2018 5th International Workshop on Compressed Sensing applied to Radar, Multimodal Sensing, and Imaging (CoSeRa)

Expansion of Dropped-Channel PolSAR CS to include a Spatial Dictionary

John Becker, Julie Ann Jackson
Dept. of Electrical and Computer Engineering
Air Force Institute of Technology
Dayton, Ohio 45433
Email: john.becker@afit.edu, julie.jackson@afit.edu

***Abstract*—Previous work introduced dropped-channel polarimetric synthetic aperture radar (PolSAR) compressive sensing assuming a point scatterer spatial response. In this paper, we expand that work to include any spatial dictionary. Of particular interest is a dictionary that can model both localized and extended scatterers to obtain a sparse representation of the scene. We simulate results for different spatial and polarization dictionaries. Dictionaries well-matched to the targets result in a sparser coefficient vector, as measured by the ℓ_1-norm.**

I. Introduction

Previous work [1], [2] suggests that polarimetric compressive sensing (CS) may utilize channel crosstalk to measure a reduced set of polarization channels and still recover full-pol imagery. In this paper, we build upon [1], [2] by also incorporating a spatial dictionary to obtain an even sparser solution. The previous spatial dictionary used in [1], [2] assumes single point scatterers, which leads to multiple coefficients and reduced sparsity when extended scatterers are present in the scene. A spatial dictionary that includes scattering phenomenology expected in a scene will result in a much sparser representation, with only one or a few coefficients per scatterer. Combined with a polarization dictionary, the spatial dictionary will provide target feature information that can be used for scene interpretation.

Recently, spatial dictionaries have been used to help with acoustic imaging [3], [4], hyperspectral imaging [5], adaptive beamforming [6], [7], and radar [8]–[12]. There have also been numerous recent applications of compressive sensing (CS) to radar imaging, ranging from autofocus [13] and scatterer persistence [14] to sub-Nyquist sampling [15] and random sampling [16]. Previously, a method was derived to use CS to reconstruct polarimetric synthetic aperture radar (PolSAR) images from a subset of measured channels [1]. In [1], single-pixel point targets are assumed, so the identity matrix is used as the spatial dictionary. The point scatterer assumption, in general, works for any scene, as any complex scatterer can be broken down into a set of point scatterers under a high frequency assumption. When extended scatterers are present, the point scatterer model leads to a less sparse solution. We prefer to have as few coefficients per scatterer as possible, which requires spatial and polarization dictionaries that are well-matched to the scene.

Our approach to combining CS and SAR is unique in that we compress over the polarization channels themselves, not the fast or slow time samples [1]. Our approach allows us to drop measured channels and reconstruct them from the cross-talk information. By extending our approach with a spatial dictionary for extended scatterers, we can obtain a sparse representation for scenes that contain multiple scatterer types.

This paper is laid out as follows. In Section II, we incorporate a general spatial dictionary into the mathematical framework of [1], [2]. In Section III, we present the results from a simple simulation involving a point target and an extended scatterer. Section IV analyzes recovery success versus scene sparisty. In Section V, we present our conclusions and suggest some further direction for this work.

II. Theory

A polarimetric stack of radar images may be decomposed into a set of scattering coefficients as follows. As in [1], [2] we assume $\mathbf{x}_{m'}$ is the unknown $N' \times 1$ reflectivity in $m' = 1, \ldots, M'$ available channels and $\mathbf{y}_m$ is the $N \times 1$ vectorized image for $m = 1, \ldots, M$ measured channels. The channels are coupled via crosstalk matrix $\mathbf{C}$, and matrix $\mathbf{J}$ indicates which $M \leq M'$channels are measured. (Dropping a channel is the PolSAR compressive sensing thesis of [1], [2].) Furthermore, the reflectivity $\mathbf{x} = [\mathbf{x}_1^T, \ldots, \mathbf{x}_{M'}^T]^T$ can be broken down into a dictionary representation $\mathbf{x} = (\mathbf{P} \otimes \mathbf{D})\mathbf{b}$, for an $M' \times Q$ polarization dictionary $\mathbf{P}$, $N' \times D'$ spatial dictionary $\mathbf{D}$, and $QD' \times 1$ coefficient vector $\mathbf{b}$. The notation $\otimes$ represents a Kronecker product. Then, using the above definitions, the noisy, channel-coupled observed image stack $\mathbf{y} = [\mathbf{y}_1^T, \ldots, \mathbf{y}_M^T]^T$ can be written in terms of the scattering coefficients $\mathbf{b}$ as

$$\mathbf{y} = \mathbf{A}\mathbf{b} + \mathbf{n} \tag{1}$$

where

$$\mathbf{A} = (\mathbf{J}\mathbf{C} \otimes \mathbf{A}_1)(\mathbf{P} \otimes \mathbf{D}) \tag{2}$$

and $\mathbf{n}$ is the $MN' \times 1$ complex additive noise vector. Equation (2) is a modified version of [1, Eq. 5d] to include the arbitrary spatial dictionary $\mathbf{D}$ in the multi-channel image formation operator $\mathbf{A}$. In (2), $\mathbf{A}_1$ is the $N' \times N'$ single channel imaging operator. As in previous work, we assume $\mathbf{A}_1$ is the same for all channels; however, (2) can be rewritten to remove such restriction. In [1], the spatial dictionary was $\mathbf{D} = \mathbf{I}_{N'}$,

an $N' \times N'$ identity matrix representing point scatterers. In this paper, we choose a generic $\mathbf{D}$ to contain both point and extended scatterers. More specific wave-based or canonical shape dictionaries may also be used [8], [11], [12]. Equation (2) is an example of Kronecker compressive sensing, as defined in [17], since the problem separates into a product of dictionaries and measurement matrices, each of which is a Kronecker product of dictionaries or measurement matrices for a specific dimension. In our case, we have polarization and spatial dimensions.

We aim to reconstruct dropped channels and find a sparse representation of the scene. Thus, we have a basis pursuit denoising (BPDN) [18] problem

$$\min_{\mathbf{b}} \|\mathbf{b}\|_1 \quad \text{s.t. } \|\mathbf{y} - (\mathbf{JC} \otimes \mathbf{A}_1)(\mathbf{P} \otimes \mathbf{D})\mathbf{b}\|_2 \leq \epsilon \quad (3)$$

which is similar to the problem in [1] with the replacement of $\mathbf{I}_{N'}$ with the spatial dictionary $\mathbf{D}$. Previously in [1], each coefficient in $\mathbf{b}$ corresponded to the polarimetric response of the target in the dictionary $\mathbf{P}$ for each pixel. Now, $\mathbf{b}$ contains coefficients for the combined spatial and polarimetric dictionaries; non-zero coefficients relate to each scatterer in the scene. Examples in the next section will compare sparsity of reconstructions with different polarimetric and spatial dictionaries. In each case, (3) is solved, resulting in recovery of the missing channel, removal of crosstalk and noise, and sparse representation of the scene.

III. Example Scene Recovery

In this section, we present simulation results from solving Equation (3) using both point scatterer and combined point and extended scatterer spatial dictionaries as well as two different polarization dictionaries. We then compare the ℓ_1-norm over the estimated scattering coefficients vector $\hat{\mathbf{b}}$ and ℓ_2 residual error of $\mathbf{b}$ and $\mathbf{x}$ for each scenario. For each test, we do not measure the HH channel directly and instead, reconstruct it from crosstalk information as in [1], [2], using the BPDN algorithm from the SPGL-1 package [19], [20].

The example scene is 5×6 pixels in size and consists of two scatterers: a single pixel point target and a three-pixel wide extended scatterer oriented in the cross-range direction. The target locations are shown in Figure 2. We assume that the extended scatterer is completely contained in the image; thus, there are 18 possible scatterer locations included in the extended target dictionary $\mathbf{D}_{ext}$. Likewise, there are 30 possible single-pixel target locations in the point scatterer dictionary $\mathbf{D}_{pt} = I_{30}$. Combining the two dictionaries into $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$, yields a full spatial dictionary of size 30×48. The full spatial dictionary structure can be seen in Figure 1, with each part labeled and highlighted for clarity. There are three stripes in $\mathbf{D}_{ext}$, corresponding to the three points in the extended scatterer. The dictionary $\mathbf{D}$ captures the spatial extent of the scatterer response in any channel. The scatterer's polarization response is fit to the polarization dictionary $\mathbf{P}$.

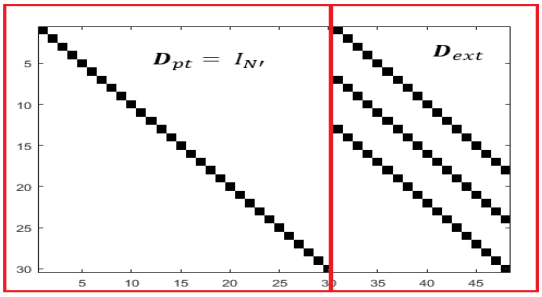

Fig. 1. Structure of the spatial dictionary $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$.

We compare two different polarization dictionaries. The first is the same canonical basis as used in [1]:

$$\mathbf{P}_{can} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix} \quad (4)$$

which corresponds to the first three columns of a Pauli basis. The second polarization dictionary is structured as

$$\mathbf{P}_{new} = \begin{bmatrix} \frac{1}{2} & \frac{1}{\sqrt{2}} \\ \frac{1}{2} & 0 \\ \frac{1}{2} & 0 \\ \frac{1}{2} & -\frac{1}{\sqrt{2}} \end{bmatrix}. \quad (5)$$

Our examples assume that the point target at $(-1, -1)$ has polarization response corresponding the first column of $\mathbf{P}_{can}$ and the extended target has polarization response corresponding to the first column of $\mathbf{P}_{new}$.

Figure 2 shows the true reflectivity of the scene in each polarization. The true reflectivity is the product of the true $\mathbf{b}$ vector with matrix $\mathbf{P} \otimes \mathbf{D}$. Observed images are formed via (1) and (2) with the crosstalk matrix the same as used in [2]. The complex crosstalk matrix picks magnitude from uniform [0,1] and phase from uniform [0, 2π], and sets the diagonal to ones. Matrix $\mathbf{J}$ corresponds to a 4×4 identity matrix with the first row omitted in order to drop the HH channel from the measurements. Noise is added at a 20 dB signal-to-noise ratio (SNR), yielding the PFA observations shown in Figure 3. The crosstalk artifacts displayed in the cross-pol images in Figure 3 enable recovery of the dropped HH channel.

Figures 4-7 show the recovered full-pol reflectivity for different spatial and polarization dictionary combinations. The four subfigures in each figure show the estimated reflectivity $\hat{\mathbf{x}} = (\mathbf{P} \otimes \mathbf{D})\hat{\mathbf{b}}$ obtained from the BPDN solution $\hat{\mathbf{b}}$. The recovered images in each case are a good representation of the target true reflectivity shown in Figure 2.

Figures 4 and 5 use the canonical polarization dictionary $\mathbf{P} = \mathbf{P}_{can}$. Figure 4 uses the point scatterer spatial dictionary $\mathbf{D} = \mathbf{D}_{pt} = \mathbf{I}_{N'}$, while Figure 5 uses the combined point and extended scatterer dictionary $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$. Figures 6 and 7 replace the canonical dictionary with $\mathbf{P} = \mathbf{P}_{new}$. It is difficult to tell by comparing Figures 4-7 the difference that was made by the change of spatial and/or polarization dictionaries. This is because the primary difference is in the sparsity

of the representation of the two scatterers. Sparsity is best shown by the difference in the ℓ_1-norm of the BPDN-estimated vector of scattering coefficients $\hat{\mathbf{b}}$. Table I shows both the ℓ_1-norm of the estimated coefficient vector and the ℓ_2-norm of the difference between true and estimated scene reflectivity. Comparing the $\mathbf{D} = \mathbf{I}_{N'}$ and $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$ cases, we can see that the spatial dictionary $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$ leads to a substantially lower ℓ_1-norm on $\hat{\mathbf{b}}$ despite similar images and ℓ_2 errors. Thus, the representation in Figure 5 is more sparse than the representation in Figure 4. Likewise, $\hat{\mathbf{b}}$ for Figure 7 is more sparse than for Figure 6. The lower ℓ_1-norm highlights the benefit of the extended scatterer spatial dictionary. By choosing a dictionary that anticipates an extended scatterer, we are able to take advantage of the spatial sparsity of the target, representing the extended scatterer as one coefficient instead of three. Likewise, by choosing a polarization dictionary that anticipates the polarimetric sparsity of the point target, we are able to further reduce the ℓ_1-norm by requiring even fewer coefficients to represent the scatterers. Specifically, if we compare row 2 and row 4 in Table I, we see that the ℓ_1-norm is smaller when the $\mathbf{P}_{new}$ dictionary is used. Since the extended scatterer can be represented by only column one of $\mathbf{P}_{new}$, as opposed to a combination of columns one and three of $\mathbf{P}_{can}$, we are able to achieve an even sparser representation of $\hat{\mathbf{b}}$. A lower ℓ_1-norm value is indicative of a more sparse representation of the scene, which is beneficial for efficient storage and communication of the scene information. On the other hand, the polarization response of the point target corresponds to the first column of $\mathbf{P}_{can}$, which cannot be represented by a combination of columns of $\mathbf{P}_{new}$, resulting in increased error on that target. Thus, it is important to choose dictionaries which are well-matched to the scene scattering phenomenology to obtain both good ℓ_1 and ℓ_2 fit.

TABLE I
COMPARISON OF ℓ_1-NORM AND ℓ_2 ERROR FOR THE RECOVERY METHODS

Combination	ℓ_1-Norm : $\hat{\mathbf{b}}$	ℓ_2 Error : $\mathbf{x}$	ℓ_2 Error : $\mathbf{b}$
$\mathbf{D} = \mathbf{D}_{pt}$, $\mathbf{P} = \mathbf{P}_{can}$	6.66026	0.51231	0.51231
$\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$, $\mathbf{P} = \mathbf{P}_{can}$	2.36957	0.56580	0.39903
$\mathbf{D} = \mathbf{D}_{pt}$, $\mathbf{P} = \mathbf{P}_{new}$	5.5679	0.28313	0.28313
$\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$, $\mathbf{P} = \mathbf{P}_{new}$	1.86978	0.48146	0.55104

IV. SPARSE RECOVERY ANALYSIS

Another import metric in testing our approach is the average error rate at each level of sparsity. To measure the error, we run a Monte Carlo simulation for each combination of spatial dictionary and polarization dictionary. The Monte Carlo simulation is run over the levels of sparsity from 5% to 60% at 5% increments. At each level, a series of random $\mathbf{b}$ vectors are generated, turned into observations via Equations (1) and (2) and run through SPGL-1. The ℓ_2-error is then generated for each result and the average is computed. The average ℓ_2-error versus number of non-zero components is shown in Figure 8 for each combination of spatial and polarization dictionary. Specifically, Figures 8(a) and 8(b) show the ℓ_2-error of the

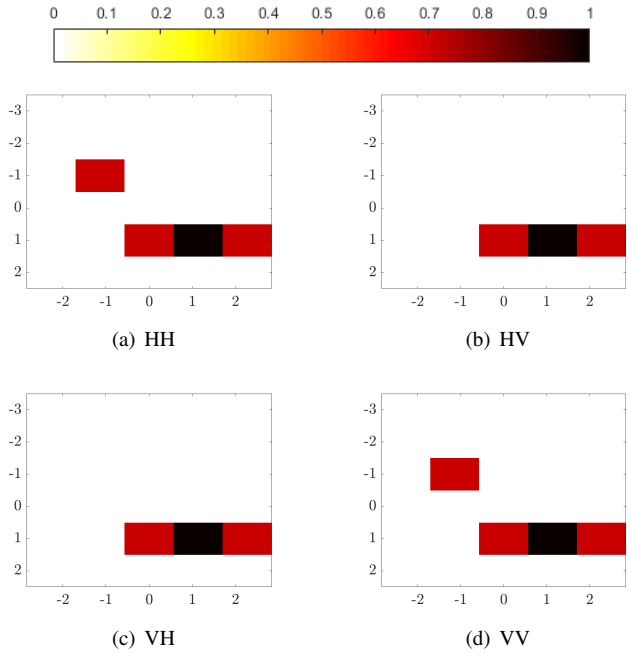

(a) HH (b) HV (c) VH (d) VV

Fig. 2. True reflectivity of $\mathbf{x}$

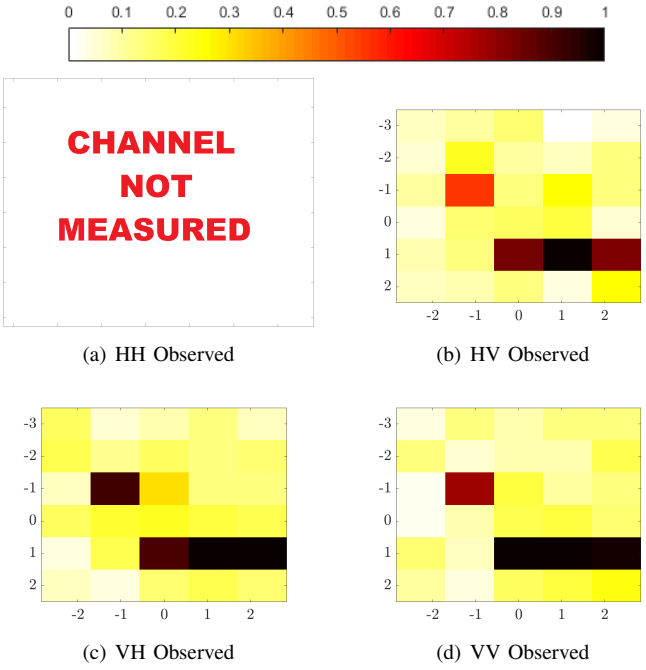

(a) HH Observed (b) HV Observed (c) VH Observed (d) VV Observed

Fig. 3. PFA Observations of Scene

$\mathbf{b}$ vectors for both polarization dictionaries when $\mathbf{D} = \mathbf{I}_{N'}$ and $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$, respectively. Figure 8(c) shows the ℓ_2-error of the $\mathbf{x}$ vectors for all four combinations of polarization and spatial dictionaries. Due to deconfliction of overlapping scatterers in the $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$ case, the sweep does not make it to 60%, the error is plotted at the average percentage achieved instead.

Figure 8 shows that the error of $\hat{\mathbf{b}}$ and $\hat{\mathbf{x}}$ for each combination generally increases as the sparsity decreases. Initial

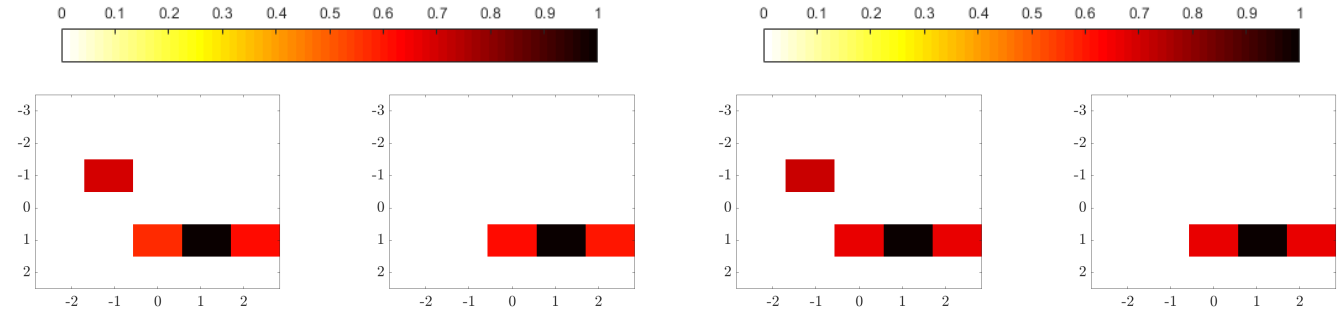

(a) HH BPDN Recovery (b) HV BPDN Recovery

(c) VH BPDN Recovery (d) VV BPDN Recovery

Fig. 4. BPDN Recovery of Scene with $\mathbf{P} = \mathbf{P}_{can}$, $\mathbf{D} = \mathbf{I}_{N'}$.

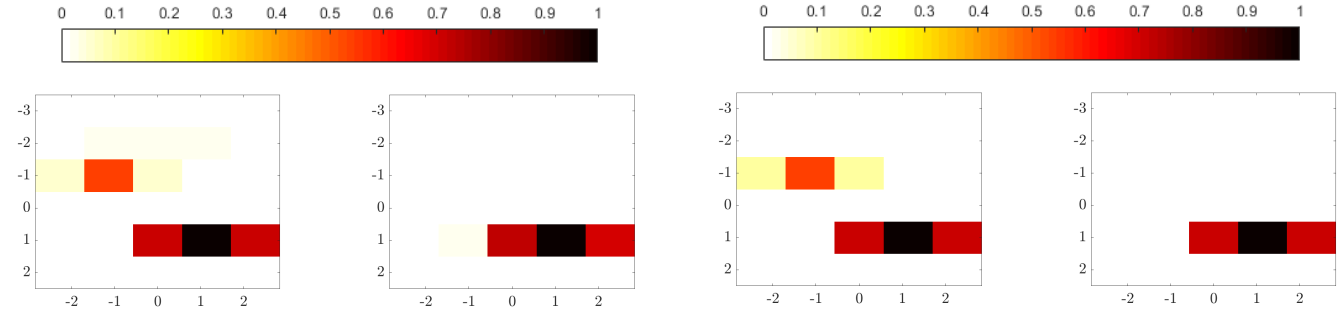

(a) HH BPDN Recovery (b) HV BPDN Recovery

(c) VH BPDN Recovery (d) VV BPDN Recovery

Fig. 5. BPDN Recovery of Scene, $\mathbf{P} = \mathbf{P}_{can}$, $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$

(a) HH BPDN Recovery (b) HV BPDN Recovery

(c) VH BPDN Recovery (d) VV BPDN Recovery

Fig. 6. BPDN Recovery of Scene, $\mathbf{P} = \mathbf{P}_{new}$, $\mathbf{D} = \mathbf{I}_{N'}$

(a) HH BPDN Recovery (b) HV BPDN Recovery

(c) VH BPDN Recovery (d) VV BPDN Recovery

Fig. 7. BPDN Recovery of Scene, $\mathbf{P} = \mathbf{P}_{new}$, $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$

results indicating that the error is much larger for the spatial dictionary $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$ may be attributable to the sparsity constraints in the BPDN algorithm. It is possible that BPDN is selecting coefficients out of the $\mathbf{D}_{ext}$ part of the spatial dictionary to represent $\hat{\mathbf{x}}$ when the true $\mathbf{x}$ is really a combination of point targets as opposed to an extended scatterer. This would yield a sparser representation of the scene while being inconsistent with the true $\mathbf{b}$ vector coefficients. Figure 8(c) also points towards this hypothesis as the errors in the $\hat{\mathbf{x}}$ are similar across all combinations of spatial and polarization dictionaries. This indicates that, despite the higher error rates in $\hat{\mathbf{b}}$ when $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$, $\hat{\mathbf{x}}$ is a close approximation of the true $\mathbf{x}$ vector. Thus, dictionary choice is an important factor for achieving both good signal recovery and sparse representation.

V. Conclusion

We have extended the work in [1], [2] to include a spatial dictionary. Our approach can extend to any spatial dictionary

but we only offer one example here. The spatial dictionary was shown to provide a more sparse representation of a scene that contains extended scatterers as compared to the point scatterer dictionary used in [1]. The increased sparsity is due to the spatial dictionary being well-matched to the scene, giving us coefficients on a per-scatterer level as opposed to a per-pixel level like in [1]. We also compared two different polarization dictionaries to explore the effects on scene error and sparsity. By achieving a sparser representation of the data through a well-matched spatial dictionary, we can store and communicate the scene information more efficiently.

An important note is that we did not change the crosstalk matrix from [2]. Since the crosstalk matrix is currently designed to give the lowest mutual coherence of $\mathbf{A}$ for an identity matrix as the spatial dictionary; it is sub-optimal for use with our spatial dictionary $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$, or any other for that matter. A more optimal crosstalk matrix could be designed for each choice of $\mathbf{D}$ and $\mathbf{P}$ as was done in [1], [2]. However, even with a sub-optimal crosstalk matrix, we are able to achieve an accurate, sparse representation of the scene. In the future, crosstalk matrix design should be considered more carefully and with the spatial dictionary in mind.

VI. Acknowledgment

The views expressed in this article are those of the authors and do not reflect the official policy of position of the U.S. Air Force, Department of Defense, or the U.S. Government. Approved for public release #88ABW-2018-3612.

References

[1] J. A. Jackson and F. Lee-Elkin, "Channel crosstalk model for fully-polarimetric SAR compressive sensing," *IEEE Radar Conf.*, pp. 1536–1541, May 2017.

[2] ——, "Polarimetric SAR compressive sensing examples," *IEEE Radar Conf.*, Apr. 2018.

[3] M. Samarawickrama, N. Epain, and C. Jin, "Super-resolution acoustic imaging using non-uniform spatial dictionaries," in *Int. Conference on Audio, Language and Image Processing*, Jul. 2014, pp. 973–977.

[4] F. Nesta and M. Fakhry, "Unsupervised spatial dictionary learning for sparse underdetermined multichannel source separation," in *IEEE ICASSP*, May 2013, pp. 86–90.

[5] W. Huang, Z. Wu, H. Liu, L. Xiao, and Z. Wei, "Spatial-spectral compressive sensing for hyperspectral images super-resolution overlearned dictionary," in *IEEE IGARSS*, Jul. 2014, pp. 4930–4933.

[6] N. Ito, S. Araki, M. Delcroix, and T. Nakatani, "Probabilistic spatial dictionary based online adaptive beamforming for meeting recognition in noisy and reverberant environments," in *IEEE ICASSP*, Mar. 2017, pp. 681–685.

[7] N. Ito, S. Araki, and T. Nakatani, "Data-driven and physical model-based designs of probabilistic spatial dictionary for online meeting diarization and adaptive beamforming," in *25th European Signal Processing Conference*, Aug. 2017, pp. 1165–1169.

[8] M. McClure and L. Carin, "Matching Pursuits with a Wave-Based Dictionary," *IEEE Trans. on Signal Processing*, vol. 45, no. 12, pp. 2912–2927, Dec. 1997.

[9] P. Bharadwaj, P. Runkle, and L. Carin, "Target identification with wave-based matched pursuits and hidden Markov models," *IEEE Trans. on Antennas and Propagation*, vol. 47, no. 10, pp. 1543–1554, Oct. 1999.

[10] R. Mazhar, J. N. Wilson, and P. D. Gader, "Use of an application-specific dictionary for matching pursuits discrimination of landmines and clutter," in *IEEE IGARSS*, Jul. 2007, pp. 26–29.

[11] K. R. Varshney, M. Cetin, J. W. Fisher, and A. S. Willsky, "Sparse representation in structured dictionaries with application to synthetic aperture radar," *IEEE Trans. on Signal Processing*, vol. 56, no. 8, pp. 3548–3561, Aug 2008.

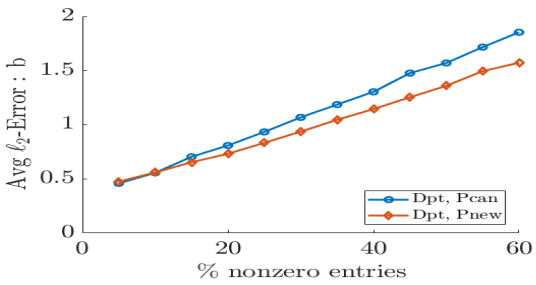

(a) $\mathbf{D} = \mathbf{D}_{pt}$

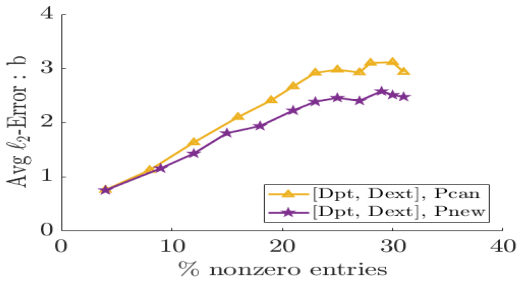

(b) $\mathbf{D} = [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$

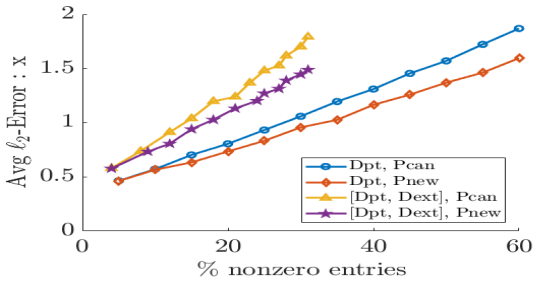

(c) $\mathbf{D} = \mathbf{D}_{pt}, [\mathbf{D}_{pt}, \mathbf{D}_{ext}]$

Fig. 8. Average ℓ_2-error rates as a function of sparsity.

[12] G. B. Hammond and J. A. Jackson, "SAR canonical feature extraction using molecule dictionaries," in *IEEE Radar Conf.*, April 2013, pp. 1–6.

[13] T. Wang, B. Liu, X. Ling, and Y. Liu, "A novel autofocus imaging method for ISAR based on compressive sensing," in *IEEE ICSPCC*, Oct. 2017, pp. 1–5.

[14] N. Sigavanam and E. Ertin, "Interrupted SAR imaging with limited persistence scattering models," in *IEEE Radar Conf.*, May 2017, pp. 1770–1775.

[15] K. Aberman and Y. C. Eldar, "Sub-Nyquist SAR via Fourier Domain Range-Doppler Processing," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 55, no. 11, pp. 6228 – 6244, Nov. 2017.

[16] W. Qui, H. Zhao, J. Zhou, and Q. Fu, "High-Resolution Fully Polarimetric ISAR Imaging Based on Compressive Sensing," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 52, no. 10, pp. 6119–6131, Oct. 2014.

[17] M. Duarte and R. Baraniuk, "Kronecker Compressive Sensing," *IEEE Trans. on Image Processing*, vol. 21, no. 2, pp. 494–504, Feb. 2012.

[18] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, 1st ed., ser. Applied and Numerical Harmonic Analysis. Birkhuser Basel, 2013.

[19] E. van den Berg and M. P. Friedlander, "SPGL1: A solver for large-scale sparse reconstruction," Jun. 2007. [Online]. Available: http://www.cs.ubc.ca/ mpf/spgl1/

[20] ——, "Probing the Pareto frontier for basis pursuit solutions," *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 890–912, Nov. 2008.