AGU100 ADVANCING EARTH AND SPACE SCIENCE

Radio Science

RESEARCH ARTICLE
10.1029/2019RS006805

Aperture-Synthesis Radar Imaging With Compressive Sensing for Ionospheric Research

D. L. Hysell[1], **P. Sharma**[2], **M. Urco**[3], **and M. A. Milla**[4]

[1]Earth and Atmospheric Sciences, Cornell University, Ithaca, NY, USA, [2]Electrical and Computer Engineering, Cornell University, Ithaca, NY, USA, [3]Leibniz Institute for Atmospheric Physics, Kuehlungsborn, Germany, [4]Jicamarca Radio Observatory, Lima, Peru

Key Points:
- Compressive sensing inverse methods have been applied to aperture synthesis radar imaging of ionospheric plasma density irregularities
- Performance of basis pursuit denoisint (BPDN) and orthogonal matching pursuit (OMP) are generally inferior to that of the maximumentropy method (MaxENT)
- Computational speed of OMP is attractive and prompts research into more suitable function library

Correspondence to:
D. L. Hysell,
dlh37@cornell.edu

Citation:
Hysell, D. L., Sharma, P., Urco, M., & Milla, M. A. (2019). Aperture-synthesis radar imaging with compressive sensing for ionospheric research. *Radio Science*, *54*, 503–516. https://doi.org/10.1029/2019RS006805

Received 28 JAN 2019
Accepted 30 APR 2019
Accepted article online 31 MAY 2019
Published online 17 JUN 2019

Abstract

Inverse methods involving compressive sensing are tested in the application of two-dimensional aperture-synthesis imaging of radar backscatter from field-aligned plasma density irregularities in the ionosphere. We consider basis pursuit denoising, implemented with the fast iterative shrinkage thresholding algorithm, and orthogonal matching pursuit (OMP) with a wavelet basis in the evaluation. These methods are compared with two more conventional optimization methods rooted in entropy maximization (MaxENT) and adaptive beamforming (linearly constrained minimum variance or often "Capon's Method.") Synthetic data corresponding to an extended ionospheric radar target are considered. We find that MaxENT outperforms the other methods in terms of its ability to recover imagery of an extended target with broad dynamic range. Fast iterative shrinkage thresholding algorithm performs reasonably well but does not reproduce the full dynamic range of the target. It is also the most computationally expensive of the methods tested. OMP is very fast computationally but prone to a high degree of clutter in this application. We also point out that the formulation of MaxENT used here is very similar to OMP in some respects, the difference being that the former reconstructs the logarithm of the image rather than the image itself from basis vectors extracted from the observation matrix. MaxENT could in that regard be considered a form of compressive sensing.

1. Introduction

Aperture-synthesis methods have been used in radio astronomy since the 1950s to form images of distant radio sources from spaced-receiver interferometry data. The signals from the receivers are stochastic and exhibit jointly normal multivariate statistics. All of the information is therefore contained in the second-order statistics or the spatial covariances or "visibilities." The transformation between the measured visibility data and the desired brightness function, which specifies the radiation intensity versus bearing, is linear and closely related to a Fourier transform (Thompson, 1986). Aperture-synthesis methods are inverse methods for computing the transformation on the basis of sparse and noisy data. The problem is generally underdetermined and poorly conditioned.

Only fairly recently have aperture synthesis methods been applied to radar observations of the upper atmosphere and ionosphere (see, e.g., Hysell, 1996; Hysell et al., 2002, 2004, 2008; Kudeki & Sürücü, 1991; Saito et al., 2006, 2008; Sommer & Chau, 2016; Urco, Chau, Milla, et al., 2018; Urco, Chau, Weber, et al., 2018). The problem is similar to the one in radio astronomy with a few important differences. For one, the number of receivers used in upper-atmospheric radar applications has so far been relatively small, spaced-receiver imaging being mainly an afterthought in radar design. For another, the required cadence of measurements in upper-atmospheric radar applications can be very high. Images typically need to be formed in tens of Doppler frequency bins and in hundreds or thousands of range gates at an experimental cadence of about once per second. This makes computational efficiency critical.

Most crucially, whereas astronomical images are typically characterized by point sources, radar targets in the upper atmosphere and ionosphere tend to be spatially extended and "blobby." The targets are much weaker at the periphery than at the center, and preserving the boundaries requires methods with high dynamic range. The faithful reconstruction of edges is also important for inferring velocities from time series. Metrics for evaluating competing methods should be designed around this requirement.

The most widely used imaging methods in upper-atmospheric research include adaptive beamforming methods like the one described by Capon (1969), iterative deconvolution methods like CLEAN (Högbom, 1974), and Bayesian optimization methods rooted in maximum entropy (e.g., Skilling & Bryan, 1984). In this paper, we consider methods arising from recent advances in the field of compressive sensing. Compressive sensing for one-dimensional imaging was evaluated recently by Harding and Milla (2013) who examined coherent backscatter from ionospheric plasma density irregularities at the magnetic equator. We pursue the same problem, this time considering imaging in two dimensions. The results should apply equally well to ionospheric scatter at low, middle, and high latitudes as well as to scatter from index-of-refraction variations in the neutral lower, middle, and upper atmosphere.

2. Compressive Sensing and Radar Imaging

Compressive sensing has occupied a central role in image processing research for more than a decade and has significantly increased the acuity of any number of experimental modalities (see Mackenzie, 2009, and references therein for review). The idea involves reproducing undersampled signals accurately by exploiting inherent sparseness of the samples in an appropriately chosen basis. Two-dimensional photographic images can often be represented very accurately by a small number of nonzero coefficients when expressed in a wavelet basis, for example. This implies not only the efficacy of image compression but also an opportunity for novel sampling strategies that incorporate sparseness into their design. The practical advantages of restoring a signal from sparse samples, as opposed to sampling the signal fully in the Nyquist sense, compressing the results, and then decompressing them later, are obvious. Moreover, the prospect of recovering complex signals from a small number of samples or from a small number of nonadaptive sensors has tremendous appeal in scientific realms where dense sensor arrays may be impractical to deploy to begin with. Aperture-synthesis radar imaging, which seeks the best estimate of the image brightness on the basis of a few measurements of the visibility, belongs in this realm.

The obvious objective function for evaluating sparseness is the $l_\circ$ pseudonorm, $||x||_\circ$, the number of nonzero entries in the vector x. If $x \in \mathbb{R}^m$ is a state vector which is s-sparse (has at most s nonzero entries) and is constrained by a data vector $y \in \mathbb{R}^n$ through an observation model $Ax = y$, where $A \in \mathbb{R}^{n \times m}$ is a noninvertible observation matrix, then a reasonable model for x could be a vector with the smallest number of nonzero values which satisfies the observation model, namely,

$$x = \underset{x}{\text{argmin}}\ ||x||_\circ\ :\ Ax - y = 0 \tag{1}$$

which is a constrained optimization problem. This model yields not just accurate but exact recovery of undersampled state vectors under the condition that A is one-to-one in all 2s-sparse vectors. However, this is prohibitively complex (hard) optimization problem in computational enumeration.

In practice, the problem in (1) may be replaced by a closely related problem known as basis pursuit (Donoho, 2006). Here, the $l_\circ$ pseudonorm is relaxed to the l_1 norm, that is,

$$x = \underset{x}{\text{argmin}}\ ||x||_1\ :\ Ax - y = 0 \tag{2}$$

where the l_1 norm enforces sparsity in a manner similar to the $l_\circ$ norm. The simple modification transforms the task from an enumeration problem to a problem in convex optimization which can be solved practically using linear programming methods.

Numerous sufficient conditions for unique recovery of an s-sparse vector x from basis pursuit and related approaches are discussed in the literature including the restricted isometry property (Candés et al., 2006; Candés & Tao, 2006), the exact recovery condition (Tropp, 2004, 2006), and the mutual incoherence condition (MIC; Donoho & Huo, 2001; Tropp, 2006). The last of these options is the most intuitive and most practical to evaluate. The mutual coherence for a matrix A with columns with unity l_2 norms is defined in terms of

$$\mu_A = \max_{i \neq j}\ |A_i^T A_j| \tag{3}$$

that is, the maximum pairwise column correlation. The sufficient condition for sparse signal recovery is that $\mu_A < (2s-1)^{-1}$. This is also the worst-case necessary condition (Tropp, 2004).

3. BPDN

Signal sparsity is approximate rather than exact in practice, and the samples themselves are contaminated by noise. In this case, the basis pursuit problem is further relaxed to the nearby basis pursuit denoising (BPDN) problem which includes the provision for a finite residual which is weighted together with the sparsity penalty in the objective function

$$x = \underset{x}{\operatorname{argmin}} \left(||Ax - y||_2^2 + \lambda ||x||_1 \right) \tag{4}$$

Equation (4) has the form of a regularized least squares regression problem with λ being the regularization parameter. While equation (4) is sometimes referred to as having Lagrangian form, λ is not an undetermined Lagrange multiplier here since the l_1 norm term does not function as a constraint. The λ parameter merely represents the trade-off between the fidelity of the experiments and the noise sensitivity.

Note that two other problems are closely related to BPDN

$$x = \underset{x}{\operatorname{argmin}} \; ||x||_1 \; : \; ||Ax - y||_2^2 \leq \epsilon \tag{5}$$

$$x = \underset{x}{\operatorname{argmin}} \; ||Ax - y||_2^2 \; : \; ||x||_1 \leq \epsilon \tag{6}$$

where ϵ is a tunable parameter like λ. All three problems have the same solutions for the appropriate choices of λ and ϵ which will depend on the data. The variant of the problem in (5) will be discussed further below. The variant in (6) incorporates the least absolute shrinkage and selection operator (Lasso; Tibshirani, 1996). An “elastic net” variant of the problem adds a regularization term based on the l_2 norm of x to the objective function.

While the observation matrix and the data in radar imaging experiments are complex, we consider the real and imaginary parts of the observing model separately, making n twice the number of interferometry baselines. This greatly simplifies the handling of the l_1 norm but necessitates an explicit treatment of the real and imaginary parts of the data in the error analysis (Hysell & Chau, 2006). The positive definite symmetric inverse error covariance matrix C_d^{-1} given there can be factored as $C_d^{-1} = C_d^{-T/2} C_d^{-1/2}$. In order to incorporate error propagation, A and y are prewhitened through scaling by the square root information matrix $C_d^{-1/2}$, that is, $(Ax - y) \rightarrow C_d^{-1/2}(Ax - y)$.

Likewise, a change of basis can be accommodated through the transformation $Ax \rightarrow (AW^{-1})(Wx)$ where W is a linear transformation from state space to a basis where sparsity is optimized. For example, as discussed below, it is commonplace in compressive sensing to transform to a wavelet basis in which imagery can sometimes be represented very sparsely. The MIC condition then applies to the new observation matrix AW^{-1}. In this case, the last step of the algorithm is the transformation of the solution vector x back to image space.

Note, finally, that the l_1 regularization problem can be cast in the form of a constrained quadratic optimization problem by defining $x = u - v$, where u and v are the parts of x that are positive and negative, respectively, and restricting the solution domain to the region $u \geq 0, v \geq 0$ so that $||x||_1 \rightarrow u + v$. This reformulation invites the application of any number of popular optimization methods rooted in linear programming. The penalty is the doubling in length of the solution vector.

A range of approaches can be applied to BPDN and related problems (see Rani et al., 2018; Yang et al., 2013 for reviews). Here, we pursue the fast iterative shrinkage thresholding algorithm (FISTA), a fast gradient algorithm which is simple to implement for aperture-synthesis imaging (Beck & Teboulle, 2009). The algorithm depends on the concept of the proximal gradient which is summarized briefly below.

Gradient algorithms must cope with the fact that l_1 norm of a function is nondifferentiable. An option for convex nondifferentiable functions is to substitute the gradient operator with the subgradient operator which bounds the gradient. The subgradient of the l_1 norm of x is simply sign(x). The subgradient method applied to BPDN can be solved iteratively using a conventional gradient descent method, but convergence will be slow.

In a proximal gradient method, which can be viewed as a generalized gradient descent method, the objective function is divided into two parts, that is, $x = \operatorname{argmin}_x \; (g(x) + h(x))$, where $g(x)$ is convex and differentiable and $h(x)$ is convex and possibly nondifferentiable. Equation (4) is obviously in this form. The idea then is to iteratively minimize the sum of h and a quadratic local model of g. Toward this end, define the proximal operator of the function h about z as

$$\mathrm{prox}_{\alpha h}(z) = \underset{x}{\mathrm{argmin}} \left(h(x) + \frac{1}{2\alpha} ||x - z||_2^2 \right) \tag{7}$$

Each iteration proceeds from the results of the last, shifted opposite the direction of the local gradient of g, namely,

$$x^k = \mathrm{prox}_{\alpha_k h}(x^{k-1} - \alpha_k \nabla g(x^{k-1})) \tag{8}$$

where the step size α_k can either be a fixed constant or determined by a line search. For the l_1-regularized quadratic regression problem, stability is guaranteed by setting α_k to a constant equal to the reciprocal of the smallest Lipschitz constant of ∇g which is twice the largest eigenvalue of $A^T A$. This in turn can be calculated using the power method with negligible computational burden.

In the case of the regularized least squares problem, $g(x)$ is quadratic, and its gradient has an elementary linear form, that is, $A^T(Ax - y)$. Furthermore, for $h(x) = \lambda ||x||_1$, the proximal operator or shrinkage function is given by the iterative soft thresholding function

$$\mathrm{prox}_{\lambda h}(y) = \begin{cases} y - \mathrm{sign}(y)\lambda & |y| \geq \lambda \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

In the event that only positive values of x are admissible, as is the case for some formulations of the radar-imaging problem, that condition can be readily incorporated in equation (9) as well.

The iterative shrinkage thresholding algorithm implied by equation (8) can be shown to converge at a rate that is linear in k like a standard gradient descent method under the condition that ∇g is continuous in the L-Lipschitz sense, for both fixed and variable step sizes. The FISTA algorithm, meanwhile, is able to achieve a quadratic rate of convergence by incorporating Nesterov's acceleration method (Nesterov, 1983). The modified algorithm for FISTA, including acceleration, is (Beck & Teboulle, 2009)

$$z^k = \mathrm{prox}_{\alpha_k \lambda h}(x^{k-1} - \alpha_k \nabla g(x^{k-1})) \tag{10}$$

$$t_k = \frac{1}{2}\left(1 + \sqrt{1 + 4t_{k-1}^2}\right) \tag{11}$$

$$x^k = z^k + \frac{t_{k-1} - 1}{t_k}(z_k - z_{k-1}) \tag{12}$$

with equation (9) giving the proximal gradient. The algorithm is iterated to convergence with convergence criteria typically being based on the norm of $x^k - z^k$.

Note that a parallelized version of FISTA has been developed by Peng et al. (2013). For the aperture-synthesis imaging application, parallelization can be exploited without reformulating the basis algorithm simply by computing images for different Doppler spectral bins and range gates in separate processes.

4. OMP

A computationally expedient and highly intuitive alternative for sparse least squares estimation is orthogonal matching pursuit (OMP; see, e.g., Cai et al., 2010; Cai & Wang, 2011; Tropp & Gilbert, 2007). This is a greedy algorithm that formulates estimates of the state vector by applying the least squares pseudoinverse to a submatrix of the original measurement matrix A. The submatrix starts from a null matrix and grows with each iteration with the addition of a column selected and copied from A. The column selected is the one with the highest correlation with the current residual. Iteration continues until a stopping criterion is met.

At this point, the state estimate is consistent with the data while being maximally sparse. In terms of the discussion above, OMP solves the compressive sensing problem stated in (5) which, like BPDN, makes allowance for observation noise. OMP begins from an essentially minimalist representation of the state vector and expands it in the l_1 norm sense until the residual falls below a specified threshold. As with BPDN, it is often expedient to transform the problem using a basis in which the solution is especially sparse.

Define $A_k \in \mathbb{R}^{n\times k}$ to be a submatrix of A with its k columns having been copied from A. At each iteration, k is increased as the column of A most highly correlated with the current residual r is appended to A_k. The residual at the kth iteration is defined as follows:

$$r^k = y - A\hat{x}^k \tag{13}$$

$$= (I - A\tilde{A}_k)y \tag{14}$$

$$\tilde{\mathrm{A}} \equiv (A^T A)^{-1}A^T \tag{15}$$

Here, $\tilde{\mathrm{A}}$ is the standard least squares pseudoinverse, and the estimator $\hat{x}^k = \tilde{A}_k y \in \mathbb{R}^m$ is defined so as to have $k < m$ nonzero elements with indices that correspond to the indices of the columns copied from A in the order in which they were copied. The index of the selected column is the index of the term in $A^T r$ with the largest modulus.

Allowances for statistical errors can be introduced through incorporation of the square root data covariance matrix $C_d^{-1/2}$ described earlier. Note also that additional regularization can be introduced here through the appropriate augmentation of the least squares pseudoinverse. Stopping criteria may be based either on the size of the residual or the anticipated sparsity of the solution. Absent the explicit introduction of regularization, OMP has no tuning parameters comparable to the λ parameter in BPDN. The choice of stopping conditions in the standard algorithm is subjective, however, particularly if neither the sparsity nor the data error covariances are known a priori (Kallummil & Kalyani, 2017). Note that the number of iterations can be no larger than the number of data n. When the number of basis functions is larger than the number of data, the problem becomes underdetermined, and the least squares pseudoinverse ceases to exist. After n iterations, the problem is even determined, and the residual is identically zero.

5. Comparison Using Synthetic Data

It is illustrative to benchmark the most common algorithms applied to the aperture synthesis imaging problem. Among them is the adaptive beamforming method often attributed to Capon (1969) and referred to as the linearly constrained minimum variance (LCMV) or the minimum variance distortionless response method (or simply as “Capon's method.”) We also consider the maximum entropy method (MaxENT) as formulated by Wilczek and Drapatz (1985). The relative performance of the two methods applied to the aperture-synthesis imaging problem has been evaluated by Yu et al. (2000). Another popular algorithm for aperture synthesis imaging is CLEAN (Högbom, 1974). In hindsight, this can be viewed as a special case of OMP and so will not be considered explicitly here.

We consider synthetic data comparable to what would be acquired by the Jicamarca Radio Observatory measuring coherent scatter from field-aligned plasma density irregularities in the equatorial ionosphere. The typical receiving antenna configuration is shown in Figure 1. Reception is performed using eight antenna submodules. Eight antennas imply 28 nonredundant interferometry baselines or 29 including the zero baseline. (The number of data, n, is therefore 58). The radar operates at 50 MHz, and so the longest interferometry baseline, 569 m, is approximately 95 wavelengths long.

Since the scattering irregularities are aligned with the geomagnetic field, the spatial correlation length of the backscatter in the north-south direction is very long and just barely measurable at Jicamarca (Farley et al., 1981). The most important information is to be found in the correlation length of the backscatter in the east-west direction, but there is utility in measuring the centroid of the backscatter in the north-south direction which varies somewhat with range.

We consider an imaging domain spanning ±0.1 rad in the plane of the magnetic equator and ±0.02 rad in the plane of the magnetic meridian. The former figure is roughly the effective field of view illuminated by the Jicamarca radar in imaging mode. The synthetic target is an elongated Gaussian ellipsoid with a half

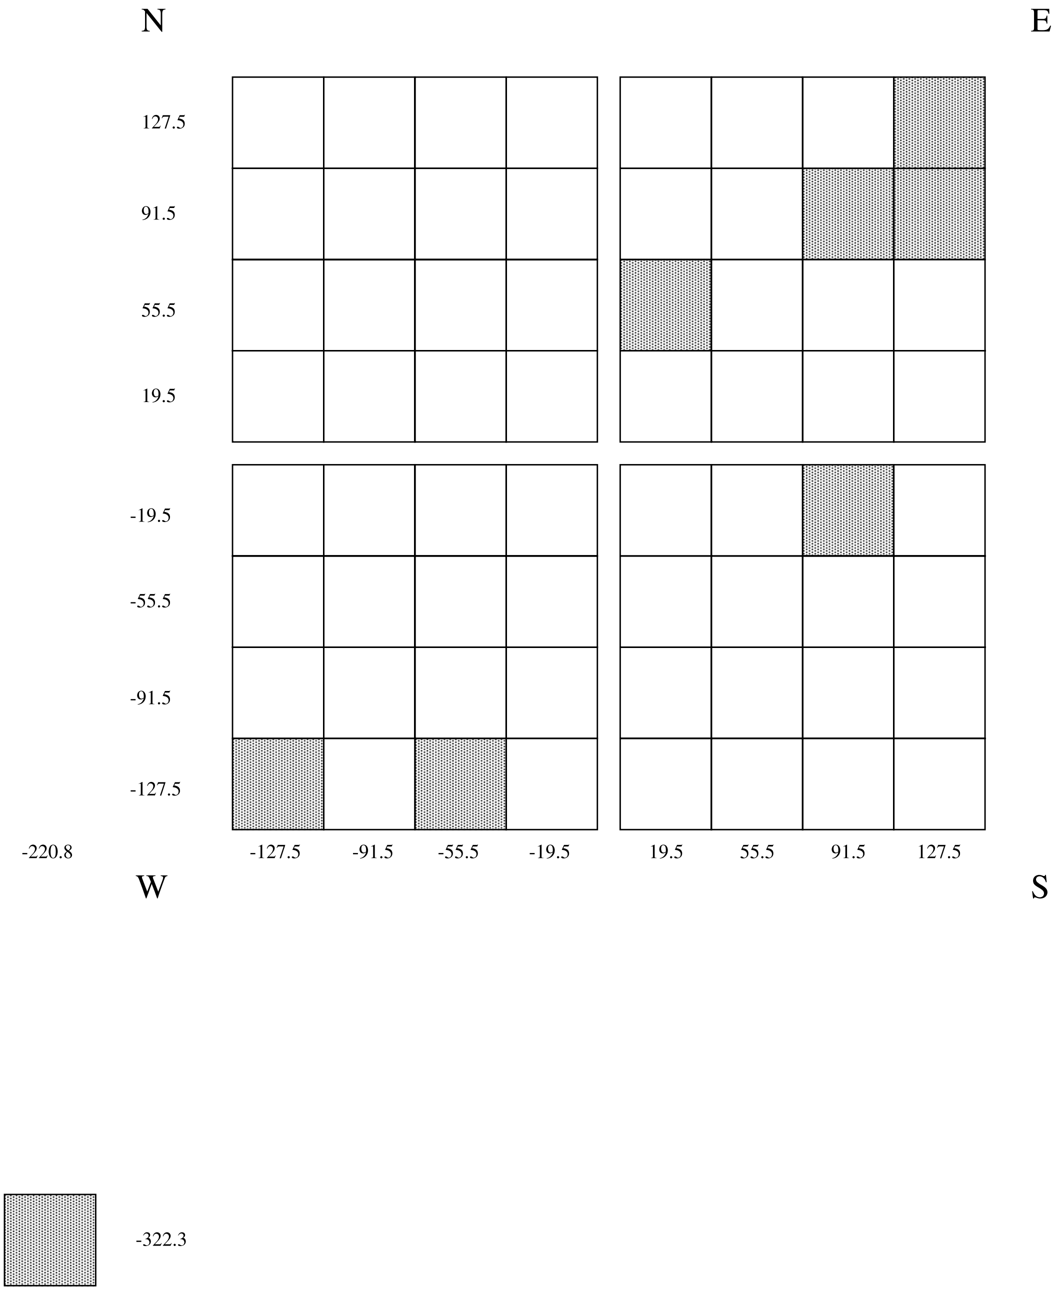

Figure 1. Configuration of the antenna array of the Jicamarca Radio Observatory in imaging mode. The corners of the main array are nearly aligned with the cardinal directions, and the north-south line is closely aligned with the direction of the geomagnetic field at present. Module coordinates with respect to the array center are indicated. The distance between adjoining modules is 36 m, and the longest baseline is 569 m. Seven modules of the main array plus an eighth outrigger array are used for imaging.

width in the plane of the magnetic equator 10 times wider than in the plane of the magnetic meridian. The image will be constructed in a domain 128 pixels in the plane of the magnetic equator by 32 pixels in the plane of the magnetic meridian. The observing matrix for the problem is given by Hysell and Chau (2006).

Normally distributed independent noise is added to the synthetic visibility data for nonzero lags at the 2% level. This is a simplified treatment for observation noise and represents an upper bound for the experimental uncertainty associated with a signal-to-noise ratio larger than unity and averages of 2,500 statistically independent samples. In actual experiments, noise bias in the zero-baseline data is estimated and removed,

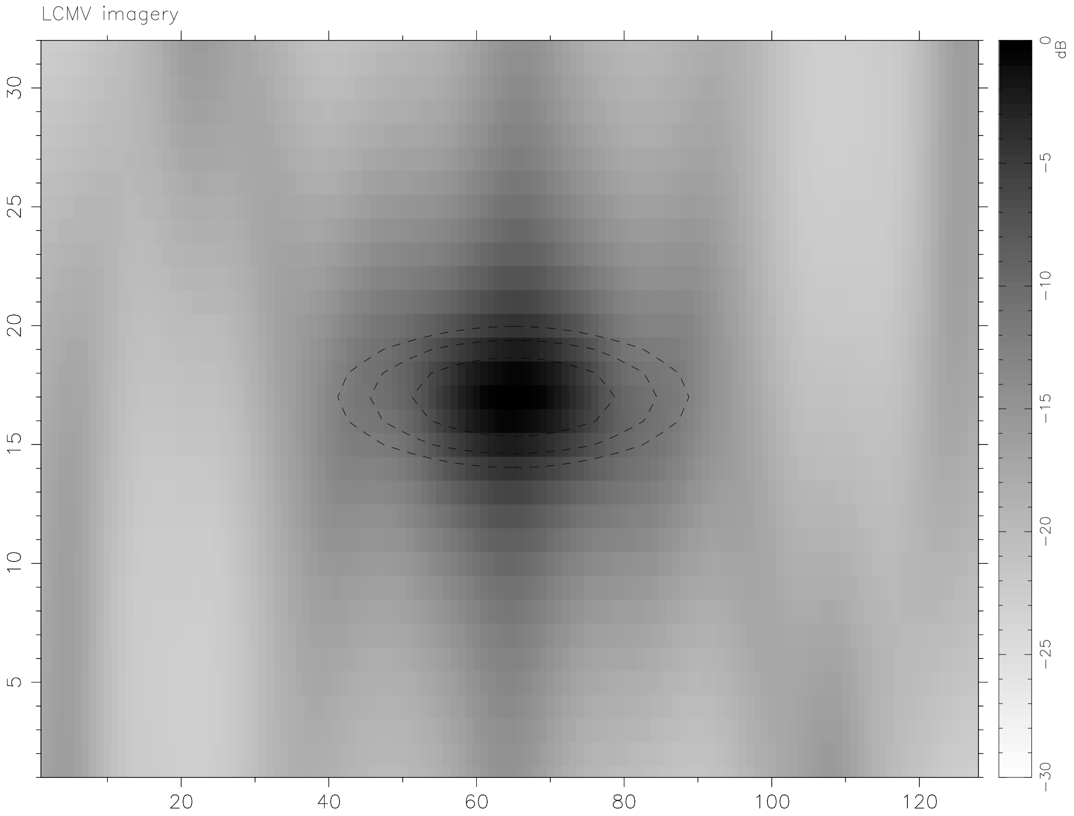

Figure 2. Image recovered from synthetic data using the LCMV algorithm. The imaging domain spans ±0.1 rad in the plane of the magnetic equator (horizontal direction) and ±0.02 rad in the plane of them magnetic meridian (vertical). The resolution of the image is 128 × 32. Gray scales denote brightness relative to the maximum in decibels. Red dashed contours represent the −10-, −20-, and −30-dB contours of the truth model, a Gaussian ellipsoid. LCMV = linearly constrained minimum variance.

and so no noise bias will be added here. For an exhaustive treatment of error analysis and propagation in aperture-synthesis imaging, see Hysell and Chau (2006).

Figure 2 shows the results of image recovery using the LCMV method which serves here as a baseline. Gray scales indicate relative backscatter intensity in decibels relative to the intensity at the center of the target. We plot 30 dB of dynamic range in the figure. In practical radar experiments, radar clutter from pulse coding is generally present at the level of about −22 to −25 dB. We desire about 30 dB of usable dynamic range from imaging methods so that pulse coding rather than radar imaging will be the limiting factor for radar clutter. The elliptical contours represent the truth model and indicate backscatter at the −10-, −20-, and −30-dB levels, respectively.

LCMV has no tuning parameters and also no means of incorporating estimates of measurement confidence levels in the analysis. The method has accurately recovered the strongest intensity region within the 10-dB contour. However, the method has done a poor job of rejecting clutter outside the 30-dB contour. (We define clutter as extraneous intensity outside the 30-dB boundary of the truth model.) Strong artifacts are present across the image. The clutter level rises sharply as the level of random fluctuations added to the synthetic data is increased.

Figure 3 shows the results of image recovery using the FISTA algorithm. For this test, we do not incorporate wavelet transforms and simply minimize the l_1 norm of the image brightness itself. Since the intensity is nonnegative, we can and do incorporate that information in the computation of the proximal gradient to speed computation.

There is one tunable parameter in FISTA—the regularization parameter λ. The larger the value of λ, the greater the tendency toward sparse solutions. Here and in the methods to follow, λ has been adjusted so as

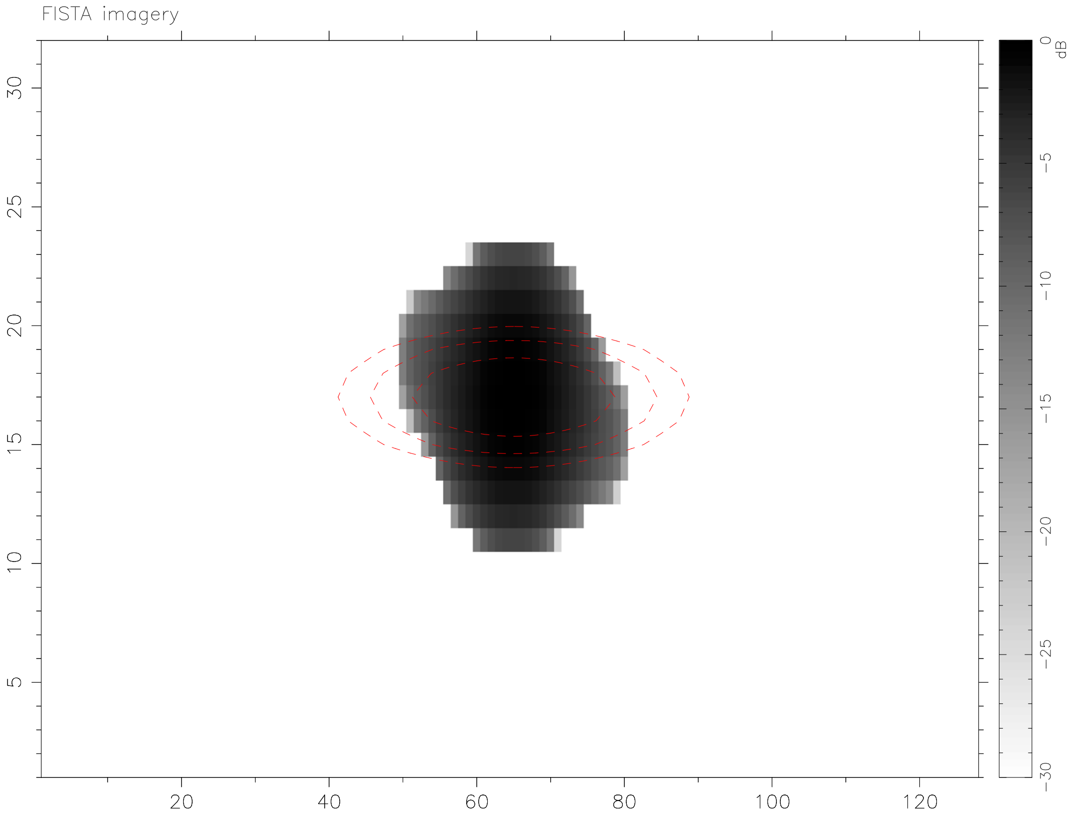

Figure 3. Image recovered from synthetic data using the FISTA algorithm. FISTA = fast iterative shrinkage thresholding algorithm.

to make the chi-square parameter equal to the number of data *n*. This has the effect of limiting clutter while adequately filling the truth model contours.

FISTA recovers an elliptical target with approximately the same width in the zonal direction as the truth model but with more than twice the width in the meridional direction. This is a consequence of the fact that the interferometry baselines in the meridional direction are relatively short, while the target itself is relatively narrow. (Note that highly field-aligned targets have long spatial correlation lengths and elongated visibilities in the direction of the magnetic meridian and, consequently, narrow angular meridional widths in brightness imagery.) This is not an important shortcoming in practice where 1-D images are normally extracted from the 2-D images either by taking a horizontal cut or by averaging over a narrow range of meridional coordinates. The imaging resolution in the meridional direction need only be sufficient to allow us to estimate the centroid of the backscatter.

Clutter is mainly limited to bleedthrough in the direction of the magnetic meridian. Clutter is largely excluded from the horizontal bisector of the image corresponding to the magnetic equator. Clutter associated with interferometry sidelobes is completely suppressed. However, the dynamic range of the FISTA image is also limited; pixels appear to be either fully set or fully unset throughout most of the image. Gradation in image intensity appears to be largely lost using this and any method rooted in compressive sensing.

Figure 4 shows the results of image recovery using the OMP algorithm. For this test, we have used a transformation *W* consistent with a two-dimensional wavelet transform. Wavelet transforms are widely used for image analysis generally and for analysis of astronomical data in particular (e.g., Starck & Bobin, 2009 and references therein). The reason is that wavelets are efficient in reproducing the kind of hierarchical, fractal-like features that typify astronomical data sets. The experience base in ionospheric imaging is much narrower, but the similarity with astronomy in terms of the targets at issue and the methods in use prompts trials with wavelets for the same reason.

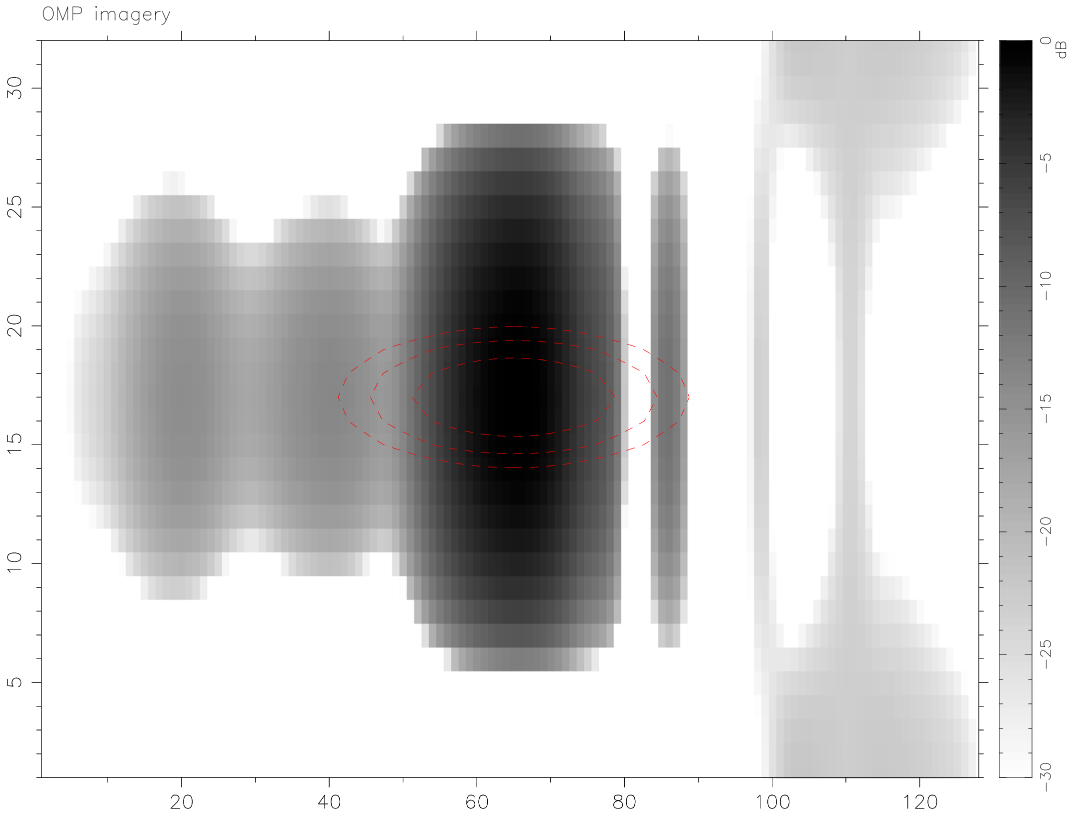

Figure 4. Image recovered from synthetic data using the OMP model. OMP = orthogonal matching pursuit.

Here, we implement the pyramidal scheme of Press et al. (1988). Daubechies d20 wavelets were used, but qualitatively similar results were obtained using d12 and d4 wavelets. This can be attributed to the fact that most of the information in the synthetic image is contained in the low-frequency components and so using larger filters makes little difference.

The tunable parameter for OMP is the number of iterations. The example shown here used 13 iterations, yielding a chi-square value close to the number of data n. Using fewer iterations leads to images composed of simple two-dimensional shapes with a blocky appearance. More iterations lead to a somewhat more elliptically shaped center image and to clutter forms with different and more complicated shapes. In the noiseless case, the additional iterations beyond the expected number of nonzero values increases the probability of signal recovery (Sahoo & Makur, 2015). However, in cases with noise, additional stopping criteria are required to avoid the selection of the zero components. Cai and Wang (2011) (equation (5)) propose a stopping condition to ensure signal recovery in the case of Gaussian noise. In this work, a chi-square value equal to the number of data is used as the stopping condition. This is readily enforced for all the methods considered here and is consistent with Morozov's discrepancy principle for regularization (Morozov, 1966).

The algorithm has done a reasonable job recovering the elliptical target although the stretching in the meridional direction is more severe than with the other methods. However, considerable clutter is evident across the image. Like LCMV and unlike FISTA (and MaxENT as we will see), the results of OMP with Daubechies wavelets are not improved significantly by taking 1-D cuts or averages since the clutter is widespread. Unlike all the other methods, OMP with Daubechies wavelets produces clutter which is not odd symmetric. When the number of iterations is increased, the clutter becomes completely asymmetric. This is particularly problematic since clutter lacking odd symmetry would be more likely to be mistaken for an interesting feature in actual imagery.

Note that qualitatively similar results were obtained using Daubechies wavelets with the FISTA algorithm albeit with much greater computational cost. Using the OMP algorithm without wavelet transforms, meanwhile, produces very unsatisfactory results. The algorithm simply selects and activated a number of pixels

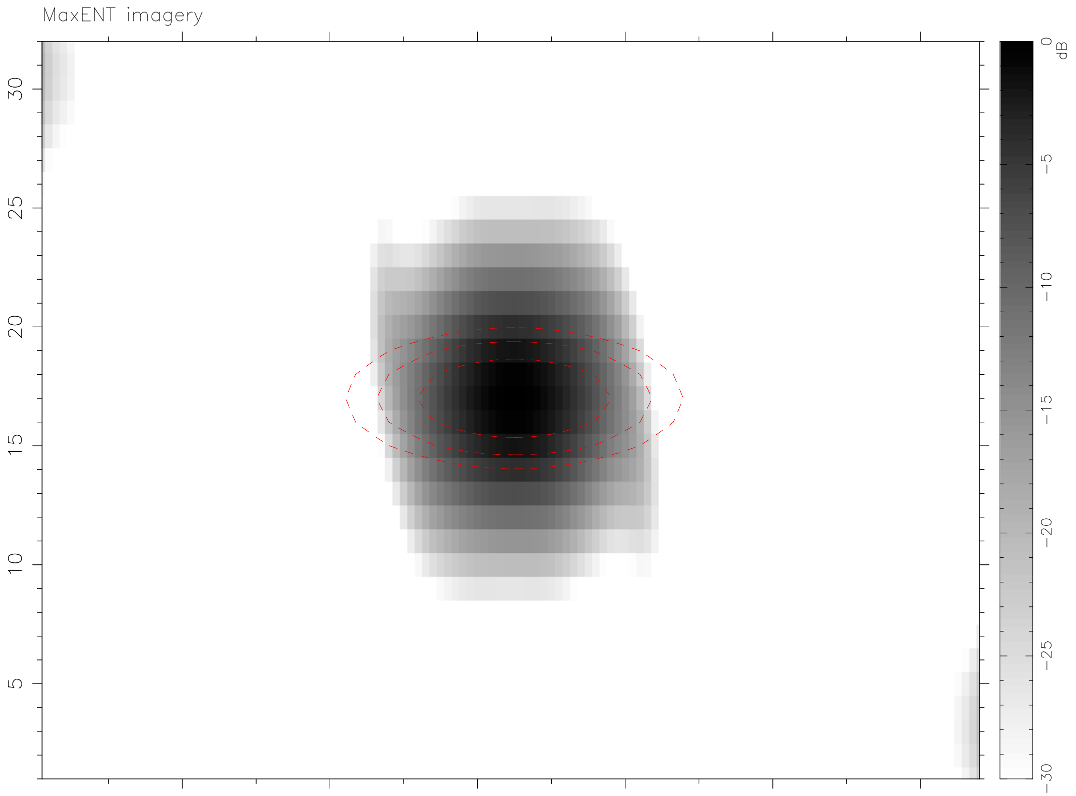

Figure 5. Image recovered from synthetic data using the MaxENT algorithm. MaxENT = entropy maximization..

equal to the number of iterations in that case. Most but not all of the pixels fall within the 30-dB contours in the figure. Given a number of iterations equal to the number of data, the residual is identically zero, but the recovery of the truth image is poor. OMP with the Dirac (or identity) basis evidently does not perform well with distributed targets. It is known that BPDN requires fewer measurements than OMP to recover a signal with the same probability. Whereas OMP needs of the order of $k\ln(m)$ measurements, BPDN needs $k\ln(m/k)$ measurements (Tropp & Gilbert, 2007). This can explain why FISTA works without transforming the domain while OMP does not. The synthetic image is not sparse enough for the original basis to apply to OMP, that is, the MIC is not met.

It is noteworthy also that the nonnegativity of the image solution is not enforced when wavelet bases are incorporated in the manner described here. For references on preserving the nonnegativity condition, see Yaghoobi et al. (2015) and Nguyen et al. (2017).

Finally, Figure 5 shows the results of image reconstruction using the MaxENT algorithm, following the prescription of Hysell and Chau (2006). The MaxENT algorithm has a tuning parameter which is the design value for the chi-square parameter. Here and in actual imaging experiments, we again force chi-square to equal the number of data n.

Like FISTA, the MaxENT algorithm recovers an elliptical target with approximately the same width in the zonal direction as the truth model but with more than twice the width in the meridional direction. Overall, the algorithm recovers the dynamic range of the truth model with a minimum of clutter. A small amount of clutter associated with the interferometry sidelobes is visible in the upper-left and lower-right corners of the image.

A quantitative comparison of the four imaging methods is made in Table 1. For each of the methods, the relative computational cost is shown. This number is the computation time in milliseconds on a single i7 960 CPU core with a 1733-MHz CPU speed. The four methods were implemented in C and compiled with gcc with full optimization. The implementations make use of LAPACK for linear algebra operations.

Table 1
Comparison of Aperture Synthesis Imaging Methods

Method	Time (ms)	RMS error (dB)	1-D cut (dB)
LCMV	20	12.73	9.99
FISTA	1,920	5.16	3.82
OMP	10	9.83	8.11
MaxENT	1,090	4.66	1.86

Note. The execution time is measured in milliseconds, and the RMS error is measured in decibels. The two RMS error figures refer to the entire 2-D image and to a 1-D cut through the horizontal bisector through the image, respectively (see text). RMS = root-mean-square; LCMV = linearly constrained minimum variance; FISTA = fast iterative shrinkage thresholding algorithm; OMP = orthogonal matching pursuit; MaxENT = entropy maximization.

Computation times vary with different algorithm parameters and truth models, and the figures shown are meant to be only representative.

Table 1 also shows the RMS discrepancy between the recovered image and the truth model in decibels. Assessing penalties in terms of decibel quantities prevents the overemphasis of just the strongest portions of the targets, a common practice that undervalues the importance of dynamic range in image construction.

The RMS figure of merit was computed by first thresholding the recovered image and the truth model and then summing the squares of the differences between the model and the image in decibels across pixels. Thresholding means setting a −30-dB floor for both the model and the image prior to computing the metric. The rationale for thresholding is to deemphasize discrepancies in regions of the image that are too weak to be meaningful. The two RMS error terms for each method refer to the result for the entire image and to a 1-D cut through the horizontal bisector through the image, respectively. The latter is the better metric for applications involving radar backscatter from field-aligned irregularities.

Most of the RMS error in the tests is due to spurious clutter, although FISTA and OMP also suffer from underpredicting brightness at the periphery of the radar target. Table 1 suggests that the imaging methods fall into two categories: slow methods that resist clutter and fast methods that are prone to clutter. In the case of the slow methods, most of the clutter in the tests is removed by considering horizontal cuts or averages through the 2-D images. This is not true for the fast methods where the clutter is more widespread in the 2-D images. Overall, neither FISTA nor OMP appear to offer improved performance over the more conventional methods, LCMV, and MaxENT, at least as they have been formulated here.

We point out, however, that our formulation of MaxENT is actually similar to OMP in some respects and so may also be considered a form of compressive sensing in that regard. OMP attempts to form an image from a superposition of basis vectors (atoms) drawn from the observation matrix such as to restrain the norm of the residual in a least squares sense. No more than n vectors may be involved. MaxENT, meanwhile, attempts to reconstruct the logarithm of the image from exactly n basis vectors also drawn from the observation matrix. The logarithm relationship derives from the form of Shannon's entropy. To reproduce a Gaussian ellipsoid, MaxENT needs therefore only construct a paraboloid from the available basis vectors. This helps to explain its comparative success in the tests posed in this study.

We conclude this section with an analysis of the performance of the various radar imaging methods under different levels of statistical uncertainty. Radar backscatter from soft targets is stochastic, and uncertainty is associated with visibility estimates based on finite time averages. The uncertainty is a function of the number of statistically independent samples, the signal-to-noise ratio, and the data themselves (Hysell & Chau, 2006). The associated fluctuations in the experimental visibility estimates contribute to imaging errors.

Regularization in imaging methods is used to limit the growth of fluctuations as they propagate from visibility estimates to brightness estimates. The penalty for regularization is bias, and the balance between fluctuations and bias is exemplified in Figure 6 which shows the 1-D RMS error parameter as a function of the relative standard deviation of the visibility estimates (standard deviation relative to mean value). The results for the four imaging methods considered in this paper are shown. In each case, bias dominates fluctuations at small values of σ/μ where the error parameter tends to reach a floor asymptotically. The situation reverses at large values where the error parameter is an increasing function of σ/μ.

The transition from bias to fluctuation is different for the four methods considered. The LCMV method shows performance comparable to FISTA and MaxENT for small values of σ/μ, but performance degrades immediately as σ/μ increases. LCMV includes no explicit regularization and so is nowhere bias limited. It functions best when fluctuations in the data are small, in the large signal-to-noise ratio limit, for example, but is prone to instability when fluctuations are large and the data covariance matrix becomes poorly conditioned. The results become erratic for large σ/μ.

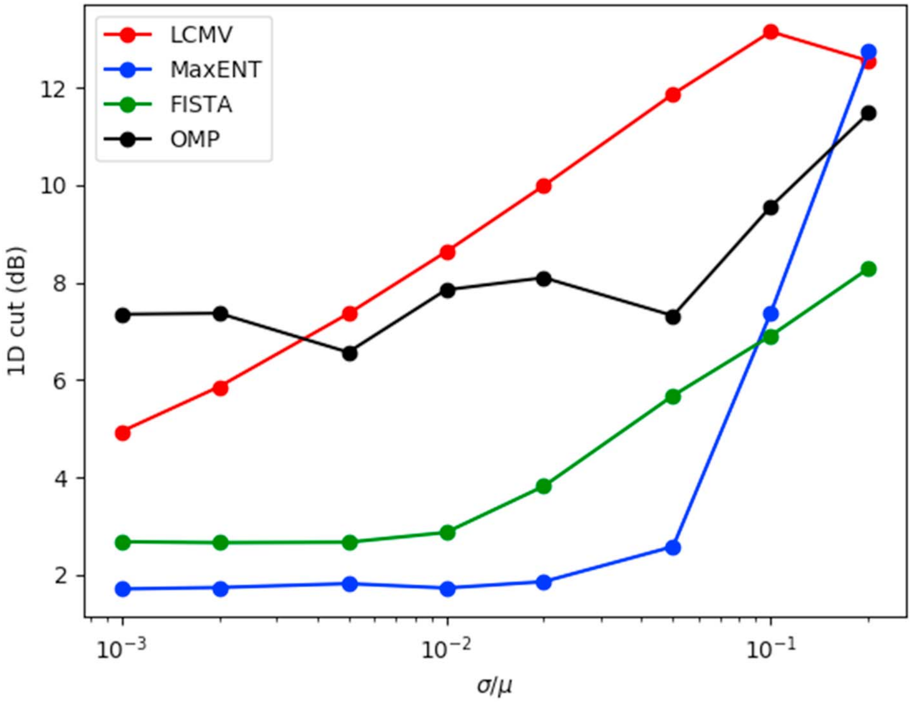

Figure 6. Root-mean-square errors for 1-D cuts versus the relative standard deviation of visibility measurements for four aperture synthesis imaging methods. LCMV = linearly constrained minimum variance; MaxENT = entropy maximization; FISTA = fast iterative shrinkage thresholding algorithm; OMP = orthogonal matching pursuit.

OMP with a wavelet basis, meanwhile, appears to be bias limited up through large values of σ/μ. This prevents the method from capitalizing on high-quality data that might be available, for example, in the high signal-to-noise ratio limit. The preference for sparse solutions appears to represent a strong bias that conceals image features even when they have support in the data. MaxENT and FISTA are superior in this regard in that their performance benefits from values of σ/μ as small as about 1%. Even smaller values would be difficult to achieve in practical experimental applications.

6. Summary

This paper sets out to compare and contrast four methods for inverting spaced-receiver radar observations of field-aligned plasma density irregularities in the ionosphere. Measured visibility data are related to desired radar images through a linear transformation similar to a Fourier transform, but the sparse and incomplete sampling of the former necessitates the use of inverse methods in forming the latter. Two new methods, BPDN (using FISTA) and OMP, were considered along with two conventional methods based in MaxENT and adaptive beamforming (LCMV). Realistic synthetic data as might be acquired by the imaging array used at the Jicamarca Radio Observatory were used for the study.

Choosing the right metric evaluation is important. The goal of aperture-synthesis imaging, an underdetermined problem, is not the minimization of the norm of the residual. For the MaxENT, FISTA, and OMP runs conducted here, the chi-square parameter was made to be equal to the number of data elements n in each case. The goal instead is the accurate recovery of the $m \ggg n$ elements of the truth model. However, the RMS discrepancy between the truth model and the recovered image would not be a very appropriate metric given the importance of dynamic range in radar imaging. All of the methods considered here would have scored well by this metric since they were all able to reproduce the strongest parts of the truth model. We instead considered the RMS discrepancy in the logarithms of the truth model and the recovered image. A threshold was set at −30 dB in the calculation to avoid overemphasizing discrepancies in parts of the images that were negligibly small, however. By this metric, MaxENT outperformed the other methods. BPDN performed adequately but did not reproduce the desired gradation at the boundary of the truth model target. Neither LCMV nor OMP provided adequate clutter suppression outside the boundary of the target.

We note again that none of the methods tested here were able to resolve the synthetic radar target in the direction of the magnetic meridian. Radar backscatter from field-aligned irregularities occurs within a very narrow range of meridional angles, and we do not expect or even attempt to measure magnetic aspect width using imaging in practice. The important metric is the ability to recover the shape of the backscatter in the appropriate 1-D cut through the 2-D imagery. There are advantages in formulating the imaging problem in 2-D, however, including the ability to measure the direction of the centroid of the backscatter.

Moreover, overestimating the width of targets in the meridional direction could pose a problem to the extent that it makes the method prone to underestimating the width in the zonal direction, a feature exhibited to some extent by all of the test cases considered here. Target widths in the two directions are coupled since the moments are approximately conserved. This observation may warrant a redesign of the imaging array used routinely at Jicamarca which presently does not include long meridional baselines.

Methods rooted in BPDN and OMP may yet be able to perform well in the aperture-synthesis imaging problem. The impressive speed of OMP warrants an examination of different basis functions which may be more suitable for extended 2-D radar targets. One intriguing possibility is the use of curvelet transforms which have proven expedient in applications similar to this one (Ma & Plonka, 2010).

Acknowledgments
This research was supported by awards AGS-1818216 from the National Science Foundation and DE-AR0000946 from ARPA-E to Cornell University. Support was also received from the Deutsche Forschunggemeinschaft (DFG, German Research Foundation) under SPP 1788 (CoSIP)-CH1482/3-1. The Jicamarca Radio Observatory is a facility of the Instituto Geofisíco del Perú operated with support from NSF award AGS-1732209 through Cornell.

References

Beck, A., & Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, *2*, 183–202.

Cai, T., & Wang, L. (2011). Orthogonal matching pursuit for sparse signal recovery with noise. *IEEE Transactions on Information Theory*, *57*, 4680–4688.

Cai, T. T., Wang, L., & Xu, G. (2010). Stable recovery of sparse signals and an oracle inequality. *IEEE Transactions on Information Theory*, *56*, 3516–3522. https://doi.org/10.1109/TIT.2010.2048,506

Candés, E. J., Romberg, K., & Tao, T. (2006). Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, *52*, 489–509.

Candés, E. J., & Tao, T. (2006). Near-optimal signal recovery from random projections: Universal encoding strategies. *IEEE Transactions on Information Theory*, *52*, 5406–5425.

Capon, J. (1969). High-resolution frequency-wavenumber spectrum analysis. *Proceedings IEEE*, *57*, 1408.

Donoho, D. (2006). Compressed sensing. *IEEE Transactions on Information Theory*, *52*, 1289–1306. https://doi.org/10.1109/TIT.2006.871, 582

Donoho, D. L., & Huo, X. (2001). Uncertainty principles and ideal atomic decomposition. *IEEE Transactions on Information Theory*, *47*, 2845–2862.

Farley, D. T., Ierkic, H. M., & Fejer, B. G. (1981). Radar interferometry: A new technique for studying plasma turbulence in the ionosphere. *Journal of Geophysical Research*, *86*, 1467–1472.

Harding, B. J., & Milla, M. A. (2013). Radar imaging with compressed sensing. *Radio Science*, *48*, 582–588. https://doi.org/10.1002/rds.20, 063

Högbom, J. A. (1974). Aperture synthesis with a non-regular distribution of interferometer baselines. *Astronomy and Astrophysics Supplement*, *15*, 417–426.

Hysell, D. L. (1996). Radar imaging of equatorial F region irregularities with maximum entropy interferometry. *Radio Science*, *31*, 1567.

Hysell, D. L., & Chau, J. L. (2006). Optimal aperture synthesis radar imaging. *Radio Science*, *41*, RS2003. https://doi.org/10.1029/2005RS003, 383

Hysell, D. L., Larsen, M. F., & Zhou, Q. H. (2004). Common volume coherent and incoherent scatter radar observations of mid-latitude sporadic E-layers and QP echoes. *Annales Geophysicae*, *22*, 3277–3290.

Hysell, D. L., Michhue, G., Larsen, M. F., Pfaff, R., Nicolls, M., Heinselman, C., & Bahcivan, H. (2008). Imaging radar observations of Farley Buneman waves during the JOULE II experiment. *Annales Geophysicae*, *26*, 1837–1850.

Hysell, D. L., Yamamoto, M., & Fukao, S. (2002). Imaging radar observations and theory of type I and type II quasi-periodic echoes. *Journal of Geophysical Research*, *107*(A11), 1360.

Kallummil, S., & Kalyani, S. (2017). Tuning free orthogonal matching pursuit. arXiv:1703.05080v1.

Kudeki, E., & Sürücü, F. (1991). Radar interferometric imaging of field-aligned plasma irregularities in the equatorial electrojet. *Geophysical Research Letters*, *18*, 41–44.

Ma, J., & Plonka, G. (2010). The curvelet transform. *IEEE Signal Processing Magazine*, *2*, 118–133.

Mackenzie, D. (2009). *What's happening in the mathematical sciences*, vol. 7. Providence, RI: American Mathematical Society.

Morozov, V. A. (1966). On the solution of functional equations by the method of regularization. *Soviet mathematics - Doklady*, *7*, 414–417.

Nesterov, Y. E. (1983). A method for solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet mathematics - Doklady*, *27*, 372–376.

Nguyen, T. T., Soussen, C., Idier, J., & Djermoune, E.-H. (2017). An optimized version of non-negative OMP. XXVIe Colloque GRETSI Traitement du Signal Images, GRETSI 2017 Juan-les-Pins, France, Sep.

Peng, Z., Yan, M., & Yin, W. (2013). Parallel and distributed sparse optimization. In *Paper presented at the 2013 Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, USA, 3–6 Nov.*

Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1988). *Numerical recipes in C*. New York: Cambridge Univ Press.

Rani, M., Dhok, S. B., & Deshmukh, R. B. (2018). A systematic review of compressive sensing: Concepts, implications, and applications. *IEEE Access*, *6*, 4875–4894. https://doi.org/10.1109/ACCESS.2018/2793,851

Sahoo, S. K., & Makur, A. (2015). Signal recovery from random measurements via extended orthogonal matching pursuit. *IEEE Transactions on Signal Processing*, *63*, 2572–2581. https://doi.org/10.1109/TSP.2015.2413384

Saito, S., Yamamoto, M., & Hashiguchi, H. (2008). Imaging observations of nighttime mid-latitude F-region field-aligned irregularities by an MU radar ultra-multi-channel system. *Annales Geophysicae*, *26*, 2345–2352.

Saito, S., Yamamoto, M., Hashiguchi, H., & Maegawa, A. (2006). Observation of three-dimensional signatures of quasi-periodic echoes associated with mid-latitude sporadic-E layers by MU radar ultra-multi-channel system. *Geophysical Research Letters*, *33*, L14109. https://doi.org/10.1029/2005GL025,526

Skilling, J., & Bryan, R. K. (1984). Maximum entropy image reconstruction: General algorithm. *Monthly Notices of the Royal Astronomical Society*, *211*, 111–124.

Sommer, S., & Chau, J. L. (2016). Patches of polar mesospheric summer echoes characterized from radar imaging observations with MAARSY. *Annales Geophysicae*, *34*, 1231–1241.

Starck, J.-L., & Bobin, J. (2009). Astronomical data analysis and sparsity: From wavelets to compressed sensing. *Proceedings of the IEEE*, *98*, 1021–1030.

Thompson, A. R. (1986). *Interferometry and synthesis in radio astronomy*. New York: John Wiley.

Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society*, *58*, 267–288.

Tropp, J. (2004). Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, *50*, 2231–2242.

Tropp, J. (2006). Just relax: Convex programming methods for identifying sparse signals in noise. *IEEE Transactions on Information Theory*, *52*, 1030–1051.

Tropp, J., & Gilbert, C. (2007). Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, *53*, 4655–4666.

Urco, J. M., Chau, J. L., Milla, M. A., Vierinen, J. P., & Weber, T. (2018). Coherent MIMO to improve aperture synthesis radar imaging of field-aligned irregularities: First results at Jicamarca. *IEEE Transactions on Geoscience and Remote Sensing*, *56*, 2980–2990. https://doi.org/10.1109/TGRS.2017.2788,425

Urco, J. M., Chau, J. L., Weber, T., & Lateck, R. (2018). Enhancing the spatio-temporal features of polar mesosphere summer echoes using coherent MIMO and radar imaging at MAARSY, *12*, 955–969. https://doi.org/10.5194/amt--2018--258

Wilczek, R., & Drapatz, S. (1985). A high accuracy algorithm for maximum entropy image restoration in the case of small data sets. *Astronomy and Astrophysics*, *142*, 9–12.

Yaghoobi, M., Wu, D., & Davies, M. E. (2015). Fast non-negative orthogonal matching pursuit. *IEEE Signal Processing Letters*, *22*, 1229–1233. https://doi.org/10.1109/LSP.2015.2393,637

Yang, A. Y., Zhou, Z., Balasubramanian, A. G., Sastry, S. S., & Ma, Y. (2013). Fast l_1-minimization algorithms for robust face recognition. *IEEE Transactions on Image Processing*, *22*, 3234–3246. https://doi.org/10.1109/TIP.2013.2262,292

Yu, T. Y., Palmer, R. D., & Hysell, D. L. (2000). A simulation study of coherent radar imaging. *Radio Science*, *35*, 1129.