view of the reconstruction of such a shape. Note that for the purposes of illustration the depression is viewable from the plane chosen. If the depression could be viewed by the vision system as in the illustration we could reconstruct the object through the use of a view perpendicular to the page. The symmetry depicted in the Figure, within the plane of the page, must be present from all view points for reconstruction to fail.

This suggests that a characterization for objects which would be suitable for backprojection reconstruction would be related to the background backprojection reconstruction lines which pare away the reconstruction volume from the object reconstruction. Such a characterization would define what class of objects are suitable for reconstruction by this form of analysis.

Two obvious examples of reconstructible and non-reconstructible objects are a cube and a drinking glass, respectively. All points on the boundary of the cube can be viewed from an angle which places the boundary point along the background. Points on the interior of the glass

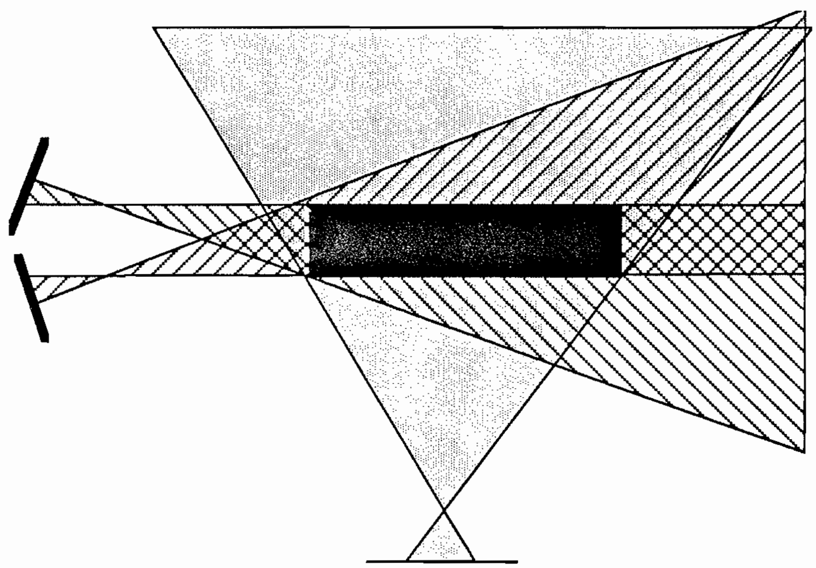

Figure 1.2 An example of the object reconstruction process using three views which succeeds

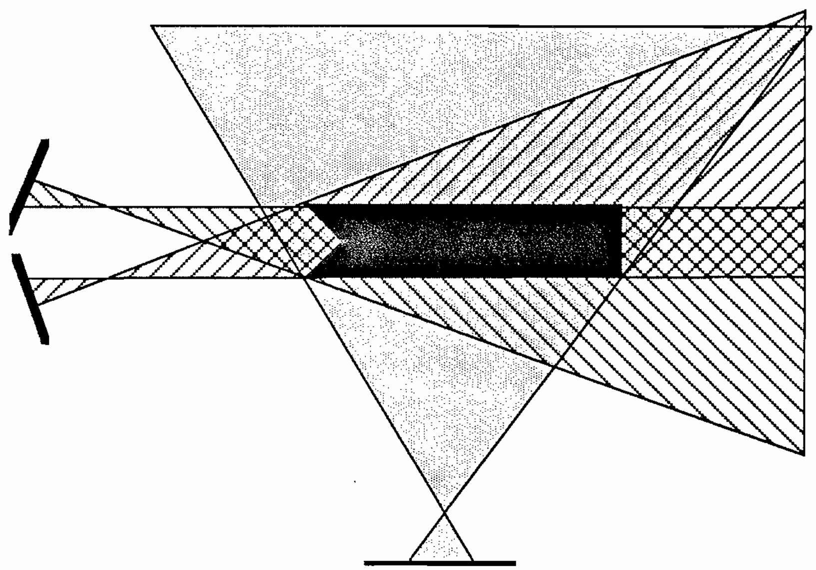

Figure 1.3 An example of the object reconstruction process using three views which fails

can not be "viewed" in this manner. Since there is no way to distinguish between object and background the glass is not reconstructible through backprojections.

Two different types of theory motivate the formalization of this concept. One was the mathematical concept of visibility sets [16, 127]. This is the mathematical study of the meaning of one set of points being visible from another set of points in an object. The other is the underlying properties of sets [50], convexity [64, 128, 129], geometry [87], and analysis [9]. Each of these areas provides a different key to defining the types of objects suitable for analysis.

In order for backprojection reconstruction to be performed, it is necessary that the object be viewed from the proper angles. Figure 1.4 shows an example of reconstruction failure due to improper view selection. In this case the views are chosen to be the worst possible views. The backprojection of each view contains a maximal cross section of the object. The same number of views, all chosen from a slightly different angle, would have a minimum cross section. In either case, the vertices of the object can be viewed as the essential points of the reconstruction. Proper selection of the view angles for these points would have enabled proper reconstruction.

One source for this information would be a CAD database. As part of the database the object would be completely specified. The crucial landmark points could be generated automatically or specified by the designer. In addition, the best angles from which to view the crucial landmark points could be specified either automatically or by the designer. Each of these options has advantages and disadvantages. Automatic generation of data requires the least amount of human interaction and can be viewed as being the most general. However, human selection of the crucial landmark points could lead to a simpler reconstruction (if fewer crucial points are needed) and a simpler set of view angles.

Links between CAD and machine vision have been proposed [47]. Here it is suggested that vision systems be information driven rather than require training for each new circumstance. Automatic sensor placement has been investigated for vision systems [28]. Sensing

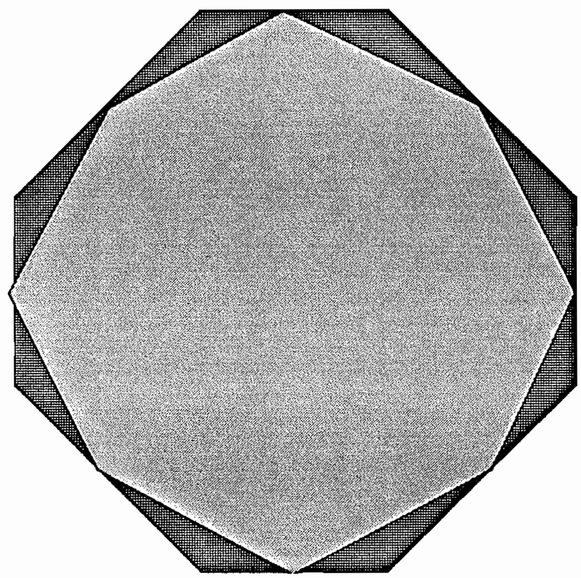

Figure 1.4 An example of the object reconstruction process which fails due to improper view selection

strategies in a robotic work cell have been investigated [54, 55]. An incremental approach where views are added as needed rather than determined all at once has been examined [27]. Aspect graphs can provide a unique method of storing the relational information on objects under examination [55, 144]. A method of computing the aspect graph for line drawings of polyhedral objects has been presented [37].

Finally, the accuracy measure of any reconstruction should take into account whether the desired information was extracted. This may or may not imply reconstruction of the entire object. In most applications complete object information would not be needed; only information about certain points or surfaces would be required. If the points of interest lie on a convex portion the object, a suitable set of views can be found to provide good results.

A number of different factors determine the accuracy of the reconstruction. The camera parameters have to be known. These include the focal length, pixel size, and optical center. Also, the relative locations of the object and camera must be known in order to prevent as much reconstruction conflict as possible.

Camera parameter estimation techniques have been examined by [42], using a rectangle projection; by [15], using a twisted cubic projection; by [84], using two planes; and by [123], using point projection. The registration of camera and object has been discussed by [123]. Error analysis in relation to stereo depth has been investigated by [88, 137]. Communication detection theory has been extensively developed by a number of authors including [96, 124, 126]. Puget [98] discusses an optimal solution for mobile camera calibration. Izaguirre [58] presents a method for the calibration of a pair of mobile cameras. Chen [23] presents an algorithm for the determination of three dimensional object position from a single calibration object. Linnainmaa [76] presents a technique for the pose determination of an object using triangle pairs. Abidi [1] presents a method for camera calibration using volumetric measurements of a tetrahedron. Ray [101] presents a maximum likelihood estimation of object pose assuming object identity and geometric models are known. Gilbert [38] presents a procedure for computing the distance between two convex objects in three space.

Werman [138] presents a distance metric for multidimensional histograms.

1.3 Scope of Investigation

There are four main areas of investigation which will be discussed. Each area will involve a different aspect of the reconstruction problem. These areas are: an object characterization for backprojection reconstruction, a limited view, goal driven, volume source method of backprojection reconstruction for limited object boundary inspection; calibration, system parameter effects and accuracy analysis in backprojection inspection, and view planning for backprojection reconstruction inspection analysis.

The first area noted is an important question which has not been addressed: the question of what type of objects are suitable for the reconstruction process. An answer in this area is essential if the usefulness of backprojection reconstruction techniques is to be determined. Chapter 2 will investigate this characterization. We will propose a characterization for objects which can be reconstructed using backprojection reconstruction, given theoretically obtainable views of the object. This characterization will be extended to include constraints for realizable backprojection reconstruction using obtainable viewpoints.

Object inspection using backprojection reconstruction techniques is the focus of this work. Chapter 3 will introduce a new method for backprojection inspection. This method will be a goal driven technique which requires reconstruction of only part of the boundary of the object. This allows for a limited set of views to be used in the reconstruction of these areas. The reconstruction technique used is a variation of volume source backprojection reconstruction. This technique projects the reconstruction volume under inspection onto the views to determine the state of the voxel. Most techniques suffer from approximating the reconstruction voxel by a single point during the projection phase. This method provides for the maximal use of the information provided in the view by not approximating the voxels during reconstruction.

Chapter 4 will investigate the accuracy analysis problem for three dimensional reconstructions. Three different areas are involved in determining the possible error of any reconstruction. These are the calibration technique, the effects of system parameters, and the techniques used to measure the accuracy in the inspection results. A clearly justifiable accuracy analysis technique is currently not available for three dimensional object examination. First, we present a technique for the calibration of an actual inspection system. Second, the parameters of this system are analyzed for their interaction and constraint on the possible accuracy of the system. Third, an accuracy analysis technique is proposed which is based upon these three areas and error measures based upon this technique are proposed.

Chapter 5 will propose a method of view selection for backprojection inspection. The view selection technique which is proposed assumes a priori CAD information about the object. This CAD information details what parts of the object are critical to the object, and therefore essential to an effective reconstruction. We shall define these critical locations to be landmarks on the object. View selection for the object can also be a considered a constant because the ideal object description does not change as the test object changes. Views are selected which provide information about these specific landmarks on the object.

Finally, Chapter 6 provides a summary of these investigations and suggested areas of future work.

1.4 Problem Summary and Statement of Contributions

Within the realm of three dimensional computer vision, industrial inspection provides some unique challenges for investigation. Most imaging techniques provide two dimensional views of an object, but three dimensional information is necessary for the inspection of three dimensional objects. Using two dimensional views and CAD information about the object, reconstructed object information can be extracted using backprojection reconstruction techniques. Accuracy analysis techniques for this reconstruction are required. Interactions between calibration, system parameters, and measurement techniques and their effects on accuracy need to be addressed. Also, it is desired that the reconstructions be performed using views which are selected in a manner which reduces the computational complexity. Finally, a characterization of objects which are suitable for inspection analysis by this method is necessary so that a statement may be made concerning the applicability of the algorithm.

Within this work, there are several specific contributions of note. These are as follows:

1) A characterization of objects for theoretical backprojection reconstruction is developed. It is extended to their practical reconstruction.
2) A new backprojection inspection technique is proposed which is goal driven, requires only limited views of the object, and uses a volume source technique to improve accuracy.
3) Calibration and system parameters are examined for their effect on inspection accuracy. A measure is proposed for inspection accuracy measurement.
4) A new view selection technique is proposed which eliminates redundancy in an original selection of views. For inspection, the original views are CAD designated.

CHAPTER 2

OBJECT CHARACTERIZATION FOR BACKPROJECTION RECONSTRUCTION

2.1 Introduction

In this chapter a characterization of objects which are suitable for backprojection reconstruction is presented. In Chapter One we presented the need for such a characterization with respect to the overall reconstruction problem. Here, this need is motivated in Section 2.2 through object examples which both can and cannot be reconstructed. From this the concept of weakly convex objects is introduced in Section 2.3. Our goal is to mathematically formalize the minimal character of the objects which can be reconstructed by backprojection reconstruction. Our definitions and theorems are presented in this section within the framework of linear, topological, set space theory in conjunction with minimal assumptions about the reconstruction environment. It shall be shown in Section 2.4 that an object being weakly convex is a necessary and sufficient condition for an object to be reconstructed through a theoretical backprojection reconstruction. Section 2.5 will extend this concept to realizable reconstruction by adding a more practical set of reconstruction environment constraints. Section 2.6 shall show the constraints of Section 2.5 provide a necessary and sufficient condition for realizable backprojection reconstruction under the assumptions presented. Corollaries of these theorems are presented in Section 2.7 for different types of backprojection reconstruction and a related three dimensional computer vision problem. Section 2.8 presents several examples of interesting objects. Section 2.9 is a discussion of the implications of this characterization for backprojection reconstruction.

The mathematical basis of this development will be Topological Linear Space theory [128]. The approach of this work was motivated by efforts at determining internal visibility characteristics of a polygon [5, 34, 72, 73, 112] from a point or line segment. This was examined in n-dimensional space for objects visible from subsets by Valentine [127]. These efforts involve the visibility of interior points of regions from other points within the region. Our efforts are directed toward determining the external visibility of a solid. External visibility of object vertices has been examined by Buchman [16]. General properties of n-dimensional set spaces are noted in [50, 64, 128, 129], among others.

2.2 Motivation for Characterization

In backprojection reconstruction algorithms, there is the tacit assumption that the two dimensional views have the ability to image the boundaries of the three dimensional object. Sometimes it is stated that these algorithms can be applied to convex objects (Figure 2.1), yet it is theoretically possible for these algorithms to reconstruct some non-convex objects (Figure 2.2), but not other non-convex objects (Figure 2.3). Obviously, the characterization of viable objects for reconstruction is not immediately apparent. It would be useful to have available a categorization against which to check objects which are candidates for reconstruction.

Towards this end we shall examine the three objects of Figures 2.1-3: a cube, reconstructible; a U channel, reconstructible; and a drinking glass, non-reconstructible. One effort to describe the difference between these objects would revolve around the cube being "solid" and the drinking glass and U channel having an "inside." Efforts to describe the difference between these two objects have led to the use of the term convex to describe suitable objects for reconstruction due to observations similar to a cube being convex and a drinking glass not being convex. This is not sufficient to include the U channel, which is reconstructible. A better description is required.

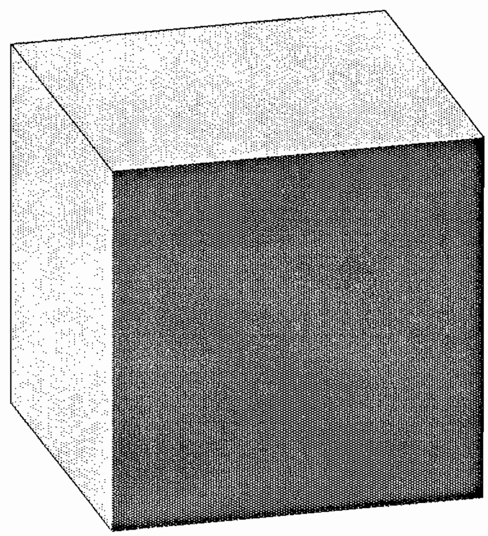

Figure 2.1 Cube: Example of a convex, reconstructible object

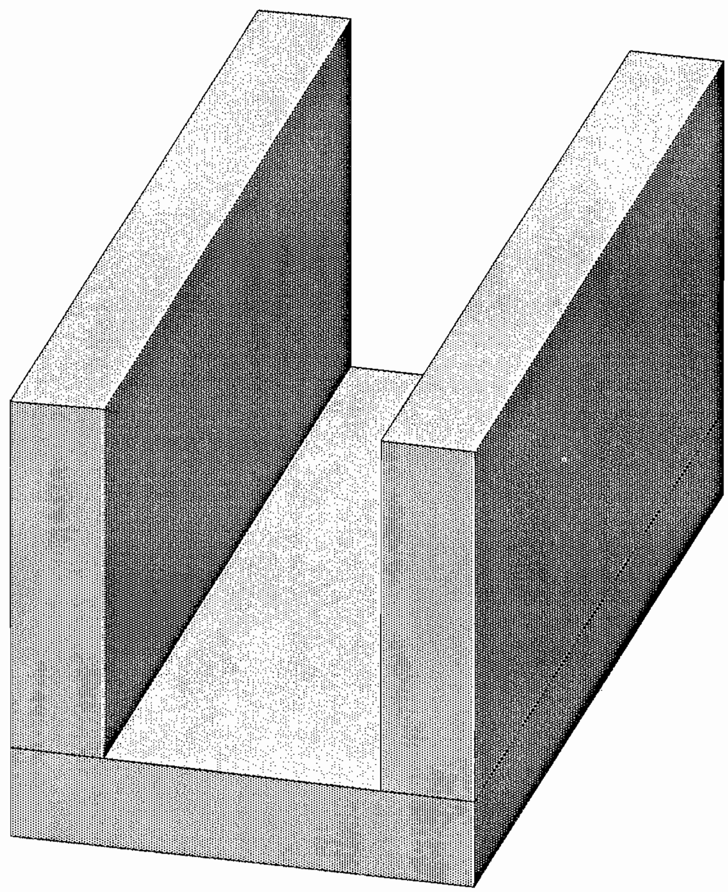

Figure 2.2 Open U Channel: Example of a non-convex, reconstructible object

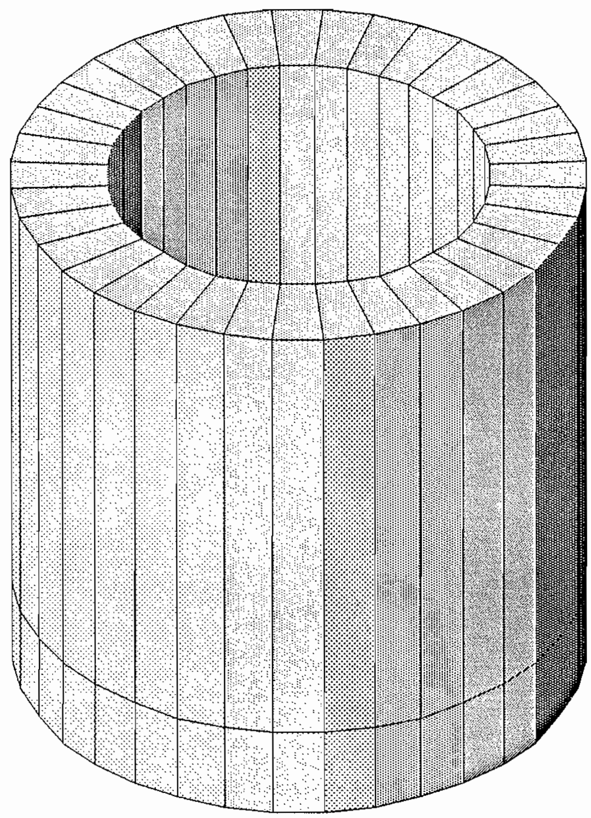

Figure 2.3 Handleless Cup: example of a non-convex, non-reconstructible object

The reconstruction problem is essentially one of determining the boundary of the object undergoing reconstruction by observing boundary point projection. This must be done for each point on the boundary of the three dimensional object for complete object reconstruction. An appropriate object characterization must define what is meant for a point to be visible for the purposes of reconstruction. This is the focus of our problem: the formalization of what it means for the backprojection reconstruction process to observe an object boundary.

This issue has been briefly addressed as different approaches to backprojection reconstruction were examined. Several authors point out that backprojection reconstruction algorithms can not reconstruct the "concavities" of objects [2, 26, 83, 93, 97, 123]. Another claims reconstruction of "convex objects without holes" [70]. However, there is no effort to describe the nature of the "concavities" or "holes" which are not reconstructible. Another description of the reconstruction result is the "convex hull" of the object [30, 100]. Figure 2.2, the U channel is an example of an object which is reconstructible from backprojection reconstruction which is not its own convex hull. Veenstra [130], presents a definition of a *representable* object. His definition is based upon silhouette intersection congruency with the object. This essentially says that those objects which we can reconstruct are those that we can construct by silhouette intersection. Our efforts differ from his in that we shall base our object characterization on the characteristics of the object and the basics of the reconstruction process.

The backprojection reconstruction process combines silhouette backprojections of object views in a manner which provides information about the object, or more importantly, the object boundaries. When we are backprojecting object view points, we do not know exactly where along the backprojection an object point was located when imaged. We only know that at least one point along the backprojection is occupied by the object. More information is actually obtained from the backprojection of a background view point. In this case we know that there are no object points anywhere along the backprojection; thus eliminating an entire line from possible object space.

For complete reconstruction to occur, it is required that we locate each object boundary point while the only points which we can positively identify are those which are not part of the object. This can only be accomplished if we are guaranteed that a background backprojection line passes through all non-object points which are adjacent to an object boundary points.

Any characterization for reconstructible objects must group objects in a manner such that we can distinguish reconstructability. The basis for such a characterization should be a mathematically description of a general object with a minimal description of backprojection reconstruction constraints. The term we shall use to refer to this characterization is *Weak Convexity* and reconstructible objects shall be described as *Weakly Convex*.

2.3 *Weak Convexity*

Thus far we have described the objects of interest by referring to the object as a whole, the boundary of the object, or each "point" on the boundary. In order to formalize the concept of *Weak Convexity* it is necessary to formalize the description of the object to be reconstructed, the reconstruction space, and the observation viewpoints.

This formalization shall be done in a linear, topological, metric space. We will define this combined space and notation for use here and in the rest of the Chapter. There are no inherent restrictions on our development to a specific dimensionality. Significant properties essential to the definition of *Weak Convexity* will then be presented. This is followed by a list of general assumptions which describe the process of backprojection reconstruction. Finally, *Weak Convexity* will be defined for this general case.

2.3.1 Introduction and Notation

The space theory we will use is based on three sets of axioms: one for set theory, one for algebraic spaces, and one for topology. This section draws heavily from Valentine [128], Millman and Parker [87], and

Bartle and Sherbert [9] for a correct set of axioms upon which to base set theory, topology, and metric norms in an integrated fashion. It should be noted that no claim is made at a complete discussion of $\mathcal{L}$ spaces and their properties.

Let us define $\mathcal{L}^n$ as a n-dimensional linear topological space over a real field $\mathcal{R}$ which has two operations, vector addition and scalar multiplication,a collection of subsets called open sets, and which have the usual axiomatic properties of both topology and algebraic spaces [128]. Unification of these ideas gives us the ability to view a given problem from either a topological view or a algebraic view and have both be equally valid. Further, it is valid to freely interchange and mix these ideas rather than view them as alternatives.

Table 2.1 shows the notation which will be used to describe our linear topological spaces. Each symbol shown represents two different qualities in our combined space: one which is normally expressed in set notation and one which is normally expressed in algebraic space notation. We will use the terms interchangeable depending upon the context of the analysis at hand.

2.3.2 Supporting Properties for *Weakly Convex* Sets

Here we state the properties which are the basis for our spaces. Most will be familiar with many of the statements. Key definitions are repeated here for clarity as we make use of the unity between the concepts of vector, set, topological and metric spaces which are essential to our development. This unity will then be specifically applied to the Euclidean space where the reconstruction occurs.

Closed Sets

The object which we shall reconstruct shall be defined to be a closed set under $\mathcal{L}^n$.

Definition A closed set is the complement of an open set $\subset \mathcal{L}^n$ [128].

Table 2.1 Notation list definition for set and metric space viewpoints

Term	Set Space	Metric Space
$\mathcal{R}$	real field	real field
$\mathcal{L}^n$	N-Dimensional set space	N-Dimensional Real arithmetic space
$\mathcal{L}^3$	3-Dimensional set space	3-Dimensional Real arithmetic space
ε	arbitrarily small scalar	arbitrarily small scalar
a	a fixed set element	a fixed point
w,x,y	arbitrary set elements	arbitrary points
∞	background element	background point
O,S	arbitrary closed object set	arbitrary closed object volume
M	e-neighborhood set	e-neighborhood volume
N	e-neighborhood subset	e-neighborhood subvolume
z	arbitrary element of N	arbitrary point in N
c,v,v'	arbitrary observation element	arbitrary observation point
V	Possible reconstruction set	Possible reconstruction volume
U	upper half plane set	upper half plane volume
B	camera boundary set	camera boundary volume

Boundary

One of the properties we will need to use is the definition of the boundary of the object to be reconstructed. Reconstruction of the boundary of the object (or set) will be central to backprojection reconstruction.

Definition Interior points of a set S are those points $x \in O \subset S$ where O is an open subset of S [128].

Definition The boundary points of a closed set S are those points which are not interior points of S [128].

Notation bd $S \equiv$ boundary of S.

Line segment

A common algebraic concept is that of a line segment. This shall now be extended to topological spaces.

Definition A line segment is a closed set between $x, y \in \mathcal{L}^n$ such

$$\text{consisting of the points} \equiv \bigcup_{0 \leq \lambda \leq 1} \left(\lambda x + (1 - \lambda) y \right) .$$

Notation $\overline{xy} \equiv$ line segment from x to y .

Metric Geometry

One of the key properties we will rely on is the property of our space having a metric geometry. Having a metric geometry allows the distance between points to be expressed. This allows for error analysis to be performed between the reconstructed object and the desired object.

Theorem The linear topological Space on $\mathcal{L}^n$ with Euclidean distance d_E is a metric geometry [87].

Definition The length of a vector $x \in \mathcal{L}^n$ is defined to be $d_E = ||x|| = \sqrt{x \cdot x}$ [128].

Betweenness

Another key property of our space is that of representing one point as being between two other points. The impact of the betweenness definition is that it gives us the ability to order points on a line in $\mathcal{L}^n$. First we shall define the concept of betweenness and then we shall define the notation we shall use to express betweenness.

Definition w is between x and y *if* w,x, and y are distinct collinear points in the metric geometry on $\mathcal{L}^3$ and $||w - x|| + ||x - y|| = ||w - y||$ [87].

Notation $\overline{w\,x\,y}$ is used to define three points where x is between w and y.

Neighborhoods

A third property is that of the neighborhood of a point. The neighborhood definition allows us to select a set of points which are arbitrarily close to a specific point. First we shall define the concept of a neighborhood and then we shall define the notation we shall use to represent this concept.

Definition Let $a \in \mathcal{L}^n$.

i) For $\varepsilon > 0$, the ε-neighborhood of a is the set $\{ x \in \mathcal{R}: || x - a || < \varepsilon$

ii) A neighborhood of a is any set that contains an ε-neighborhood of a for some $\varepsilon > 0$. [9]

Notation nh $x \equiv$ neighborhood of x

2.3.3 Reconstruction Assumptions

The assumptions used are those common to all backprojection reconstruction problems. The assumptions presented are a minimal set related to the geometry of the general backprojection reconstruction problem. These assumptions are that the object is contained in the reconstruction volume, and that the reconstruction volume is contained within the universal space. Projections are performed from a space of dimension n onto one of dimension n-1. Further, the observation point is allowed to be anywhere outside the object undergoing reconstruction. Additional assumptions are presented later to address a realizable reconstruction.

<u>Assumption 2.1</u> The reconstruction space must be a subset of the universe S, $V \subset S$.

<u>Assumption 2.2</u> The object is contained in a subset of the reconstruction space such that $O \subset V$.

<u>Assumption 2.3</u> The observation view angle can be any value such that $c \notin O$.

<u>Assumption 2.4</u> Observation range can be anything such that $c \notin O$.

<u>Assumption 2.5</u> A mapping function exists from dimension n (object) to dimension n-1 (image) which maps a line through the object to type object in the projection and a line missing the object to type background in the projection.

<u>Assumption 2.6</u> The background is distinguishable from the object when two points which are in a common ε-neighborhood map to object and background differently.

2.3.4 *Weak Convexity* Defined

The space in which the reconstruction occurs is characterized as a set which is partially filled with an object set. The space is linear, continuous, and has a metric norm defined upon it. This provides both relevance to our normal analysis techniques and an easy visualization of the reconstruction in three dimensions.

Definition 1 (*Weak Convexity*) A set $S \subset \mathcal{L}^n$, a linear topological space, is weakly convex if $\forall\ x \in \text{bd}\ S$, let M be a ε - neighborhood of x, $\forall\ N \subset M : N \cap S = \varnothing$ $\exists\ \upsilon \notin S \cup N,\ z \in N :\ \overline{\upsilon\, z\, \infty} \cap S = \varnothing\ ,\forall\ z \in N.$

Definition 1 states that a set S, being a subset of linear space $\mathcal{L}$ is weakly convex if for all points x being members of the boundary of S we let M be a ε - neighborhood of x such that $N \subset M$ is disjoint with S, and there exist point υ which is not a member of $S \cup N$, and there exists a point z which is a member of the neighborhood N of x such that the half line defined by the points υ, z, ∞ is disjoint with S for all z that is in N .

Observing the boundary as defined by the left hand side involves being able to see through to the background at all points near the boundary of the actual object (Figure 2.4). We assume that we can distinguish between the object and the background. In reality, this is a difficult problem involving segmentation, lighting, and object reflectance, etc. The assumptions insure that if we can observe all object boundaries and adjacent background we can distinguish between them. Mathematically, this implies that we can distinguish set membership without ambiguity and are justified in doing so in our analysis.

Observation lines will be required for reconstruction of a single object boundary point. This line will pass from an observation point, near a boundary point and on to the background without intersecting the object. Clearly, the observation points must not be nearby the boundary point. Another point must be nearby the boundary point but not part of the object. The establishment of what it means for one point to be nearby another point is critical. Mathematically, the concept of nearness

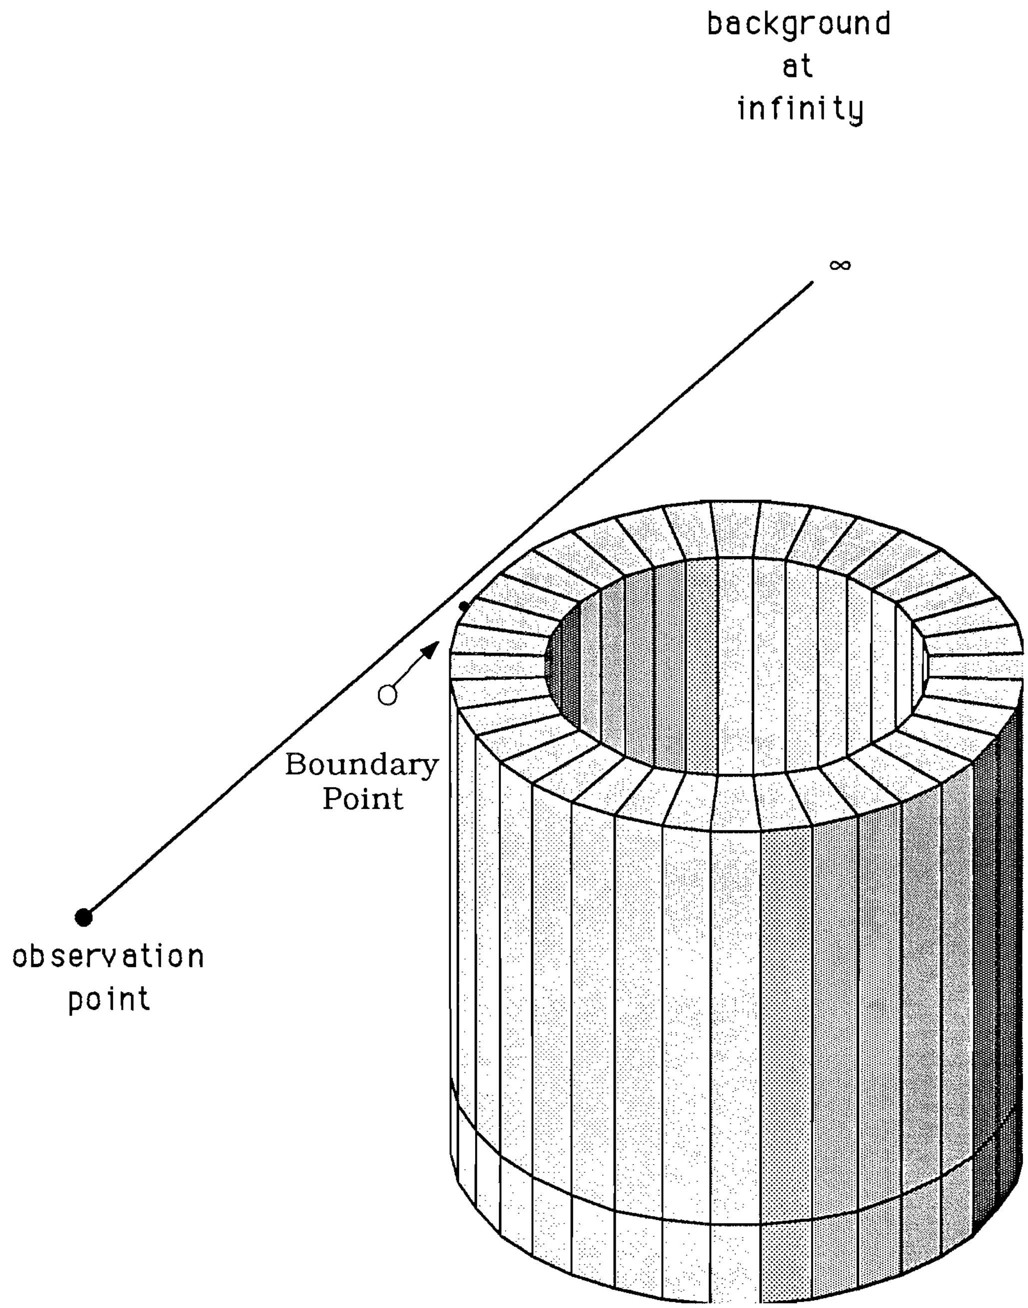

Figure 2.4 Observation Line to the background near an Object

can be established through the use of an ε - neighborhood. Refer to Figure 2.5 for an example of neighborhoods specifying nearness.

2.4 Theorem 1

Definition 1 provided a motivated description of objects which can be reconstructed through backprojection reconstruction. We will now show that Definition 1 is a necessary and sufficient condition for backprojection reconstruction to be accomplished.

Theorem 1 An object defined such that $S \subset \mathcal{L}^n$ can be reconstructed through backprojection reconstruction iff it is *Weakly Convex.*

Proof

CASE I *Weakly Convex* $\Rightarrow$ backprojection reconstruction

Let O be the weakly convex object. Assume O can not be reconstructed. Let $x \in \mathrm{bd}\ O$ be a point which can not be reconstructed, rather x can not be distinguished from the background. contradiction. x must be able to be distinguished by definition of weak convexity and assumption 2.5 and 2.6. $\therefore O$ can be reconstructed.

CASE II backprojection reconstruction $\Rightarrow$ *Weakly Convex*

Let O be the reconstructed object . Let $x \in \mathrm{bd}\ O$. Let $z \in \mathrm{nh}\ x$, $z \notin O$ and υ the point which observes x. By assumption 2.5 and 2.6 and O being reconstructed, $\overline{\upsilon\, z\, \infty} \cap S = \varnothing$.

$\therefore O$ is weakly convex.

■

2.5 Realizable Reconstruction in $\mathcal{L}^n$ Space

Thus far our effort has been aimed at determining the minimum requirements for a theoretical backprojection reconstruction to occur. This was done under the general constraint of backprojection

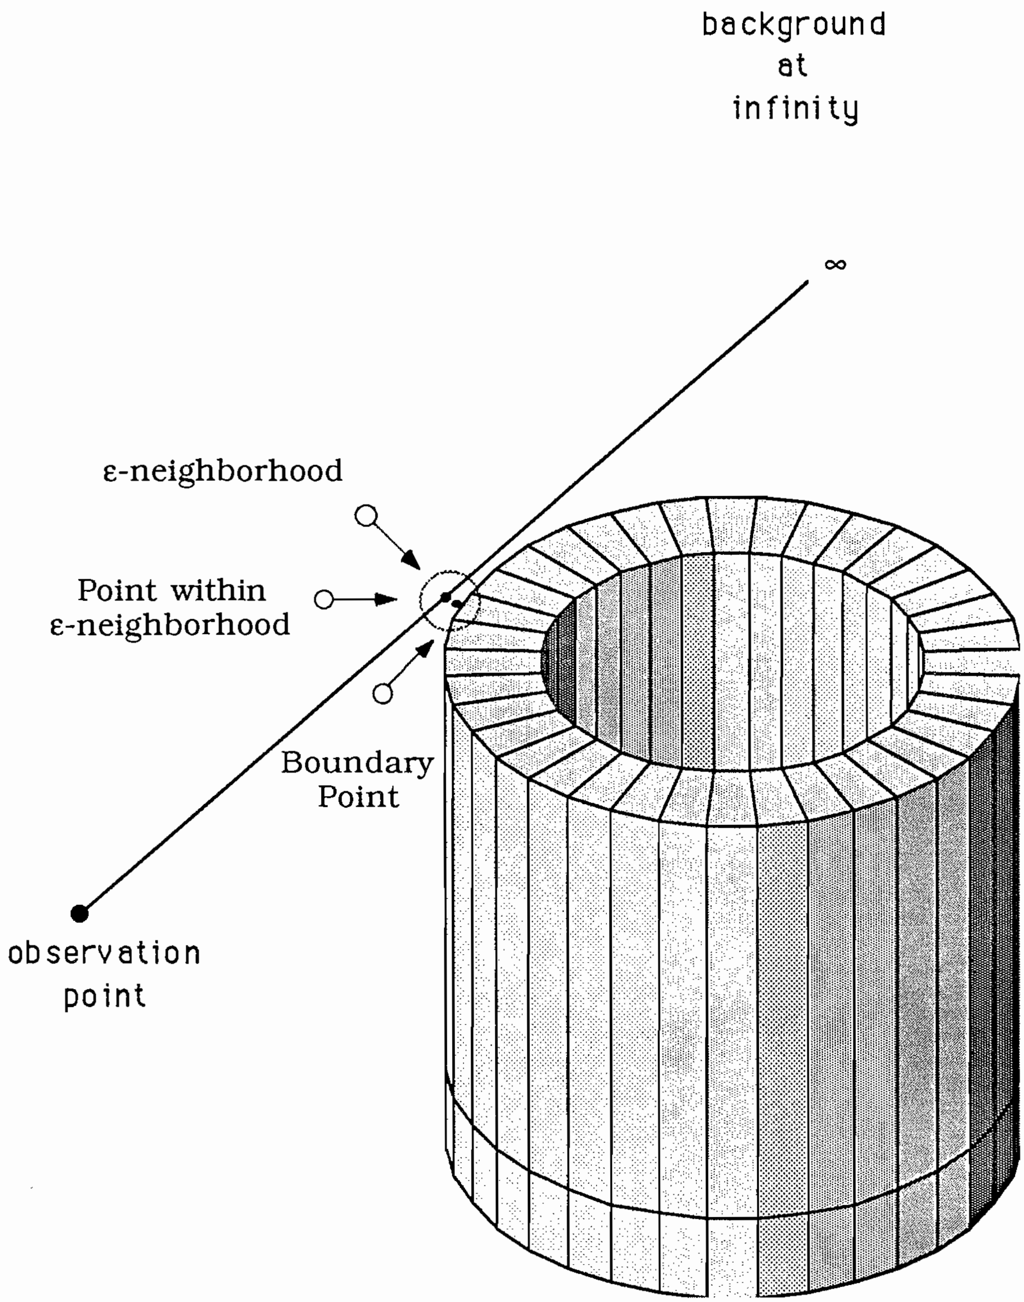

Figure 2.5 Observation line to the background through an e - neighborhood of a object boundary point

reconstruction processing, but without regard to realizability. In this section we will examine an additional set of restrictions imposed to achieve realizable.

In general, realizability will include effects due to lighting, segmentation, and aspects of the vision system. These properties will not be addressed. These properties, while affecting realizability, are limitations of specific vision systems or objects. We seek an invariant object description limited only by the geometry of backprojection reconstruction. Therefore, the assumptions used here are restricted to those which would apply to all reconstruction systems.

2.5.1 Reconstruction Assumptions Revisited

Below are listed some constraints which are due to the mechanics of typical analysis procedures and not the geometry of reconstruction. These constraints are different from those noted earlier in that even though they do restrict the analysis they are not a basic part of reconstruction analysis and may not be restrictions in the future.

<u>Assumption 2.7</u> The reconstruction occurs in three dimensional space $\mathcal{L}^3$.

<u>Assumption 2.8</u> The camera must stay outside the reconstruction space, $c \notin V$.

<u>Assumption 2.9</u> The camera must stay within some bounding volume B smaller than $\mathcal{L}^3$ but larger than V, $c \in B - V$.

<u>Assumption 2.10</u> Let $U \equiv$ upper half space whose lower boundary is the lower boundary of the object. The camera c must stay in a positive half sphere above the lower surface of the object.

2.5.2 Weak Convexity Extension for Realizable Reconstruction

It is now possible to state Definition 2, a more restricted version of Definition 1. Definition 2 has the same form as Definition 1 except that some of the more general terms of Definition 1 have been replaced by more restrictive ones which follow from the assumptions of Section 2.5.1.

Definition 2 An object $S \subset \mathcal{L}^3$ is reconstructible iff $\forall\ x \in$ bd S, M an ε - neighborhood of x, $N \subset M : N \cap S = \emptyset$ $\exists\ \upsilon \in B - (V \cup N),\ z \in N :$ $\overline{\upsilon\ z\ \infty} \cap S = \emptyset\ \forall\ z \in N\ \forall\ z \in N.$

Definition 2 differs from Definition 1 only in its restrictive nature.

2.6 Theorem 2

Definition 2 proved a description of objects which can be reconstructed through a more practical backprojection reconstruction. We will now show that Definition 2 is a necessary and sufficient condition for backprojection reconstruction to be accomplished.

Theorem 2 An object defined such that $S \subset \mathcal{L}$ can be reconstructed through backprojection reconstruction iff it is reconstructible as defined by Definition 2.

Proof

CASE I: reconstructed through backprojection reconstruction $\Rightarrow$
$\forall\ x \in$ bd S, M an ε - neighborhood of x, $N \subset M :$
$N \cap S = \emptyset\ \exists\ \upsilon \in B - (V \cup N),\ z \in N :\ \overline{\upsilon\ z\ \infty} \cap S = \emptyset$.

A realizable reconstruction implies a reconstruction occurs. By theorem 1, the object is weakly convex and therefore M is a ε - neighborhood of x, $N \subset M : N \cap S = \emptyset$
$\exists\ \upsilon \notin S \cup N,\ z \in N :\ \overline{\upsilon\ z\ \infty} \cap S = \emptyset$
Since $\overline{\upsilon\ z\ \infty} \cap (B - (V \cup N)) \neq \emptyset$ either $\upsilon \in B - (V \cup N)$ or υ' can be chosen such that $\upsilon' \in B - (V \cup N)$ and $\overline{\upsilon'\ z\ \infty} \cap S = \emptyset$. Therefore $\forall\ x \in$ bd S, M an

ε - neighborhood of x, $N \subset M : N \cap S = \varnothing$

$\exists\ v \in B - (V \cup N),\ z \in N :\ \overline{v\, z\infty} \cap S = \varnothing$

CASE II: $\forall\ x \in \text{bd}\ S$, M an ε - neighborhood of x, $N \subset M$:

$N \cap S = \varnothing\ \exists\ v \in B - (V \cup N),\ z \in N :\ \overline{v\, z\infty} \cap S = \varnothing$

$\Rightarrow$ reconstructed through backprojection reconstruction.

$B - V \not\subset S$, $v \in B - V$ implies $v \notin S$. Therefore S is Weakly Convex and has a backprojection reconstruction.

■

2.7 Corollaries of Theorems 1 and 2

Theorems 1 and 2 show that Definitions 1 and 2 are necessary and sufficient conditions for backprojection reconstruction under different sets of reconstruction assumptions. Earlier, as part of the motivation for the concept of *Weak Convexity*, several different objects and concepts were introduced. These will now be re-examined for their relation to *Weak Convexity*. Further, several other topics of related interest will be examined. First, however, the relationship between the two main definitions will be presented.

2.7.1 Relationship between *Weakly Convex* and Definition 2

Both Definition 1 and Definition 2 describe objects for the purpose of backprojection reconstruction. There we shall show that any object satisfying Definition 2 is *Weakly Convex.*

<u>Corollary</u> Any object satisfying Definition 2 is *Weakly Convex.*

<u>Proof</u> Since $B - (V \cup N) \subset \neg(S \cup N)$ any object satisfying Definition 2 is *Weakly Convex*

■

2.7.2 Convex Objects

Convex objects were the previous common description of objects which could be reconstructed. It was noted that all convex objects should be reconstructible under any new characterization scheme. First, we will define convexity.

Definition A set $S \subset \mathcal{L}^n$ is convex if for each pair of points $x \in S$, $y \in S$ it is true that $\overline{xy} \subset S$. [128].

Corollary All convex objects are *Weakly Convex.*

Proof Let $\overline{xy} \in$ bd S, z nh x such that $z \cap S = \varnothing$. Let $\overline{vz}$ be parallel to $\overline{xy}$, $\overline{vz_\infty} \cap S = \varnothing$. $\therefore$ convex objects are *Weakly Convex.*

■

2.7.3 Convex Hulls

As noted earlier, convex hulls are another earlier attempt to describe objects which can be reconstructed. A convex hull can be thought of as the smallest convex region surrounding an object which may or may not be convex. First, we shall define convex hulls and then test their suitability as a necessary and sufficient condition for reconstructability.

Definition The convex hull of a set $S \subset \mathcal{L}$ is the intersection of all convex sets in $\mathcal{L}$ containing S, and it is denoted by conv S [128].

Corollary All convex hulls are *Weakly Convex.*

Proof Since each boundary point of conv S is the boundary point of some convex set and all convex sets are weakly convex, a convex hull is *Weakly Convex.*

■

Proposition All *Weakly Convex* objects are convex hulls.

Proof Let O be the cube of Figure 2.1. O is Weakly Convex because O is convex. O is also a convex hull. Remove one line from a face of O forming a channel in O. O is still Weakly Convex because a backprojection line may pass where the line used to be, thus identifying adjacent boundary locations. O is no longer a convex hull because it no longer contains the face line.
∴ The Proposition is false.

■

2.7.4 Orthographic Backprojection

In orthographic backprojection reconstruction there is a constraint upon the relationship between the observation lines of a particular view. This constraint is expressed in the definition below. *Weak Convexity* does not express any claims about the grouping of viewpoints into view images. It is then observed and shown that *Weak Convexity* is a necessary and sufficient condition for an orthographic backprojection reconstruction computer vision system to reconstruct the object.

Definition The observation lines within a specific view of an orthographic backprojection reconstruction system must be parallel to each other and perpendicular to the view.

Corollary Any object can be reconstructed with an orthographic backprojection reconstruction system iff it is *Weakly Convex* .

Proof

CASE I *Weakly Convex* ⇒ an orthographic reconstruction

Let each view in the orthographic reconstruction contain one view point from the *Weakly Convex* reconstruction. Therefore, all view lines in each view of the orthographic reconstruction contains parallel view lines and an orthographic reconstruction exists.

CASE II an orthographic reconstruction $\Rightarrow$ *Weakly Convex*
Since a backprojection reconstruction exists the object is *Weakly Convex.*
■

This restriction on the information which may be obtained from any one view has the possible effect of increasing the number of views which may be needed to reconstruct the object. Each view is only able to contribute reconstruction information along lines perpendicular to the view. In general, *Weak Convexity* allows any combination of view lines.

2.7.5 Perspective Backprojection

In perspective backprojection reconstruction there is a constraint upon the relationship between the observation lines of a particular view. This constraint is expressed in the definition below. *Weak Convexity* does not express any claims about the grouping of viewpoints into view images. It is then observed and shown that Weak Convexity is a necessary and sufficient condition for an orthographic backprojection reconstruction computer vision system to reconstruct the object.

<u>Definition</u> The observation lines within a specific view of a perspective backprojection reconstruction system must have a perspective geometry.

<u>Corollary</u> Any object can be reconstructed with an perspective backprojection reconstruction system iff it is *Weakly Convex* .

Proof

CASE I *Weakly Convex* $\Rightarrow$ an perspective reconstruction

Let each view in the perspective reconstruction contain one view point from the *Weakly Convex* reconstruction. Therefore, all view lines in each view of the perspective reconstruction contains parallel view lines and an perspective reconstruction exists.

CASE II an perspective reconstruction $\Rightarrow$ *Weakly Convex*

Since a backprojection reconstruction exists the object is *Weakly Convex*.

■

This restriction on the information which may be obtained from any one view has the possible effect of increasing the number of views which may be needed to reconstruct the object. Each view is only able to contribute reconstruction information along lines of a perspective transformation to the view. In general, *Weak Convexity* allows any combination of view lines.

2.7.6 Object Recognition

In object recognition using silhouette shape, the geometric relationship between the vision system and the object is not known. Assuming that the object boundaries must be observable from any angle by the recognition system for proper identification, an additional restriction on an object in addition to *Weak Convexity* is the symmetry of the object.

Corollary An object can be recognized using silhouette shape if it is *Weakly Convex* with a single observation point and suitable object rotation and translation.

Proof *Weak Convexity* implies that all boundary points of the object are observable and the use of a single observation line implies that a boundary silhouette may be observed under any rotation of the object. Therefore, the object can be recognized.

■

This has several important implications on the recognition system. First, it implies that the views established for the object must be rotationally invariant. *Weak Convexity* implies that a valid reconstruction exists but it does not imply that the reconstruction is symmetric.

A realizable solution would require that a rotation exist which would allow all non-object points in the neighborhood of each boundary point to lie along the observation line. This would need to be true for all objects which might need to be recognized by a given computer vision system in addition to their having distinguishable silhouettes.

2.7.7 Local Convexity on *Weakly Convex* objects

Thus far we have discussed properties of *Weakly Convex* objects and tested classes of objects for the property of weak convexity. Local convexity is a property of points on objects. We shall define local convexity and examine *Weak Convexity* for this property.

Definition A point $x \in S$ is locally convex at x if $\exists$ N a neighborhood of x such that if $w \in S \cap N$, $y \in S \cap N$ then $\overline{wy} \subset S$. [64].

Corollary *Weakly Convex* sets are locally convex.

Proof Let x be the point of the *Weakly Convex* set O which is not locally convex. Then there exists a point in $O \cap N$ which is not part of O. $\therefore$ contradiction.

2.8 Example Objects Evaluated for *Weak Convexity*

In this section we shall first revisit our examples of Section 2.2 to check the conformity of our definition with these objects. This will be followed by several other example objects.

2.8.1 Cube

The featured characteristic of the cube is that it has no depressions in its exterior surface. It is shown in Figure 2.1. It has the property of being a convex hull.

<u>Proposition 2.1</u> The cube shown in Figure 2.1 satisfies Theorem 2.

<u>Proof</u> Let x be any point on the exterior of the cube. All of them reside on the convex hull of the cube (itself). $\therefore$ Proposition 2.1 is true by the Corollary of section 2.6.

■

2.8.2 Handleless Cup

The featured characteristic of the handleless cup is a depression in one of its faces. It is shown in Figure 2.3. It is not convex.

<u>Proposition 2.2</u> The handleless cup shown in Figure 2.3 satisfies Definition 2.

<u>Proof</u> Let x be a point on the interior surface of the cup. Regardless of the y chosen the extension of the line segment intersects the object. $\therefore$ Proposition 2.2 is false.

■

2.8.3 Open U Channel

The featured characteristic of the U channel is a depression in three of its faces. These depressions are connected in such a manner that the cross section looks like a U. It is not convex.

Proposition 2.3 The open U channel shown in Figure 2.2 satisfies Definition 2.

Proof For any point x on the exterior faces of Figure 2.4 have the same proposition 2.1. For any point x on the interior of the U shaped surface we can chose a point on an exterior edge the as $y \therefore \overline{w\,x\,y\,\infty} \cap S \neq \emptyset$ Let M be a ε - neighborhood of $x, N \subset M : z \in N$ Let $\overline{v\,z\,\infty}$ be parallel to $\overline{w\,x\,\infty}$. $\overline{v\,z\,\infty} \cap S = \emptyset$ $\therefore$ Figure 2.4 satisfies Theorem 2 and Proposition 2.3 is true.

■

2.9 *Weakly Convex* Implications

Objects or sets which are *Weakly Convex* can be reconstructed through backprojection reconstruction. While this was shown in the general n-dimensional case, several applications were shown in the 3-dimensional case. While it is true that an explicit algorithm has not been presented for object evaluation, the definition of *Weak Convexity* provides the characteristics which must be present in the object or more explicitly, at each point in the object's boundary.

Previous descriptions of reconstructible objects have been shown to be insufficiently descriptive. They did not provide a sufficient description against which to evaluate potential objects for reconstructability. Test objects used by various other authors must be *Weakly Convex* for their reconstructions to succeed.

Weak Convexity was developed within the framework of linear topological space theory. By way of corollaries of Weak Convexity, we have established its place among other types of convexity. Specifically, any object is convex is *Weakly Convex*, but objects which are *Weakly Convex* are not necessarily convex. Further, all points on a Weakly Convex object are locally convex.

CHAPTER 3

A LIMITED VIEW, GOAL DRIVEN, VOLUME SOURCE METHOD OF BACKPROJECTION RECONSTRUCTION FOR LIMITED OBJECT INSPECTION

3.1 Introduction

In this chapter an algorithm is described which will inspect a limited set of specific locations on an object for accuracy. The algorithm is goal driven in the sense that the inspection is focused on these locations only and therefore only requires a limited number of views of the object. It is able to accurately inspect any object location which is locally *Weakly Convex*. The locations of interest will be reconstructed using a volume source technique which projects elements of the reconstruction volume onto the sensory data.

In Section 3.2 we shall discuss the different factors which come together to form this goal driven inspection analysis. This will include a background discussion of machine inspection and proposed modifications to octree object representation, object views, and volume source reconstruction. Section 3.3 will present goal driven inspection. The section concludes with an example inspection trace of the algorithm. Section 3.4 will discuss the flexibility of the algorithm under different goal sets. Section 3.5 discusses aspects of our implementation of the algorithm.

3.1.1 Related Topics: Recognition and Backprojection Reconstruction

Recognition and backprojection reconstruction are two areas of three dimensional computer vision which examine similar data to that which we are using for our inspection. Aspects of these two fields are presented to form a background for the development of our inspection analysis.

Object recognition is the process of extracting information from a scene or scenes and deducing the identity of the object or objects in the scene. Often this is done using features which are significant to the object. The problem of backprojection reconstruction is a method of object reconstruction which makes use of the silhouette views of an object. Intersection techniques are used to determine information about the original object.

Backprojection Reconstruction

Three Dimensional object reconstruction is the process of extracting three dimensional information from available sensory sources. Reconstruction through backprojection has been examined from a number of approaches. Some have used perspective silhouettes [2, 25, 26, 48, 59, 83, 123]. Another approach has been to use orthographic views such as what might be found in engineering drawings [11, 14, 22, 36, 41]. At least one other has used other sensor sources, such as X-rays [121] to produce similar results. Another similar approach has been to use stereo and motion in the reconstruction [40, 52, 135, 136, 143, 146].

Tan [123] uses the approaches of backprojection reconstruction to address industrial inspection. These approaches all project the views into the reconstruction space. A different approach was presented by Potmesil [97]. He proposes projecting the reconstruction voxels onto the images and evaluating the overlap to determine the state of the voxel during the reconstruction. The voxel reconstruction phase of our inspection will follow this approach.

General backprojection reconstruction is straightforward in its implementation. An object to be reconstructed will have its image taken from a number of views. This may be accomplished by locating the object on a light table and employing a camera held in a robot arm to image the object from a variety of angles.

Since we are using silhouettes of the objects as the initial observation data, we are limited in the kind of information we can obtain about the object. This limits us to the types of objects which can be reconstructed as defined in Chapter 2.

3D Object Recognition

In object recognition the goal is to reduce data. The goal is a mapping of a quantity of sensory information into a smaller, usually individual, set of data. Obviously a lot of resolution may be lost in this process, but the information obtained is on a "higher level." The amount of information lost is a function of the how much difference there is between the data implied by the "high level" data and the actual data itself.

Often, object recognition is based upon the extraction of certain feature sets from the object. It is these features which distinguish the given object from others. Inherent in this argument is the assumption that the significant and distinguishing information about the object is found in these features. This has the added effect that the "high level" identification often implies the presence of the underlying feature set.

Several algorithms have been developed for the recognition or identification of 3D objects from images [25, 33, 49, 68, 78, 90, 91, 140, 145]. Others have emphasized the reasoning aspects which can be involved in proper identification [12, 60].

Reconstruction and Recognition Compared

Both reconstruction and recognition share the feature analysis of three dimensional objects with a goal of the extraction of information from the data. One difference is the type of data to be extracted.

Reconstruction seeks data without a loss in accuracy and recognition seeks a higher level data descriptor which may not totally describe the object in question.

Another difference is in the amount of a priori information used for processing. In object recognition, it is not assumed that the object or its orientation is known prior to processing. In reconstruction, it is assumed that both the object under examination and its orientation relative to the imaging system is known. It can be argued that there is therefore no need to perform the reconstruction. This is not true if the purpose of the processing is identification. Our goal in inspection is a verification of known object properties.

Since the goal is verification of a specific set of object properties it is appropriate to make use of the hoped for properties in the evaluation of the object under examination. In recognition, this might tend to bias decisions toward objects containing the properties assumed in the analysis.

One final difference is that recognition will often use features to distinguish one object from another. Reconstruction does not use features, but produces features as part of its processing.

Recognition and Reconstruction Implications for Inspection

There are several different areas where recognition and reconstruction analysis have suggested analysis techniques for inspection. These include data structures for analysis and data extraction techniques.

Efficient data structures for analysis are a common goal of both procedures. Object data even more than image data requires large amounts of computer memory. The octree data structure which was modified by us for inspection was initially used by both reconstruction and recognition.

Data extraction is also a common goal of both procedures. Whether you are attempting to recognize an object or reconstruct an object from two dimensional images, the mapping of image data to object data is the same. A technique similar to the one we shall employ for local reconstruction has been applied to general reconstruction problem.

3.2 Motivation for Goal Driven Inspection Analysis

In many instances it is possible that the entire surface of an object being accurate is not critical to the usefulness of the object. In those instances an analysis of the entire surface of the object is superfluous. It would be more efficient to concentrate the inspection analysis on the areas of the object which are significant to the accuracy of the object.

Such an inspection analysis routine should reconstruct the entire object as the number of significant locations on the boundary of the object reaches the limit of including the entire object. In this case, it is possible that the inspection analysis would be less efficient than a general reconstruction of the entire object. This would be due to a goal driven inspection analysis being optimized towards those times when the number of significant locations are a small fraction of the overall surface of the object.

Several different aspects of goal driven inspection shall be discussed as motivation for its use. The first of these concerns the use of object features in the inspection process. Second, we shall examine the suitability for using octree data structures similar to those used in object reconstruction to efficiently represent the information needed for a goal driven inspection. Third, the ability to limit the number of views needed for goal driven inspection will be presented. Lastly, we shall examine the use of volume source backprojection reconstruction as an efficient means to perform goal driven reconstruction of small areas of the object.

3.2.1 Motivation for the use of Object Landmarks

In many cases there are a small number of locations on the object which are significant to the usefulness of the object. These locations shall be called the object's landmarks. This is similar to the way that object recognition uses features to identify the object. Here we will use landmark verification to "identify" whether the object is "correct."

landmarks will be used to represent several different locations and types of region boundaries on the object.

One simple use of landmarks is to represent points on the boundary whose location needs to be know. An estimate of these locations could be used to determine the absolute error in an object. If this error was too great, the part could be rejected.

Another use of landmarks would be to identify the boundary of a region whose area or perimeter was significant. Figure 3.1 shows a part whose end is marked by landmarks. If the function of this part were to press against something to make an impression, the only part which would be significant might be the points indicated in the Figure. The use of landmarks introduces a representation error between the actual parameter value and the estimated value which can be computed from the estimated landmarks. This representation error will approach zero as more landmarks are used to represent a boundary.

Finally, another use of landmarks would be to sample the object surface for smoothness. The standard deviation of the distance between the actual landmark and the estimated landmark could be used for this.

3.2.2 Motivation for Octree Recursion

A key part of the recursive nature of the algorithm is a structure for representing the volume which can be reconstructed. For this case an octree representation has been chosen for its adaptability to recursion and its compactness.

Octrees are well suited for recursive analysis because of the ease with which another layer of the octree can be added to the structure. This is in fact how this algorithm is implemented. Figure 3.2 shows an example of octree representations and volumes split up for octree representation. In our case these eight vertexes represent cubes which are in turn arranged to form a cube. Each of these vertices is the root for a sub-octree. This continues on to an arbitrary representation resolution.

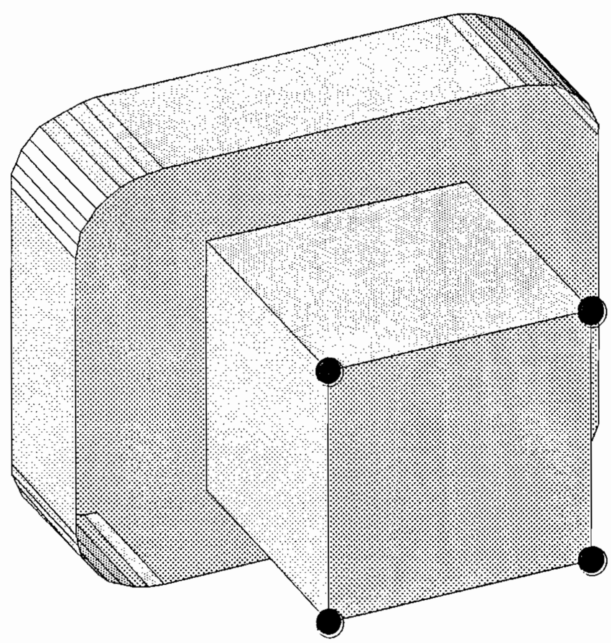

Figure 3.1 Example landmark use to identify significant locations on an object

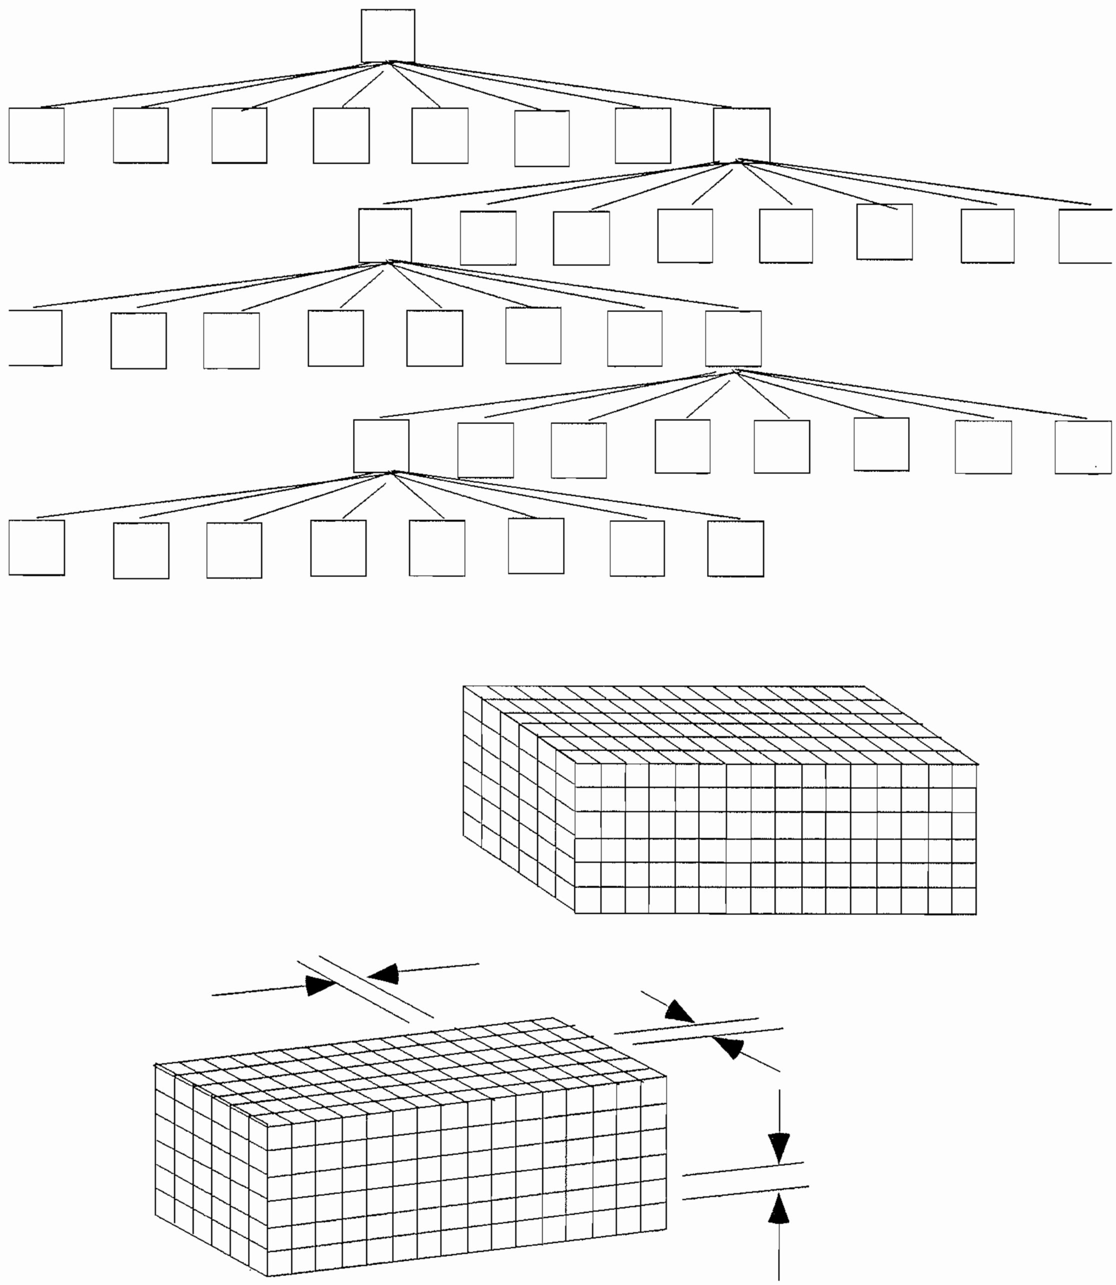

Figure 3.2 Octree representation scheme for an object/identification volume

Octrees have been extensively used for data representation. In our modified octree, it will only contain branches which lead to areas of the volume represented by landmarks to the inspection. Each lower level of the octree provides a better description of the exact location of a landmark. This allows for an adaptable representation scheme which can be modified in the detail of the description as needed.

The algorithm will recurse down the inspection octree to the level required by the accuracy and the granularity of the result. If a section of the octree is found not to have the required accuracy the inspection routine may be exited immediately. Otherwise, a complete traversal of the octree can be done and the accuracy of the object calculated.

A second reason for the selection of an octree for data representation is that it provides for a compact data set. It removes the need for the storage of the entire reconstruction volume. One of the problems with three dimensional computer vision is that there is a vast amount of information to be stored and manipulated. An octree represents information in a compact, easily manipulated manner.

3.2.3 Motivation for Limited View Processing

During the inspection process it is important that all processing be focused upon the solution to the inspection analysis. One implication of this focused effort is the limiting of the reconstruction of the object to those areas significant to the inspection. In the general reconstruction problem many views may be needed to reconstruct areas of the object which do not change the overall usefulness of the object. This reconstruction is unnecessary for inspection problems which require inspection of only a small portion of the object.

Further, processing views for the reconstructions should be limited to those views which provide the most information about the areas inspected. Chapter 2 provides a description of the minimum requirements for a boundary location to be reconstructed. It does not provide a means to group these individual view lines into views. Ideally, we would select all view lines to be part of a single view. Intuitively and in practice, a single view is not sufficient, but the goal of limiting the views is to use as few as possible.

We desire a set of views which shall efficiently describe the limited set landmarks on the object which are needed for inspection. Efficiency shall be evaluated in terms of the views being the minimal set of views needed to perform the inspection analysis presented. Chapter 5 will present an algorithm for the best way to select these limited views.

3.2.4 Motivation for Volume Source Method of Reconstruction

There are two different ways to relate the information in the three dimensional reconstruction space to the information in our two dimensional views which are consistent with backprojection reconstruction. The first way is to project the images into the reconstruction space and perform an intersection of the projections. We shall refer to all routines which use a variation of this method as volume intersection methods (Figure 3.3). The second is to project the elements of the reconstruction space onto the views and perform an intersection of the objects original projection and the projection of the reconstruction space. We shall refer to all algorithms which use a variation of this method as volume source methods (Figure 3.4).

We shall use the volume source method for several reasons. The first, and perhaps most important is that it is computationally easier to perform an intersection in two dimensional space than it is to perform one in three dimensional space. Thus, the volume source method provides a more efficient reconstruction.

Second, projection of the reconstruction volume onto the views is more easily adapted toward a limited reconstruction of the object's surface. In volume intersection reconstruction, it would be difficult to limit the intersection to the regions of the object which are of interest. Since the volume source method projects from the reconstruction space onto the image space it is possible to only perform those projections which are of help to the inspection analysis.

Third, the volume source method performs better at a variety of resolution combinations between the views and the reconstruction

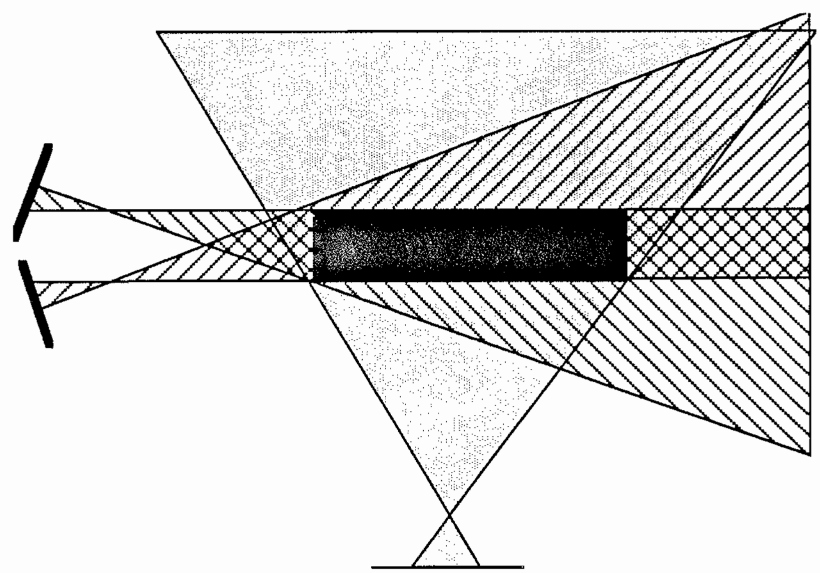

Figure 3.3 Volume intersection backprojection reconstruction

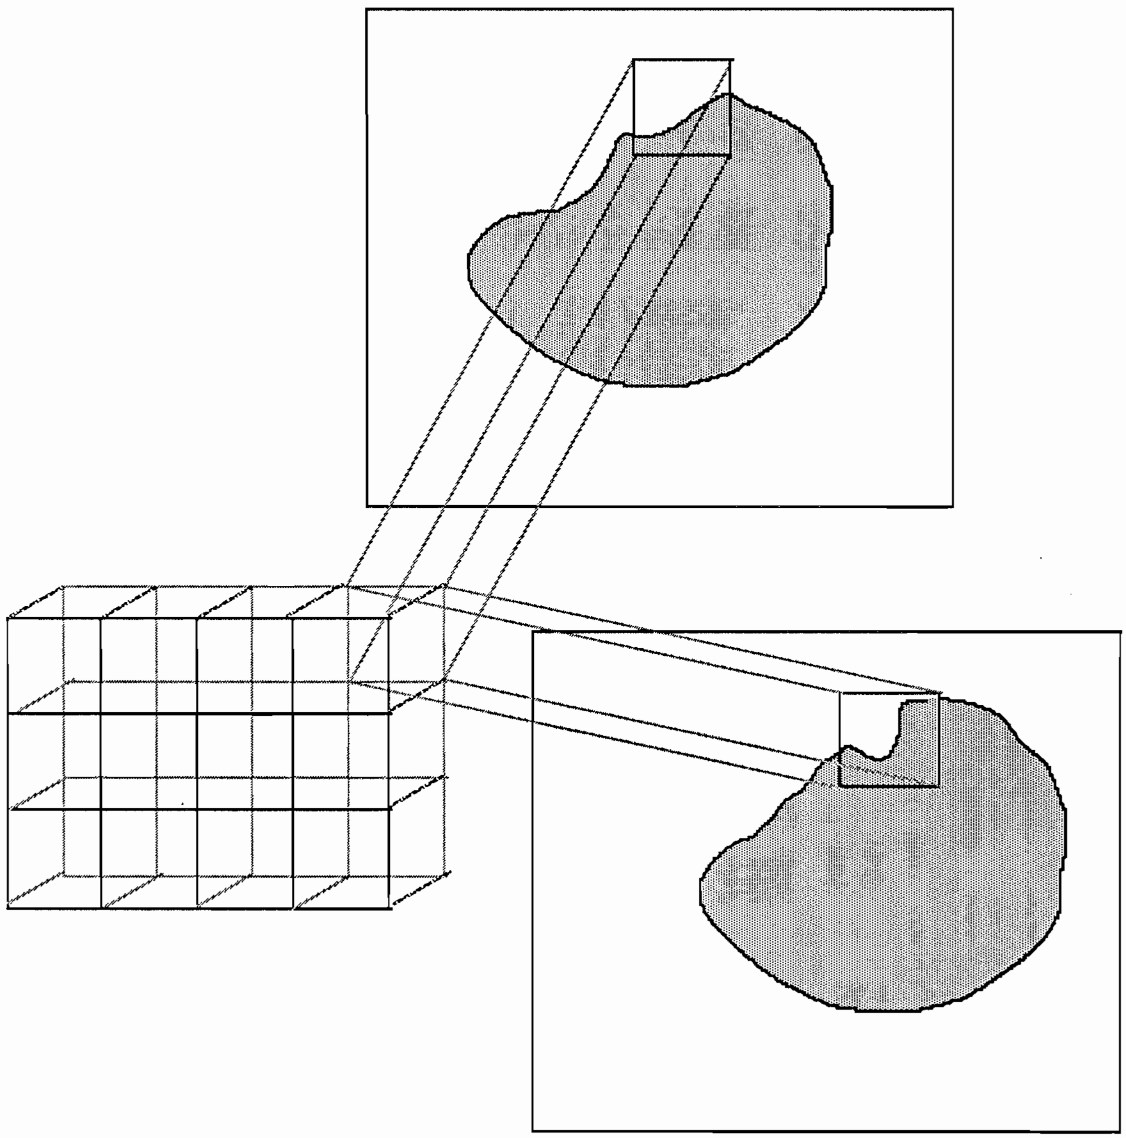

Figure 3.4 Volume source backprojection reconstruction

volume. This difference in accuracies is due to the volume source method's ability to take advantage of all the information in the overlap between the original view and the volume projection while most volume intersection methods estimate voxel location with a single point.

Finally, the volume source method can accurately label as void any volume where more than half the projection into a single view is background using only a single view. This 50% overlap rule will be discussed in more detail in Section 3.3.4 and Chapter 5. Volume intersection techniques require an examination of several views before a decision on a single voxel can be finalized.

3.3 Aspects of Goal Driven Inspection Analysis

The goal driven inspection analysis problem is one where there is no effort to inspect the entire object. Only those portions of the object which are of special significance are inspected. Here we shall examine several different parts of this problem. The first and most important part is the octree inspection map. This map provides a description of the object for the inspection algorithm to follow. Second, the information maintained about each reconstruction voxel is discussed. Third, the process of using a limited number of views of the object is presented. Finally, the reconstruction process itself is presented.

3.3.1 The Octree Inspection Map

Before the actual inspection is done, the object must be preprocessed for inspection. This phase creates a map which will guide the actual inspection process. It presupposes a knowledge of the object, as well as a decision on the landmarks of the object. This map will contain the landmarks critical to the object's usefulness by marking voxels which need to be present on the object and by marking voxels adjacent to the object which must be void of object.

Critical to the usefulness of the preprocessing phase is the assumption that the object to be inspected is registered with the inspection system. The preprocessing phase builds an exact model of object detail for the inspection. This exact model is not useful if the object is not where the model presupposes. In our analysis, we assume that the object is registered with the inspection system. In actuality, this is not a unreasonable assumption since the object was in registration with the manufacturing apparatus during manufacture. Ideally, inspection would take place immediately after each manufacture, and could thus make use of the registration already in place.

Landmarks which identify significant features could be specified for the object to be preprocessed by users during the CAD object design. These landmarks might define a region, hole, or corner which is significant to the overall usefulness of the object. Object examples with landmarks indicated are shown in Figure 3.5-7. In each figure a different type of significant feature is indicated by the landmarks. In Figure 3.5, the dimensions of a post are determined its four corner points. In Figure 3.6, the dimensions of a hole are determined by several samples along the hole's perimeter. In Figure 3.7, the smoothness of a plane face is estimated from several samples on its face.

These features are then represented as part of a sparse octree where the only branches present are those which lead to the features -- our landmarks. The associated sparse octree for the object of Figure 3.6 is shown in Figure 3.8. A sparse octree such as this can be efficiently used to represent the landmarks as well as to provide an efficient way to scan the landmarks during an inspection. One traversal of the tree will pass through all landmarks. In fact, a single satisfactory tree traversal indicates that an object has passed a minimal set of inspection criterion.

An octree representation of an object by its nature presents an increasingly detailed description of the object as tree depth is increased. Our sparse octree provides an increasingly detailed description of significant boundary locations of the object as depth is increased. We will show later that the inspection analysis can take place at a variety of levels within the octree. The level at which the inspection takes place is determined by the accuracy required for the examination of the object and the granularity that is allowable in the result of any inspection. This granularity occurs because specific object boundary locations are only known to be within the size of the boundaries of the current voxel.

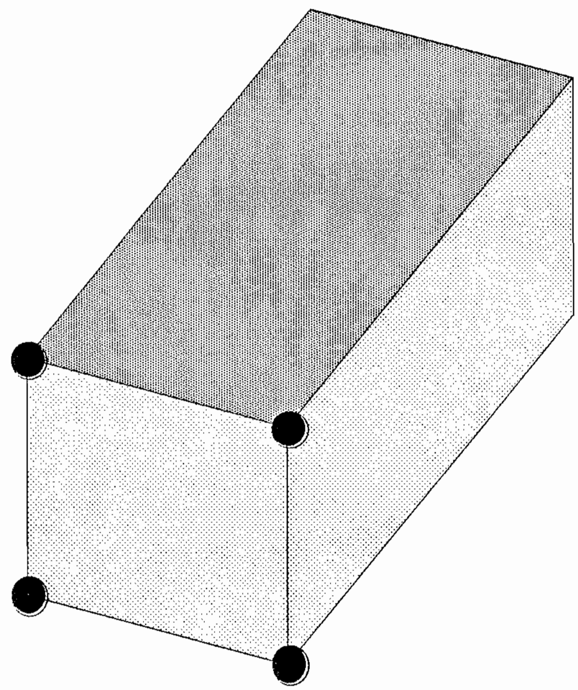

Figure 3.5 Landmark identification on a post for dimension determination using four landmarks

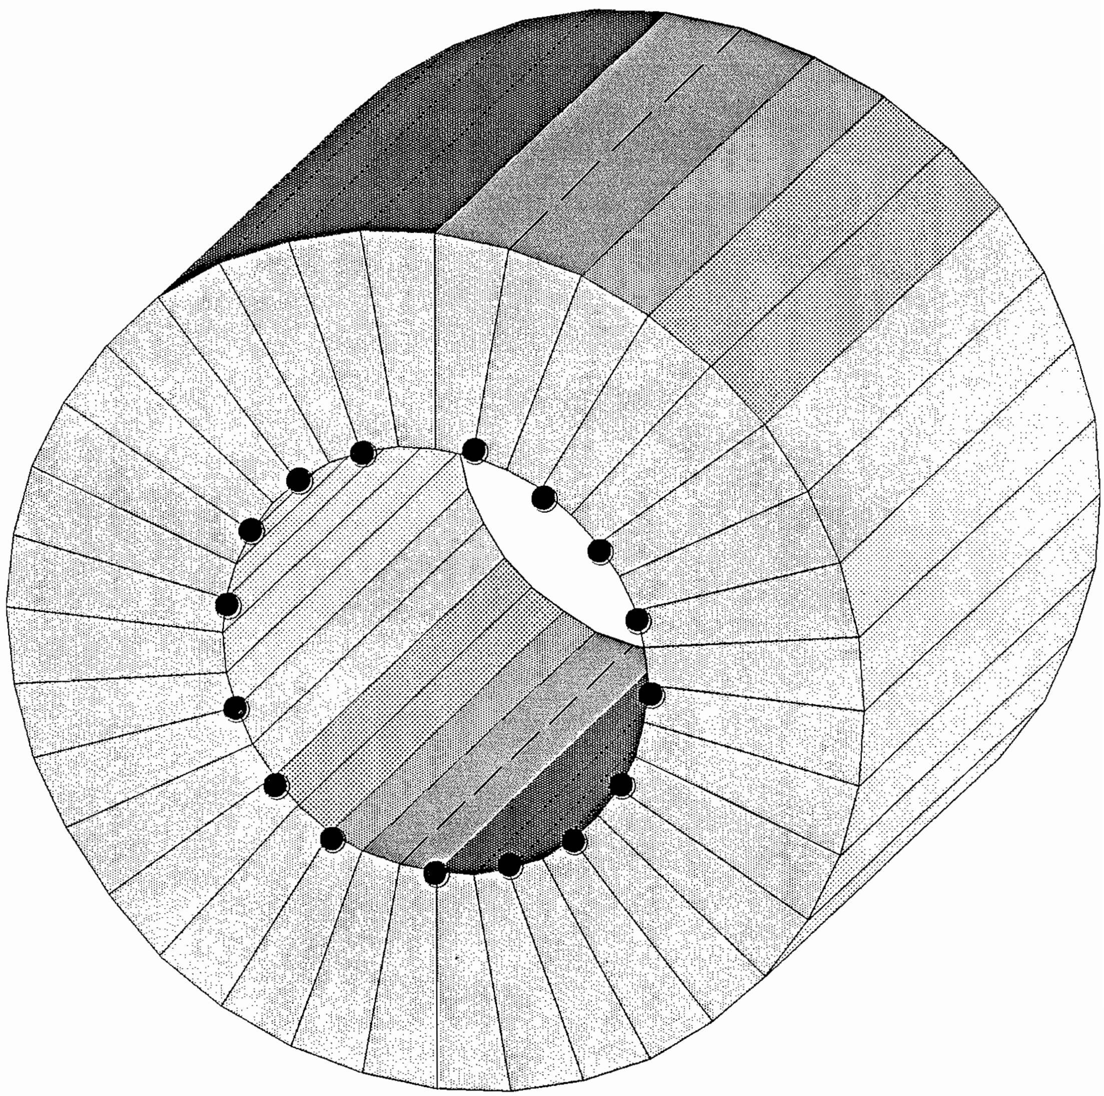

● - Landmarks at 22.5 degrees

Figure 3.6 Landmark identification on a circular hole for area determination using sixteen landmarks

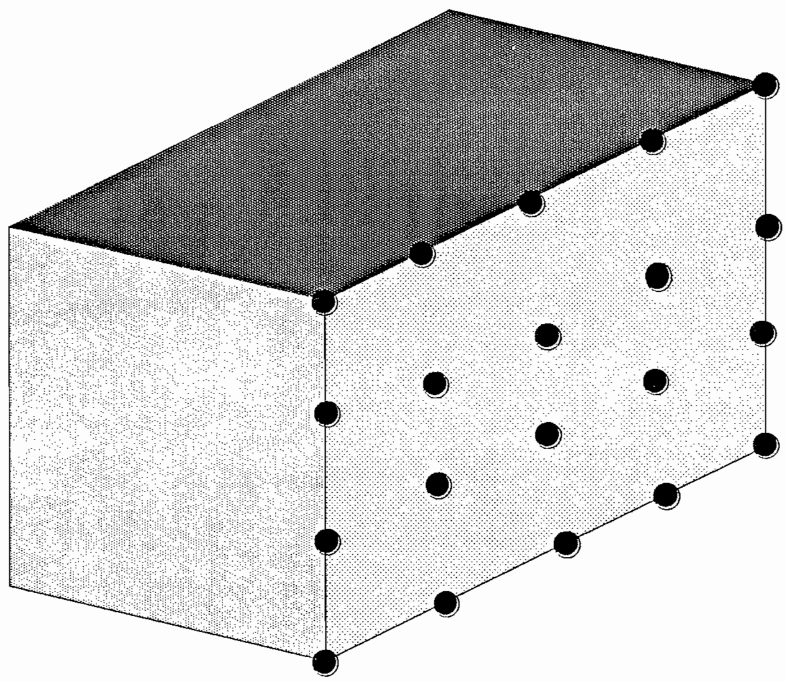

Figure 3.7 Landmark identification on a surface for smoothness determination using twenty five landmarks

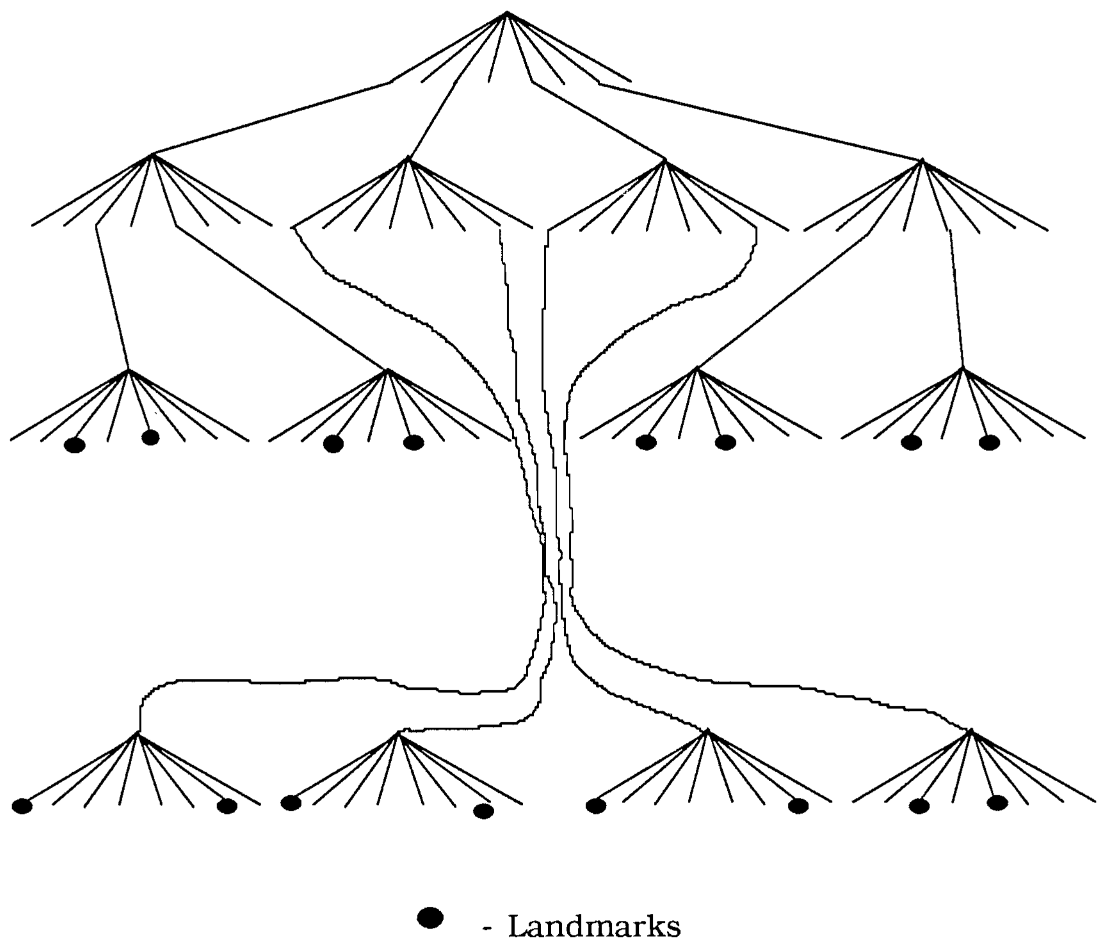

Figure 3.8 Sparse octree representation of the landmarks of Figure 3.6

Thus far, only the landmarks on the boundary of the object have been represented in the octree. From Chapter 2 we note that it is not possible to reconstruct specific boundary locations on an object; we must reconstruct locations adjacent to these boundary locations. Therefore, our octree will be extended at each level where an inspection might occur to locations adjacent to the boundary location in question. Unlike extensions to an octree which subdivide octants into increasingly smaller descriptions of a region; these extensions describe adjacent voxels of the same size.

It is possible that there may be an overlap in the representation of the adjacent voxels. This shall be represented by a divergence from a normal tree structure. A single child *adjacent* node in the tree structure may have several parent adjacent nodes. Tree structure is maintained through the various resolution levels of the octree which lead to landmarks. This modified octree structure is show in Figure 3.9. In the Figure, the octree portions which lead to landmarks are shown in dark lines with filled circles for nodes indicating increasing resolution. The octree portions which represent adjacent locations are shown with grey lines with grey circles for nodes. Note that it is possible for a voxel to be both on the path to a landmark as well as an adjacent node to another boxed location. Further, note that there are no adjacent node links among resolution levels.

Chapter 2 indicates that it is only necessary to know the adjacent locations described in the octree thus far to know that the object point is in fact a boundary point. In the inspection problem, we must allow that the boundary may not be where we want it to be. One option mentioned earlier is to reconstruct the entire reconstruction space so that all boundaries are found. Our approach is to reconstruct only those areas significant for the inspection. These significant reconstruction spaces may extend beyond a single adjacent reconstruction voxel. The extent to which this representation expands is dependant upon the accuracy required in the object before it is rejected by the inspection process. We reconstruct the surrounding region which represents the allowable error in the object.

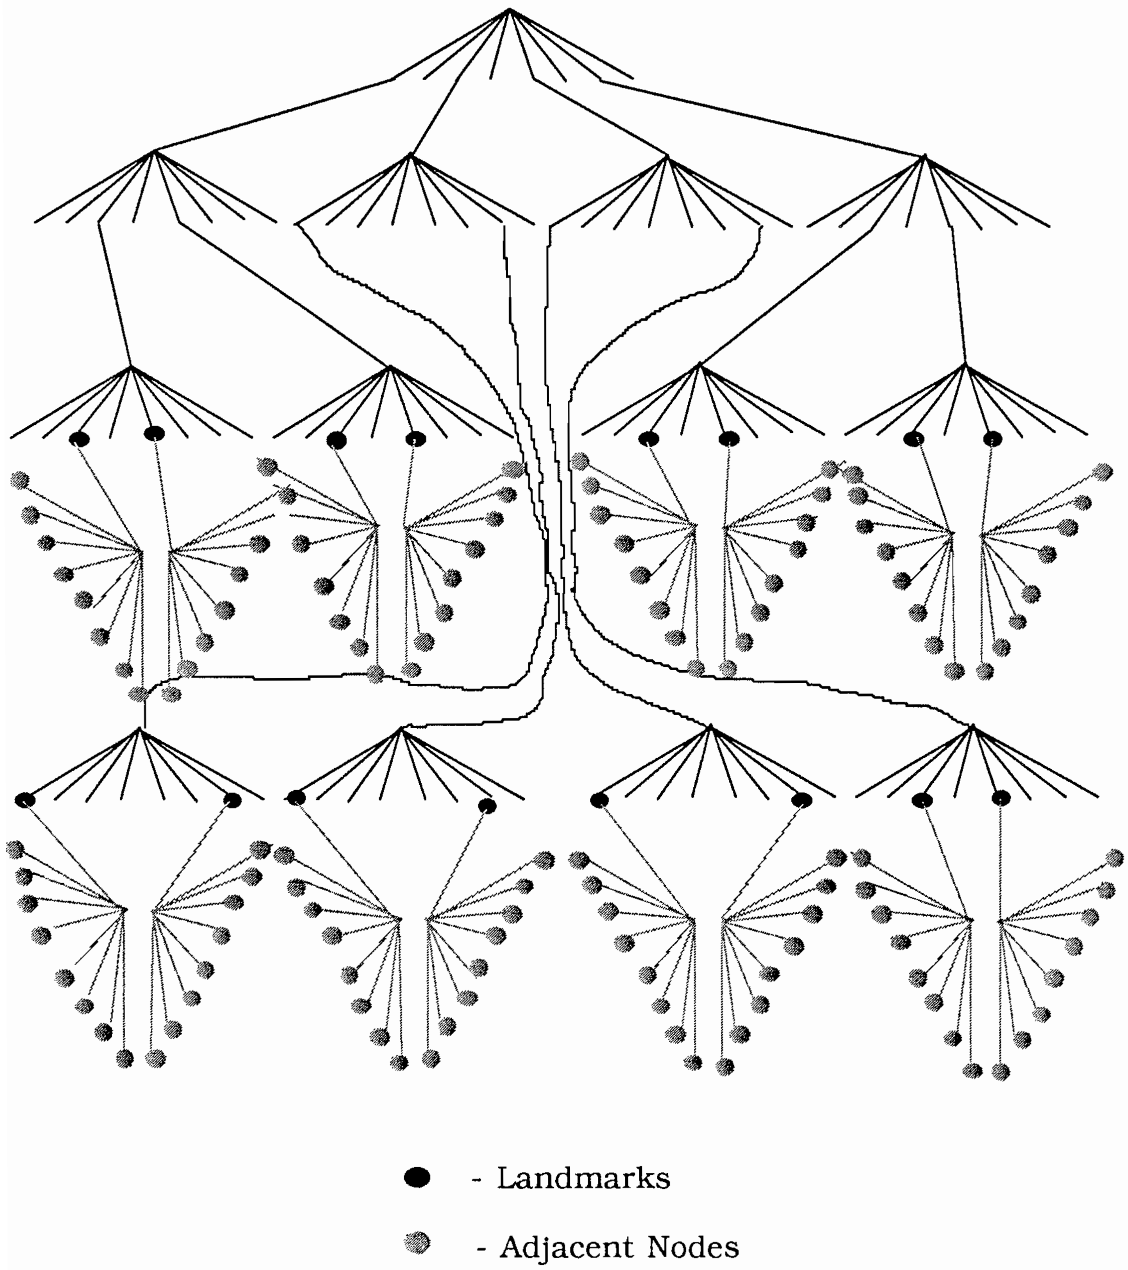

Figure 3.9 Sparse octree representation of the landmarks of Figure 3.6 with adjacent voxel sharing indicated

Figure 3.10 shows a single octree resolution level with shared adjacent nodes. At each level of octree resolution a network of adjacent nodes is formed. Each adjacent node may have other adjacent nodes attached to it. This network is searched during the inspection process for a certain distance from each boundary voxel to determine the status of the locations around the significant voxels. This process will be described later.

Each voxel will be reconstructed only once, and then only if it is needed. The result will be stored in the modified octree. During an inspection traversal to the octree, when a non-reconstructed voxel is reached it shall be reconstructed and the result both stored and passed back to the calling voxel. In this manner, each voxel will only have to be reconstructed once.

3.3.2 The Voxel Data Structure

The modified octree just described will also be carried into the data representation at the octree leaf level. There will be four types of nodes in our modified octree. The first type is the degenerate, missing node whose children lead only to portions of the reconstruction space where there are no landmarks to the reconstruction. The second type of node is one on a path leading to at least one landmark, but which is at a resolution level within the octree where the inspection would never be carried out. The third type of node is the leaf node representing such a landmark. The fourth type of node is an node which is adjacent to another adjacent node or a type three node. Each of these nodes will be described below.

The type one node is the nonexistent node. It is represented as a null pointer in the data structure at the highest level where all children below that point in the octree are insignificant to the inspection.

The type two node is singly linked downward to each of its eight children. This node provides a description of the inspection structure above the level where the inspection will occur and may be thought of as a normal octree node.

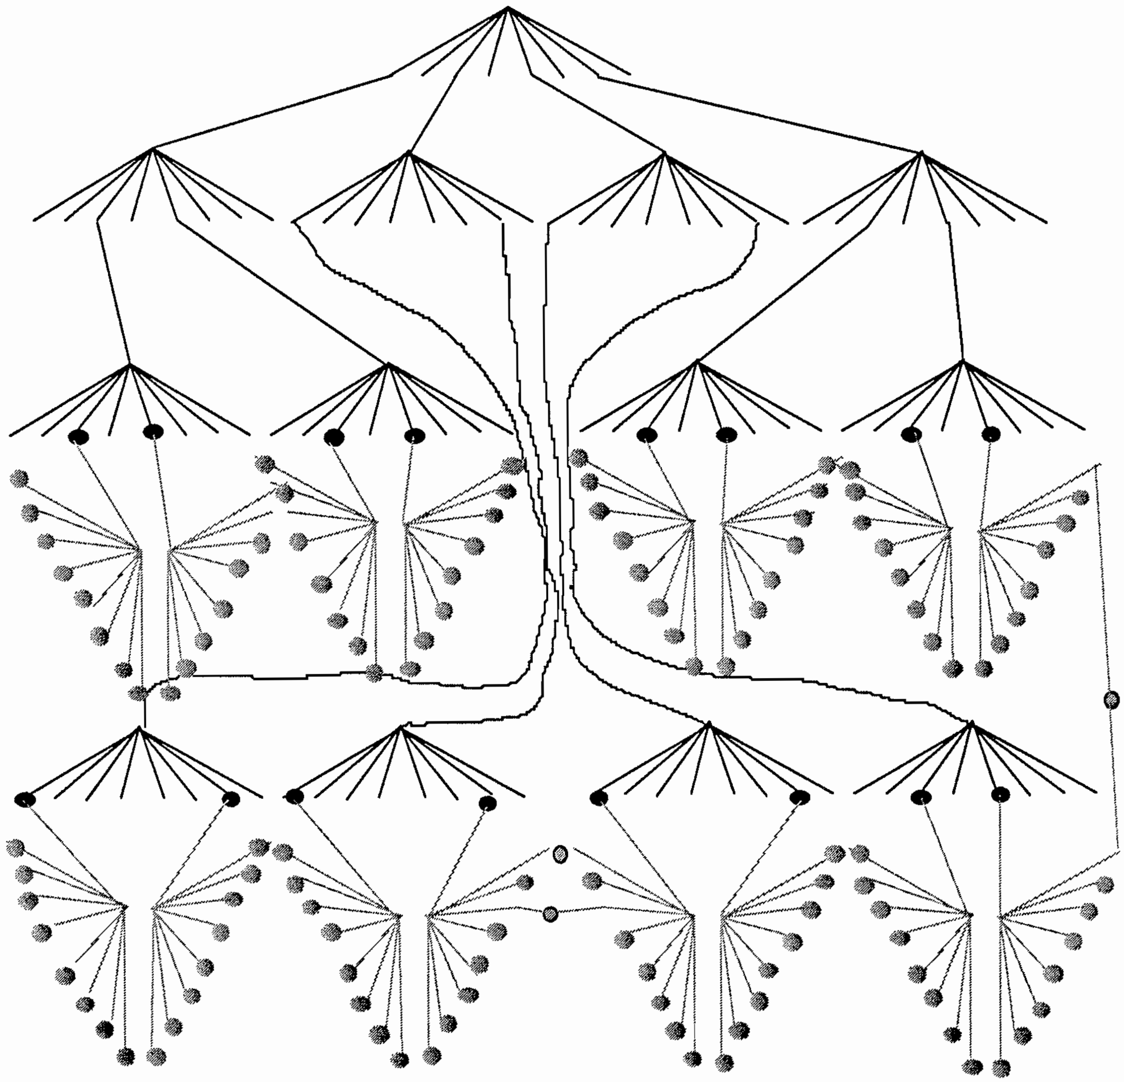

- Landmarks

- Adjacent Nodes

- Shared Adjacent Nodes

Figure 3.10 Single layer of the octree representation of the landmarks of Figure 3.6 with multiple layers of adjacent voxels

The type three node is a singly linked downward node at all levels except the leaf level of the octree. At the leaf level this node has null pointers to its children. At other levels this node provides downward links to its children providing structure to the inspection. A type three node also has pointer to its siblings. This set of twenty six pointers indicates the neighbors of the type three node at the same reconstruction resolution as the node.

These siblings shall be type four nodes. There role shall be to determine if the areas adjacent to a significant boundary location are indeed void. This, combined with the significant boundary location being filled with object will be sufficient for the inspection analysis to determine that the object is correct by the definitions of Chapter 2.

The type four node is a doubly linked node. The double linkage is to each child adjacent node and to each parent node. The second linkage is obviously a subset of the first list and the two lists are stored together. It should be noted that a type four node contains all the information of a type three node and when a landmark is also an adjacent node, it is represented as a type four node.

3.3.3 The Limited Views

One of the significant features of our Goal Driven system is its ability to use a limited set of views. The views which are needed by the goal driven system are only those required by a Chapter 2 analysis of the landmarks to be inspected. *Weak Convexity* indicates that there are points in the neighborhood of the landmark boundary point which are visibly void.

Since the Goal Driven Inspection process only requires that we reconstruct a limited number of boundary points, we only need views which will enable us to reconstruct those points. This will result in a number of views which is less than or equal to the number of views needed to reconstruct the entire object. The exact number of views will depend upon the number and arrangement of the locations to be reconstructed. Typically, this could be much less than the number of views needed for an accurate reconstruction of the entire object.

We shall assume that a set of views are collected that is such that information required for the analysis is present using the fewest number of images possible.

3.3.4 Volume Source Reconstruction

Reconstruction through backprojection implies the use of two dimensional image data to infer the shape of a three dimensional object. The volume source method of reconstruction used here relies on the forward projection of volumes from the reconstruction space onto image data to more accurately determine the shape of the object. This is different from other methods which rely upon point sources for forward projection and methods which rely on pure backprojection. This reconstruction method closely follows the work of Potmesil [97].

The advantages of the volume projection scheme are derived from the additional information present in the volume projection. Point source models assume a symmetry of the voxel which is not true. When a point source is projected onto an image it does not take into account the orientation of the voxel with respect to the image. Thus, it cannot accurately represent what happens when the voxel was originally imaged. Voxel projection presents a better representation of the actual portion of the object which was originally imaged when the data was taken. These two methods are compared in Figure 3.11.

The heavy projection line in Figure 3.11 shows a forward projection from the center of a voxel onto a 2D projection of a 3D object. In the Figure the center of the voxel projects onto non-object in the image. The projection of the entire voxel onto the image is shown by the grey lines. In this case, a significant portion of the object is within the projection. While insufficient illumination from the object is present to turn the image pixel "on", this information may be significant in conjunction with other projections. The volume source method used allows for this information to be taken into account during the reconstruction. These effects are significant whenever the resolution of the voxel representation is less than or equal to the resolution of the object images.

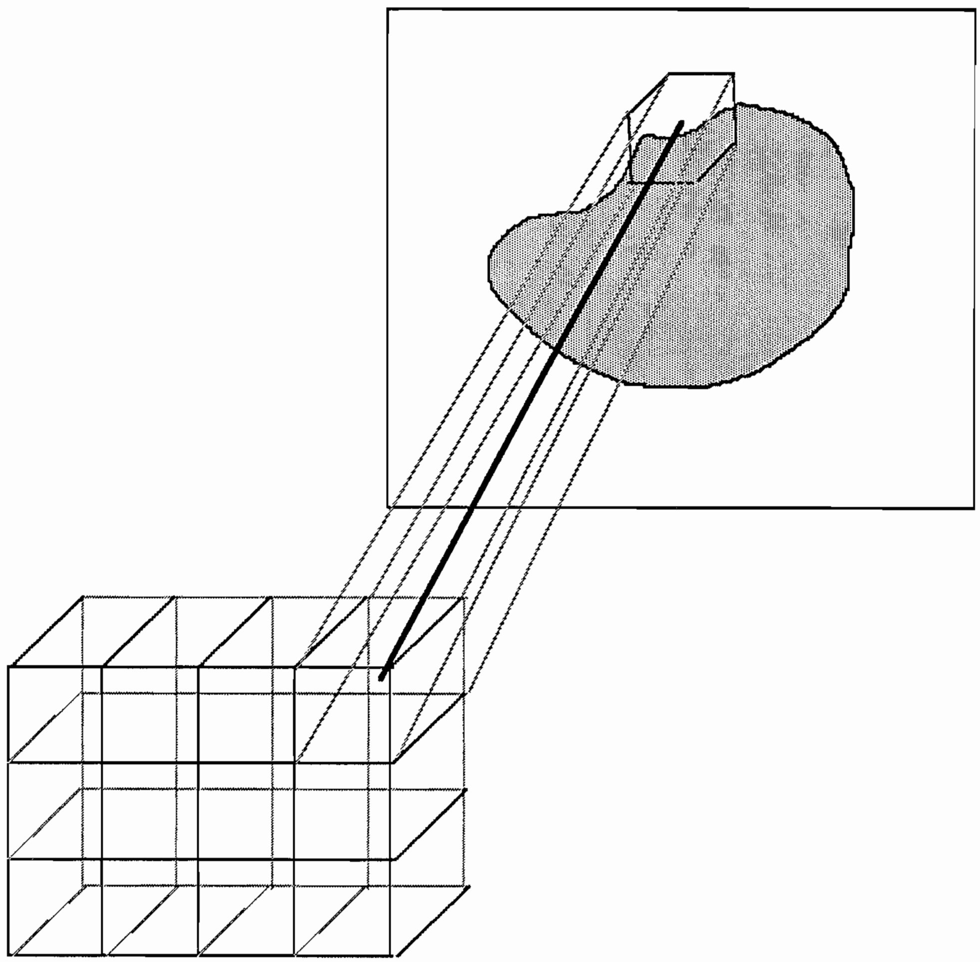

Figure 3.11 A comparison of Volume Source and Volume Intersection techniques for backprojection reconstruction

The Volume Source Method of Voxel Reconstruction

Central to the volume source method of reconstruction is voxel projection. A flowchart of the general voxel reconstruction routine is presented in Figure 3.12. Since a key part of the reconstruction is the projection of voxel vertexes this part of the algorithm will be examined first.

This point projection is a translation of coordinates from the object reference frame to the image reference frame. Several elements are involved in the transformation. The key elements are the camera focal length, pixel size, optical axis, and relative camera/object angle and position. The mathematics relating positions in these two reference frames are relatively straightforward [97, 123].

In order to project a voxel onto an image, the orientation of the image with respect to the voxel is first calculated. This is done by determining from which of twenty six quadrants surrounding the reconstruction volume the view was taken (Figure 3.13). A table look up method is used to determine the relative position between the object and the camera and the orientation of the voxel projection. This avoids the necessity of doing location coordinate calculations and comparisons.

This always uniquely determines which vertices of the voxel will be vertices of the projection onto the image and their relative orientation and connectedness. Figure 3.14 shows the relationship between the octant location of the view, the view projection, the view shape, and which vertices are projected. This allows us only to project the vertexes of interest. Since we also know the connectedness of the vertices in the projection, this provides a bounding curve in the image of the projection.

Once the projection has been determined it is possible to compare the projection of the voxel to the image of the object. The intersection between the projection and the image gives an indication of whether the voxel could have generated the image. If the intersection is empty, then the voxel is not part of the object. If the intersection is equal to the object, the voxel is part of the object. If the intersection is partial, the amount of overlap must be examined to determine if voxel identity can be established.

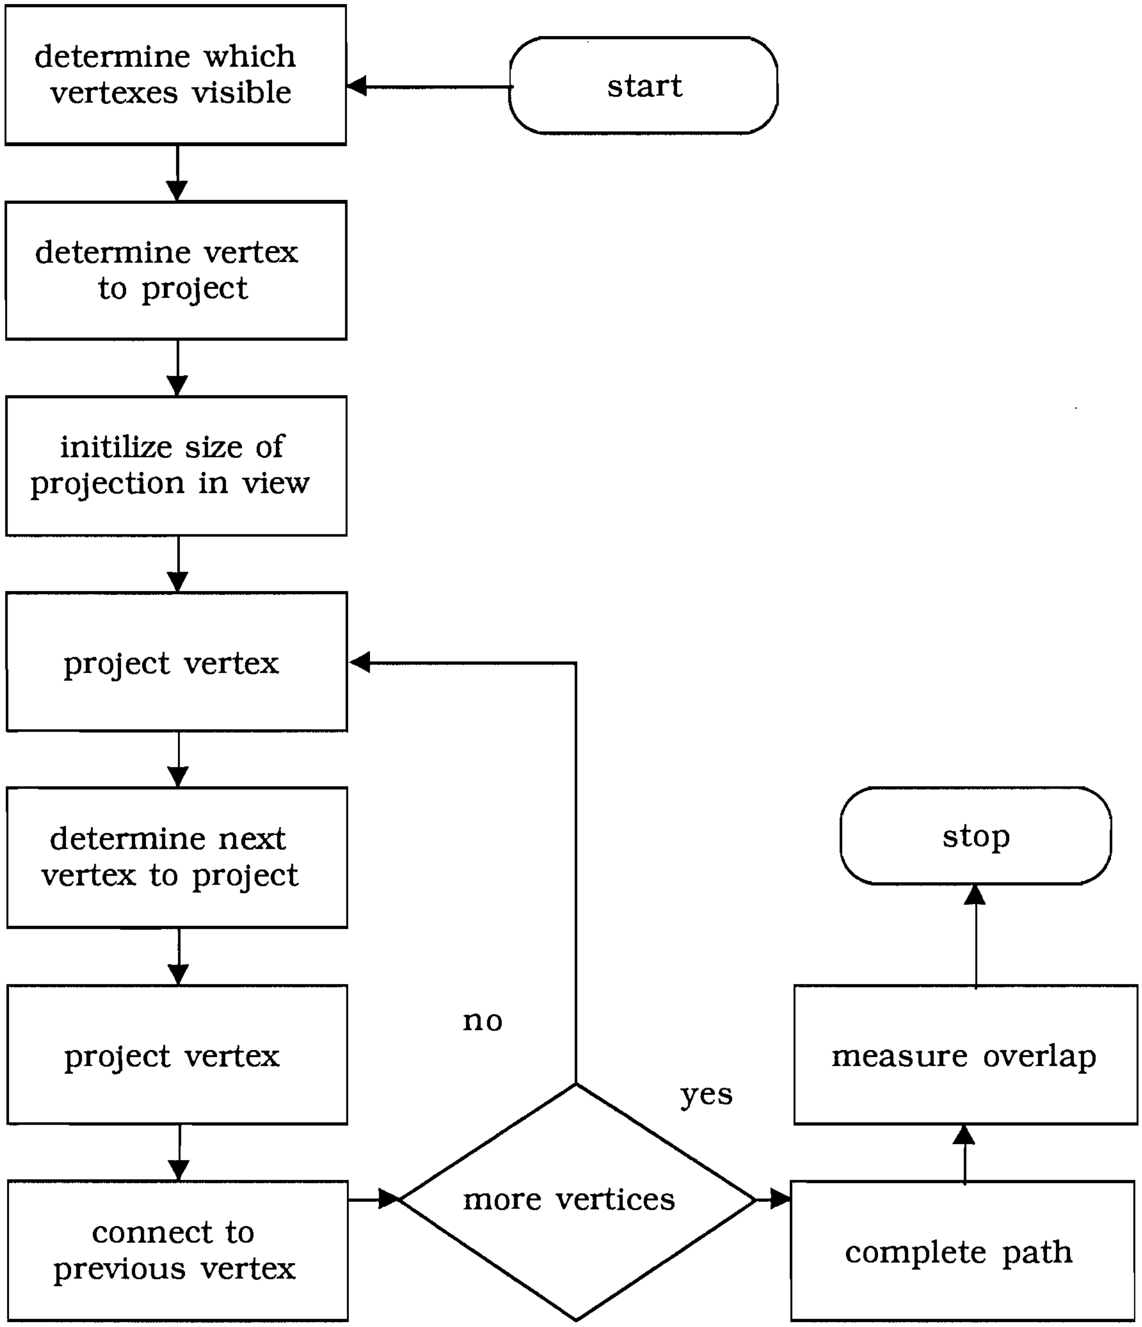

Figure 3.12 General Volume Source voxel reconstruction procedure

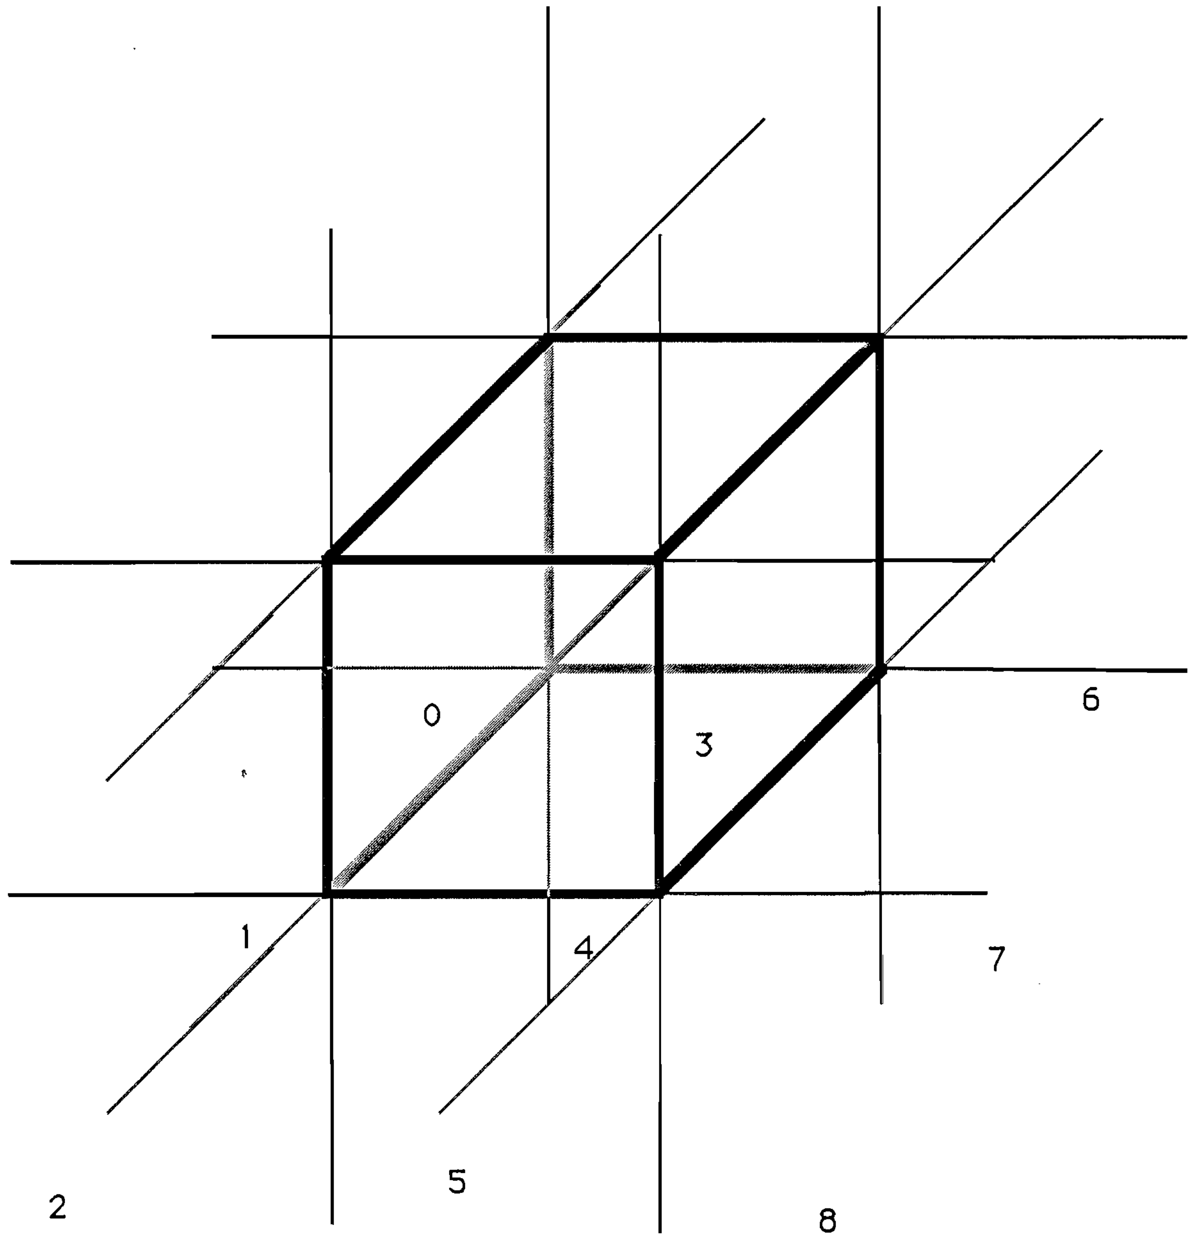

Figure 3.13 Quadrant map for view projection determination

View from
Quadrant

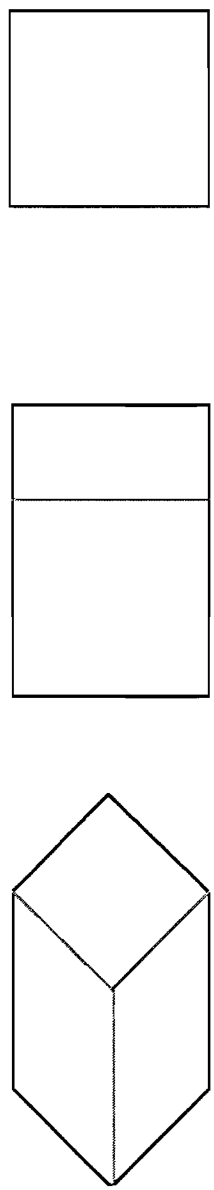

4,10,12,14,16,22

1,3,5,7,9,11,15,
17,19,21,23,25

0,2,6,8,18,20,24,26

Figure 3.14 Possible Volume Source projections onto views by quadrant

From geometric considerations, it can be shown that an intersection that is less than 50% overlap will insure that the voxel identity is always void. Figure 3.15 shows some examples of different voxel descriptions which yield 50% overlap projections. Anything less than 50% overlap is indeterminate. In practice, this means that as soon as 50% of the overlap has been evaluated to void the overlap scan can be stopped and the voxel labeled void. This will be used in Chapter 5. Here, we shall assume that be will always be able to view at least 50% overlap of any voxel which should be void. In Chapter 5, the view selection process will insure that this is correct.

3.3.5 Sources of Conflict during Reconstruction

Sources of conflict in the reconstruction process can come from two main sources. The first is inaccuracies in the view/reconstruction volume registration and calibration. The second is inaccuracies in the segmented and thresholded image itself.

Imaging setup calibration inaccuracies arise from imprecision in the knowledge of where the images are with respect to the object. This is illustrated in Figure 3.16. Different views which may have almost identical information have different projection images associated with them. This can result from camera alignment, object alignment, or the mechanism which moves the camera from view to view. As noted earlier, the purpose of this reconstruction is a comparison of known object features to the reconstructed object, therefore registration is assumed and mis-registration is treated as an error.

Images inaccuracies can arise from a number of areas. The result of all the inaccuracies is a improperly segmented image. Major causes result from the interaction of the object surface, lighting, and camera orientation. Figure 3.17 shows how improper segmentations can result in reconstruction conflict. Note that two improper segmentations have reinforced each other and have created a potential error in the boundary of the object.

This illustrates an interesting and harmful property of image inaccuracies. This is that inaccuracies can sometimes reinforce each

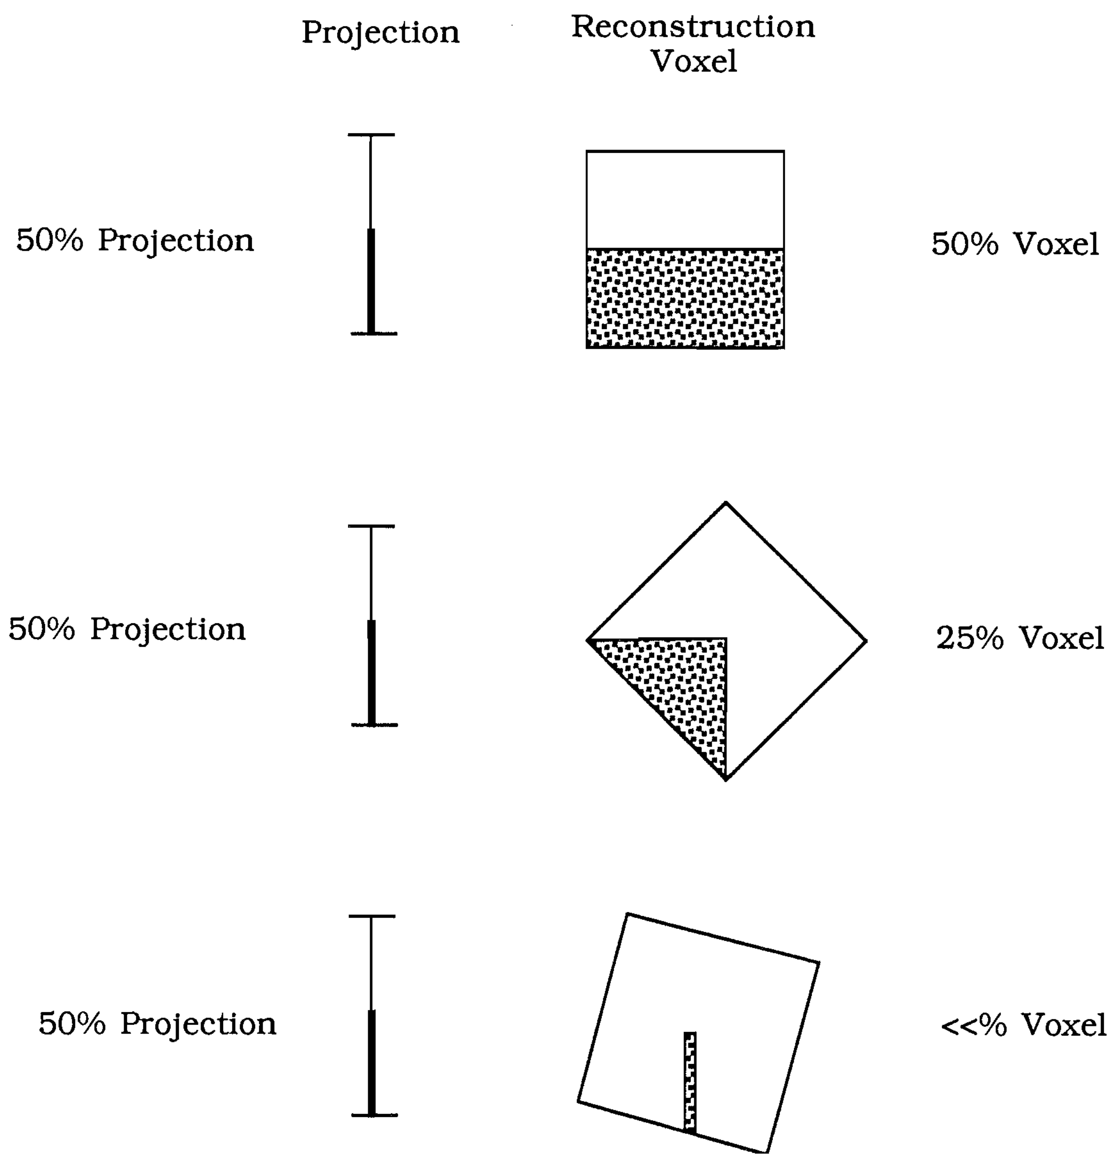

Figure 3.15 50% overlap rule examples for voxel projections

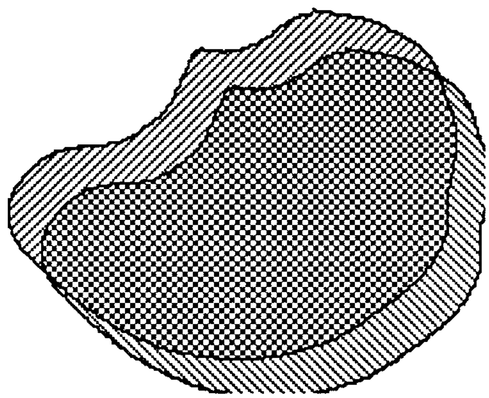

Figure 3.16 Registration conflict during reconstruction

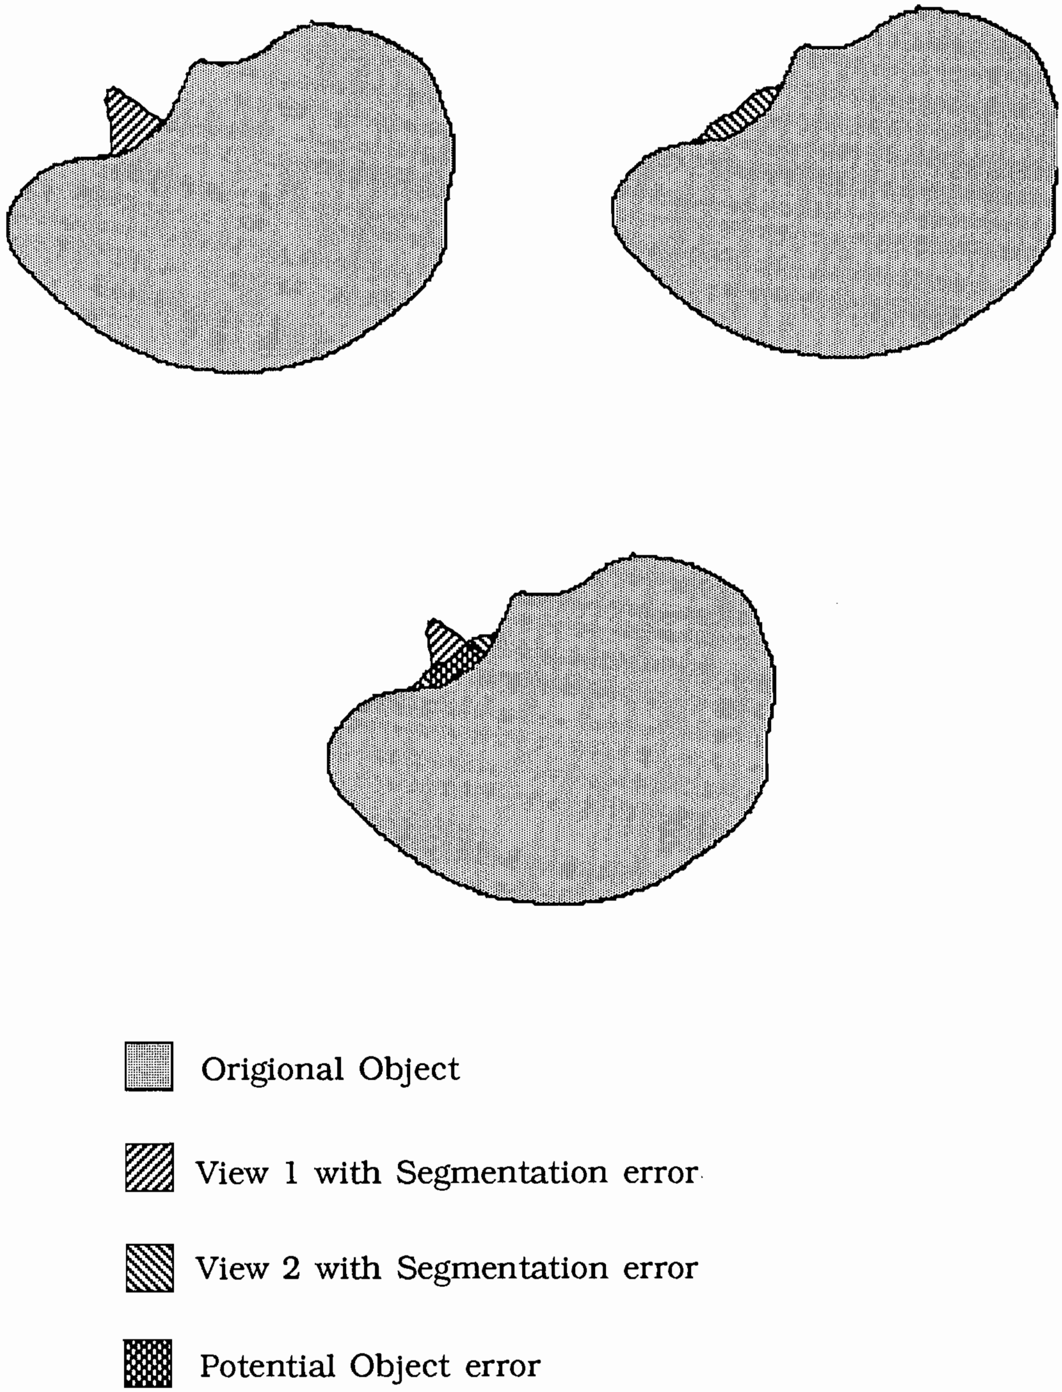

Figure 3.17 Segmentation effects during reconstruction

other such that the wrong conclusions are assumed about the reconstruction. These effects will cause the inspection analysis to incorrectly presume that the object is inaccurate.

Whether the conflict arises from improper calibration or improper segmentation, the end effect during voxel reconstruction is an improper description of the object resulting in inaccurate object measurements. This is clearly a problem since the goal of the reconstruction is measurement. Since these effects originate in the view data itself, the most straightforward way to eliminate the problem is to insure good view data. It is beyond the scope of this presentation to incorporate efforts to correct for lighting effects on the quality of the views.

3.3.6 Accuracy Calculations

The first accuracy test is the successful traversal of the inspection octree. A successful traversal of the tree indicates that the object is sufficiently close to the desired object that it passes a minimal set of criterion. In practice this means that the inspection octree could determine the boundary of the object: either in the correct location or in some area close enough to the correct location to be on the octree inspection map.

Once this has been successful it is possible that no further calculations would be required. This would be the case if the inspection problem was one of accepting or rejecting the part. Further calculations are only required if it is required to know an exact accuracy of the specific object just inspected.

This could be performed by either a second traversal of the tree or during the first traversal. A second traversal could be used if the local accuracies of a rejected part was needed. This could be done more efficiently during the first traversal if it was known a priori that this information would always be needed. This is another example of the flexibility of this general inspection scheme.

The actual error at each landmark can be measured by determining the number of voxels between where the boundary should be located and where the boundary is actually located among the adjacent nodes. For instance, if the boundary is in the correct location then the error number would be zero. If the boundary is not in the correct location the error number is equal to the number of adjacent reconstruction voxels which should be void but are reconstructed to object plus the number of object voxels which are reconstructed to void.

A minimum error could be defined similarly. Examples of reconstruction errors are shown in Figure 3.18. The error number times the size of the reconstruction voxel used in the inspection process gives a size of the volume error in the neighborhood of the landmark.

The overall error can be represented in two ways. The first is a simple sum of the volume errors. This number is an expression of the total amount of error that is present in the object being evaluated. The second error representation is using the maximum volume error.

Often, the errors in specific locations are not the goal of the inspection process. In many cases it is the relative locations of the landmarks which are significant. This might be the area or perimeter of a region, or the smoothness of a surface. In these cases the measures described earlier are not sufficient to express the error.

Since we are interested in the relative locations a second pass through the results of the inspection analysis may be required if the absolute error test is passed. Once again, this would be more efficiently performed if all tests were accomplished during one pass. During this pass the maximum and minimum estimate for the location of the landmark shall be determined. Figure 3.19 shows an example of the reconstruction error where the maximum error is two and the minimum error is zero. It is important to remember that the furthest possible location will be bound by the depth of the type four nodes in the inspection map. Further, if the error overflows the adjacency nodes, it is not possible to give upper bounds on the error estimates for the landmarks and only lower bounds if they are less than the error overflows.

This set of two points for each landmark allows us to bound the parameter calculation based upon the landmarks. For instance, if the landmarks defined a hole in a part and the significant parameter was the

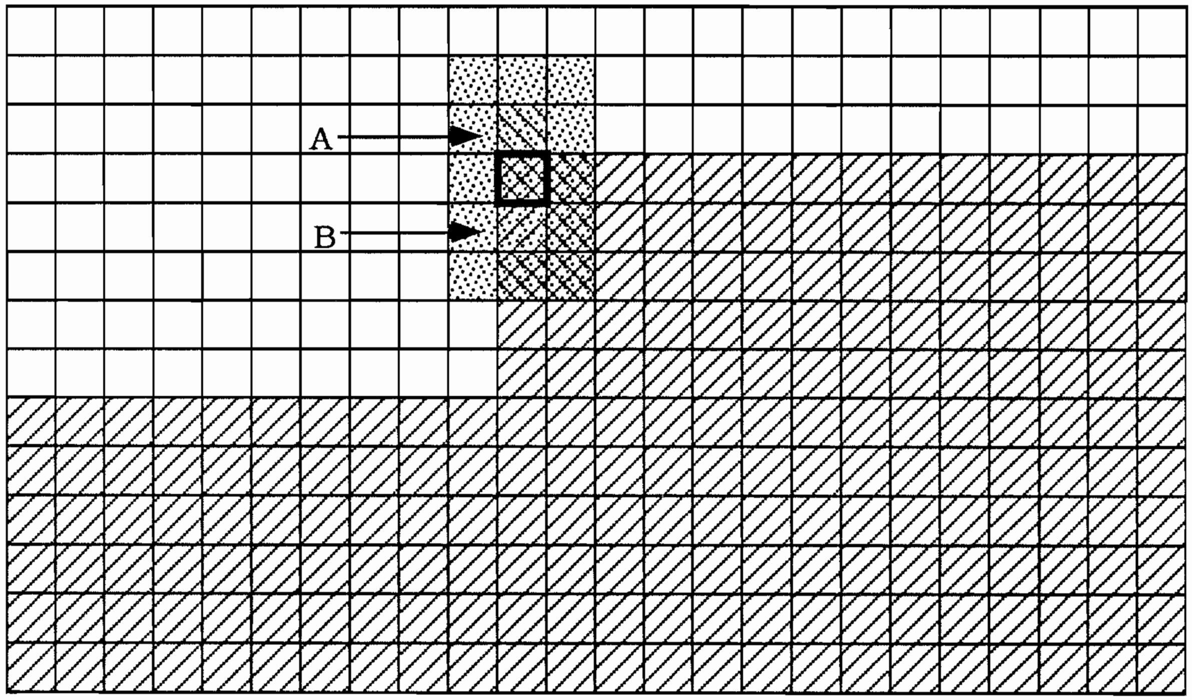

- Area Reconstructed

- Landmark

- Object

- Reconstruction as Object

A - Reconstruction error by addition

B - Reconstruction error by omission

Figure 3.18 Example reconstruction errors

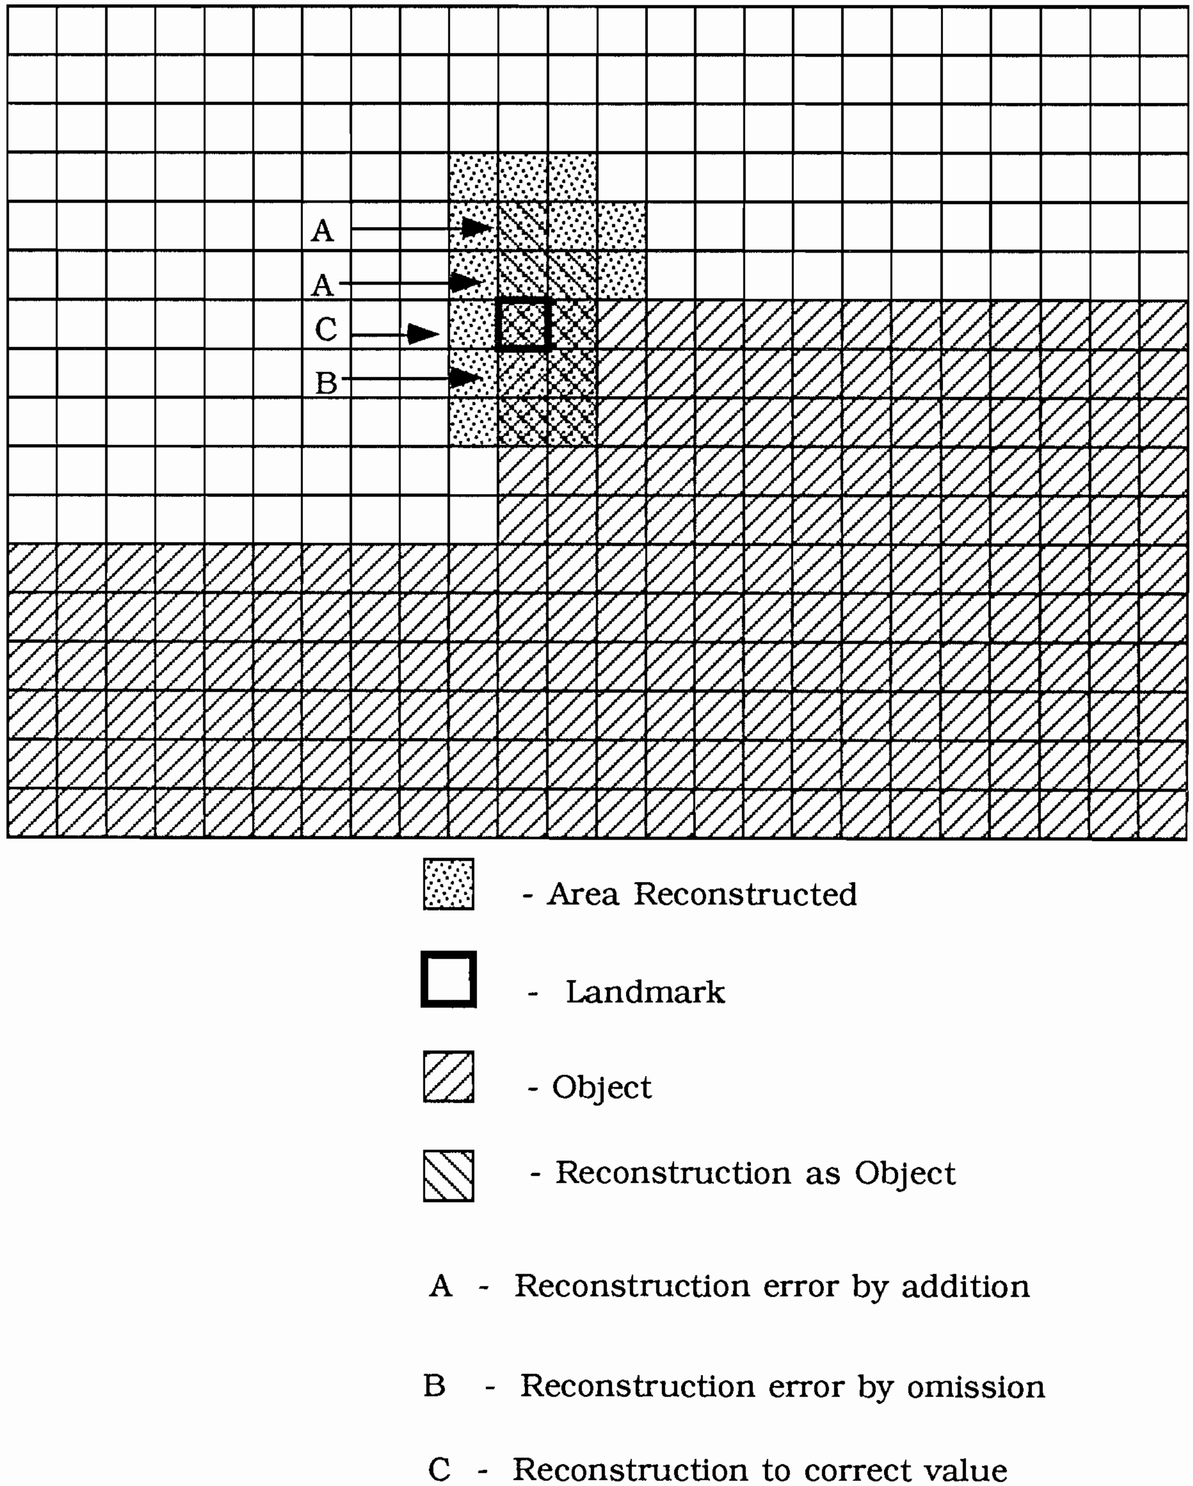

Figure 3.19 Minimum/Maximum error location for landmarks

perimeter, an upper and lower bound on the perimeter can be established for the polygonal fit of the hole.

$$(\text{Perimeter})_{max} \sum_{i=1}^{i=N-1} (\mathit{Landmark}(i+1) - \mathit{Landmark}(i))_{max} \tag{3.3.6-1}$$

A minimum bound could be defined in a similar manner. Likewise, an upper and lower bound on the area of the polygonal fit of the hole could be calculated. One method of calculating this area would be to reconstruct the hole by plotting lines through the landmark estimates, filling in the hole, and counting the pixels present.

Note that the error due to polygonal fit is due to the landmark representation of the hole for inspection analysis and not due to the inspection analysis itself. If the amount of error due to the polygonal fit is unacceptable for the inspection analysis it is only necessary to add additional landmarks along the path to be represented. This will yield a more accurate description of the computed parameter.

Similarly, the smoothness of a surface can be evaluated from the set of minimum and maximum estimates for the landmarks. This could be estimated from the standard deviation of the distance from the landmark estimates to the ideal landmark location. Bounds on other parameters could be defined similarly.

3.3.7 Resolution and Accuracy Decisions

One aspect of the octree inspection representation is that we can trade off accuracy and resolution. At higher levels of the octree there is less resolution of the object boundary because the voxels are larger. This makes it more likely that the reconstruction of these voxels will be evaluated correctly, Small errors in the object would be averaged out over these larger voxels. This analysis would also be potentially faster because of the 50% overlap rule in voxel determination completing overall evaluation of a single large voxel faster when compared to a lot of smaller voxels.

The penalty for this is that the possible accuracy measure is reduced due to the large voxel size. The boundary location on the object which is represented by our landmark is a specific location. The extent to which we can discriminate between different locations on the object is limited by the size of our reconstruction voxels. Tradeoffs between resolution accuracy and processing speed are explored more in Chapter 4.

3.3.8 A Typical Inspection

This section shall describe a typical inspection process from conception to error analysis. A flowchart of this process is given in Figure 3.20. The first step in an inspection process is the confirmation that the part can be inspected by the algorithm. This requires that the landmarks be selected which describe the significant locations on the object. This selection is most efficiently done during the part design phase by those who designed the use of the part in question. Once the landmarks are selected it is necessary to determine if the object is *Weakly Convex* at each landmark.

Chapter 2 states that if an object is *Weakly Convex* at a boundary location on the surface of the object it is possible to reconstruct that location using backprojection reconstruction. Since this reconstruction is the basis for our inspection it is also the requirement for us to be able to perform the object inspection.

Once all the landmarks have been determined to be *Weakly Convex*, it is then necessary to construct the octree inspection map of the object. First, an octree object representation is found for the object. The depth of this octree must be sufficient for each landmark to be *Weakly Convex* with a point size equal to the smallest voxel size.

After the representation has been found, type two nodes are used to pare off those sections of the octree which don't lead to landmarks. This representation is then further modified by substitution of type three nodes for each landmark in the representation. Type four nodes are then added onto the type three nodes to a depth determined by the accuracy requirements of the inspection.

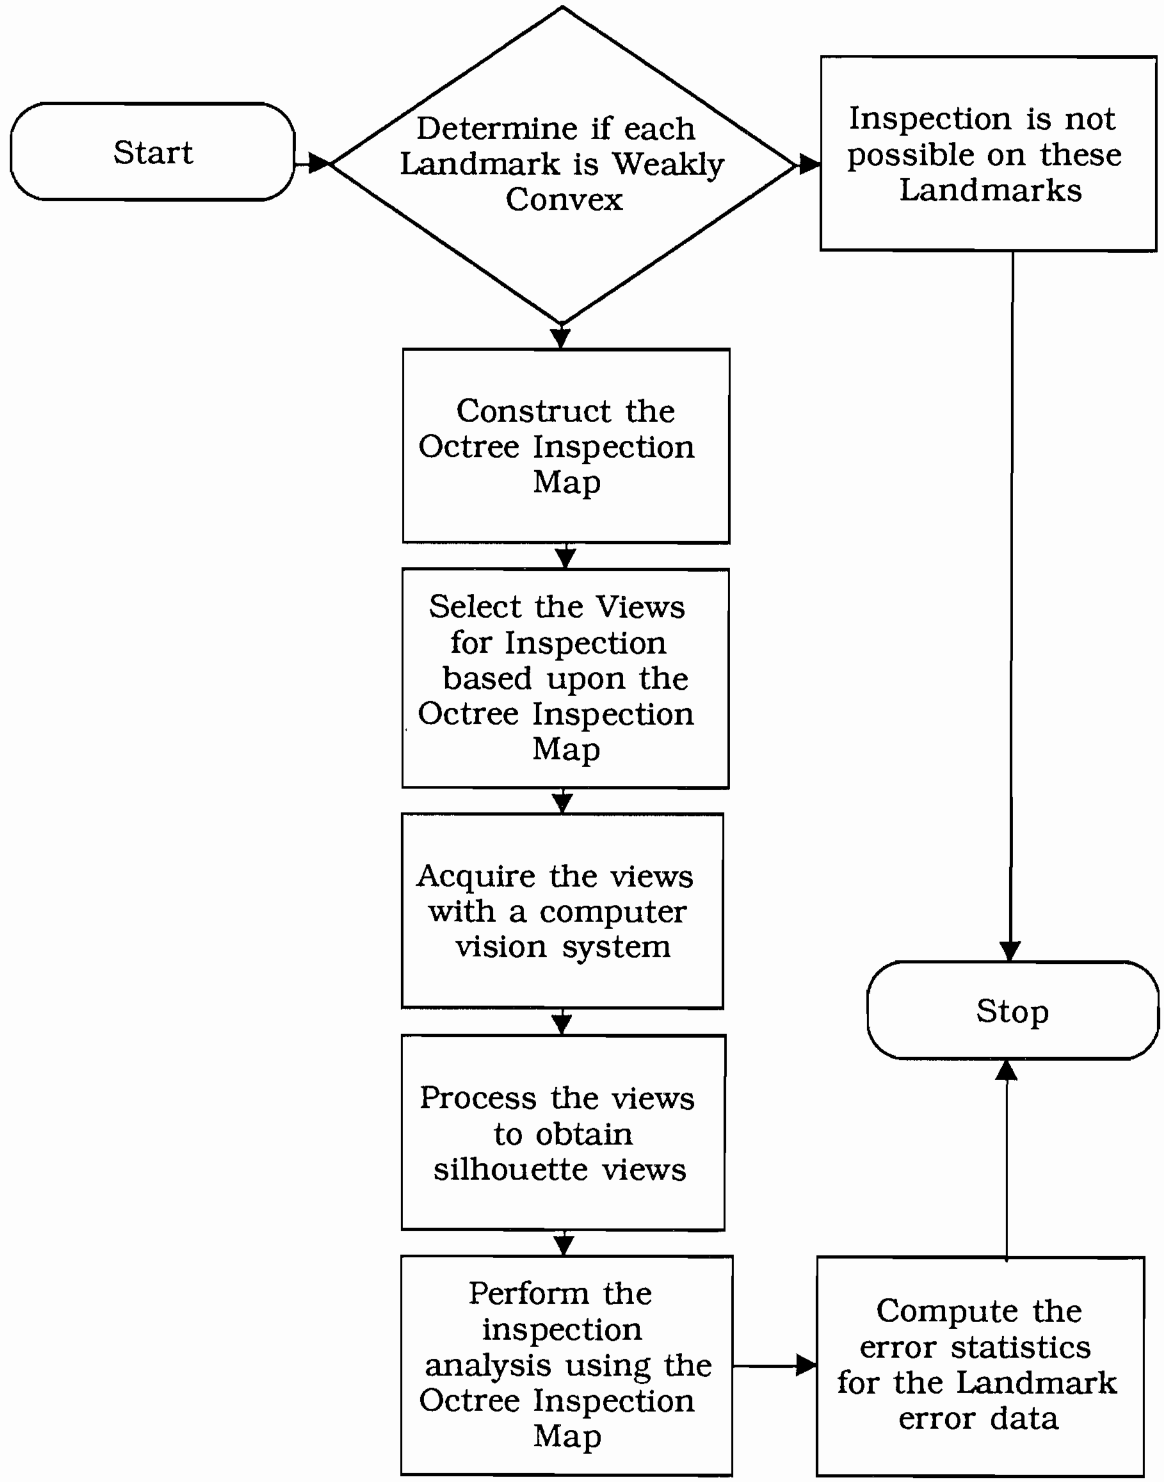

Figure 3.20 Procedure for a typical inspection

At this point it is necessary to evaluate all type four nodes which evaluate to void during reconstruction for their visibility during reconstruction. Any type four nodes which are not visible must be pared from the octree because it is not possible to evaluate an error dependent upon their reconstruction. Recall that *Weak Convexity* guarantees that there is at least one adjacent vertex to each landmark which will survive this paring.

The octree inspection map is now complete.

The next step is the selection of the views for the inspection analysis. This is dependent upon the *Weak Convexity* of the landmarks and the visibility of the type four nodes. Suitable views must be selected which allow for all the type three and four nodes in the inspection map to be visible for reconstruction. Chapter 5 explores this in more detail.

Once the views have been selected, they be acquired through the use of some form of computer vision system. This is the first step of the "on-line" inspection process. The computer vision system must be capable of both capturing an image of sufficient quality that the object silhouette can be extracted as well as identifying the location of the view with respect to the reconstruction volume and object map. Recall that we have assumed that the reconstruction volume and the object inspection map are registered.

The views now need to be processed so that the (possibly) grey scale views of the image are segmented and thresholded to yield silhouettes of the objects from the directions required by the view selection. For this, known segmentation and thresholding techniques are used. It should be noted that this is not always a straightforward process but known methods can yield useful results.

Now that we have both an inspection map of the ideal object and a set of silhouettes which describe the object to be inspected, the inspection analysis can be performed. This is done by a single traversal through the type two, three nodes and some type four nodes of the octree inspection map.

During the traversal, each type three and four node which is traversed is reconstructed. When a type three node is crossed during a traversal the inspection algorithm has detected another location whose accuracy must be determined. This accuracy is determined in a two step process. The first step is voxel reconstruction.

This is done by first reconstructing the type three node, the landmark. If this voxel is correctly identified as object, the adjacent type four nodes are reconstructed to see if the landmark is in the correct location. This is determined by the reconstruction evaluation of each voxel matching the value from the inspection map. Any type four node which is not evaluated correctly starts a recursive evaluation of its adjacent type four nodes. This continues until a set of adjacent nodes are evaluated correctly or a depth limit is reached by the recursion.

If the landmark voxel is not reconstructed to object, the inspection process has yet to find the object. The first step is to try to find the object near where it should have been located. All the adjacent type four nodes are then evaluated to determine if any of them evaluate to object. All voxels which evaluate to object are then treated as possible locations for the reconstructed landmark. Thus, the adjacent type four nodes of each of these locations is evaluated to see if their reconstruction match the inspection map. Any type four node which is not evaluated correctly has its adjacent type four nodes evaluated by reconstruction.

In either case, this recursive process is continued out to a maximum depth of type four nodes. It is necessary to maintain a check against this depth rather than examining for leaves without children because the depth limit used in the construction of the octree inspection map may have linked landmarks together in a manner which in an extreme case could link all landmarks. An example of this would be a set of landmarks chosen such that the entire object will be reconstructed.

Note that any time a type four node needs to be reconstructed, a check is made to determine if the node has already been reconstructed from an earlier part of the traversal. This limits to some extent the number of reconstructions which need to be performed as the selection of the landmarks on the object become more connected. This could happen if landmarks are in adjacent voxels or if a path of type four nodes exist between two landmarks.

As noted, the actual process of reconstructing the type four nodes is done in a recursive manner. Reconstruction of a type four node which evaluates to the correct value in the inspection map terminates the recursion which is leading away from the landmark in an effort to find a match with the inspection map. Any time a type four node is found which does not evaluate correctly to the value in the inspection map, a new layer of recursion is generated as the adjacent type four nodes to this location are reconstructed. This recursion continued our to a predetermined depth before the search is given up.

During this recursive evaluation of the type four nodes there are two measurements to be maintained. The first is the maintaining of the maximum error number of the landmark and the second is the number of incorrectly reconstructed voxels near each landmark. This latter provides a measure of local error in the neighborhood of the landmark. The result of these two measures are stored in the type three landmark node for use in later parameter evaluation.

A constant check is made of these two measurements against error bounds for the part. In some cases, we are not interested in the amount of error in a part, only in the fact that the error exceeds some bound. There are two instances when this bound could be exceeded. The first is if there is a gross error in some aspect of the part. This could be significant "extra" or "missing" object or a "bridge" between two void nodes. The second is if there is an alignment problem in some part of the manufacturing process which might result in correct landmarks in the wrong location. The inspection routine can immediately exit at a bad part evaluation in these instances.

Thus, each type three node is evaluated through the evaluation of type four nodes. A successful traversal of the tree would indicate that there were no catastrophic errors which would cause the part to be rejected. At this time the total and maximum error values are known for each landmark. Secondary traversals could now be performed to yield additional parameter values if needed. These could be the area or perimeter parameters noted earlier.

3.4 Flexibility of This Inspection Algorithm

One of the most appealing aspects of this approach to inspection analysis is its ability to be set up a priori for a very rigorous inspection and then be adapted at execution to a more appropriate level of analysis. This a priori setup of the inspection is essentially a very detailed octree inspection map. There are many different ways in which a such an inspection map could be adapted to a specific inspection problem.

One way would be to prune some of the landmarks by masking some branches of the octree. The resulting octree could then be used to examine the remaining portions of the tree.

Another adaptation of the inspection map would be to limit the depth of the tree search. This would have the effect of setting an accuracy granularity in the resulting error calculations. The best accuracy granularity achievable would be that obtainable when using the entire octree. Limits and tradeoffs in this area are examined in Chapter 4.

Another adaptation of the inspection map would be to limit the depth of the adjacent nodes which are searched. Each additional layer of adjacent nodes which are searched allows the algorithm to measure a larger error at each landmark. Limiting the depth of this search, limits the amount of error which can be measured at any location. Depth overflows at any given location could be used to reject a part as having too much error.

Further, additional inspection analysis of parameter values could be performed conditionally or during all executions of the inspection. The strength of each of these adaptations is that they can be done at the execution time of the inspection. Data processing and storage could be reduced by each adaptation. This can be done by skipping data processing which is not called for and freeing those portions of memory which are not filled by data processing during the analysis.

3.5 Implementation

There are two main aspects to the inspection analysis. These are the preprocessing before an inspection and the inspection itself. Each is comprised of two steps. The preprocessing is comprised of the generation of the octree inspection map and the selection of the views of the object. The actual inspection is comprised of the octree inspection map traversal and the subsequent error calculations.

3.5.1 Octree Inspection Structure Generation

The octree inspection map will be generated by a three part procedure. Figure 3.21 flowcharts this procedure. The first part is a complete octree representation of the ideal object. The second part is the labeling of landmarks within the octree structure and a pruning of non-landmarks. The third part is to add the type four nodes to the type three nodes.

The complete octree representation may be generated using a know technique, from CAD information about the object. This information could be preserved in case of future modification of the inspection routine at the expense of additional unnecessary tree search during inspection. Note that no additional reconstructions would be generated; just additional tree traversing through unnecessary nodes. The octree representation should be generated to a sufficient depth such that the granularity in the representation due to voxel size is below that required by the inspection analysis and such that *Weak Convexity* is obtained in the voxel representation for all landmarks.

The labeling of landmarks for inspection is done by substituting a type three node into the octree structure. This is the beginning of the octree modifications which turn the octree representation into an inspection map.

Finally, type four nodes are attached to the type three nodes to indicate areas adjacent to the landmark. Each of these nodes is labeled as to whether it should be object or background by a comparison to the original octree representation. Additional layers of type four nodes may

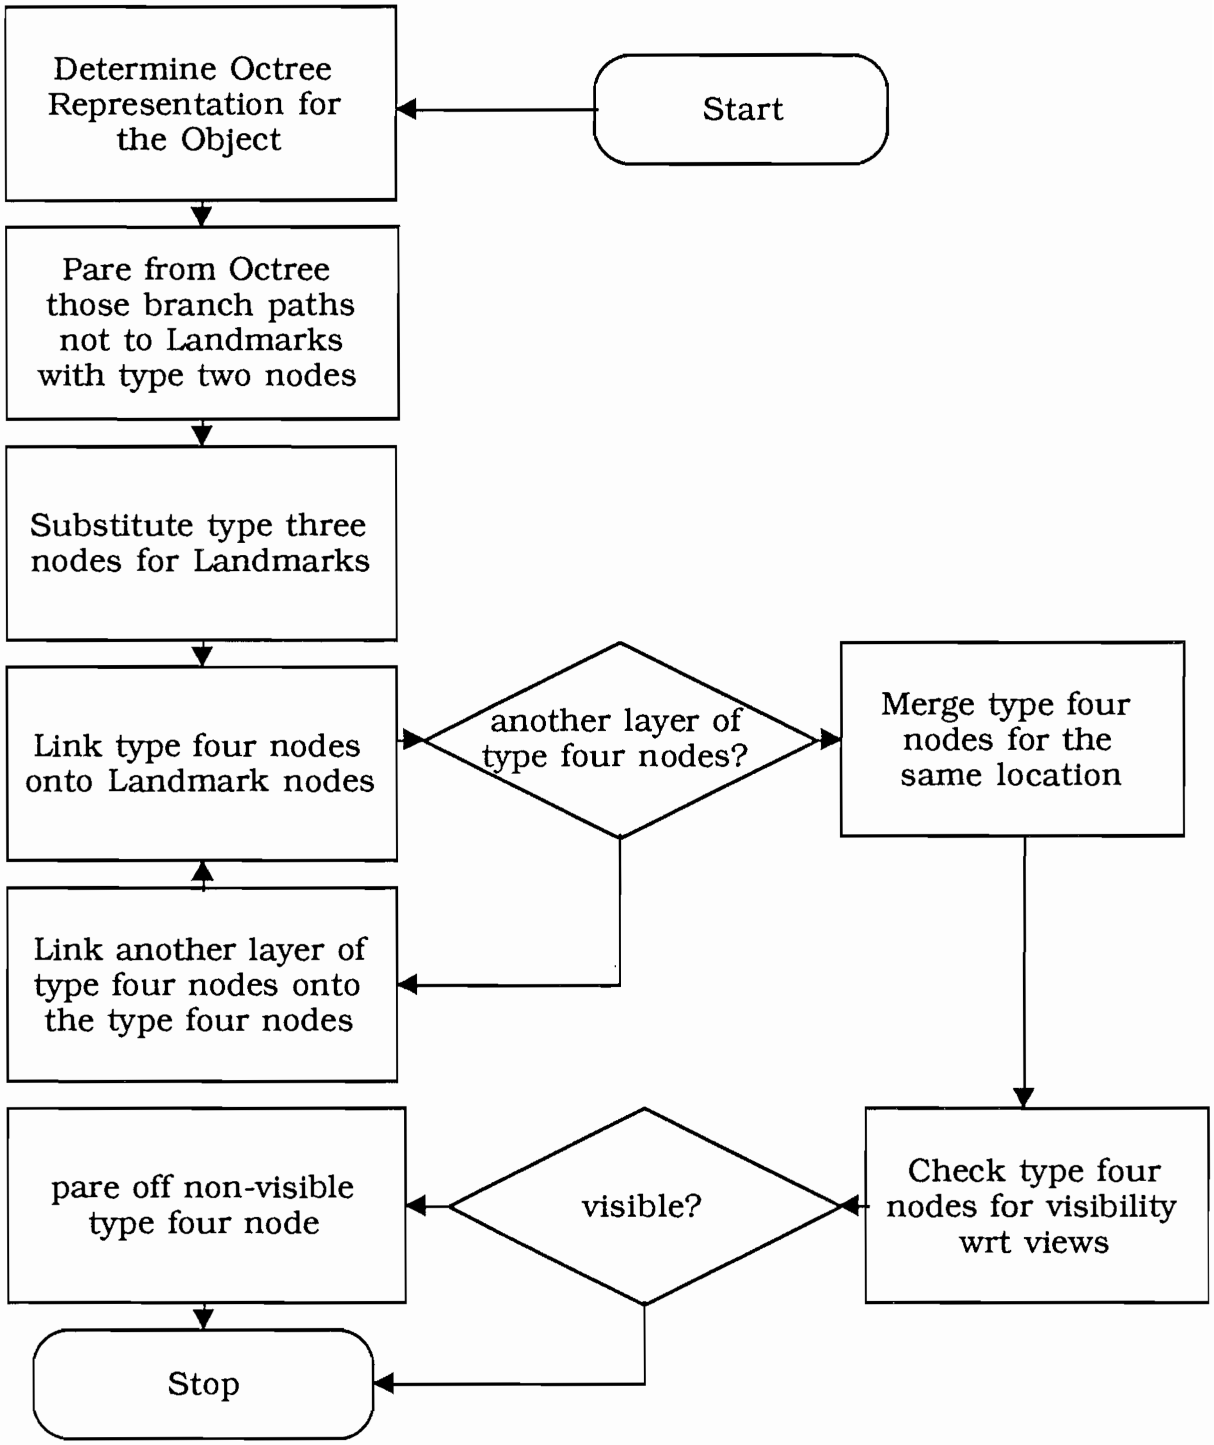

Figure 3.21 Procedure for generation of an octree inspection map

be attached to these type four nodes. This is continued until a sufficient depth of nodes is obtained which exceeds the maximum amount of error in the object which we wish to measure. Any error exceeding this depth will only be noted as exceeding the maximum measurable error.

3.5.2 View Selection

Chapter 5 will present an algorithm for the selection of an efficient set of views. In this chapter the view angles and locations will be selected by hand in a manner which will insure that the significant areas can be reconstructed. This selection manner is sufficient for the verification of the inspection algorithm itself.

A limited set of views will be used since the number of landmarks which will be reconstructed will be limited. However, the number of views chosen per significant area will be such that the landmark will be overdescribed. The algorithm presented in Chapter 5 will reduce the number of views through the 50% rule for voxel identification.

3.5.3 View Registration

An obvious problem which must be solved with this inspection analysis technique is that the views obtained during inspection must be properly registered with the octree inspection map. During the testing of the inspection algorithm, this shall be insured by manual registration.

Hand registration is obviously not possible in an automated inspection system. In all instances however, any part which is being subjected to inspection has at some time been registered with the apparatus which manufactured the part or with an apparatus which is going to use the part. Ideally, any inspection would take place either just after manufacture or just prior to use. Therefore, while the registration is a problem for our testing, it is not a problem for an actual inspection.

3.5.4 Physical Aspects of Inspection Accuracy

In our object inspection analysis, the goal is to merge data with the goal of determining the location of certain landmarks. Several sets of data are used to reconstruct this information. It is necessary for the accuracy of the view data to be propagated through this merging.

This means that the accuracy of the data taken is heavily dependent upon the physical relationships of the camera parameters, camera view location and angle, pixel shape, size and number, and camera manipulator. Chapter 4 will examine the interaction of some of these specific factors.

As might be expected, these factors are interrelated. The inspection algorithm's limiting factor has been described as the smallest resolution voxel of the inspection map. The actual minimum size of this voxel is limited by the factors mentioned above. This limits the accuracy which may be obtained by the inspection.

Other physical limitations are more general in nature. These include the table which the object is placed on for observation, the available lighting, and the manipulation range of the camera. Figure 3.22 shows some of the physical limitations on a general reconstruction environment. The object to be reconstructed must exist on some table or be held in some fixture for viewing. Limited camera angles may be available. Lighting may or may not be good or uniform.

This has the effect of limiting the number of view positions which will give good information. In essence, there are a limited number of object camera angles available for reconstruction. This limits the quality of any inspection by reducing the number of view locations. As was shown in Chapter 2, this has the effect of reducing the number or type of objects which can be reconstructed.

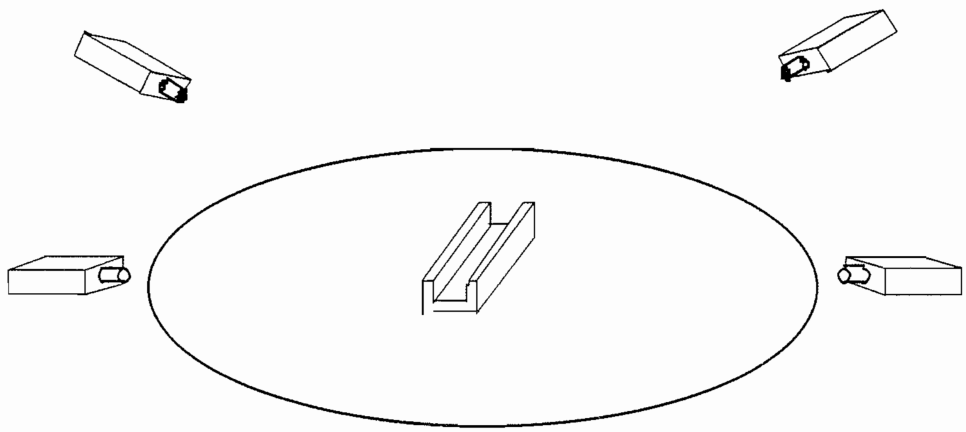

Figure 3.22 Physical limitations on backprojection reconstruction

CHAPTER 4

CALIBRATION, SYSTEM PARAMETER EFFECTS, AND ACCURACY ANALYSIS IN BACKPROJECTION INSPECTION

4.1 Introduction to Accuracy

Accuracy considerations in industrial inspection involve the determination of just how accurately an inspection process can determine object parameters well as the determination of how accurately a part is manufactured. This determination can be expressed in two ways. The first of these is a minimum a priori discrimination (MAD) specification -- the smallest location accuracy that any inspection can measure as restricted by a priori knowledge of the ideal object. This is limited by the ideal object representation. The second is a minimum experimental discrimination (MED) specification. This is an expression of the minimum discrimination of the inspection process. This is as a function of the vision system parameters in the inspection process.

The accuracy analysis presented here is restricted to analysis of the limited view, goal driven, volume source inspection algorithm of Chapter 3. Calibration of the inspection computer vision system will form the introduction to this analysis. Section 4.2 presents aspects of the calibration problem for backprojection inspection. Section 4.3 presents an algorithm for the calibration of a backprojection inspection system. This calibration procedure is then applied to an experimental setup to test the effectiveness of this procedure.

Several different elements interact to determine the actual bound on inspection accuracy for a given inspection system. Section 4.4 presents these different elements of inspection and their relationship to accuracy. Section 4.5 presents an expression for the accuracy bound on a given inspection representation. The MAD specification presented will follow from characteristics of the octree inspection map which is used to represent the object. Section 4.5 will also present a set of expressions which together express the combined accuracy bound for a computer vision inspection system. The MED specification will integrate the camera parameters and the view locations into this bound.

A general inspection error measure is presented in Sections 4.6. Here, the results of accuracy inspection will be merged into a more meaningful implication for measuring accuracy for inspection. Implications for the general inspection problem from accuracy analysis are presented in Section 4.7.

First, we shall review two topics related to the inspection analysis of Chapter 3: landmarks and backprojection inspection.

Landmarks

An object landmark is a significant location on the three dimensional object to be inspected. This concept was first introduced in Chapter 3. The location is selected as important to the overall characteristic shape or usefulness of the object. In Figure 4.1 an object is shown with several landmarks indicated. These particular features might be important if the object were a stamp whose end point dimensions were significant. In some instances, an object may only have a few landmarks. For a complex object, it would be possible to have a large number of landmarks. In the limit, the number of landmarks could cover the entire surface of an object.

One method of locating landmarks is to determine them during CAD design. During design it would be straightforward and logical for a designer to specify which points were significant for an accuracy analysis. In Chapter 5 the concept of landmarks will be used in the development of view planning. The goal of view planning is to select the views which give the best reconstruction as measured by the accuracy criterion presented in this chapter.

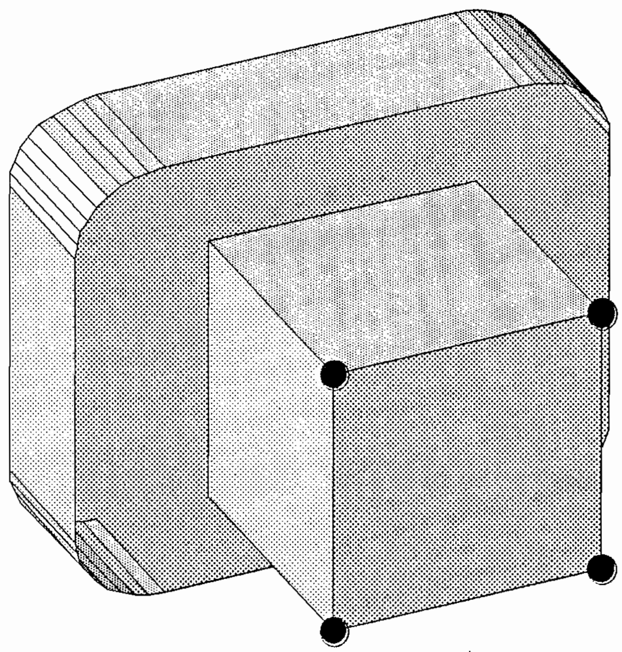

Figure 4.1 Example object with significant for a stamp application

Backprojection Inspection

Backprojection inspection is the process of determining accurate manufacture of an object through silhouette backprojections of the object to be tested. A modified octree is used to represent an inspection map of the ideal object. The particular algorithm presented in Chapter 3 uses a volume source method to accomplish the reconstruction of the landmarks which are significant. These locations are represented in the octree as voxels -- volume elements of the reconstruction space. This algorithm projects elements of the reconstruction space which are significant to the object onto the views to perform the reconstruction. Typical backprojection inspection projections are shown in Figure 4.2. Inspection is performed by a reconstruction of the areas in the neighborhood of the landmark. The pattern of this reconstruction is such that only those areas which need to be inspected are reconstructed.

The use of landmarks serves to restrict the complexity of the inspection analysis by limiting the scope of the inspection reconstruction. Inspection accuracy analysis is only valid near these landmarks. This inspection process only estimates a landmarks location in the reconstruction space near where it should be located.

4.2 Calibration and Registration for Backprojection Inspection

An essential part of backprojection inspection is the backprojection reconstruction of individual voxels near landmarks. One unique feature of backprojection reconstruction is its relatively simple method of calibration. As shall be shown the complicated part of the calibration scheme is the analysis of the camera for its characteristics. These are invariant under the analysis of different objects and need be done only once. We shall assume a pin hole camera model for this analysis.

Our inspection system is composed of two parts; the camera and the object. The calibration process for this system is twofold. First, we must determine the characteristics of the camera which are not available by direct measurement. Second, we will determine the camera/object relationship. Figure 4.3 shows the camera/object system parameters.

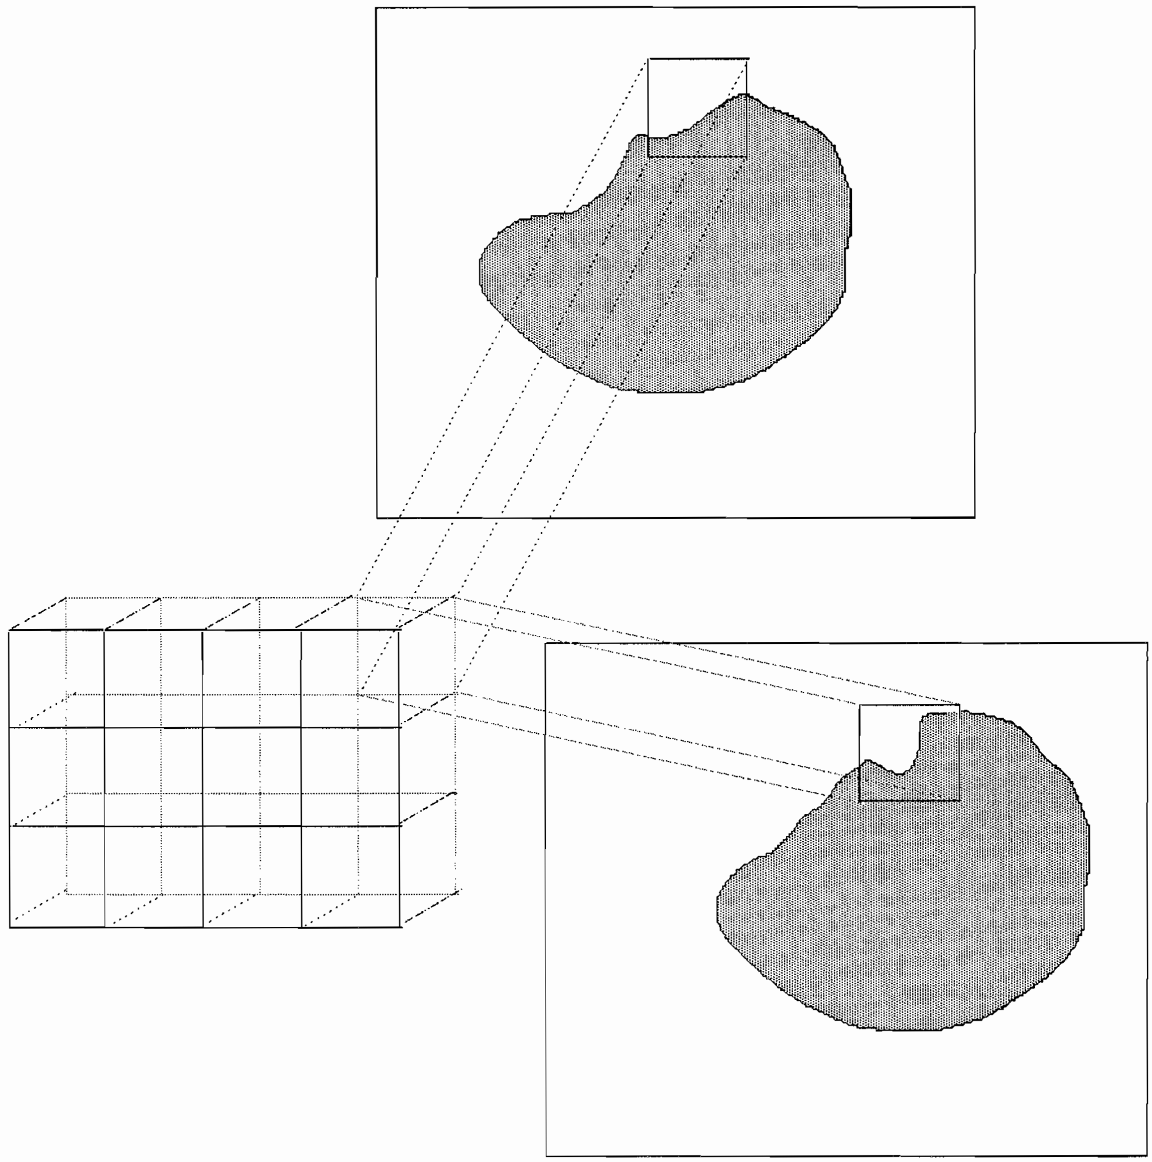

Figure 4.2 Typical backprojection inspection projections from the voxel space onto the view space

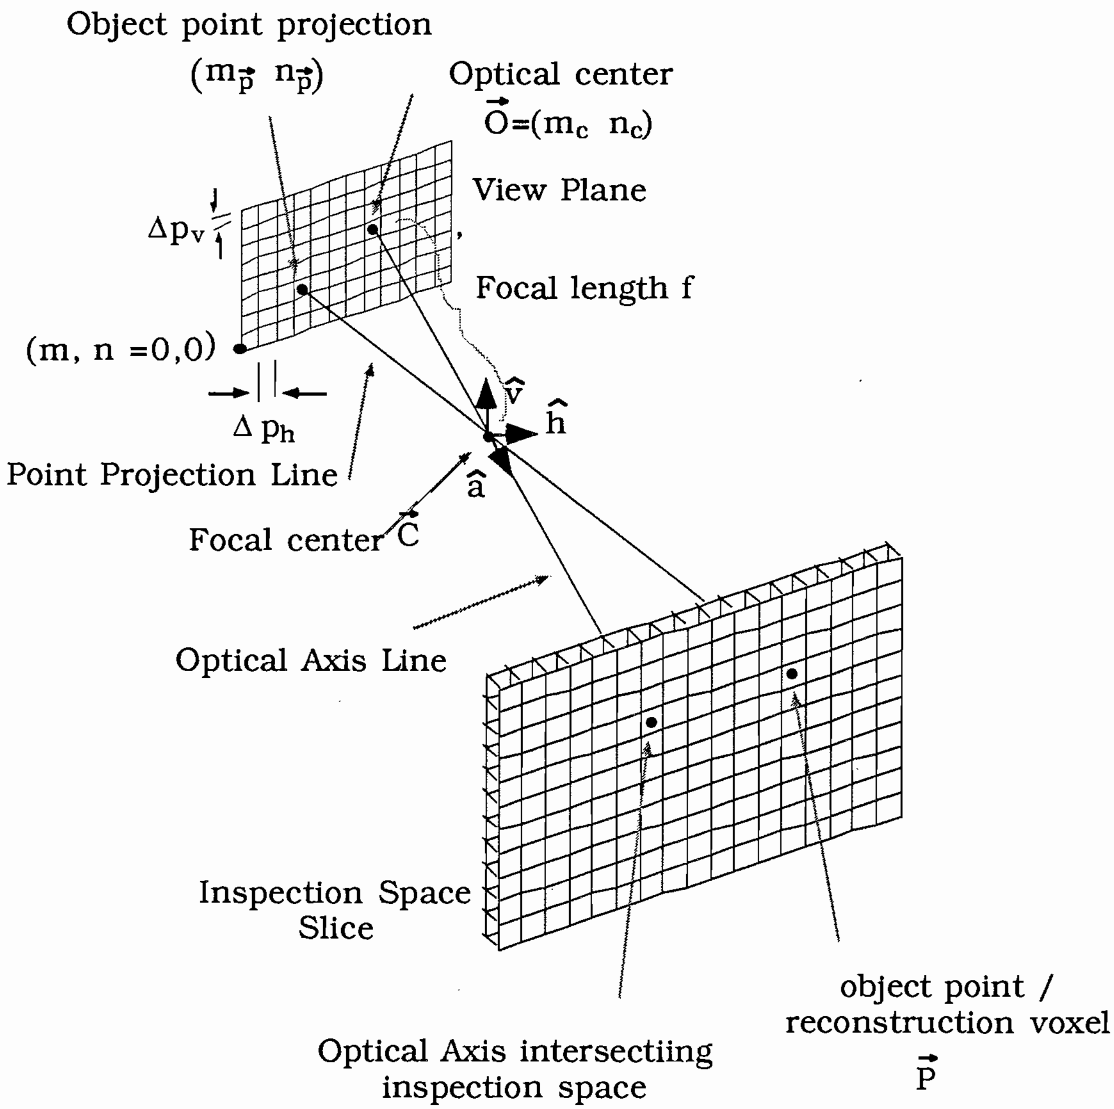

Figure 4.3 Camera/object system with variables shown

From these parameters a set of transformation equations are derived which relate the object data to the view data. The registration problem between the ideal object and our measurements is addressed. Finally, other techniques will be compared to this method.

4.2.1 The Calibration Problem

Camera/object calibration is used to describe the process of orienting the object so that its position relative to a set of calibrated views is known. There are two ways of accomplishing this task. The first method could involve the use of a robot mounted camera which moves a camera around an object on a table to view it from a number of angles. This method would be typical of an inspection cell. The second method, which we examine in our analysis, is to fix the camera and move the object by placing it in a fixture which provides for rotation and declination.

The camera parameters are not available by direct measurement. They are the optical center, the focal length, and the pixel size. These need to be determined during calibration. The optical center is a parameter of the camera optics. Its location in a view is a function of the digitizer's alignment with the camera's sensor and optics. The focal length is also a function of the camera optics. This parameter can not be directly measured because the location of the camera sensor is not available. Finally, the pixel size is a parameter of the camera scanning pattern and the digitizer. Each of these parameters affect the transformation from the test object to the view.

The other half of the calibration problem is determining, and maintaining during analysis, the relationship between the object and the camera. This includes knowledge of where the optical center passes through the object and the view distance and orientation with respect to the inspection space. Together, this is referred to as the view location of the view. This information is needed to be able to determine the orientation of the silhouette to the reconstruction space.

Knowledge of the camera parameters and the camera/object orientation give us the ability to define a set of transformation equations between the view and the inspection space. Figure 4.3 labels the parameters of the transformation on the view and in the reconstruction space. The parameter definitions follow those of Tan [123]. From the calibration and these definitions a set of transformation equations can be determined.

$$m_{\vec{p}} = \left\lfloor \frac{f}{\Delta_{p_h}} \frac{(\vec{P}-\vec{C}) \bullet \hat{h}}{(\vec{P}-\vec{C}) \bullet \hat{a}} \right\rfloor + m_c \tag{4.2.1-1}$$

$$n_{\vec{p}} = \left\lfloor \frac{f}{\Delta_{p_v}} \frac{(\vec{P}-\vec{C}) \bullet \hat{v}}{(\vec{P}-\vec{C}) \bullet \hat{a}} \right\rfloor + n_c \tag{4.2.1-2}$$

These equations describe a transformation from the point on the object onto a point in the view plane

4.2.2 The Registration Problem

The calibration problem discussed thus far describes a determination of the relationship between the views and the object to be measured before reconstruction. The related problem concerns the registration of the unknown object data and the known object inspection map after reconstruction. This completes the loop by giving us the ability to compare the data taken during inspection analysis with the know object data. The use of landmark driven inspection also assumes this registration.

Registration requires that the unknown object must have a known orientation. In general, this is difficult. Any landmarks on the object which could be used for registration might be corrupted and unsuitable for analysis. In practice, it is reasonable to assume that registration can be obtained and maintained by the manufacturing cell.

In order for an object to be manufactured, it is necessary for a machine tool to know the location of the part or raw material for manufacturing. The machine tool will then make modifications to the part or raw material which will require registration between the tool and material. The inspection process can make use of this same registration. It can either regain it in the same manner it was initially obtained or the inspection can be performed immediately after modification by the machine tool before registration is lost.

4.2.3 Calibration of Other Methods

All methods of reconstruction require some form of calibration. Some effort has gone into the examination of techniques used for this purpose. One form of automatic calibration is the comparison of lines and points in the images for correspondence between points and lines [77, 116]. Others have limited themselves to two perspective views for the analysis [137]. Other techniques for reconstruction, namely stereo, require a calibration phase for each reconstruction. This technique requires calibrating the two images using a correspondence map to determine calibration of the images [88].

4.3 The Calibration Procedure

The calibration techniques we use follow those of Tan [123]. Others have made use of similar calibration configurations [48]. In each instance, an effort is made to determine unknown camera parameters and to establish a unified reference frame for the camera and object.

In our experimental setup the camera has the ability to move up and down, and side to side in our mount. This gives us the ability to orient the camera such that the camera is aligned with the center of the object. The fixture for calibration analysis is shown in Figure 4.4.

The first step in the camera calibration procedure is to locate the optical center of the lens/digitizer system. This is the point in the image which remains fixed as the distance between the lens and the object change. The geometry of this procedure is shown in Figure 4.5. This is a side view of the experimental setup. When a point lies off the optical axis, it will move as the range between the camera and point changes. If a point lies on the optical axis, the point will not move.

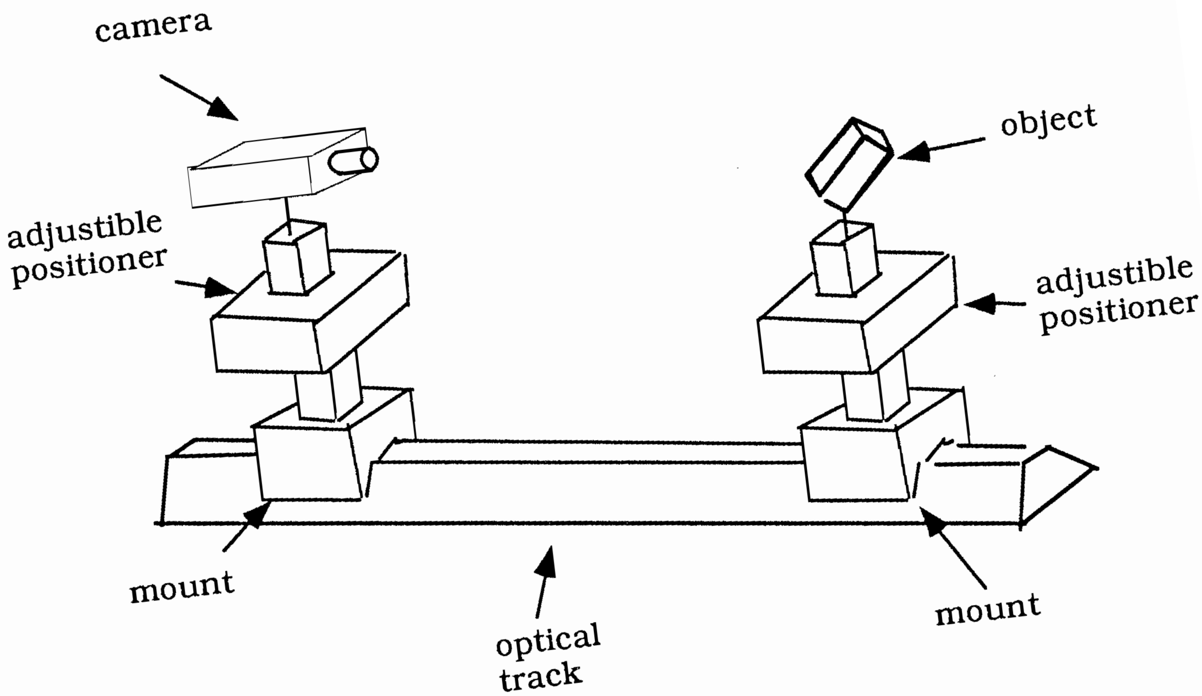

Figure 4.4 Experimental setup for calibration analysis

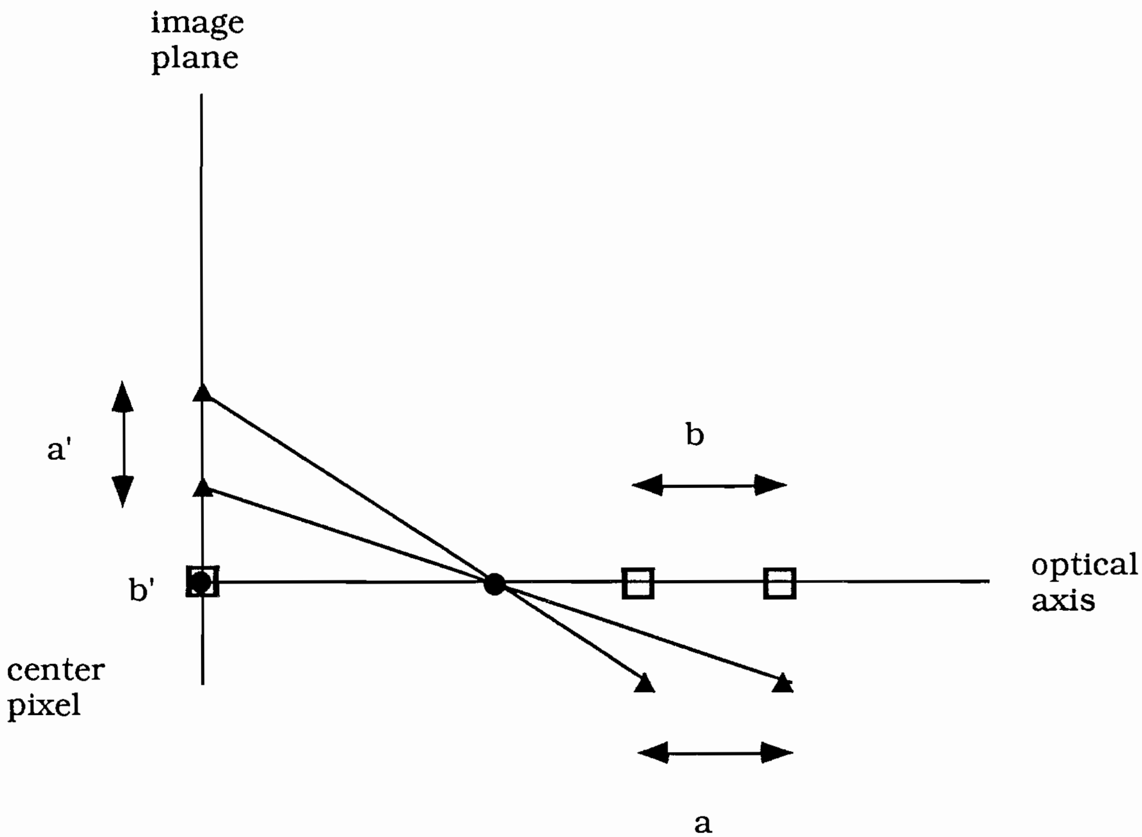

Figure 4.5 Symbolic view of the calibration setup for the determination of the optical center of the sensor

The second step in the camera calibration procedure is the determination of the focal length and pixel size. As can be observed from the transformation equations in 4.2.1-1,2, we only need the ratio of these two numbers. This ratio may be determined from the geometry of the imaging system without knowing either number specifically. Fortunately, only this ratio is required by the backprojection reconstruction inspection procedure.

The geometry of this procedure is shown in Figure 4.6. This ratio of focal length/pixel size can be determined from similar triangles. A number of points off the optical axis in both the horizontal and vertical directions are used at a number of displacements. This creates a large sample set of data for analysis to determine the system parameters.

4.3.1 Calibration Analysis for Chapter 3 Inspections

Recall Figure 4.4 showing the calibration measurement setup. A camera and object fixture are mounted on rail. The mounting for both the camera and object are such that they can be moved horizontally and vertically with respect to each other. Movement of the mounts along the rail provide for depth adjustment.

Careful adjustment of the camera and object mounts are required during both of these operations. Small adjustments of the point being imaged can result in a large changes in the point on the image. During the first step of the camera calibration procedure, we found the optical axis to be located at (264,273). The location of the optical center is in pixels relative to the lower left hand corner of our display system. This test was performed through movement of the camera along the rail.

The second step in the calibration procedure involved the movement of the camera along the rail as well as the movement of the object in the movable mount. The projection of a point onto the sensor array is then recorded for a variety of relative movements between the object and camera. Table 4.1 shows the results of this data gathering for horizontal displacements of the optical axis. Table 4.2 shows the results of this data gathering for vertical displacements of the optical axis. Note

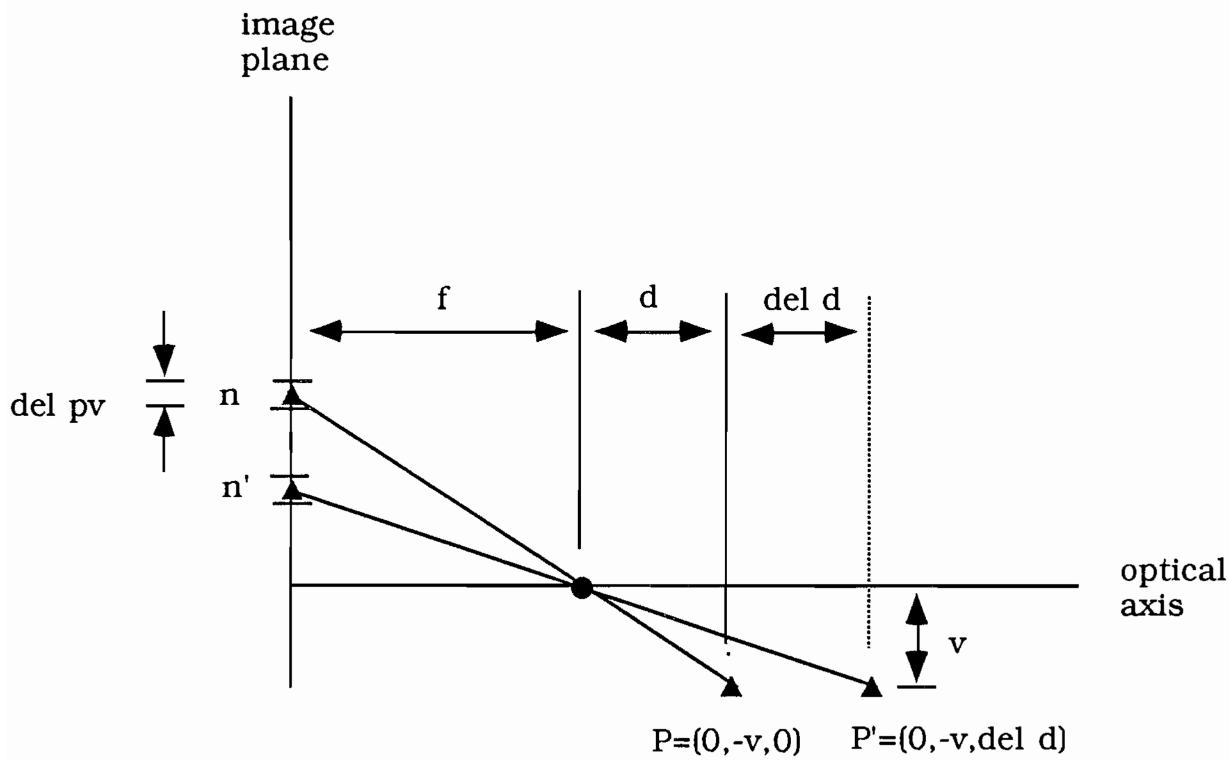

Figure 4.6 Symbolic view of the calibration setup for the determination of the focal length / pixel size

Table 4.1 Projected point object location for horizontal object displacement for the experimental setup of Figure 4.4

Horizontal Displacement

meas	0	1	2	3	4	5	6
opt axis	-0.762	-0.508	-0.254	0	0.254	0.508	0.762
z axis	25	25	25	25	25	25	25
hor	275	271	267	264	259	255	252
vert	273	273	273	273	273	273	273
z axis	31	31	31	31	31	31	31
hor	277	272	268	264	259	254	250
vert	273	273	273	273	273	273	273
z axis	37	37	37	37	37	37	37
hor	279	273	268	264	258	253	248
vert	273	273	273	273	273	273	273
z axis	43	43	43	43	43	43	43
hor	281	275	269	264	257	251	246
vert	273	273	273	273	273	273	273
z axis	49	49	49	49	49	49	49
hor	285	277	271	264	256	250	242
vert	273	273	273	273	273	273	273

Table 4.2 Projected point object location for vertical object displacement for the experimental setup of Figure 4.4

Vertical Displacement

meas		1	3	5	7	9	
opt axis		-0.04	-0.02	0	0.02	0.4	
z axis		25	25	25	25	25	
hor		264	264	264	264	264	
vert		266	269	273	275	278	
z axis		31	31	31	31	31	
hor		264	264	264	264	264	
vert		265	269	273	276	280	
z axis		37	37	37	37	37	
hor		264	264	264	264	264	
vert		264	268	273	276	281	
z axis		43	43	43	43	43	
hor		264	264	264	264	264	
vert		263	268	273	278	282	
z axis		49	49	49	49	49	
hor		264	264	264	264	264	
vert		262	267	273	279	284	

the small deviation of the pixel indices as the point being imaged is moved. In some cases displacing the point did not change the image location. This occurred in spite of the camera being moved along a large portion of the rail.

Table 4.3 and 4.4 show a calculation of the focal length/pixel size using various pairs of similar triangles as shown in Figure 4.6 and described by:

$$\frac{f}{\Delta p_v} = \frac{\Delta d}{v}\left(\frac{nn'}{n-n'}\right) \tag{4.3.1-1}$$

These pairs make use of all possible pairs of similar triangles for the data taken. Since they represent an analysis of the same camera system, they should all be the same with the bounds of experimental error. A statistical examination of these results is presented in Table 4.5 as a function of range pairings and in Table 4.6 as a function of displacement from the optical axis. In both of these cases there are no clear trends in the data and the standard deviations are high.

An overall average value for the a focal length/horizontal pixel size is found to be 903,570, and a focal length/vertical pixel size being 1,854,915. Respective standard deviations in these values are 836,409 and 1,352,084. The large value of focal length over pixel size is appropriate for a small pixel size but the large value of standard deviation raises questions about the usefulness of this data.

An examination of the original data presented in Tables 4.1 and 4.2 reveal that the data is changing with the expected slope as the independent variables are changed. An examination of slight deviations from the recorded data reveals that errors in data acquisition can not account for the high standard deviation. One possible factor is quantization due to the sensor size; the quantity we are attempting to measure. This effect can be eliminated using an experimental setup which allows observations from a distances where the quantization effects are not this profound. These errors, once eliminated during calibration, will not reintroduce themselves during an inspection analysis.

Table 4.3 Horizontal focal length to pixel size computations for Table 4.1

Displacement	Horizontal						
meas	0	1	2	3	4	5	6
opt axis	-0.76	-0.51	-0.25	-	0.254	0.508	0.762
31/25	299902	870614	1690299	-	-	765000	248031
37/31	304264	877039		-	1578472	759000	244094
43/37	308657	443356	1702961	-	1566283	375018	240189
49/43	157648	449852	861012	-	1554142	741142	117189
49/25	246850	591079	1709220	-	2088315	602362	192076
37/25	302067	873815	3380598	-	3156945	761988	246047
43/31	306445	588976	3405921	-	1572354	502000	242126
49/37	208701	446581	1143748	-	1560189	498031	157522
43/25	304232	660162	2544909	-	2358531	566973	244063
49/31	233105	533934	1715622	-	1566236	562500	178642

Table 4.4 Vertical focal length to pixel size computations for Table 4.1

Displacement	Vertical				
meas	1	3	5	7	9
opt axis	-0.04	-0.02	-	0.02	0.4
31/25	583800	2277000	-	-	1057350
37/31	1180200	-	-	2162760	1049400
43/37	1188630	1150920	-	-	1041480
49/43	600660	2326860	-	2146680	1033590
49/25	789520	2301750	-	4309380	1045380
37/25	798040	1540080	-	4293360	1037520
43/31	1184400	2301840	-	4325520	1045425
49/37	798040	1540080	-	4293360	1037520
43/25	881955	2293500	-	6488280	1049370
49/31	894600	2310120	-	3232035	1041450

Table 4.5 Statistical analysis of horizontal and vertical focal length to pixel size data as a function of the ranges compared for Tables 4.3 and 4.4

		Vertical			Horizontal	
range pr		AVE	STD		AVE	STD
31/25		1306050	873568		774769	581040
37/31		1464120	608564		752574	537860
43/37		1127010	76433		772744	672411
49/43		1526948	841692		646831	536023
49/25		2111508	1607391		904984	797316
37/25		1917250	1613971		1453577	1429223
43/31		2214296	1515682		1102970	1226765
49/37		1917250	1613971		669129	560676
43/25		2678276	2616887		1113145	1050150
49/31		1869551	1108564		798340	672327

Table 4.6 Statistical analysis of horizontal and vertical focal length to pixel size data as a function of the displacements compared for Tables 4.3 and 4.4

Vertical Displacement

meas		1	3	5	7	9	
opt axis		-0.04	-0.02	0	0.02	0.4	
AVE		1043849	3906422	-	2004683	889985	
STD		7033	1408915	-	459869	227215	

Horizontal Displacement

meas	0	1	2	3	4	5	6
opt axis	-.076	-.051	-.025	0	0.254	0.508	0.762
AVE	267187	633541	2017143	-	1889052	613401	210998
STD	53009	180281	905330	-	558249	137434	46755

A second possible factor are inaccuracies in the movable mount and camera rail. Both of these must allow for accurate movement of the relative camera/object system. Any errors in our knowledge of the object's movement as we vary its position will introduce errors in our estimate of the focal length/pixel size. Errors due to this factor would also introduce errors into any inspection analysis performed.

The outcome of accurate results from the two parts of the calibration process is a specific set of transformation equations which relate points in views to points in the inspection space. This mapping of points is an essential part of all forms of backprojection inspection and backprojection reconstruction.

4.4 System Parameter Effects in Backprojection Inspection

It is necessary to measure the accuracy of our inspection. Several different factors affect our ability to perform an accurate inspection. These include the camera parameters, the sensor resolution of the image, the ideal object representation, the range of the view from the object, and the angle of the view to the object. Each of these limits accuracy in a different manner.

Our goal is to measure the distance between a landmark on the CAD model of an object and that same point in the reconstructed object. This distance would give a measure of how much error has been introduced. Together these aspects of accuracy limit our ability to distinguish the voxels which will determine this distance.

4.4.1 Camera Parameter Effects

Some camera parameters which affect and limit accuracy were discussed earlier as part of the calibration procedure. The focal length and optical axis of the camera determine the transformation equations used to project points from the inspection space into the view space. Recalling Figure 4.4, the general transformation equations stated before are:

$$m_{\vec{p}} = \left\lfloor \frac{f}{\Delta_{p_h}} \frac{(\vec{P}-\vec{C})\bullet\hat{h}}{(\vec{P}-\vec{C})\bullet\hat{a}} \right\rfloor + m_c \tag{4.4.1-1}$$

$$n_{\vec{p}} = \left\lfloor \frac{f}{\Delta_{p_v}} \frac{(\vec{P}-\vec{C})\bullet\hat{v}}{(\vec{P}-\vec{C})\bullet\hat{a}} \right\rfloor + n_c \tag{4.4.1-2}$$

A camera parameter which does not affect the transformation equations but which does bound the inspection process is the depth of field of the camera. The depth of field limits the valid range of the transformation equations to a certain distance from the camera.

$$d_{min\ DF} \leq d \leq d_{max\ DF} \tag{4.4.1-3}$$

4.4.2 Image Plane Sensor Resolution and View Thresholding Effects

Typically, the imaging sensor which is used to acquire the view measures a grey scale image. It does this by quantizing the view into grey scale values over a pixel grid. The size of the pixels, or image plane sensor resolution, provides the smallest unit of discrimination in the view. Our accuracy, projected onto the view plane, is bounded by the effect of these values.

$$\text{MAX}\ \{\Delta p_v, \Delta p_h\} \leq \textit{Projected}\ \text{MED} \tag{4.4.2-1}$$

Once the view has been obtained it is necessary to threshold the image to obtain the silhouette of the object viewed. This process removes any information which may be present in the grey scale values in order to simplify processing of the silhouettes.

The non-linear nature of thresholding will accentuate any errors which might be present in the object view itself and remove any sub-pixel accuracy which might be present. This means that if the actual boundary between the object and non-object lies across a pixel, the pixel will only be able to represent one value after thresholding. The threshold used is usually determined from view statistics or a priori information about the view.

4.4.3 Octree Inspection Maps and the Inspection Space

The octree inspection map provides the only representation of the ideal object. All information known about a ideal object is represented by this map. Accuracy can not be measured beyond the knowledge of the ideal object, regardless of the accuracy of a computer vision system.

This map is a hierarchic structure where the top level of the octree represents a cubic, geometric real space of a specific size. The depth of the octree has a direct effect on the smallest size volume, or voxel, which is represented by the octree. At the bottom level of the octree each leaf represents a voxel which does not necessarily need to be cubic in shape. It could be any shape which can be recursively packed to describe the entire space. Usually, this shape is chosen to be cubic for simplicity.

A cubic representation of a space may always be constructed by selecting the shortest length of a voxel side as the length of the cubic side. This will always describe the space with a number of voxels greater than or equal to the number used in the non-cubic representation.

If the number of cubic voxels is greater than the number of other voxels the resolution in the representation is increased. For the cubic case, it is easy to express a relationship between the octree depth and the size of the inspection space. The relationship between the size of the reconstruction space, the depth of the octree, and the smallest cubic voxel is given by:

$$\boxed{\Delta V = \begin{matrix}\text{cubic voxel}\\ \text{side length}\end{matrix} = \frac{\begin{matrix}\text{cubic reconstruction}\\ \text{space side length}\end{matrix}}{2^{\text{octree depth}}}} \qquad (4.4.3\text{-}1)$$

This cubic voxel side length provides a bound on the MED in the object space rather than the view space as described in the preceding section. The relationship between these two spaces is examined in the following section. The bound, based upon the cubic voxel side length is:

$$\boxed{\Delta V \leq MED} \tag{4.4.3-2}$$

Thus, for a given size octree inspection map, this equation provides a bound on the resolution of our inspection. Ultimately, this size of the smallest voxel is limited by the memory size of the computer system used for the inspection. In the following Sections, a cubic representation of the inspection space is assumed.

4.4.4 View Range Magnification Effects

The distance between the test object and the views will have an effect upon the amount of information that any given area of a view can contribute to a reconstruction voxel. This is due to the magnification and perspective transformation of the voxel in the view. This was examined by Tan [123] for the case of an arbitrary set of views. Our analysis differs in its incorporation of perspective information and a tighter bound on the maximum voxel projection size. Further, this analysis shall assume that all views of an object and the reconstruction space are available.

We are interested in finding the effect of range on our ability to reconstruct regions of the inspection space. In order for this inspection to occur we must be able to distinguish between adjacent voxels during the inspection. The limiting case of this ability to distinguish voxels would use the largest possible projection of a voxel distinguished by the smallest number of view pixels. The smallest number of pixels which could be used would be one pixel per projection. The largest voxel projection would be a diagonal view of a voxel. As discussed in Section 4.4.3, we shall assume a cubic inspection space representation without loss of resolution. We shall restrict our analysis to the view which provides this largest projection into the view. Figure 4.7 shows this voxel projection with notation which we base upon Tan's analysis. This

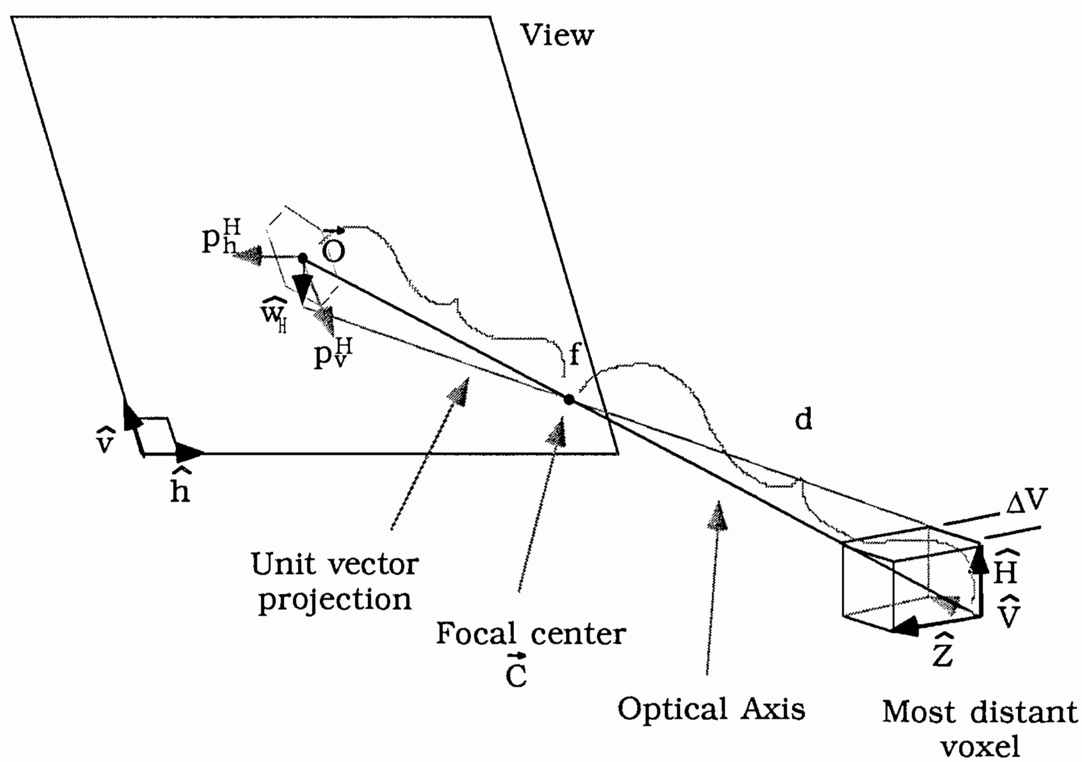

Figure 4.7 Symbolic view of large voxel projection due to range effects

projection is the largest which can occur during inspection. Therefore, we wish to find the maximum sensor dimensions (pixel size) which can distinguish this voxel from an adjacent voxel.

Let $\widehat{V}$, $\widehat{H}$, and $\widehat{Z}$ be the unit vectors of the corner of the voxel we will view diagonally as shown in Figure 4.7. Let the length of the side of a voxel be ΔV and the distance from the focal center to the view be d. Figure 4.8 shows the effect the perspective transformation has on the effective length of $\widehat{V}$. $\widehat{H}$ and $\widehat{Z}$ are similar due to symmetry about the optical axis. From similar triangles, X can be determined.

$$X = \frac{(\Delta V)^2 / 2}{d - \Delta V / \sqrt{2}} \tag{4.4.4-1}$$

This gives an effective length of ΔV to be projected between parallel planes of the sum of this and the orthogonal projection.

$$X + \frac{\Delta V}{\sqrt{2}} = \frac{d\ \Delta V / \sqrt{2}}{d - \Delta V / \sqrt{2}} \tag{4.4.4-2}$$

Figure 4.9 shows the projection of this vector onto the view plane. The length of the projection, assuming parallel vectors, can be determined from the figure using similar triangles.

$$Y = \frac{f}{d}\left(X + \frac{\Delta V}{\sqrt{2}}\right) \tag{4.4.4-3}$$

The projection of $\widehat{V}$, $\widehat{H}$, and $\widehat{Z}$ onto the view can now be determined as follows:

$$\vec{w}_v = \frac{f}{d}\left(\frac{d\ \Delta V / \sqrt{2}}{d - \Delta V / \sqrt{2}}\right)\left(-\widehat{V} \bullet \widehat{h} + -\widehat{V} \bullet \widehat{v}\right) \tag{4.4.4-4}$$

$$\vec{w}_h = \frac{f}{d}\left(\frac{d\ \Delta V / \sqrt{2}}{d - \Delta V / \sqrt{2}}\right)\left(-\widehat{H} \bullet \widehat{h} + -\widehat{H} \bullet \widehat{v}\right) \tag{4.4.4-5}$$

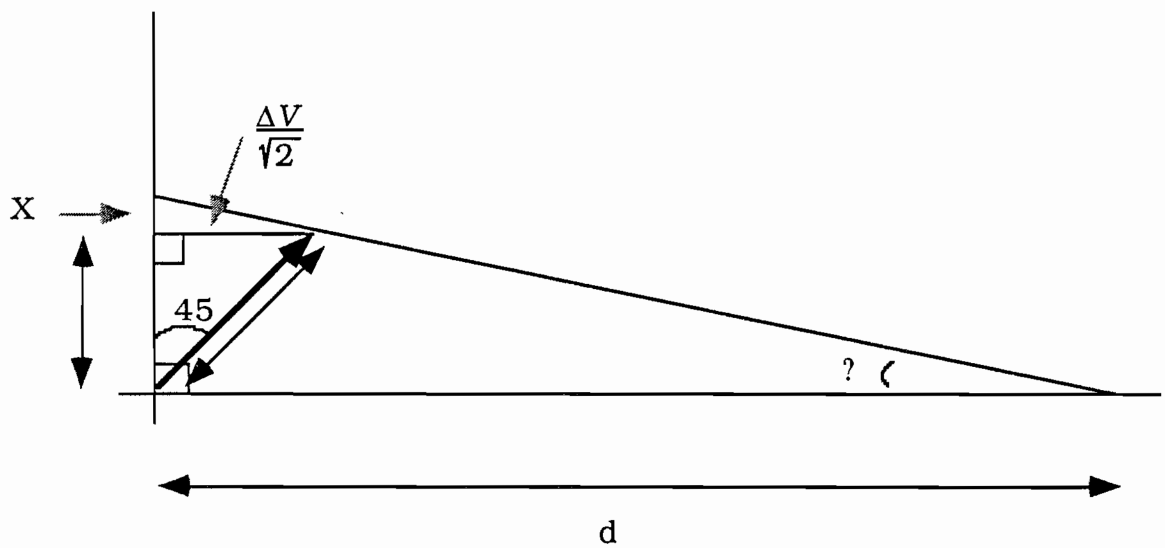

Figure 4.8 Perspective effects on a voxel edge on the optical axis, subjected to large voxel projection

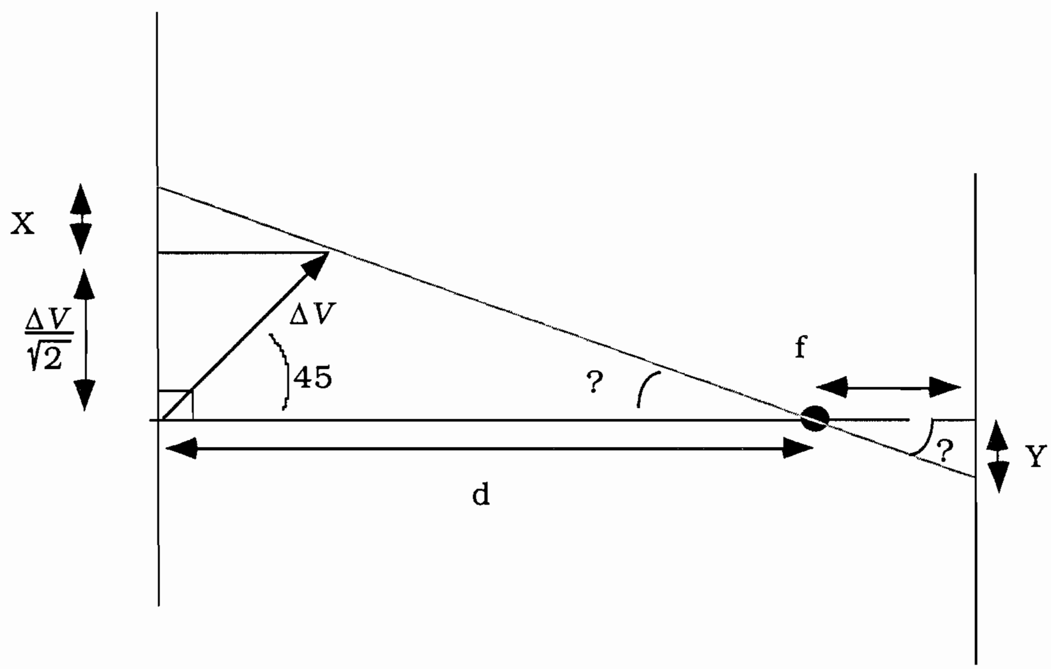

Figure 4.9 Perspective transformation on a vector parallel to the view

$$\vec{w}_z = \frac{f}{d}\left(\frac{d\ \Delta V\ /\ \sqrt{2}}{d\ -\ \Delta V\ /\ \sqrt{2}}\right)\left(-\hat{Z}\bullet\hat{h} + -\hat{Z}\bullet\hat{v}\right) \tag{4.4.4-6}$$

These projections have origin in the view plane and all have the same length. The angle between $\vec{w}_v$, $\vec{w}_h$, and $\vec{w}_z$ is 120 degrees from the symmetry of the original voxel and the optical axis being along the cube's diagonal. This is shown in Figure 4.10 with the arbitrary selection of $\left|\vec{w}_z \bullet \hat{h}\right| = 0$. We shall now project these vectors onto the view plane axes.

$$\vec{p}_v^H = \left|\vec{w}_h \bullet \hat{v}\right| \text{ and } p_h^H = \left|\vec{w}_h \bullet \hat{h}\right| \tag{4.4.4-7}$$

$$\vec{p}_v^H = \left|\vec{w}_v \bullet \hat{v}\right| \text{ and } p_h^H = \left|\vec{w}_v \bullet \hat{h}\right| \tag{4.4.4-8}$$

$$\vec{p}_v^H = \left|\vec{w}_z \bullet \hat{v}\right| \text{ and } p_h^H = \left|\vec{w}_z \bullet \hat{h}\right| \tag{4.4.4-9}$$

Form the length of each projection in the view.

$$r^H = \sqrt{(\vec{p}_v^H)^2 + (\vec{p}_h^H)^2} \tag{4.4.4-10}$$

$$r^H = \sqrt{(\vec{p}_v^V)^2 + (\vec{p}_h^V)^2} \tag{4.4.4-11}$$

$$r^H = \sqrt{(\vec{p}_v^Z)^2 + (\vec{p}_h^Z)^2} \tag{4.4.4-12}$$

Recall that all these projections are of the same length, but with different orientations.

$$r = r^H = r^V = r^Z = \left|\vec{w}_v\right| = \left|\vec{w}_h\right| = \left|\vec{w}_z\right| \tag{4.4.4-13}$$

We wish to determine the maximum pixel size which can distinguish this projection length for all three projections at the same time. For a rectangular pixel, the length of the shortest side is a lower bound on the pixel's ability to distinguish distances in a single direction. For any rotation of a pixel with respect to a voxel projection, the pixel's

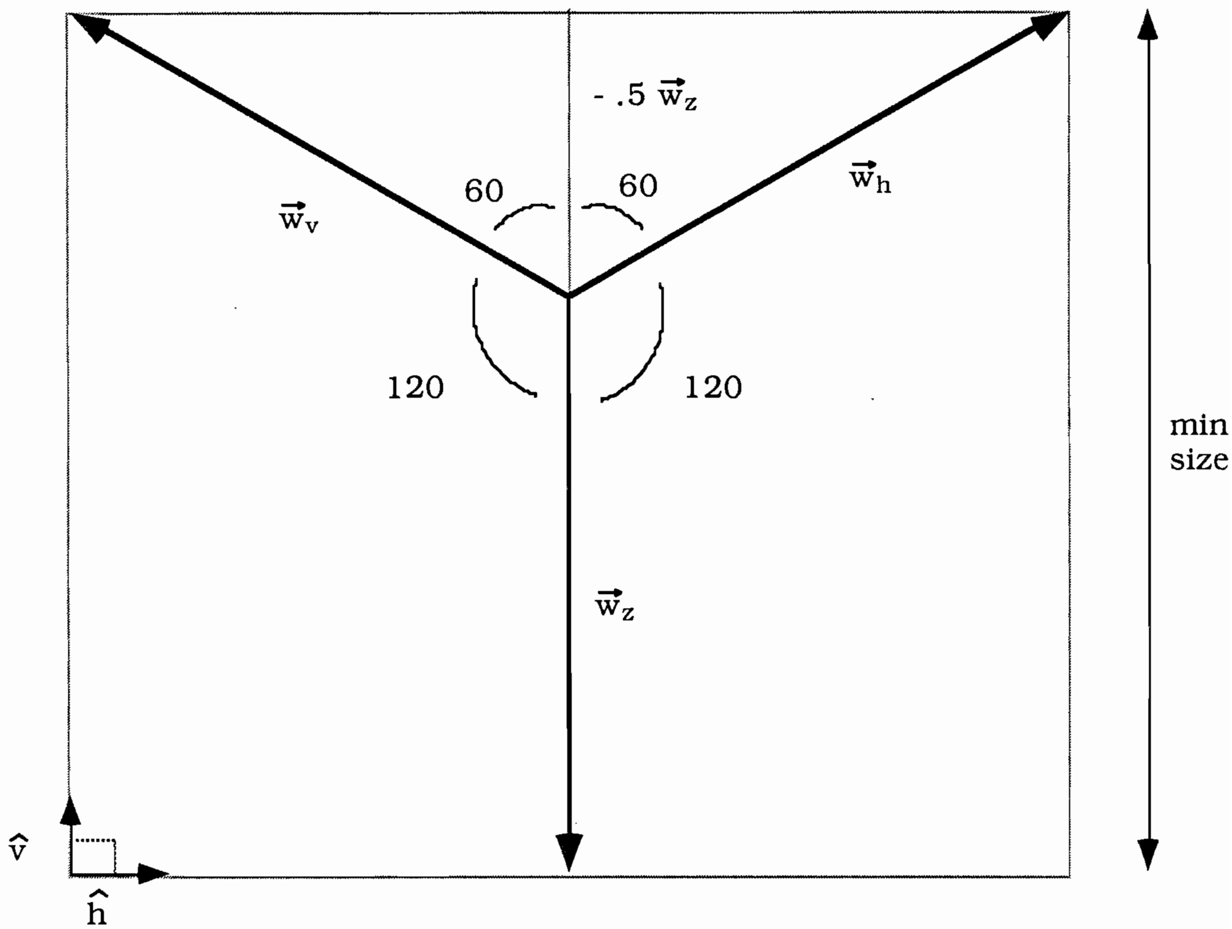

Figure 4.10 View plane projection of the edges of a voxel yielding a large voxel projection

diagonal is a lower bound on the projection length to guarantee discrimination.

$$r^d = \sqrt{(\Delta p_h)^2 + (\Delta p_v)^2} \tag{4.4.4-14}$$

This distance is the discrimination requirement stated by Tan [123]. We claim a less stringent requirement as sufficient for discrimination: The longer side of the pixel must be less than or equal to the shorter side of the smallest bounding rectangle containing the projection of the voxel edges into the view.

$$\Delta p_{v,h} < \text{shorter side of the smallest bounding rectangle} \tag{4.4.4-15}$$

From the geometry of the projection into the view shown in Figure 4.10 we find:

$$\Delta p_{v,h} < \frac{3}{2}\, r \tag{4.4.4-16}$$

Recalling that r is the magnitude of all three unit vector projections into the view; this can be stated in terms of $|\vec{w}_z|$.

$$\Delta p_{v,h} < \frac{3}{2}\, |\vec{w}_z| \tag{4.4.4-17}$$

$|\vec{w}_z|$ immediately follows from the projection into the view as:

$$|\vec{w}_z| = \frac{f}{d}\left(\frac{d\ \Delta V / \sqrt{2}}{d - \Delta V / \sqrt{2}}\right) \tag{4.4.4-18}$$

This gives us an overall expressions relating focal length, distance, and voxel size to the maximum pixel dimension.

$$\boxed{\Delta p_{v,h} < \frac{3}{2}\left(\frac{f}{d}\right)\left(\frac{d\ \Delta V / \sqrt{2}}{d - \Delta V / \sqrt{2}}\right)} \tag{4.4.4-19}$$

An examination of this equation shows that for cases where $d >> \Delta V$, this equation can be simplified to:

$$\Delta p_{v,h} \langle \frac{3}{2} \left(\frac{f}{d}\right)\left(\frac{\Delta V}{\sqrt{2}}\right) \tag{4.4.4-20}$$

This is the usual case when the view is a reasonable distance from the inspection space or when orthogonal transformations are used.

4.4.5 View Angle Effects

In a similar manner to the distance from the object, the number of voxels within an angle of the view has an effect on the number of pixels required to discriminate between adjacent voxels. This angle, Θ, shall be expressed as the greatest angle between a voxel along the optical axis and the most distant voxel from the optical axis we wish to concurrently view. The magnitude of the view angle will constrict our ability to chose the view to yield a largest projection for all voxels within this angle. The number of voxels within the view angle shall be expressed by a distance D from the optical axis to a vertex of a voxel off the optical axis. We will assume a cubic inspection representation and select a geometry which maximizes the voxel projections for a given angle.

Figure 4.11 shows these voxels, the angle describing them, and the distance between the optical axis and the most distant voxel. Once again, a maximum projection will be found when the optical axis passes through the major diagonal of a voxel. This view is selected for the voxel lying on the optical axis. Unfortunately, a view does not exist which allows more than one voxel to be viewed along the line of maximum projection. Even though it is not possible to select a view line which passes through the diagonal of the voxel a distance D off the optical axis, a view is chosen such that the view line passes as close to the diagonal as possible. This will result in projections similar to those of Figure 4.10, but without the symmetry.

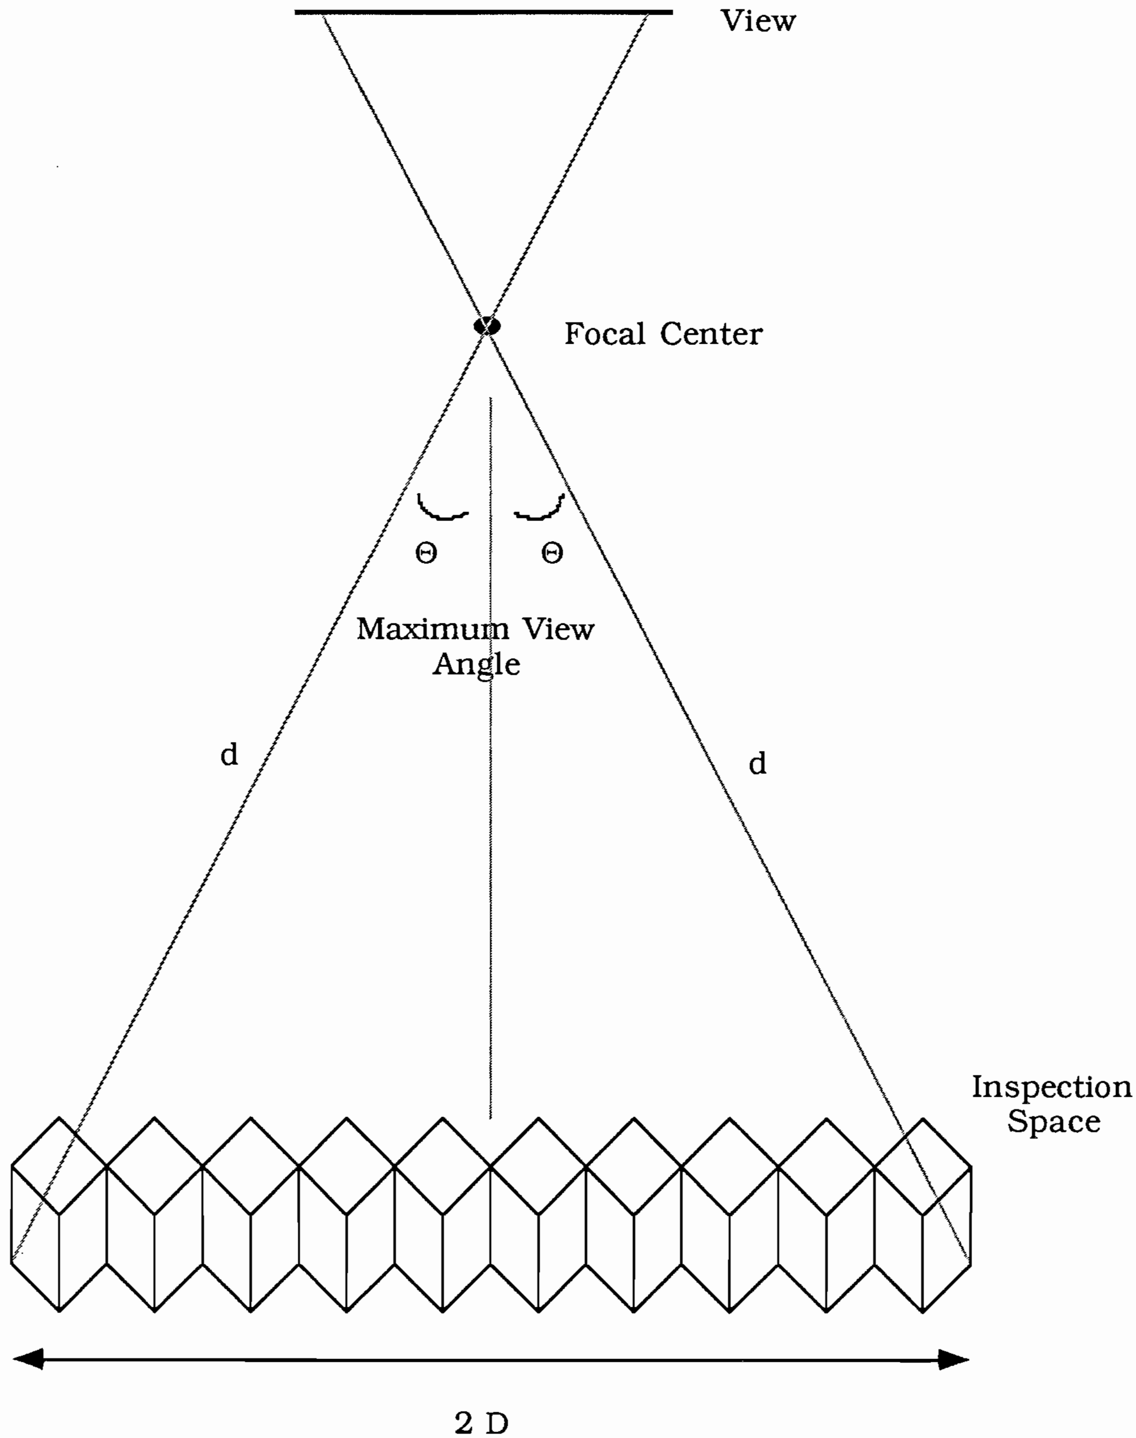

Figure 4.11 View angle observation of a line of voxels

The distance between the optical axis and the furthest voxel to be reconstructed is an measure of the amount of the inspection space which we need to be able to reconstruct with the given view angle. From this, the distance along the view line can be determined.

$$d = D \sin^{-1}\Theta \tag{4.4.5-1}$$

In this case the view line through the most distant voxel will not lay along the optical axis or through the diagonal of the voxel. The projection will take place at an angle to the optical axis and this angle passes through the reconstruction voxel at an angle determined by the view angle. Figure 4.12 shows the projection line passing through the reconstruction voxel with associated terms labeled. The effect of this angle is a perspective transformation from the object to the view due solely to the view angle.

In a similar manner to Section 4.4.4, a unit vector of the voxel at the maximum view angle can be projected into a plane parallel to the view. Unlike Section 4.4.4, the projection of a unit vector lacks the symmetry shown in Figure 4.10. Projections toward the optical axis and away from the optical axis must be performed separately. The projections will take place in a plane determined by the optical axis and the view line. All three voxel vertices will not lie in this plane. We shall use ΔV to represent the projection of the voxel edges into this plane, away from the optical axis; and ΔZ to represent the projection of the voxel edges into this plane, toward the optical axis.

The effect of the view angle, as shown in Figure 4.13, is to lengthen the projection of a voxel edge, ΔV, which points away from the optical axis in comparison with the projection of Section 4.4.4. From similar triangles X can be determined.

$$X = \frac{(\Delta V/\sqrt{2})\, d\sin\Theta + (\Delta V)^2/2}{d\cos\Theta - \Delta V/\sqrt{2}} \tag{4.4.5-2}$$

From this the effective length of the voxel edge can be determined.

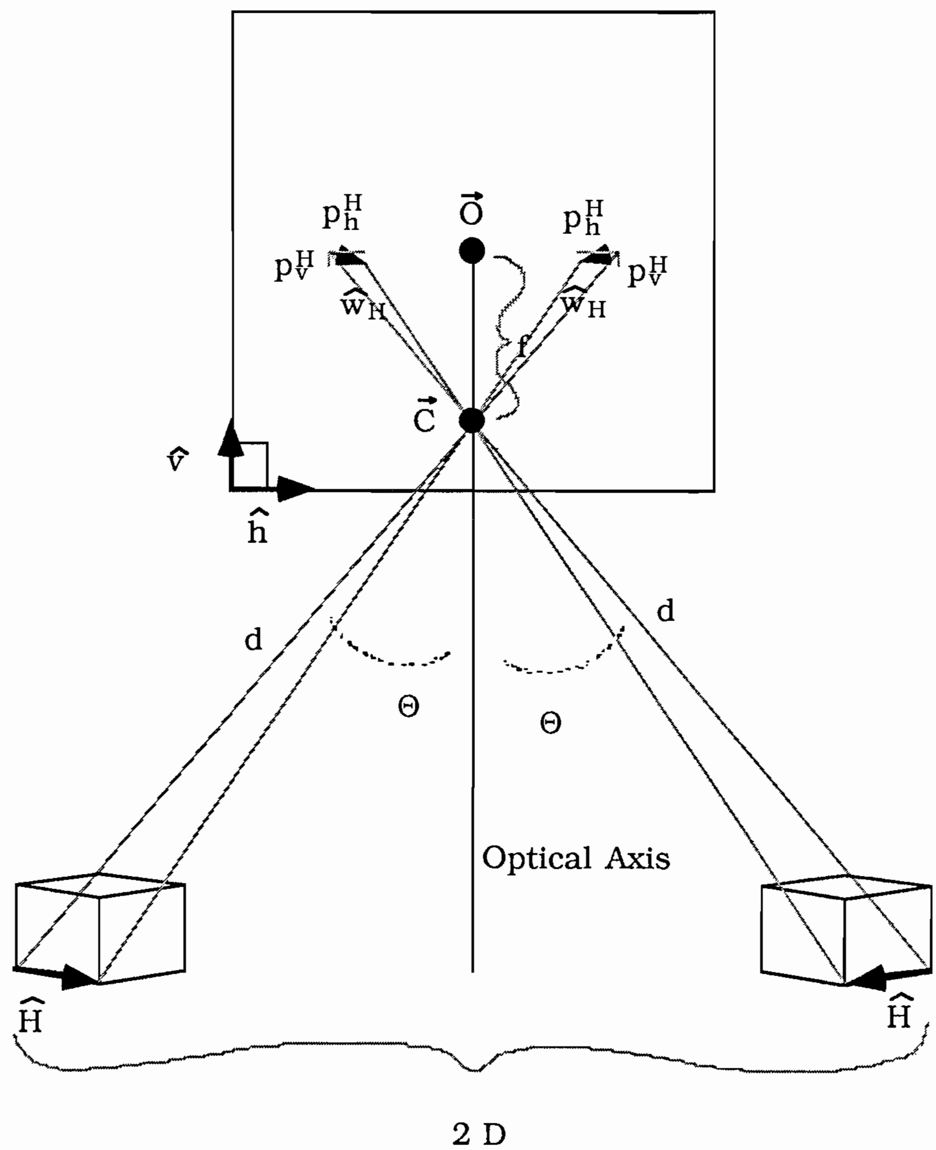

Figure 4.12 Symbolic view of large voxel projection due to angle effects

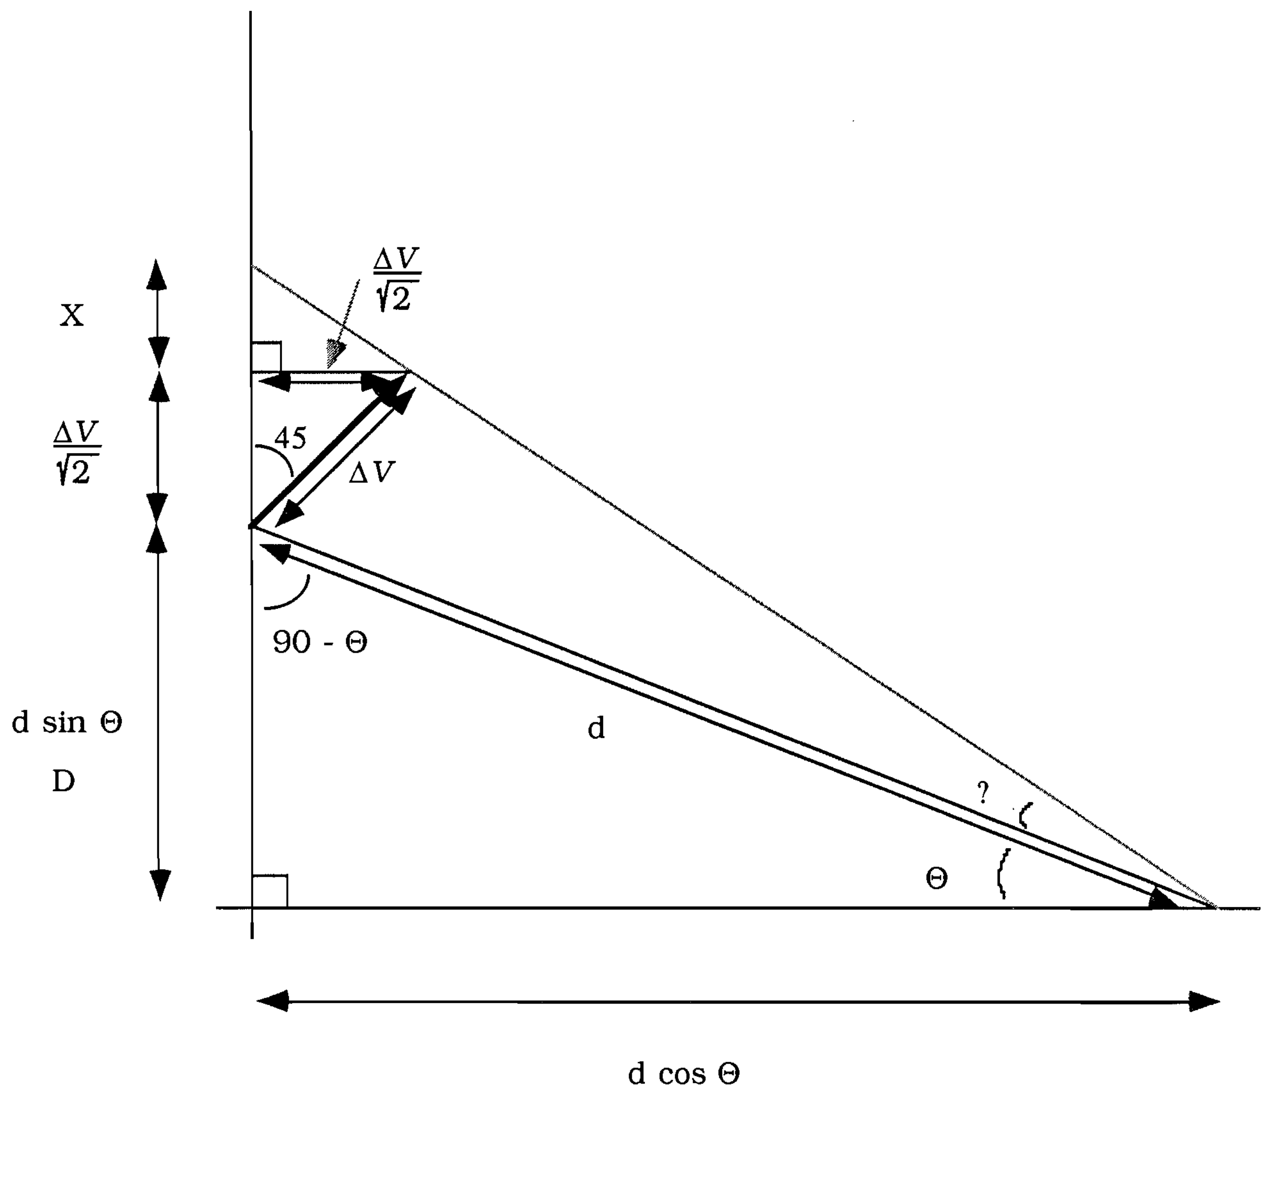

Figure 4.13 Perspective effects on a voxel edge slanted away from the optical axis but off the optical axis, for large voxel projection

$$\frac{\Delta V}{\sqrt{2}} - X = \frac{\left(\Delta V/\sqrt{2}\right) d(\cos\Theta + \sin\Theta)}{d \cos\Theta - \Delta V/\sqrt{2}} \tag{4.4.5-3}$$

The effect of the view angle, as shown in Figure 4.14, is to shorten the projection of a voxel edge, ΔW, which points toward the optical axis in comparison with the projection of Section 4.4.4. From similar triangles X can be determined.

$$X = \frac{\left(\Delta W/\sqrt{2}\right) d\sin\Theta - (\Delta W)^2/2}{d \cos\Theta - \Delta W/\sqrt{2}} \tag{4.4.5-4}$$

From this the effective length of the voxel edge can be determined.

$$\frac{\Delta W}{\sqrt{2}} - X = \frac{\left(\Delta W/\sqrt{2}\right) d(\cos\Theta - \sin\Theta)}{d \cos\Theta - \Delta W/\sqrt{2}} \tag{4.4.5-5}$$

In a similar manner to the previous section, we wish to chose a geometry for these projections which yields a maximum area in the view. The pixel size is then determined from the minimum size of the projection. The maximum projection will have a minimum distance in the plane formed by the optical axis and the view line when a voxel edge is aligned with the plane formed by the optical axis and view line as shown in Figure 4.15. The distance along the shortest side of the smallest bounding rectangle can once again provide the discrimination requirement.

$$\frac{\left(\Delta V/\sqrt{2}\right) d(\cos\Theta + \sin\Theta)}{d \cos\Theta - \Delta V/\sqrt{2}} + \frac{\left(\Delta W/\sqrt{2}\right) d(\cos\Theta - \sin\Theta)}{d \cos\Theta - \Delta W/\sqrt{2}} \tag{4.4.5-6}$$

This discrimination requirement is difficult to simplify into a form which provides either intuitive meaning or simple determination even if ΔV and ΔW are related in a simple manner. For the view angle case we propose a more stringent requirement as sufficient for discrimination: The longer side of the pixel must be less than or equal to twice the

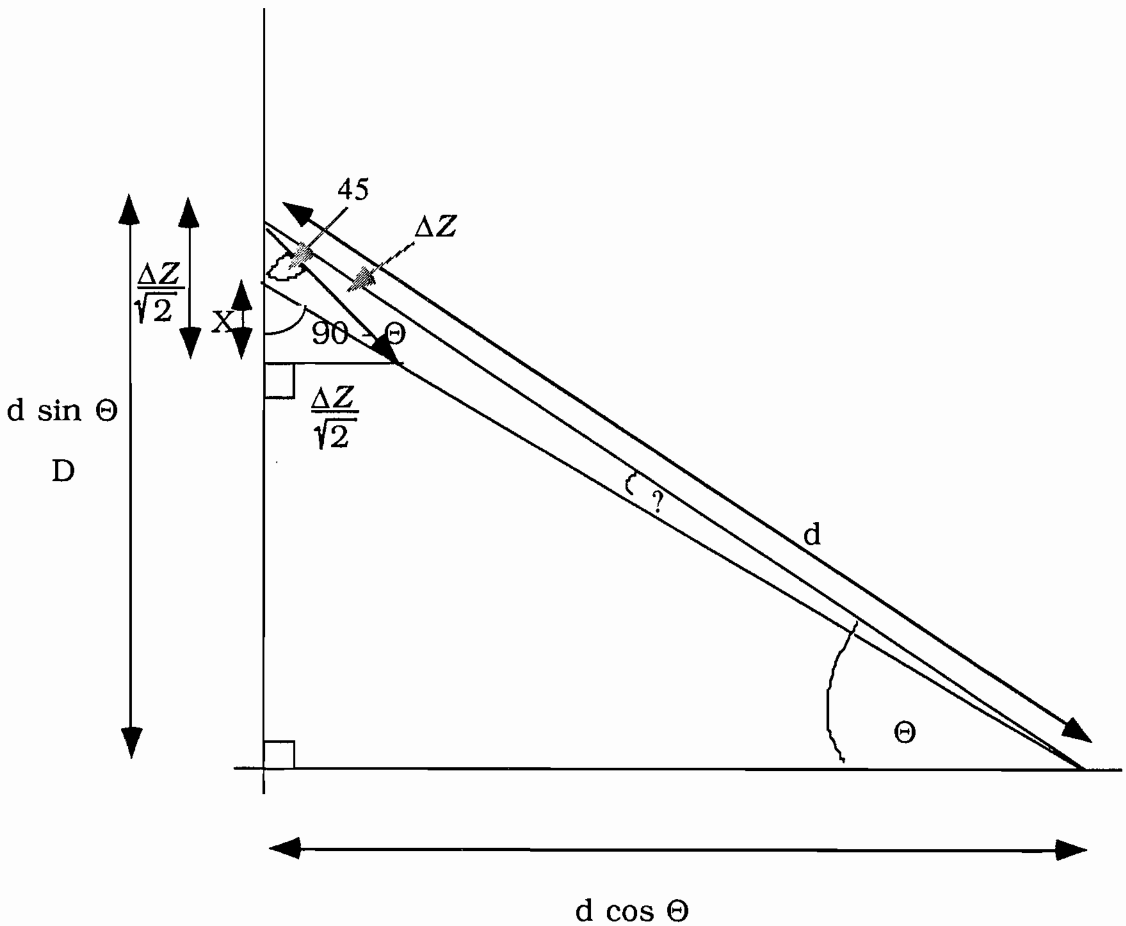

Figure 4.14 Perspective effects on a voxel edge slanted away from the optical axis but off the optical axis, for large voxel projection

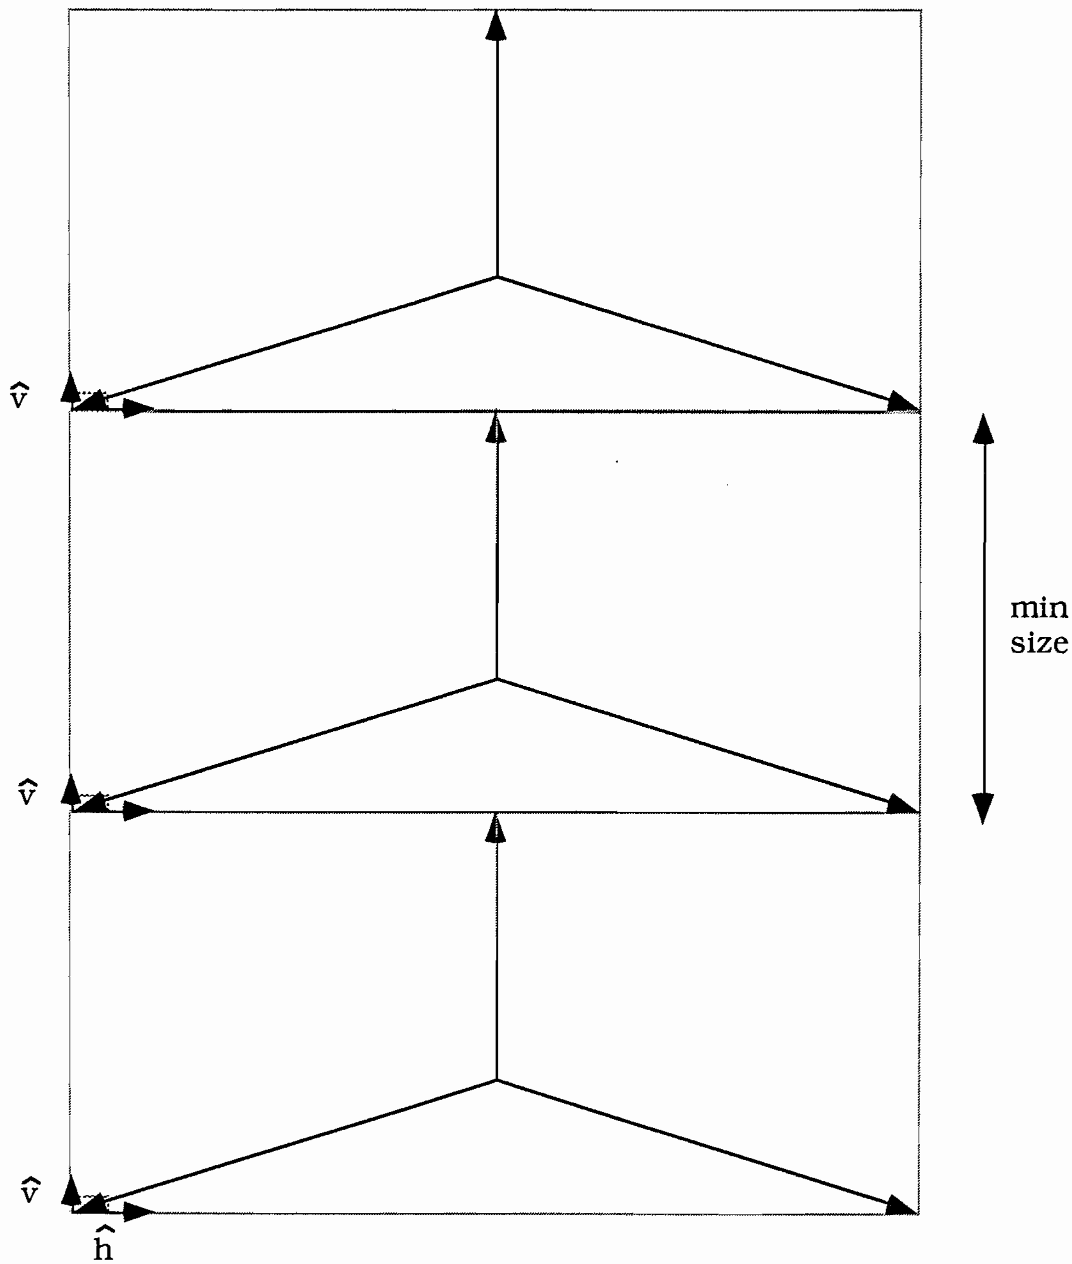

Figure 4.15 Projection of several voxels onto the view from an angle Θ

magnitude of the projection of the sum of two voxel edges toward the optical axis.

$$\Delta p_{v,h} < \text{twice the magnitude of short projection} \tag{4.4.5-7}$$

Letting r be the length of the projected edge:

$$\Delta p_{v,h} < 2\ r \tag{4.4.5-8}$$

The length of the voxel edges projected into the plane defined by the optical axis and the view line is:

$$\Delta W = \frac{\Delta V}{2} \tag{4.4.5-9}$$

Projecting this into the view yields:

$$r = \frac{\left(\Delta V/2\sqrt{2}\right)\ d(\cos\Theta - \sin\Theta)}{d\ \cos\Theta - \Delta V/2\sqrt{2}} \tag{4.4.5-10}$$

This provides an overall expression relating the view angle, distance to be examined, view distance, focal length, and pixel size of the sensor.

$$\boxed{\Delta p_{v,h} < 2\left(\frac{f}{d}\right)\left(\frac{\left(\Delta V/2\sqrt{2}\right)\ d(\cos\Theta - \sin\Theta)}{d\ \cos\Theta - \Delta V/2\sqrt{2}}\right)} \tag{4.4.5-11}$$

Unlike the limit from Section 4.4.4, this relationship provides information on how a single view can be used to view more than one voxel. Note that d can be eliminated to yield an expression which is not dependent on the view distance. This relationship is fundamental to the Chapter 5 algorithm which selects views for inspection. As Θ becomes small and $d >> \Delta V$, this relationship becomes:

$$\Delta p_{v,h} < 2 \left(\frac{f}{d}\right) \left(\frac{\Delta V}{2\sqrt{2}}\right) \tag{4.4.5-12}$$

This is consistent with our results from Section 4.4.4 for viewing along the optical axis from a reasonable distance. This bound is 33% more restrictive on the pixel values because of the looser bound on the boundary of the projected edges.

50% Overlap Rule

In the limiting cases just described, the voxel was determined by a single pixel in the view. When voxels are examined from an angle we are no longer attempting to find a view which will have a voxel projection which is all object or void. Thus, in the non-limiting cases the voxel projection may be mixed. This requires the ability to determine the identity of a voxel from a mixed projection.

What is required is a way to determine if a separating view exists given a view which does not provide a projection which separates the voxels into object and background. This is provided by the *50% Overlap Rule*. Figure 4.16 shows several views of an object which result in partial projections, but for which a separating view exists.

50% Overlap Rule: Any voxel projection which provides an overlap between voxel projection and the object in the view of greater than or equal to 50% with the void in the view could be viewed from an angle and location which will make the overlap contain only void.

This rule provides assurance that the projections along angles not along the optical axis can always distinguish voxels as required by the *Weak Convexity* of Chapter 2.

4.5 Minimum Discrimination Specification

Thus far, we have discussed the calibration of the inspection process and some bounds on our ability to measure the accuracy of the

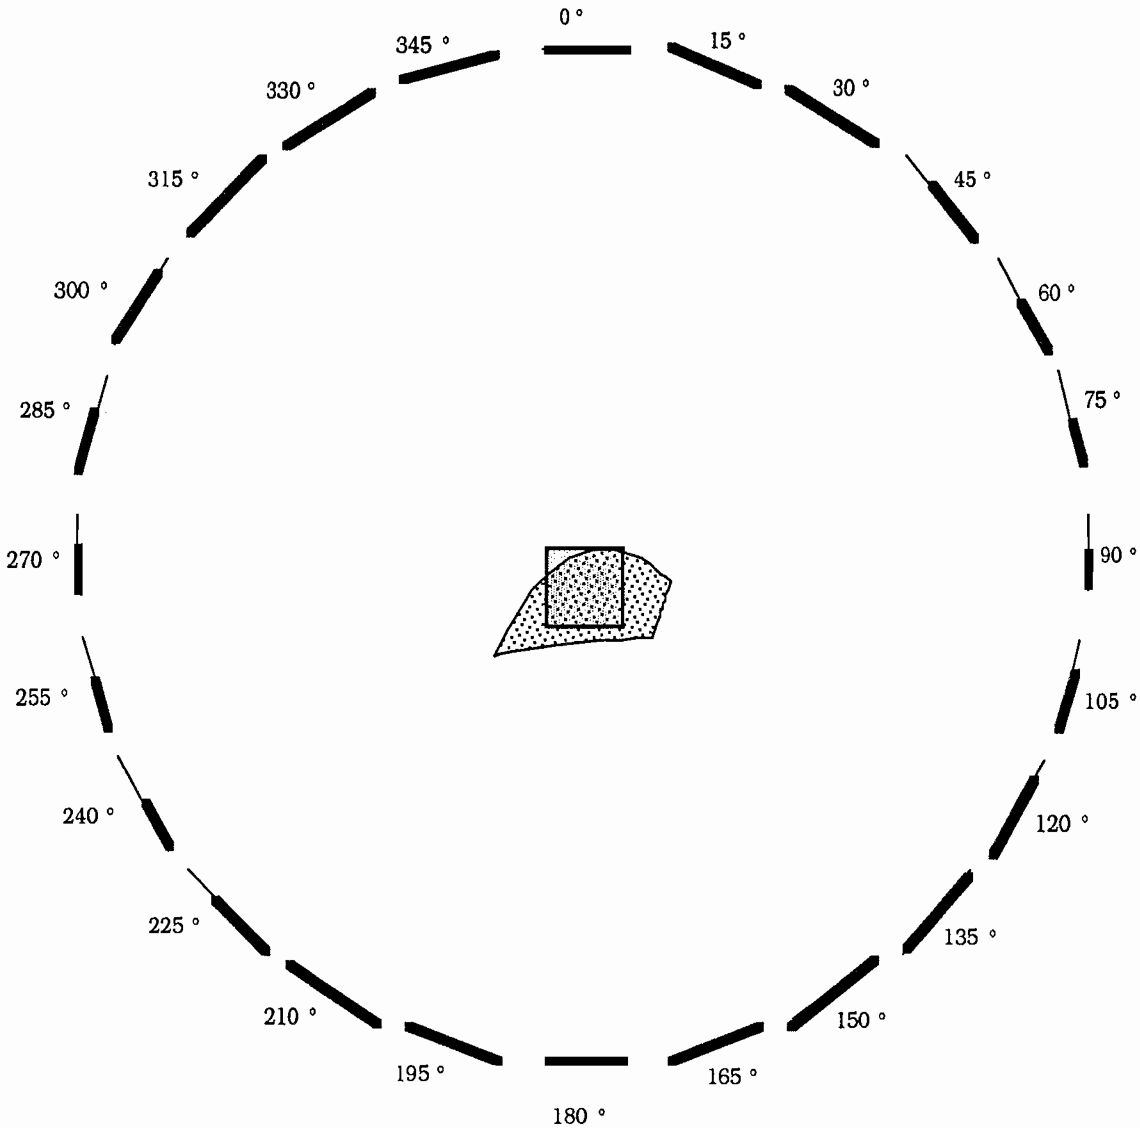

Figure 4.16 Projection views for adjacent object/void voxels with overlap shown

inspection reconstruction. These were based upon the constraints of the inspection reconstruction process itself. These bounds will now be grouped and re-expressed in a manner which will define overall bounds on the inspection system.

4.5.1 Minimum A Priori Discrimination (MAD) Specification

Recalling Section 4.4.3, the minimum a priori discrimination is the smallest measurement which can be confirmed by an inspection process. This is a function of the representation of the ideal object. It is in terms of the size of the inspection reconstruction space.

$$\boxed{\Delta V = \begin{array}{c}\text{cubic voxel}\\\text{side length}\end{array} = \frac{\begin{array}{c}\text{cubic reconstruction}\\\text{space side length}\end{array}}{2^{\text{octree depth}}}} \qquad (4.5.1\text{-}1)$$

4.5.2 Minimum Experimental Discrimination (MED) Specification

The minimum experimental discrimination is the smallest measurement which can be performed by the experimental inspection system. This amount is bound by the expressions of Sections 4.4.1-2 and 4.4.4-5. A satisfactory inspection must occur under conditions satisfied by these equations.

When performing silhouette backprojections it is not possible to achieve an accuracy better than the object representation.

$$\boxed{\Delta V \leq \text{MED}} \qquad (4.5.2\text{-}1)$$

Distance between the object and view has a magnification effect along the optical axis described by 4.5.2-2 when only one voxel needs to be observed.

$$\boxed{\Delta p_{v,h} < \frac{3}{2}\left(\frac{f}{d}\right)\left(\frac{d\ \Delta V\ /\ \sqrt{2}}{d\ -\ \Delta V\ /\ \sqrt{2}}\right)} \tag{4.5.2-2}$$

Distance and the angle of the view are related in 4.5.2-3 to the pixel size when more that one voxel needs to be observed.

$$\boxed{\Delta p_{v,h} < 2\left(\frac{f}{d}\right)\left(\frac{(\Delta V/2\sqrt{2})\ d(\cos\Theta\ -\ \sin\Theta)}{d\ \cos\Theta\ -\ \Delta V/2\sqrt{2}}\right)} \tag{4.5.2-3}$$

A satisfactory solution of these equations is necessary whenever an inspection is to be performed.

4.6 Accuracy Analysis in Backprojection Inspection

It is necessary to quantify the accuracy of our inspection. However, it is not obvious how to form the raw inspection results into meaningful information. While there are many methods for dealing with error in two dimensional data, this is not as obvious for three dimensional data. Any accuracy measure is a distance measure between a desired result and the measured result. With three dimensional data, it is not obvious how to select the measured result for comparison.

Therein lies the problem with applying error measures to three dimensional data. There is no obvious way to measure the error distance between the differences of two objects. Our method uses object landmarks to give us specific locations on an object to evaluate instead of attempting to evaluate the object as a whole. Differences can then be examined on a point level instead of on an object level. These differences can then be related to our original goals of a minimal a priori discrimination (MAD) specification and a minimal experimental discrimination (MED) specification.

Because it is not possible to characterize the error which is present in the reconstructed object we lack a straightforward method to estimate the location of a point which has been improperly reconstructed. There is an entire object boundary, many of whose points may correspond to the actual object point which was not correctly reconstructed. It is therefore necessary to make use of an error measure which non-parametrically estimates the location of the improperly reconstructed point.

None of the aspects of accuracy which we have discussed have provided a manner to express the result of the accuracy inspection itself. We shall now present some ways to describe the results of the algorithm of Chapter 3. These are algorithms for combining the results of the accuracy inspection to form some useful expression of error.

4.6.1 Previous Methods

Other error measurement methods which have been used for three dimensional analysis include volume and area calculations. The volume method involves simply comparing the volume of the two different objects; the first being the know and the second being the unknown. This has the advantage of being simple to calculate but has the disadvantage of being almost meaningless. It would be easy for the test object to be quite different from the goal object and to have complementing volume errors cancel out and yield a low error inspection.

The area method involves the determination of the area of a hole in an object. This was done by Tan [123]. There is an implied assumption that the edges of his hole were found by assuming that they were in the plane face of the inspection space. This allowed for the formation of an average of dimensions at different depths into the hole. For the specific case, this is acceptable. However, the error measure makes use of the orientation of the object in the inspection frame. This is not general and it restricts measurements unnecessarily.

Another proposed method involves defining the error as the distance along the surface normals from the know to the calculated object. This gives a distance measure along a line which is normal to one of the two surfaces. Sometimes this can lead to surprising results at places of large curvature. Further, there is no real justification behind the assumption that the surface normal provides an appropriate mapping from one surface to another.

4.6.2 A General Inspection Error Measure

We propose an error measure which is based upon the minimum and maximum distance between the ideal landmark location and estimates of its location in the inspection reconstruction. Both the minimum and maximum distance shall be noted in relation to the landmark. Figure 4.17 shows these distances. We assume that the error distances are such that they can be represented by the inspection process; i.e., the distance is less than the adjacent voxel depth in the inspection map.

This provides error distances which give the minimum and maximum displacement of the reconstructed surface from the actual object. The reconstruction error measure at each landmark will consist of the mean of these maximum and minimum error distances. This provides bounds on the overall object error as well as a specific estimate.

4.6.3 Statistical Parameter Analysis

A statistical analysis of the error in an object could be calculated by simple statistical analysis. This statistical analysis is based upon the maximum, minimum, and average of the general inspection error measure noted in Section 4.6.2. Each of these statistical error estimates seeks to find a general description for the error in the object as a whole. In the equations which follow, the parameter i varies over the N landmarks defined on the inspected object.

Object Mean Error

The overall error measures for the object could be estimated as the mean of these the landmark error bounds and averages The interpretation of this would be an expression of averages over the surface as represented by the landmarks.

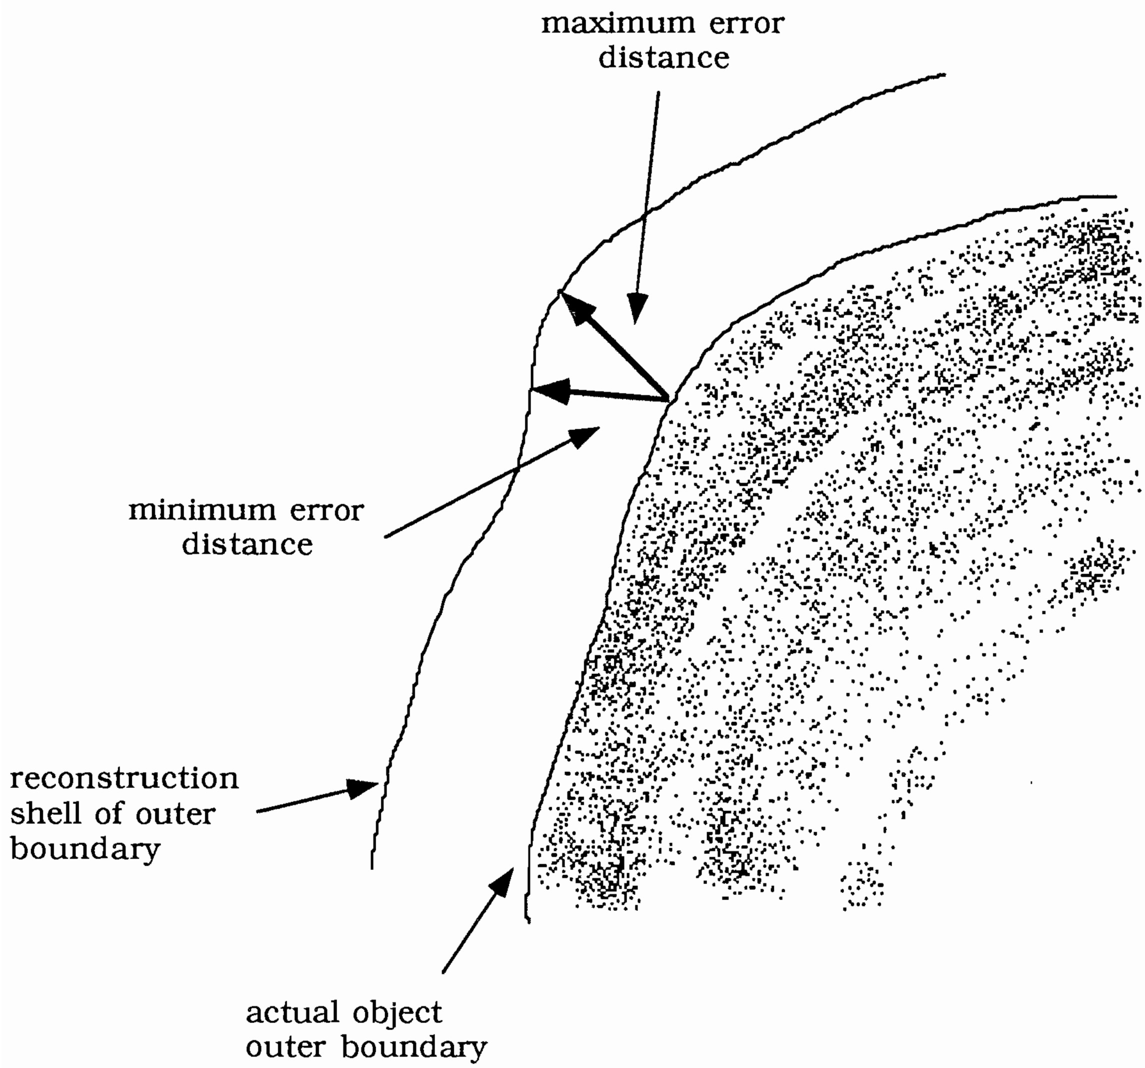

Figure 4.17 Maximum and minimum error distances for boundary reconstruction

$$ave\text{-}error_{max} = \frac{1}{N} \sum_{i=1}^{N} error_{max,i} \tag{4.6.3-1}$$

$$ave\text{-}error_{min} = \frac{1}{N} \sum_{i=1}^{N} error_{min,i} \tag{4.6.3-2}$$

$$ave\text{-}error = \frac{ave\text{-}error_{max} + ave\text{-}error_{min}}{2} = \frac{1}{N} \sum_{i=1}^{N} error_{ave,i} \tag{4.6.3-3}$$

Object Error Variance

The overall variance in the error measures for the object could be estimated as well. These are as follows.

$$std\ of\ error_{max} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (error_{max,i} - ave\text{-}error_{max})^2} \tag{4.6.3-4}$$

$$std\ of\ error_{min} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (error_{min,i} - ave\text{-}error_{min})^2} \tag{4.6.3-5}$$

$$std\ of\ ave\ error = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (error_{ave,i} - ave\ error_{ave})^2} \tag{4.6.3-6}$$

Object Maximum Error

Another expression for the error in the object could be the maximum and minimum of the individual errors. This would have the interpretation of indicating the maximum and minimum errors found at any point on the object.

$$max\ error = \operatorname*{MAX}_{i=1}^{N} (error_{max,i}) \tag{4.6.3-7}$$

$$min\ error = \operatorname*{MIN}_{i=1}^{N} (error_{min,i}) \tag{4.6.3-8}$$

4.6.4 Geometric Parameter Analysis

Geometric parameter analysis consists of interpretating the correctness of the inspected object on the basis of a parameter calculation. This parameter would be based upon the individual error measures at each landmark as described in Section 4.6.2. The use of parameters seeks to base correctness in an inspected objects ability to possess the property of a parameter.

Distance

One parameter could be the distance between two landmarks. On the ideal object, this would be a specific distance. In the inspection, this distance would be based upon the maximum and minimum estimates of the two landmark. Two approaches are possible.

The first approach would be to base the minimum distance measure upon the distance between the two minimum estimates of the landmark. Similarly, the maximum distance measure could be calculated. This approach is not precise, but it does not require a further search of the inspection map result.

Better results could be found be re-searching the inspection map for the minimum/maximum distance between the estimates of the two landmarks. It is possible that this combined minimum/maximum would be different from that found using the minimum/maximum found for each landmark individually. These could be found by examining all pairwise combinations of possible landmark estimates, selecting one from each landmark.

Perimeter

The perimeter parameter is based upon the estimated location of a set of landmarks. Combining this information presents similar problems to those described for the distance parameter. A minimum and maximum bound on the perimeter could be found from a search of the inspection map combining all possible combinations of landmarks estimates, selecting one from each landmark on the perimeter.

A computationally simpler estimate of the perimeter could be found from by calculating the minimum perimeter through all the minimum landmark estimates. A maximum perimeter estimate could be calculated through the maximum landmark estimates.

Area

Area estimates could be found in a similar manner. These could be based upon either an analysis of the relationships between the landmark estimates or the immediate use of the maximum and minimum error estimates of the landmark.

Volume

Volume estimates could be found in a similar manner. These could be based upon either an analysis of the relationships between the landmark estimates or the immediate use of the maximum and minimum error estimates of the landmark. A difference between this volume estimate and those of others is that the volume estimate can now be presented as a bounded estimate rather than a specific value.

Comments on Parameter Accuracy Measures

Parameter estimation rapidly presents computational problems in the estimate of errors for three dimensional objects. This is due to the uncertainty in the location of the landmarks which define the parameters. The computationally simpler parameter estimates provide an estimate of the parameter without the additional computational cost.

4.7 Implications of Accuracy Analysis for Backprojection Inspection

Calibration

The amount of accuracy which can be achieved during reconstruction is limited by the quality of the calibration used for reconstruction. There are two sources of error in calibration. The first of these is an inaccurate calculation of camera parameters. The second is an error in the alignment of the object/camera system.

In the case of camera parameter error, we have a statistical description of the focal length/pixel size ratio but no statistical knowledge of the accuracy of the optical center calculation. In the case of the object/camera alignment we don't know how well the alignment succeeded. Errors in either case will show up as errors in the inspection.

Inspection Accuracy

The purpose of our inspection is to analyze an unknown object which may have defects. This might be the case of a cast part with rough edges or an assembled part which was miss assembled. In any case we have no knowledge of the statistical nature of the errors. In a sense, the inspection procedure is being used to calculate data on the errors in the case of the part in question.

There are two limits on our ability to measure these errors. The first is our representation of the ideal object. We can not measure more accurately than we know the correct answer. The second bound on the error is determined by the selection of the inspection system. This is comprised of the camera and the view locations.

Within limits we can balance parameters to give an acceptable inspection system. For instance, a camera with better sensor resolution would require fewer views to achieve a constant error bound.

CHAPTER 5

VIEW PLANNING FOR BACKPROJECTION INSPECTION ANALYSIS

5.1 Introduction

The number of views processed to obtain the reconstruction is directly related to computational intensity. It is desirable to obtain the best reconstruction possible from the fewest views. It is the goal of a view selection algorithm to remove the redundancy of information in the views and to insure that a sufficient number of views are used such that the fewest number of views are used to describe the object.

We shall refer to this process as view pattern analysis. Section 5.2 will provide motivation for this analysis in relation to the inspection problem. It will also relate view selection to the more general backprojection reconstruction problem and robot vision planning systems. Section 5.3 will present a number of different aspects of the view selection problem. These areas will show different constraints which require more views or allow views to be eliminated. Section 5.4 will discuss variability in the view position, and view combination. Section 5.5 presents the View Pattern Analysis algorithm in two parts, a preprocessing section where the information about view interaction and variation is determined, and a view selection section where the views are selected based upon the preprocessing. Section 5.6 discusses an implementation of this algorithm and Section 5.7 discusses the implications of this algorithm for backprojection inspection and backprojection reconstruction.

5.2 Motivation for View Pattern Analysis

Thus far in our examination of backprojection inspection we have examined a characterization of objects which we can inspect, presented an algorithm for the inspection, and analyzed the accuracy considerations of the inspection. We shall now examine how to acquire the data for the inspection analysis. In any analysis procedure, the important consideration is collecting a sufficient quantity of data to perform the analysis. In three dimensional computer vision inspection, it is important to limit the collection to only the data which is necessary to perform the analysis.

The general inspection process is one that will be repeated over and over again. The processing of view data will be computationally intensive because of large amounts of data. View selection allows the amount of data processed during each inspection to be minimized by performing some analysis which is common to all inspections before any of the inspections. This a priori limiting of data collection for backprojection inspection is called view pattern analysis or view planning.

The problem of view planning is essentially one of selecting the views which are most appropriate for the inspection problem. First, the existence of such views comes from the object characterization of Chapter 2, assuming an appropriate object has been selected. All objects to be inspected must possess this characterization which allows them to be inspected. Second, we are limited by the accuracy considerations presented in Chapter 4. These express the interaction of the different system parameters and their effect on the accuracy of the inspection. Views need to be selected in a manner which is consistent with accuracy requirements of the inspection. The resolution of the information contained in these views will limit the accuracy which can be obtained from the inspection. Underlying this process is the need to select data which is useful for the inspection algorithm of Chapter 3.

The view planning process itself requires data for analysis. The view planning algorithm must be based upon available data which is consistent with the inspection goals. The process should be driven from available data with minor hand analysis. This data will have to contain a description of the object in sufficient detail to allow the selection of views for inspection. One obvious starting point would be CAD information.

There are two advantages to the acquiring of the information from a CAD database. The first advantage is that it allows the designer to select landmarks during design and to have the inspection process follow from these specifications. The second is that it provides the designer feedback on his design from the standpoint of its ability to be inspected. CAD driven machine vision was discussed by [47]. Several have investigated the properties of view selection. In some cases, the views were chosen sequentially [27, 55], and sometime all views were determined simultaneously [28].

The approach we will use is to initially over define the number of views required to reconstruct the object. This will be done by first defining the views required to properly view certain points on the object which are for some reason critical. These points will be referred to as the landmarks of the object. The points and the best views of these points will be hand selected. This forces the number of views up because each voxel to be reconstructed must have its best view in a view. This provides a description of the inspection which may have a large number of views, but which can have the maximum amount of accuracy at each landmark -- limited only by object representation.

This can be considered a reasonable approach to the process since there are usually a finite number of points on the object which are of concern during inspection. For a single point, the best views are often obvious. In fact, this initial selection of a set of best views of a landmark by the user will allow the user to select an accuracy bound for each landmark.

Once the initial data has been extracted from the CAD information and determined by the user, the view planning process may proceed. The general technique is one of identifying the views which contain the most redundancy and eliminating them one at a time. This has the effect of reducing both the number of views and the accuracy. Accuracy considerations from Chapter 4 provide a way to measure the interaction between the number of views, their location, and the accuracy of the inspection. Increased accuracy will require views which are closer to the object which would in turn require a greater number of views to observe the entire object.

Other areas of three dimensional computer vision can benefit from similar view planning analysis. Closely associated with backprojection inspection is the backprojection reconstruction which is performed on each voxel to be inspected. This technique is also used to perform whole object reconstruction. Another related areas is robot vision systems for scene viewing and motion.

5.2.1 Backprojection Reconstruction

Backprojection reconstruction techniques are a central tool in the inspection technique presented in Chapter 3 to reconstruct specific locations on the surface of three dimensional objects. These techniques are also used to reconstruct entire objects. In this case the emphasis is not on the accuracy of specific locations on the object examined, but on an overall accuracy of the object reconstruction. In either case, appropriate a priori selection of views can reduce the amount of processing which will be required to perform a reconstruction.

In the overall backprojection reconstruction case, views would need to be selected to reconstruct the entire surface of the object. This could be done by selecting landmarks distributed over the surface of the object. As a limiting case, all object boundary points could be selected, with the view selection algorithm removing the redundancy. While this may seem like the easiest solution, it is possible that this could still result in over processing during reconstruction.

5.2.2 Other View Planning Areas

There are two other main areas which have used planning in their analysis: data representation and a planning application area. The data representation technique is aspect graphs. Aspect graphs provide a description of an object under analysis. An application area for this is robot motion planning. Robot motion planning requires the viewing of a three dimensional scene within which the robot will move.

Aspect Graphs

Aspect graphs provide a compact information set for CAD information. Aspect graphs show relative relationships between viewing locations of the faces of an object. Aspect graphs have been used in several CAD applications. They have been used for vision sensing strategies [47, 54, 55].

Aspect graphs may be thought of a description of the space surrounding an object. Assume that the surface of an object is made up of a set of plane surface sections. If we divide 3 space up into regions by extending these plane sections as planes we end up with a set of closed and unclosed regions. We can construct the object's aspect graph by making each region a vertex of the graph and connecting those vertexes which are adjacent in space. Figure 5.1 shows an example of an aspect graph for a figure confined to a plane.

Aspect graphs provide a good measure of the general relationship between location and viewable features. They do not provide a means to describe how well a feature may be seen from a given location. A technique similar to aspect graphs is used in Chapter 3 to determine the projection of the vertices of each voxel.

Robot Motion Planning

Robot motion planning is the process of selecting a successive set of coordinate frames for a robot to move through. This could be for a mobile robot or a robot arm moving within a work cell. In either case, this requires that the scene be analyzed for collision avoidance. It requires the knowledge of locations and surfaces within an scene which must be avoided.

Once again, inspection accuracy is not required for this analysis. It is only necessary to know the location of objects and surfaces with enough accuracy to insure avoiding a collision. Views which would be

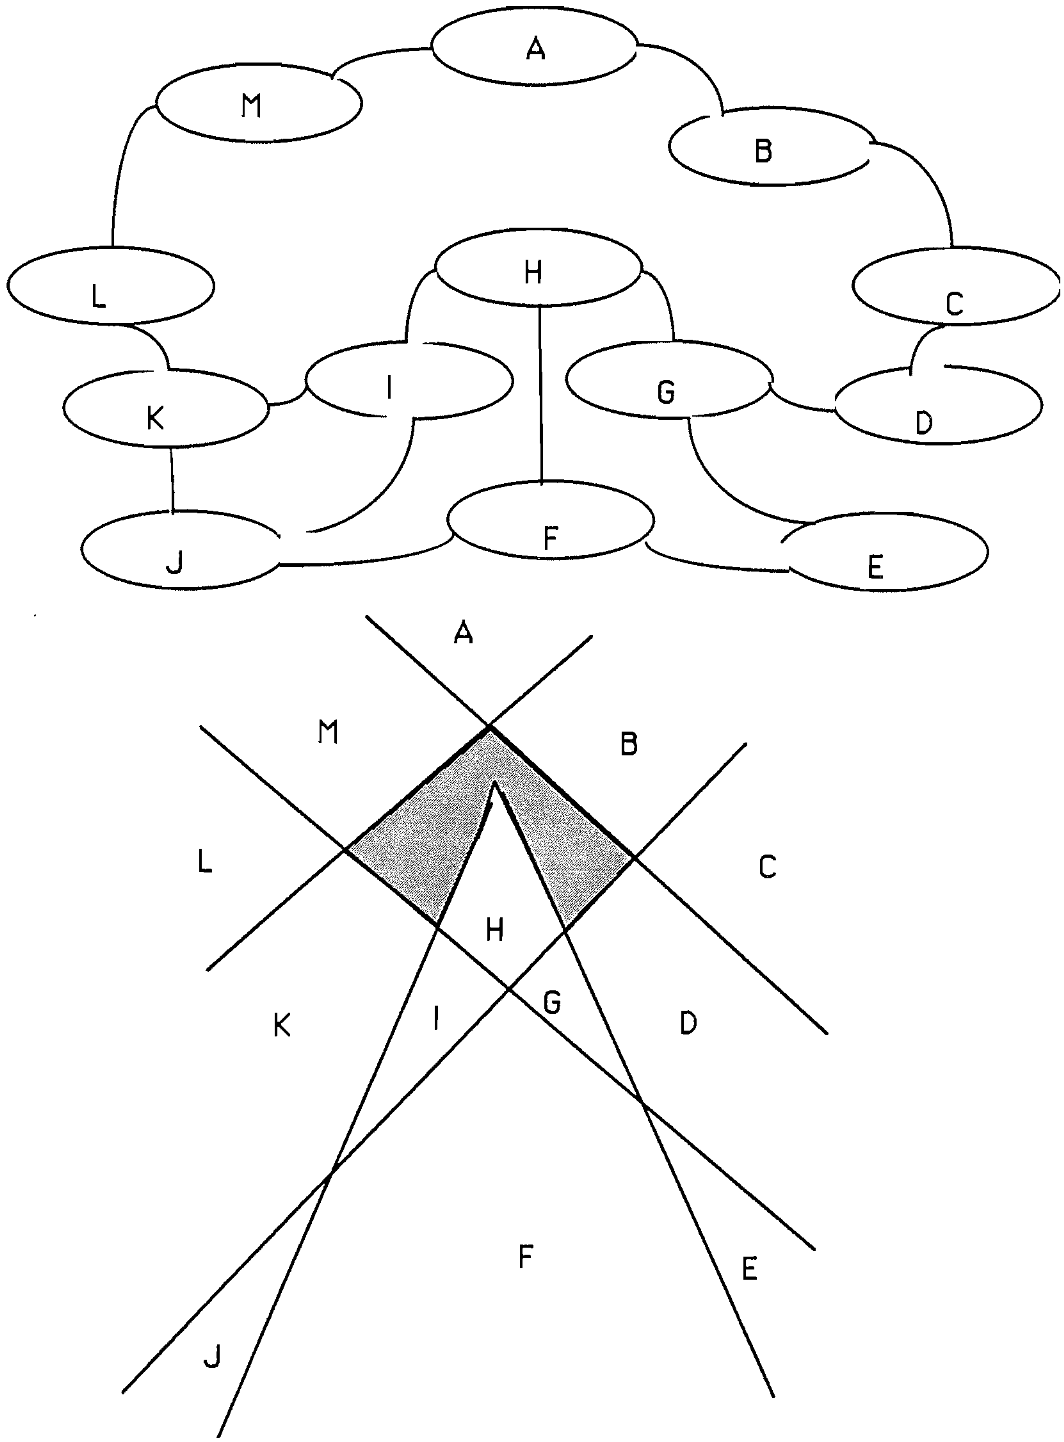

Figure 5.1 Aspect Graphs

useful in determining where obstructions are located do not necessarily provide information concerning how well a specific landmark can be seen.

5.3 Aspects of View Pattern Analysis

Several different areas impact or restrict our ability to perform view pattern analysis. Each of these areas shall be discussed separately before introduction of the view pattern analysis algorithm. Each area can be categorized as providing the view selection algorithm with object information or providing view combination rules. Finally, we shall discuss a method of combining all this information into the process of view pattern analysis.

The areas providing information about the an object include the use of landmarks as search goals for our restricted set of views, the use of best views of landmarks to indicate the needed view information, the octree inspection map for providing the requirements of the inspection, and a representation map of the object for a complete object description.

View combination rules provide information about how the views can be combined and how the inspection process makes use of information during an inspection. These include the 50% rule for the determination of a voxel's identity, the use of obstruction checking to determine variation possibilities in the best views, the effects of the inspection system itself to limit the view selection. In Section 5.4, some of these rules are combined to determine the variation possibilities of the best views such that the equivalent information is provided by the new view.

5.3.1 Landmarks as Goals

One of the inputs of the algorithm is a set of landmark locations on the object to be inspected. Any accuracy calculation is dependent on significant landmarks of the object. Figure 5.2 show an object with landmarks indicated. In this instance, the object can be thought of as a

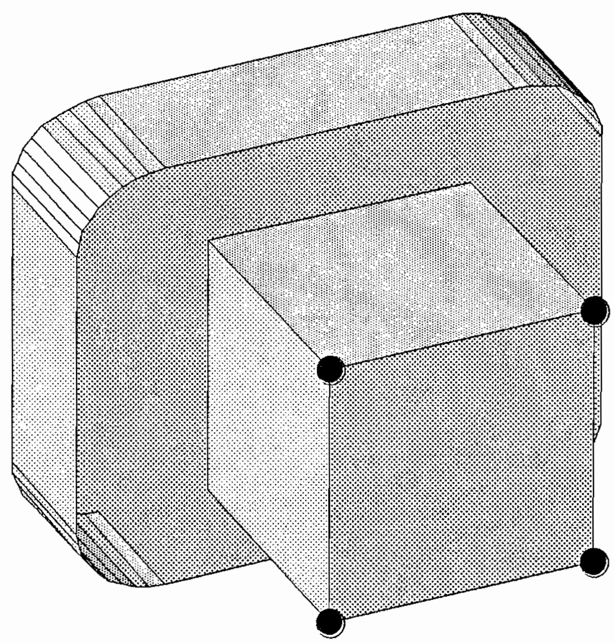

Figure 5.2 Landmarks for View Selection

stamp whose critical parameters are at the end which will make an impression.

While it might be possible for automated selection of landmarks, it might not be desirable. Automated selection of landmarks assumes that all "significant" features of the object are significant for the determination of reconstruction accuracy. As can be seen from the example just cited, this may not always the case. In most instances, we are concerned with computational intensity, and overspecification of landmarks is something to be avoided.

Only the object designer, or another designer who must use the object, can specify what the significant locations are on the object. The use of landmarks specified by the designer allows him to be involved in the inspection configuration as well. This enhances the integration of the object's design phase with the manufacturing phase of development.

5.3.2 Best Views of Landmarks

In addition to the specification of landmarks by the designer, we will also presuppose that he has specified a set of "best views" of each of these landmarks. These best views specify a set of optical axis view angles which are required to determine each landmark. A single best view could be described by a set of best sub views, any one of which provides the information of the "single" best view.

Best View A = (best sub view 1) or (best sub view 2) or ... (5.3.2-1)

This equation emphasizes that in some instances a single best view can be described by any one of several views. It is in part for this reason that the designer must describe the object by a set of best views for each landmark. Each landmark will be described by the set of best views, all of which must be present to describe the landmark

Landmark Observed = (Best View A) and (Best View B) and ... (5.3.2-2)

This description of the landmarks by best views is consistent with *Weak Convexity* of Chapter 2. *Weak Convexity* guarantees that a set of best views can be found for each landmark. Figure 5.3 shows an example of best view description of a landmark.

The best views of a landmarks must be determined a priori to the view pattern analysis. This determination provides a set of views which are known to properly describe the single specific landmark on the object. The best views are along the optical axis of a view. In most instances this provides for an overdescription of the landmarks since it is probable that more than one landmark is visible from a given view. This will be determined by the parameters of Sections 5.3.4-7.

5.3.3 Octree Inspection and Representation Maps

The octree inspection and representation maps provide descriptions of the object for the process of view pattern analysis. The octree inspection map provides a description of the landmarks of the object. The representation map provides a description of the object for obstruction checking.

The octree inspection map, developed in Chapter 3, provides all the information the inspection algorithm needs to perform the inspection. It also provides view pattern analysis with a description of the inspection requirements for view selection. The requirements provided are a listing of the landmarks, their locations on the object and the maximum resolution of the inspection process. Figure 5.4 shows an example of an octree inspection map.

The representation map is a complete description of the object in a three dimensional array. This array represents an objects location by object and non-object entries. The array resolution is equal to that of the octree inspection map. This could be a large array and may not be appropriate for any online inspection processing. It is only used during the preprocessing phase of view pattern analysis to check for obstructions. Figure 5.5 shows an example of the representation map.

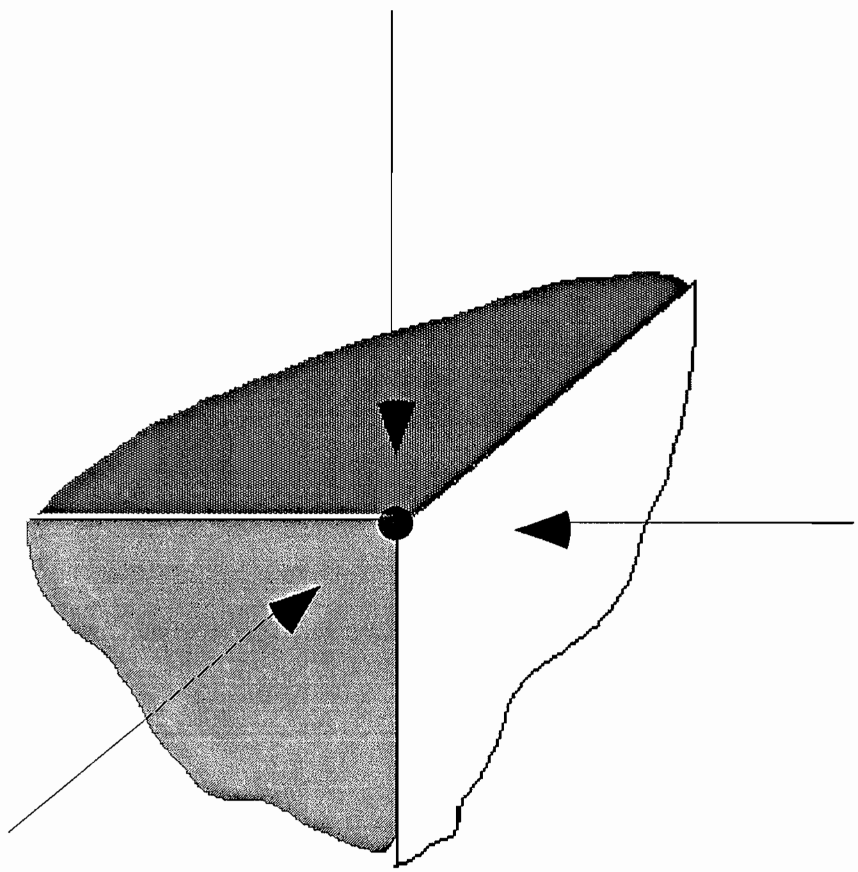

Figure 5.3 Best views of a landmark determined a priori

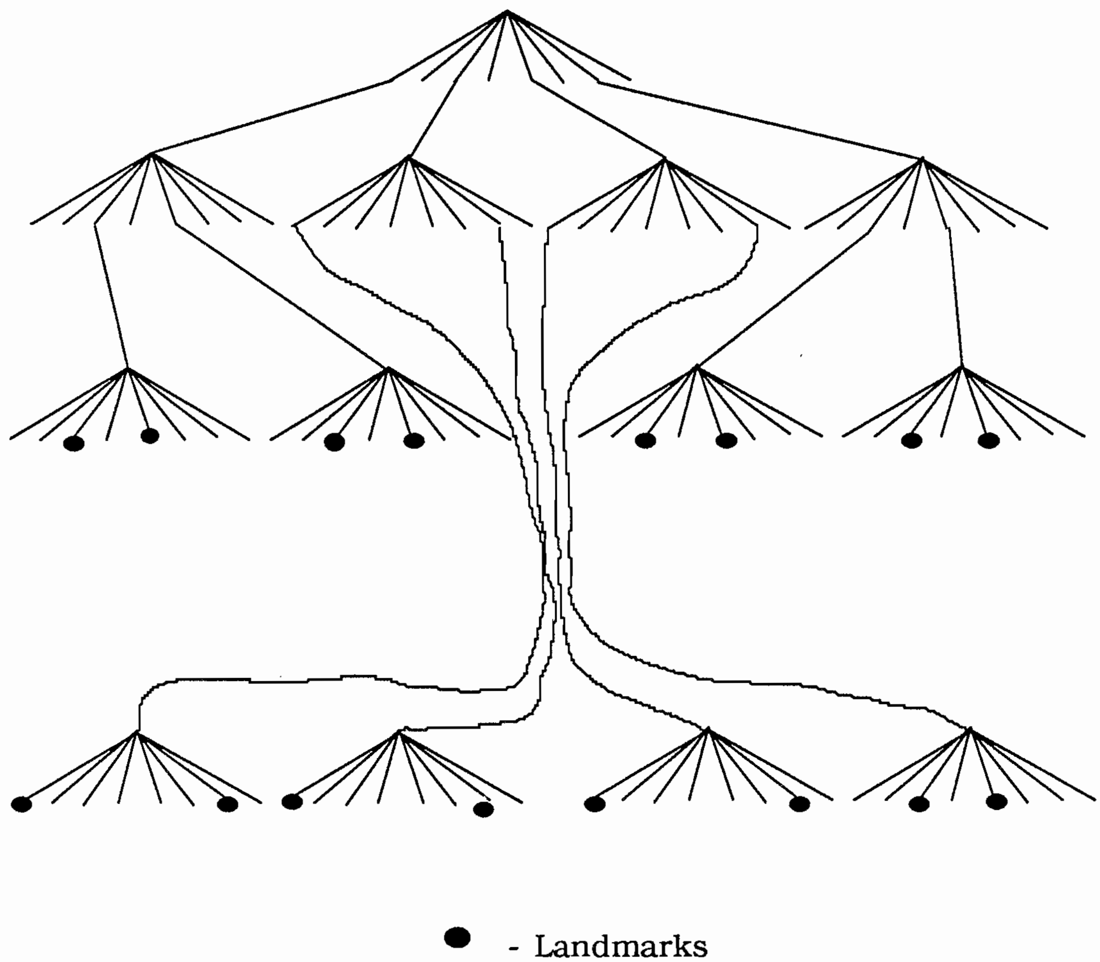

Figure 5.4 The octree inspection map with modifications for view planning

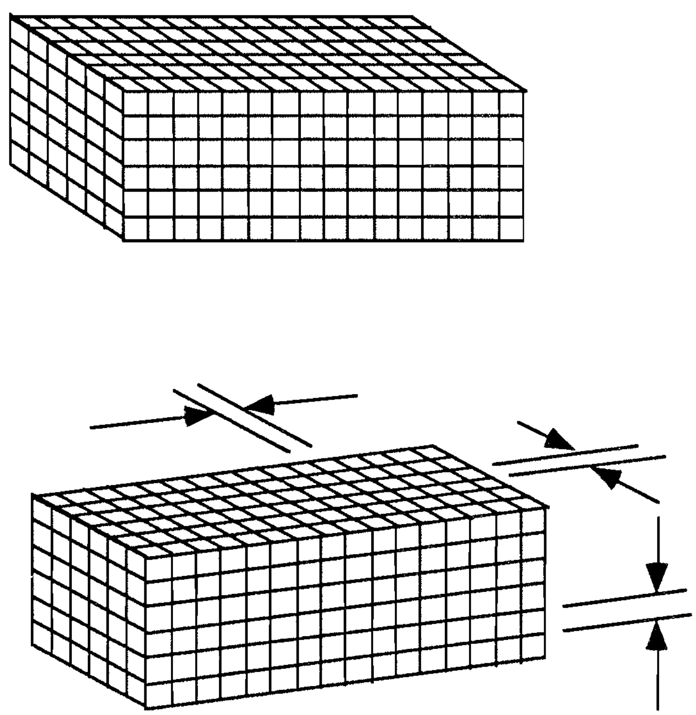

Figure 5.5 The representation map with modifications for view planning

5.3.4 50% Rule for the Determination of a Voxel

The representation scheme used for the inspection is an octree inspection map. This map is discrete. During the volume source backprojection reconstruction of the individual discrete voxels of the octree decisions must be made which discretize the inspection process. This decision process led to the 50% rule discussed in Chapter 4.

The 50% rule for voxel determination provides a measure of a voxels ability to distinguish between background and void in the determination of a voxels identity. It states that any voxel which has a projection of 50% or greater of void onto any single view during inspection must in fact be at least 50% void itself. Figure 5.6 shows an application of the 50% rule.

This provides our first rule for view pattern analysis.

<u>RULE 1</u> Valid views of an object voxel which is at least 50% void must always project to at least 50% void.

5.3.5 Obstruction Checking

Obstruction checking involves insuring that views which are selected can observe the landmarks of interest on the object. Obviously, any initial best view is able to observe its associated landmark. During the view selection process other views will probably be selected to observe a given landmark. Obstruction checking insures that the new view can observe the landmark.

There are two types of obstructions which may occur. Both obstructions restrict the view positions by disallowing certain positions and orientations. Certain bounds limiting the view selection may not be directly related to the object. These bounds are related to the universe of the reconstruction environment. Figure 5.7 shows examples of these environmental restrictions. Other bounds limiting the view selection are related to the object. These types of object limitations are discussed in Chapter 2 and examples are shown in Figure 5.8. It is necessary to check each view modification to insure that the new view chosen is not in

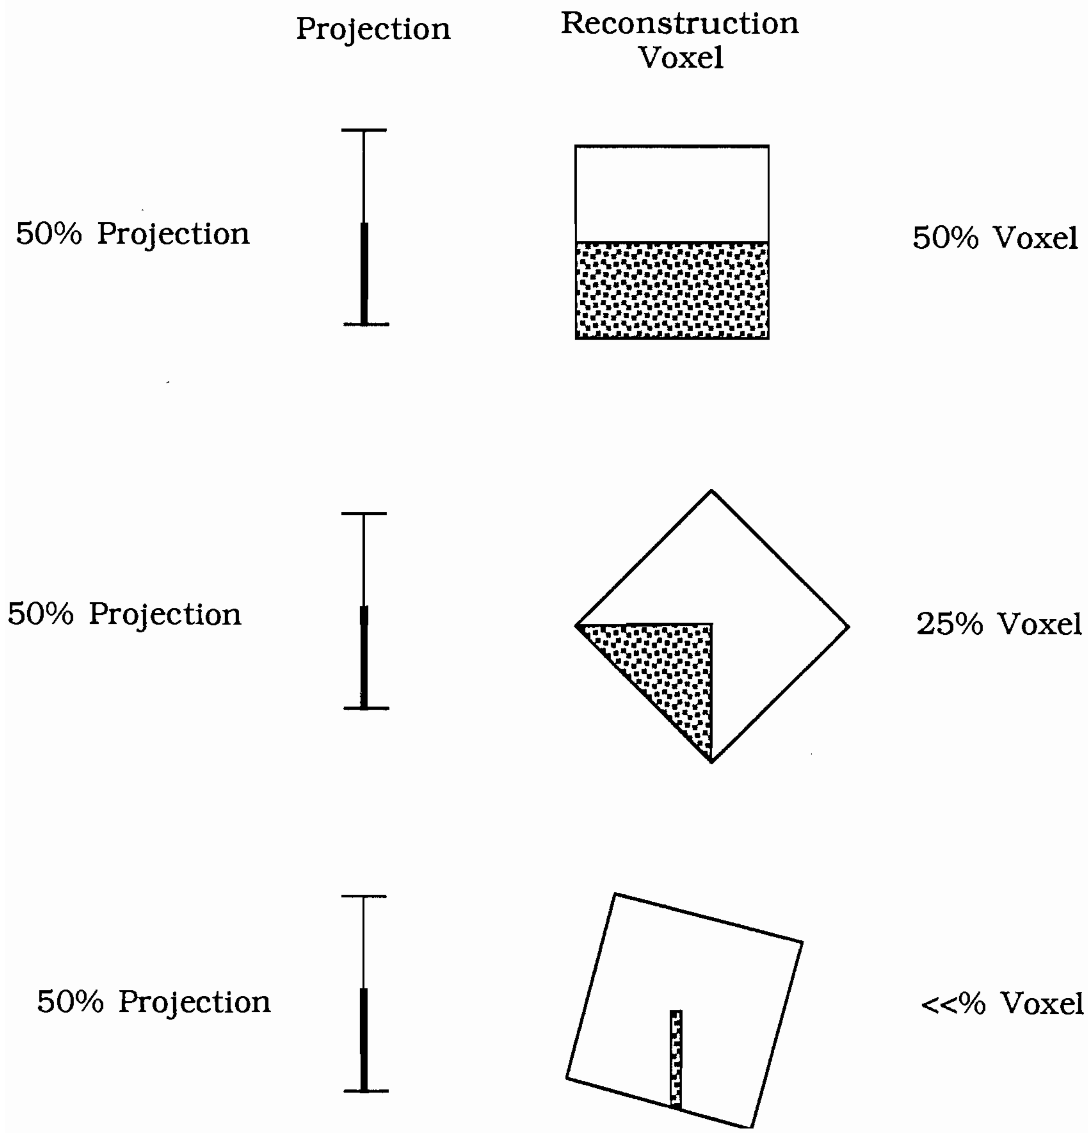

Figure 5.6 Voxel showing the 50% rule for Voxel identity determination

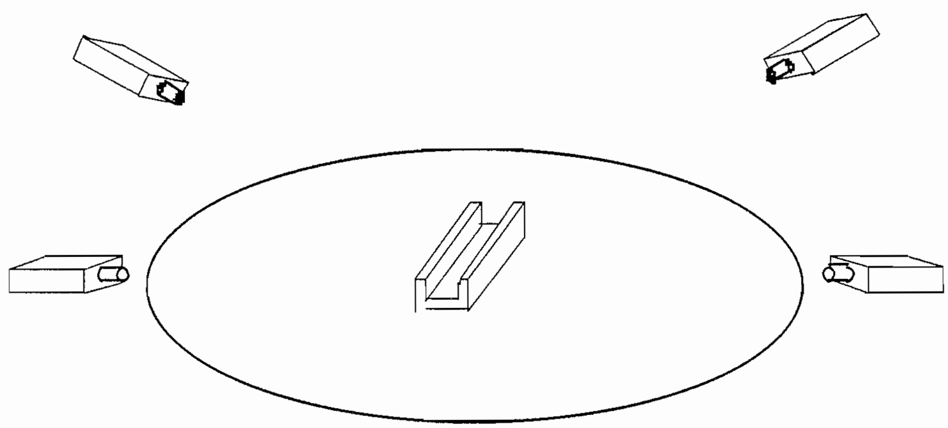

Figure 5.7 Environmental limitations on views

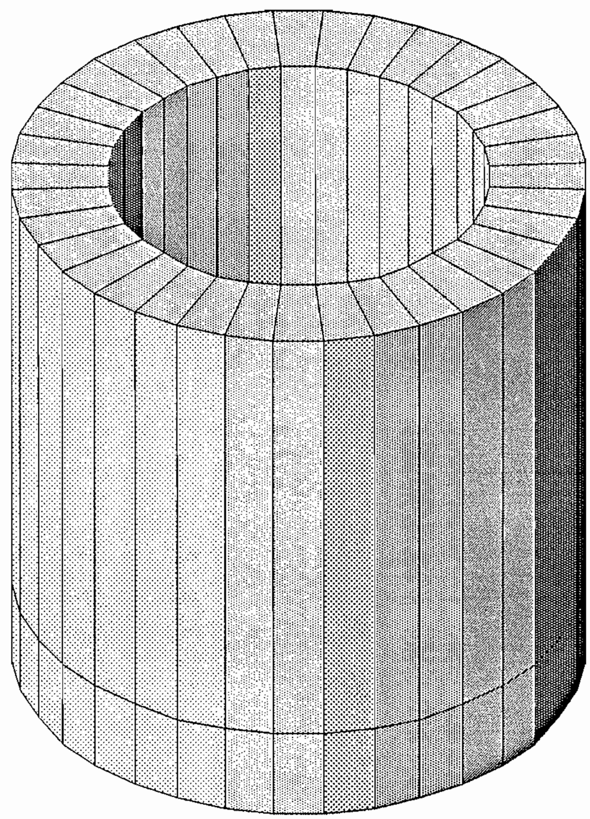

Figure 5.8 Object limitations on views

some way blocked by the object. While Weak Convexity guarantees that view locations exist that will distinguish a landmark, there are views which will not. Obstruction checking insures that the views selected maintain the observability of the landmark.

RULE 2 Valid views of an object voxel must not be obstructed by the reconstruction environment or the object.

5.3.6 Accuracy and Inspection System Parameter Effects

The accuracy analysis of Chapter 4 indicated a close interaction of the parameters of Δp_v, Δp_h, ΔV, d, f, Θ . Within the parameters of view pattern analysis we will assume that Δp_v, Δp_h, ΔV, f are fixed by the inspection system and are known. The factors which can be varied are d, the distance to the object along the observation axis; and Θ, the angle between the observation axis and the optical axis. These factors describe the varying of the position of the view with respect to the object under inspection.

One result from Chapter 4 was a lower bound on the minimum experimental discrimination (MED) specification. This bound relates the voxel size in the inspection representation to the MED. This will provide a useful manner to incorporate an accuracy requirement into the view selection procedure.

$$\Delta V \leq MED \qquad (5.3.6\text{-}1)$$

In the view planning case, Δp_v and Δp_h are fixed and provide a bound on the obtainable accuracy for any view selection. Assuming that the accuracy needed satisfies equation 5.3.6-1, the MED will provide an upper bound on the measurements of Δp_v and Δp_h for the purposes of the view selection process as described in Chapter 4. This will be used later to modify our relationships to have an accuracy bound instead of a pixel size bound.

Chapter 4 examined the inspection system parameters for their effect on accuracy. The limiting accuracy case occurs when the observation line provides the inspection system with the largest possible projection. This observation line can always be chosen. The inspection system must be able to detect the smallest dimension of this projection. This leads to a observation line near one of the major diagonals of the voxel.

This projection is not symmetrical. For simplicity, it was assumed to be symmetrical about the smaller projection. Figure 5.9 shows a cross section of this projection relating the parameters shown above to it. Chapter 4 concludes with an expression which relates all the parameters shown in this figure.

$$\Delta p_{v,h} < 2\left(\frac{f}{d}\right)\left(\frac{\left(\Delta V/2\sqrt{2}\right) d(\cos\Theta - \sin\Theta)}{d\cos\Theta - \Delta V/2\sqrt{2}}\right) \qquad (5.3.6\text{-}2)$$

A satisfactory solution of these equations is necessary whenever an inspection is to be performed and is thus a requirement for a view selected during view planning. As noted earlier, this can be modified into a bound based upon the needed accuracy of the view selection.

$$\Delta p_{v,h} < 2\left(\frac{f}{d}\right)\left(\frac{\left(MED/2\sqrt{2}\right) d(\cos\Theta - \sin\Theta)}{d\cos\Theta - MED/2\sqrt{2}}\right) \qquad (5.3.6\text{-}3)$$

As Θ becomes small and $d >> \Delta V$, this relationship becomes:

$$\Delta p_{v,h} < 2\left(\frac{f}{d}\right)\left(\frac{MED}{2\sqrt{2}}\right) \qquad (5.3.6\text{-}4)$$

Note that in each of these cases we have defined an upper bound on the sensor size. Any sensor size smaller than this will be sufficient to determine the dimensions of the projected voxel.

An examination of this relationship reveals that as the accuracy requirements of the inspection are increased (MED reduced), the distance between the view and the object must also decrease for a fixed relationship to the pixel size and focal length. This is consistent with our

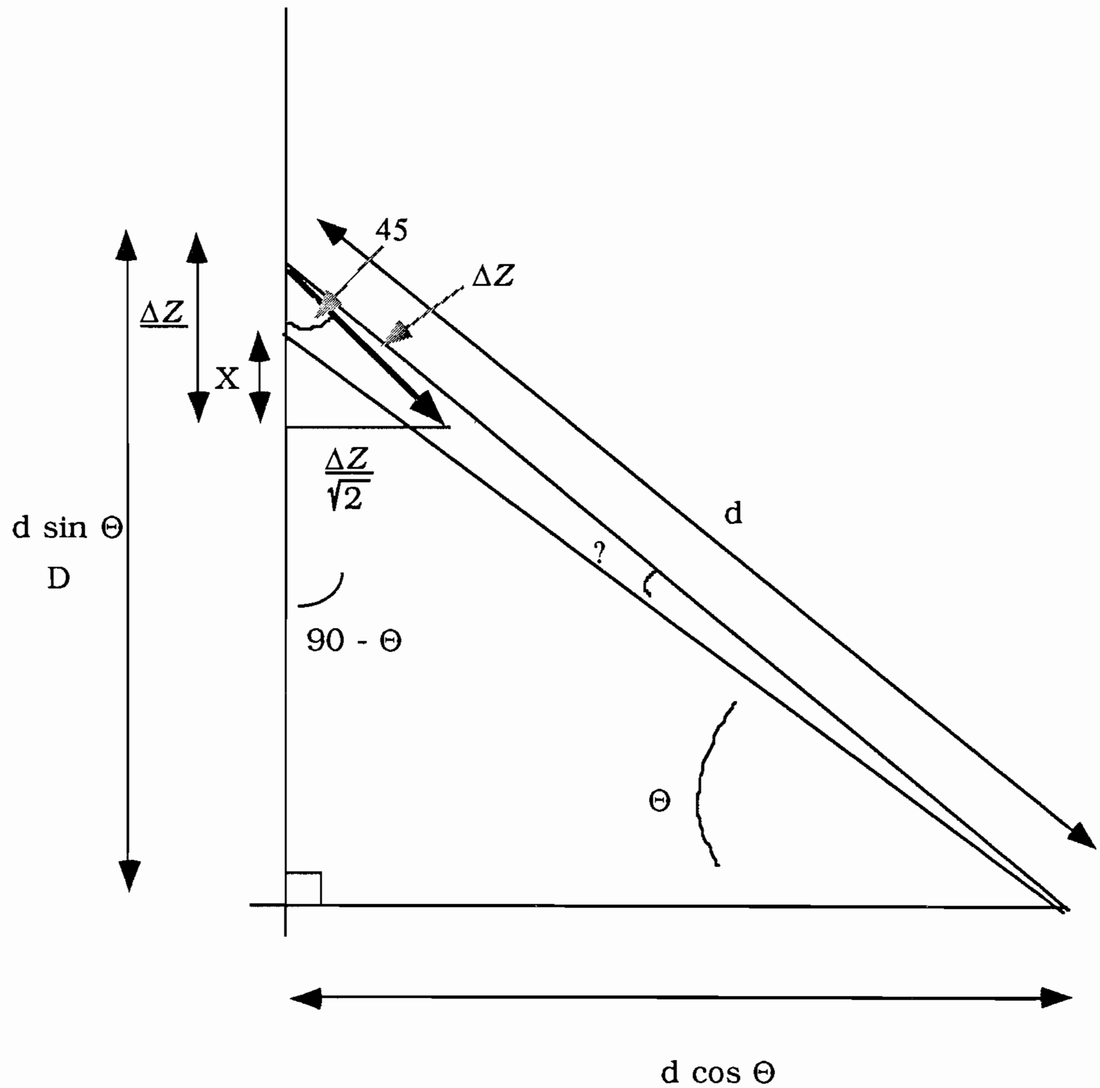

Figure 5.9 Projection of the smallest dimension of the largest projection

intuition that a more detailed analysis will require a close examination of the object.

RULE 3 Valid views of an object voxel must satisfy equation 5.3.6-3.

Another restriction on the allowable views follows directly from the camera characteristics. This relates the depth of field (DF) to the allowable distance between the view and object.

$$d_{min\ DF} \leq d \leq d_{max\ DF} \qquad (5.3.6\text{-}5)$$

This directly implies a requirement for view selection.

RULE 4 Valid views of an object voxel must be within the camera depth of field as specified by equation 5.3.6-5.

5.4 View Variation

Thus far in our discussion of the aspects of view selection we have examined specific information about the object inspection system in Sections 5.3.1-3; and rules which govern valid view determination in Sections 5.3.4-6. These sections provide information about the object inspection system and about what views we can use but they do not tell us what views to select.

Section 5.3.2 provides knowledge of certain best views of the landmarks which do provide the information which is needed for the inspection. Sections 5.3.4-6 imply that these views are not unique. In this section we shall examine the variation of a view from a know orientation for compliance with the rules of Sections 5.3.4-6. Any view thus described can provide equivalent information for the view analyzed. This will provide a set of valid views for each view which is known. These sets of views will be the basis for the view selection process to be presented.

We will begin with an examination of the single best views provided by the designer and how much these can be varied. We shall then examine a pair of landmarks viewed from a single location and the variation of this view. Finally, this shall be generalized to a view which observes more than two landmarks.

5.4.1 Best View Variation

The single best view provided by the designer can be varied in three ways. The first way is to change the range between the landmark and the view. This provides a scaling effect on the projection of the voxels in the landmark location and is governed by equation 5.3.6-3,5 with Θ fixed.

The second way this view can be varied is to translate the optical center of the view. This is shown in Figure 5.10. During this translation the view plane remains in its original orientation; only the optical center changes. This might be necessary if there were two best views which were parallel to each other and were combined. The limiting factor on the translation is the angle Θ. This angle determines the projection angle of the voxel onto the view and is limited by equation 5.3.6-3 for a fixed distance D. We can find the maximum of this type of translation.

$$C = D \tan \Theta \qquad (5.4.1\text{-}1)$$

A third way that this view can be varied is a rotation of the view plane. This is shown in Figure 5.11. In this case the optical center remains in the original view plane but the view plane itself is rotated such that the limiting case is when the optical axis of the view is at an angle Θ to the original view plane. This is once again limited by equation 5.3.6-3 for a fixed distance D. We can find the maximum of this type of translation.

$$C = D \tan \Theta \qquad (5.4.1\text{-}2)$$

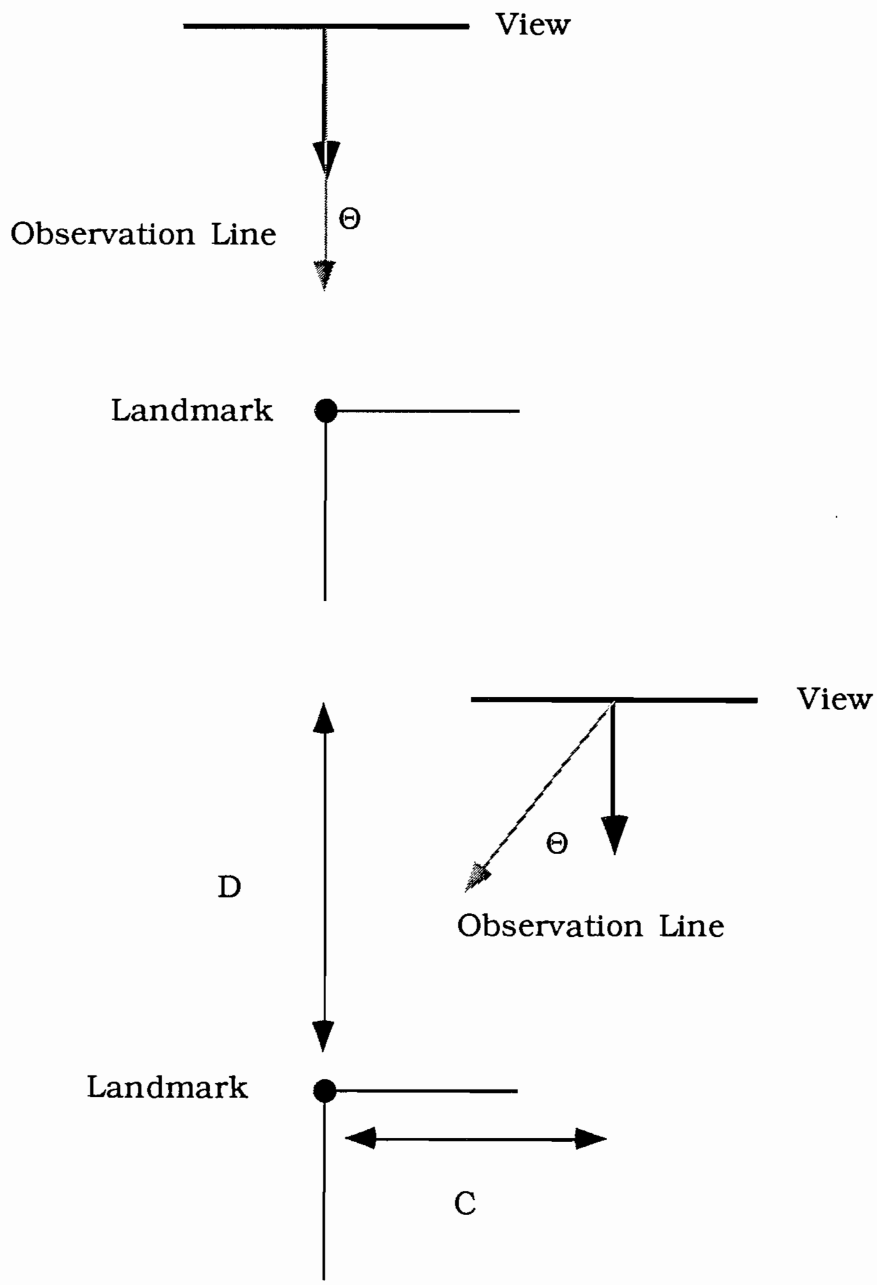

Figure 5.10 Angle effects of system parameters

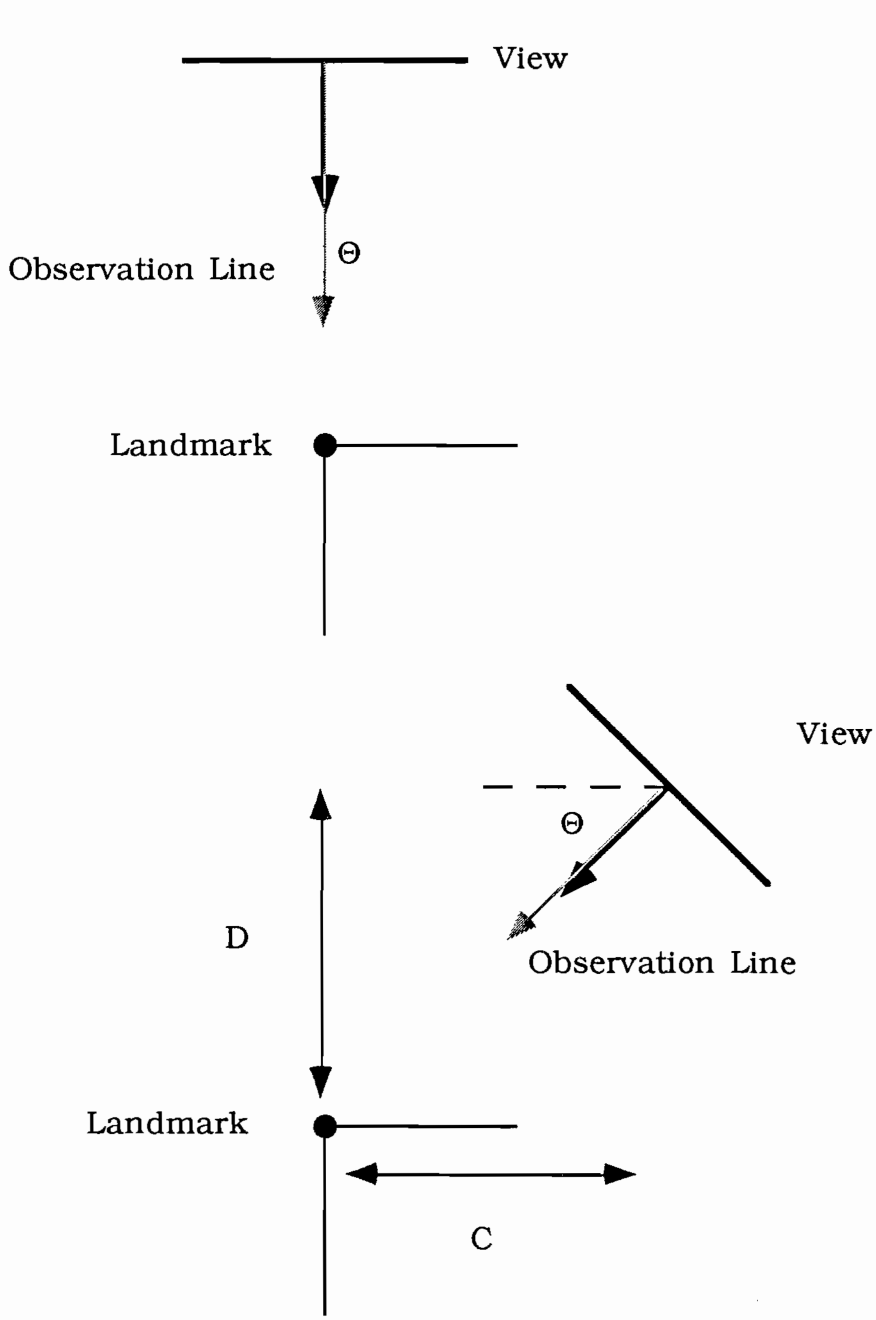

Figure 5.11 Effects of optical center displacement in view plane for a single best view of one landmark

5.4.2 Two Landmark View Variation

A single view which observes two landmarks can be moved in a similar manner to the movement of the best view above. The difference is that the symmetry of the single landmark case is lost. We shall examine the movement in a plane determined by the optical center of the view, and the two landmarks.

We will first examine movement of the optical center while the view plane must remain parallel to the original view plane. Figure 5.12 shows the original view in the center of the page and movement of the optical center in the four views surrounding it. We have assumed that the optical axis bisects the distance between the two landmarks for simplicity. If this is not the case some of the asymmetry which we will discuss later will be present in this simple case.

As the view plane is moved toward the object, the angles between the optical axis and the observation lines grow to a limiting value of Θ. Beyond this point the view can not be moved closer to the object. As the view plane is moved away from the object the angles become smaller. There is no limiting case for movement in this direction due to angles. Movement away from the object is limited by equations 5.3.6-3,5 because of either the camera depth of field or accuracy limitations.

Movement of the view plane to the "left" or "right" with respect to the figure is limited by the view angle Θ reaching a maximum. Figure 5.13 shows a cross section of this type of translation. In the figure the more general case of an angle of $b < \Theta$ between the observation line and the optical axis is shown to be moved to the maximum position at an angle Θ. The distance B is the maximum distance which the optical center may be translated and still observe the landmark.

$$B = A\left(\frac{\sin(\Theta - b)}{\cos\Theta \cos b}\right) \tag{5.4.2-1}$$

Note that as the angle b >> 0 we have the same equation as for the single best view. Thus this equation could describe both cases in our determination of view variation.

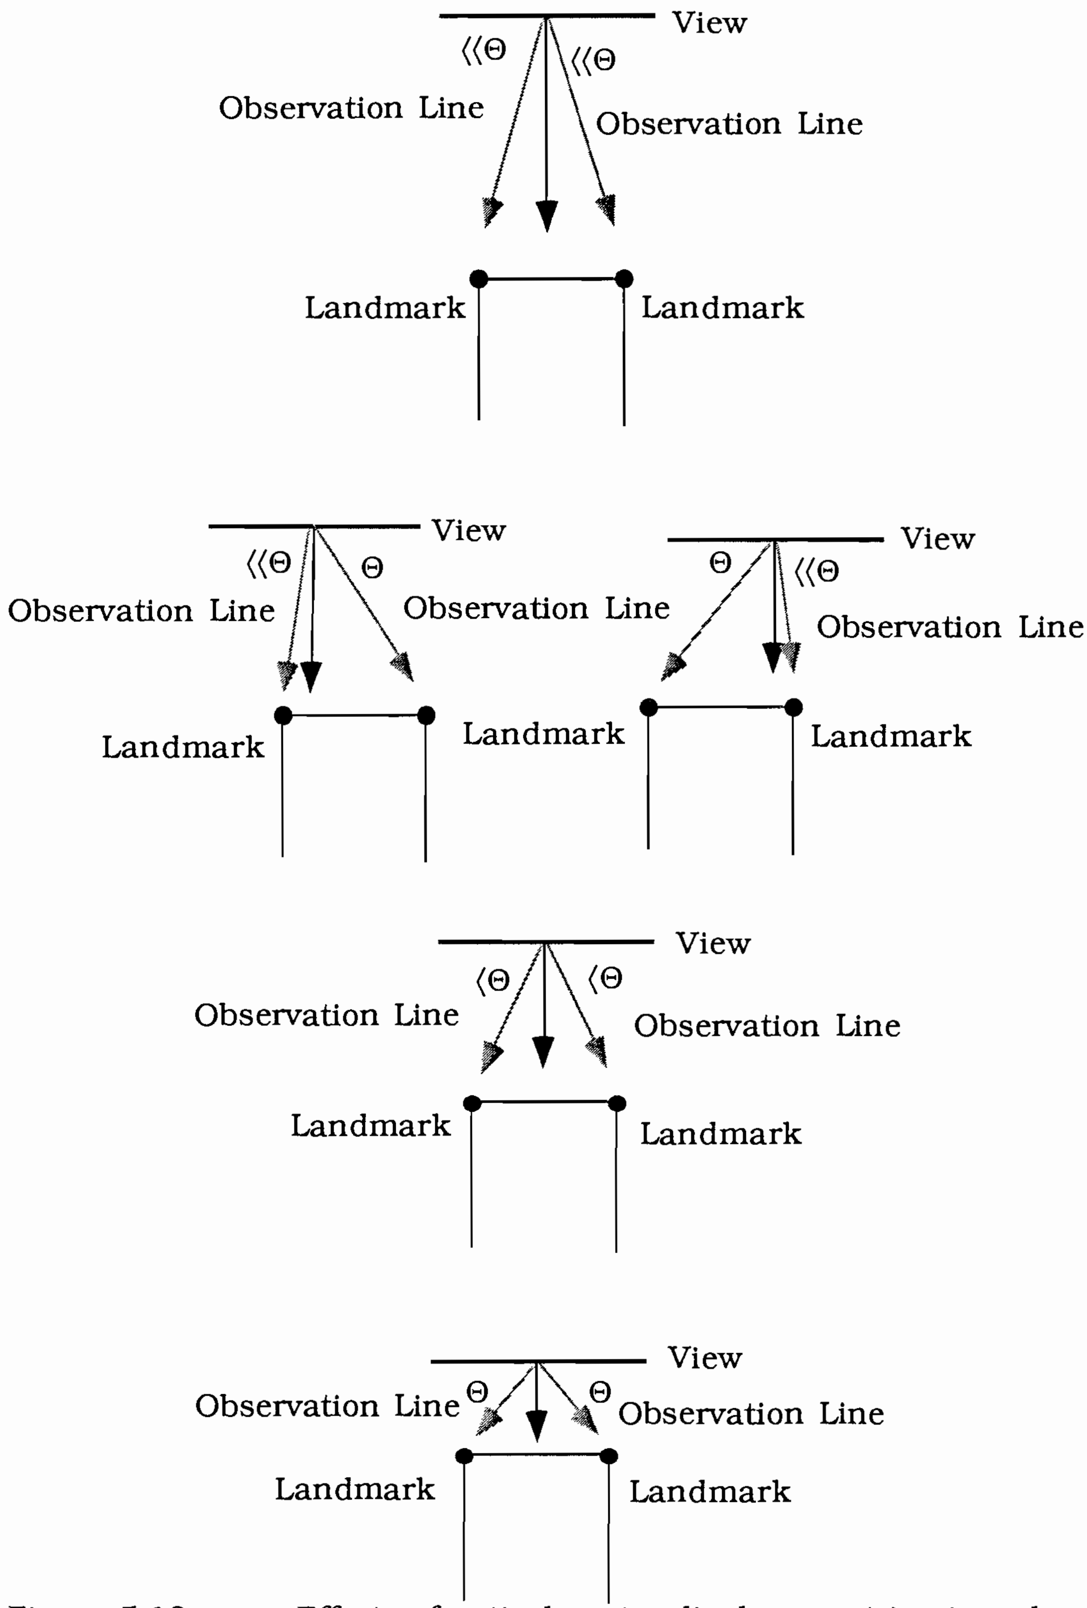

Figure 5.12 Effects of optical center displacement in view plane for two landmarks

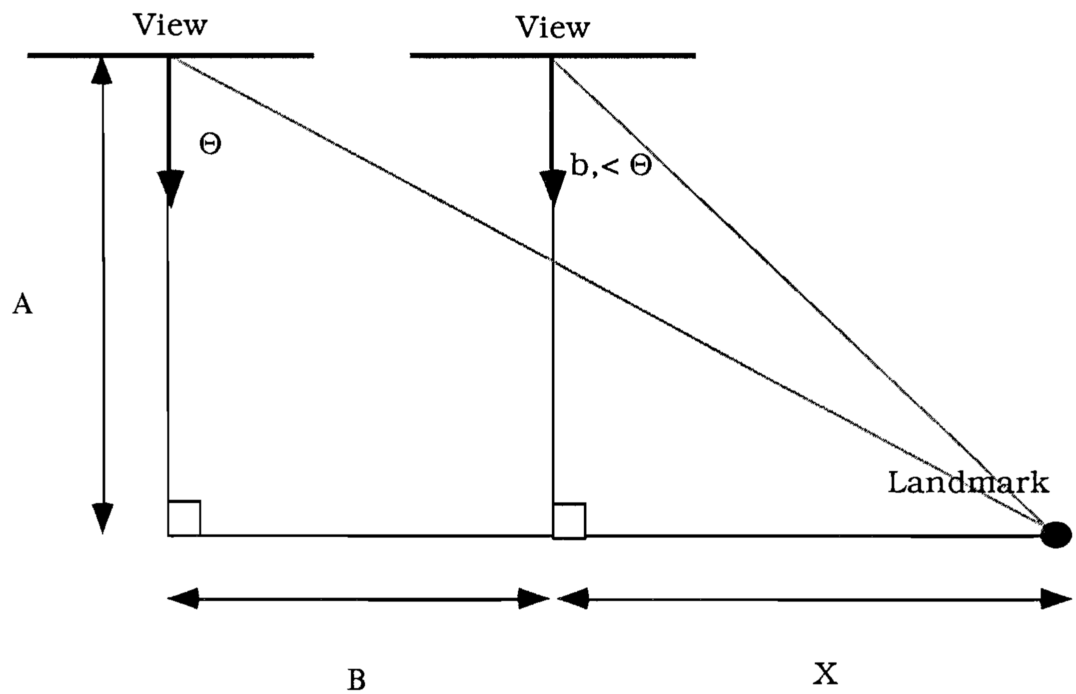

Figure 5.13 Variation about a single landmark from a best view

The other way in which the optical centers orientation with respect to the two landmarks can be varied is by a rotation of the view plane during these same translations just discussed. These translations and rotations are shown in Figure 5.14. They are similar to those of Figure 5.12 except for the rotation. The view plane could be rotated in either direction but only one direction is shown for simplicity. This consideration is valid because of symmetry considerations of the rotation.

As the optical center is moved closer to the landmarks one of the two angles between the optical axis and the observation line will reach a maximum of Θ. This provides a limit on the translation toward the object. As the optical center is moved away from the landmarks the angles diminish without bound as before. Once again, the limiting factor is the depth of field and accuracy considerations as defined by equations 5.3.6-3,5.

The asymmetry of translation to the "left" and to the "right" precludes a common treatment as done for the earlier case. Translation to the "left" is shown in Figure 5.15 and translation to the "right" is shown in Figure 5.16. An evaluation of the limiting case for translation of the optical center can be determined in each of these cases respectively.

$$B = A\left(\frac{\sin(\Theta - a)}{\cos(\Theta + \Psi)\cos(a + \Psi)}\right) \tag{5.4.2-2}$$

$$B = A\left(\frac{\sin(\Theta - b)}{\cos(\Theta - \Psi)\cos(b - \Psi)}\right) \tag{5.4.2-3}$$

Note that once again that as $\Psi >> 0$ in each of these equations we get the equation (5.4.1-1). This would allow for the general use of these equations, but the asymmetry will preclude the usefulness of this generalization.

Each of these equations helps to describe an overall limitation on allowable locations for the optical center with respect to the two landmarks which are observed with the plane to which we confined our

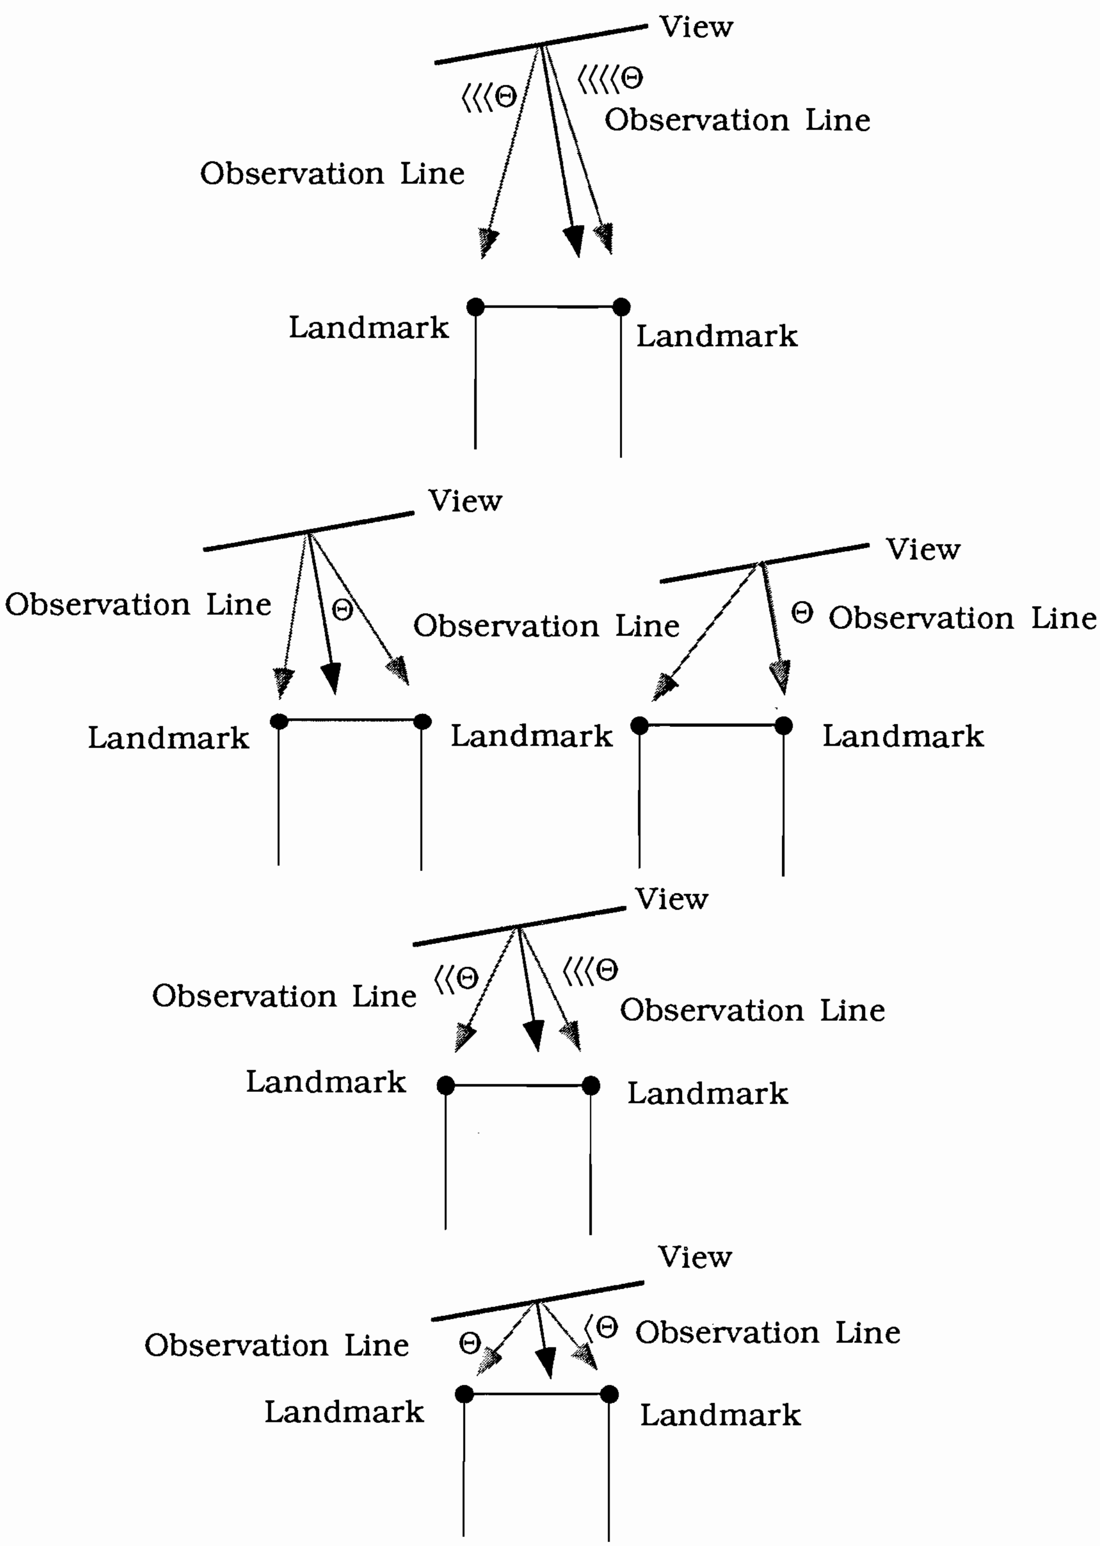

Figure 5.14 Effects of rotation of optical center about a landmark for a single best view of one landmark

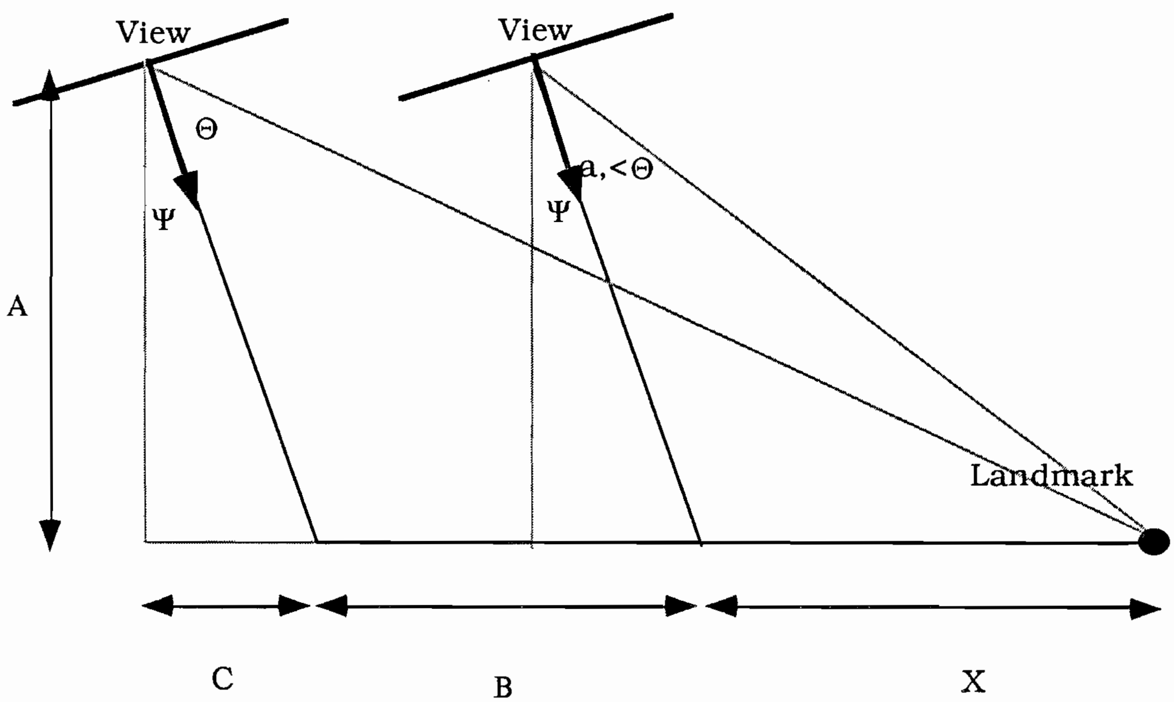

Figure 5.15 Effects of view plane rotation with respect to two landmarks

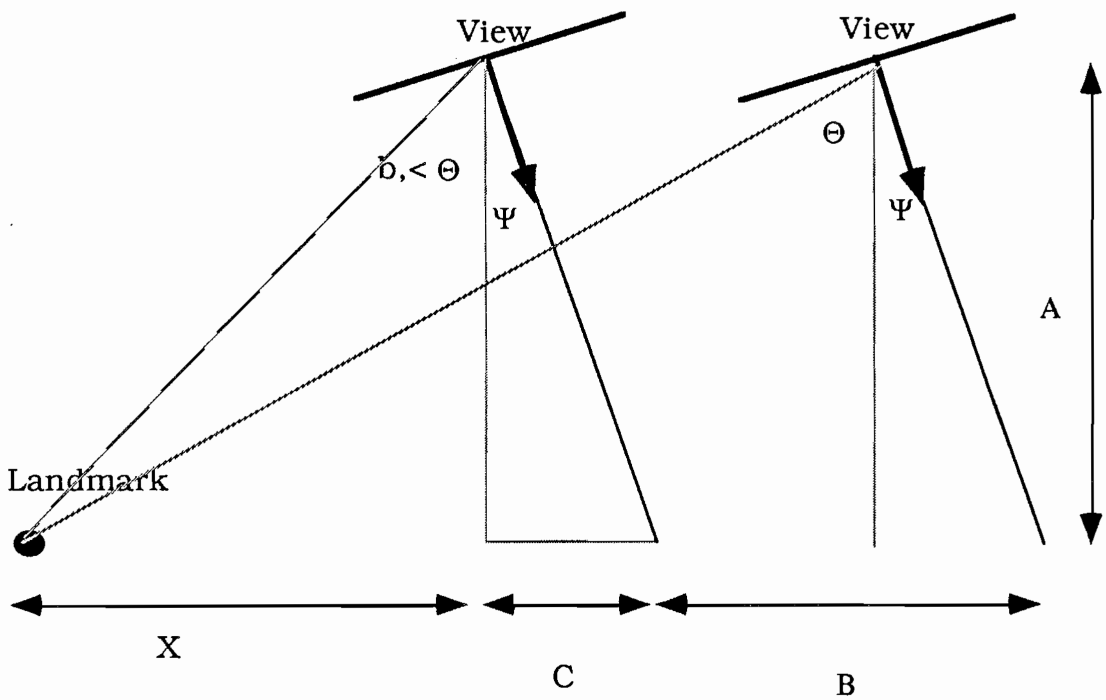

Figure 5.16 Displacement effects toward one landmark of view plane rotation with respect to two landmarks

variation. We have not addressed the variation of the optical center in locations which are not in this plane. As the optical center moves out of this plane the view angle Θ will always be larger than for its projection into the plane we are using. From this we can conclude that translation in a plane parallel to one we are using will always be less extensive than the plane just analyzed. Figure 5.17 shows this effect using plane slices. As we shall see in the next section, the introduction of more landmarks to be observed by a single optical center will increase the complexity further.

5.4.3 N Landmark View Variation

When there are more than two landmarks that are observed by a single view the valid locations are a function of the angles from all N landmarks in a similar manner to the way the valid locations were a function of the two landmarks analyzed above. The result is a volume which describes possible locations for the optical center. The complexity of this analysis precludes further objective examination. Instead, we will examine a constructive method of determining this volume in the next section.

5.4.4 Sequentially Combining Views for View Variation

The starting point for our view pattern analysis was an over description of landmarks by having each landmark have its own set of views. This causes an important part of view pattern analysis to be the combining of two views into a single view. This is done using the view variation discussed in the preceding section.

In order for two views to be combined, it is necessary that there be at least one equivalent view for each original view from the same location. By comparing the view variation lists of the two views to be combined we can determine if and where such an overlap exists. Any one of the views common to the two lists is sufficient as the combined view. Further, the view variation list for the combined view is comprised of the views which are common to both the original lists.

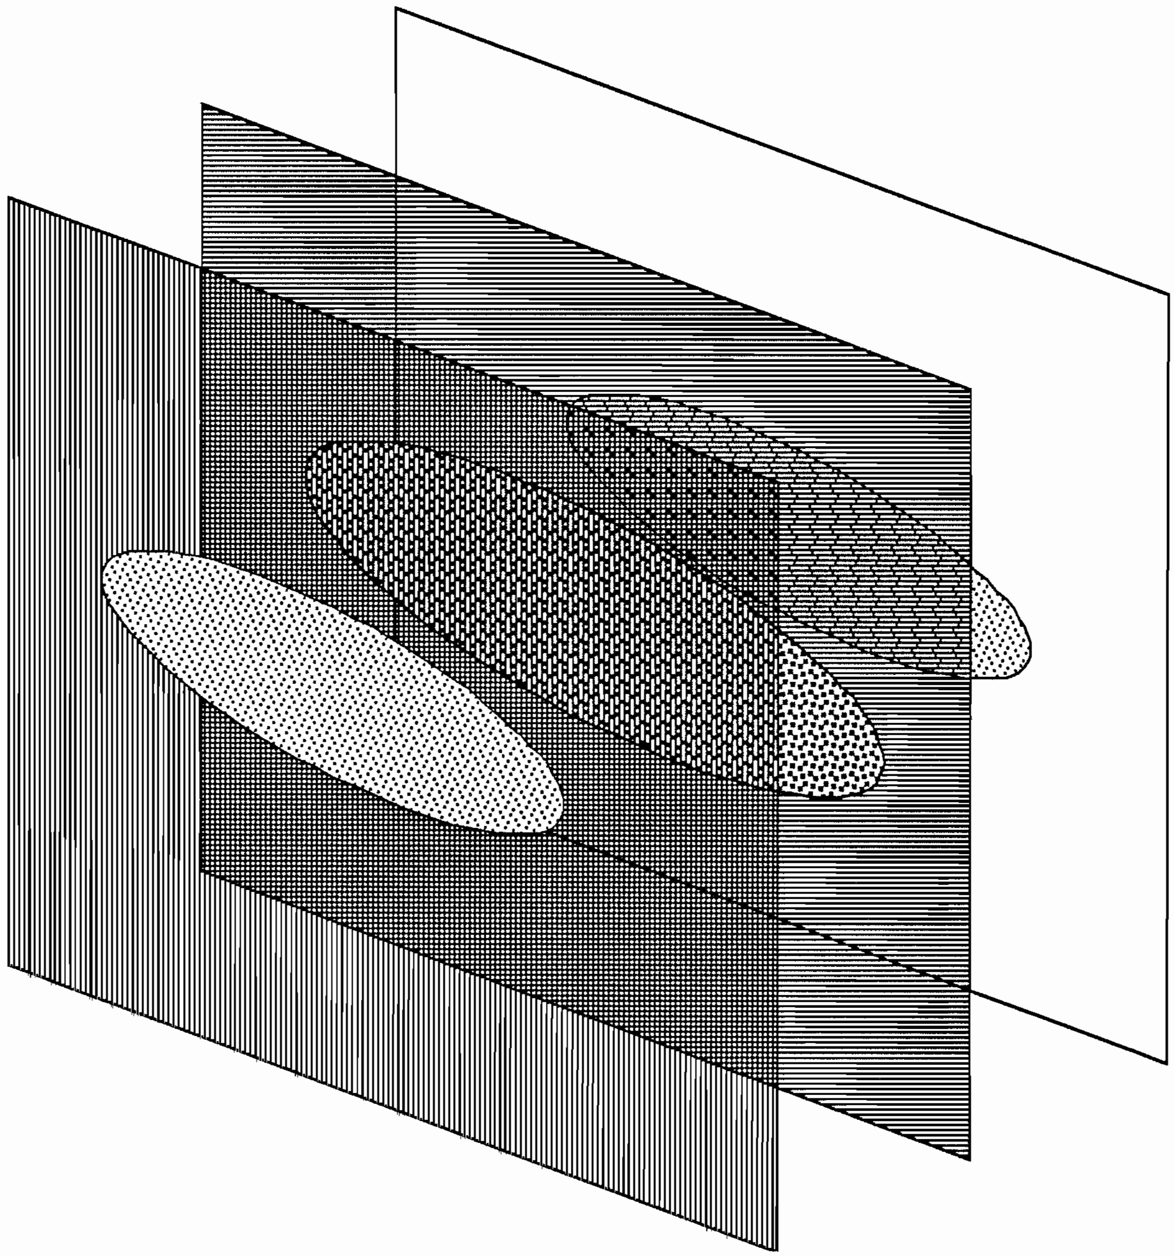

Figure 5.17 Displacement effects toward the other landmark of view plane rotation with respect to two landmarks

This intersection of the view lists of the two original landmarks precludes the need for a direct computation of the view variation of the new view. It is not possible for a view variant to be a member of the combined set and not be a member of both of the originals. This affords us the possibility of exploiting the symmetry of the original best view to determine its view variation.

The use of view variation lists also precludes the need to determine a single view for the combination of two views. This is satisfactory at intermediate stages of combining a sequence of views. A flowchart of this part of the combination routine is shown in Figure 5.18. Ultimately, a specific view will need to be selected from a set which provides equivalent information with respect to the view selection parameters.

The notion of equivalent views was sufficient to describe the views during view pattern analysis. The specific view selection after combining the view variations will need to be done using a criterion which does distinguish among the equivalent views. This final view selection is shown in the flowchart of Figure 5.19.

The criterion we will use is to select the final view which provides the most accuracy for the inspection. The equivalent views were only guaranteed to possess a minimum accuracy in their information. We desire to select a minimum number of views which provide the greatest possible accuracy. This view with the greatest accuracy will be the view which is closest to the object but can still observe all of its landmarks.

It is not immediately apparent which view is the closest to the object. After view combination there are several landmarks associated with each view. Each view-landmark distance will be different. We will use the minimum sum of the distances as our distance criterion. Even so, it is possible that the view determined in this manner may not be unique because views also possess an orientation attribute. From among the views giving the most accuracy to the inspection we will select the median of the angle orientations. This will insure that if possible none of the landmarks will be at the boundary of our ability to distinguish them from the background.

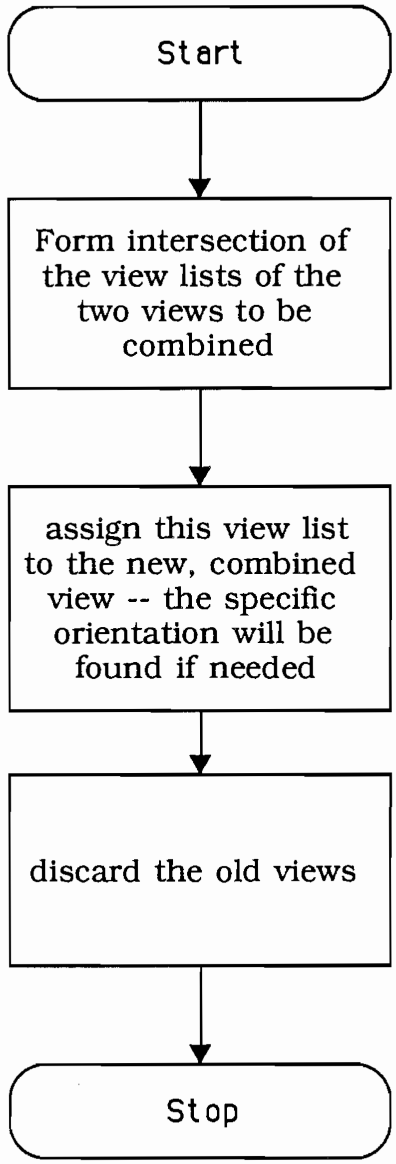

Figure 5.18 Sequential Combination of Views

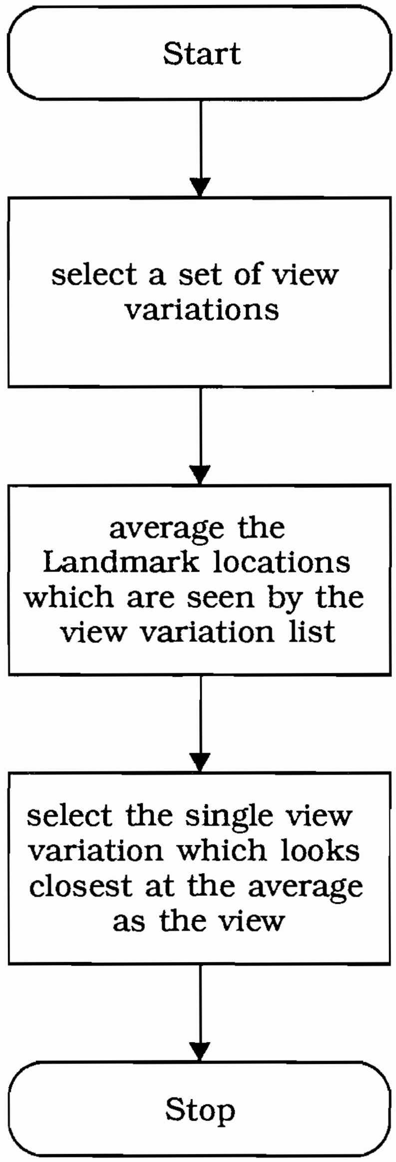

Figure 5.19 View combination algorithm

5.5 View Pattern Analysis

The purpose of view planning analysis is to make use of known information about an object to select views which are sufficient to inspect the object using the analysis of Chapter 3. Figure 5.20 shows a flowchart of this analysis from the designer specifying the items to be inspected to the selection of the final views which will be used for observation.

The view selection process needs to combine the view vectors in a manner which removes redundancy in the view information while maintaining a desired level of potential accuracy for the inspection. The algorithm will terminate when the number of views desired is obtained, with the resultant accuracy being a function of this number of views; or when the number of views is a minimum for the accuracy bound specified. The algorithm fails if it is unable to maintain the desired accuracy using some maximum number of views.

5.5.1 Preprocessing of the View Information

The preprocessing of the view information is the determination of the view variation characteristics of the best views. A flowchart of this is shown in Figure 5.21. A comparison of the view variation characteristics will provide the view selection routine with a map for selecting the best views.

The view variation data is calculated from the parameters defined in Section 5.3. Bounds are found by fixing all parameters but one and evaluating for the bound. The view variation list is formed by filling in between these bounds. The symmetry of the best view is exploited wherever possible.

5.5.2 View Selection

The preprocessing of view variation information makes the view selection process rather straightforward. It is flowcharted in Figure 5.22. We desire to use the views which provide the greatest information about

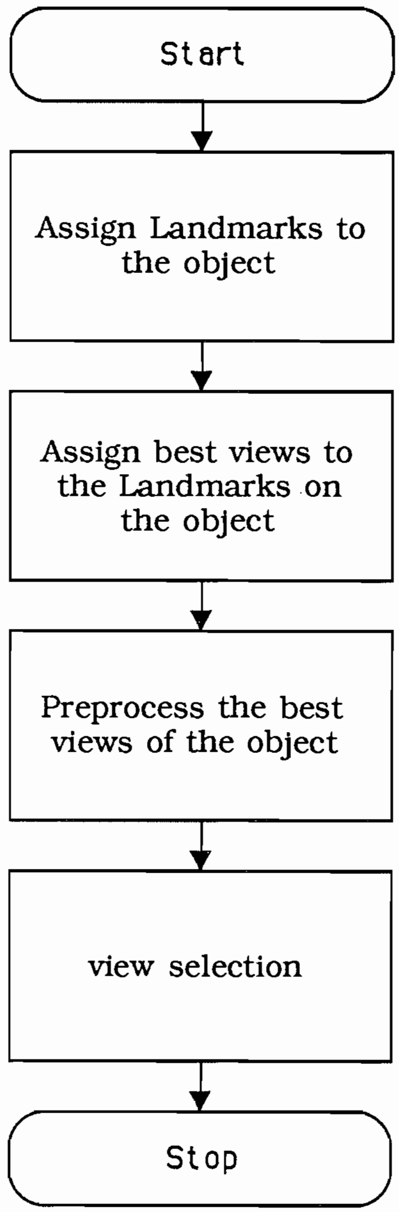

Figure 5.20 View pattern analysis

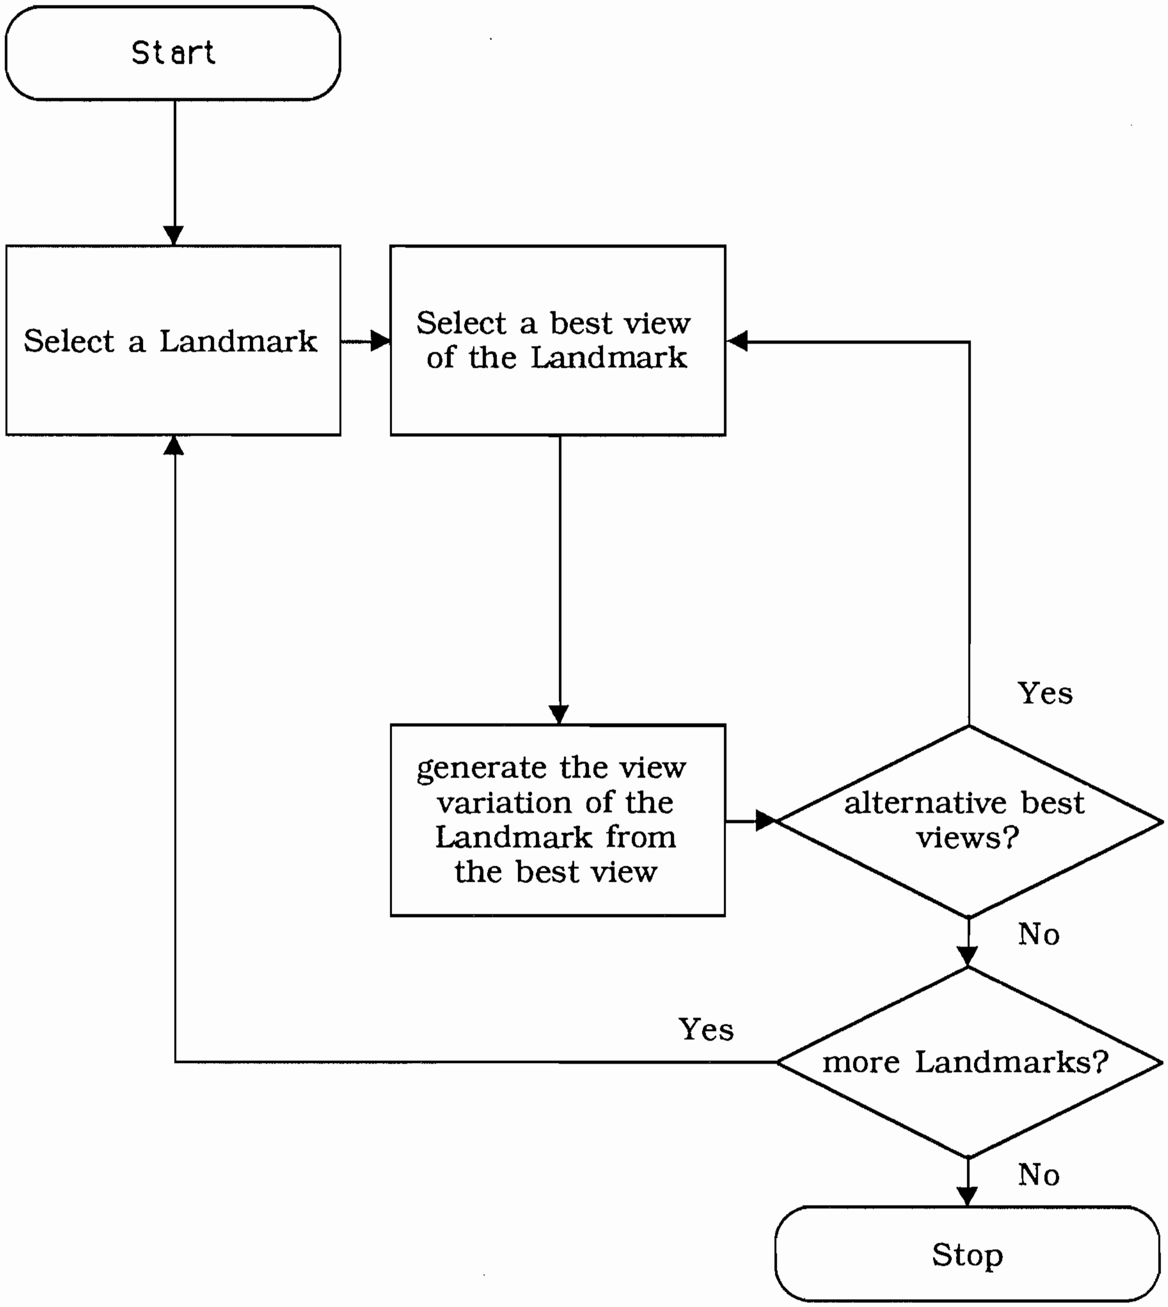

Figure 5.21 View pattern analysis preprocessing

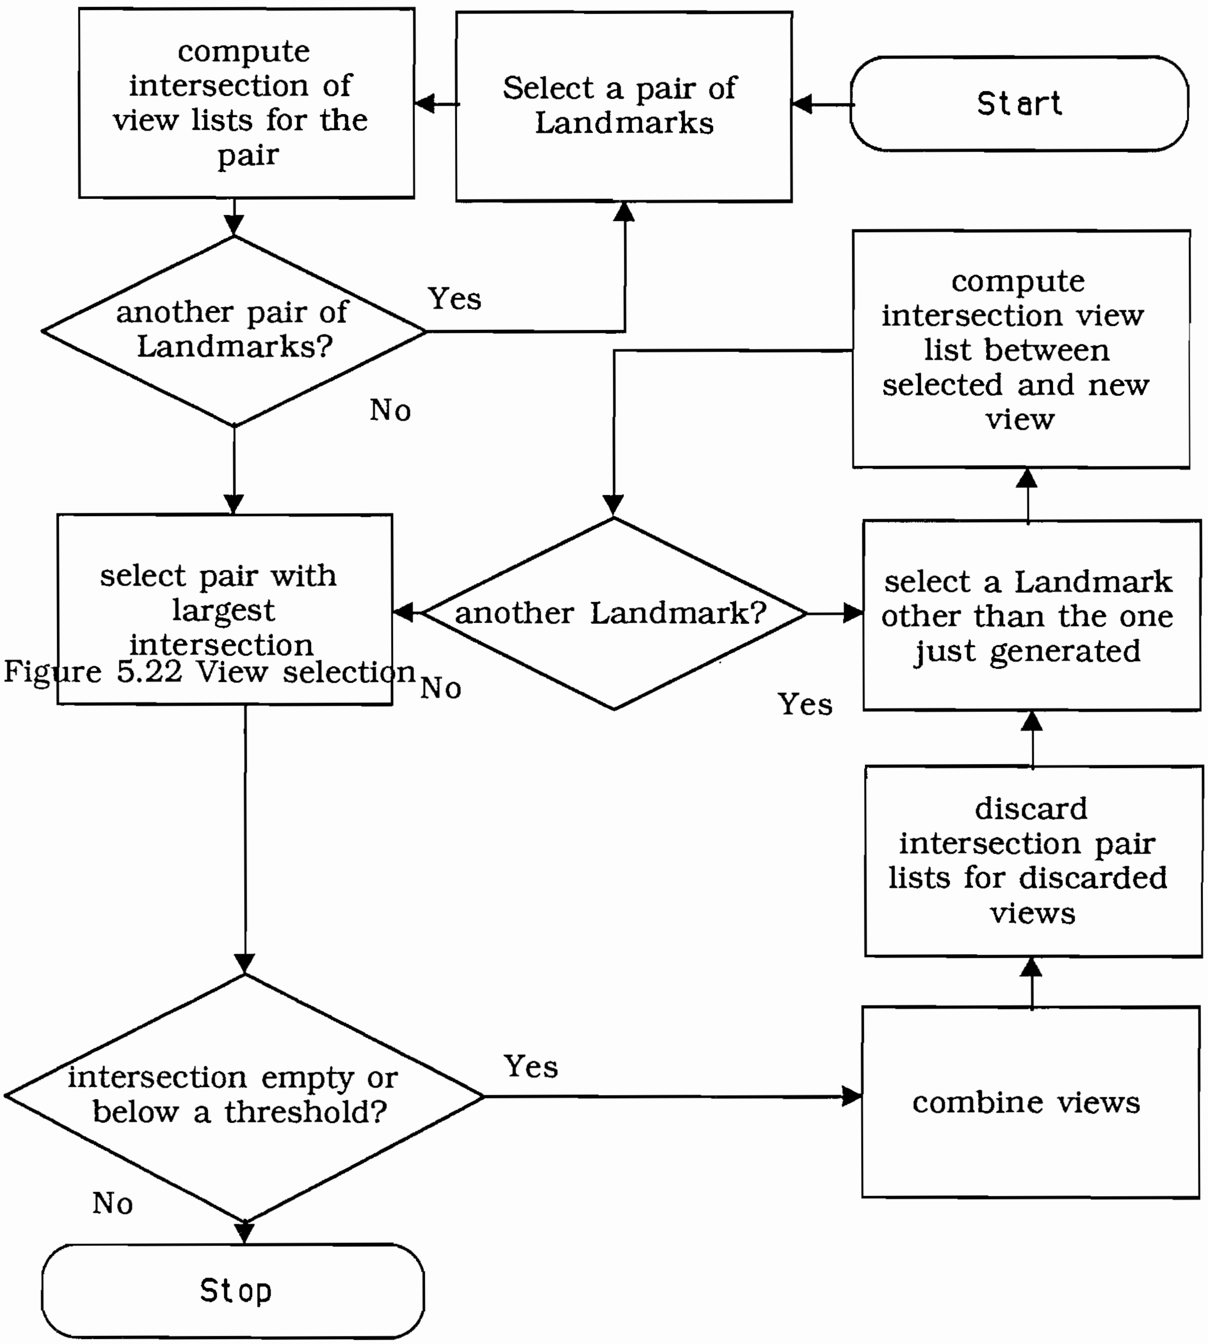

Figure 5.22 View selection

the landmarks. These are selected iteratively by intersections of the view variation lists.

Termination of the view selection process occurs when there are not views to combine or a desired number of views has been obtained. A desired accuracy is maintained by limiting the preprocessing view variation calculation to those view locations which maintain the desired accuracy.

5.6 View Pattern Analysis Implementation

The program was implemented on a Macintosh computer in Think C. C was chosen for its flexibility in handling data types. Since much of the data which is manipulated was in the form of lists, a data structure based upon LISP's list structure was used. C was chosen over LISP because the lists were more easily constructed in C and the construction of correct LISP structures in C subroutines in non-trivial.

5.6.1 Processing

The general structure of the program is shown in the flowchart of Figure 5.23. After initialization, a data input subroutine is called. This routine prompts for input or reads data from a file. The information read includes computer vision system parameters as well as landmark/best view information. The best views may be specified in sets of OR lists which will be AND related.

Next, the variation scanner routine is invoked. This routine scans through all the best views, regardless of whether the view is AND or OR related and builds a view variation list to associate with the view. This is followed by an intersection scanner. This routine scans all possible view pairs, once again regardless of type and computes view variation intersection lists and counts for each.

Now the OR redundancy in the views is removed by scanning for the single OR view which has the largest intersection with a view which is not OR related to it. This process is shown in Figure 5.24. This view should be AND related to the others because of its large intersection. All

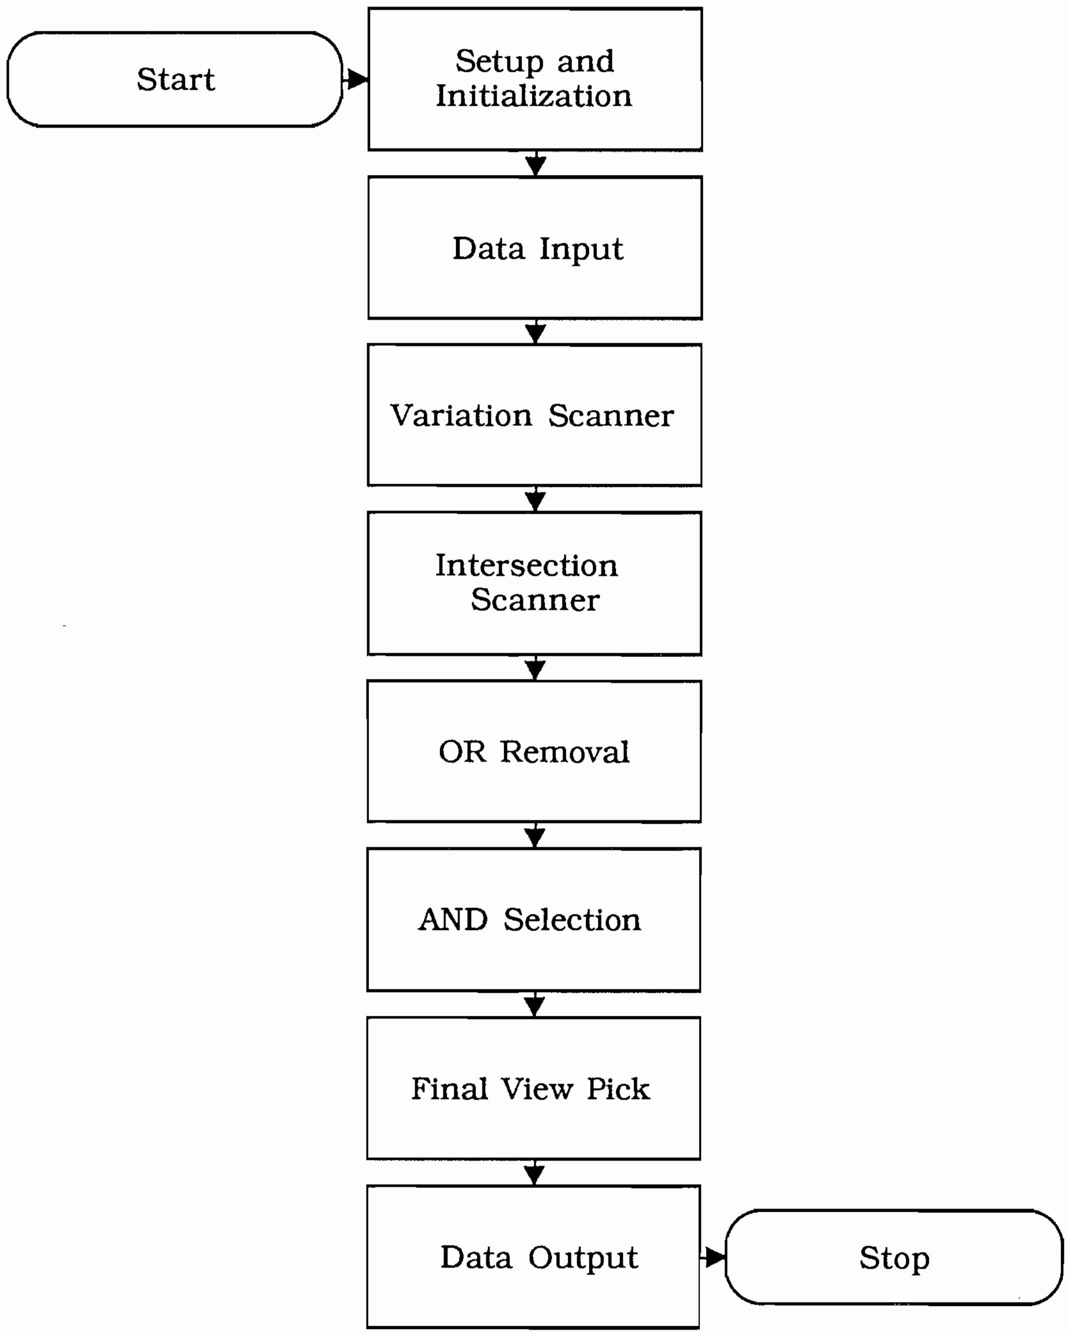

Figure 5.23 General Program Implementation Structure

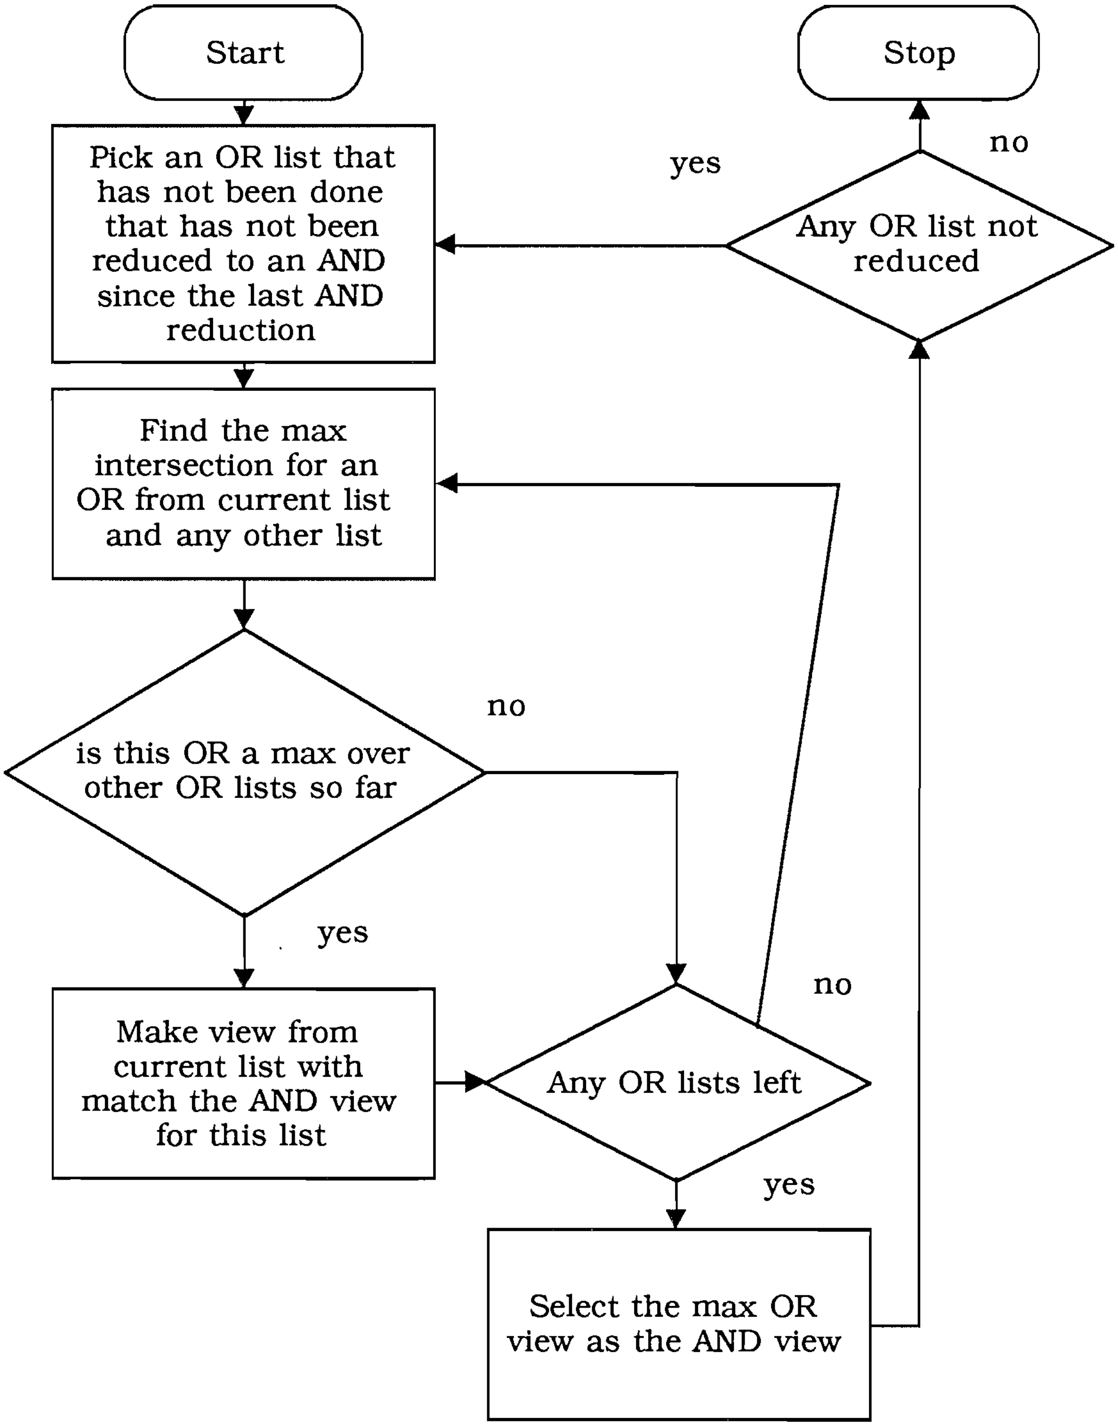

Figure 5.24 OR view redundancy removal

the views which were OR'ed with it are removed -- both themselves and their intersection possibilities with other views. Now the next largest intersection is found in the same manner as above. This is repeated until all the OR lists have been converted into AND'ed views. The selection of the intersections from largest to smallest insures that an early selection from a given OR list does not eliminate what would be a better match later.

Finally, the AND'ed views are reduced through view combination. This is unlike the OR reduction when views were eliminated, but still leaving the same number of AND predicates. The AND combination occurs by iteratively finding the largest intersection of two views. These two views are combined by having the new view have as its view variation list the intersection list from its parent views. The landmarks associated with the parents are maintained for later use in selecting the final view. This process continues until the intersections between remaining views are empty or some minimal threshold of intersection is reached.

The final view is now selected from the view variation lists. This is done as was described earlier in Figure 5.19. The landmarks associated with the each view variation list have their positions averaged. The final view is chosen as the view which looks most directly toward the average landmark location.

5.6.2 View Pattern Analysis Testing

The view selection algorithm was tested for two simple configurations of views. The view variation space was approximated using a simple cube.

The first of these combined two views which were to be AND'ed together. A diagram of this is shown in Figure 5.25. The processing for these views took approximately thirty seconds. The second test involved the combination of two views to be AND'ed together along with two pairs of views OR'ed together. A diagram of this is shown in Figure 5.26. In this case the view selection took over seven minutes.

From this we can conclude that the algorithm is capable of view combination for simple cases which verify our intuition. The time

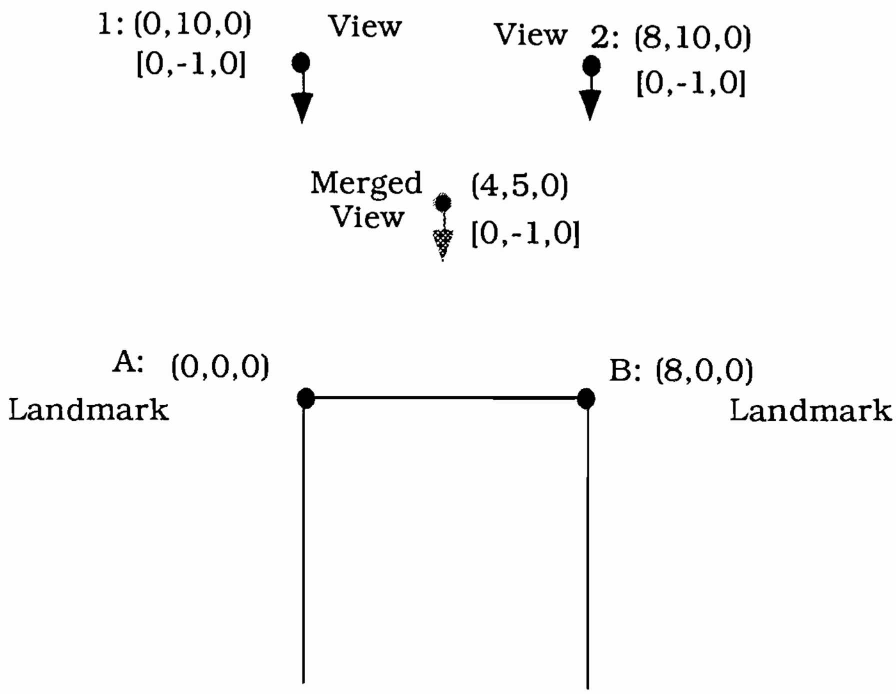

Figure 5.25 View selection of the combination of two views to be AND'ed together

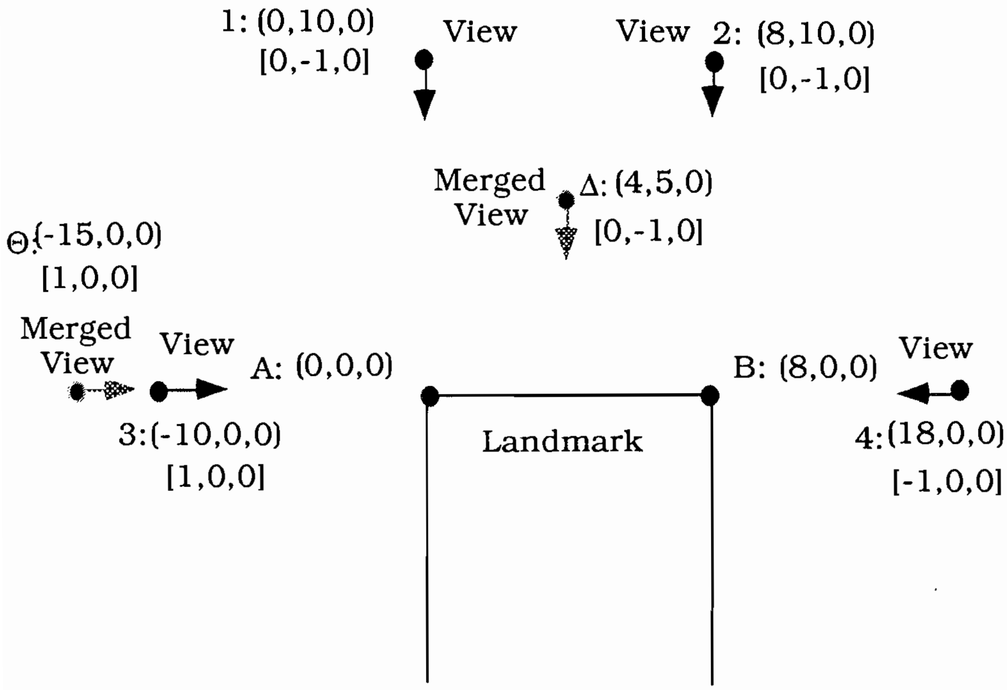

Figure 5.26 View selection of two views to be and'ed together along with views to be or'ed together

required for the analysis even in simple cases implies that the usefulness of view planning by this method is restricted to those instances where the planning will be used in a large number of inspection cases.

5.7 Implications of View Pattern Analysis for Inspection

The implications of view pattern analysis for inspection is the selection of a limited set of views which will adequately describe the landmarks of an object for the purposes of inspection. Further, the failure of the view pattern analysis implies that a set of views using the variation analysis of Section 5.3.7 can not be found which are sufficient for the inspection. Except for the limits of this variation analysis, a view pattern can not be found in any case.

View pattern analysis possesses two possible goals. The first is an inspection analysis which requires less than a specific number of views. This is analogous to having a computational time limit for the inspection. The second is obtaining a desired level of accuracy during the inspection. This limitation may force a greater number of views to be used for inspection.

The view pattern analysis presented meets both these goals by selecting a minimum number of views which provides the greatest possible accuracy subject to the constraint of a minimum accuracy for any specific landmark. The solution is a set of view locations and orientations which can be directly used to perform the inspection.

5.7.1 Implications for Backprojection Reconstruction

In a similar manner to the inspection process, backprojection reconstruction can benefit from the use of view pattern analysis. It has the goals of using a minimum number of views to provide the best possible reconstruction. It differs in its lack of landmark goals for the reconstruction.

If such a set of landmarks existed for the object to be reconstructed, the problem would be equivalent to the one for inspection. There are two different ways which this could be achieved. The object designer could specify landmarks in a similar manner to the inspection case. Here his goal is providing landmarks to describe the object for reconstruction. Another way to specify the landmarks and best views is to generate them from a description of the object. Curvature or another form of feature extraction could be used to extract landmarks. Best views could be generated using the constraints of Chapter 2 for the reconstruction of these landmarks.

Two things should be noted about such a reconstruction. First, in the limiting case, a landmark could be specified at each boundary point on the surface of the object. This would produce a computationally intensive view pattern analysis, but the analysis would remove the extreme redundancy in the view information. Second, any landmarks selected by computer analysis or by large overdescription will need to be *Weakly Convex*.

CHAPTER 6

CONCLUSIONS

6.1 Summary of Research

The purpose of this work was fourfold: to determine a characterization for objects which are suitable for backprojection inspection, to develop an accurate backprojection inspection algorithm, to examine the accuracy considerations of this inspection, and to examine the view selection process for the inspection. Together, the results provide a framework for practical backprojection inspection in conjunction with a CAD system or another system capable of providing the necessary a priori information about the objects to be examined and the views of the sample objects. In each of these areas, results were obtained which satisfy these goals.

In Chapter Two, a characterization of the types of objects suitable for reconstruction was defined. This characterization, called *Weak Convexity*, was based upon the examination of local surface characterizations at every point on the surface of the object. The characterization was shown to be a necessary and sufficient condition for the object to be reconstructed.

In Chapter Three, a backprojection inspection algorithm was presented which only required the examination of a limited portion of the surface of the object. The portion inspected was local to landmarks of the object. These landmarks indicate the locations of the object which were significant to the accurate manufacture of the object. Accuracy measurements were made based solely on the inspection of these locations.

In Chapter Four, three different, but related aspects of backprojection inspection accuracy were examined. The first of these was the calibration and a procedure for inspection calibration was presented. Second, parameters of the inspection process were examined for their interaction in determining the obtainable accuracy during an inspection. These were the object/sensor distance, the viewing angle of the sensor onto a certain size object, the focal length, and the sensor pixel size. These were found to be a highly interdependent set of parameters whose selection would be case-dependent. Finally, methods of combining the inspection results into useful inspection measures were presented. Central to this was a general, local error measure based upon the accuracy found at each landmark.

In Chapter Five, a procedure was presented for the combining of views in an effort to limit the number of views used to perform an inspection. The selection and use of the proper views is essential to a correct inspection. The procedure presented was based upon the inspection procedure of Chapter Three and the accuracy requirements of Chapter Four.

6.2 Suggestions for Future Work

There are several areas where the reconstruction results presented here can be extended. Some are related to a further examination of the results and how the data can be applied. Others relate to the underlying data; how robust the techniques are and the addition of automated preprocessing.

1) The definition of *Weak Convexity* in Chapter One provides a characterization for the types of objects which can be inspected through backprojection inspection. It does not provide an algorithm for the determination of the *Weak Convexity* of an object. It is important that a method be found for the determination of an object's *Weak Convexity*.

2) Several examples of results from the inspection algorithm of Chapter Three were presented to verify the usefulness of the algorithm. An examination of how useful the algorithm would be under actual manufacturing conditions was not presented.

3) The accuracy effects presented in Chapter Four provide a measure of the bounds which restrict backprojection inspection. A further examination of the interaction of these bounds could lead to computational limitations of backprojection reconstruction.

4) In Chapter Five, a method of combing an overdescription of the landmarks to reduce the number of views required was presented. It may be possible to incorporate information concerning how much redundancy is required by the environment to insure that reconstruction conflicts can be resolved.

5) There could be an examination of the results with respect to the determination of actual thresholds for accepting or rejecting parts in an industrial environment. Also, a further study of the camera calibration statistics which could be used for accuracy verification.

6) An examination into the automatic derivation of the best view of a point; or even an automatic determination of the significant points on an object. In the special case of convex shapes, a study could be done to determine the limits of concavity which affect the results. Further, automatic object registration with the optical system could be examined for objects within certain defect bounds.

LIST OF REFERENCES

LIST OF REFERENCES

[1] Abidi, M. A., and T. Chandra, "Pose estimation for camera calibration and landmark tracking," in *Proceedings on1990 IEEE International Conference on Robotics and Automation,* Cincinnati, Ohio, 1990, pp. 420-426.

[2] Ahuja, N., "Generating octrees from object silhouettes in orthographic views," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 2, pp. 137-149, 1989.

[3] Alvertos, Nicolas, Dragana Brzakovic, and Rafael C. Gonzalez, "Camera geometries for image matching in 3-D machine vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 9, pp. 897-915, 1989.

[4] Arman, Farshid, and J. K. Aggarwal, "Object recognition in dense range images using a CAD system as a model base," in *Proceedings on1990 IEEE International Conference on Robotics and Automation,* Cincinnati, Ohio, 1990, pp. 1858-1863.

[5] Avis, D., and G.T. Toussaint, "An optimal algorithm for determining the visibility of a polygon from an edge," *IEEE Transactions on Computers*, vol. C-30, no. 12, pp. 910-914, 1981.

[6] Ayala, D., P. Brunet, R. Juan, and I. Navazo, "Object representation by means of nonminimal division quadtrees and octrees," *ACM Transactions on Graphics*, vol. 4, no. 1, pp. 41-59, 1985.

[7] Ballard, Dana, and Christopher M. Brown, *Computer vision,* Englewood Cliffs, New Jersey: Prentice-Hall, 1982.

[8] Barraquand, Jerome, and Jean-Claude Latombe, "A monte-carlo algorithm for path planning with many degrees of freedom," in *Proceedings on1990 IEEE International Conference on Robotics and Automation*, Cincinnati, Ohio, 1990, pp. 1712-1717.

[9] Bartle, Robert G., and Donald R. Sherbert, *Introduction to real analysis*, New York: John Wiley & Sons, Inc., 1982.

[10] Besl, Paul J., and Ramesh C. Jain, "Three-dimensional object recognition," *Computing SUrveys*, vol. 17, no. 1, pp. 75-145, 1985.

[11] Bin, Ho, "Inputting constructive solid geometry representations directly from 2D orthographic engineering drawings," *Computer Aided Design*, vol. no. pp. 147-155, 1986.

[12] Bogler, Philip L., "Shafer-Dempster reasoning with applications to multisensor target identification systems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-17, no. 6, pp. 968-977, 1987.

[13] Bresler, Yoram, Jeffery A. Fessler, and Albert Macovski, "Model-based estimation techniques for 3-D reconstruction from projections," New York, NY: Springer-Verlag, pp. 115-126, 1988.

[14] Brevdo, L., S. Sideman, and R. Beyar, "A simple approach to the problem of 3-D reconstruction," *Computer Vision, Graphics, and Image Processing*, vol. 37, no. pp. 420-427, 1987.

[15] Buchanan, Thomas, "The twisted cubic and camera calibration," *Computer Vision, Graphics, and Image Processing*, vol. 42, no. pp. 130-132, 1988.

[16] Buchman, Edwin, and F. A. Valentine, "External visibility," *Pacific Journal of Mathematics*, vol. 64, no. 2, pp. 333-340, 1976.

[17] Cappellini, V., R. Casini, M.T. Pareschi, and C. Raspollini, "From multiple views to object recognition," *IEEE Transactions on Circuits and Systems*, vol. Cas-34, no. 11, pp. 1344-1350, 1987.

[18] Cernuschi-Frias, Bruno, David B. Cooper, Yi-Ping Hung, and Peter N. Belhumeur, "Toward a model-based bayesian theory for estimating and recognizing parameterized 3-D objects using two or more images taken from different positions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 10, pp. 1028-1052, 1989.

[19] Chang, Kai-Hsiung, and William G. Wee, "A planning model with problem analysis and operator hierarchy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, no. 5, pp. 672-675, 1988.

[20] Chen, Homer H., and Thomas S. Huang, "A survey of construction and manipulation of octrees," *Computer Vision, Graphics, and Image Processing*, vol. 43, no. pp. 409-431, 1988.

[21] Chen, Lih-Shyang, and Marc R. Sontag, "Representation, display, and manipulation of 3D digital scenes and their medical applications," *Computer Vision, Graphics, and Image Processing*, vol. 48, no. pp. 190-216, 1989.

[22] Chen, Zen, and Der-Baau Perng, "Automatic reconstruction of 3D solid objects from 2D orthographic views," *Pattern Recognition*, vol. 21, no. 5, pp. 439-449, 1988.

[23] Chen, Zen, Din-Chang Tseng, and Jenn-Yee Lin, "A simple vision algorithm for 3-D position determination using a single calibration object," *Pattern Recognition*, vol. 22, no. 2, pp. 173-187, 1989.

[24] Chien, Chiun-Hong, and J.K. Aggarwal, "Model construction and shape recognition from occluding contours," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 4, pp. 372-389, 1989.

[25] Chien, C. H., and J. K. Aggarwal, "Identification of 3D objects from multiple silhouettes using quadtrees/octrees," *Computer Vision, Graphics, and Image Processing*, vol. 36, no. pp. 256-273, 1986.

[26] Chien, C.H., and J.K. Aggarwal, "Volume/surface octrees for the representation of three-dimensional objects," *Computer Vision, Graphics, and Image Processing*, vol. 36, no. pp. 100-113, 1986.

[27] Connolly, C. I., "The determination of next best views," in *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, St. Louis, Missouri, 1985, pp. 432-435.

[28] Cowan, Cregg K., and Peter D. Kovesi, "Automatic sensor placement from vision task requirements," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, no. 3, pp. 407-416, 1988.

[29] Culberson, Joseph, and Robert A. Reckhow, "Orthogonally convex coverings of orthogonal polygons without holes," *Journal of Computer and System Sciences*, vol. 39, no. pp. 166-204, 1989.

[30] Cyganski, David, John A. Orr, David Cubanski, and Steven E. Shiller, "Implementation aspects and performance of 3-D object reconstruction from silhouettes," 1990, pp. 171-179.

[31] Donald, Bruce R., "A search algorithm for motion planning with six degrees of freedom," *Artificial Intelligence*, vol. no. pp. 295-353, 1987.

[32] Erdmann, Michael, "Using backprojections for fine motion: planning with uncertainty," in *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, St. Louis, Missouri, 1985, pp. 549-554.

[33] Fan, Ting-Jun, Gerard Medioni, and Ramakant Nevatia, "Recognizing 3-D objects using surface descriptions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 11, pp. 1140-1157, 1989.

[34] Fisk, Steve, "A short proof of Chvatal's watchman theorem," *Journal of Combinatorial Theory*, vol. Series B 24, no. pp. 374, 1978.

[35] Floriani, Leila De, and Bianca Falcidieno, "A hierarchical boundary model for solid object representation," *ACM Transactions on Graphics*, vol. 7, no. 1, pp. 42-60, 1988.

[36] Fukui, Y, "Input method of boundry solid by sketching," *Computer Aided Design*, vol. no. pp. 434-451, 1988.

[37] Gigus, Ziv, and Jitendra Malik, "Computing the aspect graph for line drawings of polyhedral objects," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 2, pp. 113-122, 1990.

[38] Gilbert, Elmer G., Daniel W. Johnson, and S. Sathiya Keerthi, "A fast procedure for computing the distance between complex objects in three-dimensional space," *IEEE Journal of Robotics and Automation*, vol. 4, no. 2, pp. 193-203, 1988.

[39] Grimson, William Eric Leifur, *From images to surfaces*, Cambridge, Massachusetts: The MIT Press, 1981.

[40] Grosso, Enrico, Giulio Sandini, and Massimo Tistarelli, "3-D object reconstruction using stereo and motion," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 19, no. 6, pp. 1465-1476, 1989.

[41] Gujar, Uday G., and I.V. Nagendra, "Construction of 3D solid objects from orthographic views," *Computers and Graphics*, vol. 13, no. 4, pp. 505-521, 1989.

[42] Haralick, Robert M., "Determining camera parameters from the perspective projection of a rectangle," *Pattern Recognition*, vol. 22, no. 3, pp. 225-2301, 1989.

[43] Harris, C. G., and J. M. Pike, "3D positional integration from image sequences," *Image and Vision Computing*, vol. 6, no. 2, pp. 87-90, 1988.

[44] Henderson, T., E. Weitz, C. Hansen, and A. Mitiche, "Multisensor knowledge systems: interpreting 3D structure," *The International Journal of Robotics Research*, vol. 7, no. 6, pp. 114-137, 1988.

[45] Herman, Martin, "Generating detailed scene descriptions from range images," in *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, St. Louis, Missouri, 1985, pp. 426-431.

[46] Hoff, W., and N. Ahuja, "Surfaces from stereo integrating feature matching, disparity estimation, and contour detection," *IEEE Transactions on Patterm Analysis and Machine Intelligence*, vol. 11, no. 2, pp. 121-136, 1989.

[47] Hoffman, Richard, H. R. Keshavan, and Farhad Towfiq, "CAD-driven machine vision," *IEEE Transactions on SYstems, Man, and Cybernetics*, vol. 19, no. 6, pp. 1477-1488, 1989.

[48] Hong, Tsai-Hong, and Michael O. Shneier, "Describing a robot's workspace using a sequence of views from a moving camera," *IEEE Transactions on Pattern Analysis and Machine Intellignece*, vol. PAMI-7, no. 6, pp. 721-726, 1985.

[49] Horaud, Radu, "New methods for matching 3-D objects with single perspective views," *IEEE Transactions on Pattern Analysis and Machine INtelligence*, vol. 9, no. 3, pp. 401-412, 1987.

[50] Horn, Alfred, and F. A. Valentine, "Some properties of L sets in the plane," *Duke Math Journal*, vol. 16, no. pp. 131-140, 1949.

[51] Horn, Berhold Klaus Paul, *Robot vision*, New York: McGraw-Hill Book Company, 1986.

[52] Huang, T. S., and C. H. Lee, "Motion and structure from orthographic projections," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 5, pp. 536-540, 1989.

[53] Hummel, Robert A., and Michael S. Landy, "A statistical viewpoint on the theory of evidence," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 10, no. 2, pp. 235-247, 1988.

[54] Hutchinson, S.A., *Sensor and task planning in robotic assembly*, PHD Thesis, Purdue University, 1988.

[55] Hutchinson, Seth A., and Avinash C. Kak, "Planning sensing strategies in a robot work cell with multi-sensor capabilities," *IEEE Transactions on Robotics and Automation,* vol. 5, no. 6, pp. 765-783, 1989.

[56] Ihara, Jiro, "Extension of conditional probability and measures of belief and disbelief in a hypothesis based on uncertain evidence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 4, pp. 561-568, 1987.

[57] Irani, Keki B., and Suk I. Yoo, "A methodology for solving problems: problem modeling and heuristic generation," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 10, no. 5, pp. 676-686, 1988.

[58] Izaguirre, Alberto, Pearl Pu, and John Summers, "A new development in camera calibration; calibrating a pair of mobile cameras," in *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, St. Louis, Missouri, 1985, pp. 74-79.

[59] Jackins, Chris L., and Steven L. Tanimoto, "Oct-trees and their use in representing three-dimensional objects," *Computer Graphics and Image Processing*, vol. 14, no. pp. 249-270, 1980.

[60] Jain, Anil K., and Richard Hoffman, "Evidence-based recognition of 3D objects," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, no. 6, pp. 783-802, 1988.

[61] Joshi, S., and T. C. Chang, "Graph-based heuristics for recognition of machined features from a 3D solid model," *Computer Aided Design*, vol. no. pp. 58-66, 1988.

[62] Kanatani, Kenichi, "Reconstruction of consistent shape from inconsistent data: optimization of 2 1/2 D sketches," *Internation Journal of Computer Vision*, vol. 3, no. pp. 261-292, 1989.

[63] Kanatani, K-I, "3D Euclidean Versus 2D Non-Euclidean: Two Approaches to 3D Recovery from Images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 3, pp. 329-332, 1989.

[64] Kay, David C., and Merle D. Guay, *On sets having finitely many points of local nonconvexity & property Pn*, Norman, Oklahoma: The University of Oklahoma, 1969.

[65] Kim, H.S., R. C. Jain, and R. A. Volz, "Object recognition using multiple views," vol. no. pp. 1985.

[66] Kim, Yeon C., and J. K. Aggarwal, "Rectangular parallelepiped coding: A volumetric representation of three-dimensional objects," *IEEE Journal of Robotics and Automation*, vol. RA-2, no. 3, pp. 127-134, 1986.

[67] Korf, Richard E., "Planning as search: a quantitative approach," *Artificial Intelligence*, vol. 33, no. pp. 65-88, 1987.

[68] Korn, Matthew R., and Charles R. Dyer, "3-D multiview object representations for model-based object recognition," *Pattern Recognition*, vol. 20, no. 1, pp. 91-103, 1987.

[69] Kroszynski, Uri I., Bjarne Palstroem, Erik Trostmann, and Ernst G. Schlechtendahl, "Geometric data transfer between CAD systems: solid models," *IEEE Computer Graphics & Applications*, vol. no. pp. 57-71, 1989.

[70] Lavakusha, Arun K. Pujari, and P. G. Reddy, "Linear octrees by volume intersection," *Computer Vision, Graphics, and Image Processing*, vol. 45, no. pp. 371-379, 1989.

[71] Lee, Chia-Hoang, "A comparison of two evidential reasoning schemes," *Artificial Intelligence*, vol. 35, no. pp. 127-134, 1988.

[72] Lee, D. T., and A. K. Lin, "Computing the visibility polygon from an edge," *Computer Vision, Graphics, and Image Processing*, vol. 34, no. pp. 1-19, 1986.

[73] Lee, D. T., and F. P. Preparata, "An optimal algorithm for finding the kernel of a polygon," *Journal of the Association for Computing Machinery*, vol. 26, no. 3, pp. 415-421, 1979.

[74] Lee, Sukhan, and Youngchul Kay, "An accurate estimation of 3-D position and orientation of a moving object for robot stereo vision: kalman filter approach," in *Proceedings on1990 IEEE International Conference on Robotics and Automation*, Cincinnati, Ohio, 1990, pp. 414-419.

[75] Lin, Wei-Chung, and Shiuh-Yung Chen, "A new surface interpolation technique for reconstructing 3D objects from serial cross-sections," *Computer Vision, Graphics, and Image Processing*, vol. 48, no. pp. 124-143, 1989.

[76] Linnainmaa, Seppo, David Harwood, and Larry S. Davis, "Pose determination of a three-dimensional object using triangle pairs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, no. 5, pp. 634-647, 1988.

[77] Liu, Yuncai, Thomas S. Huang, and Olivier D. Faugeras, "Determination of camera location from 2D to 3D line and point correspondences," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 1, pp. 28-37, 1990.

[78] Lowe, David G., "Three-dimensional object recognition from single two-dimensional images," *Artificial Intelligence*, vol. no. pp. 355-395, 1987.

[79] Macias-Garza, Fernando, Kenneth R. DIller, Aslan C. Bovik, Shanti J. Aggarwal, and J. K. Aggarwal, "Obtaining a Solid Model From Optical Serial Sections," *Pattern Recognition*, vol. 22, no. 5, pp. 577-586, 1989.

[80] Magee, M. J., and J. K. Aggarwal, "Using multisensory images to derive the structure of three-dimensional objects - a review," *Computer Vision, Graphics, and Image Processing*, vol. 32, no. pp. 145-157, 1985.

[81] Mao, Xiaoyang, Tosiyasu L. Kunii, Issei Fujishiro, and Tsukasa Noma, "Hierarchical representations of 2D/3D gray-scale images and their 2D/3D two-way conversion," *IEEE Computer Graphics and Applications*, vol. no. December, 1987, pp. 37-44, 1987.

[82] Marr, David, *Vision*, San Francisco: W. H. Freeman and Company, 1982.

[83] Martin, W.N., and J.K. Aggarwal, "Volumetric descriptions of objects from multiple views," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 5, no. 2, pp. 150-158, 1983.

[84] Martins, H.A., J.R. Birk, and R.B. Kelly, "Camera models based on data from two calibration planes," *Computer Graphics and Image Processing*, vol. 17, no. pp. 173-180, 1981.

[85] Martins, Joao P., and Stuart C. Shapiro, "A model for belief revision," *Artificial Intelligence*, vol. 35, no. pp. 25-79, 1988.

[86] Merickel, Michael, "3D reconstruction: the registration problem," *Computer Vision, Graphics, and Image Processing*, vol. 42, no. pp. 206-219, 1988.

[87] Millman, Richard S., and George D. Parker, *Geometry*, New York: Springer-Verlag, 1981.

[88] Mohan, Rakesh, Gerard Medioni, and Ramakant Nevatia, "Stereo error detection, correction, and evaluation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 2, pp. 113-120, 1989.

[89] Mohan, Rakesh, and Ramakant Nevatia, "Using perceptual organization to extract 3-D structures," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 11, pp. 1121-1139, 1989.

[90] Murray, D. W., "Model-based recognition using 3D shape alone," *Computer Vision, Graphics, and Image Processing,* vol. 40, no. pp. 250-266, 1987.

[91] Murray, D. W., D. A. Castelow, and B. F. Buxton, "From an image sequence to a recognized polyhedral object," *Image and Vision Computing,* vol. 6, no. 2, pp. 107-120, 1988.

[92] Nilsson, Nils J., *Principles of Artificial Intelligence,* Palo Alto, California: Tioga Publishing, 1980.

[93] Noborio, H., S. Fukuda, and S. Arimoto, "Construction of the octree approximating three-dimensional objects by using multiple views," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 10, no. 6, pp. 769-782, 1988.

[94] Phillips, Tsaiyun Hong, and Azriel Rosenfeld, "Decomposition of 3D objects into compact subobjects by analysis of cross-sections," *Image and Vision Computing,* vol. 6, no. 1, pp. 33-51, 1988.

[95] Ponce, Jean, "An object centered hierarchical representation for 3D objects: the prism tree," *Computer Vision, Graphics, and Image Processing,* vol. 38, no. pp. 1-28, 1987.

[96] Poor, H. Vincent, *An introduction to signal detection and estimation,* New York, NY: Springer-Verlag, 1988.

[97] Potmesil, Michael, "Generating octree models of 3D objects from their silhouettes in a sequence of images," *Computer Vision, Graphics, and Image Processing,* vol. 40, no. pp. 1-29, 1987.

[98] Puget, P., and T. Skordas, "An optimal solution for mobile camera calibration," in *Proceedings on1990 IEEE International Conference on Robotics and Automation,* Cincinnati, Ohio, 1990, pp. 34-39.

[99] Rao, Kashipati, and R. Nevatia, "Computing volume descriptions from sparse 3-D data," *International Journal of Computer Vision,* vol. 2, no. pp. 33-50, 1988.

[100] Raviv, Daniel, Yoh-han Pao, and Kenneth A. Laparo, "Reconstruction of three-dimensional surfaces from two-dimensional binary images," *IEEE Transactions on Robotics and Automation,* vol. 5, no. 5, pp. 701-710, 1989.

[101] Ray, L. P., "Estimation of modeled object pose from monocular images," in *Proceedings on1990 IEEE International Conference on Robotics and Automation*, Cincinnati, Ohio, 1990, pp. 408-413.

[102] Reeves, Anthony P., and Russell W. Taylor, "Identification of three-dimensional objects using range information," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 4, pp. 403-410, 1989.

[103] Richards, T. H., and G. C. Onwubolu, "Automatic interpretation of engineering drawings for 3D surface representation in CAD," *Computer Aided Design*, vol. no. pp. 156-160, 1986.

[104] Samet, Hanan, "Implementing ray tracing with octrees and neighbor finding," *Computers and Graphics*, vol. 13, no. 4, pp. 445-460, 1989.

[105] Samet, Hanan, "Neighbor finding in images represented by octrees," *Computer Vision, Graphics and Image Processing*, vol. 46, no. pp. 367-386, 1989.

[106] Samet, Hanan, and Robert E. Webber, "Hierarchical data structures and algorithms for computer graphics Part II: Applications," *IEEE Computer Graphics and Applications*, vol. no. July 1988, pp. 59-75, 1988.

[107] Samet, Hanan, and Robert E. Webber, "Hierarchical data structures and algorithms for computer graphics Part I: Fundementals," *IEEE Computer Graphics and Applications*, vol. no. May, 1988, pp. 48-68, 1988.

[108] Schwartz, J. T., and M. Sharir, "A survey of motion planning and related geometric algorithms," *Artificial Intelligence*, vol. 37, no. pp. 157-169, 1988.

[109] Shafer, Glenn, and Roger Logan, "Implementing Dempster's rule for hierarchical evidence," *Artificial Intelligence*, vol. 33, no. pp. 271-298, 1987.

[110] Shaffer, Clifford A., and Hanan Samet, "Optimal quadtree construction algorithms," *Computer Vision, Graphics, and Image Processing*, vol. 37, no. pp. 402-419, 1987.

[111] Shaffer, Clifford A., and Hanan Samet, "Algorithm to expand regions represented by linear quadtrees," *Image and Vision Computing*, vol. no. pp. 162-168, 1988.

[112] Shin, Sung Y., and Tony C. Woo, "An optimal algorithm for finding all visible edges in a simple polygon," *IEEE Transactions on Robotics and Automation,* vol. 5, no. 2, pp. 202-207, 1989.

[113] Shmuel, Amir, and Michael Werman, *Active Vision: 3D From an Image Sequence,* Jerusalem, Israel: Department of Computer Science, The Hebrew University of Jerusalem, 1989.

[114] Siebers, Gregory R., "An introduction to computer graphics," *Computer Aided Design,* vol. no. pp. 161-179, 1986.

[115] Smith, David E., "Controlling backward inference," *Artificial Intelligence,* vol. 39, no. pp. 145-208, 1989.

[116] Snyder, M. A., "The Precision of 3-D parameters in correspondence-based techniques: the case of uniform translational motion in a rigid environment," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 11, no. 5, pp. 523-528, 1989.

[117] Stephens, Mike, and Chris Harris, "3D wire-frame integration from image sequences," *Image and Vision Computing,* vol. no. pp. 24-30, 1989.

[118] Strat, Thomas M., and Martin A. Fischler, "One-eyed stereo A general approach to modeling 3-D scene geometery," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 8, no. 6, pp. 730-741, 1986.

[119] Suh, Suk-Hwan, and Albert B. Bishop, "Collision-avoidance trajectory planning using tube concept: analysis and simulation," *Journal of Robotic Systems,* vol. 5, no. 6, pp. 497-525, 1988.

[120] Suh, Suk-Hwan, and Kang G. Shin, "A variational dynamic programming approach to robot-path planning with a distance-safety criterion," *IEEE Journal of Robotics and Automation,* vol. 4, no. 3, pp. 334-349, 1988.

[121] Sun, Y. N., C. C. Li, P. R. Krishnaiah, and P. S. Reddy, "Three-dimensional reconstruction of ventricle from biplane angiocardiograms via equai-divisor surface," *IEEE Transactions on Systems, Man, and Cybernetics,* vol. 19, no. 6, pp. 1666-1671, 1989.

[122] Tamura, Shinichi, and Koji Yata, "Plan-based boundary extraction and 3D reconstruction for orthogonal 2D echocardiography," *Pattern Recognition*, vol. 20, no. 2, pp. 155-162, 1987.

[123] Tan, H.L., E. Viscito, E.J. Delp, and J.P. Allebach, "Inspection of machine parts by backprojection reconstruction," vol. no. pp.

[124] Thomas, *An introduction to applied probablility and random processes*, Huntington, New York: Robert E. Krieger Publishing Company, 1981.

[125] Tor, S. B., and A. E. Middleditch, "Convex decomposition of simple polygons," *ACM Transactions on Graphics*, vol. 3, no. 4, pp. 244-265, 1984.

[126] Trees, Harry L. Van, *Detection, estimation, and modulation theory*, New York, NY: John Wiley & Sons, 1968.

[127] Valentine, F. A., "Minimal sets of visibility," *Proceedings of the American Mathematical Society*, vol. 4, no. pp. 917-921, 1953.

[128] Valentine, Frederick A., *Convex sets*, New York: McGraw-Hill Book Company, 1964.

[129] Valentine, F. A., "Local convexity and Ln sets," *Proceedings of the American Mathematical Society*, vol. 16, no. pp. 1305-1310, 1965.

[130] Veenstra, Jack, and Narendra Ahuja, "Octree generation from silhouette views of an object," in *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, St. Louis, Missouri, 1985, pp. 843-848.

[131] Veenstra, Jack, and Narendra Ahuja, "Line drawings of octree-represented objects," *ACM Transactions on Graphics*, vol. 7, no. 1, pp. 61-75, 1988.

[132] Walker, Ellen Lowenfeld, and Martin Herman, "Geometric reasoning for constructing 3D scene descriptions from images," *Artificial Intelligence*, vol. 37, no. pp. 275-290, 1988.

[133] Wang, Y. F., and J. K. Aggarwal, "Integration of active and passive sensing techniques for representing three-dimensional objects," *IEEE Transactions on Robotics and Automation*, vol. 5, no. 4, pp. 460-471, 1989.

[134] Waxman, Allen M, and James H Duncan, "Binocular image flows: steps toward stereo fusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 715-729, 1986.

[135] Weng, Juyang, and Narendra Ahuja, "Octrees of objects in arbitrary motion: representation and efficiency," *Computer Vision, Graphics, and Image Processing*, vol. 39, no. pp. 167-185, 1987.

[136] Weng, Juyang, Thomas S. Huang, and Narendra Ahuja, "3D motion estimation, understanding, and prediction from noisy image sequences," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 9, no. 3, pp. 370-389, 1987.

[137] Weng, Juyang, Thomas S. Huang, and Narendra Ahuja, "Motion and structure from two perspective views: algorithms, error analysis, and error estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 5, pp. 451-476, 1989.

[138] Werman, Michael, Shmuel Peleg, and Azriel Rosenfeld, "A distance metric for multidimensional histograms," *Computer Vision, Graphics, and Image Processing*, vol. 32, no. pp. 328-336, 1985.

[139] Wilson, Peter R., "Geometry and topology as modeling abstractions," Detroit, MI, 1987, pp. 1-15 to 1-30.

[140] Wong, A.K.C., SI W. Lu, and M. Rioux, "Recognition and shape synthesis of 3-D objects based on attributed hypergraphs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 3, pp. 279-290, 1989.

[141] Wu, J. J., R. E. Rink, T. M. Caelli, and V. G. Gourishankar, "Recovery of the 3D location and motion of a rigid object through camera image (an exteded kalman filter approach)," *Internation Journal of Computer Vision*, vol. 3, no. pp. 373-394, 1988.

[142] Yakimovsky, Y., and R. Cunningham, "A system for extracting three-dimensional measurements from a stereo pair of TV cameras," *Computer Graphics and Image Processing*, vol. 7, no. pp. 195-210, 1978.

[143] Yamamoto, Masanobu, "A general aperture problem for direct estimation of 3-D motion parameters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 5, pp. 528-536, 1989.

[144] Yang, B., W. E. Snyder, and G. L. Bilbro, "Matching oversegmented 3D images to models using association graphs," *Image and Vision Computing,* vol. no. pp. 135-143, 1989.

[145] Yang, Meng-Chien, and Wen-Hsiang Tsai, "Recognition of single 3D curved objects using 2D cross-sectional slice shapes," *Image and Vision Computing,* vol. no. pp. 210-216, 1989.

[146] Zhuang, Xinhua, Thomas S. Huang, and Robert M. Haralick, "A simple procedure to solve motion and structure from three orthographic views," *IEEE Journal of Robotics and Automation*, vol. 4, no. 2, pp. 236-239, 1988.

VITA

VITA

Charles Scott Foshee was born in Radford, Virginia on June 10th, 1959. He attended Marietta High School in Marietta, Ohio and graduated in June of 1977. He received his Bachelor of Science in Electrical Engineering, Honors Program, in 1981; his Master of Science in Electrical Engineering in 1984; and his Doctor of Philosophy in 1990; all from Purdue University, West Lafayette, Indiana. He is an Eagle Scout, and is a member of Eta Kappa Nu, Tau Beta Pi, and Phi Eta Sigma Scholastic Honoraries, and Omicron Delta Kappa Leadership Honor Society. He is a member of the Institute of Electrical and Electronic Engineers, the American Association of Artificial Intelligence, and the National Society of Professional Engineers.