

Article

# Matrix-R Theory: A Simple Generic Method to Improve RGB-Guided Spectral Recovery Algorithms <sup>†</sup>

Graham D. Finlayson <sup>‡</sup>, Yi-Tun Lin <sup>\*‡</sup> and Abdullah Kucuk

School of Computing Science, University of East Anglia, Norwich NR4 7TJ, UK

\* Correspondence: ethanx188@gmail.com

<sup>†</sup> This article is a revised and expanded version of a paper entitled “An optimality property of Matrix-R theorem, its extension, and the application to hyperspectral pan-sharpening”, which was presented at Color and Imaging Conference (CIC31), Paris, France, 13–17 November 2023.<sup>‡</sup> Graham D. Finlayson and Yi-Tun Lin contributed equally to this work. Author order was determined alphabetically.

## Abstract

RGB-guided spectral recovery algorithms include both spectral reconstruction (SR) methods that map image RGBs to spectra and pan-sharpening (PS) methods, where an RGB image is used to guide the upsampling of a low-resolution spectral image. In this paper, we exploit Matrix-R theory in developing a post-processing algorithm that, when applied to the outputs of any and all spectral recovery algorithms, almost always improves their spectral recovery accuracy (and never makes it worse). In Matrix-R theory, any spectrum can be decomposed into a component—called the fundamental metamer—in the space spanned by the spectral sensitivities and a second component—the metameric black—that is orthogonal to this subspace. In our post-processing algorithm, we substitute the correct fundamental metamer, which we calculate directly from the RGB image, for the estimated (and generally incorrect) fundamental metamer that is returned by a spectral recovery algorithm. Significantly, we prove that substituting the correct fundamental metamer always reduces the recovery error. Further, if the spectra in a target application are known to be well described by a linear model of low dimension, then our Matrix-R post-processing algorithm can also exploit this additional physical constraint. In experiments, we demonstrate that our Matrix-R post-processing improves the performance of a variety of spectral reconstruction and pan-sharpening algorithms.

**Keywords:** spectral reconstruction; spectral super-resolution; pan-sharpening; spectral image fusion; Matrix-R



Academic Editor: Jean-Baptiste Thomas

Received: 3 October 2025

Revised: 5 December 2025

Accepted: 11 December 2025

Published: 17 December 2025

**Citation:** Finlayson, G.D.; Lin, Y.-T.; Kucuk, A. Matrix-R Theory: A Simple Generic Method to Improve RGB-Guided Spectral Recovery Algorithms. *Sensors* **2025**, *25*, 7662. <https://doi.org/10.3390/s25247662>

**Copyright:** © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Compared to RGB cameras where there are only three values per pixel [1], hyperspectral and multispectral cameras record more detailed spectral signatures from a scene. The additional information in a multi- or hyperspectral capture has been shown to be important in applications ranging from medical imaging [2,3], remote sensing [4,5], food processing [6–8] and art conservation [9,10]. However, the higher price tag, lower spatial resolution, longer integration time and/or bulkiness of spectral imagers limits their practical use.

There are many algorithms exploiting statistical regression and machine learning that attempt to recover high-quality spectral images from (or with the help of) the RGB images. In spectral reconstruction (SR), hyperspectral images are recovered directly from

their RGB image counterparts. Here, a ground-truth dataset of paired hyperspectral and RGB data is used to train the SR method. Example approaches include regression (pixel-based one-to-one mapping) [11–14] and deep learning-based algorithms (patch-by-patch mapping) [15–17].

In RGB pan-sharpening (PS), a low-resolution hyperspectral or multispectral image is upsampled to full resolution using a full-resolution RGB image as a guide. The term “sharpened” comes from the fact that if we naively upsampled the images (e.g., using bilinear upsampling), the spectral image would appear blurred relative to the RGB counterpart. When pan-sharpening works well, it looks like the low-resolution spectral image has been sharpened.

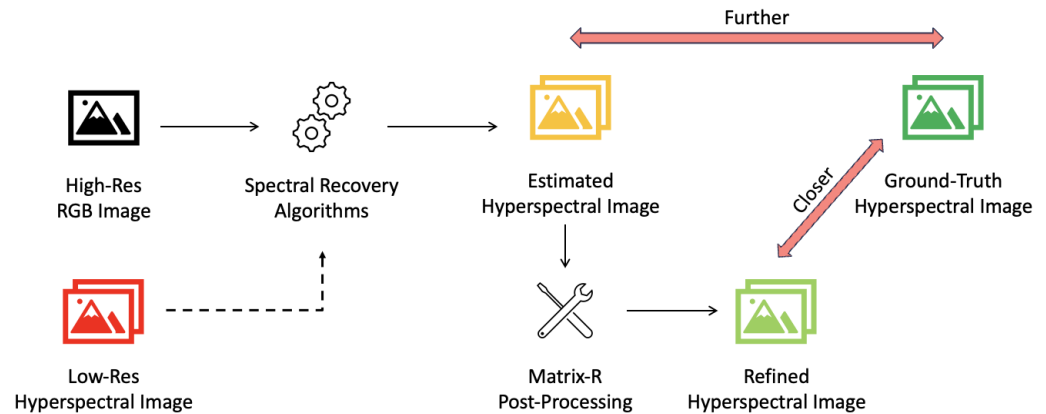
There are two variants of RGB-guided pan-sharpening. When we upsample a low-resolution hyperspectral image (where finely sampled spectra are measured at every pixel), we call it *hyperspectral* pan-sharpening. Often, however, the image we wish to upsample is still a multichannel image but with more channels than a 3-channel RGB image. In this case, we call this *multispectral* pan-sharpening. While the image data is different, the algorithms themselves can often be applied to both the hyper- and multispectral capture scenarios, e.g., [18–21]. Together, SR and PS are examples of RGB-guided spectral recovery algorithms.

Unlike most recent works in spectral recovery, in this paper, we take a step back and ask a fundamental question: “Given a recorded RGB response and assuming the camera sensitivities are known, are there fundamental properties that any recovered spectrum must adhere to?” In 1953, Wyszeski [22] 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”. The fundamental metamer integrates to the same given RGB and the black component integrates to zero RGB,  $[0, 0, 0]$  (which is where “black” comes from).

Given an RGB of a spectrum and the device spectral sensitivities, we can find the *actual* fundamental metamer defined to be the spectrum in the span of the spectral sensitivities of the camera sensors that projects to the given RGB. Then, we call the projection of a given estimated spectrum (estimated by an SR or PS algorithm) onto the same 3-dimensional spectral subspace spanned by the spectral sensitivities the *estimated* fundamental metamer. We are being careful in our definitions here as, generally, a spectral recovery algorithm—even though the actual RGB is known—will recover a spectrum where the estimated fundamental metamer is not equal to the actual fundamental metamer. One consequence of this result is that when the estimated spectrum, from a given input RGB, is numerically integrated with the camera sensitivities, the calculated output RGB will not be the same as the input [23,24]. This also means that the estimated spectrum, for most prior-art algorithms, must be the wrong answer.

As we further apply Matrix-R theory to the application of spectral recovery, we learn that a given RGB suggests its corresponding spectrum must have a unique fundamental metamer but can have different metameric blacks. This said, we would argue that the problem of spectral recovery should be about recovering the metameric black because for a given RGB, the fundamental metamer is uniquely prescribed by the, assumed known, spectral sensitivities of the camera. Yet, curiously, the vast majority of algorithms, e.g., [15,21,24–27], formulate spectral recovery as minimizing a figure of merit (e.g., RMSE) for a given dataset. And, in so doing, the individual recovered spectra can have the wrong estimated fundamental metamers. In a couple of recent works the idea that spectral reconstruction should focus on recovering the metameric black has been investigated with promising results [23,28,29]. In this paper, we show how, instead of re-architecting and retraining already deployed algorithms, we can always improve them via a simple post-processing step, as presented

in Figure 1. Here, the output of the existing spectral recovery algorithms is refined by the Matrix-R post-processing step, bringing it closer to the ground-truth hyperspectral images. Low-resolution hyperspectral images are indicated with a dashed arrow, as they are only used by the pan-sharpening algorithms; otherwise, if high-resolution RGB images alone are used, it is referred to as spectral reconstruction.



**Figure 1.** We demonstrate how the Matrix-R post-processing method works. First, either RGB images alone (in the spectral reconstruction case) or RGB images combined with low-resolution hyperspectral images (in pan-sharpening) are fed into existing spectral recovery algorithms. The output images are then enhanced by the Matrix-R post-processing algorithm. These refined images consistently achieve greater accuracy and are closer to the ground-truth compared to the initial estimates.

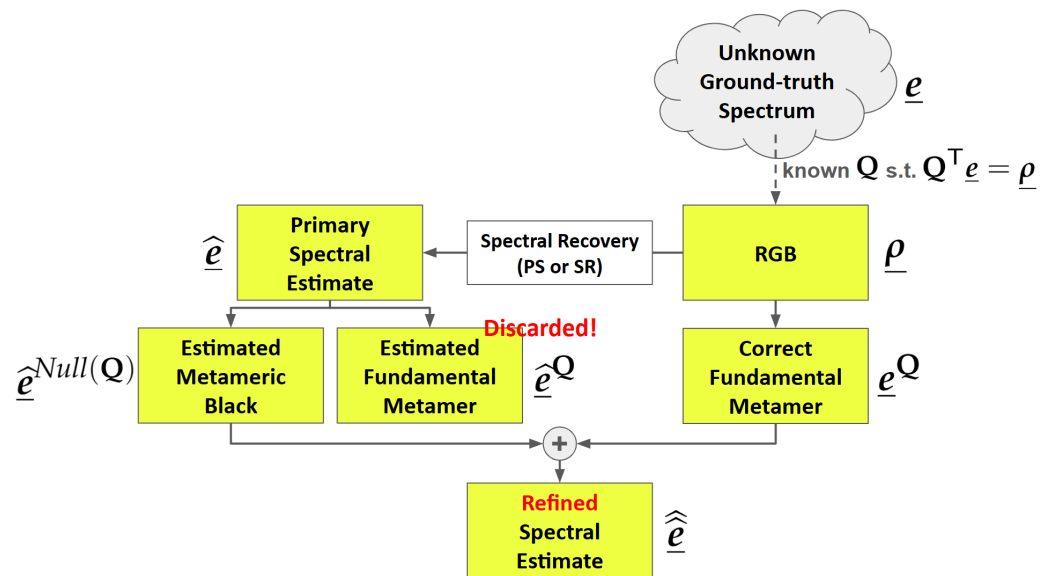
In more technical terms, Figure 2 illustrates how our post-processing method is deployed. First, we denote the targeted (but unknown) ground-truth spectrum as  $\underline{e}$ , which forms an RGB  $\underline{\rho}$  through the camera system  $\mathbf{Q}$ . Here,  $\underline{e}$  is an  $n$ -dimensional vector corresponding to the measurements made across a range of  $n$  sample wavelengths. A spectral reconstruction or pan-sharpening algorithm returns a primary estimation of the spectrum, denoted as  $\hat{\underline{e}}$ . In our method,  $\hat{\underline{e}}$  is uniquely decomposed into estimated fundamental metamer and metameric black components, respectively denoted  $\hat{\underline{e}}^{\mathbf{Q}}$  and  $\hat{\underline{e}}^{\text{Null}(\mathbf{Q})}$ . The notation  $\mathbf{Q}$  indicates that the decomposition is done with respect to the camera system  $\mathbf{Q}$ , and the metameric black component lies in the null space [30] of  $\mathbf{Q}$ .

Then, though the ground-truth  $\underline{e}$  is unknown, given the observed RGB  $\underline{\rho}$  and the camera system  $\mathbf{Q}$ , we can still derive its fundamental metamer,  $\underline{e}^{\mathbf{Q}}$  [31]. Finally, the refined spectral estimate is calculated as  $\hat{\underline{e}} = \underline{e}^{\mathbf{Q}} + \hat{\underline{e}}^{\text{Null}(\mathbf{Q})}$ . A key result of this paper is to prove that the refined estimate must be at least as close to the actual ground-truth than the primary estimate made by the SR or PS algorithms and it is, empirically, often much closer. This post-processing is generic and can be applied to all algorithms, using classical or deep learning approaches, reported in the literature.

We call our method “Matrix-R post-processing” because the algorithm illustrated in Figure 2 depends on a particular projector matrix  $\mathbf{R}$  [22,31] and the post-processing can be described in terms of simple matrix multiplications in terms of  $\mathbf{R}$ . The operation of matrix  $\mathbf{R}$  is summarized in the next section, and our post-processing algorithm and a proof of its efficacy are presented in Section 3.

We go on to develop the underlying theory when additional constraints are known about the spectra in a scene. Specifically, it is well known that spectral reflectances are smooth and are well described by low-dimensional linear models (of around 6 to 8 [32,33]). Concomitantly, the spectra, in a scene illuminated by a single dominant light, will have the same dimension as the reflectances (though they will span a different subspace as light

spectra are often not smooth). We show how we extend our method to incorporate this linear basis constraint into the Matrix-R theory.



**Figure 2.** A spectrum of light measured by a camera system  $\mathbf{Q}$  results in an RGB. In SR, an estimated spectrum is returned directly from analyzing the RGB image. In PS, low-res hyperspectral image can also guide spectral estimation. We group PS and SR algorithms in the single “spectral recovery algorithm” box. The estimated spectrum is decomposed into estimated metamer black and fundamental metamer components. Combining the correct fundamental metamer, calculated directly from the RGB [31], with this estimated metamer black returns a refined estimate of the spectrum. Refining a spectral estimate in this way is called “Matrix-R post-processing”. See the text for a description of the mathematical notation, while this figure serves as a glossary of important notations in this paper.

We empirically test our post-processing algorithm on several spectral reconstruction and spectral pan-sharpening algorithms. In all cases we improve the recovery error. When the spectra in a scene belong to a low-dimensional linear basis, our post-processing algorithm delivers an even larger reduction in the recovery error.

We also consider the case where we attempt not to recover full spectra but rather a multispectral representation (i.e., the multispectral pan-sharpening), where given an RGB guide and an  $m$  channel low-dimensional measurement of the scene (where  $m \gg 3$ ), we seek to recover the full-res  $m$ -channel multispectral image. We show how our developed theory can also be applied in this case and include a final small experimental section as a proof of concept.

This paper makes three key contributions:

- While most studies in the field of spectral recovery are dedicated to presenting new models to solve the problem, the methods in this paper are designed to improve all existing and future methods.
- Our approach is fundamental in nature—the guaranteed improvement of performance by our methods is grounded in mathematical proofs rather than empirical research.
- The methods are developed as a post-processing process, which means no change or retraining needs to be done to apply to a given algorithm, including the off-the-shelves black-box solutions. We also note that our proposed methods are *simple*; that is, it does not incur excessive processing.

## 2. Background

### 2.1. Color and Spectral Image Formation

A radiance spectrum reflected/emitted from a scene is written as the continuous spectral function  $E(\lambda)$ . 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 visible range ( $n \gg 3$ ), we get an  $n$ -dimensional vector of measurements  $\underline{e} = [E(\lambda_1), E(\lambda_2), \dots, E(\lambda_n)]^T$ . Often  $n = 31$ , where 10 nm sampling is used [15,24,25]. Here and throughout this paper,  $^T$  denotes the transpose operator.

In contrast, either an RGB or a multispectral camera uses multiple colored sensors with different spectral sensitivities. Let us denote the  $k$ th-channel spectral sensitivity as  $Q_k(\lambda)$  (a function of wavelength) [34]:

$$\int_{\Omega} E(\lambda) Q_k(\lambda) d\lambda = \rho_k. \quad (1)$$

Here,  $\rho_k$  is the  $k$ th-channel camera response depending on  $Q_k$ . For the RGB camera,  $k = 1, 2, 3$ , meaning three different color sensors are used, whereas for a multispectral camera,  $m$  sensor functions will be considered (where  $m > 3$ ). In this paper, we consider  $\Omega$ , the range of integration, to be the visible range.

A discrete variant of Equation (1) is [34]

$$\mathbf{Q}^T \underline{e} = \underline{\rho} \quad (2)$$

where the columns of  $\mathbf{Q}$  are the discretized  $Q_k(\lambda)$ 's ( $\mathbf{Q}$  is an  $n \times m$  matrix). When  $m = 3$ ,  $\underline{\rho} = [\rho_1, \rho_2, \rho_3]^T$  is an RGB vector. In the case of multispectral imaging, in this paper, we will denote the resulting camera response vector as  $\underline{c}$  instead of  $\underline{\rho}$  for distinction. The vectors  $\underline{\rho}$  and  $\underline{c}$  will always, respectively denote 3- and  $m$ -dimensional response vectors (where  $m > 3$ ).

This paper aims to study how we could improve spectral recoveries when the camera spectral sensitivities  $\mathbf{Q}$ , are "known". Therefore,  $\mathbf{Q}$  is assumed to have been measured, using techniques such as a spectral-scanning monochromator [35].

### 2.2. Matrix-R

The so-called "Matrix-R" [22,31] is the matrix that projects any spectra onto the column space of  $\mathbf{Q}$  (the projection is in the span of  $\mathbf{Q}$  and is closest in a least-squares sense). In linear algebra, this projection matrix is written as [36]

$$\mathbf{R} = \mathbf{Q}[\mathbf{Q}^T \mathbf{Q}]^{-1} \mathbf{Q}^T. \quad (3)$$

Using this matrix  $\mathbf{R}$ , we calculate the component of a given spectrum  $\underline{e}$ —the (actual) fundamental metamer  $\underline{e}^Q$ —that lies in the column space of  $\mathbf{Q}$ :

$$\begin{aligned} \underline{e}^Q &= \mathbf{R} \underline{e} \\ &= \mathbf{Q}[\mathbf{Q}^T \mathbf{Q}]^{-1} \mathbf{Q}^T \underline{e} \\ &= \mathbf{Q}[\mathbf{Q}^T \mathbf{Q}]^{-1} \underline{\rho}. \end{aligned} \quad (4)$$

It is clear that

1.  $\mathbf{Q}^T \underline{e}^Q = \mathbf{Q}^T \underline{e} = \underline{\rho}$ , meaning that  $\underline{e}^Q$  has the same RGB sensor response as  $\underline{e}$ .
2.  $\underline{e}^Q$  is fixed for all  $\underline{e}'$  (including the ground-truth  $\underline{e}$ ) satisfying  $\mathbf{Q}^T \underline{e}' = \underline{\rho}$ , i.e., all spectra that return the same color when observed by the camera sensitivities  $\mathbf{Q}$ .

3.  $\underline{e}^{\mathbf{Q}}$  can be exactly calculated given camera's spectral sensitivities  $\mathbf{Q}$  and the RGB sensor response  $\underline{\rho}$ , without the need of knowing the ground-truth  $\underline{e}$ .

Then, the residual, or metameric black, component of  $\underline{e}$  is denoted  $\underline{e}^{Null(\mathbf{Q})}$  and is calculated as

$$\underline{e}^{Null(\mathbf{Q})} = \underline{e} - \underline{e}^{\mathbf{Q}} = [\mathbf{I} - \mathbf{R}]\underline{e}. \quad (5)$$

Here,  $\mathbf{I}$  is the  $n \times n$  identity matrix. The term metameric black is used because the camera's response to this signal is zero:

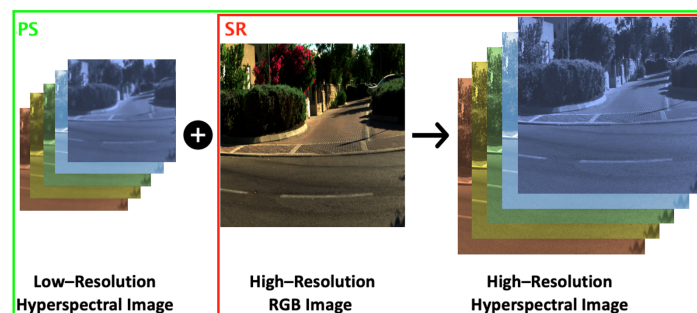
$$\mathbf{Q}^T \underline{e}^{Null(\mathbf{Q})} = [0, 0, 0]^T. \quad (6)$$

Equivalently, in the parlance of linear algebra, we say  $\underline{e}^{Null(\mathbf{Q})}$  lies in the null space of the column space of  $\mathbf{Q}$  [30] (hence the notation). 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  $\mathbf{Q}$  [36,37]. The matrix  $[\mathbf{I} - \mathbf{R}]$  is the projection matrix for this null space.

Unlike  $\underline{e}^{\mathbf{Q}}$ , which can be calculated directly from the RGB,  $\underline{e}^{Null(\mathbf{Q})}$  is unbounded by the color image formation. Indeed,  $\underline{e}^{Null(\mathbf{Q})}$  can be any vector in the  $(n - 3)$ -dimensional null space of  $\mathbf{Q}$  without altering the RGB observation  $\underline{\rho}$ .

### 2.3. Spectral Reconstruction

In spectral reconstruction (SR), high-resolution hyperspectral images are directly recovered from their RGB counterparts (Figure 3, red box). One of the simplest SR methods is linear regression [11], where a matrix transformation is found that maps the 3-dimensional RGBs to  $n$ -dimensional spectra. A simple extension to linear regression is to map each RGB to higher dimensional terms, e.g., using a polynomial expansion, [38,39]. These regression models are—in effect—a form of look-up-table since each RGB will always be mapped to the same output spectrum. Other more general one-to-one mappings reported in the literature include radial basis function regression [13], A+ sparse coding [12], and A++ sparse coding method [14]. Significantly, A++ delivers better recovery performance than many of the leading deep learning methods.



**Figure 3.** An illustration of the RGB-based hyperspectral pan-sharpening (PS; green box) and spectral reconstruction (SR; red box). The images are generated from the ICVL hyperspectral image database [40]. **Left:** The demonstration of the low-resolution hyperspectral image. **Center:** The high-resolution RGB image. **Right:** The target high-resolution hyperspectral image.

Recent SR methods are often based on machine learning and deep neural networks. Leading methods include HSCNN-D [16], HSCNN-R [16], and AWAN [17]. HSCNN-D and HSCNN-R are, respectively, the winner and runner-up of the NTIRE 2018 competition on SR [25], where the former adopted a densely-connected convolutional network [41], and the latter is based on the deep residual network architecture [42]. Then, the AWAN method is

the winner of the 2020 edition of the NTIRE competition [24], where the non-local attention mechanism [43] is incorporated.

Several classical SR methods in the literature make simplifying assumptions as an aid to recovering spectra from RGBs. Beginning with Maloney and Wandell [44], early SR approaches, e.g., Ref. [45], represented spectra with a 3-dimensional linear model. With respect to this 3-D model the spectral weights—a 3-dimensional vector—were shown to be in linear relation to the recorded RGBs. It followed that spectral recovery involved simply inverting the linear relation. Morovic and Finlayson [46] developed a Bayesian framework that is based on known spectral sensitivities and complies with Matrix-R (in the sense that Matrix-R post-processing does not further improve the estimate). According to this method, for any given input RGB, all spectra with the correct fundamental metamer are shown to form a “metamer set” that acts as a constraint for Bayesian inference. The Matrix-R post-processing was applied to spectra recovered by Zhao et al. [47] to ensure the correct fundamental metamer (without considering its optimality).

The idea that any recovered spectrum should integrate to the same (or similar) RGB has recently appeared in the deep network-based SR literature. The AWAN [17] method incorporates a color (RGB) difference term in its loss function, where the RGBs of ground-truth and reconstructed spectra are calculated via Equation (2) with the camera spectral sensitivities and compared. Despite AWAN’s effort to lower the color error, it is still not completely accurate in color, as shown in [24]. This means that the fundamental metamers of the recovered spectra by AWAN still do not match with the ones derived from the RGBs. Lin and Finlayson [23] addressed the general problem of Matrix-R non-compliance in SR by restricting all algorithms to only predict the metameric black components of the ground-truth spectra while keeping the fundamental metamer components identical to the ones derived from the RGBs. Nonetheless, this approach requires retraining of the entire algorithms, and the performance improvement is not guaranteed [23].

#### 2.4. Pan-Sharpening

In pan-sharpening, we wish to fuse low-resolution hyperspectral or multispectral images with a high-resolution RGB counterpart; see Figure 3. Here, left, a low-res hyperspectral image (or relatedly a multispectral image) is fused in some way with the full-resolution RGB image (middle) to, hopefully, produce a good estimated high-res image (right). We note that in the pan-sharpening literature (especially older algorithms), imaging systems often only have access to a grayscale full-res image instead of RGB. Here and throughout this paper, we only consider pan-sharpening problem guided by RGB input.

Interestingly, many of the prior-art methods generate their full-res outputs without formally considering whether their algorithm is accurate or not. That is because these algorithms were not trained against a benchmark ground-truth dataset. Prominent examples of this approach include the Coupled Nonnegative Matrix Factorization (CNMF) [19] and the coupled spectral unmixing methods [18]. The latter method can be seen as an improvement of the former where physical constraints are placed on the spectra that are recovered.

In more recent research, e.g., Refs. [21,48], has introduced algorithms that leverage the knowledge of ground-truth data. These methods involve the use of deep neural networks (DNNs) to map RGB images combined with low-resolution spectral input to high-resolution ground-truth outputs. Once trained, these networks can be applied to unseen data for predictions. However, one significant drawback of the DNN approach is its reliance on millions of parameters, which can make the models highly complex and computationally intensive. Moreover, there is often an insufficient amount of training data available to reliably train these networks, raising concerns about generalization [14]. Another challenge arises when new training data become available or when researchers switch to different

camera sensitivities [49]. In such cases, the entire network may need to be retrained, which can be a time-consuming and resource-intensive process.

There are hybrid methods such as the Model-Inspired Autoencoder (MIAE) [20], which are still based on finding the prior spectral and spatial prior to solve a Non-Negative Matrix Factorization problem on per-scene basis, while formulating the prior crafting as a deep learning problem. A key component of MIAE relevant to the research we report in this paper is that this algorithm exploits knowledge of the camera's spectral sensitivities. As MIAE also delivers leading performance results, we will use it here as an exemplar deep learning algorithm to benchmark against.

Finally, we note that there have been methods that incorporate knowledge of the RGB camera spectral sensitivities to refine their pan-sharpening method. Most notably, Imai and Burns [50] directly applied Matrix-R compliance as a standalone pan-sharpening algorithm: the lower-resolution hyperspectral image is first resized (upsampled) to the same image dimension as the RGB image, and then, at each pixel, we replace the fundamental metamer component of the low-resolution spectrum by the one calculated from the RGB image. Essentially, our method extends the work of Imai and Burns. We also prove that post-processing with Matrix-R must always result in improved spectral estimation for any SR or PS algorithm. This is an important point as the Imai and Burns method, viewed as the vantage point of the performance afforded by today's most effective algorithms, delivers relatively poor performance (Matrix-R alone does not suffice). Finally, we extend the Matrix-R theory so it can be more powerfully applied when we know something about the lower-dimensional spectral subspace where scene spectra lie, which is a key innovation to obtaining the best performance.

### 3. Proposed Method

#### 3.1. Matrix-R Post-Processing for Improving Hyperspectral Recovery

Let us continue to use the notation  $\underline{e}$  as the ground-truth spectrum at a pixel (measured by a hyperspectral imager) and  $\hat{\underline{e}}$  as a spectral estimate made using a PS (pan-sharpening) or SR (spectral reconstruction) algorithm. Here, we will use the convention that the overscript  $\hat{\cdot}$  denotes the primary estimation from a PS or SR algorithm. Since our Matrix-R method (see Figure 2) refines an estimate to bring an estimate closer to the actual (ground-truth) spectrum, we use the *double-hat*  $\hat{\hat{\cdot}}$  to denote the refined estimate.

A priori, we can write  $\underline{e}$  and  $\hat{\underline{e}}$  as sums of fundamental metamers and metameric blacks (see Section 2.2):

$$\begin{cases} \underline{e} = \underline{e}^{\mathbf{Q}} + \underline{e}^{\text{Null}(\mathbf{Q})} \\ \hat{\underline{e}} = \hat{\underline{e}}^{\mathbf{Q}} + \hat{\underline{e}}^{\text{Null}(\mathbf{Q})} \end{cases} \quad (7)$$

While we do not know the ground-truth spectrum  $\underline{e}$  in practice, we can still calculate  $\underline{e}^{\mathbf{Q}}$  from the input RGB  $\underline{\rho}$ , as shown in Equation (4), i.e., the *actual* unknown part of the ground-truth  $\underline{e}$  is  $\underline{e}^{\text{Null}(\mathbf{Q})}$ . Significantly, in almost all data-driven PS and SR algorithms, errors are allowed even in the fundamental metamer's part. That is,  $\hat{\underline{e}}^{\mathbf{Q}} \neq \underline{e}^{\mathbf{Q}}$ .

In this paper, we propose that using the Matrix-R theory (summarized in Figure 2) we can get a refined estimate  $\hat{\hat{\underline{e}}}$ , which is always going to be the same or *closer* to the ground-truth  $\underline{e}$  than  $\hat{\underline{e}}$  (the idea of *closer* is defined in Euclidean distance, or equivalently, lower root-mean-squared error, RMSE). In mathematical terms the refinement process is written as

$$\hat{\hat{\underline{e}}} = \hat{\underline{e}} - \hat{\underline{e}}^{\mathbf{Q}} + \underline{e}^{\mathbf{Q}}, \quad (8)$$

or equivalently,

$$\hat{\hat{\underline{e}}} = \underline{e}^{\mathbf{Q}} + \hat{\underline{e}}^{\text{Null}(\mathbf{Q})}. \quad (9)$$

**Theorem 1.** The refined output,  $\hat{\underline{e}}$ , is always as close or closer to the ground-truth  $\underline{e}$  than the initial estimate  $\hat{\underline{e}}$ , i.e.,  $\|\underline{e} - \hat{\underline{e}}\| \leq \|\underline{e} - \hat{\underline{e}}\|$  (where  $\|\cdot\|$  denotes the L-2 norm).

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

First, let us consider  $\hat{\Delta}$  with respect to the fundamental metamer and metameric black decomposition:

$$\begin{aligned}\hat{\Delta} &= \|\underline{e} - \hat{\underline{e}}\|^2 \\ &= \|(\underline{e}^{\mathbf{Q}} + \underline{e}^{\text{Null}(\mathbf{Q})}) - (\hat{\underline{e}}^{\mathbf{Q}} + \hat{\underline{e}}^{\text{Null}(\mathbf{Q})})\|^2 \\ &= \|(\underline{e}^{\mathbf{Q}} - \hat{\underline{e}}^{\mathbf{Q}}) + (\underline{e}^{\text{Null}(\mathbf{Q})} - \hat{\underline{e}}^{\text{Null}(\mathbf{Q})})\|^2 \\ &= \|\underline{e}^{\mathbf{Q}} - \hat{\underline{e}}^{\mathbf{Q}}\|^2 + \|\underline{e}^{\text{Null}(\mathbf{Q})} - \hat{\underline{e}}^{\text{Null}(\mathbf{Q})}\|^2 \\ &\quad + 2 \cdot [\underline{e}^{\mathbf{Q}} - \hat{\underline{e}}^{\mathbf{Q}}]^{\text{T}} [\underline{e}^{\text{Null}(\mathbf{Q})} - \hat{\underline{e}}^{\text{Null}(\mathbf{Q})}].\end{aligned}\tag{10}$$

Here, the cross-term is

$$[\underline{e}^{\mathbf{Q}} - \hat{\underline{e}}^{\mathbf{Q}}]^{\text{T}} [\underline{e}^{\text{Null}(\mathbf{Q})} - \hat{\underline{e}}^{\text{Null}(\mathbf{Q})}] = 0.\tag{11}$$

Indeed, because both  $\underline{e}^{\mathbf{Q}}$  and  $\hat{\underline{e}}^{\mathbf{Q}}$  lie in the spectral subspace spanned by columns of  $\mathbf{Q}$ ,  $[\underline{e}^{\mathbf{Q}} - \hat{\underline{e}}^{\mathbf{Q}}]$  is also a vector in this subspace; on the other hand,  $[\underline{e}^{\text{Null}(\mathbf{Q})} - \hat{\underline{e}}^{\text{Null}(\mathbf{Q})}]$  is a vector lies in the null-space of  $\mathbf{Q}$  [36]. Substituting Equation (11) into Equation (10), we get

$$\hat{\Delta} = \|\underline{e}^{\mathbf{Q}} - \hat{\underline{e}}^{\mathbf{Q}}\|^2 + \|\underline{e}^{\text{Null}(\mathbf{Q})} - \hat{\underline{e}}^{\text{Null}(\mathbf{Q})}\|^2.\tag{12}$$

Next, let us examine  $\hat{\hat{\Delta}}$ :

$$\begin{aligned}\hat{\hat{\Delta}} &= \|\underline{e} - \hat{\hat{\underline{e}}}\|^2 \\ &= \|(\underline{e}^{\mathbf{Q}} + \underline{e}^{\text{Null}(\mathbf{Q})}) - (\underline{e}^{\mathbf{Q}} + \hat{\hat{\underline{e}}}^{\text{Null}(\mathbf{Q})})\|^2 \\ &= \|\underline{e}^{\text{Null}(\mathbf{Q})} - \hat{\hat{\underline{e}}}^{\text{Null}(\mathbf{Q})}\|^2.\end{aligned}\tag{13}$$

Following from Equations (12) and (13), it is immediate that

$$\begin{aligned}\hat{\hat{\Delta}} &= \|\underline{e}^{\text{Null}(\mathbf{Q})} - \hat{\hat{\underline{e}}}^{\text{Null}(\mathbf{Q})}\|^2 \\ &\leq \|\underline{e}^{\mathbf{Q}} - \hat{\underline{e}}^{\mathbf{Q}}\|^2 + \|\underline{e}^{\text{Null}(\mathbf{Q})} - \hat{\underline{e}}^{\text{Null}(\mathbf{Q})}\|^2 \\ &= \hat{\Delta}.\end{aligned}\tag{14}$$

□

Equation (14) succinctly encapsulates that the recovered spectrum post-processed using the Matrix-R method is always as close or closer to the ground-truth (compared to the original recovered spectrum returned by any PS or SR algorithm).

### 3.2. Generalization of Matrix-R Post-Processing to Multispectral Recovery

In Equations (1) and (2), we proposed to sample spectra (e.g., from 400 nm to 700 nm) at  $n$  wavelengths. What if  $\mathbf{Q}$  and  $\underline{e}$  were sampled at  $2n$  or  $10n$  wavelengths? Would that change any of the arguments? No, it would not—so long as Equation (2) remains a valid physical model of how RGBs are formed. Now, suppose we think of the columns of  $\mathbf{Q}$  and  $\underline{e}$  not as discrete spectral measurements but as some other functions of wavelength that still satisfy Equation (2). Because none of our derivations depends on the physical meaning of the integration/inner product step entailed in Equations (1) and (2), all the methods developed so far continue to work. We can still find fundamental metamers and metameric blacks, which lie in the space spanned by  $\mathbf{Q}$  or in its null space, respectively. However, these metameric concepts are no longer linked to wavelength.

Let us make this abstract idea more concrete. We denote the  $m$ -dimensional measurements made by a multispectral imager ( $m > 3$ ), as  $\underline{c}$ . Now, we assume there is a linear relationship between  $\underline{c}$  and the RGB response,  $\underline{\rho}$ . Our imaging model (in direct analogy to Equation (2)) is

$$\mathbf{M}^T \underline{c} = \underline{\rho}. \quad (15)$$

It follows that for the multispectral reconstruction and pan-sharpening problem we can still apply the same Matrix-R post-processing developed thus far. All that is changed is that the matrix  $\mathbf{R}$  now depends on  $\mathbf{M}$  (rather than  $\mathbf{Q}$ ). An important detail is that  $\mathbf{M}$  is  $m \times 3$  where  $3 < m \ll n$  (a point we return to later).

Of course, the Matrix-R theory will apply if and only if Equation (15) is a good model of image formation. How might we find  $\mathbf{M}$  in practice? Let  $\mathbf{C}$  and  $\mathbf{P}$  denote, respectively,  $N \times m$  and  $N \times 3$  matrices of corresponding multispectral and RGB sensor responses for  $N$  training stimuli. We then find  $\mathbf{M}$  using a regularized least-squares regression:

$$\arg \min_{\mathbf{M}} \|\mathbf{C}\mathbf{M} - \mathbf{P}\|_F^2 + \gamma \|\mathbf{M}\|_F^2, \quad (16)$$

where  $\|\cdot\|_F$  represents the Frobenius norm [51]. The user-defined  $\gamma$ , bounds the magnitude of  $\mathbf{M}$ , effectively, mitigating overfitting [52,53]. Equation (16) is solved in closed form [52,54]:

$$\mathbf{M} = [\mathbf{C}^T \mathbf{C} + \gamma \mathbf{I}]^{-1} \mathbf{C}^T \mathbf{P}, \quad (17)$$

where  $\mathbf{I}$  is the  $m \times m$  identity matrix. The  $\gamma$  term is user-defined. We found that setting gamma to be a small fraction, say 0.01% of the mean variance of the data:  $mean(diag(\mathbf{C}^T \mathbf{C}))$ , works well for our purposes.

### 3.3. Matrix-R Post-Processing with a Low-Dimensional Spectral Representation

Now, let us suppose all spectra in a target hyperspectral image lies in a lower-dimensional space. We write

$$\underline{e} = \mathbf{B}\underline{b}, \quad (18)$$

where  $\mathbf{B}$  is an  $n \times b$  basis matrix ( $b < n$ ), and  $\underline{b}$  is a coefficient vector with  $b$  components. Here, for convenience for later derivations, we further assume the *orthonormalization* of the columns of  $\mathbf{B}$ , i.e., the columns of  $\mathbf{B}$  are normalized to unit vector and are orthogonal to each other. This can be achieved in various ways given any  $b$ -dimensional basis, e.g., using the Gram–Schmidt process [55]. To ease notation we will still write spectra as  $\underline{e}$ , where we implicitly assume the linear model assumption.

Next, the definition of fundamental metamer (and metameric black) can then be defined in terms of the interaction of the  $\mathbf{B}$  subspace and the camera spectral sensitivities (and so we will draw attention to this fact in our notation). We point out that only the part of  $\mathbf{Q}$  spanned by the basis  $\mathbf{B}$  contributes to the RGB observations of any spectra written in the form of Equation (18). Indeed, since  $\underline{e}$  lies in the column space of  $\mathbf{B}$ , the part of  $\mathbf{Q}$  perpendicular to  $\mathbf{B}$  will have no effects in the color image formation  $\mathbf{Q}^T \underline{e}$  (Equation (2)). Given this prior knowledge, we can now define a new data-dependent spectral sensitivity matrix:

$$\overline{\mathbf{Q}} = \mathbf{B}\mathbf{B}^T \mathbf{Q}, \quad (19)$$

where  $\mathbf{B}\mathbf{B}^T$  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 examine the equivalence of  $\mathbf{Q}$  and  $\overline{\mathbf{Q}}$  in color image formation by replacing  $\mathbf{Q}$  by  $\overline{\mathbf{Q}}$  in Equation (2):

$$\overline{\mathbf{Q}}^T \underline{\mathbf{e}} = \mathbf{Q}^T \mathbf{B} \mathbf{B}^T \underline{\mathbf{e}} = \mathbf{Q}^T \underline{\mathbf{e}} = \underline{\boldsymbol{\rho}}. \quad (20)$$

Here, the  $\mathbf{B} \mathbf{B}^T$  projection does not alter  $\underline{\mathbf{e}}$  because  $\underline{\mathbf{e}}$  already lies in the column space of  $\mathbf{B}$ , as shown in Equation (18). More importantly, this equivalence allows us to create a new post-processing Matrix-R method simply by following the same derivation but with  $\overline{\mathbf{Q}}$  instead of  $\mathbf{Q}$ . That is, in analogy to Equations (3) and (4), we define a modified matrix  $\overline{\mathbf{R}}$  as

$$\overline{\mathbf{R}} = \overline{\mathbf{Q}} [\overline{\mathbf{Q}}^T \overline{\mathbf{Q}}]^{-1} \overline{\mathbf{Q}}^T \quad (21)$$

and write the actual and estimated fundamental metamer,  $\underline{\mathbf{e}}^{\overline{\mathbf{Q}}}$  and  $\hat{\underline{\mathbf{e}}}^{\overline{\mathbf{Q}}}$ , respectively as

$$\begin{cases} \underline{\mathbf{e}}^{\overline{\mathbf{Q}}} = \overline{\mathbf{R}} \underline{\mathbf{e}} = \overline{\mathbf{Q}} [\overline{\mathbf{Q}}^T \overline{\mathbf{Q}}]^{-1} \underline{\boldsymbol{\rho}} \\ \hat{\underline{\mathbf{e}}}^{\overline{\mathbf{Q}}} = \overline{\mathbf{R}} \hat{\underline{\mathbf{e}}} \end{cases}. \quad (22)$$

It follows that with respect to  $\overline{\mathbf{Q}}$ —analogous to Equations (8) and (9)—our second Matrix-R post-processing algorithm with a low-dimensional data assumption is written as

$$\begin{aligned} \hat{\underline{\mathbf{e}}} &= \hat{\underline{\mathbf{e}}} - \hat{\underline{\mathbf{e}}}^{\overline{\mathbf{Q}}} + \underline{\mathbf{e}}^{\overline{\mathbf{Q}}} \\ &= \underline{\mathbf{e}}^{\overline{\mathbf{Q}}} + \hat{\underline{\mathbf{e}}}^{\text{Null}(\overline{\mathbf{Q}})}. \end{aligned} \quad (23)$$

**Theorem 2.** *Assuming that spectra are in the span of an  $m$ -dimensional linear model, the refined spectral estimate  $\hat{\underline{\mathbf{e}}}$ , calculated using Equation (23), will always be closer to the ground-truth  $\underline{\mathbf{e}}$  than the initial estimate  $\hat{\underline{\mathbf{e}}}$ , i.e.,  $\|\underline{\mathbf{e}} - \hat{\underline{\mathbf{e}}}\| \leq \|\underline{\mathbf{e}} - \hat{\underline{\mathbf{e}}}\|$ .*

We do not need to 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 that derives RGBs from spectra, which, as Equation (20) has shown, applies to both  $\mathbf{Q}$  and  $\overline{\mathbf{Q}}$ .

### 3.3.1. Determining the Basis $\mathbf{B}$

The basis  $\mathbf{B}$  can be calculated a priori, e.g., based on known reflectances and a known illuminant. For the PS application, we additionally have access to a low-res hyperspectral image of the scene, and we might extract a low-dimensional image from this image. A third alternative would be to calculate the basis from the high-res spectral reconstruction returned by a given algorithm. In all three cases, given a corpus of spectral measurements, it is easy to find the best least-squares optimal basis using techniques like characteristic vector analysis [44]. Importantly, there is a reasonable expectation that a low-dimensional basis will well describe spectral data. Indeed, most spectral reflectances are smooth functions of wavelength (and are often represented by six to eight basis functions). The dimensionality of observed smooth reflectances under a single non-smooth illuminant does not change.

### 3.3.2. Extension to Multispectral Recovery

Rather than thinking about image formation in the wavelength domain, we can instead adopt Equation (15) as our image formation model (an RGB is a linear sum of the responses from a  $m$ -sensor imager). With respect to this imager, we can again adopt a  $b$ -dimensional model for spectra. According to these assumptions, we write

$$\underline{\mathbf{c}} = \mathbf{C} \underline{\mathbf{b}}. \quad (24)$$

Here,  $\mathbf{C}$  is a  $m \times b$  orthonormal basis matrix of responses (where, to elicit any computational advantage, we need  $b < m$ ). With respect to this matrix, we can derive a new data-dependent image formation matrix:

$$\overline{\mathbf{M}} = \mathbf{C}\mathbf{C}^T\mathbf{M}. \quad (25)$$

Then, the arguments from Sections 3.2 and 3.3 all hold: we need only substitute, respectively,  $\underline{\mathbf{e}}$  for  $\underline{\mathbf{q}}$ ,  $\overline{\mathbf{M}}$  for  $\overline{\mathbf{Q}}$ , and  $\mathbf{C}\mathbf{C}^T$  for  $\mathbf{B}\mathbf{B}^T$ .

## 4. Experiments

### 4.1. Data Preparation

We will use the ICVL hyperspectral image database [40] for our experiments. ICVL consists of 201 hyperspectral images of size  $1300 \times 1392$  (though a handful of images are slightly smaller) and 31 spectral dimensions (10-nanometer sampling of the visible spectrum 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.

In our experiments, the RGB images are generated pixel by pixel from the ground-truth hyperspectral images via Equation (2). For spectral reconstruction experiments, the CIE 1964 Color Matching Functions [56] are used as the camera spectral sensitivities to generate RGB images since this is proposed in [40], and many spectral reconstruction algorithms were developed for this definition of RGB. For our hyperspectral and multispectral pan-sharpening experiments, RGB images are generated with the camera response functions of Canon 1D Mark III, and multispectral images with Spectricity's 16-channel multispectral camera sensitivity functions [57]. To generate the low-resolution spectral image input for PS algorithms, both hyperspectral and multispectral images are downsampled by a factor of 8 via bilinear interpolation for pan-sharpening experiments: we simulate a hyperspectral and multispectral *thumbnails* that are  $1/64$  the size of the original RGB image.

For both spectral reconstruction and pan-sharpening (Table 1), the root-mean-squared error (RMSE) is used as the error metric.

$$RMSE = \sqrt{\frac{1}{n} \|\underline{\mathbf{e}} - \underline{\mathbf{e}}_{\text{rec}}\|^2}, \quad (26)$$

where  $\underline{\mathbf{e}}$  is the ground-truth and  $\underline{\mathbf{e}}_{\text{rec}}$  is the recovered spectra. The recovered spectrum  $\underline{\mathbf{e}}_{\text{rec}}$  could be the primary estimate recovered by an algorithm,  $\hat{\underline{\mathbf{e}}}$ , or the refined estimate found via Matrix-R post-processing,  $\hat{\hat{\underline{\mathbf{e}}}}$ .

**Table 1.** List of considered SR and PS algorithms.

Spectral Reconstruction (SR)	Pan-Sharpening (PS)
1. A++ [14]	1. Imai and Burns [50]
2. HSCNN-R [16]	2. CNMF [19]
3. HSCNN-D [16]	3. Lanaras et al. [18]
4. AWAN [17]	4. MIAE [20]

### 4.2. Spectral Reconstruction Results

Spectral reconstruction results are summarized in Table 2. We consider four algorithms: the leading regression SR method, A++ [14], and three leading deepnets: HSCNN-R, HSCNN-D, and AWAN [16,17]. We applied the cross-validation approach [58] and reported mean and 99 percentile recovery statistics. The mean recovery error of an image is the mean error of overall image pixels. Then, the mean error shown in Table 2 is the mean of

these per image means. Similarly, the 99 percentile error is the mean of the 99 percentile recorded per image. The RMSE figures are typically small (our data is in the interval  $[0, 1]$ ), so all RMSE errors in Table 2 are multiplied by  $10^3$  for readability. In the columns in Table 2 (and all following result tables), boldface denotes the experimental condition yielding the best results.

**Table 2.** The RMSE ( $\times 10^3$ ) spectral accuracy of the Matrix-R method (“Matrix-R”) and its lower-dimensional variants (“ $m$ -dim”) on **spectral reconstruction** algorithms. The best results are shown in bold font and underlined.

	A++		HSCNN-R		HSCNN-D		AWAN	
	RMSE (Spectral) Mean	99 pt	RMSE (Spectral) Mean	99 pt	RMSE (Spectral) Mean	99 pt	RMSE (Spectral) Mean	99 pt
Original	4.10	22.26	3.98	18.83	3.57	17.77	<b><u>2.26</u></b>	12.69
Matrix-R	3.93	21.18	3.96	18.79	<b><u>3.48</u></b>	17.64	<b><u>2.26</u></b>	<b><u>12.68</u></b>
3 dim	21.24	84.07	12.35	51.92	9.51	40.09	14.53	66.55
4 dim	4.05	<b><u>19.99</u></b>	4.35	18.82	3.78	<b><u>17.37</u></b>	3.06	14.83
5 dim	4.03	20.93	4.09	18.77	3.62	17.73	2.49	13.13
6 dim	3.93	21.01	3.97	<b><u>18.73</u></b>	3.50	17.59	2.34	12.80
7 dim	3.92	21.09	3.95	18.74	3.49	17.61	2.30	12.72
8 dim	<b><u>3.91</u></b>	21.12	<b><u>3.95</u></b>	18.75	<b><u>3.48</u></b>	17.62	2.28	12.70
⋮								
21 dim	-	-	-	-	-	-	<b><u>2.26</u></b>	<b><u>12.68</u></b>

The RMSE performance of the original algorithm (no post-processing) is shown in the top row. Applying Matrix-R post-processing yields the results recorded in the second row of Table 2. We see that the performance increment is most significant for the regression-based A++, where the mean and 99 percentile error of the original method are, respectively, 4.2 and 4.9% lower. The gains for the HSCNN networks are less but are still significant. There is a very small improvement (noticeable only in the fourth decimal place) for AWAN (which is to be expected as this network was designed to approximately recover the correct fundamental metamer).

We now adopt the linear basis assumption where, per image, the best basis of a given dimension (the “ $m$ ”-dim) is found via a characteristic vector analysis of the original output spectral image recovered by the SR algorithms. Note this is a strong constraint as we are assuming we have access to the optimal linear basis that describes the spectra we are attempting to recover. Clearly, adopting too few basis vectors leads to an expected decrement in performance, as we see the generally ill-performed 3-dim results. However, in all cases—for all algorithms and error metrics—a linear model assumption exists that leads to better recovery performance.

For the A++ regression method, adopting Matrix-R post-processing together with a linear basis assumption results in, respectively, a 4.6 and 10.2% improvement in the mean and 99 percentile RMSE errors, which is a critical improvement that makes A++ outperform the much more complex HSCNN-R in mean performance.

#### 4.3. Hyperspectral Pan-Sharpener Results

Here, we have access to the full-resolution RGB image and a 1/8-resolution hyperspectral image. We wish to fuse these two images to recover a full-resolution hyperspectral image. We consider 4 algorithms. First, we only bilinearly upsample the low-resolution hyperspectral image (this is a control for our experiments). Then we benchmark against the classical algorithms: CNMF [19] and Lanaras et al. [18]. Finally, we look at the performance

of MIAE [20], one of the leading deep-net pan-sharpening algorithms. The results are summarized in Table 3.

**Table 3.** The RMSE ( $\times 10^3$ ) spectral accuracy of the Matrix-R method (“Matrix-R”) and its lower-dimensional variants (“ $m$ -dim”) on **hyperspectral pan-sharpening** algorithms. The best results are shown in bold font and underlined.

	Bilinear Upsampling RMSE (Spectral)		CNMF RMSE (Spectral)		Lanaras et al. RMSE (Spectral)		MIAE RMSE (Spectral)	
	Mean	99 pt	Mean	99 pt	Mean	99 pt	Mean	99 pt
Original	7.94	62.06	3.98	14.34	1.96	10.54	1.37	6.19
Matrix-R	3.54	26.51	2.92	11.19	1.71	8.70	1.36	6.14
3 dim	13.52	63.86	13.52	63.86	13.52	63.86	13.52	63.86
4 dim	<b><u>2.77</u></b>	<b><u>11.87</u></b>	3.04	11.59	2.54	10.32	2.52	10.55
5 dim	3.33	22.37	<b><u>2.83</u></b>	<b><u>10.72</u></b>	1.90	<b><u>8.42</u></b>	1.69	7.08
6 dim	3.33	24.02	2.85	10.91	1.76	8.49	1.46	6.32
⋮								
10 dim	3.39	25.31	2.88	11.08	<b><u>1.69</u></b>	8.62	<b><u>1.35</u></b>	<b><u>6.09</u></b>

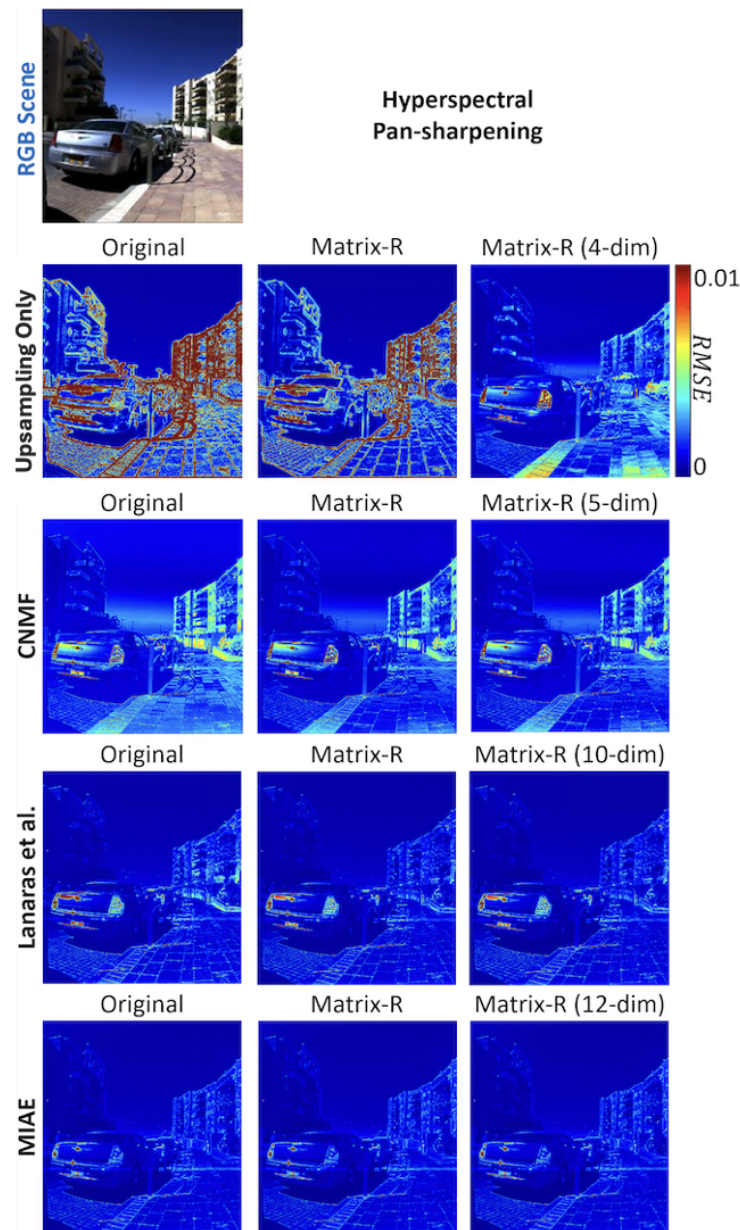
Applying our Matrix-R post-processing, for bilinear upsampling, the mean RMSE error of 7.94 is reduced to 3.54 (more than a 50% reduction), and with a 4-dimensional linear model, the mean error was further lowered to 2.77, surpassing the performance delivered by the CNMF. Matrix-R post-processing with the best performing low-dimensional linear model also significantly improves the performance of CNMF and Lanaras (their mean RMSE errors were reduced by 29 and 14%, respectively). Benchmarked against MIAE, the performance increment is much more modest. The 99 percentile error improvements follow similar trends.

Note that here, and later in the multispectral PS section, the  $m$ -dim bases are found by applying characteristic vector analysis on each input low-dimensional hyper- and multispectral image.

In Figure 4, we visualize the recovery errors for four pan-sharpening algorithms and their post-processing by the Matrix-R algorithm and the Matrix-R with the linear basis constraint. The visualizations reflect the error statistics conveyed in Tables 3 and 4. There is a large improvement for “upsampling only” and modest improvements for CNMF and Lanaras et al. For this image, it is hard to visually discern the improvement of post-processing for the MIAE algorithm.

**Table 4.** The RMSE ( $\times 10^3$ ) performance of the Matrix-R method as an **multispectral pan-sharpening** algorithm and its lower-dim variants (“ $m$ -dim”) on upsampling-only multispectral images. The best results are shown in bold font and underlined.

	Mean	99 pt
Bilinear Upsampling	7.09	52.62
Matrix-R	2.09	15.81
3 dim	32.92	167.03
4 dim	1.97	<b><u>13.93</u></b>
5 dim	<b><u>1.93</u></b>	14.54
6 dim	2.10	15.78



**Figure 4.** The Original, Matrix-R, and Matrix-R with a lower-dimensional spectral assumption results in RMSE error heat maps for the tested hyperspectral pan-sharpening algorithms.

#### 4.4. Multispectral Pan-Sharpener Results

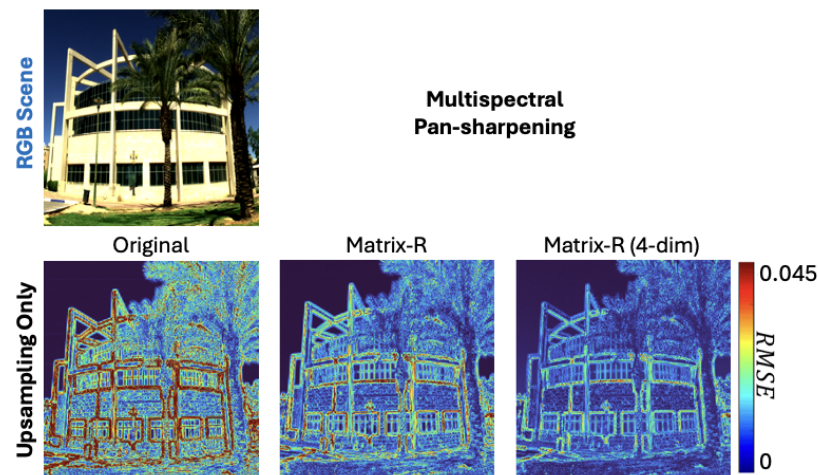
We conduct a proof-of-concept experiment on multispectral pan-sharpening, aiming to achieve high-resolution  $m$ -channel multispectral images by fusing low-resolution multispectral and corresponding high-resolution RGB images. Since this setup lacks a standard method or dataset, we perform hypothetical experiments to test Matrix-R and its lower-dimensional variants for recovering multispectral images. The RGB and multispectral images are generated by integrating the spectral images from the ICVL dataset with, respectively, the Canon 1D Mark III and Spectricity's 16-channel multispectral camera sensitivity functions [57].

Pan-sharpening for this experimental scenario is discussed in Sections 3.3 and 3.3.2, and we follow that methodology here. Importantly, to use this method, we need to know how RGBs and multispectral measurements are related to each other. The challenge is that there is no a priori, known direct mapping from multispectral to RGB (from  $\underline{c}$  to  $\underline{\rho}$  in Equation (15)), unlike in hyperspectral cases where the camera sensitivity functions directly

mapping hyperspectral data to RGB are available. Thus, we must compute the transformation matrix  $\mathbf{M}$  with regularization for multispectral pan-sharpening, i.e., Equation (17). To do that, we downscale the RGB images to match the pixels between low-resolution multispectral images for each scene. Then, we regress the multispectral images onto the RGBs and find individual matrix  $\mathbf{M}$  per image. In solving the regression in Equation (17), there is a  $\lambda$  parameter controlling the penalty term. In our experiments,  $\lambda$  equals to 0.01% of the mean variance of the data:  $\text{mean}(\text{diag}(\mathbf{C}^\top \mathbf{C}))$ .

The results are reported in Table 4, where we benchmark our multispectral pan-sharpening approach against the error found when we only bilinearly resize the multispectral image. Clearly, Matrix-R improves the upsampling-only results, and using an  $m$ -dimensional linear model consistently yields better or equivalent performance compared to standalone Matrix-R when  $m > 3$  (again, the  $m = 3$  case points to insufficient linear model representation of spectra as in the spectral reconstruction and hyperspectral pan-sharpening results). The best performance is achieved when  $m = 5$  for the mean, and  $m = 4$  for the 99th-percentile RMSE. In both cases, the errors are reduced by about 75%.

The effectiveness of our method in multispectral pan-sharpening is also illustrated in Figure 5.



**Figure 5.** The original upsampling-only, post-processing by Matrix-R, and Matrix-R with a lower-dimensional spectral assumption results in RMSE error heat maps for multispectral pan-sharpening.

## 5. Auxiliary Studies

### 5.1. Color Difference

Since the Matrix-R theorem dictates that the observed color errors result from the mismatch of fundamental metamers between ground-truth and reconstructed spectra, it is intuitive to think that the degree Matrix-R could improve the original methods should relate to the level of observed color errors. Indeed, the AWAN spectral reconstruction method [17] minimizes a color difference loss as part of its design, which also shows one of the smallest improvements when adopting Matrix-R algorithms.

In Tables 5 and 6, we show the corresponding RMSE statistics of RGB colors for the SR and PS experiments in Sections 4.2 and 4.3. Then, Figure 6 shows the correlation scatter plots between the *spectral improvement* and *color improvement*—defined as the mean-RMSE improvements in spectral and color spaces via adopting the standalone Matrix-R algorithm and denoted as  $\Delta\text{RMSE}$  (Spectral) and  $\Delta\text{RMSE}$  (RGB), respectively. Evidently, we see positive correlations between color and spectral improvements, with the pan-sharpening results, shown in the right plot of Figure 6, almost aligned in a straight line ( $R^2 = 0.995$ ). Although the spectral reconstruction (SR) result shows that it is possible that a perfect

linear correlation is not ensured, the trend is still clear that methods introducing larger color errors tend to have more benefit of adopting the Matrix-R algorithm.

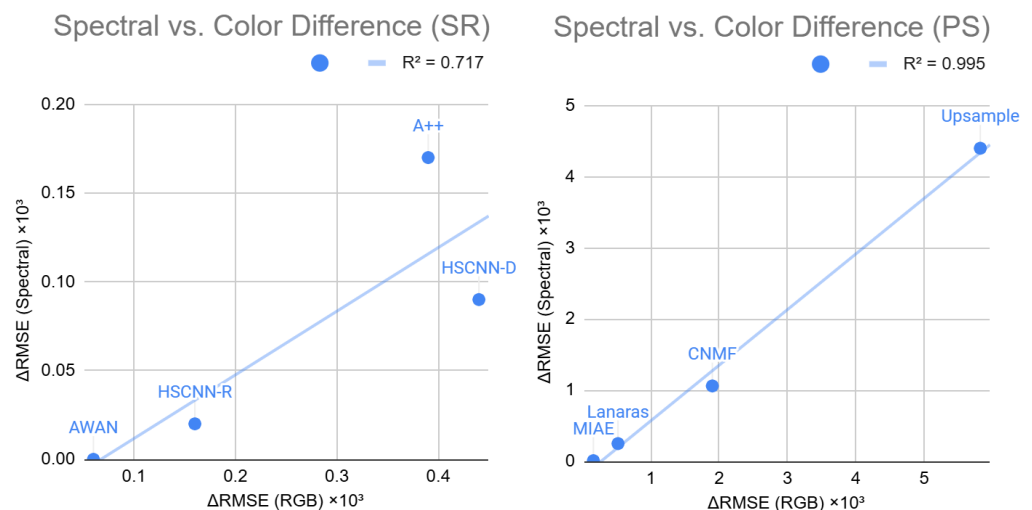
**Table 5.** The RMSE ( $\times 10^3$ ) color accuracy of the Matrix-R method (“Matrix-R”) and its lower-dimensional variants (“ $m$ -dim”) on **spectral reconstruction** algorithms.

	A++		HSCNN-R		HSCNN-D		AWAN	
	Mean	99 pt	Mean	99 pt	Mean	99 pt	Mean	99 pt
Original	0.39	3.13	0.16	0.74	0.44	1.71	0.06	0.38
Matrix-R	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
All $m$ -dims	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

**Table 6.** The RMSE ( $\times 10^3$ ) color accuracy of the Matrix-R method (“Matrix-R”) and its lower-dimensional variants (“ $m$ -dim”) on **hyperspectral pan-sharpening** algorithms.

	Bilinear Upsampling		CNMF		Lanaras et al.		MIAE	
	Mean	99 pt	Mean	99 pt	Mean	99 pt	Mean	99 pt
Original	5.82	47.92	1.90	6.82	0.52	3.57	0.16	0.96
Matrix-R	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
All $m$ -dims	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

In Tables 5 and 6, we also draw the attention to the “zero color errors” for all methods adopted lower-dimensional basis assumptions ( $m$ -dims). As we force the reconstructed spectra to lie on a particular lower-dimensional basis, we are risking that the basis might not be representative enough for the spectral data, and, subsequently, the effective spectral sensitivity matrix  $\bar{\mathbf{Q}}$  defined in Equation (19) might not be accurate in terms of color formation (Equation (20)). In contrast to this doubt, this result shows that even if the spectral basis failed to represent spectra well (in the case of  $m = 3$ , as a clear example), our low-dimensional variant of Matrix-R can still maintain perfect color fidelity.



**Figure 6.** The correlation plots of spectral and color improvements ( $\Delta\text{RMSE (Spectral)}$  and  $\Delta\text{RMSE (RGB)}$ , respectively) for spectral reconstruction (SR; **left plot**) and hyperspectral pan-sharpening (PS; **right plot**). Note that the RMSEs are scaled by  $\times 10^3$  to be consistent with the numbers in result tables.

## 5.2. CAVE Dataset

With the same experimental setup, we test the efficacy of Matrix-R methods on bilinear-upsampled images from the CAVE dataset [59]. Distinct from the ICVL dataset where

images are captured in the wild (indoor and outdoor), CAVE dataset includes 32 lab images. The mean and worst-case results are shown in Table 7. This result demonstrates our methods' generalizability over different spectral datasets where images were captured under very different settings. Indeed, the CAVE dataset results show a similar trend: while standalone Matrix-R improves the spectral accuracy of bilinear upsampling, a lower-dimensional assumption can further improve the performance in both mean and worst-case performances. It is also evident that our Matrix-R method, with or without a lower-dimensional basis in place, ensures zero color errors, consistent with the results shown in Section 5.1.

**Table 7.** The RMSE ( $\times 10^3$ ) performance of the Matrix-R method and its lower-dim variants (" $m$ -dim") on upsampling-only hyperspectral images from the CAVE dataset [59]. The best spectral accuracy results are shown in bold font and underlined.

	RMSE (Spectral)		RMSE (RGB)	
	Mean	99 pt	Mean	99 pt
Bilinear Upsampling	16.88	154.69	11.36	116.66
Matrix-R	8.53	73.68	0.00	0.00
3 dim	16.59	86.95	0.00	0.00
4 dim	11.61	<b><u>63.16</u></b>	0.00	0.00
5 dim	10.00	<u>66.47</u>	0.00	0.00
6 dim	9.20	68.84	0.00	0.00
7 dim	8.84	70.52	0.00	0.00
8 dim	8.67	71.15	0.00	0.00
9 dim	8.60	71.65	0.00	0.00
10 dim	8.54	71.94	0.00	0.00
11 dim	8.50	72.05	0.00	0.00
⋮				
16 dim	<b><u>8.43</u></b>	72.65	0.00	0.00

### 5.3. Sensitivity Analysis to Sensor Measurement Errors

A critical component of the Matrix-R framework is the reliance on the camera's spectral sensitivity matrix,  $\mathbf{Q}$ , to calculate the correct fundamental metamer. In our theoretical formulation,  $\mathbf{Q}$  is assumed to be known precisely. However, in practical applications, obtaining an exact measurement of  $\mathbf{Q}$  is challenging. There are two primary approaches to acquiring these sensitivities: direct physical measurement and indirect estimation. The direct approach involves using a monochromator to illuminate the sensor with narrow-band light across the visible spectrum, recording the response at each wavelength interval. Conversely, the indirect approach estimates the spectral sensitivities computationally from images of calibration targets (e.g., color patches) captured under known illumination [60]. Regardless of the approach, the physical characterisation of the sensor is subject to experimental uncertainties, including calibration drift, stray light, and sensor noise. Therefore, it is useful to determine the tolerance of our post-processing method to inaccuracies in the  $\mathbf{Q}$  matrix.

To investigate this, we conducted a sensitivity analysis by simulating realistic measurement errors. The magnitude of such errors might vary across different measurement setups and depends on numerous factors [61]. Consequently, we deliberately introduced random multiplicative noise to the ground-truth sensitivity matrix at two representative levels:  $\pm 5\%$  (simulating a reasonable calibration error) and  $\pm 10\%$  (representing a coarser estimation).

The results are detailed in Table 8. The first column presents the performance of the original Matrix-R post-correction using the precise camera sensitivities. The sec-

ond and third columns show the results when using the  $\pm 5\%$  and  $\pm 10\%$  perturbed sensitivities, respectively.

**Table 8.** Sensitivity analysis results for sensor measurement errors. The table reports RMSE ( $\times 10^3$ ) values for the standard Matrix-R method (“Matrix-R”) and its lower-dimensional variants (“ $m$ -dim”) applied to **hyperspectral pan-sharpening**. Performance is compared using the original camera response functions versus versions with  $\pm 5\%$  and  $\pm 10\%$  random perturbations. The best results are shown in bold font and underlined.

	Original Q		5%		10%	
	Mean	99 pt	Mean	99 pt	Mean	99 pt
Bilinear	7.94	62.06	7.94	62.06	<u>7.94</u>	62.06
Upsampling	3.54	26.51	6.62	<u>11.43</u>	11.48	33.20
Matrix-R						
4 dim	<u>2.77</u>	<u>11.87</u>	<u>6.51</u>	16.78	15.56	36.53
5 dim	3.33	22.37	6.60	24.52	12.45	31.94
6 dim	3.33	24.02	<u>6.51</u>	25.92	11.89	<u>31.70</u>

The analysis reveals that the method is moderately sensitive to calibration accuracy. With a  $\pm 5\%$  perturbation, the method still reduced the original error, although the magnitude of the improvement was approximately halved compared to the ideal case. However, at the  $\pm 10\%$  perturbation level, the method failed to improve the mean errors and, in fact, increased them. This indicates that while the Matrix-R constraint is beneficial, it requires highly accurate sensor characterisation to be effective. This requirement is particularly stringent when the method is used as a post-processing technique following advanced AI-based algorithms (although the initial estimates of models that explicitly incorporate camera sensitivities would also naturally be degraded by such measurement errors). Since these models already produce high-quality estimates, the potential margin for improvement is narrow, making the final result highly sensitive to any inaccuracies introduced by an imperfect  $\mathbf{Q}$  matrix.

#### 5.4. Incorporating Fixed General Spectral Basis

In Section 3.3.1 we suggested that, to obtain the basis  $\mathbf{B}$  in Equation (19) for a low-dimensional variant of Matrix-R, we conduct characteristic vector analysis on the reconstructed spectral image for SR and on the ground-truth lower-resolution input spectral image. Indeed, for each scene individually, this is an effective way to estimate the ground-truth spectral distribution. Nonetheless, how representative these spectra are depends on, respectively for SR and PS methods, the quality of the original spectral reconstruction and the actual resolution of the input low-resolution spectral images (imagine a 1/100 resolution difference instead of 1/8 in our experiments).

In this auxiliary study, we examine the effectiveness of lower-dimensional Matrix-R method with respect to two different methods of acquiring the basis  $\mathbf{B}$ . First, we consider the Cross-Validation (CV) setting. Under CV, we randomly separate the 201 ICVL hyperspectral images into 4 groups of 50 (or 51 for one group) images. For all images in a given group, the basis is calculated via characteristic vector analysis on all pixels of all images in the other three groups. This setting mimics a real-world application process that is to train bases based on a known set of training data and apply it directly to an unseen case. We denote the basis as  $\mathbf{B}_{K=4}$ , signifying we are adopting a  $K$ -fold cross validation with  $K = 4$ .

The second setup is more general. Assuming we have a pair of reflectance and illuminant datasets that can effectively represent the reflectances and illuminations observed in general scenes, theoretically, we can derive a basis generally for all unseen data. Here, we

use the SFU reflectance and illumination datasets [62] which consists of 1995 synthetic and natural reflectances and 102 usual illumination spectra. Note that some reflectances in the database were excluded due to the incomplete coverage of the concerned spectral range (400–700 nm), which left us with 1350 reflectances instead of 1995. Denoting reflectance as  $R(\lambda)$  and illumination spectra as  $L(\lambda)$ , we have

$$E(\lambda) = L(\lambda)R(\lambda), \quad (27)$$

The observed spectra  $E(\lambda)$  (or  $\underline{e}$  in the discrete form, i.e., the recovery target in this paper) can be derived by wavelength-by-wavelength multiplications between illumination and reflectance spectra [34]. With this equation, we form a spectral dataset by collecting all combinations of SFU reflectances and illuminations. Then, we apply characteristic vector analysis to derive the low-dimensional basis, denoted as  $\mathbf{B}_{sfu}$ .

We compare results of using  $\mathbf{B}_{K=4}$  and  $\mathbf{B}_{sfu}$  against the original per-image determination of basis  $\mathbf{B}$  for Matrix-R refined bilinear upsampling in Table 9. Here, we see that under the cross-validation setup, while we can still find a lower-dimensional basis suggesting better mean performance compared to standalone Matrix-R, it does not happen until  $m = 21$ . That is, for adopting  $\mathbf{B}_{K=4}$ , we need a much higher-dimensional basis to improve from the original Matrix-R. On the other hand, using  $\mathbf{B}_{sfu}$  trained from SFU datasets did not improve the original Matrix-R's performance. Note that since our spectral images have only 31 channels, the highest "lower" dimension is  $m = 30$ , with  $m = 31$  equivalent to stand-alone Matrix-R (spectra are kept in the original dimension).

**Table 9.** Test on using fixed spectral basis for lower-dimensional variants of our Matrix-R method (" $m$ -dim"). The table reports RMSE ( $\times 10^3$ ) values for the standard Matrix-R method ("Matrix-R") and its lower-dimensional variants (" $m$ -dim") applied to bilinear upsampling images. The best results are shown in bold font and underlined.

	Per-Image $\mathbf{B}$ (Original)		Cross Validation $\mathbf{B}_{K=4}$		SFU Dataset $\mathbf{B}_{sfu}$	
	Mean	99 pt	Mean	99 pt	Mean	99 pt
Bilinear Upsampling	7.94	62.06	7.94	62.06	7.94	62.06
Matrix-R	3.54	26.51	3.54	26.51	<b><u>3.54</u></b>	<b><u>26.51</u></b>
3 dim	13.52	63.86	8.45	23.37	14.88	43.68
4 dim	<b><u>2.77</u></b>	<b><u>11.87</u></b>	4.84	<b><u>18.43</u></b>	12.24	36.25
5 dim	3.33	22.37	4.65	26.11	12.16	36.03
6 dim	3.33	24.02	4.13	26.01	8.23	29.38
7 dim	3.37	24.73	3.92	25.89	7.96	29.12
8 dim	3.38	25.07	3.80	25.99	7.06	27.76
⋮						
21 dim	3.48	26.06	<b><u>3.49</u></b>	26.11	3.86	26.48
⋮						
30 dim	3.54	26.48	3.54	26.51	3.58	26.54

Combined, we see that under the current experimental setup, training the basis  $\mathbf{B}$  per image performs better than using a cross-scene or general basis. The SFU result implies that  $\mathbf{B}_{sfu}$  might not be representative enough for the ICVL image dataset, and the cross-validation result likely suggests that a random selection of images used for training bases might not be a good setup in practice (grouping images based on their similarity in content and/or lighting conditions might help obtaining better-performing bases).

## 6. Discussion

Our proposed post-processing Matrix-R method can be applied in a wide context: the proposed process could 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.

As for Theorem 2, in our experiment, we were not given a known lower-dimensional basis that will definitely represent all spectra in each scene. This means that the interchangeability of  $\mathbf{Q}$  and  $\overline{\mathbf{Q}}$  in Equation (20) may not hold, i.e.,  $\overline{\mathbf{Q}}$  can create color error, and subsequently, error in calculating the ground-truth fundamental metamer  $\underline{e}^{\overline{\mathbf{Q}}}$  from the RGB input  $\underline{\rho}$ . Of course, we may sensibly assume that as the assumed spectral dimension  $m$  increases (i.e., when  $m$  approaching  $n$ , the original spectral dimension), we get more accurate low-dimensional model that represents spectra. And yet, according to our results, this cannot be the only factor that affects the optimal selection of  $m$ . The determination of the optimal basis dimension  $m$  is heavily dependent on the fidelity of the initial spectral estimate. We observed that high-performing deep learning models, such as AWAN and MIAE, typically achieve peak performance with higher-dimensional subspaces. This is likely because these models successfully recover subtle spectral nuances; applying a restrictive, low-dimensional  $m$  in these cases would discard valid spectral information, effectively acting as a form of over-regularisation. Conversely, simpler methods like bilinear interpolation often produce coarse spectral estimates with significant deviations from the ground-truth. For these algorithms, using a lower  $m$  is advantageous as it enforces a stricter prior, constraining the noisy estimates to a fundamental subspace of natural spectra and thereby filtering out gross spectral errors. Consequently, we recommend determining  $m$  empirically to match the specific capacity of the chosen reconstruction algorithm.

Another observation on optimal  $m$  is that, for some algorithms, the optimal  $m$  is different for mean and worst-case (99-percentile) results, and the latter generally suggests a smaller optimal  $m$ . This is understandable as a lower-dimensional linear model has the effect of bounding the outliers from exceeding what the underlying assumed basis can explain. Conversely, some originally more accurate pixels could be overgeneralized by the basis and lose accuracy. We can observe both effects in the "Upsampling Only" result in Figure 4. Here, in the subplot labeled "Matrix-R (4-dim)", we see that a 4-dimensional spectral representation make the boundaries of buildings and the sky much more accurate, while loses accuracy in areas around the sidewalk.

We also observed that the magnitude of improvement from Matrix-R post-processing is inversely related to the baseline algorithm's accuracy. Unlike "physically blind" algorithms, advanced models such as AWAN and MIAE explicitly incorporate the camera's spectral sensitivity functions, resulting in spectral estimates with highly accurate fundamental metamers. Thus, our method offers the most significant value to algorithms that are not explicitly constrained by sensor physics.

Looking ahead, we recognize that employing real images from different cameras could present additional challenges [49,63], including image registration and varying exposure levels [24,39]. While our aim in this study was to propose a theoretical solution through simulation as a preliminary step, future research may explore these methods using actual camera setups.

## 7. Conclusions

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, and spectral reconstruction (SR) algorithms recover hyperspectral images directly from the RGBs. Yet, most of 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 spectral recovery: we simply make sure that it has the correct fundamental metamer. And, we provide mathematical proof that this improvement will always happen. Furthermore, we developed the Matrix-R method where spectra are represented by a low-dimensional linear model.

Experiments on several historic and state-of-the-art PS and SR algorithms show that our proposed post-processing Matrix-R method always improved these algorithms. In addition, the low-dimensional linear basis variant of our theorem was shown to yield the best recovery results. Finally, our exploration of multispectral pan-sharpening reaffirms the efficacy of the Matrix-R method and its lower-dimensional variant.

**Author Contributions:** Conceptualization, G.D.F., Y.-T.L., and A.K.; methodology, G.D.F., Y.-T.L., and A.K.; software, Y.-T.L. and A.K.; validation, Y.-T.L. and A.K.; formal analysis, G.D.F., Y.-T.L., and A.K.; investigation, G.D.F., Y.-T.L., and A.K.; resources, G.D.F.; data curation, Y.-T.L. and A.K.; writing—original draft preparation, Y.-T.L.; writing—review and editing, G.D.F., Y.-T.L., and A.K.; visualization, Y.-T.L. and A.K.; supervision, G.D.F.; project administration, G.D.F.; funding acquisition, G.D.F. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was funded by EPSRC, United Kingdom Grant EP/S028730/1 and Spectricity.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** Two publicly available datasets were used in this study. First, BGU ICVL Hyperspectral Dataset, which can be accessed via: <https://icvl.cs.bgu.ac.il/pages/researches/hyperspectral-imaging.html> (accessed: 22 July 2025). Second, the CAVE Multispectral Image Dataset, which can be accessed via: <https://cave.cs.columbia.edu/repository/Multispectral> (accessed: 1 December 2025).

**Conflicts of Interest:** The authors declare no conflicts of interest.

## References

1. Chakrabarti, A.; Zickler, T. Statistics of real-world hyperspectral images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Colorado Springs, CO, USA, 20–25 June 2011; pp. 193–200.
2. Lv, M.; Chen, T.; Yang, Y.; Tu, T.; Zhang, N.; Li, W.; Li, W. Membranous nephropathy classification using microscopic hyperspectral imaging and tensor patch-based discriminative linear regression. *Biomed. Opt. Express* **2021**, *12*, 2968–2978. [CrossRef]
3. Courtenay, L.; González-Aguilera, D.; Lagüela, S.; Del Pozo, S.; Ruiz-Mendez, C.; Barbero-García, I.; Román-Curto, C.; Cañueto, J.; Santos-Durán, C.; Cardeñoso-Álvarez, M.; et al. Hyperspectral imaging and robust statistics in non-melanoma skin cancer analysis. *Biomed. Opt. Express* **2021**, *12*, 5107–5127. [CrossRef]
4. Wang, W.; Ma, L.; Chen, M.; Du, Q. Joint correlation alignment-based graph neural network for domain adaptation of multitemporal hyperspectral remote sensing images. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2021**, *14*, 3170–3184. [CrossRef]
5. Torun, O.; Yuksel, S. Unsupervised segmentation of LiDAR fused hyperspectral imagery using pointwise mutual information. *Int. J. Remote Sens.* **2021**, *42*, 6461–6476. [CrossRef]
6. Chen, Z.; Wang, J.; Wang, T.; Song, Z.; Li, Y.; Huang, Y.; Wang, L.; Jin, J. Automated in-field leaf-level hyperspectral imaging of corn plants using a Cartesian robotic platform. *Comput. Electron. Agric.* **2021**, *183*, 105996. [CrossRef]
7. Gomes, V.; Mendes-Ferreira, A.; Melo-Pinto, P. Application of hyperspectral imaging and deep learning for robust prediction of sugar and pH levels in wine grape berries. *Sensors* **2021**, *21*, 3459. [CrossRef] [PubMed]
8. Pane, C.; Manganiello, G.; Nicastro, N.; Cardi, T.; Carotenuto, F. Powdery mildew caused by *Erysiphe cruciferarum* on wild rocket (*Diplotaxis tenuifolia*): Hyperspectral imaging and machine learning modeling for non-destructive disease detection. *Agriculture* **2021**, *11*, 337. [CrossRef]

9. Picollo, M.; Cucci, C.; Casini, A.; Stefani, L. Hyper-spectral imaging technique in the cultural heritage field: New possible scenarios. *Sensors* **2020**, *20*, 2843. [[CrossRef](#)] [[PubMed](#)]
10. Grillini, F.; Thomas, J.; George, S. Mixing models in close-range spectral imaging for pigment mapping in cultural heritage. In Proceedings of the International Colour Association (AIC) Conference, Online, 20–27 November 2020; pp. 372–376.
11. Heikkinen, V.; Lenz, R.; Jetsu, T.; Parkkinen, J.; Hauta-Kasari, M.; Jääskeläinen, T. Evaluation and unification of some methods for estimating reflectance spectra from RGB images. *J. Opt. Soc. Am. A* **2008**, *25*, 2444–2458. [[CrossRef](#)]
12. Aeschbacher, J.; Wu, J.; Timofte, R. In defense of shallow learned spectral reconstruction from RGB images. In Proceedings of the IEEE Conference on Computer Vision Workshops, Venice, Italy, 22–29 October 2017; pp. 471–479.
13. Nguyen, R.; Prasad, D.; Brown, M. Training-based spectral reconstruction from a single RGB image. In Proceedings of the European Conference on Computer Vision, Zurich, Switzerland, 6–12 September 2014; pp. 186–201.
14. Lin, Y.T.; Finlayson, G.D. A Rehabilitation of Pixel-Based Spectral Reconstruction from RGB Images. *Sensors* **2023**, *23*, 4155. [[CrossRef](#)]
15. Arad, B.; Timofte, R.; Yahel, R.; Morag, N.; Bernat, A.; Cai, Y.; Lin, J.; Lin, Z.; Wang, H.; Zhang, Y.; et al. NTIRE 2022 spectral recovery challenge and data set. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 19–24 June 2022; pp. 863–881.
16. Shi, Z.; Chen, C.; Xiong, Z.; Liu, D.; Wu, F. Hscnn+: Advanced cnn-based hyperspectral recovery from RGB images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, Salt Lake City, UT, USA, 18–22 June 2018; pp. 939–947.
17. Li, J.; Wu, C.; Song, R.; Li, Y.; Liu, F. Adaptive weighted attention network with camera spectral sensitivity prior for spectral reconstruction from RGB images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, Online, 14–19 June 2020; pp. 462–463.
18. Lanaras, C.; Baltasvias, E.; Schindler, K. Hyperspectral super-resolution by coupled spectral unmixing. In Proceedings of the IEEE International Conference on Computer Vision, Santiago, Chile, 11–18 December 2015; pp. 3586–3594.
19. Yokoya, N.; Yairi, T.; Iwasaki, A. Coupled nonnegative matrix factorization unmixing for hyperspectral and multispectral data fusion. *IEEE Trans. Geosci. Remote Sens.* **2011**, *50*, 528–537. [[CrossRef](#)]
20. Liu, J.; Wu, Z.; Xiao, L.; Wu, X.J. Model inspired autoencoder for unsupervised hyperspectral image super-resolution. *IEEE Trans. Geosci. Remote Sens.* **2022**, *60*, 1–12. [[CrossRef](#)]
21. Hu, J.F.; Huang, T.Z.; Deng, L.J.; Dou, H.X.; Hong, D.; Vivone, G. Fusformer: A transformer-based fusion network for hyperspectral image super-resolution. *IEEE Geosci. Remote Sens. Lett.* **2022**, *19*, 1–5. [[CrossRef](#)]
22. Wyszecki, G. *Valenzmetrische Untersuchung des Zusammenhanges Zwischen Normaler und Anomaler Trichromasie*; Technical University of Berlin: Berlin, Germany, 13 July 1953.
23. Lin, Y.T.; Finlayson, G. Physically Plausible Spectral Reconstruction. *Sensors* **2020**, *20*, 6399. [[CrossRef](#)]
24. Arad, B.; Timofte, R.; Ben-Shahar, O.; Lin, Y.T.; Finlayson, G. NTIRE 2020 challenge on spectral reconstruction from an RGB Image. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, Online, 14–19 June 2020; pp. 446–447.
25. Arad, B.; Ben-Shahar, O.; Timofte, R.; Van Gool, L.; Zhang, L.; Yang, M.-H.; Xiong, Z.; Chen, C.; Shi, Z.; Liu, D.; et al. NTIRE 2018 challenge on spectral reconstruction from RGB images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, Salt Lake City, UT, USA, 18–22 June 2018; pp. 929–938.
26. Dong, W.; Zhou, C.; Wu, F.; Wu, J.; Shi, G.; Li, X. Model-guided deep hyperspectral image super-resolution. *IEEE Trans. Image Process.* **2021**, *30*, 5754–5768. [[CrossRef](#)]
27. Zhang, L.; Nie, J.; Wei, W.; Li, Y.; Zhang, Y. Deep blind hyperspectral image super-resolution. *IEEE Trans. Neural Netw. Learn. Syst.* **2020**, *32*, 2388–2400. [[CrossRef](#)] [[PubMed](#)]
28. Stiebel, T.; Seltsam, P.; Merhof, D. Enhancing Deep Spectral Super-resolution from RGB Images by Enforcing the Metameric Constraint. In Proceedings of the VISIGRAPP (4: VISAPP), Valletta, Malta, 27–29 February 2020; pp. 57–66.
29. Lin, Y.T.; Finlayson, G.D.; Kucuk, A. An optimality property of Matrix-R theorem, its extension, and the application to hyperspectral pan-sharpening. In Proceedings of the Color and Imaging Conference, Society for Imaging Science and Technology, Paris, France, 13–17 November 2023; Volume 31, pp. 144–149.
30. Strang, G. *Introduction to Linear Algebra*, 5th ed.; Wellesley-Cambridge Press: Wellesley, MA, USA, 2016.
31. Cohen, J.; Kappauf, W. Metameric color stimuli, fundamental metamers, and Wyszecki’s metameric blacks. *Am. J. Psychol.* **1982**, *95*, 537–564. [[CrossRef](#)] [[PubMed](#)]
32. Parkkinen, J.; Hallikainen, J.; Jaaskelainen, T. Characteristic spectra of Munsell colors. *J. Opt. Soc. Am. A* **1989**, *6*, 318–322. [[CrossRef](#)]

33. Chen, Q.; Wang, L.; Westland, S. A Perceptual Study of Linear Models of Spectral Reflectance. In Proceedings of the IEEE Conference on Computer Science and Software Engineering, Wuhan, China, 12–14 December 2008; pp. 309–312.
34. Wandell, B. The synthesis and analysis of color images. *IEEE Trans. Pattern Anal. Mach. Intell.* **1987**, *PAMI-9*, 2–13. [[CrossRef](#)]
35. Jiang, J.; Liu, D.; Gu, J.; Süsstrunk, S. What is the space of spectral sensitivity functions for digital color cameras? In Proceedings of the IEEE Workshop on Applications of Computer Vision, Clearwater, FL, USA, 15–17 January 2013; pp. 168–179.
36. Carl, D. *Matrix Analysis and Applied Linear Algebra*; Society for Industrial and Applied Mathematics: Philadelphia, PA, USA, 2000.
37. Lin, Y.T.; Finlayson, G. Physically Plausible Spectral Reconstruction from RGB Images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, Online, 14–19 June 2020; pp. 532–533.
38. Connah, D.; Hardeberg, J. Spectral recovery using polynomial models. In Proceedings of the Color Imaging X: Processing, Hardcopy, and Applications, San Jose, CA, USA, 17–20 January 2005; pp. 65–75.
39. Lin, Y.T.; Finlayson, G. Exposure Invariance in Spectral Reconstruction from RGB Images. In Proceedings of the Color and Imaging Conference, Paris, France, 21–25 October 2019; pp. 284–289.
40. Arad, B.; Ben-Shahar, O. Sparse recovery of hyperspectral signal from natural RGB images. In Proceedings of the European Conference on Computer Vision, Amsterdam, The Netherlands, 11–14 October 2016; pp. 19–34.
41. Huang, G.; Liu, Z.; Van Der Maaten, L.; Weinberger, K.Q. Densely connected convolutional networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 4700–4708.
42. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
43. Xia, B.N.; Gong, Y.; Zhang, Y.; Poellabauer, C. Second-order non-local attention networks for person re-identification. In Proceedings of the IEEE/CVF Conference on Computer Vision, Long Beach, CA, USA, 15–20 June 2019; pp. 3760–3769.
44. Maloney, L.; Wandell, B. Color constancy: A method for recovering surface spectral reflectance. *J. Opt. Soc. Am. A* **1986**, *3*, 29–33. [[CrossRef](#)]
45. Drew, M.S.; Funt, B.V. Natural metamers. *CVGIP Image Underst.* **1992**, *56*, 139–151. [[CrossRef](#)]
46. Morovic, P.; Finlayson, G. Metamer-set-based approach to estimating surface reflectance from camera RGB. *J. Opt. Soc. Am. A* **2006**, *23*, 1814–1822. [[CrossRef](#)] [[PubMed](#)]
47. Zhao, Y.; Berns, R.S. Image-based spectral reflectance reconstruction using the matrix R method. *Color Res. Appl.* **2007**, *32*, 343–351. [[CrossRef](#)]
48. Hu, J.F.; Huang, T.Z.; Deng, L.J.; Jiang, T.X.; Vivone, G.; Chanussot, J. Hyperspectral image super-resolution via deep spatio-spectral attention convolutional neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *33*, 7251–7265. [[CrossRef](#)]
49. Lin, Y.T.; Finlayson, G.D. Evaluating the Performance of Different Cameras for Spectral Reconstruction. In Proceedings of the Color and Imaging Conference, Scottsdale, AZ, USA, 13–17 November 2022; pp. 213–218.
50. Imai, F.; Berns, R. High-resolution multi-spectral image archives: A hybrid approach. In Proceedings of the Color and Imaging Conference, Scottsdale, AZ, USA, 17–20 November 1998; pp. 224–227.
51. Horn, R.A.; Johnson, C.R. Norms for vectors and matrices. In *Matrix Analysis*; First paperback edition; Cambridge University Press: New York, NY, USA, 1990; pp. 313–386.
52. Tikhonov, A.; Goncharsky, A.; Stepanov, V.; Yagola, A. *Numerical Methods for the Solution of Ill-Posed Problems*; Springer: London, UK, 1995.
53. Webb, G.I. Overfitting. In *Encyclopedia of Machine Learning*; Sammut, C., Webb, G.I., Eds.; Springer: Boston, MA, USA, 2010; p. 744.
54. McDonald, G.C. Ridge regression. *Wiley Interdiscip. Rev. Comput. Stat.* **2009**, *1*, 93–100. [[CrossRef](#)]
55. Cheney, W.; Kincaid, D. *Linear Algebra: Theory and Applications*; The Australian Mathematical Society: Canberra, ACT, Australia, 2009; pp. 544, 558.
56. Commission Internationale de L'éclairage. *CIE Proceedings (1963) Vienna Session (Committee Report E-1.4.1)*; Bureau Central de la CIE: Paris, France, 1964; Volume B.
57. S1 Multispectral Image Sensor and Camera Module. Available online: <https://spectricity.com/product/> (accessed on 18 January 2025).
58. Lin, Y.T.; Finlayson, G. Reconstructing Spectra from RGB Images by Relative Error Least-Squares Regression. In Proceedings of the Color and Imaging Conference, Online, 4–19 November 2020; pp. 264–269.
59. Yasuma, F.; Mitsunaga, T.; Iso, D.; Nayar, S. Generalized assorted pixel camera: Postcapture control of resolution, dynamic range, and spectrum. *IEEE Trans. Image Process.* **2010**, *19*, 2241–2253. [[CrossRef](#)] [[PubMed](#)]
60. Tominaga, S.; Nishi, S.; Ohtera, R. Measurement and estimation of spectral sensitivity functions for mobile phone cameras. *Sensors* **2021**, *21*, 4985. [[CrossRef](#)]
61. Darrodi, M.M.; Finlayson, G.; Goodman, T.; Mackiewicz, M. Reference data set for camera spectral sensitivity estimation. *J. Opt. Soc. Am. A* **2015**, *32*, 381–391. [[CrossRef](#)]

62. Barnard, K.; Martin, L.; Funt, B.; Coath, A. A data set for color research. *Color Res. Appl.* **2002**, *27*, 147–151. [[CrossRef](#)]
63. Fu, Y.; Zhang, T.; Zheng, Y.; Zhang, D.; Huang, H. Joint camera spectral response selection and hyperspectral image recovery. *IEEE Trans. Pattern Anal. Mach. Intell.* **2020**, *44*, 256–272. [[CrossRef](#)] [[PubMed](#)]

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.