

Article

Towards ‘Fourth Paradigm’ Spectral Sensing

Forrest Simon Webler ^{1,*}, Manuel Spitschan ^{2,3,4}  and Marilyne Andersen ¹ 

¹ Laboratory of Integrated Performance in Design (LIPID), School of Architecture, Civil and Environmental Engineering (ENAC), École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland; marilyne.andersen@epfl.ch

² Translational Sensory & Circadian Neuroscience, Max Planck Institute for Biological Cybernetics, 72076 Tübingen, Germany; manuel.spitschan@tum.de

³ TUM Department of Sport and Health Sciences (TUM SG), Technical University of Munich, 80333 Munich, Germany

⁴ TUM Institute for Advanced Study (TUM-IAS), Technical University of Munich, 85748 Garching, Germany

* Correspondence: forrest.webler@epfl.ch

Abstract: Reconstruction algorithms are at the forefront of accessible and compact data collection. In this paper, we present a novel reconstruction algorithm, SpecRA, that adapts based on the relative rarity of a signal compared to previous observations. We leverage a data-driven approach to learn optimal encoder-array sensitivities for a novel filter-array spectrometer. By taking advantage of the regularities mined from diverse online repositories, we are able to exploit low-dimensional patterns for improved spectral reconstruction from as few as $p = 2$ channels. Furthermore, the performance of SpecRA is largely independent of signal complexity. Our results illustrate the superiority of our method over conventional approaches and provide a framework towards “fourth paradigm” spectral sensing. We hope that this work can help reduce the size, weight and cost constraints of future spectrometers for specific spectral monitoring tasks in applied contexts such as in remote sensing, healthcare, and quality control.

Keywords: reconstruction; symmetric non-negative matrix factorization; nonlinear dimensionality reduction; sparse sensor placement; spectral sensing



Citation: Webler, F.S.; Spitschan, M.; Andersen, M. Towards ‘Fourth Paradigm’ Spectral Sensing. *Sensors* **2022**, *22*, 2377. <https://doi.org/10.3390/s22062377>

Academic Editor: Jiayi Ma

Received: 24 December 2021

Accepted: 18 March 2022

Published: 19 March 2022

Publisher’s Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 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

Natural signals are compressible functions that represent changes in the spectrotemporal dynamics of physical phenomena [1]. Common examples of natural signals include light and sound. The information contained within a signal is encoded when it is received by an observer. Observers can be biological, such as the human eye, or mechanical (e.g., a digital camera). The most useful “observers” encode signal information into a format that can be read, copied, and shared with others via a process called quantization. Contemporary scientific discovery is increasingly dependent on encoding hardware as numerous autonomous processes require vast amounts of data. In what many are now calling the “fourth paradigm in scientific discovery” [2], efficient data transcoding is paramount. Herein, we approach the metrological process from the perspective of information science starting with quantization and concluding with reconstruction. We hypothesize that by exploiting regularities in existing datasets, we can develop optimized non-uniform protocols for spectral sensor placement and use adaptive methods to maximize reconstruction efficiency. While we focus on spectral sensing for visible light, the methods discussed are applicable to any set of signals characterized by locality and compositionality.

1.1. What Is “Spectral” Sensing?

Spectral sensing is used in a number of applied contexts including healthcare, remote sensing, and quality control [3–5]. To meet the requirements of this highly diverse

field, current compact spectral imaging solutions include systems based on the following technologies:

1. Multispectral, filter-array sensors;
2. Compressive spectral sensors;
3. Lens-encoders;
4. Spectral super-resolution;

While each application of spectral sensing has different requirements for what is considered “spectral” resolution, we can broadly identify three main classes: (1) tristimulus; (2) multispectral; and (3) hyperspectral imaging. Pixel-based, red, green, blue (RGB) is perhaps the most ubiquitous “spectral” signal and is used for color rendering and other imaging purposes. While RGB has historically been excluded from spectral sensing technologies, interest in recovering broadband spectral signals from RGB has increased over the years [6–11]. So-called “RGB-to-spectrum” approaches are, however, limited in their applicability to broadband (low-complexity) signals and generally fail for signals with low autocorrelation. Multispectral imaging is an exciting emerging field and typically refers to systems capturing between 3 and 10 channels with bandwidths greater than 20 nm [12]. Common applications are in color evaluation for quality control and remote sensing [13]. Advances in optical filters and semiconductor technology have also improved the size, weight, power, and cost constraints, making hyperspectral sensing a more attractive middle ground [14]. Hyperspectral imaging is perhaps the most diverse class reserved for devices measuring more than 10 channels [15]. For imaging applications, hyperspectral cameras can have between 512 and 2048 channels over the visible domain making them extraordinarily information-rich. While devices in each of these classes all collect “spectral” data, the range in resolution covers three orders of magnitude.

1.2. Existing Methods

Many compact spectral sensing methods exist based on the aforementioned technologies. Recent interest in multi-spectral filter-array technology has, in particular, attracted a lot of interest. Methods of multispectral demosaicking have improved low-cost “single shot” imaging [16] as well as compressive methods [17]. Alternative approaches look for statistical regularities to exploit for task-specific applications. A plethora of work in natural image statistics has motivated approaches based on scene optimization [18] and RGB images [8,19,20]. Increasingly, however, focus has shifted towards data-driven computational methods to recover as much information as possible from low-fidelity measurements. This includes improving the optimization frameworks for diffractive achromats (DAs) [21] and constructive improvements for the alternating direction method of multipliers (ADMM) optimization in solving ill-posed reconstruction problems [22].

Contemporary approaches to compressive spectral sensing rely on a diversity of sensor technologies from liquid crystal phase retarders [23–25] to stacked array spectrometers with broadband filters [26]. Emergent technologies such as quantum dot and nanowire spectrometers [27,28] also show great potential to disrupt the spectral sensing space. At the same time, the manufacturing cost of “single shot” compact filter-array spectrometers with 3–20 channels and bandwidths in the 20 nm range have dramatically decreased to a fraction of the cost of conventional scanning spectrometers [29,30]. Regardless of the underlying technology, each measurement device maps data from the “natural” dimension to a hardcoded “measurement dimension” defined by the resolution of the device. Towards this end, we are interested in applying a data-driven approach to investigate the theoretical boundaries constraining reconstruction performance in the future generation of compact sensing hardware.

1.3. Problem Statement

While determining the threshold of what constitutes “spectral sensing” is highly contextual, advances in applied math and signal processing have demonstrated that natural signals exhibit a high-degree of redundancy. Formally, this means that most natural signals

are encoded in a smaller dimension *below* the theoretical limits set by the Shannon–Nyquist sampling theorem [31]. As we are interested in investigating the magnitude of information required to reconstruct a spectral power distribution from low-dimensional data, we will try and find the best low-rank representation (i.e., *embedding*) given a set of prior observations. Determining a theoretical limit and workable latent space for visible spectra will not only help to improve how spectral data are measured, compressed, and stored, but also in how they are classified.

Efficient sensing is a two-step process requiring the informed design of an optimized encoder g_ϕ and an adaptive reconstruction algorithm characterized by f_θ . Thus, the contributions herein are two-fold. First, to determine the optimal sensor locations for the reconstruction of broadband (low-complexity) and narrowband (high-complexity) spectra. Second, to develop an a data-driven reconstruction algorithm (SpecRA) that balances simplicity and reconstruction fidelity. Together, we show that a combined workflow can account for real-world engineering and fabrication constraints on spectral sensors to determine *what are the maximally informative dimensions of visible spectra* in theory and practice.

2. Beating Nyquist

In 2005, our understanding of the theoretical limits of sampling rate for bandlimited signals were such that if a signal is sampled at a frequency f , perfect reconstruction is only guaranteed if the bandwidth $b < f/2$ [32]. This observation is a result of the Shannon–Nyquist sampling theorem and while this theorem still holds today, two seismic advances in applied math have forced engineers to contextualize it in a new light. First, in 2006, it was demonstrated that sub-Nyquist sampling was possible without violating the theorem by requiring the signals to be *sparse* (i.e., compressible) in a generic basis [31,33]. This approach is called *compressive sensing* and has since revolutionized data collection by initializing compression at the point of quantization. Second, the increase in data availability has made it possible to feature mine prior observations for statistical regularities. These regularities can then be used to exploit symmetries and other structural properties in order to make inferences that further maximize reconstruction performance. In this way, the “datafication” of our world has imparted huge implications for metrology in general and spectral sensing in particular. In this paper, we will specifically show how advanced domain knowledge can be used to optimize the measurement process.

3. What Is Reconstruction?

Simply defined, *reconstruction* refers to a process of recovering a signal from a set of limited measurements [1]. In practice, this can be done several ways depending on the application and in context, all of which result in solving the following problem:

$$\min \mathcal{K}(\hat{\mathbf{s}}, \mathbf{s}) \quad \text{subject to} \quad \hat{\mathbf{s}} = f_\theta(\mathbf{y}) \quad (1)$$

Here, $\mathbf{y} = g_\phi(\mathbf{s})$ is a measurement of a high-dimensional signal $\mathbf{s} = E(\lambda)$ by an encoder g_ϕ reconstructed by a function f_θ (c.f., Figure 1). The function \mathcal{K} being minimized can be any metric or norm capturing the dissimilarity between the reconstruction and the ground truth. If the signal is sampled at or above the Nyquist rate, reconstruction can be as simple as constructing a linear fit through the sub-sampled points (we call this the “interpolation regime”). Alternatively, if the signal is undersampled but known to exhibit unique statistical regularities, “reconstruction” can also be a process of finding a match among known prior observations. In both examples, we say the reconstruction is “naïve” in the sense that interpolation and pattern matching are inherently trivial tasks.

If the reconstruction is *lossy*, performance metrics are accompanied with a compression power score to contextualize the trade-off between complexity and descriptivity. The compression power is defined as the ratio between the uncompressed and compressed file size (e.g., defined by the vector length). What can be misleading about the compression power is that in the limit that the compressed file size obtains information saturation (i.e.,

is sampled at a lossless rate via Nyquist or CS) the compression ratio becomes meaningless when the uncompressed signal dimension is increased. This is because the features used for lossless reconstruction are fully formed at a given resolution, and artificially increasing the resolution of the uncompressed file will inflate the compression power. Towards this end, we propose a slightly amended performance metric that encapsulates the about of “work” performed by the algorithm $f_\theta : \mathbb{R}^p \rightarrow \mathbb{R}^n$ in order to reach a target error depending on the available information:

$$W = \max[p^{-1}n(\Delta - \max[\gamma, \mu]), 0] \tag{2}$$

where Δ is the target performance threshold, γ is the percent of signals matching signals in the Kanji, and μ is the percentage of signals recovered at or above the Nyquist rate. We can think of W as the penalized compression power that is set equal to 0 when the reconstruction algorithm fails to outperform either the matching problem or interpolation. Given that we reconstruct a set of observations encoded by g_ϕ , from a Kanji $\mathbf{K} \in \mathbb{R}^m$, we can define the proportion of adequately reconstructed signals as the ratio of the cardinalities of S and \mathbf{S} where $S = \{ \mathbf{s} \mid \mathbf{s} \in \mathbf{S} \wedge \mathcal{K}(\mathbf{s}, \hat{\mathbf{s}}) < \epsilon \ \forall \ \hat{\mathbf{s}} = f_\theta(\mathbf{y}) \}$ is the subset of reconstructed observations adhering to a dissimilarity score less than ϵ defined by a function \mathcal{K} appropriate for the underlying datatype. For spectral sensing applications, this is typically the spectral angle (SA) or another derivative index corresponding to the spectral information divergence [34,35].

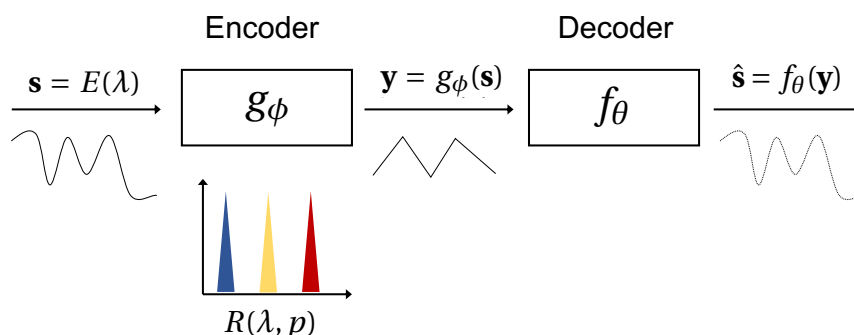


Figure 1. Spectrometry is the process of encoding spectral data to a measured dimension p from the infinite natural signal dimension. This process is visualized in the above graphic and mathematically written as $g_\phi : \mathbb{R}^\infty \rightarrow \mathbb{R}^p$. The measured data can then be reconstructed to a target high resolution $n > p$ such that $f_\theta : \mathbb{R}^p \rightarrow \mathbb{R}^n$. Here, the encoder g_ϕ is a filter-array spectrometer with some transmission functions $R(\lambda, i)$ for $i = (1, \dots, p)$ and f_θ represents a reconstruction algorithm. This structure is analogous to that of an autoencoder where the learned encoder weights would replace the response functions of the physical spectrometer and p is the dimension of the hidden layer.

When we say that the algorithm has to “work”, what we mean is that unlike in pattern matching and interpolation tasks where there is no optimization procedure taking place, reconstruction has to “add” information that is not trivially available. When signals have no corresponding match and are sampled in the sub-Nyquist regime, the reconstruction problem formulated in Equation (1) becomes:

$$\min \|\mathbf{x}\|_1 \quad \text{subject to} \quad \hat{\mathbf{s}} = \sum_{i=1}^m x_i \mathbf{k}_i \tag{3}$$

where \mathbf{x} is a coefficient vector and $\mathbf{k}_i \in \mathbf{K}$ is a prior observation contained in our “Kanji” (a Pareto-optimal library distinct from a learned dictionary). A fundamental property of signal reconstruction is that the more complex a signal is, the more basis modes are needed to accurately approximate that signal. One of the consequences of deriving \mathbf{K} from a big dataset \mathbf{L} is that the likelihood that a “new” observation has already been measured

and exists in an accessible dataset is high. Consequently, reconstruction in the era of big data is less about adhering to a specific methodology and instead about finding the most efficient process required to return the true signal from an encoded measurement. Towards this end, ensuring that \mathbf{K} comprises real spectral observations ensures that the “missing” information in an undersampled measurement can be “filled in”. That said, when we take a measurement \mathbf{y} , the only information we have is that of the encoded signal. This means that we have to “trick” the algorithm by solving:

$$\min \|\mathbf{x}'\|_1 \quad \text{subject to} \quad \hat{\mathbf{y}} = \sum_{i=1}^m x'_i g_\phi(\mathbf{k}_i) \quad (4)$$

with the assumption that $\mathbf{x}' \approx \mathbf{x}$. Given that \mathbf{K} is not a generic basis, the guarantees of compressed sensing do not hold in this case. Instead, the efficacy of this assumption is constrained by the ability of signals in \mathbf{K} to preserve a unique structure in the measured dimension.

Of course, not all measured dimensions are the same. For example, we show in Figure 2 what happens when signals are sampled with a uniform encoder. A central aim of this work is to investigate other sampling protocols by constructing g_ϕ from bases, learned via different decomposition methods. We hypothesize that non-uniform sampling procedures will be informed by the features shared by many real-world spectra and will therefore result in a greater preservation of “uniqueness” in the measured dimension. Furthermore, by extension, we hypothesize that preserving structure in the encoded output will result in greater reconstruction performance for undersampled data.

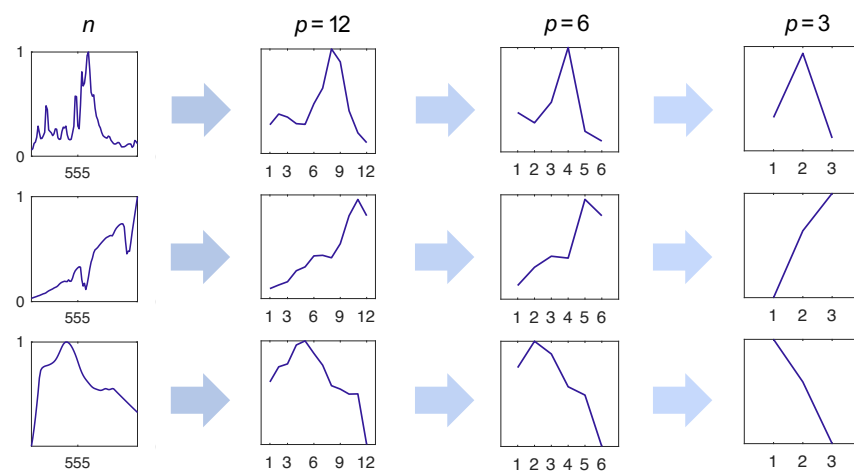


Figure 2. Unique identifying features of the signal are lost when the resolution of the sensor is decreased. Here, we show the effect of uniform sampling. Preserving unique features in the measured dimension is important for ensuring agreement between \mathbf{x}' and \mathbf{x} .

Looking at Figure 2, one may also question why we would bother forming the reconstruction problem as Equation (4) in lieu of training a neural network on \mathbf{K} , thereby learning a map between the measured and target dimensions. We refrain from taking such an approach for two reasons: first, it is well known that “black box” models do not generalize well and this has, in fact, been demonstrated to be the case in the application of recovering visible spectral distributions from encoder-array spectrometers [36,37]; second, we lose any interpretability or contextual reference of the scene. While Equation (4) may seem trivially simple, enforcing sparsity via the l_1 -norm has demonstrated incredible success in not only reconstructing low-dimensional measurements, but also in reconstructing them in a way that mimics the underlying physical system [38]. This observation has led some to claim that parsimony is the ultimate physics regularizer [39].

4. Towards Data-Driven Bases for Reconstruction

Variation in spectral distributions is created from interactions between light and matter. In order to derive a basis, we first need to compile some observable data. Here, we compile a library of illuminant and reflectance spectra from available open source datasets [40–43]. A high-level description of our library **L** is summarized in Figure 3. Since each spectrum was sampled at different frequencies, we normalized all spectra within the visible range from 380 to 780 nm and re-sampled them using standard interpolation methods. Within **L**, we have 401 spectra which were used as a representative set for color rendition studies in addition to 99 color evaluation samples (CESs) uniformly distributed within the natural color system (NCS) gamut [44,45].

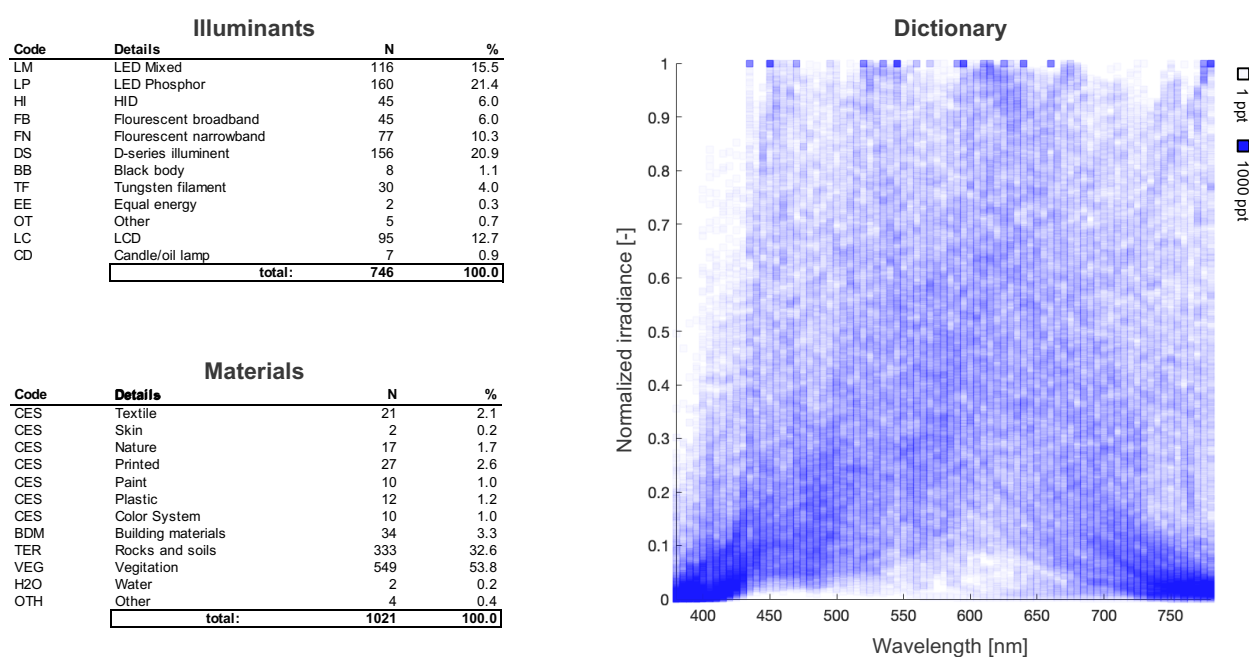


Figure 3. A high-level description of the contents of our spectral libraries for the illuminants and material reflectances. Of the 746 illuminant spectra, 146 are theoretical and 600 are real. The 1021 reflectances correspond to material samples under equal energy illumination. We can see that despite the variety of illuminants and materials, there are clear regularities in signal space (e.g., very few spectra have high relative power between 380 and 430 nm).

Additions to the illuminant class include mobile and computer screen as well as high-intensity discharge (HID) sources used in street lighting, and daylight spectra for various sun angles above the horizon in urban and rural settings [46]. We acknowledge that this set is not complete; however, we believe that from a mechanistic perspective, the space of available spectral illuminants is sufficiently sampled and is likely overcomplete for some sub-classes [47]. Like pixel space, the vastness of signal space means that natural signals are inherently rare with the vast majority of signals containing no information [48]. Even if future spectral measurements are naturally sparse in **L**, there is a lot of redundancy making **L** computationally heavy. This is where our learned Kanji can be useful. If **K** comprises the same features as **L**, then we can use **K** to derive a low-rank basis (with the default being uniform placement). Towards this end, we investigate the following sparse coding methods:

1. Singular value decomposition (SVD);
2. Symmetric non-negative matrix factorization (SymNMF);
3. Sparse dictionary learning (SDL);
4. Deep autoencoders (DAE).

4.1. Singular Value Decomposition

The simplest data-driven basis can be derived by computing the first r column vectors of the unitary matrix of the singular value decomposition [39]. Given some representative data \mathbf{K} , a basis Ψ_r can be found by solving:

$$\mathbf{K} = \mathbf{U}\Sigma\mathbf{V}^* \quad \text{where} \quad \Psi_r = \mathbf{U}_r\Sigma_r. \quad (5)$$

The main advantage of SVD is that it can be quickly and efficiently performed in most computational software packages. The resulting basis (columns of Ψ_r) are ordered with respect to the strength of their contribution in representing variance in the original data matrix. When plotted, it is evident that any similarity to real-world spectral power distributions is lost. Instead, we can think of the basis vectors as defining an abstract “feature space”. While the advantage of an SVD-derived basis is simplicity, the drawback is that the bases can only be used to find a linear map (i.e., a relatively simple relationship given the abilities of modern deep autoencoders). In fact, we can think of SVD as a special case of a DAE wherein the encoder weights describe a linear relationship.

4.2. Symmetric Non-Negative Matrix Factorization

While the SVD basis works well capturing the features of \mathbf{K} , it notoriously lacks interpretability when applied to physical systems where negative values may be meaningless. Within the context of designing a spectral imaging sensor, the response sensitivities of each channel must be positive because there is no physical way to interpret negative sensitivity. Towards this end, we implement *symmetric* non-negative matrix factorization (SymNMF) [49,50]. SymNMF overcomes the pitfalls of other algorithms insofar as it is capable of capturing *nonlinear* cluster structures (unlike standard “brother” non-negative matrix factorization). Even more interestingly, SymNMF optimization is independent of the eigenspace of the affinity matrix (unlike spectral clustering). Furthermore, the affinity matrix \mathbf{A} can be defined with respect to any appropriate distance metric given a priori knowledge of the datatype. The minimization problem for SymNMF is defined as

$$L(\Psi, r) = \underset{\Psi \geq 0}{\operatorname{argmin}} \|\mathbf{A} - \Psi\Psi'\|_F^2 \quad (6)$$

where r is the rank of Ψ . If the number of spectra in \mathbf{L} is n , then \mathbf{A} is a square $n \times n$ matrix where each element in \mathbf{A} corresponds to a measure of distance between observations. Formally, we define \mathbf{A} elementwise as

$$\mathbf{a}_{i,j} = \begin{cases} \mathcal{K}(\mathbf{k}_i, \mathbf{k}_j) & \text{if } i \neq j \\ 0 & \text{if } i = j \end{cases} \quad (7)$$

where $\mathcal{K}(\mathbf{k}_i, \mathbf{k}_j)$ is a similarity measure (e.g., Euclidean distance) between $\mathbf{k}_i, \mathbf{k}_j \in \mathbf{K}$. Here, Ψ_r has r columns corresponding to the learned basis vectors. One of the core benefits of SymNMF is that d can be selected via knowledge of the underlying datatype. In the most abstract applications, d may best be represented by information theoretic measures such as the normalized information and compression distances [51,52].

4.3. Sparse Dictionary Learning

Sparse dictionary learning (SDL) is a sub-domain of sparse representation that spans a number of algorithms, most notably: the method of optimal direction (MOD) [53,54]; k -singular value decomposition (K-SVD) [55]; and online dictionary learning (ODL) [56], which is commonly implemented for its competitive speed [57]. While it may appear that learned dictionaries are naturally superior, it is important to understand their benefits and shortcomings. First, SDL may be unnecessary if the data are naturally sparse in signal space [39]. Second, dictionary learning algorithms tend to be computationally expensive because they require multiple iterations to converge on the optimal solution. Regardless,

their applicability in sparse approximation should not be ignored and while the atoms do not necessarily retain their similarity to real-world spectra, they do share the most similarities than any of the other methods and do not produce negative values if the input data are nonnegative. We seek to find a dictionary Ψ_r by minimizing the following loss function reported in [58]:

$$L(\mathbf{y}, \Psi_t) = \operatorname{argmin}_{\mathbf{x}'} \|\mathbf{y} - \Psi_t \mathbf{x}'\|_2 + \beta \|\mathbf{x}'\|_1 \quad (8)$$

Here, β is a sparsity-promoting coefficient. This loss function is commonly referred to as the sparse-coding or LASSO regression [59]. LASSO balances sparsity with model complexity in order to promote a dictionary with low cross-validated error and is one proposed approach to solving the l_1 -minimization problem framed in Equation (3).

4.4. Deep Autoencoders

Autoencoders refer to a specific class of artificial neural networks whose aim is to learn the most efficient encoding of some data in a target low-dimensional representation (latent space) [60]. The architecture of an autoencoder is roughly represented by the sketch in Figure 1 where the encoding weights Ψ are learned via the loss function:

$$L(\mathbf{y}, \Psi) = \operatorname{argmin}_{\mathbf{x}'} \frac{1}{2} \|\hat{\mathbf{y}} - \mathbf{y}\|^2 + \beta \Omega_w + \lambda \Omega_s \quad (9)$$

where Ω_w acts as an l_2 penalty on the encoder weights and Ω_s enforces sparsity via the Kullback–Leibler (KL) divergence [61]. While the loss function for the autoencoder requires more unpacking than others, the key take-away is that the mean squared error (MSE) is minimized between the *learned* representation $\hat{\mathbf{y}} = f_\theta(\mathbf{y})$ and \mathbf{y} for some latent basis Ψ given some regularization constraints on the weights associated to each basis vector, and a sparsity constraint on the reconstruction of the output. In essence, the goals and ambitions are well aligned with the other methods but with an added degree of flexibility. Equation (9) is closely related to the sparse relaxed regularized regression (SR3) method [62] aimed at finding a less restrictive loss function.

4.5. Implementation via QU Factorization

To summarize our steps up to this point, we amassed a library $\mathbf{L} \in \mathbb{R}^n$ of available online datasets without processing it in any way. We then used a subset $\mathbf{K} \in \mathbb{R}^m$ as an input to a number of well-known sparse-coding methods to arrive at four candidate low-rank bases $\Psi_i \in \mathbb{R}^r$ for $n \gg m > r$. We now want to use these bases to design different non-uniform encoders by assigning response functions to the pivot points derived via QU factorization [63]:

$$\Psi_i \mathbf{P}_i = \mathbf{Q}\mathbf{U} \quad (10)$$

Here, a pivot matrix \mathbf{P} is derived for each of the four “data-driven” bases (n.b., the uniform basis is not subscripted). Equation (10) can be solved using preset commands in most computational suites, taking the basis Ψ_i as the only input and outputting the pivot points. The resulting pivots (non-zero entries in \mathbf{P}) correspond to the peak wavelengths used to construct our encoder. From these points, we can define a response function by fitting a Gaussian distribution consistent with most available filters used in array-type spectrometers [14]. The generic structure for a sensor with approximate Gaussian responsibility over a wavelength range $\lambda \in \Lambda$, which is defined as

$$R(\lambda, p) = \exp \left[-\frac{\lambda - \Lambda(P_{ij})}{2\sigma^2} \right] \quad (11)$$

where $\Lambda(P_{ij})$ is the peak wavelength corresponding to the non-zero element of the p -th column of \mathbf{P} (for a total of p channels), and σ is the full-width at half-maximum (FWHM). As it is infeasible to manufacture photodiode sensors with single-wavelength sensitivity

channels (the diodes themselves are made from semiconductors with limited physical properties), available filters have FWHM values such that $\sigma \in [18, 25]$ nm [14]. In Figure 4, we show example response functions constructed via Equation (11).

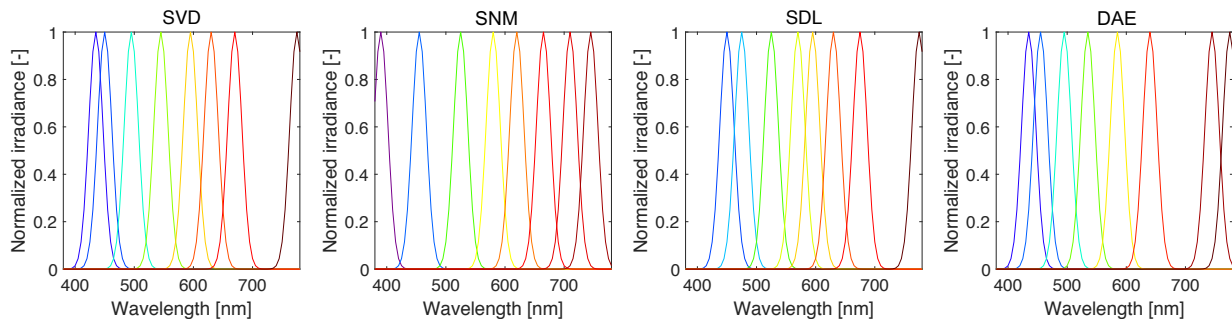


Figure 4. Response functions (shown here for $p = 8$, $\sigma = 18$ nm) constructed via Equation (11). Here, the peak responsivity corresponds to the pivot points derived via QU factorization.

An encoder, g_ϕ , with p channels, defines a measurement process:

$$g_\phi : \mathbb{R}^\infty \rightarrow \mathbb{R}^p \quad (12)$$

As with all “natural” non-bandlimited signals, $E(\lambda) \in \mathbb{R}^\infty$ requires that we make an assumption that $E(\lambda) \in \mathbb{R}^n$ where n is finite and $p \ll n$. When $p < 2K \log(n/K) + (7/5)K$, where K is the sparsity of the coefficient vector [33], then g_ϕ is a *lossy* compressor as the rate is below the minimums required for lossless reconstruction via Shannon–Nyquist and compressed sensing. We can further minimize information loss by designing g_ϕ in a way that leverages the domain knowledge that the underlying datatype g_ϕ will likely encode. To do this, we need to construct a simple mathematical model for g_ϕ that simulates the measurement process of a real-world physical sensor. This is performed by modeling the output current as proportional to the sum of the response function multiplied by the unknown spectral distribution:

$$\mathbf{y} = \epsilon(\lambda, p) + \sum_{i=1}^{\lambda_\top} R(i, p) E(i) \quad (13)$$

where $\lambda_\top = \max(\Lambda)$ and $\epsilon(\lambda, p)$ is the per-wavelength measurement error associated with the channels. If we assume the response functions are roughly Gaussian, the relative differences between measured points are only preserved if the response functions are the same width (which is not always the case in real-world sensors).

5. Spectra: An Adaptive Reconstruction Framework

Reconstruction is broadly characterized by three regimes depending on the accessibility of information: matching, reconstruction and interpolation. When signals are easily describable in \mathbf{K} , there is little risk of overfitting, but when signals exhibit unseen features, adding a greater number of elements to \mathbf{K} does not always imply greater reconstruction performance. What we seek to accomplish with SpecRA, is to develop an adaptive framework that uses the knowledge of \mathbf{K} , together with the measured signal \mathbf{y} , to triage observations as they arrive at the sensor in order to apply the reconstructive method with the highest probability of success. The core decision making comes down to the value of the following measure, which we refer to as the relative “rarity” of the measurement \mathbf{y} relative to the encoded elements in the reference set $\mathbf{k}_i \in \mathbf{K}$:

$$\mathcal{R} = \frac{1}{m} \sum_{i=1}^m \mathcal{K}(\mathbf{y}, g_\phi(\mathbf{k}_i)) \quad (14)$$

The challenge with this approach is that measured signals lose much of their unique features during the encoding process. In the previous section we outlined a workflow to improve the preservation of feature structure by deriving data-driven non-uniform sampling alternatives. Here, we aim to benefit from this groundwork in order to demonstrate the superiority of the proposed, integrated, method outlined in Figure 5.

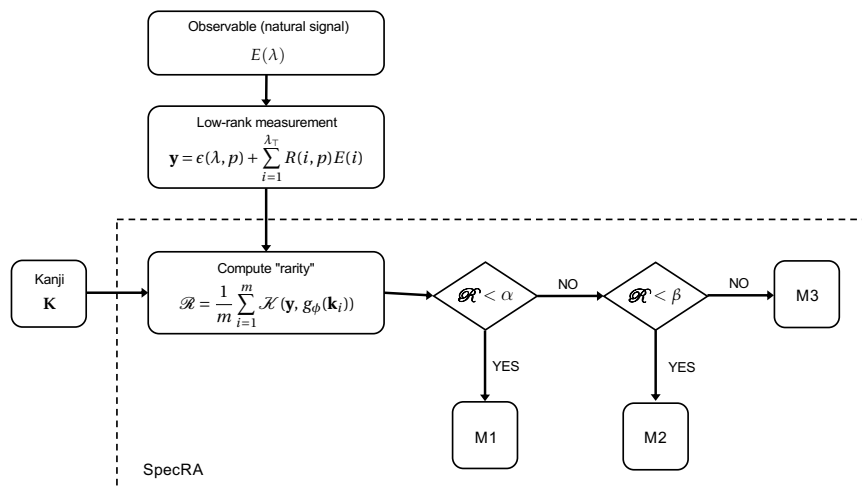


Figure 5. The SpecRA algorithm sorts raw measurements based on their rarity \mathcal{R} relative to a Kanji \mathbf{K} and applies an appropriate minimizing optimization (i.e., M1, M2, or M3). This process helps prevent both underfitting and overfitting, making for an adaptive and more generalizable framework. In general, the greater the \mathcal{R} , the greater algorithmic “work” is required.

The three minimizing optimization processes are defined as follows:

- M1 Return the nearest match \mathbf{k}_i such that $\mathcal{R} = \min(\mathcal{R})$;
- M2 Solve $\min\|\mathbf{x}\| + \|\mathbf{x} - \mathbf{w}\|_1$ for $w_{j=i} = 1; w_{j \neq i} = 0$, subject to $\mathbf{y} = \mathbf{x}g_\phi(\mathbf{K})$;
- M3 Solve $\min\|\mathbf{x}\|$ subject to $\mathbf{y} = \mathbf{x}g_\phi(\mathbf{K})$.

To determine the triage constants α and β , we incrementally increased each value while repeatedly applying the algorithm on randomized testing and validation partitions. We also realized that β should depend on the number of channels p of the encoder. This is because lower-dimensional data will more effectively “mask” the rarity of the signal (i.e., leading to greater metamerism). For this reason, we repeated this analysis for $p = (1, \dots, 25)$ channels. While we experimented with many different relationships, the ansatz that β scales with the inverse square-root of p was the most successful. Consequently, we were able to determine the following relationship for the available spectral data:

$$\beta = \frac{L_\beta}{\sqrt{p}} \tag{15}$$

where L is the target loss (of the reconstruction) determined by the spectral angle mapper [64]. In practice, $L_\beta = 0.05$. For α , we found that this relationship also holds—albeit for a smaller target loss such that $L_\alpha = 0.01$.

6. Non-Uniform Performance Dynamics

In this section, we present results from simulated data for which we compare the viability of the learned non-uniform sampling protocols against the uniform reference. We compute the mean errors and plot their distribution as a function of method and rate. Additionally, we test our hypotheses regarding the correlation between loss, signal complexity, and signal rarity.

6.1. Comparison to Existing Approaches

To obtain an idea of the differences between SpecRA and other competing approaches, we present results from the undersampled regime ($p = 3$). Here, Figure 6, we simulated the response that a tristimulus sensor would have with uniform and non-uniform responsivity. As expected, the information being fit is too coarse for the Fourier modes to find a fit in the measured dimension. LASSO fails to find a parsimonious fit and succumbs to overfitting. Because of the adaptability of SpecRA, the fit is more balanced and the ansatz made in the low-rank space is not far from the ground truth. Furthermore, we can see that the performance increases when structure is preserved by finding a more optimal low-rank encoding.

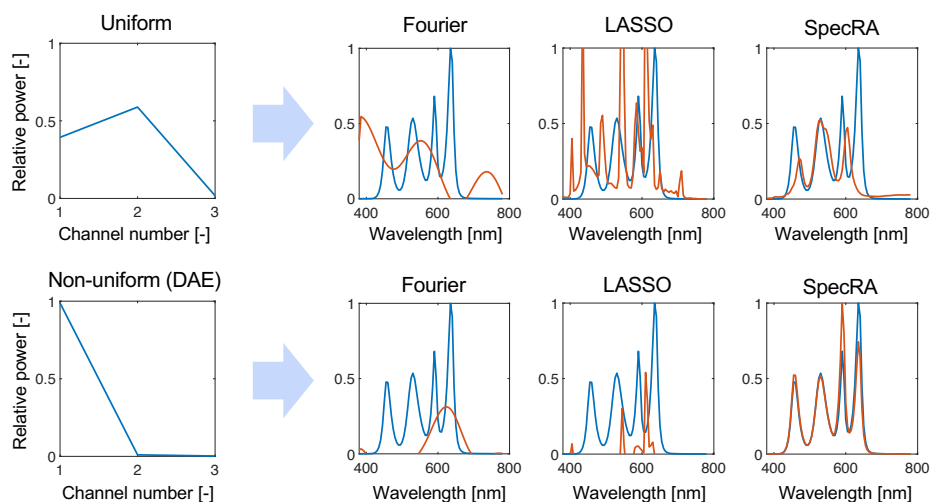


Figure 6. Comparing the performance of the algorithm in the “extreme” undersampled case of $p = 3$, we see how SpecRA outperforms the other methods, mostly by avoiding overfitting. Furthermore, the performance is improved when the signal is encoded with a learned non-uniform protocol (in this case, derived from the weights learned by training a deep autoencoder network).

While reconstruction with Fourier modes may be more appropriate in higher dimensional reconstruction (i.e., $p > 10$), the advantages are only seen when the measured signal is typologically distant from the reference set of prior observations. SpecRA takes a simple yet effective approach: maximizing the available information and not overfitting.

6.2. Loss as a Function of Method and Rate

As we are working with spectral data, we reported the reconstruction error (loss) terms of the spectral angle mapper (SAM) defined as

$$L = \mathcal{K}(\mathbf{s}, \hat{\mathbf{s}}) = \cos^{-1} \frac{\sum_{i=1}^n s_i \hat{s}_i}{\left(\sum_{i=1}^n s_i^2\right)^{1/2} \left(\sum_{i=1}^n \hat{s}_i^2\right)^{1/2}}. \quad (16)$$

In order to compare results, we first split our library \mathbf{L} into five random training, \mathbf{T} , and validation \mathbf{V} sets. The training sets comprised $n_T = 300$ spectra while the validation sets comprised $n_V = 1417$ (i.e., the remaining signals in \mathbf{L} after removing the 146 theoretical sources). We then constructed our encoder using the training set to first derive \mathbf{K} , then $\mathbf{\Psi}$, and finally \mathbf{P} used to construct the response functions \mathbf{R} for g_θ . Then, we simulated the measurement process of the spectra in the validation set by Equation (13) with $\epsilon(\lambda, p) = 0$ (i.e., for comparative analysis and applications where simulated data are used, e.g., rendering) for $p = (1, \dots, 25)$. We then reconstructed the measured spectra via the SpecRA algorithm, and computed the reconstruction loss defined above.

What we see in Figure 7 is the type of Pareto distribution we expect to see in such experiments. As the number of channels increases, eventually a plateauing effect is ob-

served wherein adding more channels to the encoder does not result in greater returns in minimizing loss. What is interesting is that the uniform approach is very clearly the exception resulting in a loss 150% that of the nearest non-uniform approach for an encoder with two channels. Aside from the mean, we can also plot the individual loss per reconstructed spectrum in the validation set. In Figure 8, we can see how plotting the uniform losses against the non-uniform losses directly evidences the efficacy of the underlying method. While it is clear from Figure 7, that non-uniform methods are competitive for $p < 4$, we can see how this competitive edge is maintained for larger p . Even after the plateauing of the mean, we can see that the distribution of losses exhibits a bias for non-uniform methods (c.f., SVD, SNM, and SDL, for $p = 12$). Interestingly, this effect is not observed for responses derived via the autoencoder weights which are presumed to be a generalization of the SVD modes demonstrating competitive results.

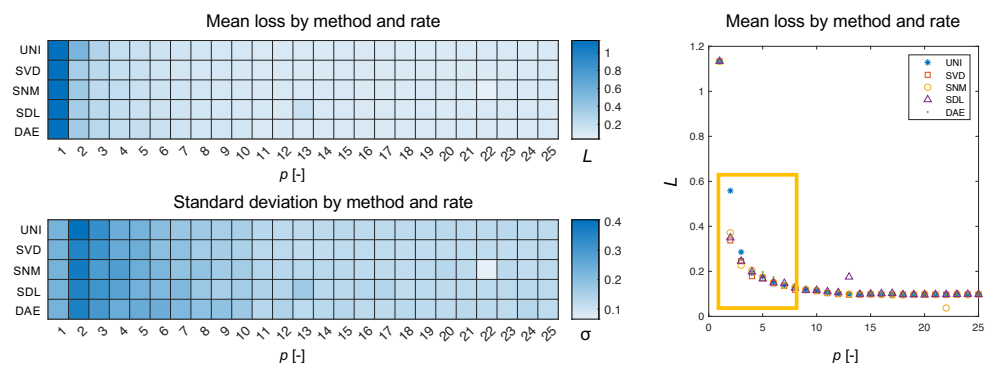


Figure 7. Here, we show the mean loss and standard deviation for each method as a function of the sampling rate. The methods are abbreviated by three-letter codes for visual clarity. The mean SAM loss, L , is displayed as a heat-map and as a Pareto plot. The difference between methods is greatest in for $p < 5$ (boxed region).

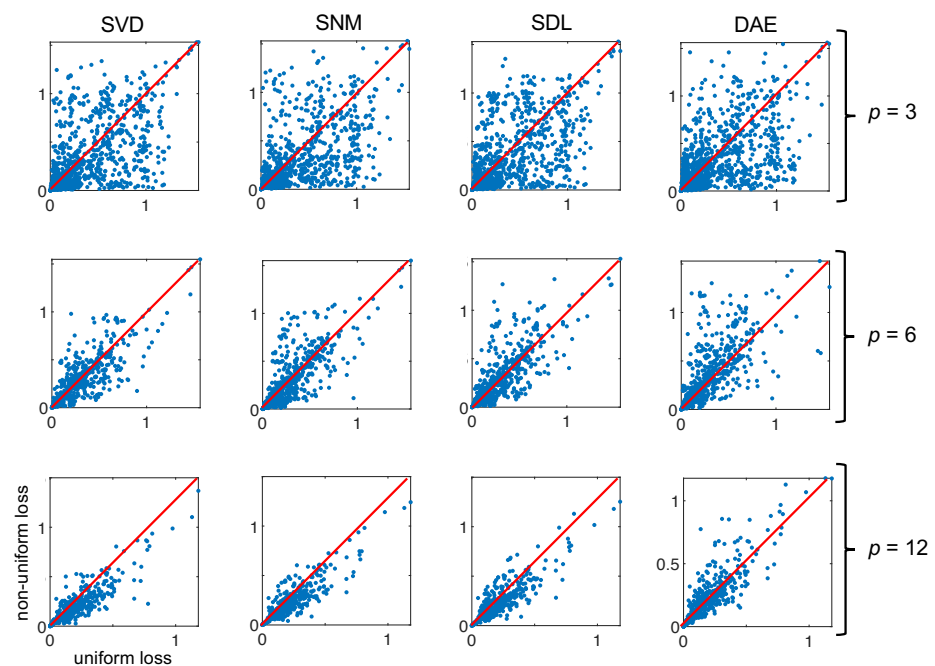


Figure 8. Here, we plot the uniform loss against the non-uniform loss for the four derived sampling protocols at three different encoder rates for all the spectra in our validation set ($n_V = 1417$). When the majority of the points fall below the $y = x$ line (shown in red), this means that the non-uniform approach outperforms the uniform method (e.g., SVD, $p = 12$).

6.3. Correlation between Loss, Signal Complexity, and Rarity

In addition to comparing methods, we had two core hypotheses: loss will be correlated with signal complexity and rarity. To test these hypotheses we computed the signal complexity by taking the standard deviation of the signal (as a proxy) and the signal rarity by computing the mean of the 10 most similar (according to Equation (16)) signals in the training set. Since SpecRA enforces sparsity, the number of spectra used in the reconstruction generally does not exceed 10 (or else it would suffer from overfitting). Towards this end, we compute these metrics and present the results in Figure 9. Interestingly, there is little to no correlation between the complexity of the signal and reconstruction via SpecRA while, we see a notable trend in the loss when plotted against the rarity of the signal relative to the training set (especially for lower rates). The lack of correlation with signal complexity could be a result of how the SpecRA switches between reconstruction methodologies and also explain why signal rarity does not exhibit a stronger correlation. If the spectra in the training library or predefined set, \mathbf{K} , are similar to the measured spectra, SpecRA will be effective at finding either a direct match or a sparse approximation, even if the rate is small. At the same time, we see detrending as the rate increases, indicating that rarity is less important when more information is available.

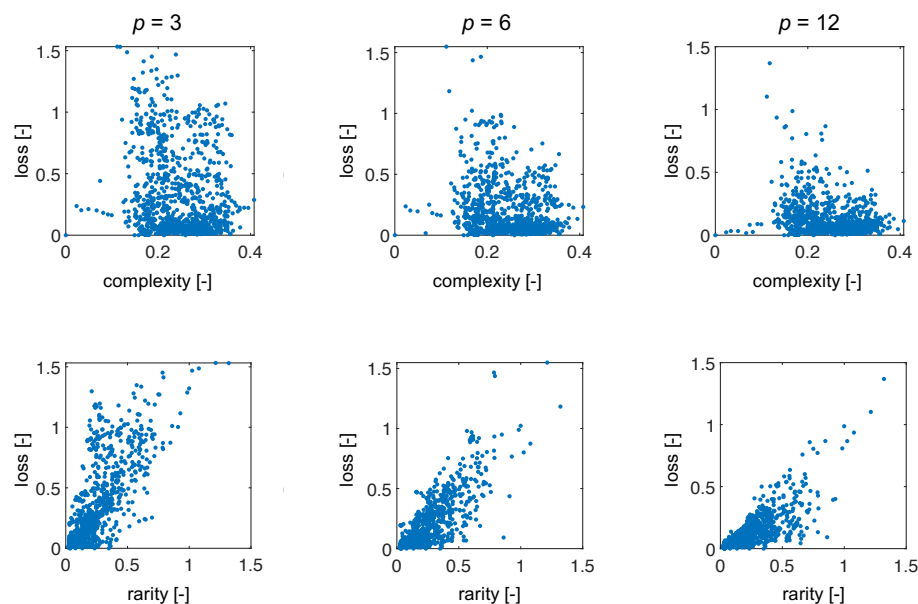


Figure 9. Here, we show the results for loss as a function of complexity (top row) and as a function of rarity (bottom row) for three representative encoder rates.

The results of our preliminary analysis are promising and generally consistent with our hypotheses. One interesting question that remains unanswered here is the degree to which a signal can differ from the reference set, \mathbf{K} , and still be effectively reconstructed from a sparse combination of prior observations.

7. Discussion and Future Integration

Formally, reconstruction, in the context of signal processing, is any process that recovers a signal from the set of points. Whether this is performed via a codebook or regression, the end goal is the same: recover information that is not readily available (i.e., via interpolation). Data-driven reconstruction methods apply some inductive bias and in the case of SpecRA, we exploit *Lex Parsimoniae* by toggling between pattern matching, l_1 -minimization via a primal-dual algorithm, and linear interpolation. As more information becomes available via online repositories, tailored bases will likely outperform generic ones such as that of Fourier. As is the case with all data-driven methods, the success is still highly linked

to the availability and tractability of the underlying data (which can be a challenge for hyperspectral imaging).

In this paper, we simulated response functions for uniform and non-uniform sampling methods based on exploiting regularities in the frequency domain. Applying real-world constraints on idealized mathematical models is a challenge in any discipline. In the case of visible spectral data, there is a clear effect of information saturation implying that finer resolution is perhaps not needed to capture sufficiently unique information about the spectrum. On the other hand, the relative rarity of a signal does play a role in constraining the possible loss. Towards this end, applying online learning algorithms to sift through repositories to construct interpretable low-rank reference libraries, \mathbf{K} , is paramount.

Furthermore, systematically selecting filter locations via a combined process of data-driven analysis and matrix factorization can dramatically improve results, especially for low-rate encoding (i.e., $p < 4$). A clear benefit of our method is its independence from fitted optimization parameters for a set of priors, and prior information in the form of a measured signal does improve results. Finally, our approach remains untested on real-world sensing hardware. It remains unclear what effect increasing channel errors will have on the reconstruction process. This will be investigated in future work. In applying this work to other domains, it is important to note that while many of these methods are generalizable, there will always be specific design constraints that need to be considered for any new class of sensors. In conclusion, spectral sensing spans many disciplines and its relevance in autonomous systems is becoming even more present. The systems and methods outlined in this paper provide a working template for research into the design and implementation of compact spectrometers and rendering software for a diversity of applications. While conventional spectrometry is seen as costly and data intensive, taking advantage of domain knowledge and sparse optimization can offer a valuable alternative to existing methods. We hope that this work provides a foundation for both theoretical and practical future developments.

Author Contributions: Conceptualization and theorization, F.S.W.; methodological approach, F.S.W.; software development and source code, F.S.W.; cross validation, F.S.W.; resource allocation, M.A.; data curation, F.S.W.; writing of the original draft preparation, F.S.W.; writing of the review and editing, F.S.W., M.A. and M.S.; figure generation, F.S.W.; supervision, M.A. and M.S.; project administration, M.A.; funding acquisition, M.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the INNOSUISSE grant number 40598.1 IP-ICT and the EPFL.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: <https://github.com/fwebler> (accessed on 17 March 2022).

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Priemer, R. *Introductory Signal Processing*; World Scientific: Singapore, 1991; Volume 6.
2. Hey, A.J.; Tansley, S.; Tolle, K.M. *The Fourth Paradigm: Data-Intensive Scientific Discovery*; Microsoft Research Redmond: Redmond, WA, USA, 2009; Volume 1.
3. Lu, D.; Batistella, M.; Moran, E.; Mausel, P. Application of spectral mixture analysis to Amazonian land-use and land-cover classification. *Int. J. Remote Sens.* **2004**, *25*, 5345–5358. [[CrossRef](#)]
4. Boukhayma, A.; Caizzone, A.; Enz, C. An ultra-low power PPG and mm-resolution ToF PPD-based CMOS chip towards all-in-one photonic sensors. *IEEE Sens. J.* **2019**, *19*, 11858–11866. [[CrossRef](#)]
5. Garcia-Allende, P.B.; Conde, O.M.; Mirapeix, J.; Cobo, A.; Lopez-Higuera, J.M. Quality control of industrial processes by combining a hyperspectral sensor and Fisher's linear discriminant analysis. *Sens. Actuators B Chem.* **2008**, *129*, 977–984. [[CrossRef](#)]
6. Glassner, A.S. How to derive a spectrum from an RGB triplet. *IEEE Comput. Graph. Appl.* **1989**, *9*, 95–99. [[CrossRef](#)]
7. Smits, B. An RGB-to-spectrum conversion for reflectances. *J. Graph. Tools* **1999**, *4*, 11–22. [[CrossRef](#)]

8. Arad, B.; Ben-Shahar, O. Sparse recovery of hyperspectral signal from natural RGB images. In Proceedings of the European Conference on Computer Vision, Munich, Germany, 8–16 October 2016; Springer: Berlin/Heidelberg, Germany, 2016; pp. 19–34.
9. Fu, Y.; Zheng, Y.; Zhang, L.; Huang, H. Spectral reflectance recovery from a single rgb image. *IEEE Trans. Comput. Imaging* **2018**, *4*, 382–394. [[CrossRef](#)]
10. Jia, Y.; Zheng, Y.; Gu, L.; Subpa-Asa, A.; Lam, A.; Sato, Y.; Sato, I. From rgb to spectrum for natural scenes via manifold-based mapping. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 4705–4713.
11. Lin, Y.T.; Finlayson, G.D. Exposure invariance in spectral reconstruction from rgb images. In Proceedings of the Color and Imaging Conference, Society for Imaging Science and Technology, Paris, France, 21–25 October 2019; pp. 284–289.
12. Van der Meer, F.D.; Van der Werff, H.M.; Van Ruitenbeek, F.J.; Hecker, C.A.; Bakker, W.H.; Noomen, M.F.; Van Der Meijde, M.; Carranza, E.J.M.; De Smeth, J.B.; Woldai, T. Multi-and hyperspectral geologic remote sensing: A review. *Int. J. Appl. Earth Obs. Geoinf.* **2012**, *14*, 112–128. [[CrossRef](#)]
13. Lapray, P.J.; Wang, X.; Thomas, J.B.; Gouton, P. Multispectral filter arrays: Recent advances and practical implementation. *Sensors* **2014**, *14*, 21626–21659. [[CrossRef](#)]
14. Crocombe, R.A. Portable spectroscopy in 2019: Smaller, cheaper and in consumer products? Next-Generation Spectroscopic Technologies XII. *Int. Soc. Opt. Photonics* **2019**, *10983*, 109830J.
15. Hagen, N.A.; Kudenov, M.W. Review of snapshot spectral imaging technologies. *Opt. Eng.* **2013**, *52*, 090901. [[CrossRef](#)]
16. Monno, Y.; Kikuchi, S.; Tanaka, M.; Okutomi, M. A practical one-shot multispectral imaging system using a single image sensor. *IEEE Trans. Image Process.* **2015**, *24*, 3048–3059. [[CrossRef](#)] [[PubMed](#)]
17. Gehm, M.E.; John, R.; Brady, D.J.; Willett, R.M.; Schulz, T.J. Single-shot compressive spectral imaging with a dual-disperser architecture. *Opt. Express* **2007**, *15*, 14013–14027. [[CrossRef](#)] [[PubMed](#)]
18. Li, Y.; Majumder, A.; Zhang, H.; Gopi, M. Optimized multi-spectral filter array based imaging of natural scenes. *Sensors* **2018**, *18*, 1172. [[CrossRef](#)] [[PubMed](#)]
19. Li, Y.; Wang, C.; Zhao, J. Locally linear embedded sparse coding for spectral reconstruction from RGB images. *IEEE Signal Process. Lett.* **2017**, *25*, 363–367. [[CrossRef](#)]
20. Li, Y.; Wang, C.; Zhao, J.; Yuan, Q. Efficient spectral reconstruction using a trichromatic camera via sample optimization. *Vis. Comput.* **2018**, *34*, 1773–1783. [[CrossRef](#)]
21. Dun, X.; Ikoma, H.; Wetzstein, G.; Wang, Z.; Cheng, X.; Peng, Y. Learned rotationally symmetric diffractive achromat for full-spectrum computational imaging. *Optica* **2020**, *7*, 913–922. [[CrossRef](#)]
22. Li, Y.; Fu, Q.; Heidrich, W. Multispectral illumination estimation using deep unrolling network. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, BC, Canada, 1–17 October 2021; pp. 2672–2681.
23. Oiknine, Y.; Arad, B.; August, I.; Ben-Shahar, O.; Stern, A. Dictionary based hyperspectral image reconstruction captured with CS-MUSI. In Proceedings of the 2018 9th Workshop on Hyperspectral Image and Signal Processing, Evolution in Remote Sensing (WHISPERS), Amsterdam, The Netherlands, 23–26 September 2018; IEEE: New York, NY, USA, 2018; pp. 1–5.
24. August, Y.; Stern, A. Compressive sensing spectrometry based on liquid crystal devices. *Opt. Lett.* **2013**, *38*, 4996–4999. [[CrossRef](#)]
25. Oiknine, Y.; August, I.; Farber, V.; Gedalin, D.; Stern, A. Compressive sensing hyperspectral imaging by spectral multiplexing with liquid crystal. *J. Imaging* **2019**, *5*, 3. [[CrossRef](#)]
26. Zhang, S.; Dong, Y.; Fu, H.; Huang, S.L.; Zhang, L. A spectral reconstruction algorithm of miniature spectrometer based on sparse optimization and dictionary learning. *Sensors* **2018**, *18*, 644. [[CrossRef](#)]
27. Bao, J.; Bawendi, M.G. A colloidal quantum dot spectrometer. *Nature* **2015**, *523*, 67–70. [[CrossRef](#)]
28. Yang, Z.; Albrow-Owen, T.; Cui, H.; Alexander-Webber, J.; Gu, F.; Wang, X.; Wu, T.C.; Zhuge, M.; Williams, C.; Wang, P.; et al. Single-nanowire spectrometers. *Science* **2019**, *365*, 1017–1020. [[CrossRef](#)] [[PubMed](#)]
29. Chang, C.C.; Su, Y.J.; Kurokawa, U.; Choi, B.I. Interference rejection using filter-based sensor array in VLC systems. *IEEE Sens. J.* **2011**, *12*, 1025–1032. [[CrossRef](#)]
30. Choi, B.I.; Lee, B.; Song, M.K. Nano-optic Filter Array Based Sensor. US Patent 9,395,473, 19 July 2016.
31. Candès, E.J. Compressive sampling. In Proceedings of the International Congress of Mathematicians, Madrid, Spain, 22–30 August 2006; Volume 3, pp. 1433–1452.
32. Shannon, C.E. Communication in the presence of noise. *Proc. IRE* **1949**, *37*, 10–21. [[CrossRef](#)]
33. Candès, E.J.; Romberg, J.; Tao, T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory* **2006**, *52*, 489–509. [[CrossRef](#)]
34. Chang, C.I. Spectral information divergence for hyperspectral image analysis. In Proceedings of the IEEE 1999 International Geoscience and Remote Sensing Symposium, IGARSS'99 (Cat. No. 99CH36293), Hamburg, Germany, 28 June–2 July 1999; Volume 1, pp. 509–511.
35. Du, Y.; Chang, C.I.; Ren, H.; Chang, C.C.; Jensen, J.O.; D'Amico, F.M. New hyperspectral discrimination measure for spectral characterization. *Opt. Eng.* **2004**, *43*, 1777–1786.
36. Botero-Valencia, J.S.; Valencia-Aguirre, J.; Durmus, D.; Davis, W. Multi-channel low-cost light spectrum measurement using a multilayer perceptron. *Energy Build.* **2019**, *199*, 579–587. [[CrossRef](#)]

37. Brown, C.; Goncharov, A.; Ballard, Z.S.; Fordham, M.; Clemens, A.; Qiu, Y.; Rivenson, Y.; Ozcan, A. Neural network-based on-chip spectroscopy using a scalable plasmonic encoder. *ACS Nano* **2021**, *15*, 6305–6315. [[CrossRef](#)]
38. Clark, E.; Brunton, S.L.; Kutz, J.N. Multi-fidelity sensor selection: Greedy algorithms to place cheap and expensive sensors with cost constraints. *IEEE Sens. J.* **2020**, *21*, 600–611. [[CrossRef](#)]
39. Brunton, S.L.; Kutz, J.N. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*; Cambridge University Press: Cambridge, UK, 2019.
40. Meerdink, S.K.; Hook, S.J.; Roberts, D.A.; Abbott, E.A. The ECOSTRESS spectral library version 1.0. *Remote Sens. Environ.* **2019**, *230*, 111196. [[CrossRef](#)]
41. Baldridge, A.M.; Hook, S.; Grove, C.; Rivera, G. The ASTER spectral library version 2.0. *Remote Sens. Environ.* **2009**, *113*, 711–715. [[CrossRef](#)]
42. David, A.; Fini, P.T.; Houser, K.W.; Ohno, Y.; Royer, M.P.; Smet, K.A.; Wei, M.; Whitehead, L. Development of the IES method for evaluating the color rendition of light sources. *Opt. Express* **2015**, *23*, 15888–15906. [[CrossRef](#)] [[PubMed](#)]
43. IES. *IES TM-30-15, IES Method for Evaluating Light Source Color Rendition*; Illuminating Engineering Society of North America: New York, NY, USA, 2015.
44. Guo, X.; Houser, K.W. A review of colour rendering indices and their application to commercial light sources. *Light. Res. Technol.* **2004**, *36*, 183–197. [[CrossRef](#)]
45. Houser, K.W.; Wei, M.; David, A.; Krames, M.R.; Shen, X.S. Review of measures for light-source color rendition and considerations for a two-measure system for characterizing color rendition. *Opt. Express* **2013**, *21*, 10393–10411. [[CrossRef](#)] [[PubMed](#)]
46. Spitschan, M.; Aguirre, G.K.; Brainard, D.H.; Sweeney, A.M. Variation of outdoor illumination as a function of solar elevation and light pollution. *Sci. Rep.* **2016**, *6*, 26756. [[CrossRef](#)]
47. Webler, F.S.; Spitschan, M.; Foster, R.G.; Andersen, M.; Peirson, S.N. What is the ‘spectral diet’ of humans? *Curr. Opin. Behav. Sci.* **2019**, *30*, 80–86. [[CrossRef](#)]
48. Manohar, K.; Kaiser, E.; Brunton, S.L.; Kutz, J.N. Optimized sampling for multiscale dynamics. *Multiscale Model. Simul.* **2019**, *17*, 117–136. [[CrossRef](#)]
49. Kuang, D.; Ding, C.; Park, H. Symmetric nonnegative matrix factorization for graph clustering. In Proceedings of the 2012 SIAM International Conference on Data Mining (SIAM), Virtual Event, 29 April–1 May 2021; pp. 106–117.
50. Kuang, D.; Yun, S.; Park, H. SymNMF: Nonnegative low-rank approximation of a similarity matrix for graph clustering. *J. Glob. Optim.* **2015**, *62*, 545–574. [[CrossRef](#)]
51. Li, M.; Badger, J.H.; Chen, X.; Kwong, S.; Kearney, P.; Zhang, H. An information-based sequence distance and its application to whole mitochondrial genome phylogeny. *Bioinformatics* **2001**, *17*, 149–154. [[CrossRef](#)]
52. Cilibrasi, R.; Vitányi, P.M. Clustering by compression. *IEEE Trans. Inf. Theory* **2005**, *51*, 1523–1545. [[CrossRef](#)]
53. Engan, K.; Aase, S.O.; Husoy, J.H. Method of optimal directions for frame design. In Proceedings of the 1999 IEEE International Conference on Acoustics, Speech, and Signal Processing, Proceedings, ICASSP99 (Cat. No. 99CH36258), Phoenix, AZ, USA, 15–19 March 1999; Volume 5, pp. 2443–2446.
54. Engan, K.; Aase, S.O.; Husøy, J.H. Multi-frame compression: Theory and design. *Signal Process.* **2000**, *80*, 2121–2140. [[CrossRef](#)]
55. Aharon, M.; Elad, M.; Bruckstein, A. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Process.* **2006**, *54*, 4311–4322. [[CrossRef](#)]
56. Mairal, J.; Bach, F.; Ponce, J.; Sapiro, G. Online learning for matrix factorization and sparse coding. *J. Mach. Learn. Res.* **2010**, *11*, 19–60.
57. Skretting, K.; Engan, K. Learned dictionaries for sparse image representation: Properties and results. In Proceedings of the Wavelets and Sparsity XIV, International Society for Optics and Photonics, Bellingham, WA, USA, 27 September 2011; Volume 8138, p. 81381N.
58. Lee, H.; Battle, A.; Raina, R.; Ng, A.Y. Efficient sparse coding algorithms. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2007; pp. 801–808.
59. Tibshirani, R. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. (Methodol.)* **1996**, *58*, 267–288. [[CrossRef](#)]
60. Kramer, M.A. Nonlinear principal component analysis using autoassociative neural networks. *AIChE J.* **1991**, *37*, 233–243. [[CrossRef](#)]
61. Olshausen, B.A.; Field, D.J. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vis. Res.* **1997**, *37*, 3311–3325. [[CrossRef](#)]
62. Zheng, P.; Askham, T.; Brunton, S.L.; Kutz, J.N.; Aravkin, A.Y. A unified framework for sparse relaxed regularized regression: SR3. *IEEE Access* **2018**, *7*, 1404–1423. [[CrossRef](#)]
63. Gander, W. Algorithms for the QR decomposition. *Res. Rep.* **1980**, *80*, 1251–1268.
64. De Carvalho, O.A.; Meneses, P.R. Spectral correlation mapper (SCM): An improvement on the spectral angle mapper (SAM). In Proceedings of the Summaries of the 9th JPL Airborne Earth Science Workshop, JPL Publication 00-18, Pasadena, CA, USA, 23–25 February 2000; JPL Publication: Pasadena, CA, USA, 2000; Volume 9.