LIGHTWEIGHT ONBOARD HYPERSPECTRAL COMPRESSION AND RECOVERY BY MATRIX COMPLETION

Grigorios Tsagkatakis[1], Leonardo Amoruso[2], Dimitrios Sykas[2], Cristoforo Abbattista[2], and Panagiotis Tsakalides[1,3]

[1]*Institute of Computer Science, Foundation for Research and Technology - Hellas (FORTH), Crete, Greece*
[2]*Planetek Italia srl, Bari, Italy*
[3]*Department of Computer Science, University of Crete, Crete, Greece*

ABSTRACT

While traditional hyperspectral imaging platforms rely on a time-consuming sequential scanning of a scene, novel snapshot approaches can acquire the full hyperspectral profile from a single exposure, employing cutting-edge technology such as Spectrally Resolvable Detector Arrays. However, faster scanning using snapshot systems comes with a price as the resulting spatial resolution is significantly reduced. In this work, we propose a new method for the enhancement of the spatio-spectral resolution of hypercubes acquired by snapshot spectral imagers by casting the problem within the novel disruptive mathematical framework of low-rank Matrix Completion. The proposed scheme can also be utilized as a lightweight hyperspectral image compression protocol where only a single spectral band is stored and transmitted for each spatial location, significantly reducing the bandwidth and power requirements, without introducing additional complexity. Experimental results suggest that high quality reconstruction is possible even at very high compression regimes.

Key words: LaTeX; Hyperspectral Imaging, Matrix Completion, Snapshot Spectral Imaging .

1. INTRODUCTION

Hyperspectral images provide critical insights into the composition of a scene and have found multiple applications in remote sensing situations. Unfortunately, acquiring, processing, storing and transmitting full spatio-spectral resolution hyperspectral images can introduce significant challenges, especially when considering the limitations of satellite resources and onboard computing [1]. To illustrate the impact of these challenges on storage and communications, some missions, even if equipped with sensors able to acquire hundreds of spectral bands, employ radical measures, like binning and band rejection, to transmit to the ground only a limited number of bands, tens over hundreds, according to their acquisition plans.

Within this context, the family of Snapshot Mosaic Multispectral Imaging architectures, also known as hyper/multi-spectral Color Filter Arrays, relies on the use of Spectrally Resolved Detector Arrays (SRDA) where each pixel is associated with a specific spectral region, thus allowing the acquisition of a full hyperspectral cube from a single exposure [2, 3]. An example of a spectral mosaic frame and the corresponding full resolution spectral bands is shown in Figure 1.

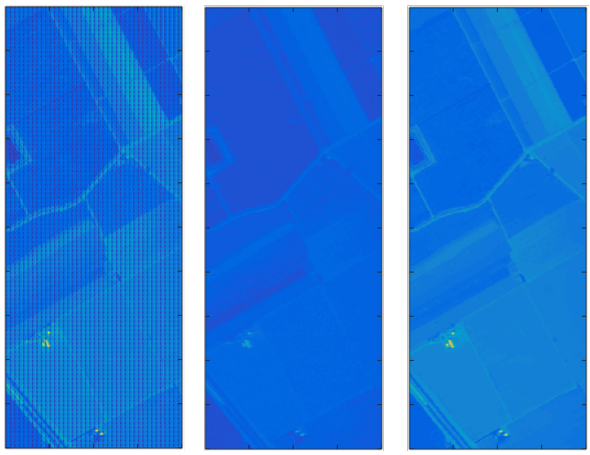

Figure 1: Spectral mosaic acquired by an SRDA imager (left) and the corresponding full resolution images of bands 4 and 7. Although the ground truth images clearly demonstrate the spatial smoothness of the scene, this information is lost during snapshot acquisition.

In order to achieve high temporal resolution imaging, SRDA architectures must sacrifice spatial resolution since only a small subset of pixels acquire energy in each specific spectral band. In practice, pixel binning is performed to group spectral-specific pixels together in full spectral resolution super-pixels. This process is depicted in Figure 2.

Despite the computational efficiency of this approach, such a binning scheme reduces the spatial resolution of the acquired hypercube by a factor of 16 or 25 for a 4×4 or a 5×5 sampling pattern, respectively. To

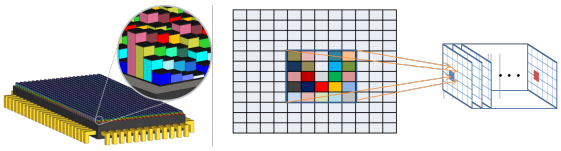

Figure 2: SRDA architecture (left), a snapshot mosaic raw frame (center), and a full spectral resolution "super-pixel" as part of the reconstruction hypercube (right). Notice that the process of producing the "super-pixels" leads to a dramatically smaller spatial resolution.

address this limitation, we propose the enhancement of the low-rank matrices generated from the acquired hyperspectral data by employing the recovery capabilities of the recently developed framework of Matrix Completion (MC), a paradigm-shift in modern multidimensional signal sampling and recovery [2, 10], which asserts that any low-rank matrix can be perfectly recovered from a limited number of randomly selected entries. To validate the potential of the proposed compression scheme, we evaluate the performance on currently available airborne and spaceborne datasets, including measurements from the AVIRIS and the Hyperion instruments.

2. LOW-RANK MATRIX COMPLETION

Our approach is based on the recently proposed framework of Matrix Completion (MC) [2, 10] which has emerged as a disciplined way of addressing the recovery of high-dimensional data from what appears to be incomplete, and perhaps even corrupted information. Low-rank MC has been utilized in a variety of image acquisition and processing tasks including the acquisition of high dynamic range imaging [9] and video denoising [17], among others. More specifically, given a $m \times n$ measurement matrix $\mathbf{M}$, recovering the (mn) entries of the matrix from a smaller number of $k << mn$ entries is not possible, in general. However, it was recently shown that the recovery of the complete set of entries in a matrix is possible, provided that both the number of missing entries and the rank of the matrix are appropriately bounded.

Formally, let $\mathcal{A}$ be a linear map from $\mathbb{R}^{m \times n} \rightarrow \mathbb{R}^k$, that selects a subset of the entries in matrix $\mathbf{M}$. The linear map $\mathcal{A}$, is defined as a random sampling operator that records a small number of entries from matrix $\mathbf{M}$, that is $\mathcal{A}(m_{ij}) = \{1 \text{ if } (ij) \in S \mid 0 \text{ otherwise}\}$, where S is the sampling set. According to the low rank MC paradigm, we can estimate $\mathbf{X}$ from the undersampled matrix $\mathbf{M}$, by solving:

$$\begin{aligned} &\underset{\mathbf{X}}{\text{minimize}}\ rank(\mathbf{X}) \\ &\text{subject to}\ \ \mathcal{A}(\mathbf{X}) = \mathcal{A}(\mathbf{M})\ . \end{aligned} \tag{1}$$

Unfortunately, rank minimization is an NP-hard problem and therefore cannot be applied in practice. Recently, a relaxation of the above problem was shown to produce accurate approximations, by replacing the rank constraint with the more computationally tractable nuclear norm, which represents the convex envelope of the rank. The relationship is manifested by the Singular Value Decomposition (SVD) of the $m \times n$ measurements matrix, into a product of an orthonormal matrix $\mathbf{U}$, a diagonal matrix $\mathbf{S}$ and another orthonormal matrix $\mathbf{V}$, such that $\mathbf{M} = \mathbf{USV}^T$.

According to the spectral theorem associated with the SVD, the number of singular values, *i.e.* the diagonal entries of $\mathbf{S}$, reveals the rank of the matrix. Low rank matrices, such as the ones produced by spatio-temporally correlated processes, are therefore characterized by a small number of singular values. Furthermore, the rank of a measurement matrix might be artificially increased, due to noise that typically follows an independent distribution. Hence, considering a lower-rank approximation of the matrix results in an implicit denoising of the sampled data. One can exploit such prior knowledge to restrict the number of singular values to a small set that accounts for most of the signal's energy by introducing a thresholding operator $\mathcal{T}$ which when applied to the SVD produces the best rank-k estimation: $\mathbf{M}_k = \mathbf{U}\mathcal{T}(\mathbf{S})\mathbf{V}^T$.

Based on the SVD analysis of a matrix, the minimization in (1) can be reformulated as:

$$\begin{aligned} &\underset{\mathbf{X}}{\text{minimize}}\ \ \|\mathbf{X}\|_* \\ &\text{subject to}\ \ \mathcal{A}(\mathbf{X}) = \mathcal{A}(\mathbf{M}), \end{aligned} \tag{2}$$

where the nuclear norm is defined as $\|\mathbf{M}\|_* = \sum |\sigma_i|$, *i.e.* the sum of absolute values of the singular values. Recovery of the matrix is possible, provided that the matrix $\mathbf{M}$ satisfies an incoherence property. The solution of (2) will converge to the solution of (1) with high probability once $k \geq Cq^{6/5} rlog(q)$ random matrix entries are obtained, where $q = max(m, n)$.

For the noisy case, an approximate version can be solved [16], by replacing the equality constraint with an inequality constraint given by $\|\mathcal{A}(\mathbf{X}) - \mathcal{A}(\mathbf{M})\|_F^2 \leq \epsilon$, where $\|\mathbf{X}\|_F^2 = \sum \lambda_i^2$ denotes the Frobenius norm and ϵ is the approximation error. The optimization is therefore formulated as:

$$\begin{aligned} &\underset{\mathbf{X}}{\text{minimize}}\ \ \|\mathbf{X}\|_* \\ &\text{subject to}\ \ \|\mathcal{A}(\mathbf{X}) - \mathcal{A}(\mathbf{M})\| \leq \epsilon. \end{aligned} \tag{3}$$

To solve the nuclear norm minimization problem of (3), various approaches have been proposed. In this work, we employ the Augmented Lagrange Multipliers (ALM) [11] approach, due to its performance with respect to both computationally complexity and recovery capabilities.

3. MC BASED SRDA RECOVERY MECHANISM

The proposed SRDA recovery mechanism is show in Figure 3 where measurements from either a 4×4 or a 5×5

mosaic SRDA imager are grouped together (ensembled) in order to estimate missing measurements, employing the Matrix Completion theory. Formally, the snapshot frame is divided into overlapping patches of side 3×3, and each of them is re-ordered into subsampled vectors containing measurements from 9 out of 16 or 25 spectral bands.

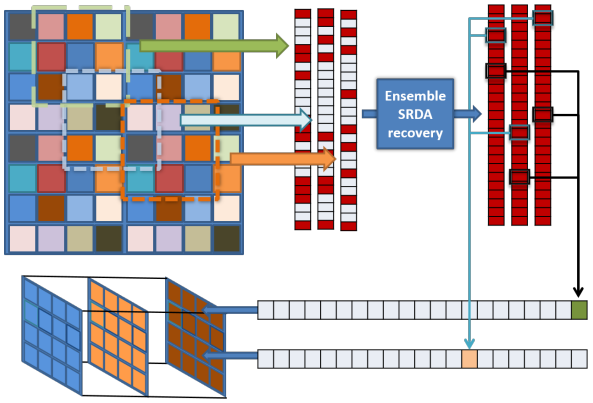

Figure 3: Proposed Matrix Completion based full hypercube recovery mechanism.

These vectors are subsequently grouped into matrices where MC recovery is applied in order to estimate the missing values. In our case, groups of 10 vectors are considered during this grouping operation. The final estimation of the complete spectral profile for each spatial pixel corresponds to the average of the different estimations produced during the recovery of different overlapping patches.

By repeating this process over all possible locations, i.e., the full spatial resolution of the detector, the complete hypercube can be estimated with high accuracy. For application in hyperspectral compression, the proposed scheme could operate by selecting a small number of spectral bands to store and transmit for each spatial location.

4. EXPERIMENTAL EVALUATION

To quantify the performance of the proposed scheme, we consider two performance metrics, namely quality of estimation and achieved compression ratio. Regarding the quality estimation, two metrics, namely Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index Metric (SSIM) are employed while for the compression ratio, we consider the compression ratio achieved by lossless compression of the full spatial resolution hypercube, the ratio achieved for compression of a low resolution hypercube generated from snapshot measurements, and the effective compression ratio which considers the full resolution recovered hypercubes.

We evaluate the performance on three exemplary hypercubes; namely, the Salinas, acquired by the AVIRIS sensor, the Botswana, acquired by the Hyperion sensor on EO-1, and the Pavia University, acquired by the ROSIS sensor [7]. For all cases, we consider the first 16 and the first 25 spectral bands, simulating the acquisition by a 4×4 and a 5×5 snapshot SRDA camera.

4.1. Reconstruction Quality

Figures 4 and 5 present the ground truth and corresponding reconstructions using the proposed method and the method of linear interpolation of spectral bands from the Salinas dataset recovered from a 4×4 and a 5×5 SRDA respectively.

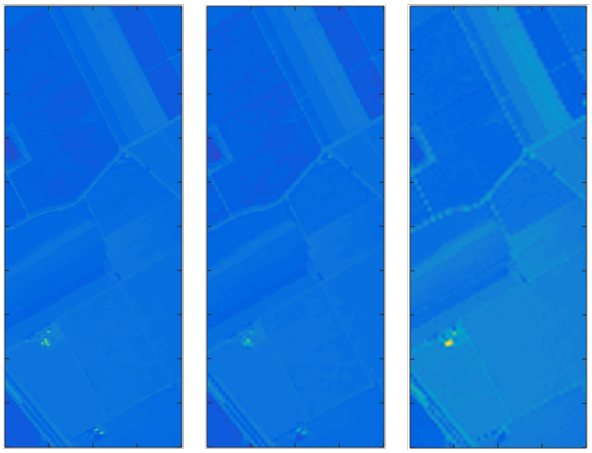

Figure 4: Salinas 7^{th} (left), recovery by proposed method (center), and recovery by interpolation (right) from a 4×4 SRDA camera.

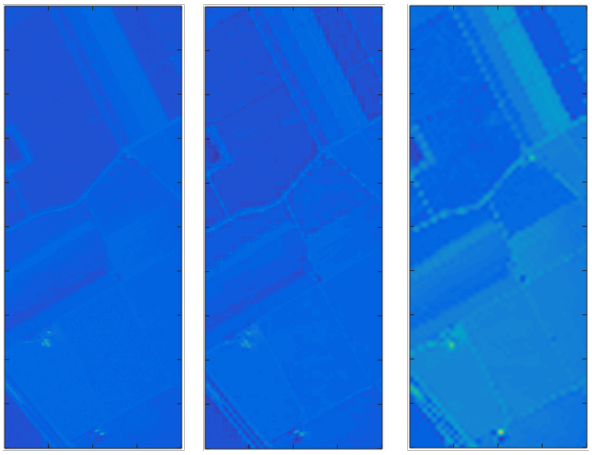

Figure 5: Salinas 4^{th} (left), recovery by proposed method (center), and recovery by interpolation (right) from a 5×5 SRDA camera.

Comparing the two recovery methods, namely interpolation and MC based, one can observe that in both cases, the visual quality is dramatically better, preserving high spatial frequency components like roads and the magnitude of regions of interest in the lower-left part of the

scene. At the same time, the performance of the interpolation based method quickly deteriorates as more spectral bands are considered while the proposed method is able to maintain the high quality of the reconstruction even under challenging conditions.

Similar observations can also be made for the Pavia University dataset as shown in Figures 6 and 7 for the 4×4 and 5×5 snapshot SRDA respectively. For the case of 16 spectral bands, we observe that the proposed scheme preserves information like cars in the parking lot and the street network, while the interpolation based approach leads to sever image quality degradation. For the 25 bands case, the proposed scheme suffers some loss in quality, which is significantly less compared to the interpolation approach.

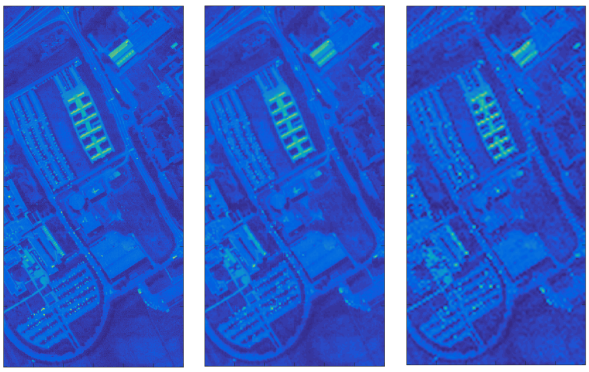

Figure 6: PaviaU 6^{th} (left), recovery by proposed method (center), and recovery by interpolation (right) from a 4×4 SRDA camera.

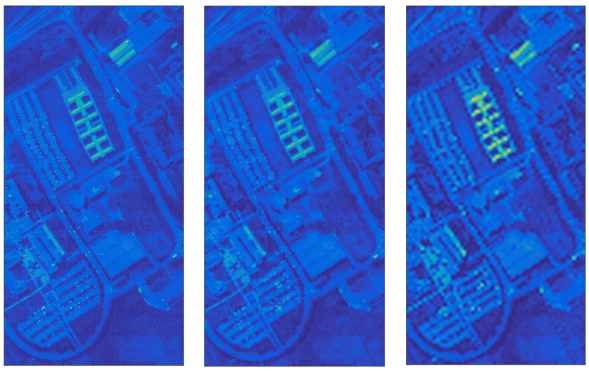

Figure 7: PaviaU 4^{th} (left), recovery by proposed method (center), and recovery by interpolation (right) from a 5×5 SRDA camera.

Last, for the Botswana set shown in Figures 8 and 9 for the two SRDA cases, we observe that for the 25 bands case, the interpolation introduces significant artifacts in both low and high spatial frequencies, while the proposed scheme suffers significantly less deterioration.

For a complete quality evaluation, Tables 1 and 2 provide the average reconstruction quality for the 4×4 and 5×5 cases, respectively. Quantitative results validate the visual observation where for certain cases, the proposed

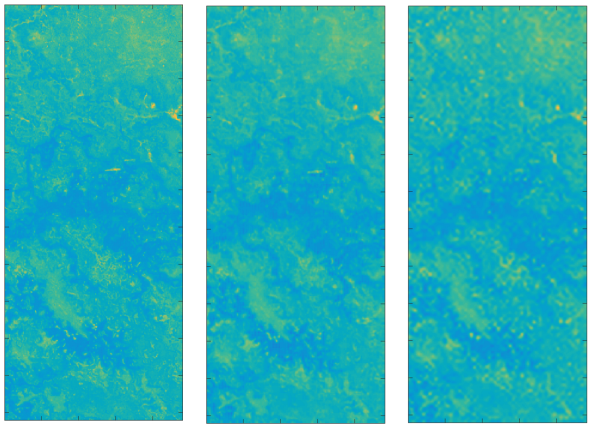

Figure 8: Botswana 6^{th} (left), recovery by proposed method (center), and recovery by interpolation (right) from a 4×4 SRDA camera.

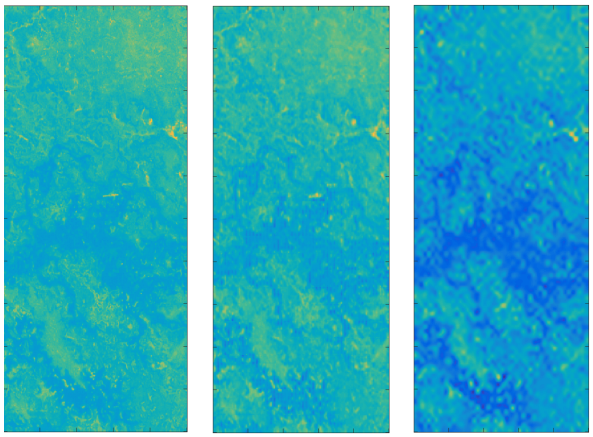

Figure 9: Botswana 5^{th} (left), recovery by proposed method (center), and recovery by interpolation (right) from a 5×5 SRDA camera.

method achieves more than 10dB gain in terms of quality and 20% increase in SSIM (which ranges from 0 to 1).

Table 1: Quality of reconstruction for 16-band hypercubes.

	Interpolated		Proposed	
Sequence	PSNR	SSIM	PSNR	SSIM
Salinas	27.2	0.81	40.3	0.96
Botswana	37.3	0.93	45.0	0.97
PaviaU	26.5	0.67	31.0	0.82

4.2. Compression Ratio

To evaluate the quality of the proposed scheme in terms of compression ratio, we follow the experimental protocol shown in Figure 10, where we distinguish three

Table 2: Quality of reconstruction for 25-band hypercubes.

	Interpolated		Proposed	
Sequence	PSNR	SSIM	PSNR	SSIM
Salinas	26.7	0.81	37.9	0.95
Botswana	22.9	0.58	31.1	0.79
PaviaU	25.1	0.64	29.7	0.82

cases. The first case corresponds to the compression ratio achieved when compression is applied on the *ideal* full resolution hypercube, and it is meant to act as a baseline. The second case assumes the realistic scenario where the measurements from the snapshot SRDA frame are rearranged into a low spatial resolution hypercube by simply binding the measurements from each 4×4 or 5×5 pattern into a single fully spectrally sampled hyperpixel. The last case, extents the second scenario by assuming that the full spatio-spectral resolution cube is recovered after the application of the proposed de-mosaicking scheme, thus the compression assumes snapshot measurements as input and the full resolution hypercube as output.

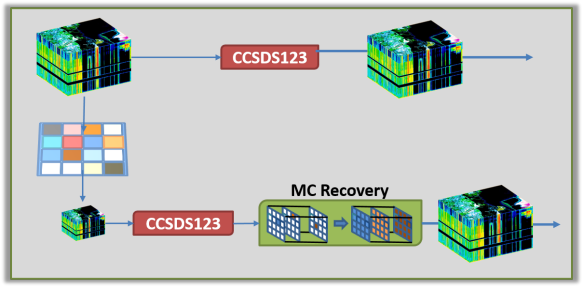

Figure 10: Experimental setup for the evaluation of the compression ratio.

We report the performance employing the CCSDS-123 [4] compression standard, a lossless encoding scheme, since introducing a lossy compression in conjunction with a model-based recovery scheme would lead to significantly lower compression quality. Furthermore, the proposed experimental results assume that only data from a snapshot SRDA sensor is available, thus no high spatial resolution hypercube is available during the acquisition stage.

Tables 3 and 4 report the compression ratios achieved by the three approaches for the case of 4×4 and 5×5, respectively. Comparing the compression ratio achieved by the full and the low resolution cases, we observe that compression of the full resolution hypercubes is significantly more efficient. This is expected since the low resolution hypercubes are produced by spatio-spectral subsampling of snapshot spectral imaging, leading to the introduction of high spatial frequencies which reduce the correlation and increase the entropy.

However, when both loss compression and the proposed recovery mechanism are combined, the *effective* compression ratio is at some cases 10 times higher compared to full resolution compression and 20 times higher compared to the low resolution hypercube compression. This dramatic increase in performance is associated to a moderate reduction in quality as demonstrated in Tables 1 and 2. Furthermore, the proposed compression-reconstruction approach is applicable to both naturally subsampling scenarios as it is the case of snapshot SRDAs, as well as a lightweight compression scheme when full spatio-spectral hypercubes are available.

Table 3: Compression Ratio achieved for the 16 spectral band case.

Resolution	PaviaU	Salinas	Botswana
Full	9.87	11.36	7.34
Low	3.99	4.53	3.47
Effective	63.84	72.48	55.52

Table 4: Compression Ratio achieved for the 25 spectral band case.

Resolution	Pavia U	Salinas	Botswana
Full	9.76	11.43	6.98
Low	3.85	4.25	3.52
Effective	96.25	106.25	88.0

5. CONCLUSIONS

A novel approach for the recovery of full spatio-spectral resolution hypercubes from snapshot spectral imaging architectures is presented. The proposed scheme exploits inherent inter- and intra- band correlations for recovering the missing values with high quality. Experimental results suggest that the considered strategy could allow high speed compact design imaging without sacrificing imaging quality.

ACKNOWLEDGMENTS

This work was partially funded by the PHySIS project (contract no. 640174) within the H2020 Framework Program of the European Commission.

REFERENCES

[1] Hagen, N. and Kudenov, M.W. Review of snapshot spectral imaging technologies. Optical Engineering, 52(9), pp.090901-090901, 2013.

[2] Geelen, B., Tack, N. and Lambrechts, A. A compact snapshot multispectral imager with a monolithically integrated per-pixel filter mosaic. In Spie Moems-Mems (pp. 89740L-89740L), 2014.

[3] Geelen, B., Jayapala, M., Tack, N. and Lambrechts, A. Low-complexity image processing for a high-throughput low-latency snapshot multispectral imager with integrated tiled filters. In SPIE Defense, Security, and Sensing, 2013.

[4] Consultative Committee for Space Data Systems (CCSDS), Lossless Multispectral & Hyperspectral Image Compression CCSDS 123.0- B-1, ser. Blue Book. CCSDS, May 2012. [Online]. Available: http://public.ccsds.org/publications/archive/123x0b1ec1.pdf

[5] Tsagkatakis G, Jayapala M, Geelen B, Tsakalides P. Non-negative Matrix Completion for the Enhancement of Snap-shot Mosaic Multispectral Imagery, SPIE Electronic Imaging, 2016.

[6] Tsagkatakis, G., and Tsakalides, P. Compressed hyperspectral sensing, SPIE/IS&T Electronic Imaging, 2015.

[7] Hyperspectral Remote Sensing Scenes, http://www.ehu.eus/ccwintco/index.php?title=Hyperspectral_Remote_Sensing_Scenes, accessed on: May 2016.

[8] Cands, E.J. and Recht, B. Exact matrix completion via convex optimization. Foundations of Computational mathematics, 9(6), pp.717-772, 2009.

[9] Tsagkatakis, G. and Tsakalides, P. Efficient high dynamic range imaging via matrix completion. In Machine Learning for Signal Processing (MLSP), IEEE (pp. 1-6), 2012.

[10] Candes, E.J. and Plan, Y. Matrix completion with noise. Proceedings of the IEEE, 98(6), pp.925-936, 2010.

[11] Lin, Z., Chen, M. and Ma, Y. The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices. arXiv preprint arXiv:1009.5055, 2010.

[12] Zhang, H., He, W., Zhang, L., Shen, H., and Yuan, Q. Hyperspectral image restoration using low-rank matrix recovery. IEEE Transactions on Geoscience and Remote Sensing, 52(8), 4729-4743, 2014.

[13] Abrardo, A., Barni, M., Bertoli, A., Grimoldi, R., Magli, E., and Vitulli, R. Low-complexity approaches for lossless and near-lossless hyperspectral image compression. In Satellite Data Compression (pp. 47-65). Springer New Yorkm, 2012.

[14] Lossless Multispectral and Hyperspectral Image Compression CCSDS 123.0-R-1, ser. Red Book (draft). CCSDS, May 2011, http://public.ccsds.org/sites/cwe/rids/Lists/CCSDS201230R1/Attachments/123x0r1.pdf, accessed on June 2013.

[15] Xu, Y., Wu, Z., Li, J., Plaza, A., and Wei, Z. Anomaly detection in hyperspectral images based on low-rank and sparse representation. IEEE Transactions on Geoscience and Remote Sensing, 54(4), 1990-2000, 2016.

[16] Candes, E.J. and Plan, Y. Matrix completion with noise. Proceedings of the IEEE, 98(6), pp.925-936, 2010

[17] Ji, H., Liu, C., Shen, Z. and Xu, Y., June. Robust video denoising using low rank matrix completion. In CVPR (pp. 1791-1798), 2010.