

Map Building with Multiple Range Measurements Using Morphological Surface Profile Extraction

Billur Barshan
Department of Electrical Engineering
Bilkent University
Bilkent, 06533 Ankara, Turkey
billur@ee.bilkent.edu.tr

Deniz Başkent
Department of Biomedical Engineering
University of Southern California
Los Angeles, CA 90089-1451 U.S.A.
baskent@scf.usc.edu

Abstract

A novel method is described for surface profile extraction based on morphological processing of multiple range sensor data. The approach taken is extremely flexible and robust, in addition to being simple and straightforward. It can deal with arbitrary numbers and configurations of range sensors as well as synthetic arrays. The method has the intrinsic ability to suppress spurious readings, crosstalk, and higher-order reflections, and process multiple reflections informatively. The essential idea of this work—the use of multiple range sensors combined with morphological processing—can be applied to different physical modalities of range sensing of vastly different scales and in many different areas. These may include radar, sonar, robotics, optical sensing and metrology, remote sensing, ocean surface exploration, geophysical exploration, and acoustic microscopy.

1. Introduction

An inexpensive, yet effective and reliable approach to machine perception is to employ multiple simple range sensors coupled with appropriate data processing. This paper deals with the determination of arbitrary surface profiles, and is completely novel in that morphological processing is applied to range data in the form of an arc map, representing angular uncertainties. The method is extremely flexible and can easily handle arbitrary sensor configurations, as well as synthetic arrays. In contrast, approaches based on geometrical or analytical modeling are often limited to elementary target types or simple sensor configurations [1, 2, 5]. A commonly noted disadvantage of range sensors is the difficulty associated with interpreting spurious read-

ings, crosstalk, higher-order, and multiple reflections. The proposed method is capable of effectively suppressing the first three of these, and informatively processing echoes returning from surface features further away than the nearest (i.e. multiple reflections).

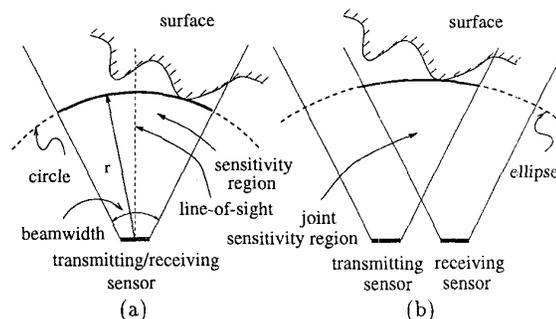


Figure 1. Transmitter and receiver are (a) the same device, (b) separate devices.

Despite the generality of the method, for concreteness, we consider simple range sensors that measure time-of-flight (TOF) t_o , which is the round-trip travel time of the pulse between the sensor and the object. Given the speed of transmission c , the range r can be easily calculated from $r = ct_o/2$. Although such devices return accurate range data, typically they cannot provide direct information on the angular position of the object from which the reflection was obtained. Thus, all that is known is that the reflection point lies on a circular arc of radius r (Fig. 1(a)). More generally, when one sensor transmits and another receives, it is known that the reflection point lies on the arc of an ellipse whose focal points are the transmitting and receiving elements (Fig. 1(b)). The arcs are tangential to the surface at the actual point(s) of reflection.

Most commonly, the wide beamwidth of the sensor is accepted as a device limitation that determines the

angular resolving power of the system, and the reflection point is assumed to be along the line-of-sight. In our method, circular or elliptical arcs, representing the uncertainty of the object location, are drawn. By combining the information inherent in a large number of such arcs, angular resolution far better than that implied by the beamwidth is obtained.

2. Morphological profile extraction

Structured sensor configurations such as linear and circular arrays as well as irregularly configured sensors have been considered in [3], where the method is also generalized to moving sensors and synthetic arrays.

To illustrate the method, Fig. 2(a) shows a surface whose profile is to be determined by using an irregular sensor configuration. A large number of arcs can be obtained with a reasonably small number of sensors because each sensor can receive pulses transmitted from all the others, provided a reflection point lies in the joint sensitivity region for that sensor pair. Fig. 2(b) shows the arcs obtained. Although each arc represents considerable uncertainty as to the angular position of the reflection point, nevertheless one can almost extract the actual curve shown in Fig. 2(a) by visually examining the arc map in Fig. 2(b). Each arc drawn is expected to be tangential to the surface at least at one point. At these actual reflection point(s), several arcs will intersect with small angles at nearby points on the surface. The many small segments of the arcs superimposed in this manner coincide with and cover the actual surface, creating the darker features in Fig. 2(b) that reveal the surface profile. The remaining parts of the arcs, not actually corresponding to any reflections and simply representing the angular uncertainty of the sensors, remain more sparse and isolated. Similarly, those arcs caused by higher-order reflections, crosstalk, and noise also remain sparse and lack much reinforcement.

In this study, morphological operators are used to eliminate the sparse and isolated segments in the arc map, leaving behind the mutually reinforcing segments that reveal the solid structure of the original surface. Erosion, dilation, opening, closing, and thinning are widely used morphological operations to accomplish tasks such as edge detection, skeletonization, segmentation, texture analysis, enhancement, and noise removal in image processing [4]. Most applications involve processing of conventional binary or gray-scale images, or range images where the range information is coded in the gray-levels of the image [7]. The present approach is completely novel in that morphological processing is applied to range measurements in the form of an arc map, representing angular uncertainties.

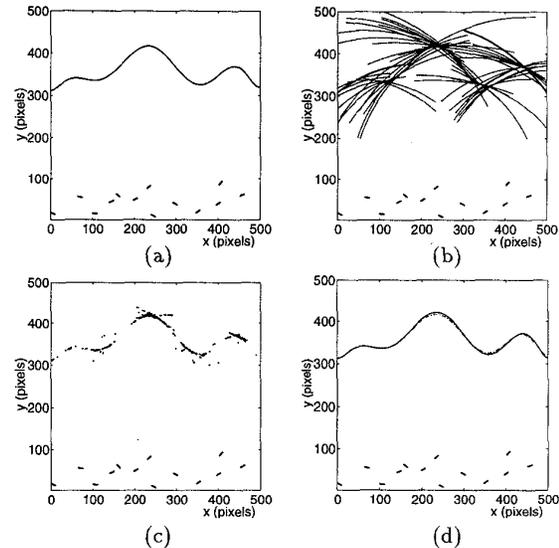


Figure 2. (a) The actual surface and the sensors, (b) the arc map obtained with 17 sensors, each of 45° beamwidth, (c) the result of $n = 6$ thinning, (d) the fitted curve (solid line) and the original surface (dashed line).

Morphological operations basically consist of a set of simple rules to modify images. A simple algorithm for *erosion* is as follows: If all eight neighbors of a pixel with value one equal one, that pixel preserves its value, otherwise its value is set equal to zero. This way, the image will be eroded or shrunk in all directions by one pixel. On the other hand, *dilation* is used to fatten an image: all eight neighbors of those image pixels which originally equal one are set equal to one.

Thinning is a generalization of erosion with a parameter n varying in the range $1 \leq n \leq 8$. In this case, it is sufficient for any n neighbors of a pixel to equal one in order for that pixel to preserve its value of one. The flexibility that comes with this parameter enables one to make more efficient use of the information contained in the arc map. Thus, *pruning* and *erosion* are the two extremes of thinning with $n = 1$ and $n = 8$.

The result of applying $n = 6$ thinning to the arc map shown in Fig. 2(b) is presented in Fig. 2(c). As a last step, a least-squares polynomial fit is obtained to compactly represent the surface profile. The curve fitted to the thinned map in Fig. 2(c) is displayed in Fig. 2(d). A root-mean-square absolute error measure $\mathcal{E} = \sqrt{\frac{1}{N} \sum_{i=1}^N [p(x_i) - y(x_i)]^2}$ is introduced, comparing the final polynomial fit with the actual curve. Here N is the total number of columns in the map matrix, $p(x_i)$ are the samples of the fitted polynomial, and $\sigma_y^2 = \frac{1}{N} \sum_{i=1}^N [y(x_i) - \frac{1}{N} \sum_i y(x_i)]^2$ is the variance of

the actual surface profile $y(x_i)$. The result of applying various morphological operators to the arc map in Fig. 2(b) are summarized in Table 1.

morphological operation	\mathcal{E} (pixels)	t_{CPU} (s)
thinning ($n = 1$: pruning)	20.59	1.21
thinning ($n = 3$)	12.53	1.07
thinning ($n = 5$)	9.19	0.99
thinning ($n = 6$)	2.75	0.98
thinning ($n = 7$)	5.29	0.97
thinning ($n = 8$: erosion)	11.75	0.96
closing & erosion	11.50	5.64

Table 1. Results of various morphological operations.

A detailed study of the performance of different sensor configurations and morphological operations has been performed. Structured arrays are often preferred in theoretical work for simplicity and ease of analysis, whereas the method presented here can handle irregular arrays equally easily. Although the problem of optimal sensor placement is a subject for future research, the large number of simulations performed indicate that it is preferable to work with irregular arrays, since the randomized vantage points of the sensors tend to complement each other better than structured ones. In a typical robotics application, these sensors may constitute an array on the same robot, or may correspond to sensors onboard different robots exploring the environment.

morphological operation	\mathcal{E} (pixels)	t_{CPU} (s)
thinning ($n = 1$: pruning)	4.98	1.91
thinning ($n = 2$)	4.84	1.88
thinning ($n = 3$)	4.07	1.73
thinning ($n = 4$)	3.28	1.82
thinning ($n = 5$)	2.58	1.66
thinning ($n = 6$)	1.96	1.72
thinning ($n = 7$)	1.63	1.80
thinning ($n = 8$: erosion)	1.42	1.77
erosion & pruning ($n = 1$)	1.41	1.89
erosion & thinning ($n = 2$)	1.50	1.93

Table 2. Experimental results for the surface in Fig. 3(a).

The method has been tested with real sonar data, experimentally obtained from a Nomad 200 mobile robot, initially using smooth cardboard surfaces. An example is shown in Fig. 3 for which additional results are provided in Table 2. Even though the method

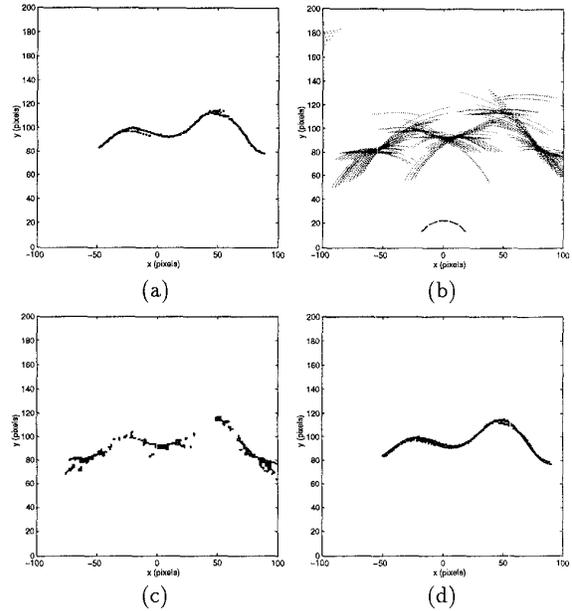


Figure 3. (a) The actual surface. (b) The arc map and the sensors. (c) Result of erosion ($n = 8$) followed by pruning ($n = 1$). (d) Part (a) superimposed with the fitted curve.

was initially developed and demonstrated for specularly reflecting surfaces, subsequent tests with Lambertian surfaces of varying roughness have indicated that the method also works for rough surfaces, with errors slightly increasing with roughness.

3. Performance of the method

Although the method is applicable to arbitrary surfaces [3], to investigate the performance of the method, sinusoidal surfaces have been considered whose parameters can be systematically varied. Simulations have been undertaken on sinusoidal surfaces of varying amplitude and periodicity, located at varying distances from the sensor array. These parameters are illustrated in Fig. 4(a). The elements of the sensor array are distributed in the box $[-35, 440] \times [0, 90]$, with their average distance from $y = 0$ being 32.7 pixels.

We investigate the dependence of the error on amplitude, period, surface distance, sensor beamwidth, and measurement uncertainty. For this purpose, the sinusoid shown in Fig. 4(a), with $A = 30$, $T = 125$, and $L = 200$ pixels, is taken as a reference and these parameters are individually varied around the reference values. The arc map generated is shown in part (b) of the same figure. The result of $n = 3$ thinning, which gives the minimum error for this example, is given in

part (c). The resulting error when various morphological operators are applied to the same arc map are summarized in Table 3. Finally, the result of curve fitting, and the comparison with the actual surface are given in part (d).

First, the period is varied by keeping the amplitude and the surface distance constant at the reference values given above. \mathcal{E} increases with decreasing period as expected (Fig. 5(a)). For periods shorter than 100 pixels, the error increases significantly. The minimum radius of curvature R_{\min} is a useful indicator of the difficulty of extracting the profile: features with smaller radii of curvature are more difficult to accurately determine. For this reason, the relation between R_{\min} and the period of the sinusoid is also plotted in Fig. 5(b).

In the next step, the amplitude is varied while keeping the period and the distance constant at the reference values. \mathcal{E} increases with increasing amplitude (Fig. 5(c)), since increasing A reduces R_{\min} . R_{\min} is plotted as a function of the amplitude in Fig. 5(d).

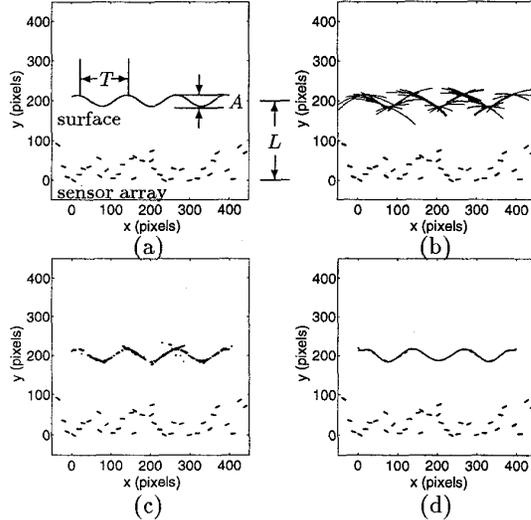


Figure 4. (a) The actual surface and the parameters A , T , and L , (b) the arc map obtained with an array of 35 sensors, each of 30° beamwidth, (c) the result of $n=3$ thinning, (d) the fitted curve (solid line) and the original surface (dashed line).

To get a better understanding of the relation between the error and curvature, the results in Fig. 5 are rearranged to generate a plot of \mathcal{E} versus R_{\min} (Fig. 6). As expected, decreasing the curvature (or increasing R_{\min}) results in lower \mathcal{E} . The fact that the solid and dashed lines (which represent varying T and A respectively) follow each other closely, suggests that what really matters is not the individual values of T and A , but the value of R_{\min} .

morphological operation	\mathcal{E} (pixels)
thinning ($n=1$: pruning)	2.41
thinning ($n=2$)	2.21
thinning ($n=3$)	2.03
thinning ($n=4$)	2.09
thinning ($n=5$)	2.46
closing & pruning ($n=1$)	2.61
closing & thinning ($n=3$)	3.02
closing & erosion ($n=8$)	3.63

Table 3. Results of various morphological operations.

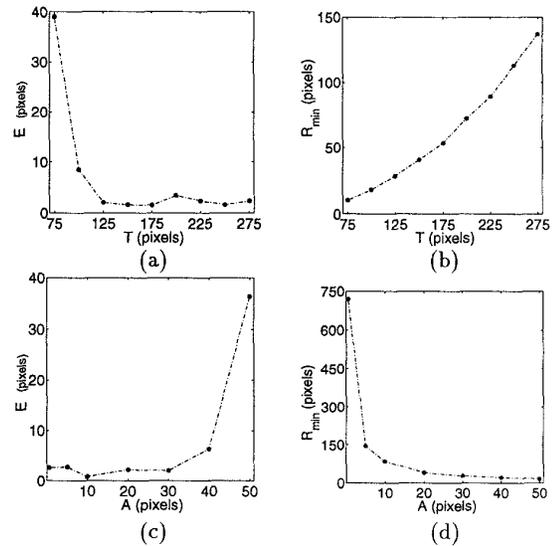


Figure 5. (a) \mathcal{E} vs. T , (b) R_{\min} vs. T , (c) \mathcal{E} vs. A , (d) R_{\min} vs. A .

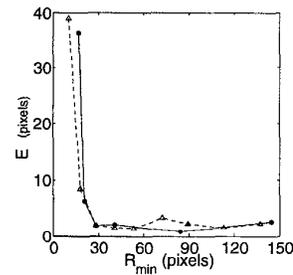


Figure 6. \mathcal{E} vs. R_{\min} . Solid dots connected by solid lines are produced by eliminating T from Fig. 5(a) and (b). Triangles connected by dashed lines are produced by eliminating A from Fig. 5(c) and (d).

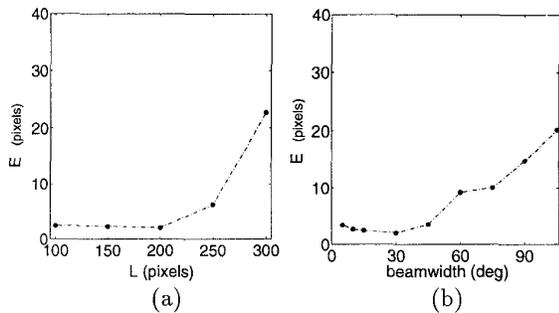


Figure 7. (a) \mathcal{E} vs. L , (b) \mathcal{E} vs. beamwidth.

Next, the distance to the surface is varied while A and T are kept constant at their reference values. \mathcal{E} increases as L increases beyond 250 pixels (Fig. 7(a)). Because the surface shape does not change, the curvature remains constant. Details about the processing involved to generate Fig. 7(a) are presented in Table 4. Since the number of arc points obtained strongly depends on L , and the most suitable morphological operation depends strongly on the density of arc points, the morphological procedure best suited to each value of L has been employed in constructing Fig. 7(a). (In addition to the alternatives shown in Table 3, $n = 6, 7, 8$ thinning, and the application of no morphological processing at all have been considered.) For a given beamwidth, when the surface is located further, the arcs become larger and uncertainty in the position of the reflection point(s) increases. In a way, the “effective” curvature of the surface increases with increasing L , resulting in larger errors. Geometrically, this is the same effect as perceiving a curved object to be flatter when we are very close to it, and more curved when further away. A distinct issue arises when the distances are very small: the arcs become very short in length and less in number, since now sensors can detect a smaller portion of the surface and there is less overlap between their sensitivity patterns. As a result, the arc map cannot cover the whole surface.

Another important parameter is the sensor beamwidth. To investigate the effect of the sensor beamwidth, the surface parameters are kept constant while the beamwidth is varied. Increasing the beamwidth results in arcs longer in length, causing a larger portion of each arc to be redundant. In other words, there is more uncertainty in the position of the reflection point(s) as compared to the case of a narrower beamwidth. As a result, the error increases as shown in Fig. 7(b). The arcs also increase in number, and these factors make it necessary to apply higher- n thinning to extract the useful information. On the other hand, when the beamwidth is very small, the arcs become very short and fewer in number,

L (pixels)	morphological operation	\mathcal{E} (pixels)
100	thinning ($n = 1$)	2.43
150	thinning ($n = 1$)	2.29
200	thinning ($n = 3$)	2.03
250	thinning ($n = 3$)	6.22
300	thinning ($n = 4$)	22.71

Table 4. Results corresponding to Fig. 7(a).

beamwidth	morphological operation	\mathcal{E} (pixels)
5°	none	3.41
10°	none	2.65
15°	none	2.43
30°	thinning ($n = 3$)	2.03
45°	thinning ($n = 5$)	3.51
60°	thinning ($n = 5$)	9.19
75°	thinning ($n = 6$)	10.07
90°	thinning ($n = 7$)	14.82
105°	thinning ($n = 8$)	20.21

Table 5. Results corresponding to Fig. 7(b).

leading to a similar situation as when L was very small. Below a beamwidth of 15°, directly fitting a polynomial to whatever few points are available in the arc map, without applying morphological processing, becomes the best choice. This customization of the applied morphological rule enables a fair comparison of the results at all beamwidth values. Smaller beamwidths result in fewer arc points and thus less reliable curve fits, leading to a slight increase in the error for very small beamwidths. Best results are obtained for a particular beamwidth (about 30° in our example). The different morphological operations applied and the resulting error values are tabulated in Table 5. Choosing beamwidths smaller than 30° does not increase the error appreciably, however using sensors with smaller beamwidths may not be desirable anyhow, since these are usually more difficult to manufacture, expensive, or entail a trade-off with some other quantity. For instance, in the case of acoustic sensors, narrower beamwidth devices must have higher operating frequencies, which imply greater attenuation in air and shorter operating range.

Now, we discuss the issue of choice of sampling resolution or pixel size: There are a couple of factors that determine the accuracy of TOF readings in a range measurement system. One of these factors is the operating wavelength of the ranging system: a TOF mea-

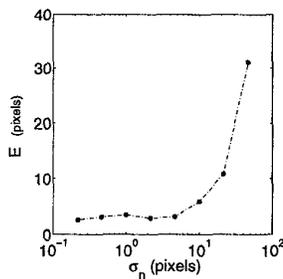


Figure 8. \mathcal{E} vs. noise standard deviation.

surement with accuracy better than the wavelength cannot be normally achieved. Other sources of uncertainty in the range measurement could be effects such as the thermal noise in the receiving circuitry or the ambient noise. Given these, it is not meaningful to choose the pixel size much smaller than the resolving limit determined by these factors, since it would increase the computational burden without resulting in a more accurate profile determination. Thus, the pixel size should be chosen comparable to the TOF measurement accuracy. Nevertheless, since the TOF accuracy may not be known beforehand, we have examined the cases where the noise or uncertainty is both smaller and larger than one pixel.

To investigate the robustness of the method to noise, zero-mean white Gaussian noise has been added to the TOF readings. The noise standard deviation σ_n is varied logarithmically to cover a broad range of noise levels. As expected, for σ_n smaller than one pixel, the performance is approximately the same as for the noiseless case. This performance can be further improved by reducing the pixel size until it becomes comparable to the TOF measurement accuracy, at the cost of greater computation time.

The error increases significantly as the noise level increases beyond one pixel (Fig. 8). Since the method relies on the mutual reinforcement of several arcs to reveal the surface, larger amounts of noise are expected to have a destructive effect on this process by moving the various arc segments out of their reinforcing positions. Consequently, the arc segments which now lack each other's mutual reinforcement tend to be eliminated by the morphological operations. A larger proportion of the arcs is eliminated, resulting in a loss of information characterizing the original curve. Nevertheless, the error growth rate is not as high as might be suggested by these arguments, and the method seems to be reasonably robust to measurement uncertainty. In Fig. 8, the performance is comparable to the noiseless case up to $\sigma_n = 10$ pixels. This is partly because the least-squares polynomial fit helps eliminate some of the noise.

4. Conclusions

A novel method is described for determining arbitrary surface profiles by applying morphological processing to data acquired by simple range sensors. The method is extremely flexible, versatile, and robust, as well as being simple and straightforward. It can deal with arbitrary numbers and configurations of sensors, including synthetic arrays. Accuracy improves with the number of sensors used and can be as low as a few pixels except when the radius of the curvature is very small. The method is robust in many aspects: it has the inherent ability to eliminate undesired TOF readings arising from higher-order reflections, crosstalk, and noise, as well as processing multiple echoes informatively.

The CPU times for the morphological operations (when implemented in the C programming language and run on a 200 MHz Pentium Pro PC) are generally about a quarter of a second [3], indicating that the method is viable for real-time applications. The method can be readily generalized to 3-D environments with the arcs replaced by spherical or elliptical caps and the morphological rules extended to 3-D [6]. In certain problems, it may be preferable to reformulate the method in polar or spherical coordinates. Some applications may involve an inhomogeneous and/or anisotropic medium of propagation. It is envisioned that the method could be generalized in such cases by constructing broken or non-ellipsoidal arcs.

References

- [1] B. Barshan and R. Kuc. Differentiating sonar reflections from corners and planes by employing an intelligent sensor. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(6):560–569, June 1990.
- [2] C. S. Chen, Y. P. Hung, and J. L. Wu. Extraction of corner-edge-surface structure from range images using mathematical morphology. *IEICE Trans. on Information and Systems*, E78D(12):1636–1641, 1995.
- [3] D. Bařkent and B. Barshan. Surface profile determination from multiple sonar data using morphological processing. *Int. J. of Robotics Research*, to appear in 2000.
- [4] E. R. Dougherty. *An Introduction to Morphological Image Processing*. SPIE Optical Engineering Press, Bellingham, WA, U.S.A., 1992.
- [5] L. Kleeman and R. Kuc. Mobile robot sonar for target localization and classification. *Int. J. of Robotics Research*, 14(4):295–318, August 1995.
- [6] K. Preston. 3-dimensional mathematical morphology. *Image and Vision Computing*, 9(5):285–295, 1991.
- [7] J. G. Verly and R. L. Delanoy. Some principles and applications of adaptive mathematical morphology for range imagery. *Opt. Eng.*, 32(12):3295–3306, 1993.