

Fast reconstruction of hyperspectral images from coded acquisitions using a separability assumption

Elizabeth Hemsley, Ibrahim Ardi, Tony Rouvier, Simon Lacroix, Hervé

Carfantan, Antoine Monmayrant

► To cite this version:

Elizabeth Hemsley, Ibrahim Ardi, Tony Rouvier, Simon Lacroix, Hervé Carfantan, et al.. Fast reconstruction of hyperspectral images from coded acquisitions using a separability assumption. Optics Express, 2022, 30 (5), pp.8174-8185. 10.1364/OE.448893 . hal-03610209v2

HAL Id: hal-03610209 https://hal.science/hal-03610209v2

Submitted on 28 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fast Reconstruction of Hyperspectral Images from Coded Acquisitions using a Separability Assumption

July 28, 2022

Elizabeth Hemsley^{1,*}, Ibrahim Ardi^{1,2}, Tony Rouvier², Simon Lacroix¹, Hervé Carfantan², An toine Monmayrant¹

7 ¹LAAS-CNRS, Université de Toulouse, CNRS, 7 Avenue du Colonel Roche, 31400 Toulouse, France

⁸ ²IRAP, Université de Toulouse, CNRS, CNES, 14 Avenue Édouard Belin, 31400 Toulouse, France

⁹ *ehemsley@laas.fr

2

3

10

Abstract

We present a fast reconstruction algorithm for hyperspectral images, utilizing a small 11 amount of data without the need for any training. The method is implemented with a dual 12 disperser hyperspectral imager, and makes use of spatial-spectral correlations by a so-called 13 separability assumption which assumes that the image is made of regions of homogenous spec-14 tra. The reconstruction algorithm is simple and ready-to-use, and does not require any prior 15 knowledge of the scene. A simple proof-of-principle experiment is performed, demonstrating 16 that only a small number of acquisitions are required, and the resulting compressed data-cube 17 is reconstructed near instantaneously. 18

19 1 Introduction

Hyperspectral (HS) imagers provide precise spectral information for every pixel in a scene, and 20 are useful for a range of applications - e.g. deep space or earth observation, material analysis, gas 21 detection, etc. The time taken to obtain the 3D HS data-cube can be long, particularly for scanning 22 type imagers, necessitating several acquisitions with a 2D sensor, and the volume of data produced 23 is large. Often much of the information in the HS datacube is redundant, given that a real world 24 scene contains a considerable amount of spectral-spatial correlations [1, 2]. These correlations lead 25 to the field of compressed sensing, which reduces the amount of data needed to reconstruct the HS 26 datacube, based on the assumption of an underlying sparse dictionary [3, 4, 5, 6, 7]. More recently, 27 neural networks have been trained to reconstruct a HS cube using a similarly small amount of 28 data [8, 9, 10]. But these methods either require significant time during the reconstruction, or a 29 learning phase which make them not straightforward and ready-to-use. In this paper we present 30 an alternative method requiring no training or prerequisite knowledge of the scene. The amount 31 of data required for this algorithm is similarly small, and the reconstruction time is very short. 32 A fundamental point is that we do not attempt to reconstruct the entire hyperspectral datacube, 33 that is to reconstruct the spectrum for each pixel of the scene. The key idea of our approach is to 34 observe that nearby pixels often share the same spectrum with only a small panchromatic intensity 35

difference, for example due to slightly different lighting conditions. If two pixels have differing spectra, the difference in the panchromatic intensity will likely be more significant. Therefore the scene can be divided into spectrally homogenous regions using information from the panchromatic image, and we need only assign a spectrum to each region.

Paper outline: the next section depicts how the reconstruction problem for a dual disperser HS imager is given by a least square minimization, augmented with a spectral regularization. Section 3 presents the experimental setup, how the scene ground truth is recovered, and the metric used to asses the performance of our approach. Section 4 analyzes experimental results for one simple and one more complex scene, and in Section 5 we discuss how to remove reliance on segmentation of the panchromatic image.

$_{46}$ 2 Methodology

Due to spectral-spatial correlations within a HS scene, adjacent pixels are likely to share the same spectra. Following this logic, we can separate any hyperspectral scene into spectrally homogenous regions, defined as spatially connected areas where the spectra for every pixel is the same, apart from a multiplication factor which depends on the intensity. The whole homogenous region can thus be described by one single spectrum and an intensity map given by the panchromatic intensity. We denote this assumption as the *separability assumption* (SA). If a homogenous region in the scene consists of K pixels indexed by $k \in \{1 \dots K\}$, the HS information \mathbf{o}_k of pixel k in the region is given by:

$$\mathbf{o}_k = P_k \mathbf{s},\tag{1}$$

- 47 where s is the normalized spectrum of the region made of W spectral bands, P_k the panchromatic
- ⁴⁸ intensity of that pixel, and the index k corresponds to all the pixels of the region. Therefore, the
- ⁴⁹ hyperspectral information \mathbf{o} of the whole region can be written:

$$\mathbf{o} = \begin{bmatrix} \mathbf{o}_1 \\ \mathbf{o}_2 \\ \vdots \\ \mathbf{o}_K \end{bmatrix} = \begin{bmatrix} P_1 \mathbf{s} \\ P_2 \mathbf{s} \\ \vdots \\ P_K \mathbf{s} \end{bmatrix} = \begin{bmatrix} P_1 \mathbf{I} \\ P_2 \mathbf{I} \\ \vdots \\ P_K \mathbf{I} \end{bmatrix} \mathbf{s}$$
(2)

⁵⁰ where **I** is the identity matrix of size W.

The measurement is implemented using a dual disperser HS imager [11], illustrated in Figure 1. 51 For this architecture each camera pixel is co-registered with a scene pixel, so the resulting measure-52 ment value on each camera pixel is a linear combination of the spectral bands present at the scene 53 pixel. The weights in this linear combination vary across the scene and are determined by a spatial 54 filter placed in the imaging plane between the two dispersive 4-f lines, as explained in [12]. Dual 55 disperser imagers are often used with reconfigurable spatial masks, for example digital micromirror 56 devices (DMDs), so with different mask configurations, we make N coded measurements of the 57 scene. 58

For each region in the scene, a measurement \mathbf{m}_n corresponding to the values of the camera



Figure 1: Simplified schematic of the dual disperser architecture and the programmable filtering **H** implemented by the DMD mask.

₆₀ pixels in the region can be written as

$$\mathbf{m}_{n} = \mathbf{H}_{n}\mathbf{o} = \mathbf{H}_{n}\begin{bmatrix} P_{1}\mathbf{I}\\ P_{2}\mathbf{I}\\ \vdots\\ P_{K}\mathbf{I}\end{bmatrix}\mathbf{s} = \mathbf{G}_{n}\mathbf{s}$$
(3)

where the measurement matrix \mathbf{H}_n of dimension $K \times W$ is given by the spatial filter for acquisition $n \in \{1 \dots N\}$, and the measurement \mathbf{m}_n is a vector of length K. In other words the measurement value corresponding to pixel k is a linear combination of the wavelength bands w for that pixel, with weights defined by the kth row of \mathbf{H}_n and multiplied the panchromatic intensity P_k .

 $_{65}$ For N such acquisitions we can concatenate the results, giving

$$\mathbf{m} = \begin{bmatrix} \mathbf{m}_1 \\ \mathbf{m}_2 \\ \vdots \\ \mathbf{m}_n \end{bmatrix} = \begin{bmatrix} \mathbf{H}_1 \\ \mathbf{H}_2 \\ \vdots \\ \mathbf{H}_n \end{bmatrix} \mathbf{o} = \begin{bmatrix} \mathbf{H}_1 \\ \mathbf{H}_2 \\ \vdots \\ \mathbf{H}_n \end{bmatrix} \begin{bmatrix} P_1 \mathbf{I} \\ P_2 \mathbf{I} \\ \vdots \\ P_K \mathbf{I} \end{bmatrix} \mathbf{s} = \begin{bmatrix} \mathbf{G}_1 \\ \mathbf{G}_2 \\ \vdots \\ \mathbf{G}_n \end{bmatrix} \mathbf{s} = \mathbf{G}\mathbf{s}$$
(4)

We obtain the panchromatic information P_k via one acquisition with the spatial filter completely open. The spectrum **s** for the homogeneous region can then be reconstructed using the measurements **m** and this panchromatic information P_k with a simple least squares procedure. However, as matrix **G** can be ill-conditioned, we propose to use a Tikhonov regularization along the spectral dimension [13, 14]. The estimated spectrum for the region is thus calculated by

$$\hat{\mathbf{s}} = \arg\min_{\mathbf{s}} ||\mathbf{m} - \mathbf{G}\mathbf{s}||_{\mathbf{\Gamma}^{-1}}^2 + \mu_{\lambda} ||\mathbf{D}_{\lambda}\mathbf{s}||^2$$
(5)

$$= (\mathbf{G}^{t} \mathbf{\Gamma}^{-1} \mathbf{G} + \mu_{\lambda} \mathbf{D}_{\lambda}^{t} \mathbf{D}_{\lambda})^{-1} \mathbf{G}^{t} \mathbf{\Gamma}^{-1} \mathbf{m}$$
(6)

where \mathbf{D}_{λ} is a first order difference operator, μ_{λ} is a regularization parameter and the weighting 71 matrix $\Gamma = \text{diag}\{\mathbf{m}\}\$ accounts for the noise variance which is proportional to the intensity for 72 a Poisson counting noise. As the system of Eq. (6) is overdetermined when $K \times N > W$, the 73 regularization parameter μ_{λ} can be set to zero if matrix $\mathbf{G}^{t} \mathbf{\Gamma}^{-1} \mathbf{G}$ is well-conditioned. Note that for 74 a constant regularization parameter μ_{λ} , the regularization is naturally weaker for large regions than 75 for small regions as the number of elements in the first term of Eq. (5) is larger, while the number 76 of elements in the second term (ensuring the regularization) stays constant. Finally, the value of 77 the estimated spectrum can take on negative values, which if necessary can be avoided by imposing 78 positive values using a non-negative least squares procedure, at an additional computational cost. 79 From a computational point of view, the solution can be directly computed using Eq. (6), as 80 the linear system to solve only has W unknowns. Note that solution of the linear system has a 81 constant computation cost regardless of the number of pixels K in the region and the number of 82 acquisitions N. Moreover, the computation costs of the matrix $\mathbf{G}^t \mathbf{\Gamma}^{-1} \mathbf{G}$ and the vector $\mathbf{G}^t \mathbf{\Gamma}^{-1} \mathbf{m}$ 83 are linear with respect to K and N. 84

Additionally, as each region can be computed independently, the algorithm is paralellizable, giving the potential for extremely fast reconstruction overall.

One important consideration for this algorithm is how the image can be divided into homo-87 geneous regions, given no prior knowledge of the scene. We can consider that if two pixels have 88 different spectra, there will likely be a difference in the panchromatic intensity, if we assume the 80 presence of adjacent metamers is unlikely. For two adjacent but spectrally distinct regions, the dif-90 ference in intensity is visible on the panchromatic as an edge, and thus some form of segmentation 91 algorithm applied on the panchromatic image can be used to divide the image into regions. As 92 segmentation is an ill-defined and non-trivial problem, we expect that this method will be most 93 reliable for simple scenes that are easily segmented and obey the separability assumption e.q.94 scene that is not highly textured, blurred and does not contain adjacent metamers or smooth spatial 95 changes in spectra. 96

3 Experiment

The algorithm was tested experimentally using the dual disperser system described in [12], with an 98 additional lens before the input. The scene has W = 110 spectral bands in the range 425-650 nm, 99 with 400 by 400 spatial pixels, and the spatial filtering is implemented using a DMD. Due to the 100 orientation of the DMD, as well as size mismatch between the DMD mirror and pixel, the matrix 101 **H** contains non-binary values. Whilst one would imagine that a fully 2D randomized DMD mask 102 gives the best conditioning of equation (6), in practice the geometry mismatches lead to significant 103 mixing of the data resulting in a less accurate reconstruction. To compensate for this effect, we 104 are obliged to use a simpler mask design with only a 1D variation in the direction of dispersion. 105 Consequently more acquisitions will be required than for an ideal 2D randomized mask, as there 106 is less variation in the measurement data. The mask is designed according to [15], where long 107 sequences of on or off mirrors are avoided to maximize variation in the information and improve 108 conditioning of equation (6). This paper does not perform an extensive study into the ideal ratio 109 of open mirrors (ROM), but we obtained good results with using ROM = 0.1, with some scope for 110 further optimization. The regularization parameter was chosen to be $\mu_{\lambda} = 10^3$. A typical procedure 111 to implement this algorithm consists of an acquisition of the panchromatic image, followed by one 112 or more coded acquisitions with a 1D random mask. 113

¹¹⁴ 3.1 Slit Scanning

To evaluate the reconstruction method, ideally one would compare to the underlying 'Ground Truth' of the scene. However, for an experimental system this is not possible, so instead we compare to the reconstruction obtained via a straightforward slit scanning approach.

The data for both Slit Scanning (SS) and the Separability Assumption (SA) approaches are 118 obtained with the same exposure time on the camera. The gray level on the DMD for the SS data 119 was at its maximum value of 255. As SA combines several bands, the acquired images are brighter 120 than for SS, so for the SA data the gray level of the DMD was reduced so as not to saturate the 121 camera. We chose to keep the acquisition time constant to ease the comparison between SS and 122 SA and avoid any nonlinearity at the camera level. In practice, the camera exposure time could 123 be reduced for the SA data acquisition which would provide a further bonus with a reduction in 124 terms of acquisition time. The SS reconstruction is limited by the dynamic range of the camera 125 and its accuracy is limited for darker regions within the scene, where the signal to noise ratio is low, 126 or where the signal is below the background level of the camera. Therefore, the SS data used for 127 comparison to SA combines two data sets, with slit width of 1 mirror and slit width of 2 mirrors, to 128 increase the dynamic range of the measurement. When the panchromatic intensity of the scene falls 129 below a given threshold, the acquisition data from the wider slit is used to reconstruct the spectrum 130 (further details in Supplemental Document). The reconstruction with slit width of 2 mirrors will 131 incidentally reduce the spectral resolution. 132

Similarly to compressed sensing approaches, the nature of the data acquisition also gives us the Fellgett's advantage: compared to the SS acquisition, the SA acquisition has a higher SNR in darker regions.

To compare the two reconstructions we use the Spectral Angle Mapper (SAM), which measures the similarity of two spectra independent of their amplitude:

$$SAM(r,c) = \left| \arccos\left(\frac{\mathbf{o}(r,c)^T \cdot \hat{\mathbf{o}}(r,c)}{\|\mathbf{o}(r,c)\| \cdot \|\hat{\mathbf{o}}(r,c)\|}\right) \right|$$
(7)

At a given position (r, c) in the scene, the SAM is equal to zero when the spectra are identical, and equal to $\pi/2$ when the spectra are exactly orthogonal. The SAM is an ideal metric for this approach as it measures the spectral correspondence, rather than the spatial correspondence which is easily preserved by the geometry of the system.

¹⁴⁰ 3.2 Segmentation of the panchromatic image

Many segmentation algorithms are available, the latest of state of the art resorting to neural net-141 works [16, 17, 18, 19]. However, the focus of this paper is not on the segmentation algorithm itself, 142 and we would ideally like that the method is robust enough that obtaining the exact segmentation is 143 not crucial for the quality of the results (see Section 5). In practice, the panchromatic image is seg-144 mented into homogeneous closed regions using the classical watershed algorithm [20], implemented 145 using MATLAB functions watershed. To alleviate noise, we first used an anisotropic diffusion algo-146 rithm [21], implemented using MATLAB function anisodiff2D, with gradient modulus threshold 147 $\kappa = 2 \times 10^{-3}$, and 50 iterations. Anisotropic diffusion denoises the image, while preserving the con-148 tours, followed by a morphological gradient to improve the region-oriented segmentation, removing 149 the gradients below a threshold T, which varies the segmentation sensitivity. The segmentation al-150 151 gorithm returns the detected homegenous regions in the scene, which are bounded by a contour that

is one pixel wide. The contours are likely to contain a spectral mixture of the two adjacent regions,
but to focus primarily on the separability assumption they are discarded during the reconstruction
of the HS image.

155 4 Results

¹⁵⁶ 4.1 Simple Scene

We performed initial tests on a simple and easily segmented scene, consisting of a plastic brick wall 157 illuminated with a white LED, shown in Figure 2a. The RGB image from the slit scanning (SS) 158 reconstruction is shown in Figure 2b, and from the SA algorithm in Figure 2c, for N = 5 coded 159 acquisitions, the region contour highlighted in white. The DMD gray level was 255 for the SS 160 acquisition and 140 for the SA acquisition. The panchromatic image was normalized between 0 161 and 1 and then segmented with a threshold of $T = 1 \times 10^{-2}$. Two examples of the reconstructed 162 spectra are shown in Figure 3, showing a good correspondence especially for the bright regions 163 (Figure 3a). Figure 4 displays the histogram and average of the pixel-by-pixel SAM values in the 164 scene, dependent on the number of acquisitions. As one can see, only a few acquisitions are required 165 for the results to converge to a good accuracy, with a low average SAM. 166

The reconstructed data cube has a total of Q = 360 regions, which means that the number 167 of data points required to represent the scene is $(Q \times W) + (R \times C) = 2 \times 10^5$, whereas the full 168 datacube has size to $R \times C \times W = 1.76 \times 10^7$, pixels. This gives a total compression factor of 88 169 for the hyperspectral data cube. In terms of data acquisitions, the SA method requires N + 1 = 6, 170 compared to a normal slit scan with W = 110 acquisitions (i.e. not high dynamic range), so the 171 compression factor for the acquisition data is $W/(N+1) \simeq 18$. In terms of overall acquisition time 172 for the whole sequence of images SA is 18 times faster. The reconstruction time for the SA method 173 was 15 seconds, using non-optimized Matlab code. 174

This example also illustrates well the two main failure modes of the SA; firstly via regions that are too small to contain enough information, and secondly via adjacent metamers.

For very small regions (such as the bottom left of Figure 2c), we can see on the RGB that the reconstructed spectra are inaccurate, as the reconstruction problem is under-determined, despite the regularization.

For adjacent metamers, Figure 5 shows spectra for two different pixels in the same region, although by looking at the RGB image we can see that they are actually from different colored bricks (points (c) and (d) in Figure 2b). When the scene contains two spectrally different zones, which are combined into a single region during segmentation, the reconstructed spectrum is a mixture of the spectra of the two individual zones. This is clear when we compare the spectrum constructed by SA (dotted red line in Figure 5) to the average spectral value for that region found by SS (solid red line in Figure 5).

¹⁸⁷ 4.2 Complex Scene

Now we examine the performance with a more complex scene. The scene consists also of plastic bricks but has a 3D structure, gradients in light and dark due to shadows, and specular reflections, so is difficult to segment accurately. The DMD gray level was 255 for the SS acquisition and 125 for the SA acquisition. The RGB reconstruction of the scene using the slit scanning data is shown



(a)



Figure 2: (a) Panchromatic image, and RGB image for (b) slit scanning and (c) the SA method with N = 5 acquisitions for the simple brick scene.



Figure 3: Comparison of reconstructed SA spectra to the SS spectra, for pixels labelled (a) and (b) in Figure 2c.



Figure 4: Histogram of SAM values with number of acquisitions.



Figure 5: Comparison of reconstructed SA spectra to the SS spectra for pixels (c) and (d) labeled in Figure 2c, and for the average SS spectrum over the region.

in Figure 6a. Figures 6b, 6c and 6d show the reconstructed RGB results for the SA algorithm with
 varying segmentation thresholds.

Figure 7a shows the SAM histogram and average SAM as a function of the number of acquisi-194 tions, similar to the simple scene, although now as there are more small regions the convergence 195 with N is somewhat slower. For N = 5, Figure 7b shows that the reconstruction accuracy does 196 not change much with the segmentation threshold, and there is a broad optimum region between 197 T = 0.001 and T = 0.015. For an over segmented scene we would expect more small regions which 198 do not contain enough information to be reconstructed accurately, whereas for the under segmented 199 scene we are more likely to have two different regions merged into one, so the resulting spectra is 200 a mixture of the two. Figure 7c shows how the average SAM of the region depends on the size of 201 the region K, showing a large increase in the value and the variation in SAM as the region gets 202 smaller. For a fully 2D mask we know the reconstruction problem can be over determined when 203 K > W/N = 22, i.e. for regions with more that 22 pixels we should have enough information to 204 accurately reconstruct the spectrum. However, as we are using only a 1D mask, we should expect 205 the limit on K to be somewhat higher, depending on the shape of the region. Of course, regular-206 ization over the spectrum reduces the limit, depending on the smoothness of the spectra and the 207 regularization parameter chosen. These results indicate that obtaining perfect scene segmentation 208 is not vital to the success of the algorithm. 209

For the best segmentation, the reconstructed data cube has a total of 1138 regions, which means that the number of data points required to represent the scene is $(Q \times W) + (R \times C) = 2.85 \times 10^6$, giving a total compression factor of 6.2 for the hyperspectral data cube. As for the simple scene, SA method permits a ×18 reduction in the amount of data needed for reconstruction compared to the SS method. The reconstruction time for the SA method was 11 seconds (Matlab code).



Figure 6: RGB images of the complex brick scene using (a) SS method, and (b-c) the SA method with N = 5 acquisitions and segmentation threshold (b) T = 0.008 (best segmentation in terms of average SAM), (c) T = 0.001 (over-segmented), and (d) T = 0.02 (under segmented).



Figure 7: Histogram of SAM values: (a) as a function or the number of acquisitions N for T = 0.008, and (b) as a function of the segmentation threshold for N = 5; (c) For N = 5, average SAM over a region as a function of the region size.

Region size = 10 x 10 pixels, average SAM= 0.24

Region size = 20 x 20 pixels, average SAM= 0.18



Figure 8: RGB reconstruction using SA methods on square regions or 'superpixels', region size (a) 10×10 pixels, and (b) 20×20 pixels, and N = 5 acquisitions. (c) the average SAM plotted against the residual error R of a region, for 10×10 regions.

²¹⁵ 5 Beyond Panchromatic Segmentation

As we have seen, the results are somewhat insensitive to the segmentation, and it may be possible to avoid panchromatic segmentation entirely. For example, Figures 8a and 8b show the results for the SA approach when we arbitrarily divide the scene into 10×10 or 20×20 square regions. We can gain information on how well each region is reconstructed spectrally using the residual error, given by

$$R = \left\langle \left| \frac{\mathbf{m} - \mathbf{G}\hat{\mathbf{s}}}{\mathbf{m} + \mathbf{G}\hat{\mathbf{s}}} \right| \right\rangle \tag{8}$$

Where the average $\langle . \rangle$ is over all the pixels of the region and over all the acquisitions. As Figure 8c shows, the residual gives us information on how well a region is reconstructed, although there is a large variation in the results. We could use the residual error to post-process the reconstruction and improve the segmentation and overall accuracy of the reconstruction. We could also rely on the residual error to devise a split/merge strategy to refine the SA reconstructed HS cube, starting from some arbitrary division into regions. This paves the way towards implementation of SA that do not suffer from the limitations induced by segmentation of the panchromatic image.

228 6 Conclusion

The reconstruction algorithm is based on the particular properties of the dual disperser hyperspec-229 tral imager, namely the property of co-registration, and relies on spectral-spatial correlations in the 230 scene. The algorithm reconstructs simple real-world scenes accurately and quickly using a small 231 amount of data, also with an acquisition time much shorter than for a slit scanning acquisition. 232 Full parallelization of the computation could theoretically reduce the computation time by a fac-233 tor equal to the number of regions. Experimental improvements allowing a 2D mask, as well as 234 investigation into the optimum ROM could reduce the number of acquisitions necessary or allow 235 smaller regions to be reconstructed. Currently, due to the reliance on panchromatic segmentation, 236 this method may not be well suited to challenging scenes that are difficult to segment e.g. highly 237 textured scenes or with lots of metamers. However, preliminary tests using arbitrary division of the 238 scene into regions show good results, and demonstrates that we can use the residual error as a met-239 ric on reconstruction accuracy. Future work involves avoiding reliance on imperfect segmentation 240 algorithms, instead utilizing the residual errors to improve the reconstruction and the segmentation 241 itself. 242

- Funding. This work was funded by the Agence Nationale de la Recherche (ANR-18-ASTR-0012-01
 (ImHypAd project))
- 245 Disclosures. The authors declare no conflicts of interest.

246 Data availability. Data underlying the results presented in this paper are not publicly available at this 247 time but may be obtained from the authors upon reasonable request.

248 Supplemental document. See Supplement 1 for supporting content.

249 **References**

- [1] Sérgio MC Nascimento, Kinjiro Amano, and David H Foster. Spatial distributions of local
 illumination color in natural scenes. *Vision Research*, 120:39–44, 2016.
- [2] Ayan Chakrabarti and Todd Zickler. Statistics of real-world hyperspectral images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 193–200, 07
 2011.
- [3] Ying Fu, Yinqiang Zheng, Imari Sato, and Yoichi Sato. Exploiting spectral-spatial correlation
 for coded hyperspectral image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3727–3736, 2016.
- [4] Ashwin Wagadarikar, Renu John, Rebecca Willett, and David Brady. Single disperser design for coded aperture snapshot spectral imaging. *Applied Optics*, 47(10):B44–B51, 2008.
- [5] Ashwin A Wagadarikar, Nikos P Pitsianis, Xiaobai Sun, and David J Brady. Video rate spectral
 imaging using a coded aperture snapshot spectral imager. *Optics Express*, 17(8):6368–6388,
 2009.
- [6] David Kittle, Kerkil Choi, Ashwin Wagadarikar, and David J Brady. Multiframe image estimation for coded aperture snapshot spectral imagers. *Applied Optics*, 49(36):6824–6833, 2010.
- [7] Gonzalo R Arce, David J Brady, Lawrence Carin, Henry Arguello, and David S Kittle. Compressive coded aperture spectral imaging: An introduction. *IEEE Signal Processing Magazine*, 31(1):105–115, 2013.
- [8] Ziyi Meng and Xin Yuan. Perception inspired deep neural networks for spectral snapshot
 compressive imaging. In 2021 IEEE International Conference on Image Processing (ICIP),
 pages 2813–2817, 2021.
- [9] Lizhi Wang, Tao Zhang, Ying Fu, and Hua Huang. Hyperreconnet: Joint coded aperture opti mization and image reconstruction for compressive hyperspectral imaging. *IEEE Transactions* on Image Processing, 28(5):2257–2270, 2018.
- [10] Hao Zhang, Xu Ma, Xianhong Zhao, and Gonzalo R Arce. Compressive hyperspectral image
 classification using a 3d coded convolutional neural network. *Optics Express*, 29(21):32875–
 32891, 2021.
- [11] Michael E Gehm, Renu John, David J Brady, Rebecca M Willett, and Timothy J Schulz.
 Single-shot compressive spectral imaging with a dual-disperser architecture. *Optics Express*, 15(21), 2007.
- [12] Elizabeth Hemsley, Simon Lacroix, Hervé Carfantan, and Antoine Monmayrant. Calibration
 of programmable spectral imager with dual disperser architecture. Optics Communications,
 page 125767, 2020.
- [13] Jérôme Idier. Bayesian approach to inverse problems. John Wiley & Sons, 2013.
- [14] Andrey N. Tikhonov and Vasiliy Y. Arsenin. Solutions of ill-posed problems. V. H. Winston
 & Sons, 1977.

- [15] Elizabeth Hemsley, Ibrahim Ardi, Simon Lacroix, Hervé Carfantan, and Antoine Monmayrant.
 Optimized coded aperture for frugal hyperspectral image recovery using a dual-disperser system. JOSA A, 37(12):1916–1926, 2020.
- [16] Nikhil R Pal and Sankar K Pal. A review on image segmentation techniques. Pattern Recognition, 26(9):1277–1294, 1993.
- [17] HP Narkhede. Review of image segmentation techniques. International Journal of Science and Modern Engineering, 1(8):54–61, 2013.
- [18] Catalin Amza. A review on neural network-based image segmentation techniques. De Montfort
 University, Mechanical and Manufacturing Engg., The Gateway Leicester, LE1 9BH, United
 Kingdom, pages 1–23, 2012.
- [19] Nima Tajbakhsh, Laura Jeyaseelan, Qian Li, Jeffrey N Chiang, Zhihao Wu, and Xiaowei
 Ding. Embracing imperfect datasets: A review of deep learning solutions for medical image
 segmentation. *Medical Image Analysis*, 63:101693, 2020.
- ²⁹⁹ [20] J. B. T. M. Roerdink and A. Meijster. The watershed transform: Definitions, algorithms and ³⁰⁰ parallelization strategies. *Fundamenta Informaticae*, 41(1, 2):187–228, January 2000.
- [21] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.