
High-Definition Vector Imaging

Gerald R. Benitz

■ High-definition vector imaging (HDVI) is a data-adaptive approach to synthetic-aperture radar (SAR) image reconstruction based on superresolution techniques originally developed for passive sensor arrays. The goals are to produce more informative, higher-resolution imagery for improving target recognition with UHF and millimeter-wave SAR and to aid the image analyst in identifying targets in radar imagery. Algorithms presented here include two-dimensional minimum-variance techniques based on the maximum-likelihood method (Capon) algorithm and a two-dimensional version of the MUSIC algorithm. We use simulations to compare processing techniques, and we present results of wideband Rail SAR measurements of reflectors in foliage, demonstrating resolution improvement and clutter rejection. Results with airborne millimeter-wave SAR data demonstrate improved resolution and speckle reduction. We also discuss the vector aspect of HDVI, i.e., the incorporation of non-pointlike scattering models to enable feature detection. An example of a vector image is presented for data from an airborne UHF radar, using the broadside flash model to reveal greater information in the data.

HIGH-DEFINITION vector imaging (HDVI) is the application of superresolution digital signal processing techniques to radar imaging. HDVI is built on the expertise developed in research on the passive array processing problem, which provides resolution beyond the limits of conventional digital signal processing methods, and provides the ability to cancel unwanted interference. The benefits of HDVI to radar imaging are subpixel resolution, and sidelobe and clutter cancellation. Additional benefits come from exploiting the richness of information in the radar data, which provides analysts and target-recognition systems with the ability to categorize scattering mechanisms. Thus HDVI replaces conventional processing by providing us with a more accurate and higher-resolution image, and allowing us to extract additional information about the composition of radar backscatter.

The goal of HDVI is twofold, namely, to improve the automated recognition of targets, and to aid the image analyst in identifying targets in radar imagery. It is well known that resolution is the single most important factor in the ability to recognize targets via radar. Radar resolution is limited by bandwidth and

integration time, two expensive commodities in a system architecture. Image resolution, however, can be improved in a less expensive manner via digital signal processing, due to continuing advances in computers. This is the approach of HDVI.

The most promising aspect of HDVI is the ability to distinguish scatterers according to their scattering mechanisms. While a conventional image provides a single value for each pixel—i.e., the radar cross section (RCS)—HDVI can provide multiple values. This is the “vector” aspect of HDVI. Elements of the HDVI vector can include RCS, trihedral likeness, dihedral likeness, plate likeness, height, and polarization. HDVI accomplishes this vector decomposition by incorporating models of the scattering mechanisms in the image-formation process. In contrast, conventional radar imaging employs only a point-scatterer model. Furthermore, the adaptive algorithms of HDVI attenuate scatterers that deviate from the model, thus providing discrimination of scattering types.

HDVI is an entirely new approach to image formation. It employs both amplitude and phase information; it is not a post-processing technique applied

to intensity-only images. HDVI recasts image formation as a spectrum-estimation problem, treating individual pixels as beamformer outputs. It employs algorithms such as Capon's maximum-likelihood method (MLM), which is an adaptive-beamforming algorithm, and multiple signal classification (MUSIC), a powerful direction-finding algorithm. HDVI is a fully two-dimensional approach, not a successive application of one-dimensional algorithms that separately resolve an image in range and cross-range.

Previous attempts at applying superresolution processing techniques to radar images had varying degrees of success. Bandwidth extrapolation, a linear-prediction technique, provided superresolution in images with high signal-to-noise ratios and high target-to-clutter ratios [1]. Bandwidth extrapolation is a one-dimensional approach providing superresolution in range, and can be repeated in the cross-range, or Doppler, dimension. It is less than optimal because it is a one-dimensional approach, and because the linear-prediction model does not fit the data very well. A total least-squares approach (maximum likelihood in the case of white Gaussian noise) by S.R. DeGraaf exhibits good results, but is too computationally intensive [2], and suffers from model mismatch.

Adaptive-beamforming techniques include spatially variant apodization [3], adaptive sidelobe reduction [4], and an adaptive-filtering approach by J. Li and P. Stoica [5]. Both spatially variant apodization and adaptive sidelobe reduction attempt to perform sidelobe nulling without first computing a covariance matrix, resulting in cancellation of weak scatterers. Li and Stoica's adaptive-filtering approach preserves weak scatterers but provides wider main lobes, i.e., lower resolution. Also, J.W. Odendaal et al. [6] have implemented fully two-dimensional MUSIC in a manner similar to HDVI, but with inherently lower resolution due to a less effective covariance formation method. In contrast to all the methods listed above, HDVI is fully two-dimensional, does not require a high signal-to-noise ratio or target-to-clutter ratio, preserves gain on weak scatterers, and provides a more effective trade-off between robustness and resolution.

This article provides a detailed description of HDVI as applied to synthetic-aperture radar (SAR) data. The following section, entitled "Creating a SAR

Image: A Beamforming Approach," describes SAR image formation and shows how classical superresolution processing techniques such as Capon's adaptive beamformer have been modified to function with SAR data. This section describes the formation of the covariance matrix and the robustness constraints for Capon's technique, as well as assumptions and models such as the point-scatterer model and the extensions to this model. This section also shows how HDVI is efficiently applied to an image. Later sections show results of HDVI processing for airborne UHF SAR radar and the Ka-band Advanced Detection Technology Sensor (ADTS) radar. Articles in this issue by K.W. Forsythe [7] and L.M. Novak [8] provide descriptions of the superresolution algorithms and the improvements in automatic target recognition due to HDVI.

Creating a SAR Image: A Beamforming Approach

This section provides a detailed summary of HDVI. We begin with an overview of adaptive-beamforming concepts and then we describe the application of these concepts to SAR data. We also present the details of two-dimensional MLM and two-dimensional MUSIC, and we compare two-dimensional MLM with the technique of bandwidth extrapolation. The point-scatterer model and flash model are also discussed, along with how they are incorporated into the algorithms.

Beamforming Concepts

We begin a presentation of HDVI by comparing it with conventional SAR imaging techniques. Because the output of SAR processing is usually an image, it is best to consider the version of HDVI that replaces the conventional SAR image. Since this version is an intensity-only image, it is sometimes referred to as a "scalar" image rather than a vector image.

Beamforming is the process of integrating separate data samples to estimate the amplitude of a desired signal. The term arises from the application to adaptive arrays where we can view the beam, or antenna pattern, that arises from the combination of antenna elements. A simple beamformer is the steered array, generally accomplished via the Fourier transform in the case of a uniform linear antenna array. In this

case, the varying time delays are simply removed to provide maximum integration gain for the direction of interest.

The limitation of the conventional beamformer is its susceptibility to interference. Because an antenna pattern exhibits gain in directions other than the direction of interest, undesired signals are not always rejected. Interference from a signal widely separated from the direction of interest is called sidelobe interference. Conventional beamforming addresses sidelobe interference by *a priori* shaping of the beam through the use of a windowed fast Fourier transform (FFT). Interference from a direction near the direction of interest is called main-lobe interference. The conventional beamformer is ineffective against main-lobe interference because the main-lobe width, which is also called the beamwidth, or resolution of the array, is a function of the antenna aperture and cannot be reduced. The sidelobes and the main lobe can both be adaptively controlled by processing techniques that are based on the environment, e.g., Capon's MLM technique, which is a processing method of HDVI. Examples with SAR data are given below.

Conventional SAR Image Formation

The conventional SAR data-collection and image-formation process is a simple beamformer with sidelobe control, as illustrated in Figure 1. Radar reflection coefficients over a band of frequencies are collected for various viewing angles of the targets. The antenna positions that provide these viewing angles constitute the synthetic aperture, which is a viewing aperture synthesized over time. Information about the target location is encoded in the phase of the data samples $\{z_{mn}\}$. The range of the target results in a certain phase delay versus frequency, and the azimuth of the target results in a certain phase delay versus antenna position. The phase and amplitude characteristics of the target are encoded in the reflection coefficient ρ , illustrated in Figure 1(b) for a tank, a corner reflector, and environmental clutter, as a function of frequency f and look angle α . Image formation is the process of undoing these delays, integrating the data, and forming the RCS at the desired location. Returns from the desired pixel add coherently, while returns from other pixels add incoherently.

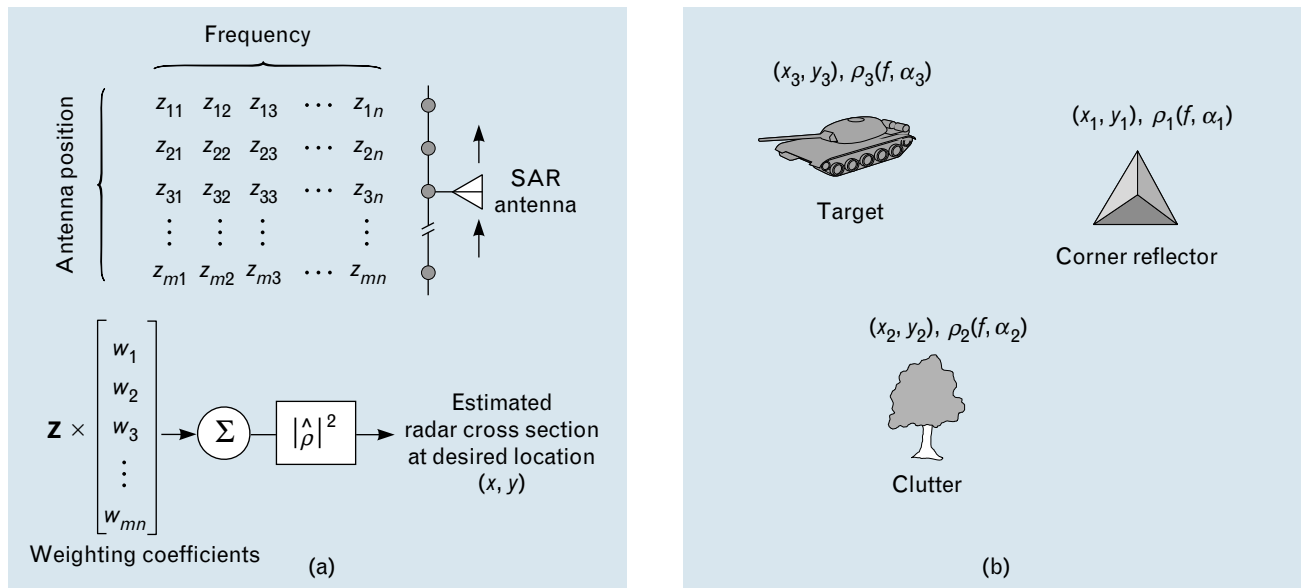


FIGURE 1. Synthetic-aperture radar (SAR) data collection and image formation. (a) The radar collects backscatter coefficients $\mathbf{Z} = \{z_{mn}\}$ over a band of frequencies at each of several antenna positions. (b) The backscatter contains a superposition of reflection coefficients ρ (which depend on the radar frequency f and look angle α), but shifted in phase because of the target location (x, y) . Image formation, which is a beamforming process, is accomplished through a weighted integration of the data $\{z_{mn}\}$ with unique weighting coefficients for each pixel location in the image (for example, a windowed fast Fourier transform, or FFT, provides sidelobe control). The squared magnitude of $\hat{\rho}$ provides an estimate of the radar cross section, or RCS, of the target.

The various methods of SAR image formation employed in operational systems are computationally affordable approximations to the simple beamformer illustrated in Figure 1. The book by W.J. Carrara et al. provides abundant details [9]. The operational image-formation process usually employs some data resampling, such as polar-to-rectangular conversion, to make the phase linear and enable the use of a two-dimensional FFT. In other words, raw SAR data have point responses with second-order and higher-order phase components that must be removed if an FFT is to be used to compute the RCS. Operational SAR image formation also employs an automatic focusing routine, called autofocus, to remove residual phase errors caused by uncompensated platform motion.

The Point-Scatterer Model

The underlying assumption in SAR image formation is that the data consist of a superposition of discrete point scatterers in stationary noise. The ideal point scatterer exhibits no variation in reflection coefficient with look angle or frequency. The ideal SAR response that results from an ideal point scatterer is a constant amplitude with a phase delay determined solely by geometry (the delay due to distance) and frequency. This phase delay is $4\pi rf$ radians for an object at range r and frequency f (the factor of 4π occurs because of the radar signal's round trip).

Application of Adaptive Beamforming to SAR Data

Adaptive beamforming applied to SAR data entails the derivation of a unique set of weighting coefficients, or weights, to estimate the RCS at each pixel in the output image. These weights replace the FFT used in conventional image formation. (For an illustration of the differences between conventional image formation and the adaptive beamforming accomplished with HDVI, see the sidebar entitled "Comparison of Conventional Image Formation with High-Definition Vector Imaging.") A critical step in the derivation of weights in adaptive beamforming is the estimation of an autocorrelation matrix, also called the covariance matrix. Given a covariance matrix \mathbf{R} , the weights \mathbf{w} are determined from the Capon algorithm as $\mathbf{w} \propto \mathbf{R}^{-1}\mathbf{v}$ [7, 10], where \mathbf{v} is the desired scatterer response, or steering vector.

The sections below describe the generation of the covariance matrix and the associated steering vectors. Details about Capon's MLM algorithm are described later in the article in the section entitled "Capon's Technique."

Generation of the Sample Covariance Matrix

The purpose of adaptive beamforming is to provide an accurate estimate of the cross section and location of a scatterer in the presence of neighboring scatterers that cause interference. The sample covariance matrix, which is necessary for this estimation task, is a second-order statistic that provides information on how to adapt the nulling pattern to minimize the influence of the interfering scatterers. The covariance matrix arises naturally in minimum-variance adaptive-filtering problems, such as the Capon algorithm.

In a communications application, the sample covariance matrix is computed as a time average of pairwise correlations of antenna outputs. To determine the covariance matrix from SAR data, where time averaging is not an option, another computational method must be used in place of time averaging. SAR data collection doesn't provide multiple renditions of the scene, or *looks*, as is customary in the communications application, in which samples are continually arriving over time. SAR data collection provides only a single two-dimensional sample of the two-dimensional scene. To get looks for SAR that are analogous to looks in the communications application, we would have to retrace the SAR flight path exactly, and we would have to change reflection coefficients to emulate modulation. This is clearly impossible.

One way to generate looks from SAR data is to form images from subsets of the given data. For example, two images could be formed by sectioning the synthetic aperture into halves, and computing an image with each half. An analogous sectioning could be done in frequency. This sectioning would provide four looks, but each of these looks would possess only half the original resolution.

A better way to form looks is to section the SAR data into overlapping subsets that do not sacrifice so much resolution, such as subsets using 80% of the synthetic aperture. For example, if the full set of SAR data in Figure 1 is represented by 100×100 samples,

every 80×80 contiguous subset will provide a unique image of the scene. Though not independent, such looks are sufficient to define a covariance matrix and support adaptive beamforming.

A drawback of using looks that contain more than 50% of the synthetic aperture and 50% of the frequency band is that there are fewer looks than dimensions in the covariance matrix. Having an insufficient number of looks results in a rank-deficient matrix that is not invertible. There are two advantages, however, of employing such higher-resolution looks. They relieve the burden on the algorithm to make up for the loss in resolution that occurs with lower-resolution looks, and they maintain a higher target-to-clutter ratio, which is critical in the imaging of ground targets.

In practice, HDVI begins with polar-formatted data in the frequency, or pre-image, domain. That is, the data have been resampled and refocused so that the point-response function has linear phase, which allows image formation via FFTs. This linear phase response of a point scatterer is then

$$\phi_{m,n} = \omega_x x m + \omega_y y n, \quad (1)$$

where (m, n) is the SAR data index, (x, y) is the point-scatterer location, and ω_x and ω_y are constants. Note that if the SAR data are sectioned into equally sized subsets, each subset will exhibit this same phase function, but with a unique overall phase offset. That is, the phase response of a point scatterer is preserved in the formation of looks. Hence an adaptive beamformer will exhibit identical gain patterns to each of the looks. The problem is now analogous to that of applying adaptive beamforming to a two-dimensional phased-array antenna.

Let us examine the process of generating the HDVI covariance matrix in more detail. HDVI begins with the complex-valued SAR image, transforms it to the frequency domain via a two-dimensional FFT (the reverse of the conventional image formation), removes any sidelobe control (for example, the Hamming weighting applied in conventional image formation), and decimates the data into 12×12 data sets. This decimation is accomplished by a conventional filter bank that effectively sections the SAR image into overlapping tiles. Looks are formed as $10 \times$

10 contiguous subsets of the 12×12 data, thus providing a total of nine distinct looks. Another nine looks are obtained via forward-backward averaging, a process that exploits the symmetry of the point-scatterer phase response. In Equation 1, note that phase conjugation is equivalent to a negation of the data index (m, n) . Hence a backward look (i.e., reversing the data indices and taking the conjugate) preserves the point-scatterer phase response (within a phase constant). The advantage of using the backward looks is improved noise averaging and scatterer decorrelation, and thus a more accurate covariance matrix. (Note: backward looks are applicable for point-scatterer analysis but not for azimuthal variations such as broadside flash.) The resultant covariance matrix has dimension 100 (from the 10×10 look), but has rank of only 18 (from the eighteen subset looks).

In more detailed notation, let \mathbf{z}_i denote a look vector, i.e., a 100×1 vector containing the samples of a 10×10 subset. Then the 100×100 covariance matrix is determined to be

$$\hat{\mathbf{R}} = \frac{1}{18} \sum_{i=1}^{18} \mathbf{z}_i \mathbf{z}_i^H. \quad (2)$$

This form of $\hat{\mathbf{R}}$ reveals that its elements contain the pair-wise correlations of the elements of \mathbf{z}_i , which is exactly the information required to implement adaptive beamforming. In practice, $\hat{\mathbf{R}}$ is not explicitly formed; rather, its eigendecomposition is computed directly from the looks. The information is the same but the number of computations is significantly reduced by using the decomposition.

Generation of the Steering Vector

The next step in the adaptive-beamforming algorithm is computing the response function of the desired scatterer; this response function is called the steering vector (the name derives from its use in electronically steered directional arrays). It contains both the location of the scatterer and the model of the desired scattering. Furthermore, the steering vector enables HDVI to superresolve a scatterer in location and discriminate according to scattering type (an essential and important result of vector imaging).

The simplest steering vector is the point-scatterer

COMPARISON OF CONVENTIONAL IMAGE FORMATION WITH HIGH-DEFINITION VECTOR IMAGING

THE FIGURES IN this sidebar illustrate the image improvements produced by high-definition vector imaging (HDVI) in main-lobe reduction and sidelobe cancellation, along with how HDVI

adapts the beamforming coefficients to produce these improvements. The simulated targets are two point scatterers separated by 1.5 pixels, each with high signal-to-noise ratio. Figure A shows im-

ages of the two point scatterers, as produced by the two-dimensional fast Fourier transform (FFT), the Taylor-weighted FFT, the quadratically constrained Capon's technique used in HDVI. The

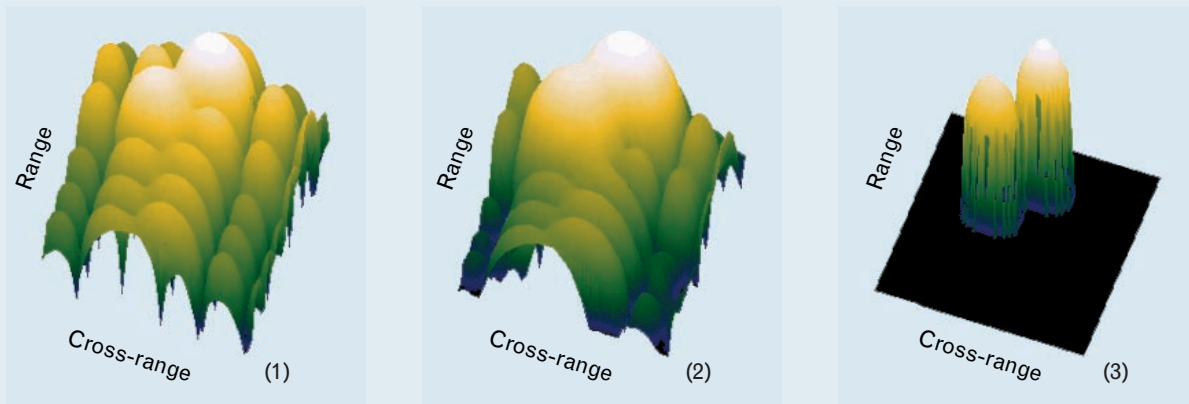


FIGURE A. Comparison of three image-reconstruction techniques. Two point scatterers separated by 1.5 pixels at high signal-to-noise ratio are imaged via (1) two-dimensional fast Fourier transform (FFT), (2) Taylor-weighted FFT (conventional image), and (3) HDVI.




FIGURE B. Spatial leakage patterns for the integration coefficients used to image the center of the patch (at the + sign): (1) two-dimensional FFT, (2) Taylor-weighted FFT, and (3) HDVI. The circles show the locations of the scatterers. HDVI shifts the null pattern (dark areas) onto the scatterers, thus cancelling their contribution (i.e., cancelling the sidelobes).

two-dimensional FFT image shows separation of the scatterers but exhibits high sidelobes, and the Taylor-weighted FFT sacrifices resolution of the scatterers to achieve a reduction in sidelobes. In contrast, HDVI provides dramatic sidelobe cancellation along with significantly improved scatterer resolution (i.e., narrower main lobes).

The improvements in HDVI are accomplished by adaptive beamforming. As Figure 1 in the main text shows, an image is formed from a weighted integration of synthetic-aperture radar data. In FFT processing, these coefficients are chosen to provide

maximum gain at the location of interest, namely, the location of the pixel in the output image. Because of the physical limitations of finite aperture and bandwidth, the beamforming coefficients allow some of the energy from neighboring pixels to leak into the integrated output, resulting in the phenomenon of sidelobes. The goal of adaptive beamforming is to minimize this leakage.

Figure B shows the spatial leakage patterns for the three processing techniques shown in Figure A. The location of interest is the center of the patch, indicated by the plus sign. In Figure B(1), the leakage pattern in the traditional two-

dimensional FFT technique allows energy from both scatterers, indicated by the circles, to bias the estimate of the radar cross section. In Figure B(2), the Taylor-weighted FFT decreases the leakage from the distant scatterer but increases the leakage from the nearby scatterer. In Figure B(3), however, HDVI simultaneously reduces the leakage from both scatterers while maintaining unit gain at the patch center. Hence the corresponding image in Figure A(3) evidences no sidelobes at the patch center. By repeating the adaptive beamforming at every pixel, HDVI produces an image without sidelobes.

response shown in Equation 1. This steering vector emulates a point scatterer in that the amplitude is constant, both in azimuth and in frequency, and the phase is linear. The location of the scatterer (x, y) is encoded in the steering vector as the phase change from sample to sample. A steering vector is computed for each pixel in the desired output image by recomputing Equation 1 for each pixel location (x, y) . Steering vectors can be generated explicitly as in Equation 1, or they can be generated implicitly via a two-dimensional FFT in special cases.

For non-pointlike scattering, an amplitude and phase function are superimposed on the point-scatterer model. A simple model of non-pointlike scattering is the azimuthal flash, which we discuss in more detail in a later section. Steering vectors for this model are computed as for point scatterers, but with an amplitude profile applied in the dimension corresponding to azimuth.

Description of HDVI Algorithms

This section presents the particulars of the application of the techniques of adaptive processing to SAR data. The enabling steps—generation of the covari-

ance matrix and the corresponding steering vectors—are computed as described in the preceding section, and are used in conjunction with the techniques described below.

Capon's Technique

Capon's technique is a well-known approach to the problem of adaptive beamforming for the purpose of estimating a power spectrum [10]. In the application of this technique to radar, the power being estimated is the RCS of the scatterer, and the power spectrum is the image. In other words, Capon's technique provides RCS estimates for each pixel location (x, y) ; the estimates are then displayed as the brightness of each pixel, hence forming an image. Capon's technique is also known as the maximum-likelihood method (MLM), which was Capon's original name for the technique; minimum-variance distortionless response (MVDR); and reduced-variance distortionless response (RVDR) [11]. Capon's technique produces a SAR image that appears similar to the conventional image, but with improved resolution. The classical form of this technique is not applicable here, however, because the covariance matrix is rank deficient.

The modified form of Capon's technique is presented below.

The fundamental idea in Capon's technique is to determine beamforming coefficients that minimize the detected energy of interferers while simultaneously keeping unit gain on the location of interest (x, y) . That is,

$$\text{RCS}(x, y) = \min_{\mathbf{w}} \mathbf{w}^H \hat{\mathbf{R}} \mathbf{w}$$

such that

$$\mathbf{w}^H \mathbf{v}(x, y) = 1,$$

where $\mathbf{v}(x, y)$ is the steering vector for a point scatterer at location (x, y) , and $\|\mathbf{v}\| = 1$. If the covariance matrix $\hat{\mathbf{R}}$ were full rank (in other words, invertible), the solution would be the classical form of Capon's technique given by


$$\text{RCS}(x, y) = (\mathbf{v}^H \hat{\mathbf{R}}^{-1} \mathbf{v})^{-1}.$$

In the case of SAR data, $\hat{\mathbf{R}}$ is not full rank, and Capon's technique is not applicable, as stated above. The solution to this difficulty is to constrain the allowable set of weighting coefficients \mathbf{w} . We consider two constraints on \mathbf{w} : a norm constraint and a subspace constraint. The norm constraint, or quadratic constraint, limits the ability to resolve closely spaced scatters, effectively limiting the amount of adaptivity, that is, the amount of deviation of the adaptive beamformer from the conventional. The subspace constraint has the additional benefit of retaining the clutter background in the resultant image.

Quadratically Constrained Coefficients. The quadratic constraint [12] restricts the amount of adaptivity; this constraint is achieved by limiting the norm of the weights \mathbf{w} , where

$$\|\mathbf{w}\| = \sqrt{\mathbf{w}^H \mathbf{w}}.$$

Figure 2 visualizes this concept and illustrates the behavior of \mathbf{w} in Capon's technique. Observe the unit-norm steering vector, $\mathbf{v}(x, y) = \mathbf{v}_{\text{model}}$, and the locus of vectors \mathbf{w} such that $\mathbf{w}^H \mathbf{v} = 1$ (Capon's gain constraint), shown by the dashed line. The conventional beamformer is simply $\mathbf{w} = \mathbf{v}_{\text{model}}$. Adaptivity is illustrated as deviations of \mathbf{w} from $\mathbf{v}_{\text{model}}$.



$\ \mathbf{w}\ _{\text{dB}}$	θ	$\ \mathbf{w}\ _{\text{dB}}$	θ
0.5	20°	Hamming	31°
1.0	27°	Taylor	24°
3.0	45°		

FIGURE 2. Illustration of the quadratically constrained beamformer. The unconstrained Capon beamformer \mathbf{w}_{free} preserves gain on $\mathbf{v}_{\text{model}}$ while nulling \mathbf{v}_{true} . The quadratic constraint prevents the nulling of \mathbf{v}_{true} by restricting \mathbf{w}_{free} to be within angle θ of $\mathbf{v}_{\text{model}}$. The table lists the equivalence between the norm of \mathbf{w} and the angle θ , and the effective angle for Taylor and Hamming weights.

The goal of Capon's technique is to find the \mathbf{w} on the dashed line such that this \mathbf{w} minimizes output energy. Suppose the vector \mathbf{v}_{true} in Figure 2 is the response of a scatterer at a subpixel separation from $\mathbf{v}_{\text{model}}$. With no constraint, Capon's technique chooses \mathbf{w}_{free} orthogonal to \mathbf{v}_{true} , thus nulling this scatterer. In practice, it may be undesirable to allow a null to fall so close to \mathbf{v} , especially when we need to allow for modeling errors. It is then desirable to restrict the angular separation between \mathbf{w} and $\mathbf{v}_{\text{model}}$, which is effectively accomplished by restricting the norm of $\mathbf{w}_{\text{constrained}}$ (shown in green).

One method of constraining the norm of \mathbf{w} is to add the norm, multiplied by a weighting factor α , into the minimized quantity; i.e.,

$$\min_{\mathbf{w}} \mathbf{w}^H \hat{\mathbf{R}} \mathbf{w} + \alpha \mathbf{w}^H \mathbf{w},$$

for which the minimizing \mathbf{w} is

$$\mathbf{w} \propto (\hat{\mathbf{R}} + \alpha \mathbf{I})^{-1} \mathbf{v}.$$

In this expression, α is added to the diagonal elements of $\hat{\mathbf{R}}$. This method is referred to as *diagonal loading* of the covariance matrix, and is the method of H. Cox and R. Zeskind [11]. Diagonal loading is a soft con-

straint because the norm of \mathbf{w} can still become large. It is not often employed as a technique for HDVI, however, because there is no reliable means of choosing α .

The constraint generally used in HDVI is a hard constraint on the norm of \mathbf{w} ; i.e.,

$$\text{RCS}(x, y) = \min_{\mathbf{w}} \mathbf{w}^H \hat{\mathbf{R}} \mathbf{w},$$

such that

$$\mathbf{w}^H \mathbf{w} \leq \beta,$$

where $\sqrt{\beta}$ is the length of the vector $\mathbf{w}_{\text{constrained}}$. This inequality is called a quadratic constraint because the quantity $\mathbf{w}^H \mathbf{w}$ is of quadratic (second order) form. The form of $\mathbf{w}_{\text{constrained}}$ is

$$\mathbf{w}_{\text{constrained}} \propto [\hat{\mathbf{R}} + \alpha(\beta, \hat{\mathbf{R}}, \mathbf{v})\mathbf{I}]^{-1} \mathbf{v}$$

for some value of α . Though this solution has the form of diagonal loading, now the loading is optimized by the algorithm on a pixel-by-pixel basis.

Subspace Constraints. The subspace constraint was devised to preserve the background, or clutter, regions of the image. The quadratic constraint alone is not sufficient to prevent loss of background information, because it allows the null solution $\text{RCS}(x, y) = 0$. In Figure 2 this null solution can be seen when $\mathbf{w}_{\text{constrained}}$ and \mathbf{w}_{free} coincide, which occurs, for example, if \mathbf{v}_{true} is further separated from $\mathbf{v}_{\text{model}}$. Then $\mathbf{w}_{\text{constrained}}$ will be orthogonal to \mathbf{v}_{true} so that

$$\mathbf{w}_{\text{constrained}}^H \cdot \mathbf{v}_{\text{true}} = 0.$$

The resultant pixel will be black, signifying $\text{RCS} = 0$. In practice, this null solution generally occurs in clutter regions of the image and is not representative of the actual RCS. The null solution is an artifact that can be disturbing to the image analyst, however, and needs to be mitigated.

Figure 3 visualizes the problem of the null solution in a more general fashion, and visualizes the subspace-constraint solution. The null solution occurs when the quadratically constrained weight vector \mathbf{w}_q (shown in red) becomes orthogonal to the columns of the covariance matrix, denoted by $\text{span}(\hat{\mathbf{R}})$. The notation $\text{span}(\hat{\mathbf{R}})$ denotes the subspace spanned by the columns of $\hat{\mathbf{R}}$ (see Equation 2), which is also called




FIGURE 3. Illustration of the subspace-constrained beamformer. The subspace constraint on the weighting coefficients prevents the null solution \mathbf{w}_q (in red). The plane denotes the space spanned by the covariance matrix, or $\text{span}(\hat{\mathbf{R}})$. The constrained weights \mathbf{w}_s (in green) result from perturbing \mathbf{v} by a vector ϵ derived from $\text{span}(\hat{\mathbf{R}})$.

the signal subspace. The goal of the subspace constraint is to prevent \mathbf{w} from becoming orthogonal to $\text{span}(\hat{\mathbf{R}})$ by constraining the allowable directions in which \mathbf{w} can lie with respect to \mathbf{v} .

We can understand this concept by viewing the adaptive beamformer as a perturbation of the conventional beamformer. The conventional beamformer assigns $\mathbf{w} = \mathbf{v}$; that is, \mathbf{w} is not adapted to the data. The adaptive beamformer allows modification of \mathbf{w} such that

$$\mathbf{w}^H \mathbf{v} = 1,$$

which is shown as the dashed line in Figure 2. The weight vector \mathbf{w} can be viewed as a perturbed version of \mathbf{v} , such that the perturbation is orthogonal to \mathbf{v} . Mathematically, $\mathbf{w} = \mathbf{v} + \epsilon$ such that ϵ is orthogonal to \mathbf{v} . The length of the perturbation vector ϵ is an indicator of the amount of adaptivity; that is, a smaller ϵ indicates little adaptivity and a larger ϵ indicates more. In addition to the length of ϵ , there is also a direction to ϵ . The direction of ϵ is the subject of the subspace constraint.

Returning to Figure 3, we consider constraining the direction of ϵ such that it must be derived from $\text{span}(\hat{\mathbf{R}})$ while also being orthogonal to \mathbf{v} . This set of allowable vectors in $\text{span}(\hat{\mathbf{R}})$ is illustrated by the top edge of the dashed parallelogram. (The bottom edge of the parallelogram lies in $\text{span}(\hat{\mathbf{R}})$ and is orthogonal to the projection of \mathbf{v} into $\text{span}(\hat{\mathbf{R}})$, denoted $\text{proj}_{\hat{\mathbf{R}}} \mathbf{v}$,

to satisfy the orthogonality constraint.) In this case, ε is prohibited from pointing away from $\text{span}(\hat{\mathbf{R}})$, and therefore \mathbf{w} is prevented from approaching \mathbf{w}_q . The subspace-constrained solution will then be $\mathbf{w}_s = \mathbf{v} + \varepsilon$ (shown in green) and will not be orthogonal to $\text{span}(\hat{\mathbf{R}})$. Thus the subspace constraint on ε effectively prevents the null solution.

The subspace-constrained and quadratically constrained version of Capon's technique is given by

$$\text{RCS}(x, y) = \min_{\mathbf{w}} \mathbf{w}^H \hat{\mathbf{R}} \mathbf{w},$$

such that

$$\begin{aligned} \mathbf{w} &= \mathbf{v} + \varepsilon \\ \varepsilon &\perp \mathbf{v} \\ \varepsilon &\in \text{span}\{\hat{\mathbf{R}}\} \\ \|\mathbf{w}\|^2 &\leq \beta. \end{aligned}$$

This subspace-constrained technique (sometimes referred to as the confined-weight technique) is employed in HDVI at higher frequencies, i.e., X-band and above. It has the property of preserving the background information and reducing the speckle, which is a by-product of the averaging inherent in the formation of the covariance matrix.

Coherent versus Incoherent Summing of Looks. The Capon technique produces an RCS estimate that is computed as the incoherent sum of the individual looks averaged in the covariance matrix. This incoherent sum is

$$\mathbf{w}^H \hat{\mathbf{R}} \mathbf{w} = \frac{1}{L} \sum_{i=1}^L \left| \mathbf{w}^H \mathbf{z}_i \right|^2,$$

where the \mathbf{z}_i denote the looks and L is the number of looks (e.g., $L = 18$ in Equation 2). Incoherent averaging is necessary if the relative phasing of the desired signal is unknown from look to look. With SAR data, however, the relative phasing is known because the looks were created from one coherent set of collected data. This knowledge of the relative phasing can be exploited to provide an image of the target with enhanced sidelobe and clutter rejection.

Coherent summing exploits the known relative

phase of the looks. It was originally proposed by L.L. Horowitz as a means of estimating scatterer phase in addition to RCS. Let u_i be the $L \times 1$ vector of the phase progression of the looks, where u is a function of $\mathbf{v}(x, y)$. The coherent sum of looks can then be expressed as

$$z(x, y) = \frac{1}{L} \sum_{i=1}^L \mathbf{w}^H \mathbf{z}_i u_i^*,$$

where $z(x, y)$ is the complex output, and \mathbf{w} is derived as in the incoherent case. Note that the above summation is computed only over the forward looks; i.e., backward looks must be excluded.

Computational Load. Current implementations of Capon's technique for the point-scatterer model require approximately two thousand operations per output sample, where an operation is one real-valued multiplication or addition. This operation requirement is not dependent on the size of the image. For comparison, a conventional image requires one thousand to two thousand operations per pixel, including polar formatting and autofocus.

The MUSIC Algorithm

The MUSIC algorithm is a powerful direction-finding algorithm that exploits properties of the signal subspace of the covariance matrix [13]. It does not produce an estimate of the RCS, but rather an estimate of how closely the data match the scattering model. If the point-scatterer model is employed, a MUSIC image shows how pointlike the scattering is.

Let $\{\mathbf{e}_i\}$ be a set of orthonormal vectors that span the column space of $\hat{\mathbf{R}}$. These vectors can be derived from an eigendecomposition of $\hat{\mathbf{R}}$ as the eigenvectors that correspond to the nonzero eigenvalues. The MUSIC output is then

$$\text{MUSIC}(x, y) = -10 \log_{10} \left[1 - \sum_{i=1}^L \left| \mathbf{e}_i^H \mathbf{v}(x, y) \right|^2 \right].$$

We can recognize the summation as the norm of the projection of $\mathbf{v}(x, y)$ into the span of the eigenvectors. If the norm of the projection of $\mathbf{v}(x, y)$ is large (near one), then the MUSIC output is large. Although MUSIC usually employs an estimate of the

number of signals, HDVI simply uses the rank of the covariance matrix as an estimate of the number of signals because there are so few looks.

Other Related Techniques

While investigations in HDVI have employed MLM (Capon's technique) and MUSIC, there are a variety of modern spectrum-estimation techniques that may be applied to the problem, such as maximum likelihood, linear prediction (so-called bandwidth extrapolation), adaptive event processing [7], ESPRIT [14], and a sinusoid estimation technique adapted to imaging by Li [5]. Two of these techniques are addressed further.

The maximum-likelihood algorithm attempts a joint estimation of scatterer locations, amplitudes, and phases by assuming that the data are a sum of ideal scatterers in noise [2]. This algorithm has not been employed in practice because of its computational burden and its sensitivity to model mismatch.

Bandwidth extrapolation is an application of linear prediction to SAR data [1]. It derives a prediction filter to estimate out-of-band data samples, which effectively increases resolution, and it employs the point-

scatterer model to extend the dominant sinusoidal components in the data. Also, it is an inherently one-dimensional technique that is applied successively in range and cross-range to produce the effect of two-dimensional superresolution.

Figure 4 shows a comparison of bandwidth extrapolation with the MLM technique employing the soft constraint (defined earlier in the section on quadratically constrained coefficients). The simulated target consists of four scatterers in one resolution cell (pixel), which is a difficult superresolution example. The result is that MLM does a better job than bandwidth extrapolation in resolving the scatterers because MLM is a fully two-dimensional technique that is generally better suited to SAR data.

Vector Imaging

Vector imaging refers to the incorporation of multiple scatterer models into the image-formation process. The goal of vector imaging is to categorize and classify radar backscatter according to the phenomenon that caused the reflection. Figure 5 shows a notional vector image, where each element of the vector is a separate high-definition image formed by using a dif-




FIGURE 4. A comparison between bandwidth extrapolation (left) and Capon's maximum-likelihood method (MLM) using the soft quadratic constraint (right). The simulated target is four point scatterers within one resolution cell, and the area imaged is 1.5×1 resolution cells. Capon's technique clearly provides resolution of the four scatterers.

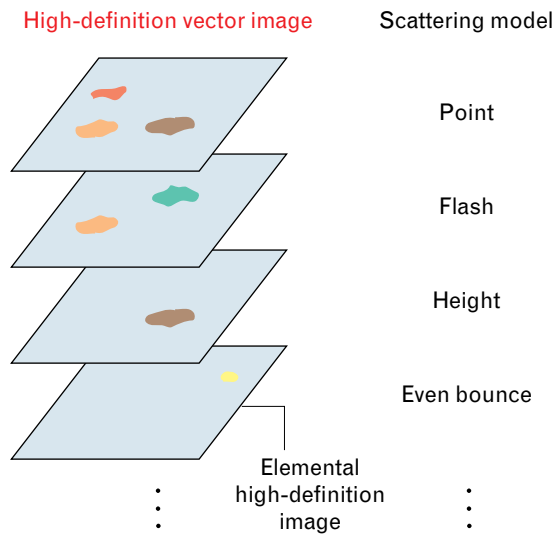


FIGURE 5. Illustration of a vector image, where a separate elemental high-definition image is produced for each scattering model. Each pixel in the vector image contains a vector of parameters that analyzes the received backscatter according to its similarity to the scattering model. This vector structure provides a more comprehensive understanding of the target.

ferent scattering model. Vector imaging requires the user to define the scattering models that are most representative and useful, which unfortunately is not a simple task. Some notional vector images include the point-scatterer and flash phenomena, discussed in this article, as well as other scattering phenomena such as scatterer height (an extension of interferometric SAR) or polarization.

Figure 6 illustrates the computation of a vector image. The vector capability is achieved in the selection of the steering vector $\mathbf{v}(x, y)$ because the steering vector incorporates the model of the desired scattering. Also, some changes may be required in the generation of the covariance matrix; these changes are described below, along with a description of the flash model and the polarimetric model.

The Azimuthal-Flash Model

Azimuthal flash is a scattering phenomenon caused by a flat reflecting surface several wavelengths long. Though this phenomenon can occur at any frequency

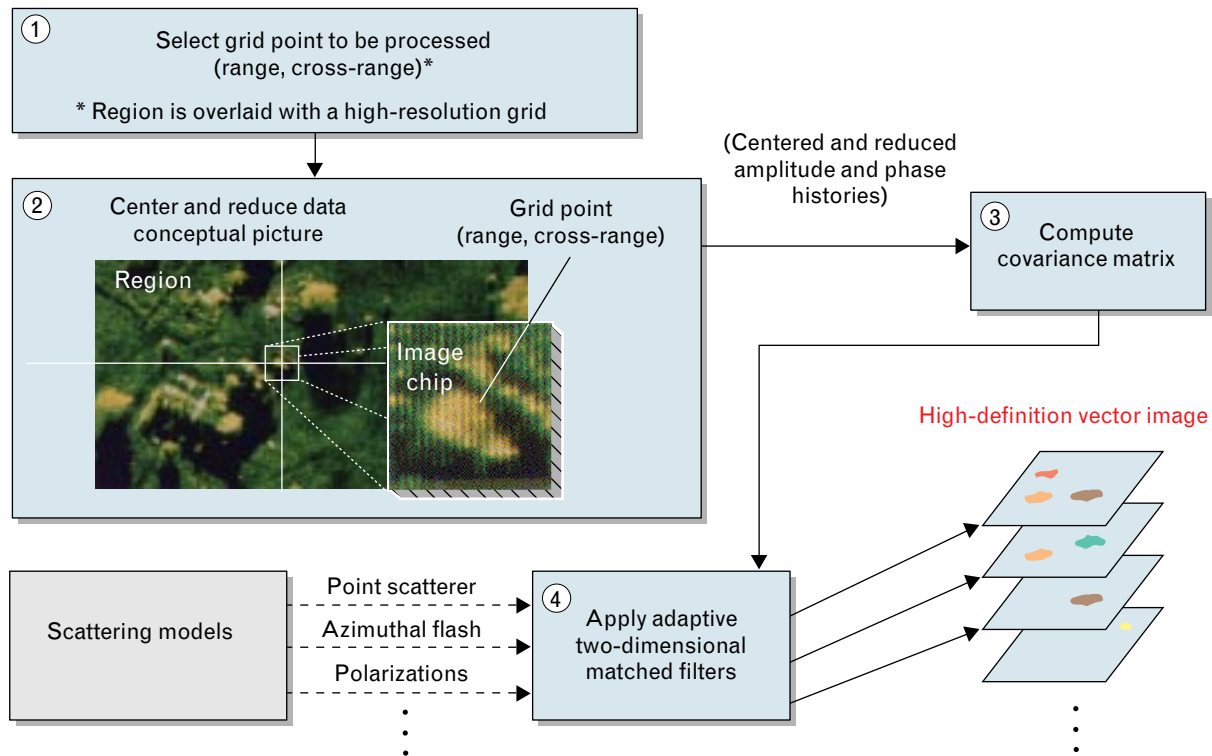


FIGURE 6. The processing flow in HDVI. The image is processed in small image chips, and a covariance matrix is formed for each chip. The adaptive-filtering technique (e.g., Capon's technique) employs the covariance matrix and the desired scattering model to produce the high-definition vector image.

with many types of target, a common example is observed at UHF when a vehicle is illuminated from broadside. This broadside flash occurs when the entire side of a vehicle acts as a reflector, causing the RCS to be dependent on the look angle, with the brightest return occurring when the radar is normal to the side of the vehicle.


The model for the azimuthal flash used in HDVI was generated heuristically as a coherent sum of point scatterers in adjacent resolution cells. This form of the model results in a sum-of-cosines amplitude profile across the rows (the azimuthal observations) of the data matrix (i.e., the flash was approximated by a Fourier series). Broad flashes require low-order cosines, while narrow flashes require higher-order cosines. The underlying phase of the azimuthal-flash model is identical to that of the linear-phase point scatterer defined in Equation 1.

The azimuthal-flash model does not exhibit the shift-invariance property of the point-scatterer model assumed in the generation of the covariance matrix. For azimuthal-flash processing, looks must be limited in the amount of azimuthal shifting allowed, and forward-backward averaging cannot be used.


The processing requirements for image formation are increased because it is now necessary to search for the location of the azimuthal flash within the synthetic aperture. The benefit of this search, however, is that we gain an estimate of the orientation of the reflecting object.

Polarimetric Models

Both the point-scatterer model and the azimuthal-flash model are easily extended to include polarizations [15]. Polarimetric data can be viewed as an extension of the two-dimensional SAR data discussed above, wherein each data sample is replaced with a three-element vector containing the HH, HV, and VV polarization information. The looks and covariance matrix are formed in the usual fashion, except that forward-only looks are employed. The resultant covariance matrix has a dimension three times that of the covariance matrix for a single polarization. We can choose the steering vectors to model a desired response, e.g., trihedral or dihedral. The importance of this vector processing is that nulls can be generated in



(a)



(b)

FIGURE 7. (a) The ultra-wideband Rail SAR—consisting of a movable antenna mounted on a rail—provides *in situ* measurements of small areas over a wide band of frequencies. (b) Multiple antennas allow SAR coverage over a range of frequency bands from 100 MHz to 18 GHz.

polarization space as well as in range and azimuth. A few preliminary results have been achieved and described elsewhere [15].

Rail SAR Results

The early results of HDVI were achieved with the Lincoln Laboratory ultra-wideband Rail SAR. Shown

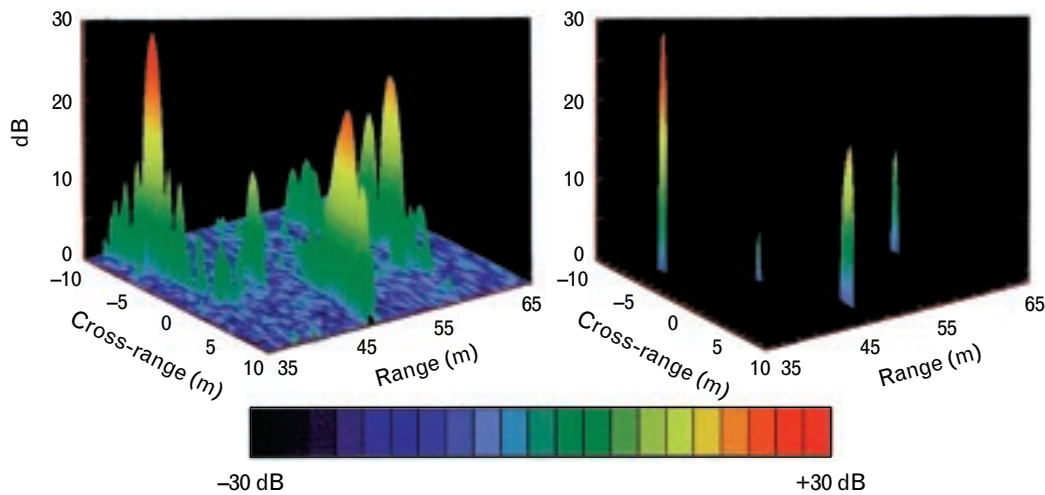


FIGURE 8. Comparison of conventional image processing (left) and HDVI (right) at L-band frequencies for four reflectors in foliage. The quadratically constrained MLM algorithm used to produce the HDVI image reveals the reflectors and suppresses the foliage.

in Figure 7, the Rail SAR is an instrumentation-quality radar designed for *in situ* measurements of single targets, thus avoiding the errors introduced by platform motion and waveform limitations. For the results that follow, the Rail SAR measured the complex reflection coefficients of a target from 500 MHz to 2.0 GHz in 801 frequency steps, and traversed a 4.5-m synthetic aperture in 31 steps.

Clutter Rejection Examples

This example shows the ability of the Capon algorithm to discriminate reflectors from foliage. Four reflectors—three eight-foot trihedral reflectors and one four-foot dihedral reflector—were placed at various locations in foliage, with the farthest reflector placed at a range of sixty meters, and thirty meters deep in

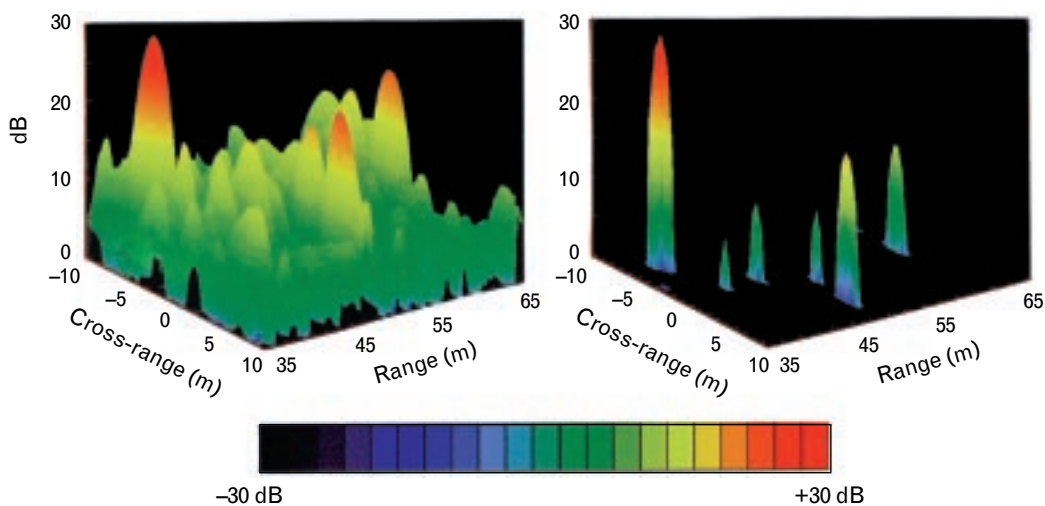


FIGURE 9. Comparison of conventional image processing (left) and HDVI (right) at UHF frequencies for four reflectors in foliage. The quadratically constrained MLM algorithm used to produce the HDVI image reveals the reflectors and suppresses the foliage, except for two tree trunks that appear as scatterers (which are seen in this HDVI image but not in the HDVI image in Figure 8).




FIGURE 10. Superresolution of target features for VV (vertical) polarization (left) and HH (horizontal) polarization (right) at half-beamwidth separation. The two-dimensional MUSIC algorithm resolves two scattering centers within one pixel. The target is a circular disk that exhibits a single scattering center when illuminated with VV polarization but two resolvable scattering centers when illuminated with HH polarization.

foliage. We compare conventional image processing (without azimuthal sidelobe control) and quadratically constrained MLM processing without the subspace constraint. A three-dimensional display reveals the shape of the sidelobes.

Figure 8 shows the result of this discrimination test at L-band, 1.0-to-1.44-GHz, VV polarization. The nominal range resolution is 0.34 m, and cross-range resolution is 1.1 m at a range of forty meters. The four reflectors are easily discerned in the HDVI image; the weakest discrimination is the four-foot dihedral. The reduced gain on the reflector at a range of sixty meters in the HDVI image is due to the degradation of its response because of the foliage. MLM is designed to preserve gain on ideal point scatterers, and reduces gain on non-pointlike scatterers such as trees and corrupted reflectors.

Figure 9 shows the result of this discrimination test at UHF, 500-to-720-MHz, VV polarization. The nominal range resolution is now 0.68 m, and cross-range resolution is 2.1 m at a range of forty meters. Reflector RCS is smaller at this lower frequency, and the lower resolution allows more clutter per resolution cell. Because of the lower target-to-clutter ratio, the MLM constraint had to be relaxed, resulting in the appearance of two clutter returns (tree trunks), which are not visible in the HDVI image in Figure 8.

The improvement over the conventional image, however, is readily apparent.

Superresolution Example

This example demonstrates the ability of the MUSIC algorithm to resolve the edges of a circular disk, which appears as two scattering centers separated by a distance of half a resolution cell. The SAR data were collected for a disk illuminated at near-normal incidence. Electromagnetic theory predicts that a circular disk, when illuminated with linear polarization, will exhibit flux concentrations at its physical extremes, namely, in the vertical direction for vertical polarization and in azimuth for horizontal polarization. Because the Rail SAR has a horizontal aperture, we can expect to resolve scatterers exhibiting horizontal separation but not scatterers exhibiting only vertical separation. Hence we can expect to resolve the circular disk into two scatterers when HH polarization is employed but not when VV polarization is employed. Figure 10 shows the results of the MUSIC algorithm for VV polarization and HH polarization, at 1.2 to 1.5 GHz. Note that the VV polarization reveals one peak while the HH polarization reveals two, as expected. The separation of the scattering centers in the figure is approximately equal to the diameter of the disk.



FIGURE 11. SRI International ultra-wideband UHF SAR. The antenna is mounted under the wing of the aircraft. The SAR parameters are listed in Table 1.

Airborne UHF SAR Results

Lincoln Laboratory has done extensive studies on the detection of targets in foliage. Foliage renders microwave radar useless, forcing the radar to employ lower frequencies where resolution is limited. One promising approach to penetrating foliage has been to employ a wide-bandwidth UHF radar and to integrate the data over a wide synthetic aperture. The resultant radar signatures, however, present a unique detection problem for which HDVI may provide a solution.

The results presented here employ data from a se-

Table 1. Ultra-Wideband SAR Parameters

Waveform	Impulse
Bands	200–400 MHz, 100–300 MHz
Resolution	1 m × 1 m
Polarization	HH
Power	50-kW peak, 0.1-W average

ries of SAR data collections that were performed in Cedar Swamp, Maine, using the SRI International ultra-wideband UHF SAR, shown in Figure 11. The radar waveform is a 200-MHz-bandwidth impulse with a carrier frequency of 300 MHz, and the polarization is HH (see Table 1). The synthetic aperture is a 30° intercept with respect to the region of interest, resulting in an image with 1-m × 1-m nominal resolution. This resolution is sufficient to spatially separate vehicles from tree-trunk returns (the dominant clutter source), but the resultant target-to-clutter levels are not high enough to provide reliable detection.

The goal of HDVI is to employ models that more closely resemble the backscatter from man-made ob-



FIGURE 12. Aerial photo (top) and SAR image (bottom) of the Cedar Swamp, Maine, deployment. The radar line of flight in the SAR image was across the top of the photo. The vehicles in the center were not visible from the point of view of the radar because of the thick foliage.




FIGURE 13. Overview of the Cedar Swamp deployment shown in Figure 12. The nine vehicles in this illustration are described in Table 2.

jects, rather than backscatter from trees, thus providing a way of discriminating the man-made objects. The models employed are the point-scatterer model and the azimuthal-flash model, both of which were described earlier. Point-scatterer responses are regularly observed with man-made objects. Azimuthal-flash responses are the result of the interaction of the broadside of the vehicle and the ground, forming a dihedral-type reflection. The extent of this dihedral is several wavelengths; e.g., a vehicle ten meters in length is ten wavelengths at 300 MHz, which is visible only over a narrow range of angles. Trees, on the other hand, are neither pointlike nor flashlike at these frequencies. Though closer to point scatterers, trees exhibit random amplitude and phase variations, which tends to separate them from ideal point signatures. HDVI is thus able to reject more of the tree responses through the use of these models.

Point-Scatterer Results

Figure 12 shows an aerial photograph of the SAR data-collection site in Cedar Swamp, Maine. The radar path was parallel to the top of the photo, at a range of about six hundred meters, and at a 45° elevation with respect to the vehicles along the road. Vehicles in the middle to right part of the photo are parked next to tall pine trees, and are thus concealed from the view of the radar. Figure 13 shows a diagram

Table 2. Cedar Swamp Vehicle Descriptions

Vehicle Label	Vehicle Type	Length (m)
V1	Logging truck, loaded	24
V2	Logging truck, empty	18
V3	Five-ton truck	8.4
V4	Lincoln Laboratory test target	15
V5	Logging truck, loaded	24
V6	Logging truck, empty	18
V7	Two-ton truck	6.7
V8	HEMTT (truck)	10
V9	Five-ton truck	8.4

of the locations of the nine vehicles at Cedar Swamp, and Table 2 lists their types and lengths.

Figure 14 presents a comparison of the conventional image and the HDVI image using the point-scatterer model. A dynamic range of 30 dB is displayed with a rainbow color scale, and peaks above a threshold are displayed three-dimensionally. The thresholds are chosen for each image such that eight of the nine vehicles exceed it; i.e., the probability of detection is 0.89 in each image. HDVI has reduced the number of false detections by a factor of four, making it much easier to identify the returns from the vehicles along the top edge of the roadway. The largest return in the center is from an eight-foot trihedral corner reflector, also concealed in trees.

We employed the MUSIC algorithm here to provide a measure of point likeness, that is, the similarity between the measured data and a true point-scatterer response. The output of MUSIC was compared to a threshold, and a matched-filter radar-cross-section estimate was inserted for the display. A false background was added to provide a visual reference for the roadway.

Vector-Imaging Results

We processed the same data by using two sizes of the azimuthal-flash model, a 10° flash and a 3° flash. Recall that the synthetic aperture is a 30° intercept with

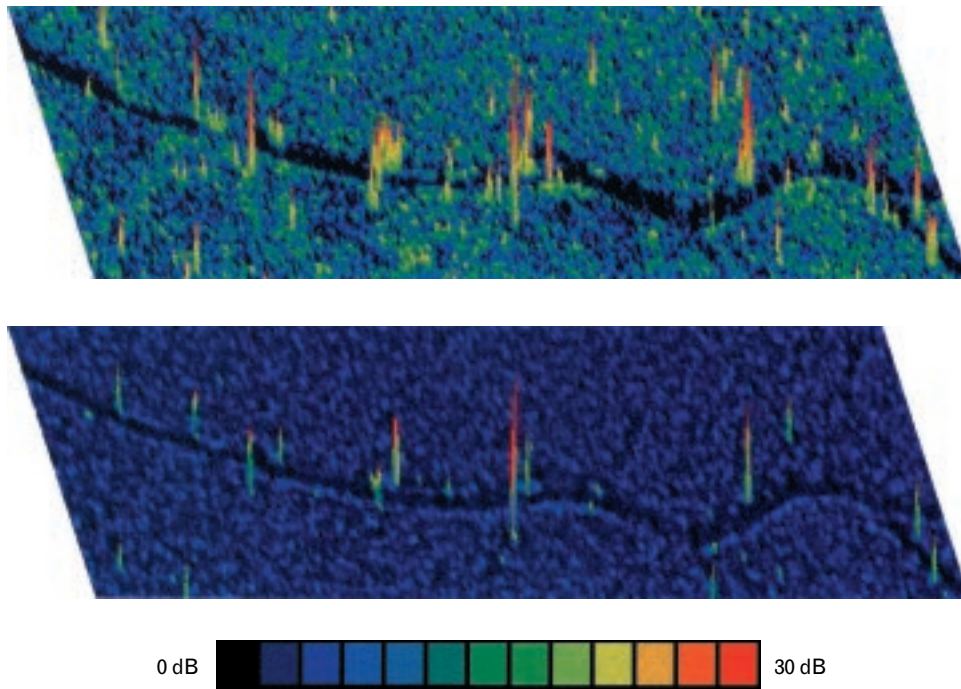


FIGURE 14. Comparison of conventional image processing (top) with HDVI (bottom). HDVI employs the point-scatterer model, which reduces clutter false alarms in the image by a factor of four. A smaller number of false alarms reduces the burden on subsequent automatic-target-recognition algorithms because fewer detections need to be reviewed. Both images are normalized to detect eight of the nine vehicles, with detections displayed in three-dimensional relief.

respect to the targets. Figure 15 shows the HDVI results for the azimuthal-flash models, using the MUSIC algorithm as a detector. At each pixel, the MUSIC output was computed for several shifted positions of the azimuthal flash within the synthetic aperture, to allow for the unknown orientation of the vehicle. The largest value was reported and compared to a detection threshold. Positive detections are displayed with a radar-cross-section estimate computed with a flash-matched filter.

The point-scatterer and two flash-scatterer images constitute a three-level vector image, which is significantly more useful than a conventional image in an analysis of the received backscatter. The vector image also provides improved clutter rejection beyond that of the point-scatterer image alone, especially for the 3° flash, where all bright returns are targets. Figure 16 shows an overlay of the MUSIC outputs with vehicle returns outlined and clutter excised. Note how flash size corresponds to the inverse of vehicle length. Several vehicles exhibit multiple signatures. For example,

the two logging trucks on the left exhibit a 3° flash from the entire body and a point return from the cab end. Not shown is the orientation estimate derived from the flash timing; i.e., the peak of the flash occurs when the radar is broadside to the vehicle.

The vector image clearly provides information that conventional radar image processing did not reveal. The next step is to discover a method to jointly exploit the multiple levels of the vector image for target detection and recognition.

Advanced Detection Technology Sensor Results

To achieve better than $1\text{-m} \times 1\text{-m}$ resolution with SAR data, we must operate at microwave frequencies. The Advanced Detection Technology Sensor (ADTS) Ka-band radar operated by Lincoln Laboratory provides $1\text{-ft} \times 1\text{-ft}$ resolution and fully polarimetric output [16]. It has been used to collect more than two hundred square kilometers of data to support Lincoln Laboratory's automatic target-recognition studies.

An important issue in target recognition is the

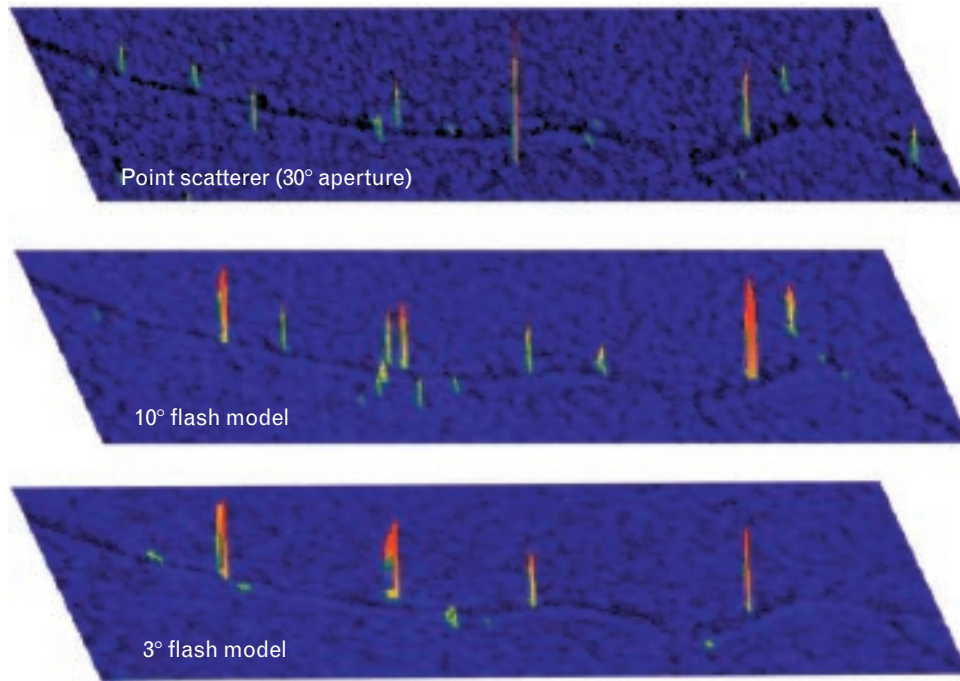


FIGURE 15. Vector image of the Cedar Swamp deployment. By employing 10° and 3° flash models, HDVI further reduces clutter false alarms. The flash models reject the corner reflector, which is the large return in the center of the point-scatterer image.

resolution required for reliable performance. ADTS studies demonstrate that a resolution of one foot is sufficient for reliable target recognition. Operational radars, however, often provide coarser resolution and, correspondingly, markedly reduced recognition performance. For example, the resolution of the SAR

mounted on the TIER II+ unmanned air vehicle is one meter in wide-area (stripmap) mode. One of the motivations, then, for HDVI was to improve the resolution of a one-meter radar with the hope of improving target recognition. This goal has been achieved [12].

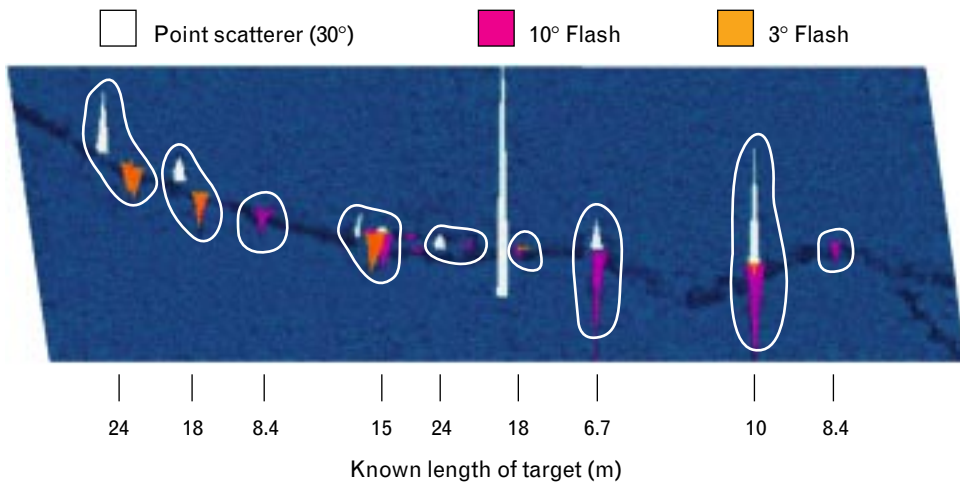


FIGURE 16. Vector-image signature analysis. The vehicles exhibit multiple signatures, and vehicle length is inversely proportional to flash size. Hence the vector image reveals more information about the targets than the conventional image shown in Figure 14.




FIGURE 17. Comparison of conventional imaging (left) and HDVI (right) of two vehicles at Ka-band. HDVI provides improved resolution, reduced sidelobes, and reduced speckle. The HDVI algorithm employed here is Capon's technique with both the subspace and quadratic constraints.

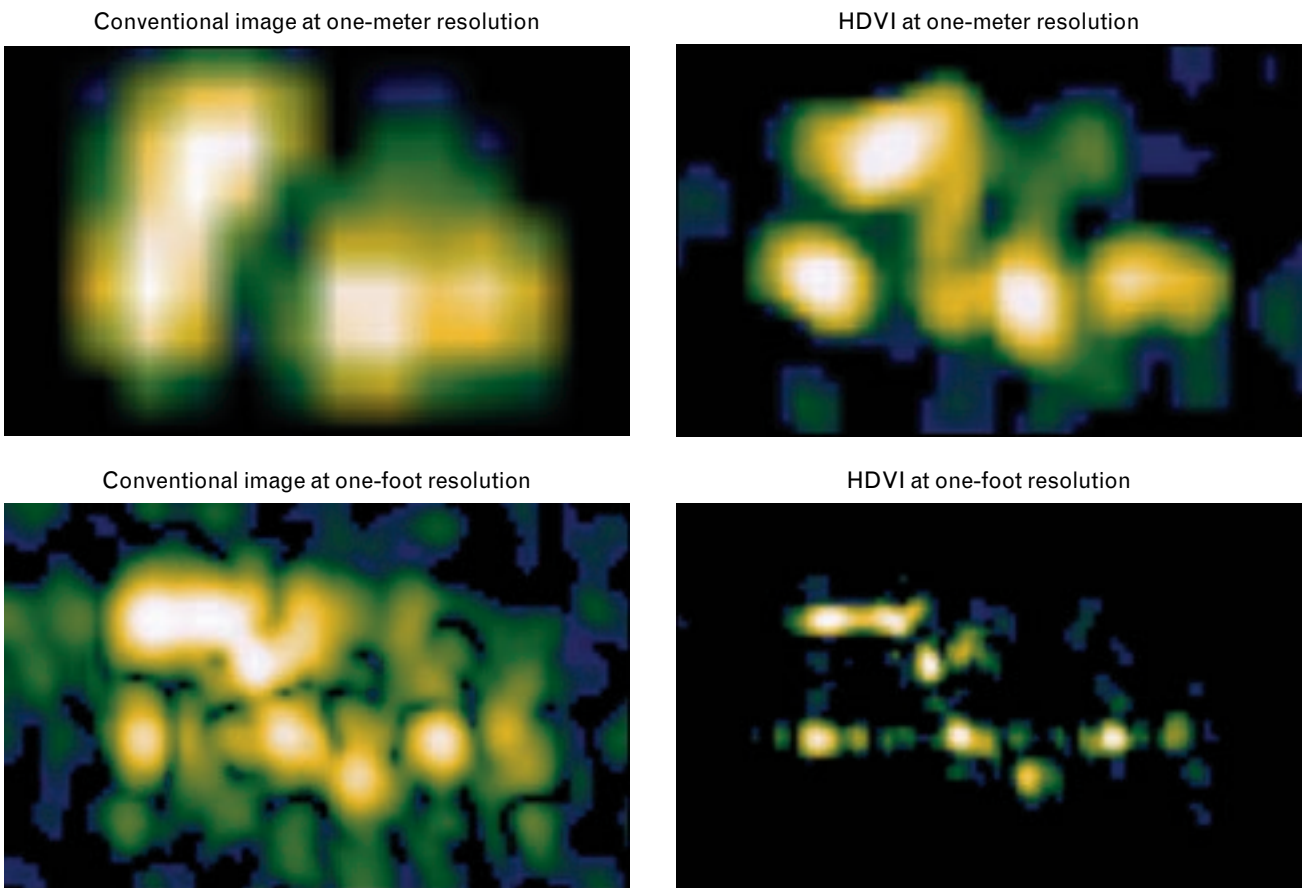


FIGURE 18. Comparison of imaging techniques at one-meter and one-foot resolution, showing the rightmost vehicle in Figure 17. HDVI at one-foot resolution clearly improves on the conventional image at one-foot resolution. HDVI, using one-meter data, approaches the one-foot conventional image in appearance, but with lower sidelobes.




FIGURE 19. Comparison of imaging techniques at one-meter resolution: (a) conventional, (b) incoherent-look two-dimensional MLM, and (c) coherent-look two-dimensional MLM. The vehicles, from top to bottom, are a Howitzer, a personnel carrier, and a tank. Quantitative measurements showing improvement in the lobe widths, target-to-clutter ratio, and speckle appear in Table 3.

The first result with the ADTS radar was for two targets of opportunity at Westover Air Reserve Base, Massachusetts, as shown in Figure 17. This figure compares conventional imaging using Hamming weighting with HDVI employing Capon's technique with norm and subspace constraints, and with an incoherent combination of looks. (The sidelobes run diagonally because north is toward the top of the image and the radar is at about two o'clock.) This comparison demonstrates resolution improvement,

sidelobe reduction, and speckle reduction. Speckle reduction results from the incoherent combination of looks in Capon's technique.

Figure 18 provides a close-up of the rightmost vehicle shown in Figure 17. This figure compares conventional imaging with HDVI (coherent look) at resolutions of one foot and one meter. The one-meter data were extracted from the one-foot data through a coherent spoiling process. While one-meter data processed with HDVI clearly do not achieve the same re-

Table 3. Statistics for Images in Figure 19			
	<i>Lobe Width</i>	<i>Target-to-Clutter Ratio</i>	<i>Speckle*</i>
Conventional	1.04 m	31.8 dB	5.8 dB
Incoherent MLM	0.59 m	31.9 dB	3.8 dB
Coherent MLM	0.58 m	33.5 dB	5.8 dB

* Standard deviation

sult as one-foot data processed conventionally, the general appearance of results from the two techniques is the same, and the sidelobes in HDVI are lower.

Figure 19 shows results for military vehicles at one-meter resolution, and provides a comparison of the coherent-look and incoherent-look versions of HDVI (Capon's technique with norm and subspace constraints). The vehicles, from top to bottom, are a Howitzer, a personnel carrier, and a tank. The improvement in resolution is indicated by the 3-dB lobe width, which decreases from 1.04 m to 0.58 m (see Table 3). Target-to-clutter ratios are calculated as the ratio of the brightest return to the average of the surrounding clutter; these ratios show a slight improvement with coherent-look HDVI. Speckle reduction is indicated by the standard deviation of the dB-valued images evaluated over a patch of clutter: the conventional image shows 5.8 dB and incoherent-look HDVI shows 3.8 dB—a marked reduction.

The two versions of HDVI are visibly and statistically different, with the coherent-look technique providing slightly higher target-to-clutter ratio and higher speckle. It is worth noting that the incoherent-look technique provides the best performance with a template-matching classifier, while the coherent-look technique provides the sharper image preferred by image analysts.

Conclusion

High-definition vector imaging is a new approach to processing SAR data. It employs modern spectrum-estimation techniques to derive more information from SAR data. Benefits include improved resolution, reduced sidelobes, reduced speckle, reduced clutter, and better discrimination based on scattering phenomenology. HDVI provides improved images for exploitation by image analysts and more information for automatic target recognition.

We describe and demonstrate several HDVI techniques. Versions of Capon's adaptive-beamforming technique have demonstrated improved resolution, and reduced sidelobes, speckle, and clutter. A version of the MUSIC algorithm proved useful in discriminating targets from clutter for UHF SAR data. We also discuss modifications of the techniques, the application of these modified techniques to SAR data,

and the trade-offs between algorithm robustness and resolution.

The implication of HDVI technology is that radar performance can be significantly improved via digital signal processing. Existing systems can achieve improved performance, and new system designs could be less expensive. Further extensions of HDVI could prove beneficial to more advanced problems such as change detection, moving-target exploitation, and interference cancellation. Continual advances in computer design make HDVI a practical and affordable option.

Acknowledgments

I would like to acknowledge the assistance of Dennis Blejer in providing me with Rail SAR data, and the assistance of the Surveillance Systems group at Lincoln Laboratory for providing me with ADTS and FOPEN data.

REFERENCES

1. S.L. Borison, S.B. Bowling, and K.M. Cuomo, "Super-Resolution Methods for Wideband Radar," *Linc. Lab. J.* 5 (3), 1992, pp. 441–462.
2. S.R. DeGraaf, "Parametric Estimation of Complex 2-D Sinusoids," *4th Annual Workshop on Spectrum Estimation and Modeling, Minneapolis, 3–5 Aug. 1988*, pp. 391–396.
3. H.C. Stankwitz, R.J. Dailaire, and J.R. Fienup, "Nonlinear Apodization for Sidelobe Control in SAR Imagery," *IEEE Trans. Aerosp. Electron. Syst.* 31 (1), 1995, pp. 267–279.
4. S.R. DeGraaf, "Sidelobe Reduction via Adaptive FIR Filtering in SAR Imagery," *IEEE Trans. Image Proc.* 3 (3), 1994, pp. 292–301.
5. J. Li and P. Stoica, "An Adaptive Filtering Approach to Spectral Estimation and SAR Imaging," *IEEE Trans. Signal Process.* 44 (6), 1996, pp. 1469–1484.
6. J.W. Odendaal, E. Bernard, and C.W.I. Pistorius, "Two Dimensional Superresolution Radar Imaging Using the MUSIC Algorithms," *IEEE Trans. Antennas Propag.* 42 (10), 1994, pp. 1386–1391.
7. K.W. Forsythe, "Utilizing Waveform Features for Adaptive Beamforming and Direction Finding with Narrowband Signals," *Linc. Lab. J.* 10 (2), 1997, in this issue.
8. L.M. Novak, G.J. Owirka, W.S. Brower, and A.L. Weaver, "Automatic Target Recognition with Enhanced Polarimetric SAR Data," *Linc. Lab. J.*, in this issue.
9. W.J. Carrara, R.S. Goodman, and R.M. Majewski, *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms* (Artech House, Boston, 1995).
10. J. Capon, "High-Resolution Frequency-Wavenumber Spectrum Analysis," *Proc. IEEE* 57 (8), 1969, pp. 1408–1419.
11. H. Cox and R. Zeskind, "Reduced Variance Distortionless Response (RVDR) Performance with Signal Mismatch," *25th Asilomar Conf. on Signals, Systems & Computers 2, Pacific Grove, Calif., 4–6 Nov. 1992*, pp. 825–829.
12. B. Van Veen, "Minimum Variance Beamforming," chap. 4 in *Adaptive Radar Detection and Estimation*, S. Haykin and A. Steinhardt, eds. (Wiley, New York, 1992), pp. 161–236.
13. R. Schmidt, "Multiple Emitter Location and Signal Parameter Estimation," *Proc. RADC Spectrum Estimation Workshop, Rome, N.Y., 3–5 Oct. 1979*, pp. 243–258.
14. S. Barbarossa, L. Marsili, and G. Mungari, "SAR Super-Resolution Imaging by Signal Subspace Projection Techniques," *AEÜ Int. J. Electron. Commun.* 50 (2), 1996, pp. 133–138.
15. D.F. DeLong and G.R. Benitz, "Extensions of High Definition Imaging," *SPIE* 2487, 1995, pp. 165–180.
16. L.M. Novak, M.C. Burl, R.D. Chaney, and G.J. Owirka, "Optimal Processing of Polarimetric SAR Automatic Target Recognition System," *Linc. Lab. J.* 3 (2), 1990, pp. 273–290.



GERALD R. BENITZ is a staff member in the Advanced Techniques group. He received a B.S. degree in electrical engineering from Purdue University, and an M.S. degree in electrical engineering, an M.A. degree in mathematics, and a Ph.D. degree in electrical engineering, all from the University of Wisconsin, Madison. His thesis work was an investigation of the asymptotic theory of optimal quantization and detection. His areas of research at Lincoln Laboratory have included analysis, development, and testing of algorithms for direction finding and waveform estimation with adaptive antenna arrays. His current research includes applications of adaptive techniques to image processing. He is a member of the Information Theory Society of the IEEE.