https://doi.org/10.2352/CIC.2023.31.1.27

An optimality property of Matrix-R theorem, its extension, and the application to hyperspectral pan-sharpening

Yi-Tun Lin, Graham D. Finlayson, and Abdullah Kucuk
University of East Anglia, Norwich, United Kingdom

Abstract

The Matrix-R decomposition teaches that, with respect to a set of sensors, any spectrum can be written as the sum of its fundamental metamer—in the subspace spanned by the spectral sensitivities of the camera—and a metameric black (orthogonal to the camera spectral sensitivities). In any RGB-to-spectra recovery algorithm we might expect a good spectral recovery to have the property that, when projected onto the RGB sensors, it equals the RGB from which it was estimated. Or, equivalently, the fundamental metamers of the ground-truth and algorithm-recovered spectra should be equal.

In this paper, we make this expectation more concreted and present an elementary proof that this "Matrix-R-compliance" post-processing step must always improve the RMSE (root-mean-squared error) accuracy of any RGB-to-hyperspectral recovery algorithms. Further, we consider and rework the proof for the case where the spectral data is known to live in a basis of small finite dimension. Experimental results are presented for three historic RGB-guided hyperspectral pan-sharpening algorithms. Here the algorithm input includes a high-spatial-resolution RGB image and a low-resolution hyperspectral counterpart. We show the evidence that the RMSE accuracy of all tested algorithms are improved by this Matrix-R-compliant process, while we also find the best recovery results from adopting a low-dimensional linear model that models the variation of spectra in a scene.

1. Introduction

An RGB camera uses 3 types of color sensor to capture radiance spectral signals coming from the scene, resulting in 3-channel RGB responses. Nowadays, this RGB imaging technique is the most common way of observing the light spectra, and yet it loses much information of the spectra by only recording 3 numbers per pixel [6]. In many practical applications, including medical imaging [21, 11], remote sensing [28, 26], food processing [7, 13, 23] and art conservation [24, 14], it is more useful to measure high-spectral-resolution signals at all pixels using the *hyperspectral cameras*.

Due to the high price tag of hyperspectral cameras and also their limited spatial resolution, mobility and temporal resolution, there are numerous works focus on recovering the hyperspectral images from the RGB images. Examples include the direct RGB-to-hyperspectral mappings (i.e., spectral reconstruction) [3, 4] and RGB-based pan-sharpening [17, 19, 18, 30, 15]. We will focus on the latter approach, where the low-spatial-resolution hyperspectral images are super-resolved using their RGB image counterparts which have much higher spatial resolution. But, unlike most recent works in pan-sharpening, in this paper we take a step back to re-investigate a fundamental question: "Given the RGB observation, what can we be *certain* about the spectral signal?"

Back in 1953, Wyszescki [29] first described that each radiance spectrum is composed of a fundamental component intrinsic to its RGB tristimulus response (later called the "fundamental metamer") and its "metameric black" which returns [0, 0, 0] for the RGB response. The formulation of how these two components are derived was proposed later by Cohen and Kappauf [9], namely the Matrix-R theorem. In Matrix-R, the fundamental metamer is derived by projecting the spectrum onto the 3-dimensional spectral subspace spanned by the spectral sensitivities of the 3 color sensors. Equally, for a given set of RGB spectral sensitivities, the fundamental metamer can be calculated from the measured RGB response of the spectrum (which does not require the spectrum directly). And, by subtracting the fundamental metamer from the spectrum we derive the metameric black component, which is provably "zero-colored" because it is perpendicular to all 3 sensors' spectral sensitivities.

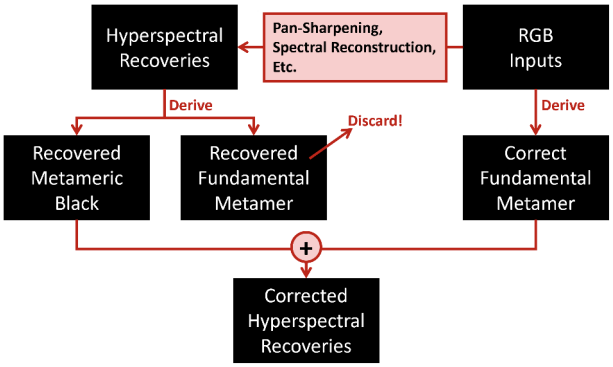

Figure 1. *Based on the Matrix-R theorem, the fundamental metamer of the ground-truth spectra can be exactly calculated from their RGBs. This flowchart shows that the Matrix-R-compliance process post-corrects the fundamental metamers of the spectra recovered by learning algorithms, by replacing them with ones calculated from the RGB inputs.*

Since the fundamental metamer of a spectrum can be uniquely derived from its RGB values, all spectra having the same RGB values, i.e., the *metamers* [9, 12], will also have the same fundamental metamer component. On the other hand, the metameric black is what distinguishes the metamers from each other and—save the requirement of returning zeros for RGB values—the metameric black is arbitrary. In other words, the Matrix-R theorem teaches that given the RGB observation of a spectrum, we can directly derive, and thus 100% certain about, its fundamental metamer. The only uncertainty lies in its metameric black component.

Indeed, in 1998, Imai and Burns [15] found that Matrix-R itself is a useful RGB pan-sharpening algorithm: by replacing the fundamental metamer at each pixel of the low-resolution hyperspectral image (resized to the size of the RGB image) by the one derived from the high-resolution RGB image, they obtained

the sharpened hyperspectral image with tolerable error. Nevertheless, thereafter most proposed pan-sharpening algorithms (e.g., [17, 19, 18, 30]) do not explicitly ensure the compliance of Matrix-R theorem, i.e., the fundamental metamers of the recovered spectra might not be the same as ones derived from the RGBs.

As shown in Figure 1, we propose that all hyperspectral recovery approaches and algorithms—not limited to spectral reconstruction and RGB-based pan-sharpening—should conduct a post-processing step where the fundamental metamers of the recoveries are replaced by ones computed from the input RGBs. Significantly, we prove that this procedure will always improve or retain the Root-Mean-Squared Errors (RMSE) between the ground-truths and recoveries, where RMSE is already a common performance metric for evaluating spectral recoveries.

Further, we rework the proof for a world where measured spectra are well modeled by a low-dimensional linear basis. Relative to this assumption, the low-dimensional-basis variant of our proof suggests an analogous "must improve" feature of the Matrix-R-compliance process: it prescribes a spectral component substitution that will always take you closer to the ground truth.

We test our theorem on several historic RGB-based pan-sharpening algorithms and show that the proposed Matrix-R-compliance post-processing procedure always—as we knew it must—improve the performance of the algorithms. Then, empirically, the best performance for each pan-sharpening algorithm is found for our low-dimensional-basis construct.

2. Background

2.1. Color image formation

The light's radiance spectrum coming from the scene can be written as a continuous spectral function $E(\lambda)$. Then, given an RGB camera, we have 3 types of color sensors with different spectral sensitivities, denoted as $Q_k(\lambda)$, such that [27]:

$$\int_{\Omega} E(\lambda)Q_k(\lambda)d\lambda = \rho_k \; ; \quad k = 1,2,3 \; . \tag{1}$$

Here, ρ_k is the camera response with respect to Q_k, and the range of integration, Ω, is the *visible range* of wavelengths, within which the RGB sensors are responsive. Ω runs roughly from 400 to 700 nanometers (nm).

Using a hyperspectral imaging device, we can measure $E(\lambda)$ at finely-sampled wavelengths. Assuming we sample n points within the 400 to 700 nm range ($n \gg 3$), we get an n-dimensional vector of measurements $\underline{\mathbf{e}} = [E(\lambda_1), E(\lambda_2), \cdots, E(\lambda_n)]^{\mathsf{T}}$. Then, we can also measure $Q_k(\lambda)$'s (i.e., the spectral sensitivity functions of the 3 color sensors) at the same discrete wavelengths using, e.g., a spectrally-scanning monochromator [16]. With these discretized measurements, we derive a vectorized form of Equation (1) [27]:

$$\mathbf{Q}^{\mathsf{T}}\underline{\mathbf{e}} = \underline{\rho} \; , \tag{2}$$

where the columns of $\mathbf{Q}$ are the discretized $Q_k(\lambda)$'s ($\mathbf{Q}$ is an $n \times 3$ matrix), and $\underline{\rho} = [\rho_1, \rho_2, \rho_3]^{\mathsf{T}}$ is the RGB vector.

2.2. Matrix-R

The so-called "Matrix-R" is in effect a matrix that projects any spectra onto the column space of $\mathbf{Q}$. In linear algebra, this *projection matrix* is written as [5]:

$$\mathbf{R} = \mathbf{Q}[\mathbf{Q}^{\mathsf{T}}\mathbf{Q}]^{-1}\mathbf{Q}^{\mathsf{T}} \; . \tag{3}$$

Using this $\mathbf{R}$ matrix, we can calculate the component of a given spectrum $\underline{\mathbf{e}}$ that lies in the column space of $\mathbf{Q}$:

$$\begin{aligned} \underline{\mathbf{e}}_{fm} &= \mathbf{R}\underline{\mathbf{e}} \\ &= \mathbf{Q}[\mathbf{Q}^{\mathsf{T}}\mathbf{Q}]^{-1}\mathbf{Q}^{\mathsf{T}}\underline{\mathbf{e}} \\ &= \mathbf{Q}[\mathbf{Q}^{\mathsf{T}}\mathbf{Q}]^{-1}\underline{\rho} \; . \end{aligned} \tag{4}$$

$\underline{\mathbf{e}}_{fm}$ is called the *fundamental metamer* of $\underline{\mathbf{e}}$. It is clear that:

1. $\mathbf{Q}^{\mathsf{T}}\underline{\mathbf{e}}_{fm} = \mathbf{Q}^{\mathsf{T}}\underline{\mathbf{e}} = \underline{\rho}$, meaning that $\underline{\mathbf{e}}_{fm}$ has the same RGB sensor response as $\underline{\mathbf{e}}$.
2. $\underline{\mathbf{e}}_{fm}$ is fixed for all spectra satisfying $\mathbf{Q}^{\mathsf{T}}\underline{\mathbf{e}} = \underline{\rho}$ (all spectra that returns the same color when observed by the camera sensitivities $\mathbf{Q}$).
3. $\underline{\mathbf{e}}_{fm}$ can be exactly calculated given camera's spectral sensitivities $\mathbf{Q}$ and the RGB sensor response $\underline{\rho}$.

On the other hand, the residual component:

$$\underline{\mathbf{e}}_{mb} = \underline{\mathbf{e}} - \underline{\mathbf{e}}_{fm} \; , \tag{5}$$

is called the *metameric black*. If we seek to calculate the color of $\underline{\mathbf{e}}_{mb}$, we get:

$$\mathbf{Q}^{\mathsf{T}}\underline{\mathbf{e}}_{mb} = \underline{\rho} - \underline{\rho} = [0,0,0]^{\mathsf{T}} \; . \tag{6}$$

In the parlance of linear algebra, we say $\underline{\mathbf{e}}_{mb}$ lies in the *null space* of the column space of $\mathbf{Q}$. And, this null space is in fact the residual $n-3$ dimensions in the spectral space that are perpendicular to the 3-dimensional camera sensor subspace spanned by columns of $\mathbf{Q}$ [5].

Unlike $\underline{\mathbf{e}}_{fm}$ which can be calculated directly from the RGB, $\underline{\mathbf{e}}_{mb}$ is *unbounded* by the color image formation. Indeed, $\underline{\mathbf{e}}_{mb}$ can be any vector in the $(n-3)$-dimensional null space of $\mathbf{Q}$ without altering the RGB observation $\underline{\rho}$. Despite the theoretical uncertainty of $\underline{\mathbf{e}}_{mb}$ (and thus $\underline{\mathbf{e}}$) given an RGB observation, in practice, we can statistically bound this uncertainty given a dataset of concern.

2.3. RGB-based pan-sharpening

In pan-sharpening, we wish to fuse hyperspectral images captured in low spatial resolution with their higher-resolution panchromatic (gray scale) or RGB counterparts (Figure 2), aiming at producing hyperspectral images with the higher spatial resolution [20]. In this paper, we consider the following three historic RGB-based pan-sharpening algorithms.

First, in Imai and Burns [15], the lower-resolution hyperspectral image is first resized (upsampled) to the same image dimension as the RGB image. Then, at each pixel, we calculate the fundamental metamer component ($\underline{\mathbf{e}}_{fm}$) from the RGB and the metameric black component ($\underline{\mathbf{e}}_{mb}$) from the hyperspectral measurement. The final spectral recovery is simply the sum of the two components. In effect, this method uses the Matrix-R theorem to correct only the part of spectra that are theoretically certain, i.e.,

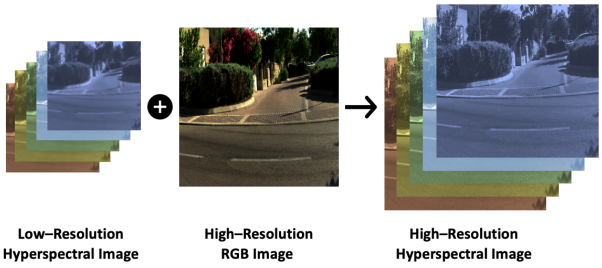

Figure 2. *The illustration of the RGB-based pan-sharpening. The images are generated from ICVL hyperspectral image database [2]. Left: the demonstration of the low-resolution hyperspectral image. Center: the high-resolution RGB image. Right: the pan-sharpened high-resolution hyperspectral image.*

fundamental metamers, without further bounding the spectral uncertainty using data.

Second, the Coupled Nonnegative Matrix Factorization (CNMF) method [30] emerged as a prominent approach in pan-sharpening. CNMF alternately unmixes the hyperspectral and multispectral (RGB) images using the nonnegative matrix factorization technique. This iterative process enables accurate estimation of *end-member* spectral signatures and high-resolution abundance maps, leading to enhanced spatial resolution in the final hyperspectral data.

Lastly, Lanaras et al. [19] combines two input images and jointly separates them into pure reflectance spectra of the observed materials and the corresponding mixing coefficients. By incorporating these physical constraints on the spectra when formulating a coupled matrix factorization problem, Lanaras et al. achieves further improved results.

3. Proposed Method

Let us denote $\underline{\mathbf{e}}$ as the ground-truth spectrum at a pixel and $\widehat{\underline{\mathbf{e}}}$ the spectral recovery from a pan-sharpening algorithm (or straight up the spectrum in the resized low-resolution hyperspectral image for Imai & Burns [15]). *A priori*, we can write both $\underline{\mathbf{e}}$ and $\widehat{\underline{\mathbf{e}}}$ as sums of fundamental metamer and metameric black:

$$\begin{cases} \underline{\mathbf{e}} = \underline{\mathbf{e}}_{fm} + \underline{\mathbf{e}}_{mb} \\ \widehat{\underline{\mathbf{e}}} = \widehat{\underline{\mathbf{e}}}_{fm} + \widehat{\underline{\mathbf{e}}}_{mb} \end{cases} . \tag{7}$$

While we do not know the ground-truth spectrum $\underline{\mathbf{e}}$ in practice, we can still calculate $\underline{\mathbf{e}}_{fm}$ from the input RGB $\underline{\rho}$, as shown in Equation (4). Nevertheless, most algorithms in data-driven pan-sharpening do not ensure $\widehat{\underline{\mathbf{e}}}_{fm} = \underline{\mathbf{e}}_{fm}$. Hence, in this paper, the following Matrix-R-compliant post-processing of $\widehat{\underline{\mathbf{e}}}$ is proposed:

$$\begin{aligned} \widehat{\underline{\mathbf{e}}}' &= \widehat{\underline{\mathbf{e}}} - \widehat{\underline{\mathbf{e}}}_{fm} + \underline{\mathbf{e}}_{fm} \\ &= \underline{\mathbf{e}}_{fm} + \widehat{\underline{\mathbf{e}}}_{mb} , \end{aligned} \tag{8}$$

where $\widehat{\underline{\mathbf{e}}}'$ is the corrected pan-sharpening output.

Theorem 1. *The adjusted output, $\widehat{\underline{\mathbf{e}}}'$, calculated using Equation (8), will always be closer to the ground truth $\underline{\mathbf{e}}$ than the initial estimate $\widehat{\underline{\mathbf{e}}}$, i.e., $||\underline{\mathbf{e}} - \widehat{\underline{\mathbf{e}}}'|| \leq ||\underline{\mathbf{e}} - \widehat{\underline{\mathbf{e}}}||$ (where $||\cdot||$ denotes the L-2 norm).*

Proof. Let us denote $\Delta = ||\underline{\mathbf{e}} - \widehat{\underline{\mathbf{e}}}||^2$ and $\Delta' = ||\underline{\mathbf{e}} - \widehat{\underline{\mathbf{e}}}'||^2$. Clearly, the theorem will be proved if we prove $\Delta' \leq \Delta$.

First, let us consider Δ with respect to the fundamental metamer and metameric black decomposition:

$$\begin{aligned} \Delta &= ||\underline{\mathbf{e}} - \widehat{\underline{\mathbf{e}}}||^2 \\ &= ||(\underline{\mathbf{e}}_{fm} + \underline{\mathbf{e}}_{mb}) - (\widehat{\underline{\mathbf{e}}}_{fm} + \widehat{\underline{\mathbf{e}}}_{mb})||^2 \\ &= ||(\underline{\mathbf{e}}_{fm} - \widehat{\underline{\mathbf{e}}}_{fm}) + (\underline{\mathbf{e}}_{mb} - \widehat{\underline{\mathbf{e}}}_{mb})||^2 \\ &= ||\underline{\mathbf{e}}_{fm} - \widehat{\underline{\mathbf{e}}}_{fm}||^2 + ||\underline{\mathbf{e}}_{mb} - \widehat{\underline{\mathbf{e}}}_{mb}||^2 \\ &\quad + 2 \cdot [\underline{\mathbf{e}}_{fm} - \widehat{\underline{\mathbf{e}}}_{fm}]^{\mathsf{T}} [\underline{\mathbf{e}}_{mb} - \widehat{\underline{\mathbf{e}}}_{mb}] . \end{aligned} \tag{9}$$

Here, the cross-term:

$$[\underline{\mathbf{e}}_{fm} - \widehat{\underline{\mathbf{e}}}_{fm}]^{\mathsf{T}} [\underline{\mathbf{e}}_{mb} - \widehat{\underline{\mathbf{e}}}_{mb}] = 0 . \tag{10}$$

Indeed, because both $\underline{\mathbf{e}}_{fm}$ and $\widehat{\underline{\mathbf{e}}}_{fm}$ lie in the spectral subspace spanned by columns of $\mathbf{Q}$, $[\underline{\mathbf{e}}_{fm} - \widehat{\underline{\mathbf{e}}}_{fm}]$ is also a vector in this subspace; on the other hand, $[\underline{\mathbf{e}}_{mb} - \widehat{\underline{\mathbf{e}}}_{mb}]$ is a vector lies in the null-space of $\mathbf{Q}$, which is perpendicular to all vectors lie in the column space of $\mathbf{Q}$ [5]. Substituting Equation (10) into Equation (9), we get:

$$\Delta = ||\underline{\mathbf{e}}_{fm} - \widehat{\underline{\mathbf{e}}}_{fm}||^2 + ||\underline{\mathbf{e}}_{mb} - \widehat{\underline{\mathbf{e}}}_{mb}||^2 . \tag{11}$$

Next, let us examine Δ':

$$\begin{aligned} \Delta' &= ||\underline{\mathbf{e}} - \widehat{\underline{\mathbf{e}}}'||^2 \\ &= ||(\underline{\mathbf{e}}_{fm} + \underline{\mathbf{e}}_{mb}) - (\underline{\mathbf{e}}_{fm} + \widehat{\underline{\mathbf{e}}}_{mb})||^2 \\ &= ||\underline{\mathbf{e}}_{mb} - \widehat{\underline{\mathbf{e}}}_{mb}||^2 . \end{aligned} \tag{12}$$

Jointly considering Equation (11) and (12), it is immediate that:

$$\begin{aligned} \Delta' &= ||\underline{\mathbf{e}}_{mb} - \widehat{\underline{\mathbf{e}}}_{mb}||^2 \\ &\leq ||\underline{\mathbf{e}}_{fm} - \widehat{\underline{\mathbf{e}}}_{fm}||^2 + ||\underline{\mathbf{e}}_{mb} - \widehat{\underline{\mathbf{e}}}_{mb}||^2 = \Delta . \end{aligned} \tag{13}$$

□

Equation (13) encapsulates succinctly that the recovered spectrum corrected via our proposed post-processing step is always as close or closer to the ground truth (compared to the original spectrum returned by any pan-sharpening algorithm).

3.1. Low-dimensional spectral representation

Suppose spectral data lies in a lower-dimensional space, we write:

$$\underline{\mathbf{e}} = \mathbf{B}\underline{\alpha} , \tag{14}$$

where $\mathbf{B}$ is an $n \times m$ basis matrix ($m \leq n$), and $\underline{\alpha}$ is a coefficient vector with length m. Here, we further *orthonormalize* the columns of $\mathbf{B}$, i.e., the columns of $\mathbf{B}$ are normalized to unit vector and orthogonalized to each other. This can be achieved in many different ways, e.g. using the Gram-Schmidt process [8].

Also note that the optimal basis matrix $\mathbf{B}$ can be calculated per image (i.e., modeling the spectral variation within the query image), from either the input low-resolution hyperspectral image or the algorithm-recovered full-resolution image (ground-truth target full-resolution spectra are not available in practice). In

Table 1. The average of the per-image mean and per-image 99-percentile (99pt) hyperspectral pan-sharpening accuracy was measured in RMSE. For each algorithm, the best results are shown in bold and underlined. §: The direct application of Matrix-R on the resized low-resolution hyperspectral image refers to Imai & Burns [15].

	Resized Only [15]		**CNMF** [30]		**Lanaras et al.** [19]	
	Mean	**99pt**	**Mean**	**99pt**	**Mean**	**99pt**
Original	0.00875	0.06274	0.00395	0.01428	0.00196	0.01054
Matrix-R§	0.00469	0.03357	0.00296	0.01142	0.00172	0.00896
3 dim	0.01625	0.06872	0.01625	0.06872	0.01625	0.06872
4 dim	**0.00388**	**0.02112**	0.00287	0.01153	0.00259	0.01085
5 dim	0.00451	0.03053	**0.00284**	**0.01105**	0.00194	0.00898
6 dim	0.00458	0.03222	0.00287	0.01115	0.00178	**0.00880**
7 dim	0.00460	0.03265	0.00289	0.01123	0.00174	0.00884
8 dim	0.00462	0.03280	0.00291	0.01128	0.00172	0.00885
9 dim	0.00462	0.03294	0.00293	0.01132	**0.00171**	0.00887
10 dim	0.00462	0.03303	0.00293	0.01133	**0.00171**	0.00889

our experiment, we set $\mathbf{B}$ as the first m singular vectors returned by Singular Value Decomposition (SVD) [25] (since it is the m-dimensional space that—over all possible m-dimensional spaces spanned by the singular vectors—accounts for the greatest percentage of the variation of the data) on the low-resolution input hyperspectral image.

Then, we point out that only the part of $\mathbf{Q}$ spanned by the basis $\mathbf{B}$ contributes to the RGB observations. Indeed, since $\underline{\mathbf{e}}$ lies in the column space of $\mathbf{B}$ (Equation (14)), the part of $\mathbf{Q}$ perpendicular to $\mathbf{B}$ will have no effects in the color image formation $\mathbf{Q}^{\mathsf{T}}\underline{\mathbf{e}}$ (Equation (2)). Let us define a new data-dependent spectral sensitivity matrix:

$$\overline{\mathbf{Q}} = \mathbf{B}\mathbf{B}^{\mathsf{T}}\mathbf{Q}\,, \tag{15}$$

where $\mathbf{B}\mathbf{B}^{\mathsf{T}}$ is the projection matrix with respect to $\mathbf{B}$ (plugging $\mathbf{B}$ into Equation (3) returns this projector because $\mathbf{B}$ has orthonormal columns). We can now replace $\mathbf{Q}$ by $\overline{\mathbf{Q}}$ in color image formation (Equation (2)):

$$\overline{\mathbf{Q}}^{\mathsf{T}}\underline{\mathbf{e}} = \mathbf{Q}^{\mathsf{T}}\mathbf{B}\mathbf{B}^{\mathsf{T}}\underline{\mathbf{e}} = \mathbf{Q}^{\mathsf{T}}\underline{\mathbf{e}} = \underline{\rho}\,. \tag{16}$$

Note that the $\mathbf{B}\mathbf{B}^{\mathsf{T}}$ projection does not alter $\underline{\mathbf{e}}$ because $\underline{\mathbf{e}}$ already lies in the column space of $\mathbf{B}$, as shown in Equation (14).

With respect to $\overline{\mathbf{Q}}$, we can formulate an alternative Matrix-R-compliance-type correction process analogous to Equation (8):

$$\widehat{\underline{\mathbf{e}}}'' = \widehat{\underline{\mathbf{e}}} - \widehat{\underline{\mathbf{e}}}^{\overline{\mathbf{Q}}}_{fm} + \underline{\mathbf{e}}^{\overline{\mathbf{Q}}}_{fm}\,. \tag{17}$$

Here, the superscript $\overline{\mathbf{Q}}$ indicates that these fundamental metamer terms are calculated following Equation (4), but using $\overline{\mathbf{Q}}$ in place of $\mathbf{Q}$.

Theorem 2. *Assuming the spectra in a scene live in an m-dimensional subspace, the adjusted output, $\widehat{\underline{\mathbf{e}}}''$, calculated using Equation (17), will always be closer to the ground truth $\underline{\mathbf{e}}$ than the initial estimate $\widehat{\underline{\mathbf{e}}}$, i.e., $||\underline{\mathbf{e}} - \widehat{\underline{\mathbf{e}}}''|| \leq ||\underline{\mathbf{e}} - \widehat{\underline{\mathbf{e}}}||$ (where $||\cdot||$ denotes the L-2 norm).*

We needn't formally prove the second theorem, as the original Matrix-R theorem does not limit us to use any particular spectral sensitivity matrix $\mathbf{Q}$ for the Matrix-R decomposition. In fact, the theorem holds for any $n \times 3$ matrix with linearly independent columns. However, for our proposed correction to work, we need to ensure the ground-truth fundamental metamer can be calculated from the RGB (which applies to both $\mathbf{Q}$ and $\overline{\mathbf{Q}}$). The additional *power* of the second theorem is its adoption of the m-dimensionality assumption. Empirically, at least, it is this assumption that leads to the best pan-sharpening results (see next section).

4. Experimental Results

We wish to evaluate the three pan-sharpening algorithms introduced in Section 2.3 with our proposed Matrix-R-compliant post-processing methods (Equation (8) and (17)). Since Imai & Burns [15] itself includes a Matrix-R process, we regard the resized low-resolution hyperspectral images as $\widehat{\underline{\mathbf{e}}}$. For CNMF [30] and Lanaras et al. [19], $\widehat{\underline{\mathbf{e}}}$ refers to the outputs from these algorithms.

To evaluate the algorithms, we use the ICVL hyperspectral image database [2] for our experiments. ICVL consists of 201 hyperspectral images of size 1300×1392 (some are slightly smaller) and 31 spectral dimensions. The spectral dimensions correspond to 10-nanometer sampling intervals between 400 and 700 nanometers. The original images are encoded in 12 bits, i.e., the maximal pixel value is 4095. We re-scale the encoding range to [0, 1] by dividing 4095 from the original pixel values.

To evaluate pan-sharpening algorithms, we are to create high-resolution RGB and low-resolution hyperspectral image pairs (as inputs to the algorithms) from the original hyperspectral image (ground-truth target).

For the RGB images, we use the CIE 1964 Color Matching Functions [10] as the camera spectral sensitivities, i.e., the $\mathbf{Q}$ matrix in our derivations. The hyperspectral-to-RGB generation is calculated per-pixel using Equation (2). The low-resolution hyperspectral images are generated by downsizing the original hyperspectral images by a factor of 8 via bilinear interpolation (analogously, it is as if we have a hyperspectral *thumbnail* 1/64 the size of the RGB image).

Experimental results are reported in Table 1. Here, we use the root-mean-squared error (RMSE) as the evaluation metric, which is defined as:

$$RMSE = \sqrt{\frac{1}{n}||\underline{\mathbf{e}} - \underline{\mathbf{e}}_{\text{rec}}||^2}\,, \tag{18}$$

where $\underline{\mathbf{e}}$ is the ground-truth and $\underline{\mathbf{e}}_{\text{rec}}$ is the recovered spectra (either

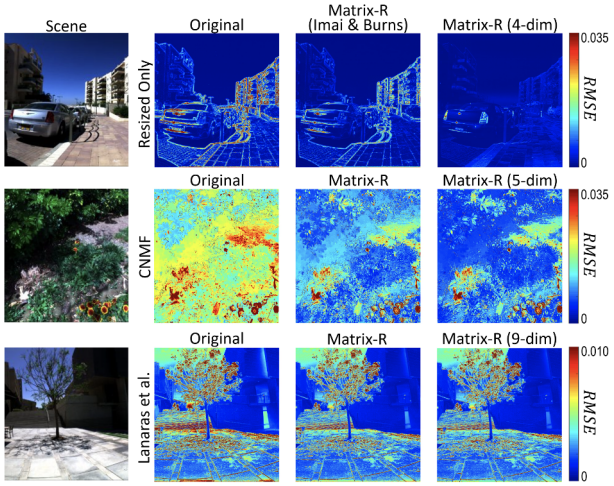

Figure 3. *The Original, Matrix-R, and Matrix-R with a lower-dimensional spectral assumption results in RMSE error heat maps for the three tested pan-sharpening algorithms.*

the original algorithm-recovered one, $\hat{\underline{\mathbf{e}}}$, or our corrected ones, $\hat{\underline{\mathbf{e}}}'$ and $\hat{\underline{\mathbf{e}}}''$). n is the length of the spectral vectors. Clearly, the RMSE metric is proportional to the L-2 distances considered in our theorems.

For each image we calculate RMSE at each pixel and record the per-image mean and 99-percentile RMSE. Then, the presented numbers in Table 1 are these two per-image statistics averaged over the whole testing image set in ICVL. For all three tested pan-sharpening algorithms, we present:

- Original: the performance of the original algorithms
- Matrix-R: performance after applying Equation (8)
- "m dim" (Matrix-R with m-dimensional spectral assumption): performance after applying Equation (17).

A visualisation of the results is shown in Figure 3.

First, let us compare the first two rows of Table 1. Evidently, we see that by applying the Matrix-R-compliant correction step, we improve both CNMF and Lanaras et al. from their original performance, in both mean and 99-percentile results. It is also shown that the direct application of Matrix-R on resized low-resolution hyperspectral images (i.e., Imai & Burns [15]) is a readily capable pan-sharpening algorithm, with 46% mean RMSE improvement from the resized images. Next, we see that Matrix-R with a 3-dimensional assumption for spectra ensures all three algorithms perform relatively poorly. In fact all approaches return the same errors. This is because the 3-dimensional assumption renders the mapping from RGBs to spectra to be one-to-one [22, 1], and the recovery performance is bounded by the extent which 3-dimensional models account for spectral variation.

We find that as we continue to increase the assumed dimensionality of the spectral data, we can always achieve results even better than the standalone Matrix-R correction. Indeed, we see that Imai & Burns with a 4-dimensional assumption delivers mean performance even better than the original performance of the much more developed CNMF algorithm. Then, while with Matrix-R correction alone we have improved the mean performance of CNMF by 25%, further with a 5-dimensional assumption, we push the improvement to 30%. Finally, even for Lanaras et al. which was originally the best-performing algorithm of the three, with a 9- or 10-dimensional assumption, we bring in 12% improvement from its original mean and 16% on the worst-case (99-percentile) performance (though compared to Matrix-R without lower-dimensional assumption the improvement is much less for this algorithm).

4.1. Discussion

Our proposed Matrix-R-compliant correction procedure can be applied in a wider context. Indeed, other RGB-to-hyperspectral recovery approaches, e.g., spectral reconstruction [3, 4], can also adopt this correction. Advantageously, the proposed process can also be used to enhance the performance of off-the-shelves "black box" algorithms where the algorithm source code is not available. Indeed, our theorems do not require the knowledge of the algorithm itself—we need only the input RGB and camera sensitivity information for the Matrix-R decomposition. Theorem 1 also gives the user *comfort*. It does no evil: it will *always* either improve the performance of any algorithm or, failing that, it will not reduce the algorithm's performance (for Theorem 2 it depends on how well a lower-dimensional assumption works for particular data).

There remain important further questions to be explored in the future work. First and foremost, it is yet to explain why adopting a lower-dimensional spectral assumption can further improve the efficacy of Matrix-R and what it means if a particular dimension returns the best result. We also observe in Table 1 that for different algorithms the optimal dimensionality assumption varies, ranging from 4 to 10 dimensions. Lastly, the power of the proposed approach will likely depend on the spectral sensitivity of the RGB camera used. This is another interesting avenue for investigation.

5. Conclusion

The Matrix-R theorem teaches that, given the RGB observation and the spectral sensitivity functions of the sensors, we can certainly calculate the *fundamental metamer* component of the ground-truth spectrum, leaving the residual *metameric black* component to be uncertain. On the other hand, hyperspectral pan-sharpening algorithms seek to super-resolve low-spatial-resolution hyperspectral images given their high-spatial-resolution RGB counterparts. Yet, these algorithms do not guarantee the exact reproduction of the fundamental metamers.

In this paper, we showed how the Matrix-R method can be used to always improve the performance of pan-sharpening and, indeed, any spectral recovery algorithm: we simply make sure that it has the correct fundamental metamer. Furthermore, we developed the Matrix-R method where spectra are represented by a low-dimensional linear model.

Experiments using three hyperspectral pan-sharpening algorithms show that our proposed Matrix-R-compliant post-processing always improved these algorithms (often significantly). In addition, the low-dimensional linear basis variant of our theorem was shown to yield the best recovery results.

References

[1] F. Agahian, S.A. Amirshahi, and S.H. Amirshahi. Reconstruction of reflectance spectra using weighted principal component analysis. *Color Research & Application*, 33(5):360–371, 2008.

[2] B. Arad and O. Ben-Shahar. Sparse recovery of hyperspectral signal from natural RGB images. In *European Conference on Computer Vision*, pages 19–34. Springer, 2016.

[3] B. Arad, R. Timofte, O. Ben-Shahar, Y. Lin, G.D. Finlayson, et al. NTIRE 2020 challenge on spectral reconstruction from an RGB image. In *Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, June 2020.

[4] B. Arad, R. Timofte, R. Yahel, N. Morag, A. Bernat, Y. Cai, J. Lin, Z. Lin, H. Wang, Y. Zhang, et al. Ntire 2022 spectral recovery challenge and data set. In *Conference on Computer Vision and Pattern Recognition*, pages 863–881, 2022.

[5] D.M. Carl. *Matrix analysis and applied linear algebra*. Society for Industrial and Applied Mathematics, 2000.

[6] A. Chakrabarti and T. Zickler. Statistics of real-world hyperspectral images. In *Conference on Computer Vision and Pattern Recognition*, pages 193–200. IEEE, 2011.

[7] Z. Chen, J. Wang, T. Wang, Z. Song, Y. Li, Y. Huang, L. Wang, and J. Jin. Automated in-field leaf-level hyperspectral imaging of corn plants using a cartesian robotic platform. *Computers and Electronics in Agriculture*, 183:105996, 2021.

[8] W. Cheney and D. Kincaid. *Linear Algebra: Theory and Applications*, volume 110. The Australian Mathematical Society, 2009.

[9] J.B. Cohen and W.E. Kappauf. Metameric color stimuli, fundamental metamers, and wyszecki's metameric blacks. *The American journal of psychology*, pages 537–564, 1982.

[10] Commission Internationale de L'eclairage. *CIE Proceedings (1964) Vienna Session, Committee Report E-1.4. 1*, 1964.

[11] L.A. Courtenay, D. González-Aguilera, S. Lagüela, S. Del Pozo, C. Ruiz-Mendez, I. Barbero-García, C. Román-Curto, J. Cañueto, C. Santos-Durán, M.E. Cardeñoso-Álvarez, et al. Hyperspectral imaging and robust statistics in non-melanoma skin cancer analysis. *Biomedical Optics Express*, 12(8):5107–5127, 2021.

[12] G.D. Finlayson and P. Morovic. Metamer sets. *Journal of the Optical Society of America A*, 22(5):810–819, 2005.

[13] V. Gomes, A. Mendes-Ferreira, and P. Melo-Pinto. Application of hyperspectral imaging and deep learning for robust prediction of sugar and ph levels in wine grape berries. *Sensors*, 21(10):3459, 2021.

[14] F. Grillini, J. Thomas, and S. George. Mixing models in close-range spectral imaging for pigment mapping in cultural heritage. In *International Colour Association Conference*, pages 372–376, 2020.

[15] F.H. Imai and R.S. Berns. High-resolution multi-spectral image archives: a hybrid approach. In *Color and imaging conference*, volume 1998, pages 224–227. Society for Imaging Science and Technology, 1998.

[16] J. Jiang, D. Liu, J. Gu, and S. Süsstrunk. What is the space of spectral sensitivity functions for digital color cameras? In *Workshop on Applications of Computer Vision*, pages 168–179. IEEE, 2013.

[17] C. Jin, L. Deng, T. Huang, and G. Vivone. Laplacian pyramid networks: A new approach for multispectral pansharpening. *Information Fusion*, 78:158–170, 2022.

[18] R. Kawakami, Y. Matsushita, J. Wright, M. Ben-Ezra, Y. Tai, and K. Ikeuchi. High-resolution hyperspectral imaging via matrix factorization. In *Conference on Computer Vision and Pattern Recognition*, pages 2329–2336. IEEE, 2011.

[19] C. Lanaras, E. Baltsavias, and K. Schindler. Hyperspectral superresolution by coupled spectral unmixing. In *International Conference on Computer Vision*, pages 3586–3594, 2015.

[20] L. Loncan, L.B. De Almeida, J.M. Bioucas-Dias, X. Briottet, J. Chanussot, N. Dobigeon, S. Fabre, W. Liao, G.A. Licciardi, M. Simoes, et al. Hyperspectral pansharpening: A review. *IEEE Geoscience and Remote Sensing Magazine*, 3(3):27–46, 2015.

[21] M. Lv, T. Chen, Y. Yang, T. Tu, N. Zhang, W. Li, and W. Li. Membranous nephropathy classification using microscopic hyperspectral imaging and tensor patch-based discriminative linear regression. *Biomedical Optics Express*, 12(5):2968–2978, 2021.

[22] L.T. Maloney and B.A. Wandell. Color constancy: a method for recovering surface spectral reflectance. *Journal of the Optical Society of America A*, 3(1):29–33, 1986.

[23] C. Pane, G. Manganiello, N. Nicastro, T. Cardi, and F. Carotenuto. Powdery mildew caused by erysiphe cruciferarum on wild rocket (diplotaxis tenuifolia): Hyperspectral imaging and machine learning modeling for non-destructive disease detection. *Agriculture*, 11(4):337, 2021.

[24] M. Picollo, C. Cucci, A. Casini, and L. Stefani. Hyper-spectral imaging technique in the cultural heritage field: New possible scenarios. *Sensors*, 20(10):2843, 2020.

[25] G. Strang. *Introduction to Linear Algebra, 5th Edition*. Wellesley, MA: Wellesley-Cambridge Press, 2016.

[26] O. Torun and S.E. Yuksel. Unsupervised segmentation of lidar fused hyperspectral imagery using pointwise mutual information. *International Journal of Remote Sensing*, 42(17):6461–6476, 2021.

[27] B.A. Wandell. The synthesis and analysis of color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (1):2–13, 1987.

[28] W. Wang, L. Ma, M. Chen, and Q. Du. Joint correlation alignment-based graph neural network for domain adaptation of multitemporal hyperspectral remote sensing images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:3170–3184, 2021.

[29] G. Wyszecki. Valenzmetrische untersuchung des zusammenhanges zwischen normaler und anomaler trichromasie. 1953.

[30] N. Yokoya, T. Yairi, and A. Iwasaki. Coupled nonnegative matrix factorization unmixing for hyperspectral and multispectral data fusion. *IEEE Transactions on Geoscience and Remote Sensing*, 50(2):528–537, 2011.

Author Biography

Yi-Tun (Ethan) Lin is a Research Associate in the Colour & Imaging Lab, University of East Anglia (UEA), UK. He received an M.Sc. degree in Colour Science from the Erasmus+ COSI programme in 2018. His main research areas are spectral image reconstruction and colour constancy.

Graham Finlayson is a Professor of Computer Science at the UEA, where he leads the Colour & Imaging Lab. Professor Finlayson is interested in 'computing how we see' and his research spans computer science (algorithms), engineering (embedded systems) and psychophysics (visual perception). Significantly, some of Graham's research is implemented and used in commercial products including photo processing software, dedicated image processing hardware and in embedded camera software.

Abdullah Kucuk received an M.Sc. degree in Artificial Intelligence at the University of Nottingham, UK, in 2019, and has several years of industrial experience as a software engineer. Currently, he is a Ph.D. student in Colour & Imaging Lab at the UEA. His research interests include colour correction, hyperspectral imaging and artificial intelligence.