

Spectral Super-Resolution of Hyperspectral Imagery Using Reweighted ℓ_1 Spatial Filtering

Adam S. Charles, *Student Member, IEEE*, and Christopher J. Rozell*, *Senior Member, IEEE*

Abstract—Sparsity-based models have enabled significant advances in many image processing tasks. Hyperspectral imagery (HSI) in particular has benefited from these approaches due to the significant low-dimensional structure in both spatial and spectral dimensions. Specifically, previous work has shown that sparsity models can be used for spectral super-resolution, where spectral signatures with HSI-level resolution are recovered from measurements with multispectral-level resolution (i.e., an order of magnitude fewer spectral bands). In this paper we expand on those results by introducing a new inference approach known as reweighted ℓ_1 spatial filtering (RWL1-SF). RWL1-SF incorporates a more sophisticated signal model that allows for variations in the SNR at each pixel as well as spatial dependencies between neighboring pixels. The results demonstrate that the proposed approach leverages signal structure beyond simple sparsity to achieve significant improvements in spectral super-resolution.

Index Terms—Hyperspectral imagery, sparse approximation, reweighted ℓ_1

I. INTRODUCTION

Hyperspectral Imagery (HSI) captures detailed terrestrial information with high resolution in both the spatial and spectral dimensions [1]. While this level of detail is important in many remote sensing tasks such as object detection and material discrimination, acquiring HSI can be more prohibitive than acquiring multispectral imagery (MSI). For example, sensor fabrication cost and image acquisition time (for comparable SNR per spectral bin) increases as the sensor’s bandwidth narrows [1]. The financial cost differential leads to MSI being more accessible than HSI, both in terms of currently flown instruments and archived data. In applications where high temporal resolution is required (e.g., due to high velocity of the imaging platform), the decreasing SNR per bin due to shorter acquisition times might make it more desirable to combine HSI bins to collect imagery with MSI-level spectral resolution.¹

Recent results have shown the potential to use MSI to obtain HSI-level resolution images by performing *spectral super-resolution* [2]. The task of super-resolution is to use prior knowledge of the signal statistics in post-processing to infer the content of the signal at a finer resolution than the original observations. In photographic images, signal models based on the notion of *sparsity* (i.e. images can be described by a small

number of atoms in a potentially large dictionary) have been very successful in spatial super-resolution applications [3].

In previous work we have demonstrated the applicability of sparsity-based methods in spectral super-resolution for HSI [2]. Specifically, by learning a dictionary of spectral signatures that sparsely decompose the spectral response in each pixel, we learn an approximation to the data manifold that captures rich higher-order statistical structure in HSI data. This model can then be used to perform spectral super-resolution from MSI-level data to HSI-level resolution with very high accuracy [2]. In this work we improve on these previous results by proposing a reweighted ℓ_1 spatial filtering algorithm to incorporate spatial regularity to improve spectral super-resolution. This approach closely follows recent work in dynamic filtering where temporal correlations have been used to improve recovery of time-varying signals in a reweighted ℓ_1 framework [4]. The main contribution of this work is to show that more advanced recovery algorithms can produce significant improvements in the spectral super-resolution results for scenes with significant spatial regularity, with most of the improvement coming from pixels that are not well-modeled by a basic sparsity model.

II. SUPER-RESOLUTION VIA SPARSE CODING

A. The Sparse Coding Model

The sparse coding model for HSI represents the observed spectrum at each pixel $\mathbf{x}_{i,j} \in \mathbb{R}^M$ by a small number of terms in a linear generative model

$$\mathbf{x}_{i,j} = \sum \phi_k a_{i,j,k} + \epsilon_{i,j} = \Phi \mathbf{a}_{i,j} + \epsilon_{i,j},$$

where $\Phi \in \mathbb{R}^{M \times N}$ is a matrix of (non-negative) dictionary elements, $\mathbf{a}_{i,j} \in \mathbb{R}^N$ is the (non-negative) coefficient vector for pixel $\{i, j\}$ and $\epsilon_{i,j} \in \mathbb{R}^M$ is a noise term. The coefficient values can be recovered for a given pixel by solving an ℓ_1 regularized least-squares optimization problem (termed Basis Pursuit De-Noising (BPDN))

$$\hat{\mathbf{a}}_{i,j} = \arg \min_{\mathbf{a}} \|\mathbf{x}_{i,j} - \Phi \mathbf{a}\|_2^2 + \gamma \|\mathbf{a}\|_1,$$

under positivity constraints ($\mathbf{a}_{i,j,k} \geq 0$), where γ is a parameter that trades off between the data fidelity (least-squares) term and the sparsity based regularizer (the ℓ_1 norm) [5]. The BPDN optimization can be interpreted as finding a maximum a-priori (MAP) estimate of the coefficients under assumptions of Gaussian noise and an independent identically distributed (i.i.d.) Laplacian prior distribution on the coefficients [2]. This prior probability distribution for the coefficients is chosen to encourage only a small fraction of them to be non-zero

Manuscript received September 06, 2012. This work was supported in part by NSF grant CCF-0830456.

* Corresponding author. The authors are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, 30332-0250 USA (e-mail: {acharles6, crozell}@gatech.edu). The authors are grateful to C. Bachmann at the Naval Research Laboratory for generously providing the Smith Island HSI data set.

¹In the remainder of the paper we will generically refer to MSI data, including the case of MSI-level spectral resolution from an HSI sensor.

simultaneously (e.g., the prior distribution is chosen to have high kurtosis so it has a peak around zero).

The ℓ_1 regularization approach has previously been shown to be effective for un-mixing in HSI [6]–[8]. Furthermore, we have previously shown that a completely unsupervised approach can be used with this model to learn effective dictionary elements [2]. While these learned dictionary elements are not guaranteed to be material spectra (and so are not called endmembers), previous work has shown that they are often highly correlated with known materials in the scene. There is no assumption that the coefficients in each pixel sum to one, meaning that the coefficients can be thought of as relative contributions of each dictionary to the spectral signature (but the total sum of the coefficients will vary depending on the total radiance at that pixel). A more detailed examination of the dictionary elements learned under this model can be found in [2]. Note that although the generative model above is linear, the sparsity constraint introduces a nonlinear component to the model that can capture data statistics that are not well represented in a typical linear mixture model (e.g., forming a local approximation to manifold structures in the data) [2].

B. Spectral Super-Resolution

The main approach to spectral super-resolution will be to learn a dictionary for HSI using a set of training data, and then use this dictionary (which captures higher order statistics in the HSI data) to infer HSI-level resolution from MSI data. Since each MSI band corresponds to a weighted pooling of measurements taken over several HSI bands, we concisely write the relationship between the (unobserved) HSI spectrum and the (observed) MSI spectrum using a matrix multiplication

$$\mathbf{y}_{i,j} = \mathbf{B}\mathbf{x}_{i,j} = \mathbf{B}\Phi\mathbf{a}_{i,j} + \tilde{\epsilon}_{i,j},$$

where $\mathbf{y}_{i,j} \in \mathbb{R}^P$ is the MSI spectrum, $\tilde{\epsilon}_{i,j} = \mathbf{B}\epsilon_{i,j}$ is the MSI resolution measurement error, and $\mathbf{B} \in \mathbb{R}^{P \times M}$ is the “blurring” matrix which pools neighboring bands. We note here that in some cases $\mathbf{y}_{i,j}$ completely omits some spectral bands that are in HSI but not measurable with some MSI sensors, but these bands must still be inferred. Figure 1 shows an example of simulated MSI bands and their relationship to HSI ranges. Note that while we will typically consider pooling operations that have flat spectral responses in bands that do not overlap for simplicity, other more realistic models tailored to a specific sensor could be used as well.

Using the sparse coding model, we infer sparse coefficients in the HSI-level spectral signatures from the MSI data by performing an inverse problem using BPDN

$$\hat{\mathbf{a}}_{i,j} = \arg \min_{\mathbf{a}} \|\mathbf{y}_{i,j} - \mathbf{B}\Phi\mathbf{a}\|_2^2 + \gamma \|\mathbf{a}\|_1. \quad (1)$$

The HSI spectrum can be recovered from these coefficients by $\hat{\mathbf{x}}_{i,j} = \Phi\hat{\mathbf{a}}_{i,j}$. In words, the approach above solves an optimization problem that seeks the HSI coefficients that are both consistent with the model (i.e., they are sparse) and that explain the measured MSI spectrum. In previous work, the above approach is demonstrated with different choices for the matrix \mathbf{B} (representing MSI sensors and HSI sensors run at MSI-level spectral resolution), and with scenes collected

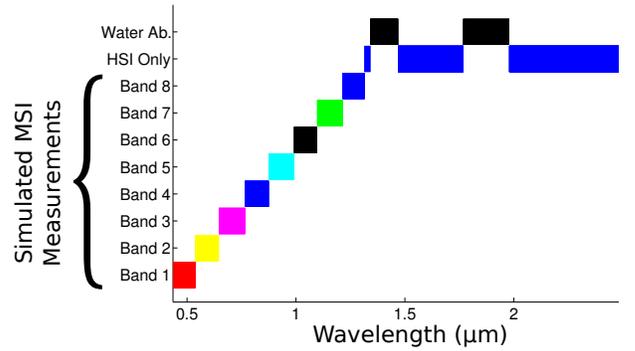


Fig. 1. Simulated MSI responses are comprised of eight spectral bands that pool HSI measurements. The bars show the wavelengths included in each band for the simulated MSI measurements. The top row shows the water absorption bands that are not included in any data. The second row shows the spectral ranges that are present in the HSI data but not included in the simulated MSI measurements (i.e., these bands must be inferred with no data).

at different times of the year (i.e., containing differences in the statistics due to vegetation changes, etc.) [2]. The reconstruction errors between the true (oracle) HSI and the inferred HSI spectra were very low, in the range of 2-3% relative mean squared error (rMSE) depending on the case. While this prior work is encouraging, the results also showed that the relatively small subset of pixels inconsistent with the sparsity model (e.g., mixtures of many types of vegetation) can be outliers with much worse reconstruction performance.

III. REWEIGHTED ℓ_1 FOR SUPER-RESOLUTION

A. Reweighted ℓ_1 (RWL1)

As a first step to improving super-resolution performance, we generalize the sparsity model to allow the SNR for each coefficient to be an unknown parameter that is estimated as part of the inference process. In BPDN (equation (1)), the tradeoff parameter γ depends on the SNR (the ratio of the variance in the sparse coefficients to the noise variance [2]) and is the same for each coefficient. In contrast, the reweighted ℓ_1 (RWL1) framework [9], [10] allows each coefficient $a_{i,j,k}$ its own parameter $\gamma_{i,j,k}$, where a and γ are inferred concurrently. Specifically, RWL1 is equivalent to using the iterative Expectation-Maximization (EM) algorithm to find a joint estimate of a and γ assuming that γ has an i.i.d. Gamma hyperprior distribution. While more technical details of the model and algorithm can be found in [10], the RWL1 algorithm applied to the super-resolution problem can be stated succinctly as alternating a weighted BPDN optimization and an analytic update to the weights until convergence:

$$\hat{\mathbf{a}}_{i,j}^n = \arg \min_{\mathbf{a}} \|\mathbf{x}_{i,j} - \mathbf{B}\Phi\mathbf{a}\|_2^2 + \gamma_0 \sum_k \hat{\gamma}_{i,j,k}^{n-1} a_k,$$

$$\hat{\gamma}_{i,j,k}^n = \frac{\alpha}{|\hat{\mathbf{a}}_{i,j,k}^n| + \beta},$$

where α , β and γ_0 are parameters related to the hyperprior on γ and n is the iteration number.

One way to intuitively understand the RWL1 algorithm is to understand the effect each $\gamma_{i,j,k}$ has on the weighted ℓ_1

optimization problem. Lowering a given $\gamma_{i,j,k}$ value makes it easier for the corresponding coefficient to be activated in the next BPDN iteration. By iteratively recalculating the weights, coefficients that are activated in the initial optimization become more easily activated in future iterations (via smaller weights) and unused coefficients are more difficult to activate in future iterations (via higher weights). Additional literature has linked RWL1 to approximating solutions to ℓ_p regularized least squares problems for $p < 1$ [11] and asymptotic theoretical guarantees in other inverse problems (e.g., compressed sensing) [12].

B. Reweighted ℓ_1 Spatial Filtering (RWL1-SF)

While spectral statistics are informative enough to perform super-resolution in many cases, spatial regularity can often be leveraged in some types of scenes to improve performance (especially when the sparsity model is not a good fit for a given pixel). Spatial regularity was also used recently in the context of material classification, indicating its utility in HSI [13]. Therefore, as a second step to improving super-resolution performance, we further generalize the RWL1 model to incorporate spatial information into the inference process. Specifically, in our proposed reweighted ℓ_1 spatial filtering (RWL1-SF), we update the weights for a given coefficient using a combination of information from the previous iteration on neighboring pixels (similar to the reweighted ℓ_1 dynamic filtering algorithm developed in [4]). In this way, even weak evidence from individual pixels in a local neighborhood can be aggregated to improve the inference in cases that would be particularly difficult when just considering individual pixels independently.

To be precise, consider the matrix of all coefficients for the k^{th} dictionary element, $[\mathbf{A}_k]_{i,j} = a_{i,j,k}$. In each iteration of RWL1-SF, the weight for the k^{th} coefficient at the pixel in row i and column j is set by a weighted pooling of the previous estimates for the k^{th} coefficient at the neighboring pixels. While there are many potential ways to implement this spatial aggregation and weight updating, in this paper we use a simple linear weighted average:

$$\gamma_{i,j,k} = \frac{\alpha}{[\Psi * \mathbf{A}_k]_{i,j}} + \beta$$

where the term $[\Psi * \mathbf{A}_k]_{i,j}$ represents the $\{i,j\}^{\text{th}}$ term of the kernel $\Psi \in \mathbb{R}^{L \times P}$ convolved with the spatial field of previous estimates for the k^{th} coefficient. Note that while this spatial regularization can accumulate weak evidence spread over several neighboring pixels to perform inference, the model does not force spatial homogeneity so that single-pixel (or sub-pixel) objects are missed. In other words, rather than low-pass filtering the estimates of interest (the $a_{i,j,k}$ variables), the spatial averaging is applied to a second order variable ($\gamma_{i,j,k}$) that simply biases a sparse inference process. In fact, though an explicit test with single-pixel anomalies is beyond the scope of this letter, previous work using this approach for dynamic filtering [4] showed that this method of stochastic filtering is particularly robust to model mismatch.

The kernel Ψ incorporates the knowledge that dependencies should have a limited spatial extent and will be modulated

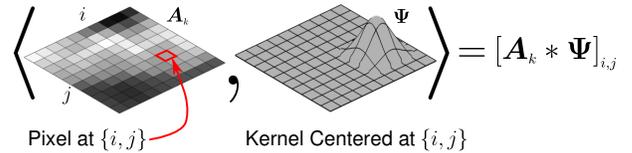


Fig. 2. The kernel Ψ determines the influence from neighboring pixels on coefficient inference at a given location. When the $L \times P$ kernel is centered on the $\{i,j\}^{\text{th}}$ pixel it describes the weighted summation of neighboring coefficient estimates that influence the next coefficient estimate in that pixel.

depending on the distance between the pixels, as depicted in Figure 2. The value in the $\{l,p\}^{\text{th}}$ entry of Ψ indicates the amount which the $\{i+l-L/2, j+p-P/2\}^{\text{th}}$ element of \mathbf{A}_k influences the $\{i,j\}^{\text{th}}$ element of \mathbf{A}_k in the next iteration of the inference. Typically, the center (0,0) value of Ψ should be unity and the kernel values should taper off towards the edges to represent the decaying dependence with distance. In this work we use the same 5×5 pixel Gaussian kernel shape for all parts of the estimation, but in general each coefficient or pixel location could have a different kernel if there was advanced knowledge of the spatial and spectral dependencies in the data. Indeed, in scenes with very different statistics than the HSI used as an example here (e.g., urban scenes), the spatial regularization process may benefit from a specialized treatment of edges in the image.

IV. PERFORMANCE COMPARISONS

We test the performance of RWL1 and RWL1-SF against previous results on segments of HSI from Smith Island, VA. These two HSI images were taken by the PROBE2 sensor on October 18, 2001 and August 22, 2001 and have 113 usable spectral bands spanning the $0.44\text{-}2.486\mu\text{m}$ range (after removal of water absorption bands and applying atmospheric correction to estimate reflectance) and a spatial resolution of approximately 4.5m^2 . We simulate MSI measurements by creating a matrix \mathbf{B} to represent a response function that entirely eliminated measurements in higher wavelength regions and pooled the remaining HSI measurements into eight spectral bands shown in Figure 1 (each row of \mathbf{B} has ones over bands included and zeros otherwise). We learn a 44-element dictionary Φ on the October 18, 2001 image as in [2], and test recovery on both images. Of particular note is that the two images were taken several months apart, and the statistical changes with the seasonal variations made the recovery of the August image the most challenging test case in prior work [2]. We estimate the original 113 bands from the 8 simulated MSI bands for both images via BPDN, RWL1 and RWL1-SF.

For testing purposes we recover a contiguous 68×288 pixel region (omitting 11 pixels with severe sensor errors) from the Smith Island dataset, shown in Figure 3. This region yielded particularly poor performance when using BPDN for super-resolution in prior work [2]. As shown in Table I, the previous

²More details about this dataset can be found in [14].

TABLE I
SUPER-RESOLUTION FROM SIMULATED MSI MEASUREMENTS IN TERMS
OF RELATIVE MSE AND SPECTRAL ANGLE (SA).

October 18 (Same Day)				
	rMSE		SA (degrees)	
	Mean	Median	Mean	Median
BPDN	2.33%	0.35%	5.838°	3.205°
RWL1	0.85%	0.24%	3.817°	2.683°
RWL1-SF	0.68%	0.23%	3.447°	2.575°
August 11 (Different Day)				
	rMSE		SA (degrees)	
	Mean	Median	Mean	Median
BPDN	6.25%	6.25%	11.812°	13.587°
RWL1	3.34%	3.02%	8.824°	9.439°
RWL1-SF	2.45%	1.89%	7.492°	7.382°

mean rMSE was 6.3% and the median rMSE was 3.3% for this region on the August image, which is considerably worse than the performance seen on sets of pixels randomly selected throughout the entire image (nearly triple the 2.456% mean and an order of magnitude higher than the 0.1219% median rMSE observed on the full dataset [2]). As stated in [2], BPDN super-resolution resulted in the highest error in portions of the scene that are expected to have more heterogeneous compositions, therefore making the basic sparsity model a poorer fit than it is in more homogeneous regions. To illustrate this, Figure 3 shows the distribution of BPDN reconstruction errors (measured in spectral angle) for the same day dataset, highlighting the difference in performance in distinct regions of heterogeneous materials on the ground. Unsurprisingly, the higher errors are also concentrated in the HSI spectral bands that are not measured in the MSI data. Previous work [2] shows that if the same number of measurements are taken over the whole HSI spectral range (corresponding to an HSI sensor operating in a lower spectral resolution mode for higher temporal resolution), this ambiguity is reduced and performance increases significantly.

Table I provides mean and median recovery results, illustrating significant performance improvements when using RWL1 instead of BPDN, and further substantial improvements when using RWL1-SF. Figure 4 illustrates two example pixels that are representative of the easiest and most challenging performance for the October image. For the best case, the spectra are nearly indistinguishable from the true HSI. For the worst case reconstruction we note that the errors are clearly concentrated in the unmeasured (high wavelength) spectral ranges and that the proposed algorithms make substantial improvements in the recovery over the previous results using BPDN. Figure 5 illustrates that the overall statistics of the data in the August image are also better preserved when using RWL1-SF instead of BPDN, with first four principal components of the reconstructed data (accounting for 99.99% of the energy in the image segment) much more closely approximating the principal components of the HSI when using RWL1-SF.

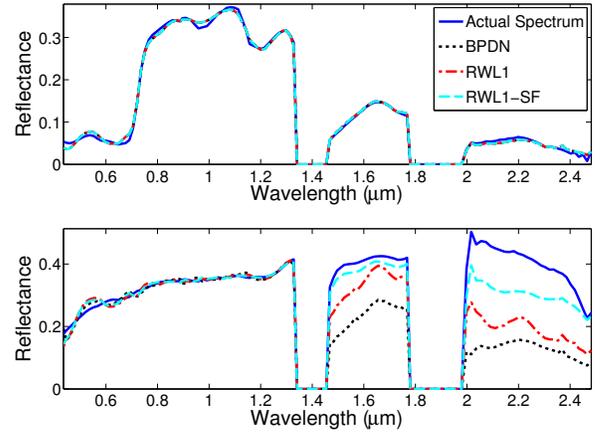


Fig. 4. Two example spectra super-resolved from MSI-level data. Top plot is representative of best-case performance and bottom plot is representative of worst-case performance for previous approaches [2]. Note that errors are highly concentrated in unmeasured bands.

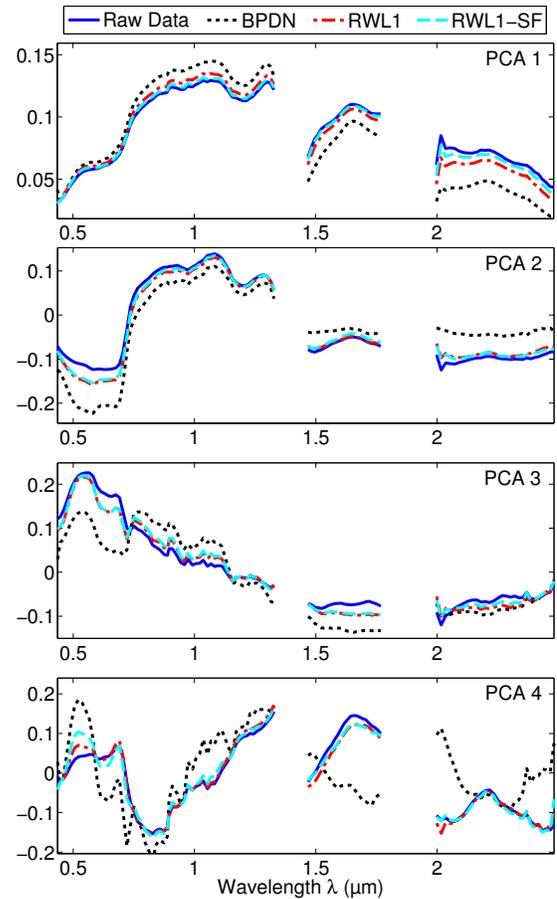


Fig. 5. The first four principal components the recovered HSI spectra compared to the principal components of the original HSI data.

V. CONCLUSIONS

Super-resolving MSI data to HSI-level spectral resolutions is a technique that is of particular importance given the value of high resolution spectral information. The proposed algorithms leverage both more advanced sparsity models in

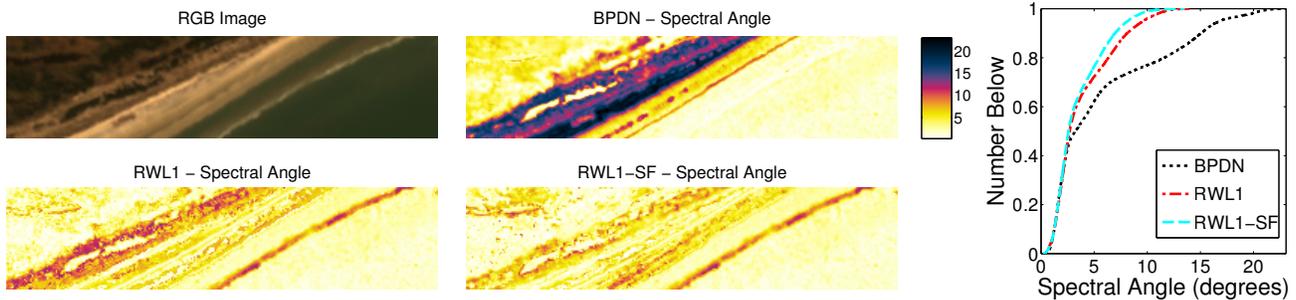


Fig. 3. Left: The RGB image of the October region being tested and the heat maps depicting the spectral angle errors throughout the region using BPDN, RWL1 and RWL1-SF. The largest improvements over BPDN occur along the shoreline where the material mixture is very heterogeneous (e.g., water, sand, vegetation) and the sparsity model alone is insufficient. Right: The cumulative distribution function (CDF) of the spectral angle errors. Note that the BPDN CDF has a heavy tail, indicating many pixels with poor performance. RWL1 improves performance significantly. RWL1-SF uses a model of spatial dependence to further reduce the outliers and improve performance, with 90% of the pixels having spectral angle errors less than 6.9536 degrees.

each pixel, as well as spatial regularity between pixels. This increased model structure improves on previous super-resolution results significantly, especially in the pixels that were outliers in previous results due to their poor super-resolution performance [2]. Specifically, using additional intra-pixel structure in RWL1 yielded a 35.62% and 16.29% improvement in the mean and median SA, respectively. Incorporating spatial dependencies in RWL1-SF boosted these results further, giving a total of 40.96% improvement in the mean SA and 19.66% improvement in median SA. While the value of this super-resolution technique will ultimately need to be verified in terms of performance in specific applications, we note that 90% of the recovered pixels in the current dataset had a spectral angle error less than 7 degrees. While future improvements may continue to be made, this level of error is well within the class spectral width of some classifiers currently in use (e.g., 7 degrees to 30 degree in [15]). Again, we note that the presented data includes some of the most challenging problem aspects from the previous work (i.e., difficult pixels and MSI measurements with no data from some HSI bands).

While the previous approach using BPDN achieved very good performance in many cases, the RWL1-SF enhancement in this work shows substantial improvements in the most challenging test cases. The increased performance demonstrated here leads us to conclude that important structure beyond simple spectral sparsity exists in some types of HSI data and can be exploited for significant gains in super-resolution performance. We also note that this performance improvement comes at moderate additional cost to the previously reported results (no more than a few EM iterations solving BPDN).

While these results are encouraging, there are several further avenues to explore. Most importantly, we have evaluated the proposed algorithms on simulated MSI data to facilitate oracle evaluation with ground-truth HSI data (all in the reflectance domain). A more realistic test with an oracle evaluation would require registered MSI and HSI data collected simultaneously, and a thorough exploration of whether super-resolution is best performed before or after atmospheric compensation. Additionally, the form of Ψ will obviously have a significant impact on super-resolution performance. While the simple Gaussian kernel worked well on the Smith Island dataset,

further work is required to evaluate this spatial regularization on different types of HSI (especially urban scenes where edges are more prominent) and to determine better regularization approaches (including possibly learning Ψ from data).

REFERENCES

- [1] J. P. Kerkes and J. R. Schott, "Hyperspectral imaging systems," in *Hyperspectral Data Exploitation: Theory and Applications*, C.-I. Chang, Ed. John Wiley & Sons, Inc., 2007, pp. 19–45.
- [2] A. S. Charles, B. A. Olshausen, and C. J. Rozell, "Learning sparse codes for hyperspectral images," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 5, pp. 963–978, September 2011.
- [3] M. Elad, M. Figueiredo, and Y. Ma, "On the role of sparse and redundant representations in image processing," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 972–982, 2010.
- [4] A. Charles and C. Rozell, "Re-weighted ℓ_1 dynamic filtering for time-varying sparse signal estimation," 2012, submitted.
- [5] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Review*, vol. 43, no. 1, pp. 129–159, 2001.
- [6] A. Szlam, Z. Guo, and S. Osher, "A split Bregman method for non-negative sparsity penalized least squares with applications to hyperspectral demixing," *Proceedings of the IEEE International Conference on Image Processing*, Feb 2010.
- [7] M. D. Iordache, J. M. Bioucas-Dias, and A. Plaza, "Sparse unmixing of hyperspectral data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 6, pp. 2014–2039, June 2011.
- [8] J. Greer, "Sparse demixing of hyperspectral images," *IEEE Transactions on Image Processing*, vol. 21, no. 1, pp. 219–228, 2012.
- [9] E. Candès, M. Wakin, and S. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5, pp. 877–905, 2008.
- [10] P. Garrigues and B. Olshausen, "Group sparse coding with a Laplacian scale mixture prior," *Advances in Neural Information Processing Systems*, pp. 1–9, 2010.
- [11] D. Wipf and S. Nagarajan, "Iterative reweighted ℓ_1 and ℓ_2 methods for finding sparse solutions," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 317–329, 2010.
- [12] D. Needell, "Noisy signal recovery via iterative reweighted ℓ_1 -minimization," in *Forty-Third Asilomar Conference on Signals, Systems and Computers*, 2009, pp. 113–117.
- [13] A. Castrodad, Z. Xing, J. Greer, E. Bosch, L. Carin, and G. Sapiro, "Learning discriminative sparse representations for modeling, source separation, and mapping of hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 11, pp. 4263–4281, 2011.
- [14] C. M. Bachmann, T. F. Donato, G. M. Lamela, W. J. Rhea, M. H. Bettenhausen, R. A. Fusina, K. R. D. Bois, J. H. Porter, and B. R. Truitt, "Automatic classification of land cover on smith island, va, using hmap imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 40, no. 10, pp. 2313–2330, Oct 2002.
- [15] A. P. Crosta, C. Sabine, and J. V. Taranik, "Hydrothermal alteration mapping at bodie, california, using aviris hyperspectral data," *Remote Sensing of Environment*, vol. 65, no. 3, pp. 309–319, 1998.