Use of an opacity constraint in three-dimensional imaging

Richard G. Paxman, John H. Seldin, James R. Fienup, and Joseph C. Marron

Optical & Infrared Science Laboratory, Advanced Concepts Division
Environmental Research Institute of Michigan
P.O. Box 134001, Ann Arbor, Michigan 48113-4001
internet email: paxman@erim.org

ABSTRACT

Three-dimensional imaging provides profile information not available with conventional two-dimensional imaging. Many three-dimensional objects of interest are opaque to the illuminating radiation, meaning that the object exhibits surface, as opposed to volume, scattering. We investigate the use of an opacity constraint to perform three-dimensional phase retrieval. The use of an opacity constraint in conjunction with frequency-diverse pupil-plane speckle measurements to reconstruct a three-dimensional object constitutes a novel unconventional-imaging concept. This imaging modality avoids the difficulties associated with making phase measurements at a cost of increased computations.

1. INTRODUCTION

Marron and Schroeder [1,2] have demonstrated that three-dimensional imaging can be accomplished by sequentially illuminating an object with different laser frequencies and measuring the far-field speckle pattern for each of the illuminating frequencies [1]. In-phase and quadrature measurements provide optical-field values for the laser-speckle patterns. Properly formatted, these data form a three-dimensional Fourier-volume (or Fourier-aperture) representation of the illuminated object. A simple three-dimensional Discrete Fourier Transform (DFT) can be performed to provide a three-dimensional representation of the object. This three-dimensional lensless-imaging concept is referred to as Holographic Laser Radar (HLR).

The collection of HLR data could be considerably simplified if instead of collecting field measurements for the speckle images, intensity measurements were collected. Intensity measurements are straightforward and eliminate the need for interference with a reference beam with precision alignment, tracking, and phase stability. A multiple-frequency speckle-intensity data set would provide three-dimensional Fourier-magnitude information or, equivalently, the three-dimensional object autocorrelation. In order to recover a literal three-dimensional image, a phase-retrieval algorithm is required. Therefore, this proposed imaging concept trades complexity and cost in hardware for increased computing.

Phase retrieval requires some type of *a priori* information about the object. Two-dimensional complex-valued objects have been recovered using phase retrieval with a support constraint [3], although this is a challenging problem. It is well known that the uniqueness properties of two-dimensional phase retrieval are much better than for the one-dimensional problem [4]. We conjecture

0-8194-1545-6/94/$6.00

that three-dimensional phase retrieval with a support constraint is better conditioned than its two-dimensional counterpart.

An additional constraint that has great promise in the three-dimensional imaging case is an *opacity* constraint. An opaque object is one that exhibits only surface scattering and no volume scattering (over volumes that extend beyond the desired range resolution). The reflectivity function for an opaque object is confined to a two-dimensional curved (possibly discontinuous) surface, embedded in a three-dimensional space. In this paper we explore the use of an opacity constraint to perform phase retrieval. We have previously used an opacity constraint to perform superresolution [5].

The opacity constraint is a special type of support constraint. It is a "quality of support" constraint – the actual location of the support is not given, although the object is known to be confined to a two-dimensional curved surface. This constraint promises to be very powerful since it greatly reduces the class of feasible objects from which to choose an estimate. Moreover, there are many imaging applications in which the objects will be known with confidence to be opaque. For example, space objects, balistic missiles, aircraft, and a multitude of industrial-inspection parts qualify as opaque objects. Most objects in our everyday experience satisfy the opaque condition. The constraint is invalid for objects with distributed volume scatterers such as translucent or fog-like objects.

The use of an opacity constraint in conjunction with frequency-diverse pupil-plane speckle measurements to reconstruct a three-dimensional object constitutes a novel unconventional-imaging concept.

2. REPRESENTATION OF OPAQUE OBJECTS

Consider a three-dimensional opaque object, $f(x,y,z)$, defined on an object-centered coordinate system. Let x and y be the cross-range (angle-angle) coordinates and let z be the range coordinate that is co-aligned with the illumination direction. Since the collected data will be sampled and since object reconstructions must be performed with a digital computer, we will adopt a discrete representation for the object. Accordingly, we require the coordinates (x,y,z) to take on integer values. Because of the opaque nature of the object, only radiation from reflecting sources in the illuminated surface will contribute to the received signal. Hidden surfaces do not contribute. Let $h(x,y)$ denote the height of the object in the z-dimension. Because the object is confined to a two-dimensional curved surface embedded in three-dimensional space, it can be represented with delta-function notation [5,6],

$$f(x,y,z) \;=\; r(x,y)\delta[z-h(x,y)]\;, \qquad x,y,z \in \{0,1,2,\ldots N-1\}\;, \tag{1}$$

where $r(x,y)$ is the complex surface reflectivity, N is the total number of samples in each dimension, and $\delta(\cdot)$ is used to represent the Kronecker delta,

$$\delta(z-z_o) \;\equiv\; \begin{cases} 1, & z=z_o \\ 0, & z\neq z_o \end{cases}\;. \tag{2}$$

Recall that the collected Fourier intensity data can be transformed to compute a three-dimensional object autocorrelation. In order to avoid aliasing, we require the object to have finite support with sufficient zero padding,

$$f(x,y,z) \;=\; 0\;, \qquad x,y,z \in \left\{\frac{N}{2}, \frac{N}{2}+1, \ldots, N-1\right\}\;. \tag{3}$$

The rectilinear object-support bound given in Eq. (3) is not likely to be a "tight" bound. More restrictive or tighter bounds on object support can be determined from the known object autocorrelation by using a generalization of methods developed for two-dimensional phase retrieval [7]. Any three-dimensional support bound can be expressed in terms of a two-dimensional angle-angle support bound and a one-dimensional range support bound that depends upon angle-angle position.

3. STATEMENT OF PROBLEM

The Fourier representation of the object is found with the DFT,

$$F(u,v,w) \;=\; \sum_{x=0}^{N-1}\sum_{y=0}^{N-1}\sum_{z=0}^{N-1} f(x,y,z)\exp\{-i2\pi(ux+vy+wz)/N\} \tag{4}$$

$$\;=\; \sum_{x=0}^{N-1}\sum_{y=0}^{N-1}\sum_{z=0}^{N-1} r(x,y)\delta[z-h(x,y)]\exp\{-i2\pi(ux+vy+wz)/N\} \tag{5}$$

$$\;=\; \sum_{x=0}^{N-1}\sum_{y=0}^{N-1} r(x,y)\exp\{-i2\pi[ux+vy+wh(x,y)]/N\}\;, \tag{6}$$

where we have used the Kronecker delta to eliminate the summation over z. Since we detect intensities in the Fourier domain, a noiseless measurement would be given by the squared modulus of the DFT of the object, $|F(u,v,w)|^2$. The actual detected data will be corrupted by noise. The detected data are represented by

$$D(u,v,w) \;=\; \mathcal{N}\{|F(u,v,w)|^2\}\;, \tag{7}$$

where the noise operator $\mathcal{N}\{\cdot\}$ corrupts the argument according to an appropriate noise model. Additive Gaussian or Poisson noise models are appropriate when detector readout noise or photon noise, respectively, are the dominant noise sources.

The problem that we wish to address can now be stated. Given the data, $D(u,v,w)$, estimate the object height function, $h(x,y)$, and the object surface reflectivity function, $r(x,y)$, over the angle-angle object-support bound. Performing such an estimate implicitly provides an estimate for the object's Fourier phase. Therefore this problem may be properly viewed as a phase-retrieval problem.

4. MAXIMUM-LIKELIHOOD ESTIMATION

When detector readout noise dominates, the noise operator has the effect of adding a realization of a Gaussian noise process,

$$D(u,v,w) \;=\; |F(u,v,w)|^2 + n(u,v,w)\;. \tag{8}$$

Assuming a common noise variance and statistical independence between differing samples of the noise process, it is straightforward to express the probability density function for the recorded data [8] . The associated log-likelihood function is readily found to be [8]

$$L[\hat{r}(x,y),\hat{h}(x,y)] = -\sum_{u=0}^{N-1}\sum_{v=0}^{N-1}\sum_{w=0}^{N-1}\left[D(u,v,w) - |\hat{F}(u,v,w)|^2\right]^2 \tag{9}$$

$$= -\sum_{u=0}^{N-1}\sum_{v=0}^{N-1}\sum_{w=0}^{N-1}\left[D(u,v,w) - \left|\sum_{x=0}^{N-1}\sum_{y=0}^{N-1}\hat{r}(x,y)\exp\{-i2\pi[ux+vy+w\hat{h}(x,y)]/N\}\right|^2\right]^2 , \tag{10}$$

where the caret symbol indicates an estimated quantity. Recall that the object surface reflectivity is complex valued. It is convenient to explicitly denote the real and imaginary parts of the reflectivity,

$$r(x,y) \;=\; r_r(x,y) + ir_i(x,y)\;. \tag{11}$$

We can use standard nonlinear optimization methods to search for $\hat{r}_r(x,y)$, $\hat{r}_i(x,y)$, and $\hat{h}(x,y)$ ($(x,y) \in$ the angle-angle object-support bound) that maximize the objective function, L, to yield a maximum-likelihood estimate. If the angle-angle object support bound is comprised of $N/2$ by $N/2$ samples, then L is defined on a parameter space of dimension $3N^2/4$. The range support bound that derives from the three-dimensional support bound could be used to constrain the search over $\hat{h}(x,y)$.

There are a variety of nonlinear optimization algorithms that could be used to maximize the log-likelihood function. One choice is the well-known conjugate-gradients algorithm [9]. As its name suggests, this algorithm makes use of the gradient of the objective function. We have derived closed-form expressions for the partial derivatives that constitute the gradient. These expressions afford a gradient computation that is more accurate and computationally much more efficient than would be rendered using the method of finite differences. The partial derivatives of the log-likelihood function with respect to the real and imaginary parts of the surface reflectivity and with respect to the height function, at the angle-angle location (x_o, y_o), are given by

$$\frac{\partial L}{\partial \hat{r}_r(x_o,y_o)} =$$

$$4\mathrm{Re}\left\{\sum_{u=0}^{N-1}\sum_{v=0}^{N-1}\sum_{w=0}^{N-1}\left[D(u,v,w) - |\hat{F}(u,v,w)|^2\right]\hat{F}^*(u,v,w)\exp\{-i2\pi[ux_o+vy_o+w\hat{h}(x_o,y_o)]/N\}\right\}, \tag{12}$$

$$\frac{\partial L}{\partial \hat{r}_i(x_o,y_o)} =$$

$$-4\mathrm{Im}\left\{\sum_{u=0}^{N-1}\sum_{v=0}^{N-1}\sum_{w=0}^{N-1}\left[D(u,v,w) - |\hat{F}(u,v,w)|^2\right]\hat{F}^*(u,v,w)\exp\{-i2\pi[ux_o+vy_o+w\hat{h}(x_o,y_o)]/N\}\right\}, \tag{13}$$

$$\frac{\partial L}{\partial \hat{h}(x_o, y_o)} = \frac{8\pi}{N} \mathrm{Im} \left\{ \hat{r}(x_o, y_o) \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} \sum_{w=0}^{N-1} \left[D(u,v,w) - |\hat{F}(u,v,w)|^2 \right] \right.$$
$$\left. \hat{F}^*(u,v,w)\, w \exp\{-i2\pi[ux_o + vy_o + w\hat{h}(x_o, y_o)]/N\} \right\}, \quad (14)$$

where the operators $\mathrm{Re}\{\cdot\}$ and $\mathrm{Im}\{\cdot\}$ take the real and imaginary parts of the argument, respectively. Notice that these partial-derivative expressions each have the form of a two-dimensional DFT (over u and v) followed by a summation over w for each (x_o, y_o). This affords efficient evaluation the entire gradient using the Fast Fourier Transform (FFT). The three-dimensional function $\hat{F}(u,v,w)$ must be computed to evaluate the objective function and the gradient. We find that N two-dimensional FFTs are needed to compute $\hat{F}(u,v,w)$. Careful examination of Eqs. (12)–(14) reveals that the entire gradient can be computed with an additional N two-dimensional FFTs followed by $2SN^2$ summations over w, where S is the proportion of pixels in the angle-angle object support relative to the number of samples in the total angular field of view. In the case of the rectilinear support expressed in Eq. (3), $S = 1/4$. Of course additional computational overhead is required to perform complex multiplies and sums. Closed-form expressions for the log-likelihood and its gradient can also be derived for the case of photon noise.

5. SIMULATIONS

We have performed two simulations that demonstrate three-dimensional phase retrieval using an opacity constraint. For both of the simulations we chose $N = 32$ and a maximum object extent of 16 samples in each dimension, as required by Eq. (3). This guarantees that the Fourier-intensity data are adequately sampled. In addition, for this initial demonstration we used noiseless Fourier-intensity data, giving $n(u,v,w) = 0$ in Eq. (8).

The first simulation utilized a simple object consisting of 6 separated points of equal reflecting strength. The angle-angle (x-y) view of the magnitude of the surface reflectivity is shown in Figure 1(A). The phases of these points are independent, identically-distributed samples from a probability distribution function (PDF) that is uniformly distributed on the interval $[0, 2\pi]$. The height distribution and reflectivity of these points are summarized in Table 1. Coordinates are referenced from $(0,0)$ in the upper left corner. Despite the discrete representation of the object in the x and y dimensions, the z-dimension heights, $h(x,y)$, are not restricted to integer values.

The phase-retrieval problem consists of estimating a set of real-valued quantities, $r_r(x,y)$, $r_i(x,y)$, and $h(x,y)$, defined on a discrete (x,y) coordinate system. The conjugate-gradients algorithm, used to estimate the object parameters, is initiated with a first guess of the unknown parameters. We began with an initial estimate of $h(x,y) = 0$ and independent, identically distributed samples from a PDF that is uniform on the interval $[0,1]$ for $r_r(x,y)$ and $r_i(x,y)$. The magnitude and phase of the initial reflectivity guess are shown in Figures 1(B) and 1(C), respectively. The magnitude of the initial reflectivity guess ranges between 0.083 and 1.34, and the phase varies from 0.001 to 1.567 radians. By restricting the initial guess of $r_r(x,y)$ and $r_i(x,y)$ to positive

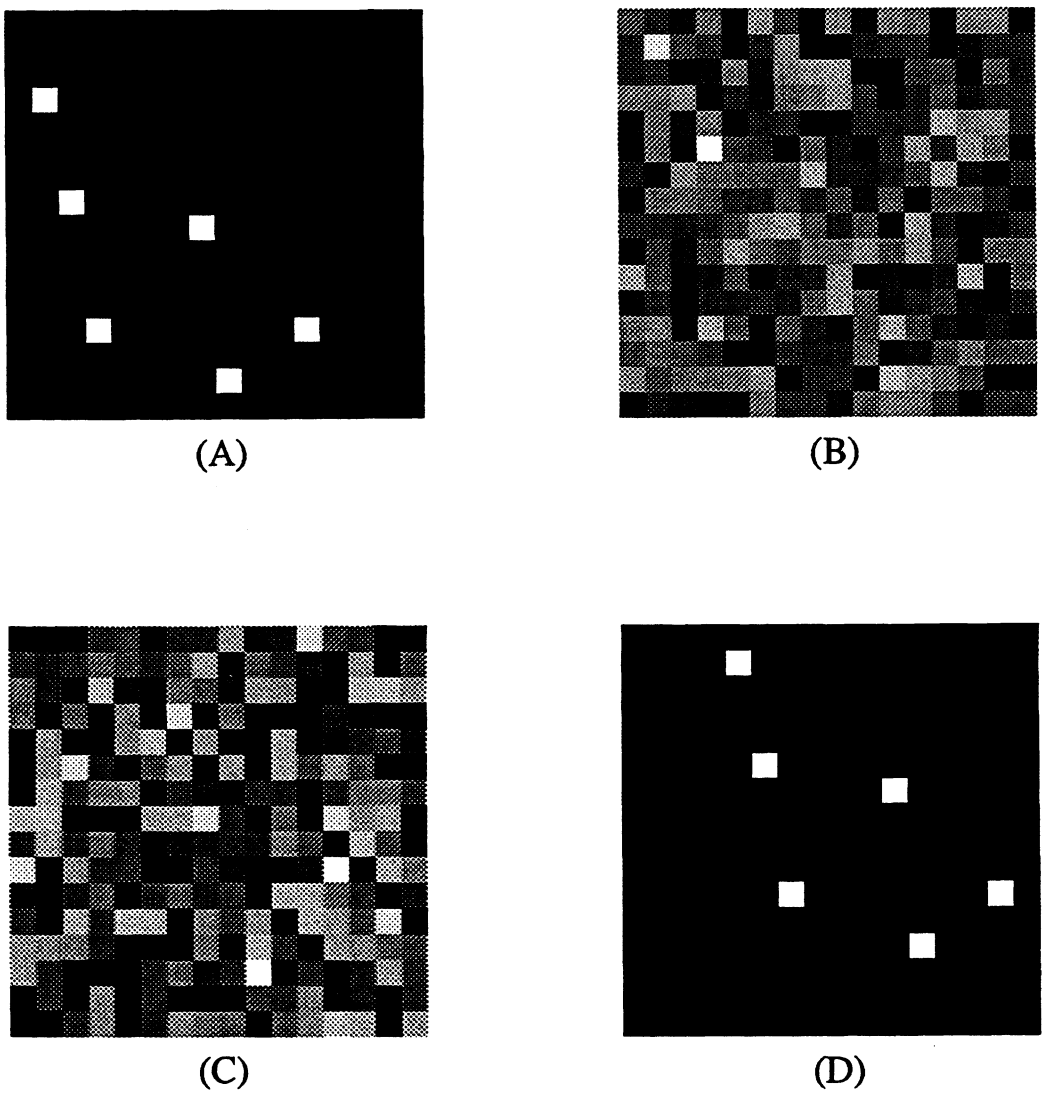

Figure 1. 6-point object simulation. (A) Angle-angle (x-y) view of the magnitude of the reflectivity of a 6-point, complex-valued object; (B) Magnitude of the initial reflectivity estimate; (C) Phase of the initial reflectivity estimate; (D) Reconstructed object.

(x,y)	$r_r(x,y)$	$r_i(x,y)$	$h(x,y)$
(1,3)	-0.976	0.216	0.0
(2,7)	0.411	0.911	0.3
(7,8)	0.994	0.111	0.9
(3,12)	-0.988	0.153	1.5
(8,14)	0.745	-0.666	2.1
(11,12)	0.089	0.996	3.0

Table 1: 6-Point Object Description.

values, we have reduced the chance of biasing the simulation by starting too close to the true solution. No *a priori* knowledge of the support of the unknown object was used during the estimation of its reflectivity and height distribution, so a total of $16 \times 16 \times 3 = 768$ parameters were estimated simultaneously.

The angle-angle view of the estimated reflectivity magnitude after about 400 iterations of the conjugate-gradients algorithm is shown in Figure 1(D). The reconstruction is virtually perfect, and is limited only by computational precision. There are several features of the reconstruction to note. First, the reconstructed object is shifted from its true location. A feature of all phase retrieval algorithms is a translation ambiguity among the set of valid object solutions. A translation of the object adds a linear phase term to the complex-valued Fourier transform, but does not affect the Fourier intensity. There is also a 180-degree object-rotation ambiguity. Such a rotation gives rise to a Fourier transform that is the complex-conjugate of the original Fourier transform, but which leaves the Fourier intensity unchanged.

The height estimates were also virtually perfect. It is of interest that the non-integer height values were estimated very accurately, indicating that in addition to phase retrieval, we were accomplishing superresolution in range. Not only is the object estimate translated in (x,y), but it is also translated in the z direction by a constant value of -1.2. Also, whereas the magnitude of the estimated reflectivity matches the true reflectivity, the phase of the estimate differs by the constant value of 0.64 radians at each point. This is yet another ambiguity associated with the phase retrieval of complex-valued objects. A constant phase bias across the object, however, is typically irrelevant.

For the case of noiseless Fourier-intensity measurements we have shown that an array of points with a modest height distribution can be reconstructed perfectly with no *a priori* knowledge about the support or height distribution of the points, except for the use of an opacity constraint. We have, however, found that the conjugate-gradients algorithm is subject to stagnation in what we believe to be local extrema in the log-likelihood function, and a successful reconstruction was obtained only after several random initial guesses for the reflectivity. Also, we have observed that the relative distribution of the points in the z-dimension has a bearing on the success of the reconstruction. We chose a distribution over a range of 3 samples, and we have noted that as this range is expanded, stagnation becomes more likely when using a naive initial guess of $h(x,y) = 0$. Thus, we believe

that the development of methods that yield improved initial estimates and support bounds that will constrain the range of height parameters will be important for successful reconstructions for surfaces that have large variation in depth. Our second simulation provides an example of how such a height constraint could be derived and utilized.

The second simulation example involves an object that is a curved surface with an angle-angle (x-y) support that is a 5×5-pixel right triangle. To simplify the height distribution, we made the height of the object a function of the x-dimension alone. The support of the object projected in each of the 3 directions is shown in Figure 2. When the object is projected in the y-direction, we obtain a single pixel of z-direction support for each position in the x dimension, as illustrated in Figure 2(A). We conclude from this that the height is constant in y for each x and is symmetric about the $x = 8$ plane. The height of the object varies over a 5-pixel interval, with a 2-pixel height gap between the x-dimension edges and their neighbors.

If no support constraint is used along with a naive initial guess of $h(x, y) = 0$, the parameter estimation stagnates and the object cannot be reconstructed. Thus, the reconstruction of this more complicated object might benefit from a support constraint. Techniques for bounding the object support using the autocorrelation support have been developed for two-dimensional phase-retrieval problems [7]. We can apply these same techniques to the support of the three-dimensional autocorrelation projected into 2 dimensions in each of the 3 directions. Unfortunately, in doing so we lose registration information between the 3 resulting 2-D support constraints and, thus, cannot merge these results into a tight 3-D support constraint. However, for this particular object for which the height is constant along the y-direction, the triple-intersection rule [7] applied to the support of the y-projected autocorrelation yields the exact z-support shown in Figure 2(A). This somewhat surprising result suggests that opacity could be extremely useful in finding support bounds. The derived height support in turn yields the true solution for the height and indicates that the x-dimension support is 5-pixels wide. Merging this information with the support of the autocorrelation projected along the z-direction into the x-y plane yields the x-y object support bound shown in Figure 3(A). This bound could further be reduced using a two-point intersection rule, but we found that this was not necessary.

The magnitude of the reflectivity of the true object is shown in Figure 3(B). The initial guess for the reflectivity is found in Figure 3(C), and the initial guess for the height was set to the values yielded from the x-z support derived from the triple-intersection of the projected autocorrelation. The initial guess of the height was set constant in the y-direction over the support defined in Figure 3(A). The x-y support constraint reduces the number of parameters to be estimated to 114. After approximately 800 iterations of the conjugate-gradients algorithm, the estimates were perfect to within the limits of computational precision. Despite our knowledge of the height from the support estimate, we still allowed the height estimates to vary during minimization. The final estimates actually moved a fixed distance from their initial settings, but remained constant in the y-direction. The final estimate of the reflectivity is shown in Figure 3(D), and is essentially identical to the true reflectivity in Figure 3(B).

Although it is unlikely that the height of the object can be estimated directly from the projected

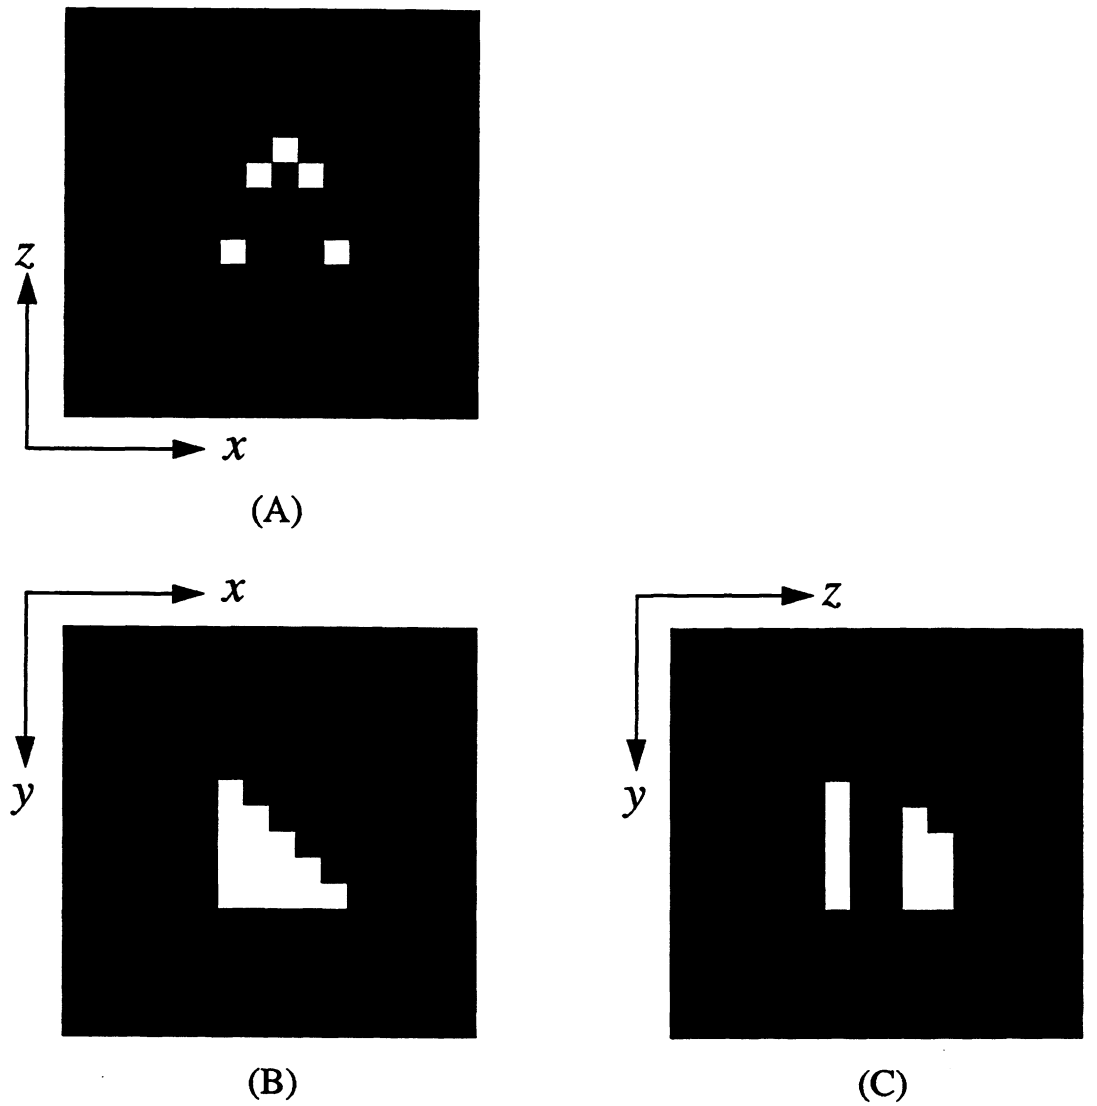

Figure 2. Three views of the support of a triangular curved surface. (A) x-z view of the surface. From this view we see that the height distribution is the same for every value of y and is symmetric about the x=8 plane; (B) x-y view of the surface; (C) y-z view of the surface.

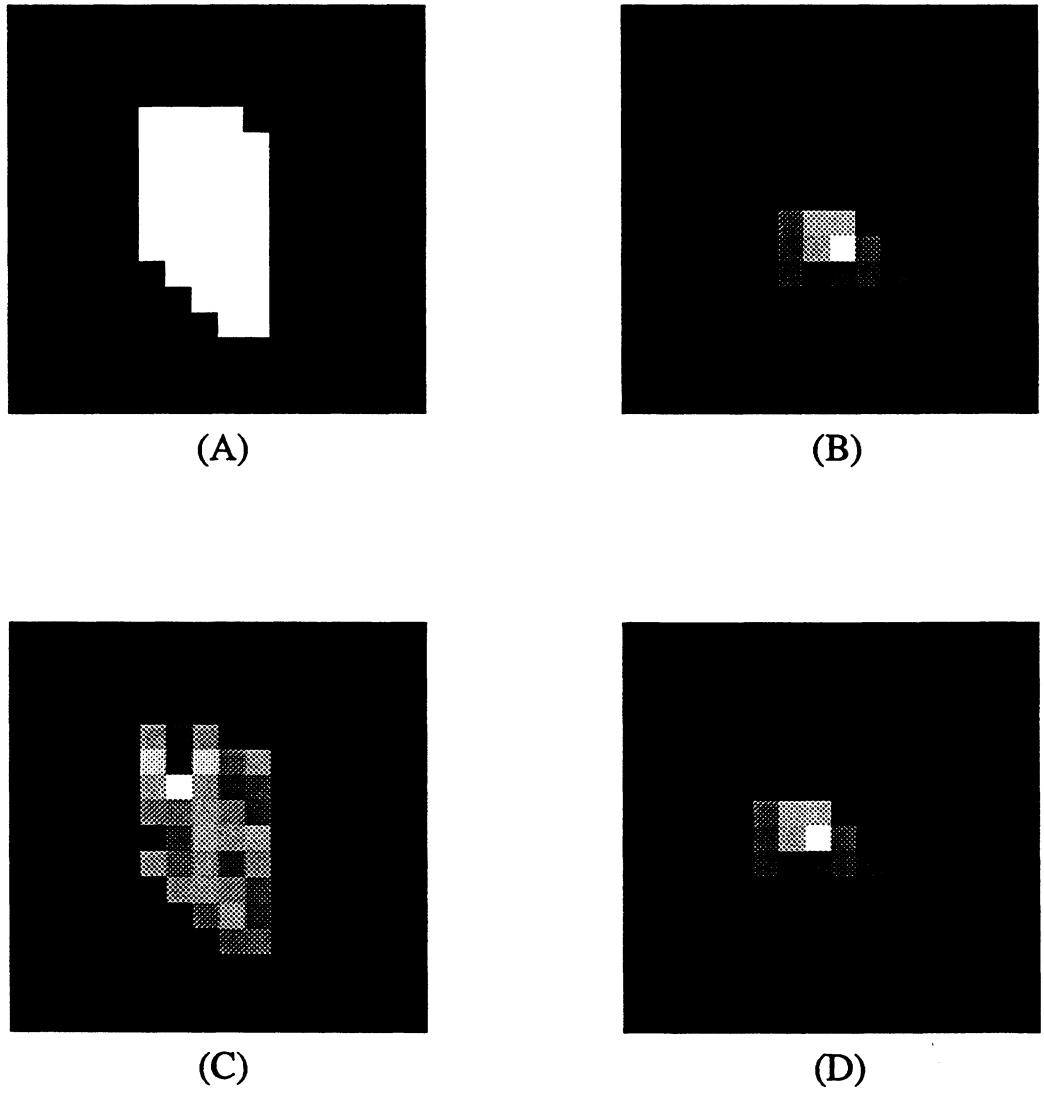

Figure 3. Reconstruction of triangular curved surface. (A) x-y support bound. The bound was obtained by intersecting the estimated x-z support bound [same as shown in Figure 2(A)] with the x-y autocorrelation support; (B) Magnitude of the true reflectivity; (C) Magnitude of the initial guess on the support of (A); (D) Magnitude of reconstructed reflectivity.

autocorrelation in practical scenarios involving measurement noise and 2-D height variation (instead of the 1-D variation of this example), this example demonstrates that a filled surface reflectivity can be estimated with a loose angle-angle support.

6. CONCLUSIONS

We have reported the first demonstration of three-dimensional phase retrieval with an opacity constraint. These results make feasible the possibility of an entirely new three-dimensional imaging modality. Further experimentation is required to determine how loose the height constraint can be made and the importance of the role of the initial height estimate. We believe that the theory of two-dimensional support bounds [7] can be generalized to the three-dimensional case to give good support bounds. We also conjecture that opacity will play a role to help tighten range-dimension support bounds. Historically modes of stagnation have been studied in the two-dimensional phase-retrieval problem and appropriate algorithms to avoid such modes have been developed. A similar process needs to be undertaken in the three-dimensional case. Finally, the sensitivity of three-dimensional phase retrieval to noise needs to be investigated.

7. ACKNOWLEDGEMENTS

We wish to recognize Stanley R. Robinson who originally suggested that the opacity constraint be considered for use in phase-retrieval. We are also grateful for useful discussions with Brian J. Thelen. This research was supported by the Innovative Science and Technology office of the Ballistic Missile Defense Organization of the Office of Naval Research.

REFERENCES

1. J.C. Marron and K.S. Schroeder, "Three-dimensional lensless imaging using laser frequency diversity," Appl. Opt. **31**, 255-262 (1992).
2. J.C. Marron and K.S. Schroeder, "Holographic Laser Radar," Opt. Lett. **18**, 385-387 (1993)
3. J.N. Cederquist, J.R. Fienup, J.C. Marron, and R.G. Paxman, "Phase retrieval from experimental far-field speckle data," Opt. Lett. **13**, 619-621 (1988).
4. J.C. Dainty and J.R. Fienup, "Phase Retrieval and Image Reconstruction for Astronomy," in Image Recovery: Theory and Application, H. Stark, ed. (Academic Press, New York, 1987).
5. R.G. Paxman, "Superresolution with an opacity constraint," *Topical Meeting on Signal Recovery and Synthesis III*, Technical Digest Series **15**, (Optical Society of America, Washington DC, 1989), PD1, North Falmouth, Cape Cod, MA, June 1989.
6. N. Bleistein, *Mathematical Methods for Wave Phenomena* (Academic Press, New York, 1984).
7. T.R. Crimmins, J.R. Fienup, and B.J. Thelen, "Improved bounds on object support from autocorrelation support and application to phase retrieval," J. Opt. Soc. Am. A **7**, 3-13 (1990).
8. R.G. Paxman, T.J. Schulz, and J.R. Fienup, "Joint estimation of object and aberrations by using phase diversity," J. Opt. Soc. Am. A **9**, 1072-1085 (1992).
9. D.C. Luenberger, *Linear and Nonlinear Programming* (Addison-Wesley, Reading, MA, 1984).