MATHEMATICAL MORPHOLOGY APPLIED TO RANGE IMAGE PROCESSING

Colin C. Archibald
Machine Vision International Ltd., 280 Albert St., Ottawa, Canada. K1P 5G8

Stanley R. Sternberg
Machine Vision International Corp., 325 E. Eisenhower Pway Ann Arbor, MI U.S.A. 48104

ABSTRACT

Methods of low level image processing have predominantly been derived by generalizing one dimensional signal processing methods to 2 dimensions. Although much progress has been made using Fourier domain analysis, it is still unintuitive and inflexible for many types of applications. This paper reviews mathematical morphology as a basis for low level image processing, and demonstrates that image domain transformations can be applied usefully to various types of images. Specifically, range images are processed using grey scale morphology to extract various features used in three dimensional analysis of a scene. This type of scene analysis has broad potential for applications in quality control and automated manufacturing.

RESUME

Le traitement primaire d'image a traditionellement été inspiré des méthodes de traitement de signaux en les généralisant pour deux dimensions. Le traitement basé sur l'analyse dans le domaine de Fourier a permis d'important progrès mais défie l'intuition et demeure trop peu flexible pour nombre de problèmes.

Nous faisons ici un compte rendu des applications de la morphologie en tant que base d'analyse primaire de l'image en l'appliquant a diverses classes d'images. Plus précisement, des images tri-dimensionelles sont traitées par la méthode morphologique en tons de gris pour extraire divers éléments propres à l'analyse de scenes tri-dimensionelles.

Cette classe d'analyse ouvre un large éventail de possibilités dans les applications de contrôle de qualité et de fabriquation assistée par ordinateur.

I. INTRODUCTION

The arrival of mathematical morphology for image processing has allowed North Americans to reconsider their approach to image transformation. Transform domain techniques, like the Fourier transform, used to decompose an image into constituent spatial frequencies, grew out of classical one dimensional signal processing where the dominant analytical theme is linear transformation. Conventional image processing extended the subject base to two or more dimensions, but the paradigm remained one of filter design for removal of noise and enhancement of visual fidelity.

A morphological approach allows the designer of vision systems a ubiquitous tool to perform image transforms in the image domain, using the algebra of shapes. Although the morphological treatment of images has been studied for more than 50 years, its recent popularization is due mostly to Serra (1), and Sternberg (2) (3), the latter introducing greyscale morphology.

The use of mathematical morphology has overcome several obstacles in the development of industrial vision systems. Mathematical morphology is first of all a mathematics of image transformation and analysis, thus it forms the basis for a powerful language of image processing. This language is not only powerful, it is intuitive, and can be understood visually while developing image processing algorithms interactively. As a general purpose method of image transformation, all types of images can be processed regardless of the sensor used to collect the image. Satellite, microscope, x-ray, T.V., ultrasound, tactile array sensors, laser range finder images, etc., can all be transformed productively using mathematical morphology. This paper describes the basic operations of how to transform an image with a shape, calling on the readers intuition instead of the mathematical basis of the transformations which are exhaustively described in (1) and (4). We will demonstrate, using figures, transformations on binary and grey scale images, including images acquired using a laser range finder. It will be pointed out that these morphological transforms can be used for low level (data driven), and high level (model driven) aspects of industrial applications.

II. RANGE IMAGE PROCESSING

Currently range finding cameras are becoming commercially available. These devices use passive or active light sources to record two dimensional arrays of depth, i.e., the distance from the camera to the surface being imaged. Active light sources (often lasers) are used in two ways to collect depth information. Time of flight range finders are similar in concept to radar where the time required for the laser to be reflected from the scene is used as a measure of distance. Triangulation based range finders depend on the location of the reflected light to calculate the distance. A triangulation range finder developed in the Division of Electrical Engineering at the National Research Council was used in this research, and is described in (5). (A survey of range finding methods can be found in (6).) The elements in a range image are referred to as "surfels" or surface elements, and must be processed using algorithms designed specifically for this type of image. In some respects a range image is easier to analyze than intensity images because the information in this type of image is a function of only one scene property, namely depth. Pixels in intensity images, such as those from T.V. cameras, are dependent on reflectivity, lighting intensity and direction, colour, etc., making robust algorithms elusive for all but the most controlled environments.

Range image processing research has centered around extraction of information contained in edges and/or surfaces. Although it can be argued that one leads to the other, the various methods for extracting edge features differ significantly from region detection. Each method has virtues, depending on the ultimate goal of the research. Edge finding methods have been used in (8) (9) (10). Gil et al., (11) used registered range and intensity images to build a more complete edge map. Planar and/or quadratic surfaces are extracted in (7) (12) (13) (14). An interesting approach to collecting the maximum information from the image is to combine the edge and region detection methods. An operator for this purpose is described in (10). A more indepth survey of related literature can be found in (15). Generally, both edge and region finding methods have grown from extending and manipulating ideas used in intensity images. A critical assumption which can be made about range images has been ignored. The information in a range image is not predominantly in the edges or the regions, but in the three dimensional shape represented by the image data. Instead of processing range images with edge or region operators, we suggest that shapes should be used. These shapes exist in the same coordinate system as the image data making the concept of shape transformations easy to visualize. A mathematics of shape transformation called mathematical morphology is described in the following section.

III. MATHEMATICAL MORPHOLOGY

Mathematical morphology is the algebraic treatment of shapes. A transformation based on a shape (sometimes referred to as a structuring element) is surprisingly easy to understand with the use of a simple visual example. There are two classical transformations that are most often used; erosion and dilation. Consider first the two dimensional, binary image application of these operations. What does it mean to erode and dilate a binary image with a shape? To illustrate the answer to this question we have inserted three figures of watch gears. Fig. 1 is the original binary image, Fig. 2 is a dilated image, and Fig. 3 is an eroded image. Notice the small disk in Fig. 1. This shape is used to dilate and erode the image into Figs 2 and 3 respectively. The disk is not a part of the image, but is placed in the figure to help the reader visualize the transformations. In the following textual description, Fig. 1 is the original image, and Fig. 2 is the resultant image. To perform a dilation of the original image, place the structuring element (in this case a disk) so that the centre point falls on a pixel i, j in the original image. If the value of the pixel i, j in the original image is 1, then each pixel which is covered by the structuring element will become 1 in the resultant image. This step is performed for all pixels in the image. The reader is encouraged to verify that this is indeed how Fig. 2 resulted from Fig. 1.

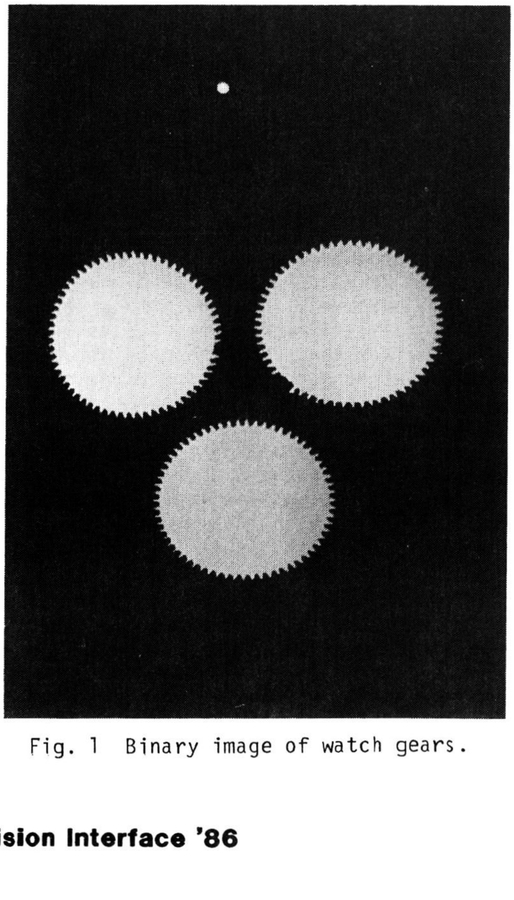

Fig. 1 Binary image of watch gears.

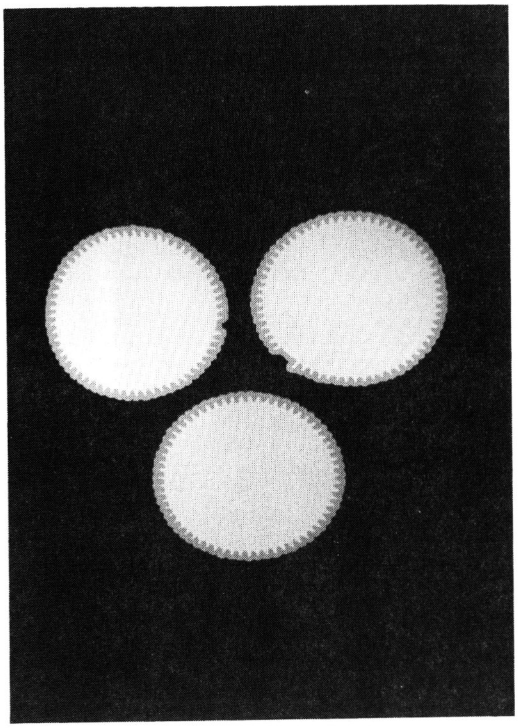

Fig. 2 Dilated watch gears.

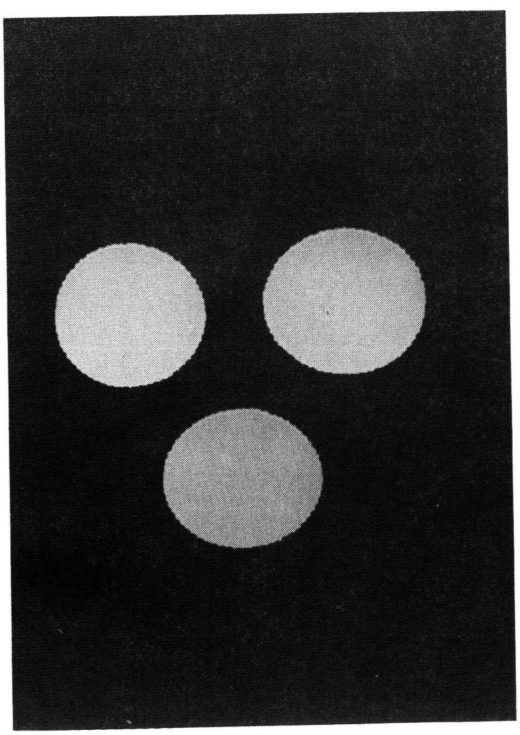

Fig. 3 Eroded watch gears.

The erosion transformation can be visualized in two ways. If the dilation is clearly understood, then the erosion transformation can be considered a dilation of the background, i.e., dilate the 0's in the image instead of the 1's. Notice that Fig. 3 is the result of eroding Fig. 1. with the same structuring element used previously. The second way to understand an erosion, is fundamentally the same as the description of the dilation. Place the structuring element at each pixel i, j in the original image. If the value of pixel i, j is 0 then, each pixel which is covered by the structuring element will become 0 in the resultant image.

The structuring element of a small disk was used in this explanation to demonstrate the utility of such a transform in an industrial application. In this case, a priori knowledge of the size of the gear and spacing of the gear teeth allowed the choice of the correct size structuring element which could be used to identify the broken teeth on the gears. This is an example of a model driven transform, where the model was a priori knowledge of the size and shape of the gear.

Next consider what it means to transform a greyscale image, either range or intensity, with a three dimensional shape such as a ball. To facilitate this description we point out that a grey scale image can be thought of as a function f(x,y) on the points of Euclidean 2-space. In 3-space a grey scale image is a set of points x,y, f(x,y) where f(x,y) is the pixel value. In a range image f(x,y) = z since the image is a three dimentional representation of the scene. This can be visualized as a thin, not necessarily continuous sheet. Fig. 4 is an oblique graphical representation of a range image, demonstrating how an image can be thought of as a sheet. (This image is further discussed in the following section.)
Greyscale erosions and dilations on these sheets are most often used in conjunction with each other. An erosion followed by a dilation is called an opening, and a dilation followed by an erosion is called a closing. The opening or closing of a sheet by a geometrical solid is a grey scale transformation which treats different portions of an image uniquely, depending on how well the local grey level topology is matched by the shape.

Fig. 5 illustrates a greyscale closing operation, using a ball. The ball is rolled over every pixel (surfel) on the image. Where the topology of the sheet prevents the ball from touching each pixel, the value of the pixel is raised to the surface of the ball. A greyscale opening is the mathematical dual. The ball is rolled on the underside of the sheet, and the pixel values are lowered to touch the ball. (To see this turn Fig. 5 upside down.)

Fig. 6(a), 6(b), and 6(c) are images of a man's face, where (a) is the original image, (b) is closed using a ball, and (c) is the residue image, i.e., Fig. 6(a) - Fig. 6(b). Fig. 6(b) is a version of Fig.(a) with some information removed. The removed information is considered to be where the ball would not fit into the sheet that represents the original image. A pixelwise subtract of Fig. 6(a) - Fig.6(b) actually extracts this information and is shown in Fig. 6(c). This type of residue image is a common method of extracting useful information in industrial machine vision algorithms.

The power of these openings and closings is realized in the fact that these operations can be performed in real time using any shape and size of that shape that the application requires.

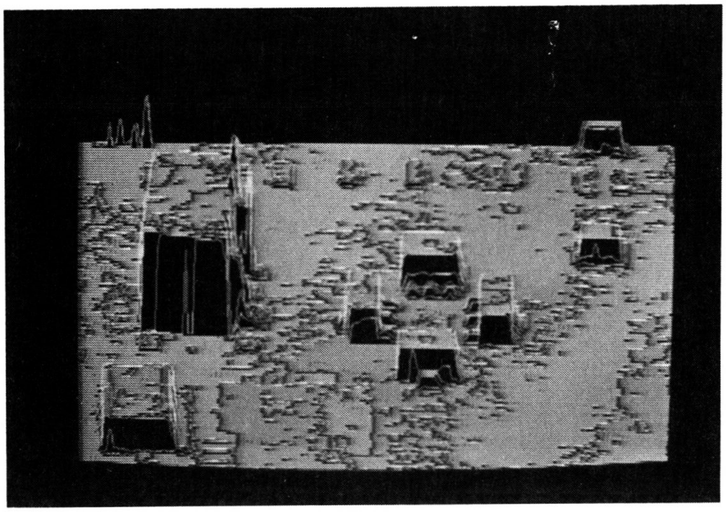

Fig. 4 Oblique graphical representation of a range image demonstrating how an image can be thought of a sheet.

IV. MORPHOLOGICAL PROCESSING OF RANGE IMAGES

Three dimensional information is explicit in a range image. Applying three dimensional operators to this data is demonstrated as a useful approach to transforming three dimensional data. Transforming data in the same coordinate system in which it is originally represented becomes feasible for real time applications with the advent of specialized hardware capable of 3-D morphology. The ability to choose the size and shape of the operator used to transform the image implies some a priori knowledge. Depending on the objective of the processing, this a priori knowledge may be a low level of characterization of the image properties, such as signal to noise ratio, up to a complete 3-D model of what is expected in the scene. To illustrate this we examine the application of component placement verification on printed circuit boards. With automated component placement, it is necessary to inspect the boards visually to ensure that the components are present and properly aligned. Fig. 4 is a graphical representation of a range image that shows components placed on a printed circuit board. This image is shown in Fig. 7(a) displayed as intensity, i.e., the bright areas are nearer the viewer than the dark background. If the range imaging process were ideal, an exact 3-D representation of the scene would be contained in the data. There are some clearly defined reasons why this is not the case (5). The shadow effect, and specular reflectance, are inherent in this technology and cause errors in the data as can be seen in Fig. 4 and Fig. 7(a). It is often the case that the errors are considerably smaller than the objects or features of interest to the application. This presents the familiar problem of processing the image to remove the error that is relatively small, while preserving the information of interest. Typically this low level processing would be followed by model driven, higher level processing. A single morphological operation can accomplish both of these tasks.

Returing to the component inspection application; we know a priori the size of the smallest component expected in the scene. Thus, we can choose a structuring element that is known will fit inside the data representing the smallest component expected in the scene. Using this structuring element, a 3-D morphological closing is performed, which in effect removes all peaks in the data that are smaller than the smallest object of interest. Fig. 7(b) is the result of this closing operation. A threshold is used to create a binary image from this processed image. Fig. 7(c) is the binary image representing the location of the components. Well known connectivity analysis of this binary image quantifies the results which can be compared to the expected data.

This processing is so simple in concept that automatic image processing programming based on a priori knowledge becomes feasible for some highly structured environments. It is not difficult to generate the steps for the application described above, based on a simple data base representation of the printed circuit boards. It is less clear that this is useful for more loosely structured applications, but is worth considering one application at a time.

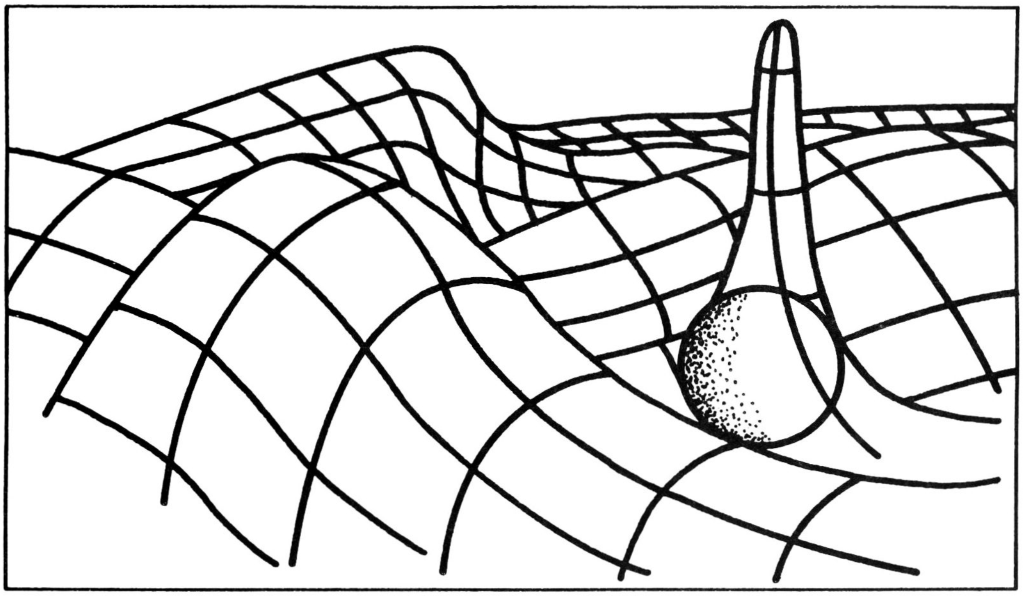

Fig. 5 The rolling ball algorithm illustrated graphically.

Fig. 6(a) A digital image of a man's face.

Fig. 6(b) The image closed using the rolling ball algorithm.

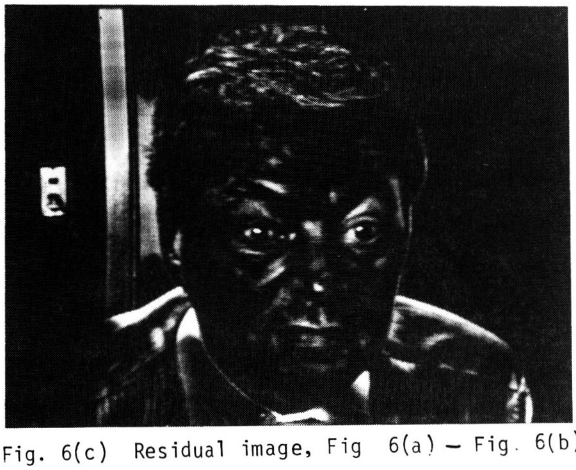

Fig. 6(c) Residual image, Fig 6(a) – Fig. 6(b)

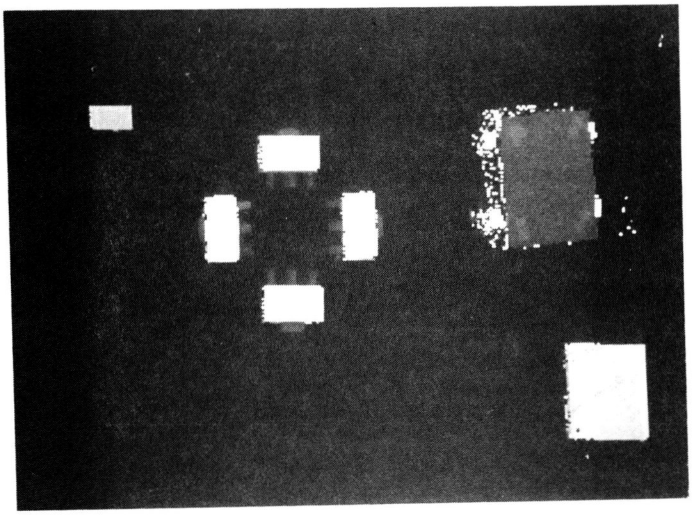

Fig. 7(a) Range image of a surface mount technology board, displayed as intensity. The bright areas are components, raised from the substrate.

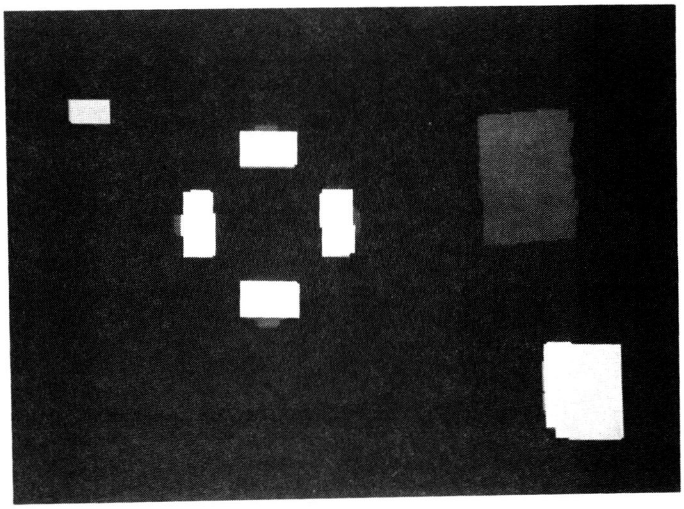

Fig. 7(b) The transformed image, using a morphological closing with a structuring element that is known to fit inside the smallest component.

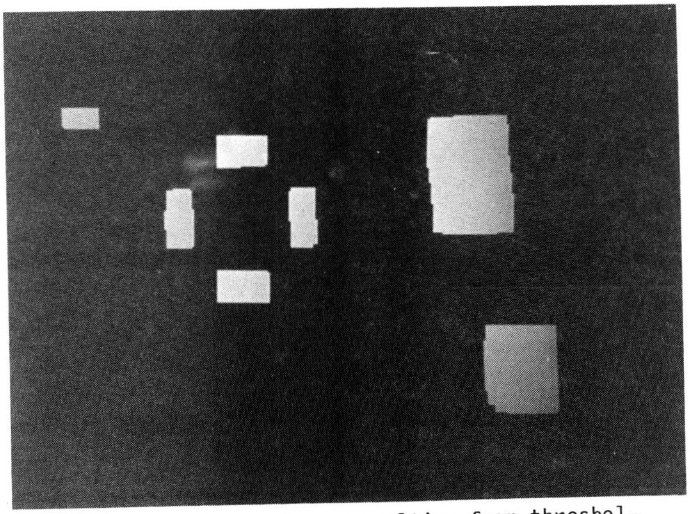

Fig. 7(c) A binary image resulting from thresholding Fig. 7(b).

V. CONCLUSION

Mathematical morphology can be used to transform images in both data driven and model driven processes. Choosing shapes (structuring elements) for data driven transforms is usually based on characteristics of the imaging device, and may accomplish such things as noise removal, edges effect removal, etc. In the model driven part of the algorithm, shapes are chosen based on a priori knowledge of what is expected in the scene. The morphological transforms are applied using these shapes to locate expected attributes in the image.

Morphological transformations are ideal for processing of range images. Industrial applications for verification vision are plentiful, and using the methods described in this paper it is possible to totally automate the inspection task with sufficient a priori knowledge combined with highly reliable range images containing explicit 3-D information. Although this scenario is the most appealing, there is a broad spectrum of applications varying on how much a priori knowledge is available, and the quality and type of images acquired. Mathematical morphology is sufficiently general to be used at all levels of image and scene analysis over this spectrum.

REFERENCES

(1) Serra, J., Omage Analysis and Mathematical Morphology, Academic Press, London,1982.

(2) Sternberg, S., Industrial Morphology, Proc. Interobot West, Long Beach Cal, Oct. 10-12, 1984.

(3) Sternberg, s., Industrial Computer Vision by Mathematical Morphology, Proc. of the 13th Interna- tional Symposium on Industrial Robots, Robot 7, Chi- cago, Illinois, April 18-20, 1983.

(4) Sternberg, S., Grayscale Morphology, in Algorithms in Mathematical Morphology, Special issue of IEEE Computer, J.P. Serra and S.R. Sternberg, editors (in preparation).

(5) Rioux, M., Laser Range Finder Based on Synchronized Scanners, Applied Optics, v. 23, no. 21, Nov. 1984.

(6) Jarvis, R.A., A Perspective on Range Finding Techniques for Computer Vision, IEEE Transactions on Patterns Recognition and Machine Intelligence, Vol. PAMI-5, No. 2.

(7) Archibald, C and Rioux, M., WITNESS: A System for Object Recognition Using Range Images, National Research Council of Canada, Technical Report ERB-986, 1985.

(8) Nitzan, D., Brain, A., and Duda, R., The measurement and Use of Registered Reflectance and Range Data in Scene Analysis, Proc. of the IEEE, Vol. 65, No. 2.

(9) Mitiche, A. and Aggarwal, J.K., Detection of Edges Using Range Information, IEEE PAMI, Vol. PAMI-5, No. 2, 1982.

(10) Inokuchi, S., Nita T., Matsuda, F., and Sakurai, Y., A Three Dimensional Edge-Re- gion Operator for Range Pictures, Proc. of the 6th International Conf. on Pattern Recognition, Munich, Germany, 1982.

(11) Gil, B., Mitiche, A. and Aggarwal, J.K., Experiments Combining Intensity and Range Edge Maps, CVGIP21, 1983.

(12) Duda, R.O., Nitzan and Barett, P., Use of Range and Reflectance Data to Find Planar Surface Regions, IEEE PAMI, Vol. PAMI-1(3), 1979.

(13) Hebert, M. and Ponce, J. A New Method for Segmenting 3-D Scenes into Primitives, Proc. of the 6th International Conf. on Pattern Recognition, Munich, Germany, 1982.

(14) Henderson, T.C. and Bhanu, B., Three Point Seed Method for Extraction of Planar Surfaces from Range Data, Conference Record of the 1982 Workshop on Industrial

(15) Besl, P.J. and Jain, R.C., Three Dimensional Object Recognition, ACM Computing Surveys, Vol. 17, No. 1, March 1985.