Differentiable Programming for Hyperspectral Unmixing using a Physics-based Dispersion Model

John Janiczek¹, Parth Thaker¹, Gautam Dasarathy¹, Christopher S. Edwards², Philip Christensen¹, and Suren Jayasuriya¹

Arizona State University
 Northern Arizona University

Abstract. Hyperspectral unmixing is an important remote sensing task with applications including material identification and analysis. Characteristic spectral features make many pure materials identifiable from their visible-to-infrared spectra, but quantifying their presence within a mixture is a challenging task due to nonlinearities and factors of variation. In this paper, spectral variation is considered from a physics-based approach and incorporated into an end-to-end spectral unmixing algorithm via differentiable programming. The dispersion model is introduced to simulate realistic spectral variation, and an efficient method to fit the parameters is presented. Then, this dispersion model is utilized as a generative model within an analysis-by-synthesis spectral unmixing algorithm. Further, a technique for inverse rendering using a convolutional neural network to predict parameters of the generative model is introduced to enhance performance and speed when training data is available. Results achieve state-of-the-art on both infrared and visible-to-near-infrared (VNIR) datasets, and show promise for the synergy between physicsbased models and deep learning in hyperspectral unmixing in the future.

Keywords: hyperspectral imaging, spectral unmixing, differentiable programming

1 Introduction

Hyperspectral imaging is a method of imaging where light radiance is densely sampled at multiple wavelengths. Increasing spectral resolution beyond a traditional camera's red, green, and blue spectral bands typically requires more expensive detectors, optics, and/or lowered spatial resolution. However, hyperspectral imaging has demonstrated its utility in computer vision, biomedical imaging, and remote sensing [38, 8]. In particular, spectral information is critically important for understanding material reflectance and emission properties, important for recognizing materials.

Spectral unmixing is a specific task within hyperspectral imaging with application to many land classification problems related to ecology, hydrology, and

mineralogy [25, 28]. It is particularly useful for analyzing aerial images from aircraft or spacecraft to map the abundance of materials in a region of interest. While pure materials have characteristic spectral features, mixtures require algorithms to identify and quantify material presence.

A common model for this problem is linear mixing, which assumes electromagnetic waves produced from pure materials combine linearly and are scaled by the material abundance. Mathematically this is expressed as $\mathbf{b} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}$ where **b** is the observed spectra, **A** is a matrix whose columns are the pure material spectra, $\boldsymbol{\eta}$ is the measurement noise, and **x** is the abundance of each pure material. The model assumes that the pure material spectra, referred to as endmember spectra, is known before-hand. Nonlinear effects are known to occur when photons interact with multiple materials within a scene for which we refer readers to the review by Heylen et al. for techniques to account for the non-linear mixing [25].

A key challenge that affects both linear and nonlinear mixing models is that pure materials have an inherent variability in their spectral signatures, and thus cannot be represented by a single characteristic spectrum. Spectral variability of endmembers is caused by subtle absorption band differences due to factors such as different grain sizes [39, 48, 43, 44] or differing ratios of molecular bonds [7, 51] as shown in Figure 1. Since variability causes significant errors in unmixing algorithms, it is an active area of research [60, 16, 65].



Fig. 1. Endmember Variation: Several spectra of olivine are plotted to demonstrate it's spectral variability. The olivine mineral is a solid solution with continuous compositional variation of Iron (Fe₂) and Magnesium (Mg₂) bonds. This ratio of bonds (indexed by the Fo number), causes absorption bands to shift in frequency and strength

Recently, differentiable programming has become a popular research area due to its potential to bridge gaps between physics-based and machine learningbased techniques for computer vision and graphics [53, 21, 1]. Our key insight is to leverage differentiable programming by modeling the variation of spectra with a physics-based dispersion model, and incorporating this differentiable model into an end-to-end spectral unmixing algorithm. Such an approach has the capacity to unmix scenes with a large amount of variability, while constraining the predictions to be physically plausible. These physically plausible variations of endmember spectra also provide additional science data as the variation of absorption bands can reveal properties about the composition and history of the material. To our knowledge, we are the first to use a generative physics model to account for spectral variability in an unmixing algorithm.

Contributions: Our specific contributions in this paper are the following:

- We introduce a physics-based dispersion model (first presented in [49, 30, 55]) to generate and render spectral variation for various pure materials.
 We provide an efficient optimization method via gradient descent to find dispersion model parameters for this spectral variation.
- We incorporate this dispersion model into an end-to-end spectral unmixing algorithm utilizing differentiable programming to perform analysis-bysynthesis optimization. Analysis-by-synthesis is solved via alternating minimization optimization and requires no training data.
- We further design an inverse rendering algorithm consisting of a convolutional neural network to jointly estimate dispersion model parameters and mineral abundances for spectral unmixing. This method requires training data, but is computationally efficient at test time and outperforms analysisby-synthesis and other state-of-the-art methods.

We provide extensive analysis of our proposed methods with respect to noise and convergence criteria. To validate our contributions, we test on both synthetic and real datasets using hyperspectral observations in the visible and near infrared (VNIR), and mid to far infrared (IR). The datasets also span three different environments from laboratory, aircraft, and satellite based spectrometers. Our methods achieve state-of-the-art across all datasets, and we compare against several baselines from literature. Our code is openly available and accessible here: https://github.com/johnjaniczek/InfraRender. We hope this work inspires more fusion between physics models and machine learning for hyperspectral imaging and computer vision more generally in the future.

2 Related Work

Optimization-based Approaches. Standard optimization techniques for linear unmixing include projection, non-negative least squares, weighted least squares, and interior point methods [43, 24, 47, 12, 23]. Further, sparsity-based optimization can improve abundance prediction [63, 9]. However, most optimization have not leveraged physics priors as we do in our model.

Spectral Variability. Spectral variability has been a topic of recent interest [60, 6]. One approach is to augment **A** with multiple variations or spectra for each endmember. To do this, multiple endmember spectral mixture analysis (MESMA) [46] and multiple-endmember linear spectral unmixing (MEL-SUM) [15] both require labeled data of the spectral variation for each endmember. In contrast, unsupervised techniques learn endmember sets from unlabelled hyperspectral images, including semi-automated techniques [3], k-means clustering [3], and the sparsity promoting iterated constrained endmember algorithm (SPICE) [58, 59] which simultaneously finds endmember sets while unmixing

for material abundances. These techniques are limited by the amount of sets in the endmember library, and computational complexity increases with more additions. Our method by contrast finds an efficient parameter set to physically model the spectral variation.

Another category of endmember variability techniques models the endmember spectral variation as samples from a multivariate distribution $\mathbf{P}(\mathbf{e}|\theta)$ where \mathbf{e} is the endmember spectra, and θ are the distribution parameters. Common statistical distributions proposed include the normal compositional model [50], Gaussian mixture models [65], and the beta compositional model [16]. These distribution models have large capacity to model spectral variations, however sometimes they can render endmember spectra that are not physically realistic.

Deep Learning for Hyperspectral Classification and Unmixing. Deep learning has recently improved many hyperspectral imaging tasks [62, 33]. In particular, networks process hyperspectral pixel vectors using both deep belief networks [36] and CNNs [26, 10, 56]. For spatial hyperspectral data, CNNs [11], joint spectral-spatial feature extraction [64], and 3D CNNs [35] are used. All these methods require large hyperspectral datasets that are annotated correctly. One of our methods uses analysis-by-synthesis and differentiable programming to avoid low training data issues, but our technique can be made complementary to deep learning architectures as we show in our inverse rendering CNN.

Differentiable Programming and Rendering. Differentiable programming refers to the paradigm of writing algorithms which can be fully differentiated end-to-end using automatic differentiation for any parameter [53, 54, 4]. This has been applied for audio [18] and 3D geometry processing [45]. In graphics, differentiable rendering has improved ray tracing [34, 40, 37, 61], solved analysis-by-synthesis problems in volumetric scattering [22, 21], estimated reflectance and lighting [1], and performed 3D reconstruction [52]. In our paper, we write a forward imaging model utilizing the physics of dispersion in spectral variation to allow our pipeline to be differentiable end-to-end.

3 Method

Our approach to hyperspectral unmixing features two main components: (1) use of a physically-accurate dispersion model for pure endmember spectra, and (2) a differentiable programming pipeline to perform spectral unmixing. This approach has synergistic benefits of leveraging prior domain knowledge while learning from data. Our first algorithm solves spectral unmixing in a self-supervised fashion using analysis-by-synthesis optimization with the dispersion model as the synthesis step. Further, we show how inverse rendering via a convolutional neural network (CNN) can learn parameters of this model to help speed up our end-to-end pipeline and improves performance when training data is available.

3.1 Dispersion Model

We first describe the dispersion model for generating endmember spectra. Endmember and/or endmember spectra is what we call the spectral curve for emissivity ϵ as a function of wavenumber ω . Each pure material has a characteristic endmember spectrum, although spectra can vary, which is the problem we are trying to solve/disambiguate. Let $\varepsilon^{\text{measured}}(\omega)$ be endmember spectra we have measured, typically in a lab or in the field, whose emissivity is sampled at different wavenumbers: $[\varepsilon^{\text{measured}}(\omega_1), \cdots \varepsilon^{\text{measured}}(\omega_N)]^T$. Our goal is to propose a model $\varepsilon^{\text{model}}(\boldsymbol{\Lambda};\omega)$ with parameters $\boldsymbol{\Lambda}$ such that the following loss is minimized: $L(\boldsymbol{\Lambda}) = \sum_{i=1}^{N} (\varepsilon^{\text{measured}}(\omega_i) - \varepsilon^{\text{model}}(\boldsymbol{\Lambda};\omega_i))^2$. That is, we fit the model emissivity of an endmember spectrum to the measured spectrum. In practice, we need to add regularization and constraints to this endmember loss for better fitting which we describe after the derivation of the dispersion model.

Derivation of the Dispersion model: Our model of endmember spectra is derived from an atomistic oscillator driven by electromagnetic waves impinging on the molecular structure of the pure material [49, 30]. In Figure 2, we show a conceptual diagram of this model, and how it generates emissivity curves as a function of wavelength. For the full derivation of the model from first principles, we refer the reader to Appendix A in the supplementary material. Instead, we outline the model below based on the equations derived from that analysis.



Fig. 2. Dispersion Model Concept Figure: The insight of the dispersion model is that optical properties can be related to molecular structure through first principles via an atomistic oscillator model. We use this generative model for the formation of spectral variation in our spectral unmixing algorithm

Let $\Lambda = [\rho, \omega_o, \gamma, \epsilon_r]$ be a matrix of parameters, where $\rho, \omega_0, \gamma, \epsilon_r \in \mathbb{R}^K$ and K is a model hyperparameter corresponding to the number of distinct massspring equations used to model the emissivity. ρ is the band strength, ω_o is the resonant frequency, γ is the frictional force (dampening coefficient), and ϵ_r is relative dielectric permeability. Please see the supplemental material for the physical significance of these parameters to the atomistic oscillator model, and their control over the shape of spectral absorption bands. Note: usually ϵ_r is a constant vector which does not vary with K. Thus $\Lambda \in \mathbb{R}^{K \times 4}$. The refractive index terms n, k are given as follows [49, 30]:

$$n(\boldsymbol{\Lambda};\omega) = \sqrt{\frac{\theta+b}{2}}, \quad k(\boldsymbol{\Lambda};\omega) = \frac{\phi}{n(\boldsymbol{\Lambda};\omega)},$$
 (1)

where the expressions for θ, b, ϕ are given as follows:

$$\theta = \epsilon_r + \sum_{k=1}^{K} 4\pi \rho_k \omega_{0_k}^2 \frac{(\omega_{0_k}^2 - \omega^2)}{(\omega_{0_k}^2 - \omega^2)^2 + \gamma_k^2 \omega_{0_k}^2 \omega^2},$$
(2)

$$b = \sqrt{\theta^2 + 4\phi^2}, \quad \phi = \sum_{k=1}^{K} 2\pi \rho_k \omega_{0_k}^2 \frac{\gamma_k \omega_{0_k} \omega}{(\omega_{0_k}^2 - \omega^2)^2 + \gamma_k^2 \omega_{0_k}^2 \omega^2}.$$
 (3)

We note that subscript k denotes the k-th coordinate of the corresponding vector. Also there is another useful relation (derived in Appendix A) that $n^2 - k^2 = \theta$, $nk = \phi$. We then define the complex refractive index as $\hat{n}(\boldsymbol{\Lambda}; \omega) = n(\boldsymbol{\Lambda}; \omega) - i \cdot k(\boldsymbol{\Lambda}; \omega)$, where $i = \sqrt{-1}$ is the imaginary number. Hence, we can calculate the emissivity as follows:

$$\epsilon(\boldsymbol{\Lambda};\omega) = 1 - R(\boldsymbol{\Lambda};\omega), \qquad \text{where } R(\boldsymbol{\Lambda};\omega) = \left|\frac{\hat{n}(\boldsymbol{\Lambda};\omega) - 1}{\hat{n}(\boldsymbol{\Lambda};\omega) + 1}\right|^2.$$
(4)

When considering minerals, we introduce $M \in \mathbb{N}$, the number of optical axes of symmetry in crystal structures, (eg. 2 axes of symmetry in quartz [49, 55]), to define the full model:

$$\varepsilon^{\text{model}}(\boldsymbol{\Lambda};\omega) = \sum_{m=1}^{M} \alpha_m \cdot \epsilon(\boldsymbol{\Lambda}_m;\omega) \quad \text{such that } \sum_{m=1}^{M} \alpha_m = 1, \alpha_m \ge 0, \quad (5)$$

where we use a different parameter matrix A_m and weight α_m for each optical axis of symmetry.

The dispersion model has been primarily used to analyze optical properties of materials to determine n and k, which then can be subsequently applied to optical models like radiative transfer [49, 55]. After n and k are found, spectra such as reflectance, emissivity, and transmissivity can be generated. In particular, we notice that fine-grained control of the dispersion model parameters can realistically render spectral variation that occurs in hyperspectral data. Our contribution is to leverage these properties in a differentiable programming pipeline for spectral unmixing.

Endmember Fitting. Using the dispersion model presented above, we want to robustly estimate the model parameters to fit the spectra $\varepsilon^{\text{measured}}$ captured in a lab or in the field. To fit the model, we wish to perform gradient descent to efficiently find these parameters. Using chain rule on the loss function, we see that $\frac{\partial L}{\partial A_{ij}} = \frac{\partial L}{\partial \varepsilon^{\text{model}}} \frac{\partial \varepsilon^{\text{model}}}{\partial A_{ij}}$, where (ij) corresponds to that element of the parameter matrix, and all expressions are scalars once the coordinate is specified. While the partial derivatives can be calculated explicitly via symbolic toolboxes, the resulting expressions are too long to be presented here. For simplicity and ease of use, we use the autograd function [4, 41] in PyTorch [42] to automatically compute derivatives for our model as we are performing backpropagation.

One main challenge in performing endmember fitting is that the dispersion model is not an injective function, and hence is typically not identifiable, that is more than one Λ can result in the same fit. This can be solved, in-part, through regularization to enforce sparsity, especially since a preference for fewer dispersion parameters has been suggested in the literature [49, 55]. In our implementation, we initialize our model with K = 50 rows of the parameter matrix. Since the parameter ρ controls the strength of the absorption band, small values of ρ do not contribute much energy to the spectra (unnecessary absorption bands), and can be pruned. After performing sparse regression by penalizing the L_1 norm of ρ , K is typically around 10-15 in our experiments.

Thus, our modified sparse regression problem may be written as

$$\underset{\boldsymbol{\Lambda}_{\min} \leq \boldsymbol{\Lambda} \leq \boldsymbol{\Lambda}_{\max}}{\arg\min} \sum_{i=1}^{N} \left(\epsilon^{real}(\omega_i) - \epsilon^{model}(\boldsymbol{\Lambda}; \omega_i) \right)^2 + \lambda_{\rho} ||\rho||_1, \tag{6}$$

where Λ_{\min} and Λ_{\max} restrict the variation of the dispersion parameters to a plausible range. In addition, endmembers (particularly minerals) can have multiple optical axes of symmetry described by separate spectra, which has been noted in the literature [49, 55]. Without prior knowledge of the number of axes for every material we encounter, we run this optimization for a single and double axes, and pick the one with the lowest error. See Section 4 for results on endmember fitting and Figure 6 for examples of modelled vs. measured spectra.

Despite the fact that this regression problem is non-convex, we solve it using gradient descent with a random initialization; this is known to converge to a local minimum with probability 1 [31]. A global minimum is not necessary at this stage, since we use endmember fitting to provide a good initialization point for the subsequent alternating minimization procedure introduced in the next subsection.

3.2 Differentiable Programming for End-to-End Spectral Unmixing

Analysis-by-Synthesis Optimization. In Figure 3, we show our full endto-end spectral unmixing pipeline. Here, $\varepsilon^{\text{model}}(\Lambda; \omega)$, which is initially fit to $\varepsilon^{\text{measured}}$, is then aggregated into the columns of **A**. Then, the observed spectra **b** is linearly unmixed by solving a regularized least-squares optimization: $\operatorname{argmin}_{\mathbf{x}} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_p$ subject to sum-to-one and non-negativity constraints $\|\mathbf{x}\|_1 = 1, \mathbf{x} \ge 0$. Given these constraints, one cannot impose sparsity with the usual L_1 norm. Instead, we use the L_p norm to induce sparsity for the predicted abundances; this has been proposed before for spectral unmixing [9].

The key to our pipeline is that everything is fully differentiable, and thus we can actually minimize the following equation:

$$\underset{\mathbf{x}, \boldsymbol{\Lambda} \in [\boldsymbol{\Lambda}_{min}, \boldsymbol{\Lambda}_{max}]}{\arg\min} \|\mathbf{b} - \mathbf{A}(\boldsymbol{\Lambda})\mathbf{x}\|_{2}^{2} + \lambda \|\mathbf{x}\|_{p} \text{ such that } \|\mathbf{x}\|_{1} = 1, \mathbf{x} \ge 0.$$
(7)

with respect to both the parameters of the dispersion model Λ and the abundances **x**. This gives us our recipe for hyperspectral unmixing: first, perform endmember fitting to initialize $\mathbf{A}(\Lambda)$, then, solve Equation 7 in an alternating

8 J. Janiczek et al.



Fig. 3. Analysis-by-Synthesis: The analysis-by-synthesis algorithm uses differentiable programming to find optimal dispersion parameters and abundances. The initial dispersion parameters and the target spectra are fed as inputs, and the algorithm alternates between optimizing the abundances in the least squares sense and updating the dispersion parameters with respect to the gradient

fashion for \mathbf{x} and $\boldsymbol{\Lambda}$. One could also solve this equation jointly for both unknowns, however, we found that the alternating optimization was faster and converged to better results.

The optimization problem established in equation (7) is an alternating minimization problem and is unfortunately not convex [27]. One popular approach to tackle nonconvex problems is to find a good initialization point [17, 5], and then execute a form of gradient descent. Inspired by this, we first initialize $\mathbf{A}(\boldsymbol{\Lambda})$ by performing endmember fitting using Equation 6 as described in the previous subsection. Our experiments indicate that this provides a useful initialization for our subsequent step. We then perform alternating minimization on Equation 7 for \mathbf{x} and $\boldsymbol{\Lambda}$. Note that each iterate of the resulting alternating minimization involves the solution of a subproblem which has a convergence rate which depends on the condition number of the matrix $\mathbf{A}(\boldsymbol{\Lambda})$. For more details on this, we refer the reader to Appendix C where we discuss on the properties of $\mathbf{A}(\boldsymbol{\Lambda})$ across multiple runs.

In the ideal scenario, this initial matrix $\mathbf{A}(\boldsymbol{\Lambda})$ would consist of the endmember spectra that fully characterizes the mixed spectra **b**. However, since spectra for the same material can significantly vary [39, 48, 43, 44, 7, 51] (see Figure 1), the initialization can be slightly off and we follow up with (7) to obtain a better fit. Note that this optimization problem is solving for the maximum likelihood estimator under a Gaussian noise model. Our optimization technique is performing **analysis-by-synthesis**, as given a single observation **b**, the dispersion model synthesizes endmember variation until a good fit is achieved.

Inverse Rendering of Dispersion Model Parameters. The previous analysis-by-synthesis optimization does not require training data (labeled abundances in spectral mixtures) in order to perform spectral unmixing. However, there is room for even more improvement by using labeled data to help improve the parameter fitting of the model in the synthesis step. We train a CNN to predict the parameters for a generative model, known as inverse rendering in other domains [57]. In Figure 4, we show this inverse rendering conceptually, and how it can be fed into our differentiable programming pipeline for end-to-end spectral unmixing.

Our CNN architecture consists of convolutional layers followed by a series of fully-connected layers. We refer the reader to the supplemental material for the exact network structure and implementation details. Using a CNN for inverse rendering is significantly faster at test time as compared to the analysis-bysynthesis optimization. However, it does have a drawback of requiring training data which is unavailable for certain tasks/datasets.



Fig. 4. Inverse Rendering: A CNN is trained to "inversely render" pixels of the hyperspectral image, by predicting both the dispersion parameters that control the spectral variability, and the abundances that control the mixing model. During training, the reconstruction error is back-propagated through the differentiable dispersion model to boost the performance of the network at making physically realistic predictions

4 Experimental Results

Datasets. We utilize three separate datasets to validate our spectral unmixing algorithms. In Figure 5, we visually represent these datasets and their exemplar data. For specific implementation details and dataset pre-processing, please see Appendix D in the supplemental material.

Feely et al. Dataset. We utilize 90 samples from the Feely et al. dataset [19] of thermal emission spectra in the infrared for various minerals measured in the lab. Ground truth was determined via optical petrography [19], and a labeled endmember library is provided. The limited amount of data is challenging for machine learning methods, so we utilize the dispersion model to generate 50,000 additional synthetic spectra for dataset augmentation.

Gulfport dataset. The Gulfport dataset from Gader et al. [20] contains hyperspectral aerial images in the VNIR along with ground truth classification labels segmenting pixels into land types (e.g. grass, road, building). Although the dataset is for spectral classification, it can also be used to benchmark unmixing algorithms by creating synthetic mixtures of pure pixels from the Gulfport dataset with random abundances as done by [16, 65]. We perform both spectral

classification (Gulfport) and unmixing (Gulfport synthetic) tasks in our testing. Both datasets are split into a train and test set (although some methods do not require training data), and the training data is augmented with 50,000 synthetically generated mixtures from the dispersion model.

One main difficulty of this dataset is the endmembers identified correspond to coarse materials such as grass and road as opposed to pure materials. Such endmembers can significantly vary across multiple pixels, but this spectral variation is not physically described by the dispersion model. To solve this problem, we utilize K-means clustering to learn examplar endmembers for each category (e.g. grass, road, etc). Then the resulting centroid endmember can be fit to the dispersion model to allow further variation such as absorption band shifts in the spectra. We found that K = 5 worked the best for the Gulfport dataset.

TES Martian Dataset. The Thermal Emission Spectrometer (TES) [13] uses Fourier Transform Infrared Spectroscopy to measure the Martian surface. We utilize pre-processing from Bandfield et al. [2], and the endmember library used by Rogers et al. to analyze Mars TES data [47]. There is no ground truth for this dataset, as the true abundance of minerals on the Martian surface is unknown, so other metrics such as reconstruction error of the spectra are considered.



Fig. 5. Datasets: This figure shows representative data and instrumentation for the three datasets considered in this paper. Data includes laboratory, aircraft, and satellite measurements, and ground truth ranges from detailed abundance analysis under a microscope (Feely [19]) to pure pixel labels of land type for spectral classification (Gulfport [20]) to no ground truth for the Martian data (TES [13])

Baselines. We compare against several state-of-the-art baselines in the literature. The basic linear unmixing algorithm is Fully Constrained Least Squares (FCLS) [24] which solves least squares with sum-to-one and non-negativity constraints on the abundances. We also implement two state-of-the-art statistical methods for modelling endmember variability as distributions: the Normal Com-

11

positional Model (NCM) [50] and the Beta Compositional Model (BCM) [16]. NCM and BCM use a Gaussian and Beta distribution respectively, perform expectation-maximization for unmixing, and require a small amount of training data to determine model parameters.

We also compare against two state-of-the-art deep learning networks by Zhang et al. [63]. The first network utilizes a 1D CNN (CNN-1D) architecture, while the second network utilizes a 3D CNN (but with 1D convolutional kernels) (CNN-3D). CNN-3D is only applicable to datasets with spatial information, and not testable on the Feely and Gulfport synthetic data. We further created a modified CNN architecture (CNN-1D Modified) to maximize the performance on our datasets by changing the loss function to MSE, removing max-pooling layers, and adding an additional fully connected layer before the output. In the supplemental material, we provide information about the parameters, network architectures, and training procedures we used for these baselines as well as details for our own methods. We also have all of our code available here: https://github.com/johnjaniczek/InfraRender.

Endmember Fitting Results. To bootstrap both the analysis-by-synthesis and inverse rendering algorithms, good initial conditions for the dispersion parameters need to be input to the model. Determining dispersion parameters typically required detailed molecular structure analysis or exhaustive parameter searching methods [49, 55, 32]. One main advantage of our method is that we utilize gradient descent to efficiently find parameter sets for different materials. In Appendix B of the supplemental material, we share some of these parameter sets and our insights using the dispersion model for the scientific community.

In Figure 6, we show qualitative results of our endmember fitting by minimizing the loss in Equation 6 using gradient descent. The reconstructed spectra achieves a low MSE with the measured spectra with an average MSE of 0.016 for the TES library, 0.0019 for the Feely library, and an MSE of 2.6e-5 on the Gulfport cluster centroids. Note that there is noise in the measurements, and so MSE is not an absolute metric of the fit to the true unknown spectra.



Fig. 6. Endmember Fitting: (Left) Measured and modelled spectra for a quartz sample in the IR. (Right) Cluster centroids found for pixels labelled as grass in the Gulfport dataset, and the model fit to these centroids. Note the high fidelity of fit via the dispersion model for both these cases

truth for	Feely,	Gulfport,	and	Gulfport	$\operatorname{synthetic}$	datasets.	The	bold	entries	indicate
top performance										
Dataset	FCI	S NCM	BCM	CNN-1	D CNN-3	D CNN-1	DA	nalvs	is- In	verse

Table 1. Results: Table - Mean squared error of the abundance predictions vs. ground

Databet	I O DO	110111	DOM		CIT OD		111101 y 515	Inverse
	[24]	[50]	[16]	[63]	[63]	Modified	by-	Rendering
							synthesis	
Feely	0.121	0.119	0.131	0.469	N/A	0.205	0.052	0.188
[19]								
Gulfport	0.75	0.799	0.800	1.000	0.497	0.297	0.45	0.272
[20]								
Gulfport	0.911	0.471	0.136	0.824	N/A	0.148	0.147	0.059
Synthetic								

Spectral Unmixing Results. In Table 1, we show results on the Feely, Gulfport, and the Gulfport synthetic mixture datasets. For Feely, the analysisby-synthesis method achieved a MSE of 0.052, with the next closest method (NCM) achieving 0.119. Due to the Feely dataset only containing 90 test samples, the machine learning methods were trained on synthetic data which explains their lower performance as data mismatch. Thus, the low error of analysis-bysynthesis shows the utility of the dispersion model for modelling endmember variability, particularly in cases with low training data.

For the Gulfport dataset, the task was to predict the material present since the labeled data is for single coarse materials (e.g. road, grass, etc) at 100% abundance per pixel. Here, the deep learning methods of CNNs and Inverse Rendering have the highest performance. This is expected as there exists a large amount of training data to learn from. Note that Inverse Rendering performs the best at 0.272 MSE, demonstrating that the addition of a generative dispersion model to the output of the CNN improves performance over purely learned approaches. Also note that our analysis-by-synthesis method still has relatively high performance (0.45 MSE) without using any training data at all.

For the Gulfport synthetic mixture dataset, Inverse Rendering achieves the lowest MSE of 0.059, leveraging both physics-based modeling for spectral mixing as well as learns from available training data. The BCM and the analysis-by-synthesis methods both outperform the CNN methods, even though they do not have access to the training data. In fact, BCM even slightly outperforms the analysis-by-synthesis method, which could be because the sources of variation in this data are well-described by statistical distributions.

Speed of Methods. The additional capacity of adding statistical and physical models usually has a cost of speed in implementation. Averaged over 90 mixtures, the convergence for a single operation was FCLS - 10ms, BCM - 1.23s, NCM - 18ms, CNN - 33ms, Inverse Rendering - 39ms, and analysis-by-synthesis - 10.2s. Future work could potentially increase the speed of analysis-by-synthesis with parallel processing.

Noise analysis. Prior to spectral unmixing, emissivity is separated from radiance by dividing out the black-body radiation curve at the estimated tem-

13

perature [43, 14]. In general, a Gaussian noise profile in the radiance space with variance $\sigma_{\text{radiance}}^2$ results in wavenumber dependent noise source in the emissivity space with the profile $\sigma^2(\omega) = \sigma_{\text{radiance}}^2 \cdot 1/B(\omega, T)$ where B is the black-body function given by Planck's law. In our noise experiments we use a black-body radiation curve for a 330K target, which is the approximate temperature the Feely dataset samples were held at. In Figure 7 left, we see that the emissivity noise is higher where the radiance signal is lower.

We simulated varying the noise power to determine the methods' robustness tested on 30 samples from the Feely dataset. In Figure 7, you can see that analysis-by-synthesis still has the best performance in the presence of noise, and is relatively flat as noise increases compared to other methods. We note that statistical methods, while having higher average error, seem to be robust to increased noise as they can handle random perturbations of each spectral band statistically. CNN and Inverse Rendering methods perform the worst for high noise, as these methods were trained on data without noise.



Fig. 7. The left plot shows the radiance profile of a spectra perturbed by Gaussian noise and the resulting emissivity profile after separating out the blackbody radiance. The right figure shows the robustness of the algorithms to increasing amounts of noise

TES Data. The Mars TES data was unmixed using our analysis-by-synthesis method to demonstrate it's utility on tasks where zero training data is available. The method produces mineral maps which correctly finds abundances of the mineral hematite at Meridiani Planum in Figure 8. This is an important Martian mineral which provides evidence for liquid water having existed at some point on Mars, and has been verified by NASA's Opportunity Rover [29]. Note how FCLS predicts many sites for hematite, while our method narrows down potential sites on the Martian surface, which is useful for planetary scientists. By allowing for spectral variation through our physics-based approach, our method has lower RMS reconstruction error than previous analysis of TES data. FCLS, which was previously used on TES because of the zero training-data problem, has an average RMS reconstruction error of 0.0043 while analysis-by-synthesis has an average of 0.0038. This is an exciting result as our methods could provide a new suite of hyperspectral analysis tools for scientists studying the Martian surface.



Fig. 8. Martian Surface Map: The images show the mineral map for hematite of the Martian surface produced by FCLS (left) and analysis-by-synthesis (right) using TES data. Both algorithms find the known deposit of hematite on Meridiani Planum, but analysis-by-synthesis predicts a sparser map which matches expected distributions.

5 Discussion

This paper incorporated generative physics models into spectral unmixing algorithms via differentiable programming. We adopt a physics-based dispersion model to simulate spectral variability, and show how this model can realistically fit several real measured spectra via gradient descent. We further show how to jointly optimize for the dispersion parameters and material abundances with an analysis-by-synthesis optimization. A second algorithm is introduced for tasks where additional data is available by training a CNN to "Inversely Render" a hyperspectral image with the differentiable dispersion model in the loop.

We validate these contributions extensively on three datasets ranging from mid to far IR and VNIR, and compared against state-of-the-art optimization, statistical and deep learning benchmarks. From these experiments we observe that analysis-by-synthesis has the best performance when training data is not available, and that Inverse Rendering has the best performance when training data is available. We also see that analysis-by-synthesis is noise resilient, and reconstructs Mars spectra with lower error than previous techniques.

There are still limitations for the methods proposed. First, analysis-by-synthesis has a large computational cost compared to other methods, although this could be mitigated through parallelization. Secondly, the spectral unmixing community is limited by the lack of quality training datasets, and it is not easy for experts to label remote sensing datasets from prior knowledge alone. Future work could investigate generating realistic synthetic data suitable for training machine learning based algorithms for better performance. We hope using generative physics-based models inspires others to produce realistic synthetic data as well as differentiable programming methods which require low training data.

Acknowledgements. This work was supported by NSF grant IIS-1909192 as well as GPU resources from ASU Research Computing. We thank Dr. Alina Zare, Christopher Haberle, and Dr. Deanna Rogers for their helpful discussions, and Kim Murray (formerly Kim Feely) for providing the laboratory measurements and analysis contributing to this paper.

References

- Azinovic, D., Li, T.M., Kaplanyan, A., Niessner, M.: Inverse path tracing for joint material and lighting estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2447–2456 (2019)
- Bandfield, J.L.: Global mineral distributions on mars. Journal of Geophysical Research: Planets 107(E6), 9–1 (2002)
- Bateson, C.A., Asner, G.P., Wessman, C.A.: Endmember bundles: A new approach to incorporating endmember variability into spectral mixture analysis. IEEE Transactions on Geoscience and Remote Sensing 38(2), 1083–1094 (2000)
- Baydin, A.G., Pearlmutter, B.A., Radul, A.A., Siskind, J.M.: Automatic differentiation in machine learning: a survey. The Journal of Machine Learning Research 18(1), 5595–5637 (2017)
- Bhojanapalli, S., Kyrillidis, A., Sanghavi, S.: Dropping convexity for faster semidefinite optimization. In: Conference on Learning Theory. pp. 530–582 (2016)
- Borsoi, R.A., Imbiriba, T., Bermudez, J.C.M., Richard, C., Chanussot, J., Drumetz, L., Tourneret, J.Y., Zare, A., Jutten, C.: Spectral variability in hyperspectral data unmixing: A comprehensive review. arXiv preprint arXiv:2001.07307 (2020)
- Burns, R.G.: Crystal field spectra and evidence of cation ordering in olivine minerals. American Mineralogist: Journal of Earth and Planetary Materials 55(9-10), 1608–1632 (1970)
- Chang, C.I.: Hyperspectral imaging: techniques for spectral detection and classification, vol. 1. Springer Science & Business Media (2003)
- Chen, F., Zhang, Y.: Sparse hyperspectral unmixing based on constrained lp-l 2 optimization. IEEE Geoscience and Remote Sensing Letters 10(5), 1142–1146 (2013)
- Chen, Y., Jiang, H., Li, C., Jia, X., Ghamisi, P.: Deep feature extraction and classification of hyperspectral images based on convolutional neural networks. IEEE Transactions on Geoscience and Remote Sensing 54(10), 6232–6251 (2016)
- Cheng, G., Li, Z., Han, J., Yao, X., Guo, L.: Exploring hierarchical convolutional features for hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing 56(11), 6712–6722 (2018)
- Chouzenoux, E., Legendre, M., Moussaoui, S., Idier, J.: Fast constrained least squares spectral unmixing using primal-dual interior-point optimization. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 7(1), 59–69 (2014)
- Christensen, P., Bandfield, J., Hamilton, V.E., Ruff, S.W., Kieffer, H.H., Titus, T.N., Malin, M.C., Morris, R.V., Lane, M.D., Clark, R.L., Jakosky, B.M., Mellon, M.T., Pearl, J.C., Conrath, B.J., Smith, M.D., Clancy, R.T., Kuzmin, R.O., Roush, T., Mehall, G.L., Gorelick, N., Bender, K., Murray, K., Dason, S., Greene, E., Silverman, S., Greenfield, M.: Mars global surveyor thermal emission spectrometer experiment: investigation description and surface science results. Journal of Geophysical Research: Planets **106**(E10), 23823–23871 (2001)
- Christensen, P.R., Hamilton, V.E., Mehall, G., Pelham, D., ODonnell, W., Anwar, S., Bowles, H., Chase, S., Fahlgren, J., Farkas, Z., Fisher, T., James, I., Kubik, I., Lazbin, M., Miner, M., Rassas, L., Schulze, K., Shamordola, T., Tourville, G., West, R.and Woodward, D., Lauretta: The osiris-rex thermal emission spectrometer (otes) instrument. Space Science Reviews 214(5), 87 (2018)

- 16 J. Janiczek et al.
- 15. Combe, J.P., Le Mouélic, S., Sotin, C., Gendrin, A., Mustard, J., Le Deit, L., Launeau, P., Bibring, J.P., Gondet, B., Langevin, Y., Pinet, P.: Analysis of omega/mars express data hyperspectral data using a multiple-endmember linear spectral unmixing model (melsum): Methodology and first results. Planetary and Space Science 56(7), 951–975 (2008)
- Du, X., Zare, A., Gader, P., Dranishnikov, D.: Spatial and spectral unmixing using the beta compositional model. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 7(6), 1994–2003 (2014)
- Duchi, J.C., Ruan, F.: Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval. Information and Inference: A Journal of the IMA 8(3), 471–529 (2019)
- Engel, J., Hantrakul, L.H., Gu, C., Roberts, A.: Ddsp: Differentiable digital signal processing. In: International Conference on Learning Representations (2020)
- Feely, K.C., Christensen, P.R.: Quantitative compositional analysis using thermal emission spectroscopy: Application to igneous and metamorphic rocks. Journal of Geophysical Research: Planets 104(E10), 24195–24210 (1999)
- Gader, P., Zare, A., Close, R., Aitken, J., Tuell, G.: Muufl gulfport hyperspectral and lidar airborne data set. Univ. Florida, Gainesville, FL, USA, Tech. Rep. REP-2013-570 (2013)
- Gkioulekas, I., Levin, A., Zickler, T.: An evaluation of computational imaging techniques for heterogeneous inverse scattering. In: European Conference on Computer Vision. pp. 685–701. Springer (2016)
- Gkioulekas, I., Zhao, S., Bala, K., Zickler, T., Levin, A.: Inverse volume rendering with material dictionaries. ACM Transactions on Graphics (TOG) **32**(6), 162 (2013)
- Goudge, T.A., Mustard, J.F., Head, J.W., Salvatore, M.R., Wiseman, S.M.: Integrating crism and tes hyperspectral data to characterize a halloysite-bearing deposit in kashira crater, mars. Icarus 250, 165–187 (2015)
- Heinz, D.C., et al.: Fully constrained least squares linear spectral mixture analysis method for material quantification in hyperspectral imagery. IEEE Transactions on Geoscience and Remote Sensing 39(3), 529–545 (2001)
- Heylen, R., Parente, M., Gader, P.: A review of nonlinear hyperspectral unmixing methods. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 7(6), 1844–1868 (2014)
- Hu, W., Huang, Y., Wei, L., Zhang, F., Li, H.: Deep convolutional neural networks for hyperspectral image classification. Journal of Sensors **2015** (2015)
- Jain, P., Kar, P., et al.: Non-convex optimization for machine learning. Foundations and Trends[®] in Machine Learning **10**(3-4), 142–336 (2017)
- Keshava, N., Mustard, J.F.: Spectral unmixing. IEEE Signal Processing Magazine 19(1), 44–57 (2002)
- Klingelhöfer, G., Morris, R.V., Bernhardt, B., Schröder, C., Rodionov, D.S., De Souza, P., Yen, A., Gellert, R., Evlanov, E., Zubkov, B., et al.: Jarosite and hematite at meridiani planum from opportunity's mössbauer spectrometer. Science 306(5702), 1740–1745 (2004)
- 30. Larkin, P.: Infrared and Raman spectroscopy: principles and spectral interpretation. Elsevier (2017)
- Lee, J.D., Simchowitz, M., Jordan, M.I., Recht, B.: Gradient descent converges to minimizers. arXiv preprint arXiv:1602.04915 (2016)
- 32. Lee, S., Tien, C.: Optical constants of soot in hydrocarbon flames. In: Symposium (international) on combustion. vol. 18, pp. 1159–1166. Elsevier (1981)

- 33. Li, S., Song, W., Fang, L., Chen, Y., Ghamisi, P., Benediktsson, J.A.: Deep learning for hyperspectral image classification: An overview. IEEE Transactions on Geoscience and Remote Sensing 57(9), 6690–6709 (2019)
- Li, T.M., Aittala, M., Durand, F., Lehtinen, J.: Differentiable monte carlo ray tracing through edge sampling. ACM Trans. Graph. (Proc. SIGGRAPH Asia) 37(6), 222:1–222:11 (2018)
- Li, Y., Zhang, H., Shen, Q.: Spectral–spatial classification of hyperspectral imagery with 3d convolutional neural network. Remote Sensing 9(1), 67 (2017)
- Liu, P., Zhang, H., Eom, K.B.: Active deep learning for classification of hyperspectral images. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 10(2), 712–724 (2016)
- Loubet, G., Holzschuch, N., Jakob, W.: Reparameterizing discontinuous integrands for differentiable rendering. Transactions on Graphics (Proceedings of SIGGRAPH Asia) 38(6) (Dec 2019). https://doi.org/10.1145/3355089.3356510
- Lu, G., Fei, B.: Medical hyperspectral imaging: a review. Journal of Biomedical Optics 19(1), 010901 (2014)
- Moersch, J., Christensen, P.R.: Thermal emission from particulate surfaces: A comparison of scattering models with measured spectra. Journal of Geophysical Research: Planets 100(E4), 7465–7477 (1995)
- Nimier-David, M., Vicini, D., Zeltner, T., Jakob, W.: Mitsuba 2: A retargetable forward and inverse renderer. Transactions on Graphics (Proceedings of SIGGRAPH Asia) 38(6) (Dec 2019). https://doi.org/10.1145/3355089.3356498
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, highperformance deep learning library. In: Advances in Neural Information Processing Systems. pp. 8024–8035 (2019)
- Ramsey, M.S., Christensen, P.R.: Mineral abundance determination: Quantitative deconvolution of thermal emission spectra. Journal of Geophysical Research: Solid Earth 103(B1), 577–596 (1998)
- 44. Ramsey M., C.P.: Mineral abundance determination: Quantitative deconvolution of thermal emission spectra: Application to analysis of martian atmospheric particulates. Journal of Geophysical Research: Solid Earth (2000)
- Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Pytorch3d. https://github.com/facebookresearch/pytorch3d (2020)
- 46. Roberts, D.A., Gardner, M., Church, R., Ustin, S., Scheer, G., Green, R.: Mapping chaparral in the santa monica mountains using multiple endmember spectral mixture models. Remote Sensing of Environment 65(3), 267–279 (1998)
- Rogers, A., Aharonson, O.: Mineralogical composition of sands in meridiani planum determined from mars exploration rover data and comparison to orbital measurements. Journal of Geophysical Research: Planets 113(E6) (2008)
- Salisbury, J.W., D'Aria, D.M., Sabins Jr, F.F.: Thermal infrared remote sensing of crude oil slicks. Remote Sensing of Environment 45(2), 225–231 (1993)
- Spitzer, W., Kleinman, D.: Infrared lattice bands of quartz. Physical Review 121(5), 1324 (1961)
- Stein, D.: Application of the normal compositional model to the analysis of hyperspectral imagery. In: IEEE Workshop on Advances in Techniques for Analysis of Remotely Sensed Data, 2003. pp. 44–51. IEEE (2003)

- 18 J. Janiczek et al.
- Sunshine, J.M., Pieters, C.M.: Determining the composition of olivine from reflectance spectroscopy. Journal of Geophysical Research: Planets 103(E6), 13675– 13688 (1998)
- 52. Tsai, C.Y., Sankaranarayanan, A.C., Gkioulekas, I.: Beyond volumetric albedo A surface optimization framework for non-line-of-sight imaging. In: IEEE Intl. Conf. Computer Vision and Pattern Recognition (CVPR) (2019)
- 53. Wang, F., Decker, J., Wu, X., Essertel, G., Rompf, T.: Backpropagation with callbacks: Foundations for efficient and expressive differentiable programming. In: Advances in Neural Information Processing Systems. pp. 10180–10191 (2018)
- Wang, F., Zheng, D., Decker, J., Wu, X., Essertel, G.M., Rompf, T.: Demystifying differentiable programming: Shift/reset the penultimate backpropagator. Proceedings of the ACM on Programming Languages 3(ICFP), 1–31 (2019)
- Wenrich, M.L., Christensen, P.R.: Optical constants of minerals derived from emission spectroscopy: Application to quartz. Journal of Geophysical Research: Solid Earth 101(B7), 15921–15931 (1996)
- Yang, X., Ye, Y., Li, X., Lau, R.Y., Zhang, X., Huang, X.: Hyperspectral image classification with deep learning models. IEEE Transactions on Geoscience and Remote Sensing 56(9), 5408–5423 (2018)
- Yu, Y., Smith, W.A.: Inverse rendernet: Learning single image inverse rendering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3155–3164 (2019)
- Zare, A., Gader, P.: Sparsity promoting iterated constrained endmember detection in hyperspectral imagery. IEEE Geoscience and Remote Sensing Letters 4(3), 446– 450 (2007)
- Zare, A., Gader, P.: Hyperspectral band selection and endmember detection using sparsity promoting priors. IEEE Geoscience and Remote Sensing Letters 5(2), 256–260 (2008)
- Zare, A., Ho, K.: Endmember variability in hyperspectral analysis: Addressing spectral variability during spectral unmixing. IEEE Signal Processing Magazine 31(1), 95–104 (2013)
- Zhang, C., Wu, L., Zheng, C., Gkioulekas, I., Ramamoorthi, R., Zhao, S.: A differential theory of radiative transfer. ACM Transactions on Graphics (TOG) 38(6), 1–16 (2019)
- Zhang, L., Zhang, L., Du, B.: Deep learning for remote sensing data: A technical tutorial on the state of the art. IEEE Geoscience and Remote Sensing Magazine 4(2), 22–40 (2016)
- Zhang, S., Li, J., Li, H.C., Deng, C., Plaza, A.: Spectral-spatial weighted sparse regression for hyperspectral image unmixing. IEEE Transactions on Geoscience and Remote Sensing 56(6), 3265–3276 (2018)
- 64. Zhao, W., Du, S.: Spectral–spatial feature extraction for hyperspectral image classification: A dimension reduction and deep learning approach. IEEE Transactions on Geoscience and Remote Sensing 54(8), 4544–4554 (2016)
- Zhou, Y., Rangarajan, A., Gader, P.D.: A gaussian mixture model representation of endmember variability in hyperspectral unmixing. IEEE Transactions on Image Processing 27(5), 2242–2256 (2018)