

MULTISPECTRAL SNAPSHOT DEMOSAICING VIA NON-CONVEX MATRIX COMPLETION

Giancarlo A. Antonucci*, Simon Vary*, David Humphreys[‡], Robert A. Lamb[‡], Jonathan Piper[†], Jared Tanner^{*,‡}

* Mathematical Institute, University of Oxford, Oxford, OX2 6GG

‡ The Alan Turing Institute, NW1 2DB

† DSTL, Sensor Systems, Sensors and Countermeasures Department, Porton Down, Salisbury, SP4 0JQ

‡ Leonardo MW Ltd., Edinburgh, EH5 2XS

ABSTRACT

Snapshot mosaic multispectral imagery acquires an under-sampled data cube by acquiring a single spectral measurement per spatial pixel. Sensors which acquire p frequencies, therefore, suffer from severe $1/p$ undersampling of the full data cube. We show that the missing entries can be accurately imputed using non-convex techniques from sparse approximation and matrix completion initialised with traditional demosaicing algorithms. In particular, we observe the peak signal-to-noise ratio can typically be improved by 2 dB to 5 dB over current state-of-the-art methods when simulating a $p = 16$ mosaic sensor measuring both high and low altitude urban and rural scenes as well as ground-based scenes.

Index Terms— Snapshot multispectral imaging. Sparse approximation. Compressed sensing. Matrix completion. Demosaicing.

1. INTRODUCTION

Multispectral imaging is the process of recording 2D arrays of information at multiple spectra (frequencies). Having access to such a rich, three-dimensional data cube allows different materials to be distinguished due to their differing spectral emission profiles. As a result, multispectral imaging is used in applications ranging from landmine detection, precision agriculture, and medical diagnosis to name but a few of its application domains; for a partial survey see the January 2014 special issue of IEEE Signal Processing Magazine [1]. The increased sensor size and acquisition time are some of the central obstacles to the more widespread use of multispectral imagery.

Snapshot mosaic multispectral sensors allow for a compact video rate multispectral imaging by acquiring only a

fraction of the multispectral cube. For example¹, the IMEC SNM4x4 records 16 bands at a rate of 340 frames per second on a spatial two-dimensional 2048×1088 pixel domain by only acquiring a single spectrum per pixel; specifically this is achieved by tiling the two-dimensional domain by 4×4 pixel supercells where each supercell acquires the spectra independently.

Herein we demonstrate the efficacy of multiple methods for interpolating the missing values in the above snapshot mosaic data cube by simulating the undersampling from complete three-dimensional data cubes provided by DSTL as well as AVIRIS [2], Stanford SCIEN [3], Nascimento [4], Foster [5], IEEE GRSS Data Fusion Contest [6]. In addition to reviewing the existing state-of-the-art interpolation methods in Sec. 2.1, we demonstrate sparse approximation and matrix completion methods in Sec. 2.2 and 2.3 respectively, which we observe to substantially outperform the prior state-of-the-art. Over the above diverse data sets, we observe that non-convex compressed sensing and matrix completion methods initialised with traditional interpolations methods typically improve the peak signal-to-noise ratio by 2 dB to 5 dB, see Table 1.

2. ALGORITHMS FOR DEMOSAICING

Demosaicing is the process by which the undersampled three-dimensional multispectral data cube has the missing entries approximated so as to simulate a full data acquisition. While most three-dimensional interpolations methods would be directly applicable, we consider a few methods previously used by the multispectral community, such as direct interpolation as described in Sec. 2.1 as well as sparse approximation regularisation methods in Sec. 2.2 and low-rank structure as presented in Sec. 2.3.

This publication is based on work partially supported by: the EPSRC Centre For Doctoral Training in Industrially Focused Mathematical Modelling (EP/L015803/1) in collaboration with DSTL, the EPSRC I-CASE studentship (voucher 15220165) in partnership with Leonardo, The Alan Turing Institute through EPSRC (EP/N510129/1) and the Turing Seed Funding grant SF019, Institut Henri Poincaré (UMS 839 CNRS-Sorbonne Université), and LabEx CARMIN (ANR-10-LABX-59-01).

¹We illustrate the architecture through the IMEC sensor, but note there are numerous similar sensors such as the S 137 system by CUBERT.

2.1. Direct interpolation methods

Brauers et al. [7] developed methods to estimate the missing values in the multispectral cube based on extending a spatial bilinear interpolation of the measured values to include any spectral correlation. The weighted bilinear interpolation (WB) for the 4×4 pixel regular mosaic filter follows by padding the missing entries with zeros and convolving with the cartesian product of a discrete seven-pixel width filter $\frac{1}{4}[1 \ 2 \ 3 \ 4 \ 3 \ 2 \ 1]$. Then, the spectral correlation is included in the spectral difference (SD) method by a) taking the output of WB to independently compute, for each band, say k , the difference between the values of the measured pixels for spectrum k and the WB interpolated values of every other band restricted to the support of the measured pixels of spectrum k , then b) applying WB to these spectral differences c) to form an approximation of the full spectrum k by adding the output of step (b) to the difference with l at the location of the measured pixels for spectrum l .

Mihoubi et al. [8] extended the SD approach to consider alternative ways to build correlations between the bands. In intensity difference (ID) they build spectral correlation by first constructing a spatial intensity map whose value at a pixel is the measurement for whichever spectra was measured at that spatial location, then this intensity map is averaged using a weighting based on the number of pixels per spectra contained in the averaging width. See Sec.3.2 for details on the choice of averaging used here and Mihoubi et al. [8] for a more general discussion. Hence, the difference between this averaged intensity map and the measurements for each spectrum is computed, the unknown values zero padded, and each band averaged such as in WB.

Interpolation methods designed in transform domains have been considered by Miao et al. [9] in the binary tree-based edge-sensing (BTES) method, which has the additional benefit of allowing for variable sampling densities per frequency band. However, we observe it is inferior to SD and ID described above in the setting of snapshot imaging. Pseudo-panchromatic image difference (PPID) [10] builds upon BTES and ID. However, due to the applicability of PPID to only some specific mosaic arrangements we leave comparison with our algorithms for a later time.

2.2. Sparse approximation and compressed sensing inpainting

Sparse approximation inpainting allows one to easily consider the interpolation of the under-complete snapshot data cube in transforms more general than the linear interpolation of (WB). In particular, one can assume that the image is well approximated by a sparse representation in a suitable image domain and exploit this structure to reconstruct it from under-sampled measurements [11–13], e.g. by solving

$$\min_x \|y - P_\Omega \Psi^{-1} x\|_2, \quad \text{s.t. } \|x\|_0 \leq k, \quad (1)$$

where Ψ represents the transform in which the data is known to be compressible and y is the full data cube projected by P_Ω to the undersampled locations. Degraux et al. [14] apply this model to a reconstruction of multispectral imagery acquired by mosaic snapshot cameras.

The primary challenge lies in two aspects: (i) the significant subsampling ratio of $1/K$, where K is the number of spectral bands, and (ii) the selection of the suitable transform Ψ . The first problem can be overcome by initialising the state-of-the-art sparse approximation algorithms for solving (1) with sufficiently accurate initial estimates, such as those from the classical interpolation methods described in Sec. 2.1. As it was pointed out in [14], we find that, for natural scenes captured by snapshot multispectral imaging, a Kronecker product of 2D wavelet transform spatially and the discrete cosine transform for the spectral dimension is an effective choice for the representation Ψ . In particular, the 2D wavelet transform spatially includes elements of nearly global support to allow broad correlations as well as local elements to express fine detail and the discrete cosine transform models the slowly varying values in the spectral dimension.

2.3. Matrix completion

Rather than using local correlations, matrix (and tensor) completion exploit the correlation in the data cube through a low-rank structure, e.g. by solving

$$\min_x \|y - P_\Omega X\|_2, \quad \text{s.t. } \text{rank}(X) \leq r, \quad (2)$$

where y is the observed data, X is a matrix corresponding to an unfolding of the complete three-dimensional data cube, and P_Ω is a restriction to the measured values as before.

Although the low-rank matrix completion problem is NP-hard in general (see [15] for a recent survey) there is a number of computationally fast solvers for the problem with provable convergence guarantees [16–18]. In fact, matrix completion has been previously applied to the reconstruction of subsampled multispectral imagery [19] by the Coded Aperture Snapshot Spectral Imager (CASSI) [20]. Here we show that matrix completion can be used also in the case of a more severe subsampling by snapshot mosaic camera designs if provided with suitable initialisation.

While there are many non-convex methods for matrix completion, we showcase two exemplary cases but expect that other non-convex methods would perform similarly well. We apply conjugate gradient iterative hard thresholding (CGIHT) [18] and alternating steepest descent (ASD) [17] to solve (2), providing them with an initial guess from either SD or ID. We show that both CGIHT and ASD can improve on the classical interpolation methods. This differs substantially from prior work both in terms of using more recent algorithms for matrix completion which have been shown to be more effective and initialising them with prior state-of-the-art interpolation methods, and moreover in that unlike [19] which

treat each spectral band separately with 30% undersampling, we vectorise the spatial dimensions to create a matrix of size 2,228,224 by 16 with 1/16 undersampling. We observe that this unfolding which allows full correlation between the spectral information is particularly effective, often resulting in reconstructions which are visually indistinguishable from fully acquired data.

3. NUMERICAL SIMULATIONS

In this section, we show and explain the numerical results obtained by applying the methods discussed above.

3.1. Data sets

We consider the efficacy of the algorithms for demosaicing on the following data sets:

- High altitude airborne images from the AVIRIS [2] and 2018 IEEE GRSS Data Fusion contest [6]. AVIRIS line-scan captures 224 spectral bands in the 380 nm to 2500 nm and the GRSS images have 48 spectral bands in the range of 380 nm to 1050 nm.
- Low altitude airborne images acquired at DSTL Porton Down, in August 2014, from which we selected 10 representative radiance images of fields from a HySpex VNIR-1600 line-scan camera in the range 400 nm to 1000 nm.
- Ground-based images from the Stanford SCIEN [3], Nascimento [4] and Foster [5]. The Stanford SCIEN images come from the line-scan HySpex VNIR-1600 camera.


We processed these data sets to simulate the spectrum measures by the IMEC SNm4x4 snapshot sensor with access to the complete data cube. Then, we undersampled the data cube following the sensor sampling pattern and the following simulations conducted.

3.2. Simulation setup

We implement and test recovery by two interpolation methods ID and SD, two matrix completion methods ASD and CGIHT and a compressed sensing version of CGIHT with a sparsifying transform as a Kronecker products of 2D Daubechies wavelets (W2) in the spatial and 1D Discrete Cosine Transform (DCT) in the spectral domain which we reference as W2×DCT.

The iterative algorithms are terminated once the error in iteration t is $\|P_{\Omega}X^{(t)} - y\|_2/\|y\|_2 \leq 10^{-7}$ or at the 500th iteration.

We report the quality of an image approximated by demosaicing by the peak signal-to-noise ratio (PSNR), defined as the log of the ratio between the maximum possible power of



(a) PSNR throughout spectral bands.



(b) W2×DCT [33.2 dB].



(c) CGIHT [35.3 dB].

Fig. 1: Results for reconstruction of N06. W2×DCT (CS) and CGIHT (MC) initialised from SD. Wavelet based CS method smooths out the image while CGIHT MC is able to better preserve sharp edges.

(the slide of) an image and the power of corrupting noise that affects the fidelity of its representation, computed in terms of the average squared difference (or mean squared error, MSE) between the reference image and its reconstruction:


$$\text{PSNR}_k = 10 \log_{10} \left(\frac{(\max_{p \in \mathcal{P}} I_p^k)^2}{\frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} (I_p^k - \hat{I}_p^k)^2} \right), \quad (3)$$

where I^k and \hat{I}^k are the k -th band slices of the reference cube and the reconstruction, respectively, and \mathcal{P} denotes the set of all pixels.


We also employ the structural similarity (SSIM) index [21], which is a decimal value between -1 and 1 , with 1 being reachable only in the case of two identical sets of data. SSIM is a perception-based model that considers image degradation as perceived change in structural information, while also incorporating important perceptual phenomena, including both luminance masking and contrast masking terms.

3.3. Results


Figure 1a shows the PSNR of each spectrum given its band centre, for a sample image from Nascimento [4], for the classical interpolation methods SD and ID as well as the compressed sensing (CS) and matrix completion techniques initialised with SD. Notice that, except for the first



(a) D2301, ID [37.9 dB PSNR and SSIM of 0.99992].



(b) F06, W2×DCT from SD [39.3 dB PSNR and SSIM of 0.99955].



(c) GD, CGIHT from SD [36.2 dB PSNR and SSIM of 0.99988].

Fig. 2: Left: Colour renderings of image reconstructions. Right: MSE of reconstructions (log-scale).

band, the matrix completion algorithms outperform SD and ID. On the other hand, the compressed sensing approach does not improve on the interpolation methods. In particular, note the overall incorrect contrast level resulting in yellowing of Figure 1b. Moreover, we lose the sharpness of the edges in the balcony when employing CS W2×DCT (Figure 1b), while we recover it with the matrix completion variant of CGIHT (Figure 1c).

To further emphasise how the different algorithms differ from each other we show the reconstructions PSNR from the corresponding reference images in Figure 2. Note in Figure 2c how ID accurately reconstructs the field image, taking into account the spectral correlation between the bands, but smooths the sharp details. The compressed sensing approach does a better job in identifying the edges (Figure 2b), but suffers from the same problem overall. On the other hand, the matrix completion CGIHT outperforms the other methods due to the presence of field-like uniformities. Finally, as shown in Table 1, in the majority of cases the best performing algorithms are the matrix completion CGIHT and ASD initialised with SD, with improvements over both SD and ID from 2 dB to 5 dB. StCh from Stanford SCIEN [3] seems to

IMAGE		INIT	ASD	CGIHT	W2×DCT
D0201	ID	34.5 ± 0.8	36.3 ± 1.1	37.9 ± 1.1	34.3 ± 0.5
	SD	35.9 ± 0.7	38.1 ± 0.9	39.5 ± 1.2	35.8 ± 0.6
D0301	ID	36.0 ± 1.8	38.3 ± 2.2	39.2 ± 1.7	35.1 ± 1.1
	SD	37.3 ± 1.3	39.7 ± 1.4	40.5 ± 1.5	36.3 ± 1.0
D0303	ID	39.1 ± 1.0	41.2 ± 1.2	43.1 ± 1.4	38.8 ± 0.7
	SD	40.8 ± 0.8	43.4 ± 1.1	45.5 ± 1.3	40.5 ± 0.7
D0307	ID	34.1 ± 1.8	36.1 ± 2.3	36.9 ± 2.0	33.2 ± 1.1
	SD	35.2 ± 1.1	37.2 ± 1.2	37.8 ± 1.5	34.3 ± 0.8
D2301	ID	37.9 ± 0.3	39.4 ± 0.7	41.1 ± 1.3	37.8 ± 0.3
	SD	39.5 ± 0.3	41.5 ± 1.0	43.6 ± 1.5	39.7 ± 0.4
AvLF	ID	33.7 ± 0.7	35.6 ± 1.2	36.7 ± 2.3	33.5 ± 0.7
	SD	35.3 ± 0.7	37.8 ± 1.6	39.3 ± 2.6	34.7 ± 0.6
StCh	ID	36.2 ± 1.0	36.7 ± 1.7	36.6 ± 1.8	36.0 ± 1.0
	SD	37.3 ± 1.1	37.5 ± 1.0	37.2 ± 0.9	37.4 ± 1.0
N04	ID	37.3 ± 3.1	38.6 ± 2.6	38.2 ± 2.8	35.3 ± 1.9
	SD	35.6 ± 0.8	36.9 ± 0.8	35.4 ± 2.2	34.9 ± 0.9
N06	ID	33.3 ± 0.8	34.5 ± 0.6	35.2 ± 0.5	32.6 ± 0.4
	SD	33.5 ± 0.3	34.8 ± 0.4	35.3 ± 0.8	33.2 ± 0.2
N08	ID	33.5 ± 0.4	34.4 ± 0.6	35.6 ± 1.0	32.7 ± 0.2
	SD	33.2 ± 0.4	34.4 ± 0.7	35.5 ± 1.3	32.9 ± 0.3
F05	ID	36.1 ± 2.4	36.6 ± 1.9	36.5 ± 1.9	34.0 ± 1.4
	SD	35.0 ± 0.4	36.1 ± 0.8	35.7 ± 1.6	34.1 ± 0.7
F06	ID	40.0 ± 0.9	40.2 ± 0.8	39.9 ± 0.9	38.6 ± 0.7
	SD	38.9 ± 0.4	39.6 ± 0.9	39.3 ± 0.9	39.1 ± 0.5
F07	ID	34.9 ± 1.2	36.3 ± 1.5	36.1 ± 1.6	34.4 ± 0.9
	SD	34.6 ± 0.8	36.3 ± 1.2	35.1 ± 1.8	34.6 ± 0.8
GB	ID	34.3 ± 1.1	36.0 ± 0.9	36.7 ± 1.1	33.3 ± 0.5
	SD	34.9 ± 0.3	37.2 ± 0.9	37.5 ± 1.5	34.7 ± 0.7
GD	ID	34.1 ± 1.4	35.8 ± 1.1	36.2 ± 1.2	33.0 ± 0.5
	SD	34.6 ± 0.4	36.9 ± 0.9	36.6 ± 1.1	34.3 ± 0.7
GR	ID	35.3 ± 1.2	37.0 ± 1.0	37.3 ± 1.2	34.4 ± 0.5
	SD	36.0 ± 0.4	38.0 ± 0.9	38.1 ± 1.7	35.7 ± 0.6

Table 1: Average PSNR over the 16 bands, with standard deviations. The best results for each image are highlighted in bold.

be the only outlier, with an improvement of just 0.2 dB.

Being directly related to the number of bands, the rank of the spectral unfolding seems to be effective in capturing the spectral information of the analysed datasets. Our results suggest a high correlation between the frequency bands and a low-rank structure of the spectral unfolding of our images, in which most of the information is contained in the first 3 singular values of the spectral unfolding.

4. CONCLUSIONS

We provide a numerical comparison of multispectral demosaicing by traditional interpolation, sparse approximation and matrix completion methods. Our experiments demonstrate that non-convex matrix completion typically improves reconstruction by 2 dB to 5 dB over the current state-of-the-art methods. This differs substantially from prior work in terms of employing matrix completion on the spectral unfolding of the image in the context of demosaicing, initialising it with classical interpolation methods and using more recent non-convex matrix completion algorithms.

5. REFERENCES

- [1] W. Ma, J. M. Bioucas-Dias, J. Chanussot, and P. Gader, "Signal and image processing in hyperspectral remote sensing," *IEEE Signal Processing Magazine*, vol. 31, no. 1, pp. 22–23, 2014.
- [2] G. Vane, R. O. Green, T. G. Chrien, H. T. Enmark, E. G. Hansen, and W. M. Porter, "The airborne visible/infrared imaging spectrometer (aviris)," *Remote Sensing of Environment*, vol. 44, no. 2-3, pp. 127–143, 1993.
- [3] T. Skauli and J. Farrell, "A collection of hyperspectral images for imaging systems research," in *Proc. SPIE 8660, Digital Photography IX*, 2013, p. 86600C.
- [4] S. M. C. Nascimento, F. P. Ferreira, and D. H. Foster, "Statistics of spatial cone-excitation ratios in natural scenes," *Journal of the Optical Society of America A*, vol. 19, no. 8, pp. 1484, 2002.
- [5] D. H. Foster, S. M. C. Nascimento, and K. Amano, "Information limits on neural identification of colored surfaces in natural scenes," *Visual neuroscience*, vol. 21, no. 3, pp. 331–6, 2004.
- [6] "IEEE GRSS Data Fusion Contest," <http://www.grss-ieee.org/community/technical-committees/data-fusion/data-fusion-contest/>, 2018, [Online; accessed 29-Oct-2018].
- [7] J. Brauers and T. Aach, "A color filter array based multispectral camera," Tech. Rep., Institute of Imaging and Computer Vision, RWTH Aachen University, 2006.
- [8] S. Mihoubi, O. Losson, B. Mathon, and L. Macaire, "Multispectral demosaicing using intensity-based spectral correlation," in *2015 International Conference on Image Processing Theory, Tools and Applications (IPTA)*, 2015, pp. 461–466.
- [9] L. Miao, H. Qi, R. Ramanath, and W. E. Snyder, "Binary tree-based generic demosaicking algorithm for multispectral filter arrays," *IEEE Transactions on Image Processing*, vol. 15, no. 11, pp. 3550–3558, 2006.
- [10] S. Mihoubi, O. Losson, B. Mathon, and L. Macaire, "Multispectral demosaicing using pseudo-panchromatic image," *IEEE Transactions on Computational Imaging*, vol. 3, no. 4, pp. 982–995, 2017.
- [11] G. Kutyniok, E. King, and W. Lim, "Image inpainting: Theoretical analysis and comparison of algorithms," in *SPIE Proceedings, Wavelets and Sparsity*, 2013, vol. XV(0), pp. 1–11.
- [12] B. Dong, H. Ji, J. Li, Z. Shen, and Y. Xu, "Wavelet frame based blind image inpainting," *Applied and Computational Harmonic Analysis*, vol. 32, no. 2, pp. 268 – 279, 2012.
- [13] M. Elad, J.-L. Starck, P. Querre, and D. L. Donoho, "Simultaneous cartoon and texture image inpainting using morphological component analysis (mca)," *Applied and Computational Harmonic Analysis*, vol. 19, no. 3, pp. 340 – 358, 2005, Computational Harmonic Analysis - Part 1.
- [14] K. Degraux, V. Cambareri, L. Jacques, B. Geelen, C. Blanch, and G. Lafruit, "Generalized inpainting method for hyperspectral image acquisition," in *2015 IEEE International Conference on Image Processing (ICIP)*, 2015, pp. 315–319.
- [15] M. A. Davenport and J. Romberg, "An overview of low-rank matrix recovery from incomplete observations," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 4, pp. 608–622, 2016.
- [16] Z. Wen, W. Yin, and Y. Zhang, "Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm," *Mathematical Programming Computation*, vol. 4, no. 4, pp. 333–361, 2012.
- [17] J. Tanner and K. Wei, "Low rank matrix completion by alternating steepest descent methods," *Applied and Computational Harmonic Analysis*, vol. 40, no. 2, pp. 417–429, 2016.
- [18] J. D. Blanchard, J. Tanner, and K. Wei, "Cgiht: Conjugate gradient iterative hard thresholding for compressed sensing and matrix completion," *Information and Inference: A Journal of the IMA*, vol. 4, no. 4, pp. 289–327, 2015.
- [19] T. Gélvez, H. Arguello, and H. Rueda, "Coded aperture design for hyper-spectral image recovery via matrix completion," in *2015 20th Symposium on Signal Processing, Images and Computer Vision (STSIVA)*, 2015, pp. 1–7.
- [20] M. E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schulz, "Single-shot compressive spectral imaging with a dual-disperser architecture," *Opt. Express*, vol. 15, no. 21, pp. 14013–14027, 2007.
- [21] Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.