

Non-negative Matrix Completion for the Enhancement of Snapshot Mosaic Multispectral Imagery

Grigorios Tsagkatakis¹, Murali Jayapala², Bert Geelen², Panagiotis Tsakalides^{1,3}

¹ ICS - Foundation for Research and Technology - Hellas (FORTH), Crete, Greece

² IMEC, Leuven, Belgium

³ Computer Science - University of Crete, Crete, Greece

Abstract

Multi- and Hyperspectral Imaging (HSI) are characterized by the discrepancy between the dimensionality of hyperspectral image and video data and the dimensionality of the spectral detectors. This issue has been addressed by various architectures, including the Snapshot Mosaic Multispectral Imaging architecture, where each pixel (or group of pixels) is associated with a single spectral band. An unavoidable side effect of this design is the hard trade-off between the spatial and the spectral resolution. In this work, we propose a formal approach for overcoming this trade-off by formulating the problem of full resolution recovery as a low rank Matrix Completion problem. Furthermore, we extend the traditional formulation of Matrix Completion by introducing non-negativity constraints during the recovery process, thus significantly enhancing the reconstruction quality. Experimental results suggest that the Non-Negative Matrix Completion (NN-MC) framework is capable of estimating a high spatial and spectral resolution hypercube from a single exposure, surpassing state-of-the-art schemes like nearest-neighbors as well as the unconstrained Matrix Completion technique.

Introduction

A fundamental issue that hyperspectral imaging sensors have to address is how to collect the four dimensional HSI data, two spatial, one spectral, and one temporal, using a single detector, an 1D array, or 2D plane detectors. The discrepancy between the requested and the available dimensionality of detectors has sparked different philosophies in hyperspectral image acquisition designs, leading to spatial, spectral, and frame scanning architectures [1]. A shortcoming shared by these approaches is the scanning requirements for constructing the complete 3D hyperspectral datacube, where in the case of spatial/spectral scanning, multiple lines/pixels have to be scanned, while for 2D frame scanning systems, multiple frames have to be acquired in order to obtain the complete spectral profile of the scene [2].

These limitations are responsible for a number of issues that hinder the HSI performance, including slow acquisition time and motion artifacts. Furthermore, the need for miniaturization of the imaging systems implies that novel designs should strive to be free of mechanical parts, such as moving mirrors, since they limit the temporal resolution and increase the complexity of the designs. Recent approaches address these limitations by employing novel hardware and sophisticated signal processing techniques to achieve similar performance and imaging capabilities. Snapshot (or Simultaneous) Spectral Imaging (SSI) systems acquire the complete spatio-spectral cube from a single or a few captured

frames, i.e., during a single or a few integration periods, without the need for successive frame acquisition [3]. While earlier approaches relied on additional hardware, such as coherent fiber bundles and mirror slicers to satisfy the requirements for SSI, more recent paradigms employ novel light manipulation components and state-of-the-art signal processing to achieve this task.

One such prominent case is the Snapshot Mosaic Multispectral Imaging architectures, also known as hyper/multi-spectral Color Filter Arrays. This paradigm 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 cubes from a single exposure [4, 5]. Unfortunately, to achieve high temporal resolution imaging, SRDA architectures must sacrifice spatial resolution since only a small subset of pixels acquire images from a specific spectral band. In practice, pixel binning is performed where groups of spectral-specific pixels are grouped together in full spectral resolution super-pixels. The process is depicted in Figure 1.




Figure 1: SRDA architecture (left), a snapshot mosaic raw frame (center), a full spectral resolution "super-pixel" as part of the reconstruction hypercube (right). Notice that the process of producing the "super-pixels" leads to dramatically smaller spatial resolution.

Objectives and state-of-the-art

The objective of this work is to provide a formal method for addressing the challenging spatio-spectral trade-off that characterizes Snapshot Mosaic Multispectral Imaging architecture relying on SRDA detectors. More specifically, SRDA detectors perform spatial subsampling of each spectral band by producing a two-dimensional array of "super-pixels" where each such "super-pixel" corresponds to a binned group of physical pixels, containing measurements from multiple spectral bands. As a consequence, the effective spatial resolution for each spectral band is given by the total number of pixels divided by the number of binned pixels in each super-pixel, as seen in Figure 1.

According to our approach, to address this issue we exploit the inherent redundancies that exist in high dimensional hyperspectral data in order to accurately estimate the missing spectral

bands from binned groups of pixels. Formally, the end goal is to generate a full spatial resolution hypercube $\mathcal{M} \in \mathbb{R}_+^{m \times n \times b}$ from a single exposure image $\mathbf{M} \in \mathbb{R}_+^{m \times n}$, thus allowing imaging of highly dynamic scenes or when imaging takes place in moving platforms such as UAV and satellites. The objective is visually depicted in Figure 2 where a raw snapshot, the corresponding lifted hypercube and the completed hypercube are illustrated.




Figure 2: Illustration of recovery process. A raw snapshot (left) corresponds to an undersampled lifted hypercube (center) which must be completed to obtain the complete spectral profile of the scene (right).

During the last decade, the concept of signal compressibility and sparsity has raised a lot of attention in the mathematics and signal processing communities which have treated novel concepts, including Compressed Sensing (CS) and Matrix Completion (MC), as part of a disruptive new framework which has revolutionized the way we efficiently sense, compress, and process visual information, *e.g.*, [6, 7, 8, 9]. In a nutshell, according to the CS framework, a signal can be perfectly recovered from a severely under-sampled set of measurements provided the signal is sparse in some basis, and the basis on which the signal is sampled is incoherent with the basis on which the signal is sparse [10]. A prominent example of a CS based imaging architecture is the Single Pixel Camera (SPC), employing a Digital Micromirror Device to acquire scene information using a single detector element [11].

A Hyperspectral Single Pixel Camera corresponds to an extension of the typical SPC where the single detector element is replaced by a spectrometer [12]. The Compressive HS Imaging by Separable Spatial and Spectral operators (CHISS) is an alternative spatial scanning architecture which employs a DMD placed before a grating which itself is modulated by a Coded Aperture before acquiring a two-dimensional measurement [13]. Similarly, the Coded Aperture Snapshot Spectral Imaging (CASSI) is a snapshot spectral imaging architecture which employs a DMD for spatially modulating the incoming light before it is dispersed by a grating and imaged by a 2D detector array [14].

More recently, a novel SSI architecture termed Spatial-Spectral Encoded Compressive HS Imager (SSCSI) was proposed combining a spectral dispenser with a random shearing mask, extending the single wavelength computational light field acquisition architecture [15]. The CS framework has also been considered for the compression of multi and hyperspectral imaging without resorting to any modification in hardware, thus maintaining the limitations of current hyperspectral imagers [16]. Recently, the authors in [17] formulated a recovery method for SRDA HSI architectures based on a generalized inpainting approach, while a spatio-spectral Compressed Sensing based acquisition and recovery approach was proposed for HSI data acquisition [18].

Low Rank Matrix Completion

Our approach is based on the recently proposed framework of Matrix Completion (MC) [19, 20] 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 [21] and video denoising [22] 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 $k \ll mn$ 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 the matrix \mathbf{M} . The linear map \mathcal{A} , is defined as a random sampling operator that records a small number of entries from the 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:

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

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:

$$\mathbf{M} = \mathbf{U}\mathbf{S}\mathbf{V}^T . \quad (2)$$

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 Eq. (1) can be reformulated as:

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} \|\mathbf{X}\|_* \\ & \text{subject to } \mathcal{A}(\mathbf{X}) = \mathcal{A}(\mathbf{M}) , \end{aligned} \quad (3)$$

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 (3) will converge to the solution of (1) with high probability once $k \geq Cq^{6/5}r \log(q)$ random matrix entries are obtained, where $q = \max(m, n)$.

For the noisy case, an approximate version can be solved [23], by replacing the equality constraint with an inequality constraint given by $\|\mathcal{A}(\mathbf{X}) - \mathcal{A}(\mathbf{M})\|_F^2 \leq \varepsilon$, 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:

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} \quad \|\mathbf{X}\|_* \\ & \text{subject to} \quad \|\mathcal{A}(\mathbf{X}) - \mathcal{A}(\mathbf{M})\| \leq \varepsilon. \end{aligned} \quad (4)$$

To solve the nuclear norm minimization problem of Eq. (4), various approaches have been proposed. In this work, we employ the Augmented Lagrange Multipliers (ALM) [24, 25] approach, due to its performance with respect to both computationally complexity and recovery capabilities.

In this work, we also extent the ALM approach to MC by considering the situation where the measurements are non-negative, a case that is very common among signal recovery scenarios including the recovery of hyperspectral cubes from SRDA data. To achieve this, the formulation in Eq. (4) is enhanced with an additional non-negativity constraint, leading to:

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} \quad \|\mathbf{X}\|_* \\ & \text{subject to} \quad \|\mathcal{A}(\mathbf{X}) - \mathcal{A}(\mathbf{M})\| \leq \varepsilon \\ & \quad \quad \quad \mathbf{X} \geq 0. \end{aligned} \quad (5)$$

To solve this minimization problem, we begin by introducing the augmented Lagrangian form:

$$\begin{aligned} \mathcal{L}(\mathbf{X}, \mathbf{Y}_1, \mu) = & \|\mathbf{X}\|_* + \text{tr}(\mathbf{Y}_1^T (\mathcal{A}(\mathbf{X}) - \mathcal{A}(\mathbf{M}))) \\ & + \frac{\mu}{2} (\|\mathcal{A}(\mathbf{X}) - \mathcal{A}(\mathbf{M})\|_F^2) \end{aligned} \quad (6)$$

The Lagrangian form encodes all the constraints into a single unconstrained equation which we can solve iteratively by minimizing \mathcal{L} with respect to one variable at each step. The specific algorithmic steps are described in Algorithm 1.

Ensemble Recovery of SRDA data

In order to apply the NN-MC recovery, pixel responses must be reformulated into appropriate matrices. We consider the patch selection operator $\mathcal{S}(x, y, s)$ which collects measurements corresponding to a window of size s centered at location (x, y) from the acquired frame. Concatenating these measurements into a vector $\mathbf{m}_{x,y}$, we repeat this process for a range of p spatial locations x, y generating the undersampled spectral matrix $\mathbf{M} \in \mathbb{R}_+^{s \times b}$, where b is the total number of spectral bands which is completed according to Eq. (5).

In addition to the non-negativity constraint introduced in this work, we also investigate the benefits of an *ensemble* recovery paradigm. We assume that a particular spatial pixel at location (x, y) may belong in up to r different matrices \mathbf{M}_i^* , $i \in [1, r]$. To obtain the final spectral estimation of each pixel, we consider averaging the different estimations such that $\hat{\mathbf{M}}(x, y) = \frac{1}{r} \sum \mathbf{M}_i^*(x, y)$. The process is shown in Figure 3 where one can observe that multiple completed spatial locations are considered during the finalization of the estimation.

Algorithm 1: Non-negative Augmented Lagrange Methods for recovery of Low Rank matrices

Input: The subsampled matrix $\mathbf{X}^0 = \mathcal{A}(\mathbf{M})$, the penalty update parameter α , the error tolerance *threshold* and/or maximum number of iterations *limit*.

Output: The estimated matrix $\hat{\mathbf{X}}$.

- 1: **initialization** $e^0 = 0, k = 0$,
- 2: **while** $error \geq threshold$ or $k \leq limit$ **do**
- 3: Minimize with respect to \mathbf{L}

$$(\mathbf{U}, \mathbf{S}, \mathbf{V}) = \text{SVD}(\mathbf{X}^0 + \mathbf{Y}_1/\mu)$$

$$\mathbf{L}^{(k+1)} = \max(\mathbf{0}, (\mathbf{U}\mathcal{T}(\mathbf{S})\mathbf{V}^T))$$

- 4: Minimize with respect to \mathbf{X}

$$\mathbf{X}^{(k+1)} = \mathbf{M} + \tilde{\mathcal{A}}(\mathbf{L}^{k+1} - \mathbf{I} - \mathbf{Y}_1/\mu)$$

- 5: Update Lagrangian multipliers

$$\mathbf{Y}_1^{(k+1)} = \mathbf{Y}_1^{(k)} + \mu^{(k)}(\mathbf{M}^{(k+1)} - \mathbf{L}^{(k+1)})$$

- 6: Update penalty term

$$\mu^{(k+1)} \leftarrow \alpha \mu^{(k)}$$

- 7: Estimate error

$$e^{(k+1)} = \|\mathcal{A}(\mathbf{M}^0 - \mathbf{M}^{k+1})\|_2$$

set $k \leftarrow k + 1$

- 8: **end while**
-

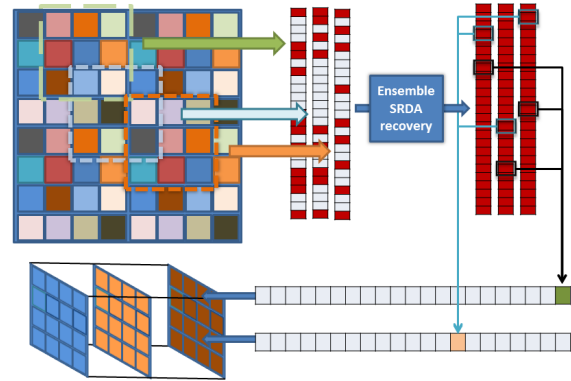


Figure 3: Ensemble SRDA recovery process. Patches are extracted from different locations of the row mosaic frame before they are concatenated into undersampled matrices. The proposed method produces fully populated matrices where the spectral content of each pixel is encoded into multiple locations. The different spectral estimations are averaged in order to estimate the full hypercube.

Experimental Results

To validate the merits of the proposed approach, we explored the enhancement of images acquired using a Ximea cam-

era, equipped with the IMEC Snapshot Mosaic sensor capturing images in the 600 – 875 nm range, according to a 5×5 spectral pattern structure. We consider three recovery algorithms, namely the k-nearest neighbor imputation (KNN) methods, the typical ALM based MC approach and the proposed Non-Negative Matrix Completion (NN-MC), while the error is measured per band by the Peak Signal to Noise Ration (PSNR) metric, in dB. Figure 4 presents the raw input mosaic and the resulting spectral bands where a subset of pixels in the raw input associated with a specific band is selected.




Figure 4: Raw input image (top), the 1st (608 nm), 8th (790 nm) and 16th (644 nm) spectral bands (bottom). The images are presented in scale which demonstrates the differences in spatial resolution.

Based on our problem formulation, the two key parameters that control the recovery performance are the window size s , and the number of spatial locations p that are considered during the generation of the undersampled matrix \mathbf{M} . The window size parameter directly controls the “super-pixel” grouping and ranges from 0 where a single pixel is considered to the extreme case of 5 where a 5×5 region is considered providing the full spectral content. The parameter of the number of spatial locations is directly related to the size of the undersampled matrix that is processed and can lead to different recovery performance as will be seen later.

Effect of super-pixel grouping

A natural way of solving the problem of SRDA spatial resolution enhancement via MC is to consider the unfolding of the hypercube into groups of spatial “super-pixels” encoded in matrix forms. The question we seek to answer is whether a sequential selection of the super-pixels is preferred over a random selection. On the one hand, a sequential selection suggests that the rows of the matrices that will be completed will be composed of spatially neighboring location thus offering the potential of exploiting inherent spatial correlations. On the other hand, such a sampling scheme violates the requirements for incoherent sampling imposed by the theoretically justified recovery capabilities of MC.

We first consider the effects of the window size, and thus the “super-pixel” construction, on the recovery performance. Figure 5 presents the case of no grouping ($s = 0$), Figure 6 presents the $s = 3$ case where 9 spectral measurements are available for each “super-pixel”. In this subsection, we assume that 25 spatial lo-

cations are considered leading to 25×25 undersampled spectral matrices, where the undersampling rate is controlled by s .




Figure 5: Recovery using sequential mapping and 1×1 “super-pixels”.




Figure 6: Recovery using sequential mapping and 3×3 “super-pixels”.

Considering the results, two key observations can be made. First, increasing the “super-pixel” size has a positive impact on recovery performance of all methods. This behaviour is expected since in the second case, more measurements are assumed to be available, which make the recovery processes easier. The second observation is that this problem has considerable more positive impact on the MC based architectures, namely typical MC and the proposed NN-MC. Based on these results, we selected the 3×3 window in the following experiments.

In addition to the effects of the “super-pixel” size, we also investigated if selecting sequential versus random regions during the construction of the spectral matrix also affects the recovery performance. Figure 7 presents the recovery performance for each band that is achieved through a random selection of patches for the generation of the spectral matrix. These results suggest that this zero computational cost step of randomization can have a dramatic impact of the behavior of the recovery process.

Spectral matrix generation

In this subsection, we explore the effects of super-pixel grouping which is introduced in order to transform the hypercube data into matrices. One dimension of these matrices is fixed and is equal to the requested number of spectral bands b , 25 in our case. The other dimension however is controlled by the number of spatial “super-pixels” that will be considered during the generation of the spectral matrices. In general, the grouping number is lower bounded by 25 which leads to 25×25 spectral matrices while we consider up to 400×25 spectral matrices.

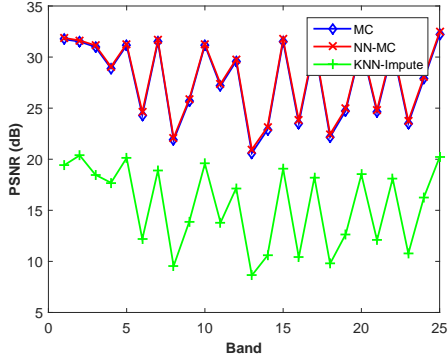


Figure 7: Recovery using random mapping and 3×3 “super-pixels”.

Figure 8 presents the mean recovery performance with respect to the group size, while Figure 9 provides information regarding the processing time required for different group sizes. The results suggest that increasing the group size does not have any noticeable effect as far as the KNN approach is concerned. On the other hand, MC based methods seem particularly affected by this value. More specifically, the results suggest that increasing the grouping size to 50, *i.e.* a spatial to spectral ratio of 2 : 1 introduces more correlated data, leading to lower reconstruction error. The performance of the Non-Negative MC remains relatively stable above 50, however, the typical MC approach experiences a reduction in performance at higher grouping sizes. Considering the processing requirements encoded in the processing time shown in Figure 9, one can observe that increasing the grouping size leads to lower processing times, since for a particular sensor spatial resolution, a smaller number of undersampled spectral matrices have to be processed.

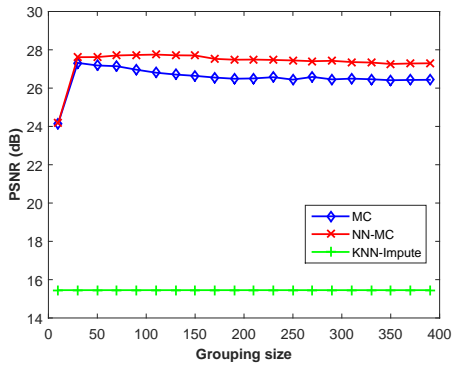


Figure 8: Average recovery performance as a function of grouping size, *i.e.* dimensions of spectral matrices.

Effect of ensemble recovery

In the previous subsection of experimental results, we observed that generating 50×25 spectral matrices with 9 spectral measurements per location, due to the 3×3 “super-pixel” groupings leads to the best performance especially when the locations that are considered during the generation of the spectral matrices are selected randomly instead of sequentially. This subsection goes a step further by exploring the experimental evidence associated to the benefits of randomized ensemble recovery.

We explore various degrees of randomized re-sampling rates, starting from 0 which encodes the case where no-resampling and a sequential ordering is considered during the spectral matrix gener-

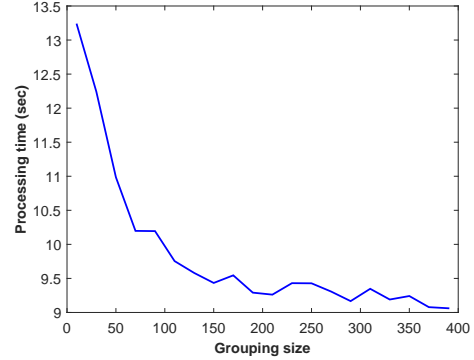


Figure 9: Processing time for full hypercube recovery as a function of grouping size.

eration, to 20 where the same pixel is evaluated over 20 different spectral matrices and the final value corresponds to the average. Figure 10 presents the performance for 25-dimensional spectral matrices, while Figure 11 presents the 50-dimensional spectral matrix recovery case.

Observing these results, one can easily notice that there is no gain in terms of performance for the KNN-impute method, as expected. On the contrary, there is a significant performance gain when the typical MC and the proposed NN-MC are considered. More specifically, we observe that increasing the randomization rate offers a dramatic gain in improvements, especially when going from 0, *i.e.* no randomization, to moderate randomization rates. Furthermore, the results also indicate that the introduction of the non-negativity constrain can play a dramatic role in the recovery, especially when large resampling rates are considered.

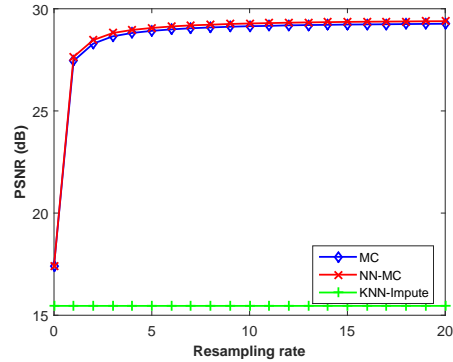


Figure 10: Average recovery performance as a function of resampling rate for 25-dimensional spectral matrices.




Figure 11: Average recovery performance as a function of resampling rate for 50-dimensional spectral matrices.

Figure 12 provides a visual illustration of the reconstruction using the KNN-imputation method and the proposed Non-Negative MC. The results clearly demonstrate the superiority of the proposed approach in high quality recovery.




Figure 12: Reconstruction of 8th (790 nm) band using the KNN-impute (top) and the proposed Non-Negative MC (bottom).

Conclusions

The novelty of this work is twofold. On the one hand, we provide a mathematically sound approach for the estimation of missing spectral measurements, thus enhancing the high temporal resolution imaging capabilities that characterize SRDA based Snapshot Mosaic Multispectral Imaging architectures. As a consequence, high quality imaging of dynamic phenomena can be achieved from extremely low volume and weight hyperspectral imagers. On the other hand, this work proposes a novel approach in low rank matrix estimation through the development of the Non-Negative Matrix Completion framework and the ensemble recovery approach. Experimental results suggest that NN-MC can indeed be utilized for the estimation of full hypercubes, while ensemble recovery can have a dramatic impact in performance.

References

- [1] Bioucas-Dias, J.M., Plaza, A., Camps-Valls, G., Scheunders, P., Nasrabadi, N.M. and Chanussot, J., 2013. Hyperspectral remote sensing data analysis and future challenges. *Geoscience and Remote Sensing Magazine, IEEE*, 1(2), pp.6-36.
- [2] Gupta, N., Dahmani, R., Bennett, K., Simizu, S., Suhre, D.R. and Singh, N.B., 2000, July. Progress in AOTF hyperspectral imagers. In *AeroSense 2000* (pp. 30-38).
- [3] Hagen, N. and Kudenov, M.W., 2013. Review of snapshot spectral imaging technologies. *Optical Engineering*, 52(9), pp.090901-090901.
- [4] Geelen, B., Tack, N. and Lambrechts, A., 2014, March. A compact snapshot multispectral imager with a monolithically integrated per-pixel filter mosaic. In *Spie Moems-Mems* (pp. 89740L-89740L).
- [5] Geelen, B., Jayapala, M., Tack, N. and Lambrechts, A., 2013, May. Low-complexity image processing for a high-throughput low-latency snapshot multispectral imager with integrated tiled filters. In *SPIE Defense, Security, and Sensing* (pp. 87431E-87431E).
- [6] Tsagkatakis, G., Woiselle, A., Tzagkarakis, G., Bousquet, M., Starck,

- J.L. and Tsakalides, P., 2015. Multireturn compressed gated range imaging. *Optical Engineering*, 54(3), pp.031106-031106.
- [7] Fotiadou, K., Tsagkatakis, G. and Tsakalides, P., 2014. Low Light Image Enhancement via Sparse Representations. In *Image Analysis and Recognition* (pp. 84-93). Springer International Publishing.
- [8] Tsagkatakis, G., Woiselle, A., Tzagkarakis, G., Bousquet, M., Starck, J.L. and Tsakalides, P., 2012, November. Active range imaging via random gating. In *SPIE Security+ Defence* (pp. 85420P-85420P).
- [9] Starck, J.L., Murtagh, F. and Fadili, J.M., 2010. *Sparse image and signal processing: wavelets, curvelets, morphological diversity*. Cambridge university press.
- [10] Donoho, D.L., 2006. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4), pp.1289-1306.
- [11] Duarte, M.F., Davenport, M.A., Takhar, D., Laska, J.N., Sun, T., Kelly, K.E. and Baraniuk, R.G., 2008. Single-pixel imaging via compressive sampling. *IEEE Signal Processing Magazine*, 25(2), p.83.
- [12] Sun, T. and Kelly, K., 2009, October. Compressive sensing hyperspectral imager. In *Computational Optical Sensing and Imaging* (p. CTuA5). Optical Society of America.
- [13] August, Y., Vachman, C., Rivenson, Y. and Stern, A., 2013. Compressive hyperspectral imaging by random separable projections in both the spatial and the spectral domains. *Applied optics*, 52(10), pp.D46-D54.
- [14] Wagadarikar, A., John, R., Willett, R. and Brady, D., 2008. Single disperser design for coded aperture snapshot spectral imaging. *Applied optics*, 47(10), pp.B44-B51.
- [15] Lin, X., Liu, Y., Wu, J. and Dai, Q., 2014. Spatial-spectral encoded compressive hyperspectral imaging. *ACM Transactions on Graphics (TOG)*, 33(6), p.233.
- [16] Barducci, A., Guzzi, D., Lastrì, C., Nardino, V., Pippi, I. and Raimondi, V., 2014, October. Compressive sensing for hyperspectral Earth observation from space. In *International Conference on Space Optics* (Vol. 7, p. 10).
- [17] Degraux, K., Cambareri, V., Jacques, L., Geelen, B., Blanch, C. and Lafruit, G., 2015. Generalized inpainting Method for Hyperspectral Image Acquisition. *arXiv preprint arXiv:1502.01853*.
- [18] Tsagkatakis, G. and Tsakalides, P., 2015, March. Compressed hyperspectral sensing. In *IS&T SPIE Electronic Imaging* (pp. 940307-940307).
- [19] Cands, E.J. and Recht, B., 2009. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6), pp.717-772.
- [20] Recht, B., Fazel, M. and Parrilo, P.A., 2010. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3), pp.471-501.
- [21] Tsagkatakis, G. and Tsakalides, P., 2012, September. Efficient high dynamic range imaging via matrix completion. In *Machine Learning for Signal Processing (MLSP), 2012 IEEE* (pp. 1-6).
- [22] Ji, H., Liu, C., Shen, Z. and Xu, Y., 2010, June. Robust video denoising using low rank matrix completion. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE* (pp. 1791-1798).
- [23] Candes, E.J. and Plan, Y., 2010. Matrix completion with noise. *Proceedings of the IEEE*, 98(6), pp.925-936.
- [24] Lin, Z., Chen, M. and Ma, Y., 2010. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055*.
- [25] Yang, J. and Yuan, X., 2013. Linearized augmented Lagrangian and alternating direction methods for nuclear norm minimization. *Mathematics of Computation*, 82(281), pp.301-329.