# Joint Characterization of Multiscale Information in High Dimensional Data

**Daniel Sousa** Department of Geography, San Diego State University, San Diego CA 92182 USA. **Christopher Small** Lamont Doherty Earth Observatory, Columbia University, Palisades, NY 10964 USA.

Corresponding Author: Daniel Sousa.

Copyright © 2021 Daniel Sousa and Christopher Small. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

#### Abstract

High dimensional feature spaces can contain information on multiple scales. At global scales, spanning an entire feature space, covariance structure among dimensions can determine topology and intrinsic dimensionality. In addition, local scale information can be captured by the structure of low-dimensional manifolds embedded within the high-dimensional feature space. Such manifolds may not easily be resolved by the global covariance structure. Analysis tools that preferentially operate at one scale can be ineffective at capturing all the information present in cross-scale complexity. We propose a multiscale joint characterization approach designed to exploit synergies between global and local approaches to dimensionality reduction. We illustrate this approach using Principal Components Analysis (PCA) to characterize global variance structure and t-distributed Stochastic Neighbor Embedding (t-SNE) to characterize local manifold structure, also comparing against a second approach for characterization of local manifold structure, Laplacian Eigenmaps (LE). Using both low dimensional synthetic images and high dimensional imaging spectroscopy data, we show that joint characterization is capable of detecting and isolating signals which are not evident from either algorithm alone. Broadly, t-SNE is effective at rendering a randomly oriented low-dimensional map of local manifolds (clustering), and PCA renders this map interpretable by providing global, physically meaningful structure. LE provides additional useful context by reinforcing and refining the feature space topology found by PCA, simplifying structural interpretation, clarifying endmember identification and highlighting new potential endmembers which are not evident from other methods alone. This approach is illustrated using hyperspectral imagery of agriculture resolving crop-specific, field scale, differences in vegetation reflectance. The fundamental premise of joint characterization could easily be extended to other high dimensional datasets, including image time series and nonimage data. The approach may prove particularly useful for other geospatial data since both robust manifold structure (due to spatial autocorrelation) and physically interpretable global variance structure (due to physical generative processes) are frequently present.

PCA, t-SNE, Laplacian Eigenmaps, Imaging Spectroscopy, Dimensionality Keywords: reduction, Multiscale.

203

DAN.SOUSA@SDSU.EDU

CSMALL@COLUMBIA.EDU

## **1. INTRODUCTION**

The challenge of high dimensional data analysis has been recognized for decades as critical for scientific progress [1]. Physical, chemical, biological, geological, and astronomical instruments routinely collect measurement streams which rapidly reach terabytes in volume. The Large Synoptic Survey Telescope (under construction at time of writing) is expected to collect 20 terabytes of raw data each night using a 3.2 gigapixel, 5-channel imaging system[2]; a single future NASA mission for the Surface Biology and Geology designated observable is planned to collect – and downlink to Earth – roughly 100 terabytes of >200-channel imagery per day [3]. Similar missions are planned by others, including the European Space Agency and the private sector. As such data continue to grow in size and complexity, novel algorithms are needed to extract relevant signals and mitigate the effects of noise.

In high dimensional data, information can exist at different scales. Some information manifests as *global patterns* – the overall structure of the feature spacebased on all the observations. Additional information can be found in *local manifold structure* – often conceptualized as connectivity or clustering relations among similar observations. In such cases, a meaningful signal may be well-represented as a low-dimensional manifold embedded in a much higher dimensional feature space. Given the wide range of scales at which meaningful information may exist within high dimensional feature spaces, it is plausible that the same dataset may possess useful information at both global and local scales. This situation, which we refer to as multiscale structure, may even be quite common.

Different analysis tools are designed to capture different scales of variability. The well-established tool of Principal Component Analysis (PCA) [4], can be considered to reflect the global variance structure of a dataset because it uses the covariability (correlation or covariance) structure of all observations to apply a single, linear transform to the entire dataset. In contrast, more recent manifold learning algorithms have sought to capture or retain local structure (e.g. Euclidean proximity), for instance by using nearest neighbor-type approaches [5]. Both approaches have their merits. Both can be useful analytic tools in some circumstances, and less useful in other circumstances, depending on the information content of the data.

We propose a new way to consider high dimensional feature space structures – through *joint characterization*. Rather than only using one or another type of dimensionality reduction, joint characterization seeks to exploit synergies among multiple scales of variance and manifold structure to capture multiscale information that is not well-represented by any single tool on its own. This information can be manifest as structure in the space of bivariate (global vs local) distributions.

While the fundamental idea can be generalized, for the purposes of this study we illustrate the approach using a progression of low to high dimensional image data. To do so, we begin by choosing two commonly used dimensionality reduction tools: PCA as our global projection; and t-distributed Stochastic Neighbor Embedding (t-SNE) as our local algorithm [6]. We explore a series of increasingly complex examples to illustrate the strengths of this approach, starting with highly simplified synthetic images and progressing to real-world airborne imaging spectrometer data.

Here, joint characterization is shown to consistently harness the strengths of both methods, using the interpretability of low-order PCs to orient the well-separated but disordered clusters that are the hallmark of t-SNE. The utility of joint characterization becomes increasingly apparent as data become more complex and multiscale structure becomes more important. In the imaging spectroscopy dataset, joint characterization is shown to be capable of identifying subtle differences in reflectance spectra of green vegetation at biophysically meaningful wavelengths. The structure captured by joint characterization is not readily apparent in either the PC or t-SNE spaces alone, indicating clear potential utility of the approach for imaging spectroscopy applications that seek to make full use of hyperspectral information content. Further, we extend these results using another manifold learning algorithm, Laplacian Eigenmaps (LE) [7]. Contrasting the results of PC+t-SNE and PC+LE joint characterizations illustrates the point that additional approaches to dimensionality reduction may yield complementary insight by simplifying endmember identification mixing relations, clarifying topological structure, and potentially identifying new endmembers not evident from other approaches alone.

# 2. METHODOLOGICAL SUMMARY

The joint characterization methodology is briefly outlined in the following steps:

- 1. Compute one set of feature space basis vectors using a global variance metric
  - a. Evaluate hyperparameter sensitivity, if necessary
- 2. Compute another set of feature space basis vectors using a local variance metric
  - a. Evaluate hyperparameter sensitivity, if necessary
- 3. Evaluate topology of global and local feature spaces separately
- 4. Characterize joint feature space formed by combination of both global and local basis vectors

# **3. ANALYSIS**

#### 3.1 A Simple Synthetic Image

Simplified examples can help illustrate the approach. We begin with a synthetic image comprised only of shades of gray (FIGURE 1). This 100 x 100 pixel 8-bit image is divided into a 10 x 10 grid of monoshade chips, varying linearly from black (data values [0, 0, 0]) to white (data values [255, 255, 255]). Every pixel within each chip possesses identical data values. Because each pixel's data values are identical across the three channels, each 3-D [r,g,b] vector can be represented by a single scalar value with no loss of information.

The geometric implication of all 3 channels of the image possessing equal values for all pixels is that the data lie on a 1-D line within a redundant 3-D feature space. Standard PCA reflects this geometric arrangement by partitioning 100% of the total variance (within numerical noise) into the first PC. As a result, only one dimension of PCA is needed to exactly replicate the original image data.

The t-SNE output in this simplified example (bottom center) illustrates several aspects of the method. First, the algorithm separates the data into approximately equally spaced clusters. Each cluster of

points in the t-SNE scatterplot corresponds to a homogenous diagonal of the matrix of color chips - i.e., all points in the image with a single [r,g,b] value.

This leads to the second important point illustrated in FIGURE 1 - that the relative positions of the clusters in t-SNE space are arbitrary. This property is a significant departure from PCA, where position along each axis corresponds to relative contribution of a specific eigenvector. Not so for t-SNE.



Figure 1: A synthetic grayscale image illustrates how PCA can orient t-SNE clusters. In the case of a grayscale image (top left), continuous variation of one principal component (top center) captures all the information in the data. 2D t-SNE (top right) correctly assigns each equal-brightness matrix diagonal to a small range of t-SNE values. The t-SNE space (bottom center) has well-separated clusters corresponding to discrete brightness values, but conveys no information about the global spatial structure (brightness gradient) of the matrix. The joint PC + t-SNE space (lower left and lower right) retains the local cluster structure, but orients the global relationships among clusters to align them along the diagonal brightness gradient.

The third and fourth salient properties are less obvious at first glance: clusters of identical values have nonzero dispersion in t-SNE space, and different t-SNE runs with identical hyperparameters give slightly different results. These are also in contrast to PCA, where data values correspond to one and only one set of PC values.

We show results using the scikit-learn default hyperparameter settings here (most importantly, perplexity = 30), but note that the experimentation suggests the general structure of the output is robust to a wide range of hyperparameter choices.

Examining the PC + t-SNE outputs together – joint characterization – retains important information from each. PCA effectively orients the data along its global variance structure – in this case, from lowest to highest brightness values. T-SNE effectively partitions similar (in this case, identical) values into well-separated clusters, but does not structure these clusters in a meaningful way. By combining the global and local information captured by PCA and t-SNE, both sets of information are retained and we are able to illustrate patterns which would be hidden by one or the other tool on its own. In the 1-D case of FIGURE 1 there is no local structure, and the  $1^{st}$  PC captures all the global structure, so joint characterization provides no new information. This is not the case for the datasets with greater dimensionality examined below.

#### 3.2 A More Complex Synthetic Image: Discrete 3D

Further insight can be obtained by examining a slightly more complex synthetic dataset (FIGURE 2). Here, three additional submatrices are introduced which illustrate binary mixing for each of the (red, green); (red, blue); and (green, blue) combinations of primary additive colors.



Figure 2: Synthetic binary color mixtures illustrate local and global information. A synthetic image (A; left) comprised of submatrices with both binary color mixtures and grayscale brightness variability is analyzed with both PCA (A; center) and t-SNE (A; right). Three PCs are now required to describe the RGB data. In PC space (B), the binary color mixtures form a hexagonal manifold (PC 2 vs 3) arranged by hue and a perpendicular gray axis to represent brightness. In t-SNE space (B; lower right), each equal-valued diagonal of the grayscale submatrix forms a large cluster in the center while each distinct color of the binary mixing submatrices forms a denser cluster in the surrounding cloud. The joint PC+t-SNE space (C) aligns the randomly oriented t-SNE clusters along directions of global variance (brightness), with the PC 1 projection clearly separating the grayscale and color submatrices, and PC 2 and 3 projections orienting clusters in the color submatrix on the basis of hue.

PCA in this case yields markedly more structure to the data. PC 1 corresponds to overall pixel brightness and PCs 2 and 3 form a discrete, hexagonal approximation of the color wheel, where each of the 6 vertices corresponds to either a (red, green, blue) primary color or a (yellow, cyan, magenta) binary mixture. The gray axis protruding from the center of the color wheel (white (W)) towards high PC 1 values (black (K)) corresponds to the grayscale submatrix we have seen in FIGURE 1, which again can be completely represented by a single dimension.

The t-SNE result for this more complex synthetic image bears notable similarities to the t-SNE result for the grayscale-only image. Again, each set of pixels with identical values forms a cluster. The increased number of unique [R,G,B] values is reflected in the increased number of clusters in t-SNE space. Again, the clusters form an approximately equidistant lattice-like structure; again, the relative position of clusters within this structure is arbitrary; again, pixels with exactly identical [R,G,B] values (gray) form clusters with nonzero variance. The grayscale submatrix occupies the central portion of the t-SNE plot and differs from the binary mixing submatrices because it has a greater number of pixels with identical data values (entire submatrix diagonals, rather than individual matrix elements).

As the complexity and dimensionality increase from FIGURE 1 to FIGURE 2, the value added by joint characterization becomes apparent. The global variance structure captured by PCA serves to orient the local structure captured by t-SNE. However, in this image, the homogeneity of each color chip in this image limits the usefulness of local structural metrics like t-SNE. As we shall see below, local information can be much more important as image complexity increases further.

### 3.3 Greater Complexity: Continuous 3D with Compression Artifacts

While simple datasets comprised of clusters with precisely equal data values can provide valuable insight, most real-world images contain continuously varying data values. FIGURE 3 illustrates PCA, t-SNE, and the joint characterization approach for a third example image with both discrete and continuous variability. Importantly, the image has also undergone lossy compression using the commonly used JPEG algorithm. Lossy compression is typically avoided for datasets with scientific applications because it changes underlying data values. However, it does so in a way designed to be difficult for the human eye-brain system to perceive. For this reason, ability to detect these compression artifacts can be an indicator of the power of an analysis approach to extract subtle, low-variance information which might be difficult for humans to visualize in the raw data.

Again, as image complexity increases further, the relevance of the joint characterization approach becomes clearer. PCA again reflects the global variance structure of the data – in this case, a 3-D color cube. Locally, t-SNE again partitions the sets of identical or nearly identical values of the color chips in the bottom half of the image to tight clusters in the periphery of the plot. The continuum in the top half of the data image is expressed as a more tortuous central point cloud. T-SNE appears to be sufficiently sensitive to coherent small amplitude variability to capture JPEG compression artifacts which are not obvious to the naked eye. Together, the joint space retains benefits of both PC and t-SNE: global ordering and local cluster separability, respectively.



Figure 3: Joint characterization reveals continuous and discrete structure in a more complex 3D example. Both continuous and discrete variability is present in these discrete and continuous [R,G,B] color tables (A; left). PCA (A; center) retains the global color gradients while t-SNE (A; right) separates colors discontinuously on the basis of local structure. In PC space (B) distinct [R,G,B] values fill out a color cube with a gray axis spanning the diagonal between black and white. The t-SNE space (B) finds subtle nonlinear manifolds that appear to correspond to compression artifacts in the continuous hue scale, surrounded by tight clusters that correspond to specific [R,G,B] triplets on the discrete array. The joint PC+t-SNE characterization (C) introduces complexity by disrupting the interpretable geometric structure of the PC color space with the apparently random structure of the t-SNE space.

#### 3.4 Greater Dimensionality and Complexity: Airborne Imaging Spectroscopy

One promising application for joint characterization is in the examination of Earth-observing images from sensors on aerial and satellite platforms. Imaging spectroscopy is a particularly interesting use case because it allows us to explore images with much greater dimensionality than the images recorded by typical cameras. Whereas a digital camera images the visible spectrum using 3 color channels (red, green and blue primary colors), imaging spectrometers image a much wider range of visible to Shortwave Infrared (SWIR) wavelengths with hundreds of narrowband spectral channels. The result is a hyperspectral image cube where each geographic pixel provides a very detailed optical reflectance spectrum capable of resolving subtle molecular absorptions characteristic of

different chemical compositions of the material being imaged. The broad scale global structure of each spectrum is characteristic of the material type (e.g. soil, vegetation, rock, water, snow) while the finer scale absorption features allow differentiation within each material type (e.g. soil composition, vegetation type and condition, rock mineral content). The resulting spectral feature space generally has >100 dimensions corresponding to each of the narrow spectral channels in the image cube. PCA is commonly used to characterize the lower-D global structure of the hyperspectral feature space with respect to different types of materials, but fine scale absorption features exist which contribute much less variance than the overall spectral shape and so tend to form amorphous continua embedded within the low-D subspace of the low order PCs. The situation is analogous to the continuous color scale in FIGURE 3 in which the primary colors control the overall cubic form of the 3D PC space, while the continuum of additive colors form continua on the manifold between the primary color apexes. The main difference is the vastly greater information content of imaging spectroscopy, resulting from narrower spectral channels (e.g. ~40 channels in the visible spectrum alone) and much broader spectral range (e.g. 400 nm (visible blue) to 700 nm (visible red) to 900 nm (Near Infrared) to 2500 nm (Shortwave Infrared)).

In the context of this expository analysis, industrial-scale agriculture is a particularly well-posed problem because of the easily comprehensible geographic patterns that occur at spatial scales coarser than pixel resolution. Whereas individual fields are generally planted with a single crop type having a similar overall spectral shape, spatial variations in soil moisture and nutrients within individual fields can result in subtle differences in the reflectance of the aggregate spectra corresponding to the area imaged by each pixel. The ability to image these subtle spectral features is part of the basis for the great potential of imaging spectroscopy to precision agriculture and crop monitoring. To date, over 23 TB of such imagery has been collected by NASA's Airborne Visible / Infrared Imaging Spectrometer (AVIRIS) and are publicly available free-of-charge at: https://aviris.jpl.nasa.gov/dataportal/.

We use an example AVIRIS hyperspectral cube of agriculture in the California's Central Valley, home to extensive industrial agriculture with total cash receipts exceeding \$50 billion/year (CDFA, 2019). The area shown in FIGURE 4 is a 500 x 500 pixel ( $10 \times 10 \text{ km}$ ) subset of flightline f140606t02p00r04. Each pixel represents a spatially integrated optical reflectance spectrum at roughly 16.5 m spatial resolution and 10 nm spectral resolution over the range 365 to 2496 nm. Bands 1-14, 49-64, 75-84, 90-120, 150-180, and 210-224 were excluded due to atmospheric absorptions and/or sensor noise. 107 bands were retained for analysis.

The AVIRIS dataset clearly contains real-world complexity. However, the same properties which characterize the synthetic datasets are also evident even here. The 3 low-order PCs effectively capture global variance structure, capturing >97% of variance in this dataset and clearly grading between distinct materials like soils, senescent vegetation, green vegetation, and standing water. This single subset has similar global variance features to those found in previous, more extensive observations of agricultural landscapes across the region in [8].

Even in the face of real-world complexity, t-SNE can be remarkably effective at capturing local manifold structures in imaging spectroscopy data. Clusters in t-SNE space unambiguously correspond to spatially coherent patterns of intra- and inter-field variability. However, as in the synthetic images, the t-SNE space is clustered but unstructured. Adjacent clusters generally have no clear relation to each other, and clusters in disparate corners of the t-SNE space can possess great similarity. This



Figure 4: Joint characterization reveals novel structure in imaging spectroscopy data. The data here (A; left) show an agricultural landscape in California's Central Valley as imaged by the 224-channel AVIRIS hyperspectral sensor.PCA (A; center) represents the global variance continuum, and t-SNE (A; right) captures local manifold structure embedded within the global feature space. The structure of the data in PC space (B) distinguishes differences in vegetation, soil and water reflectance while the structure in t-SNE space (B; lower right) clusters the data but does not orient according to any global structure. Joint characterization (C) contains both the global variance structure represented by PCA and the local structure found by t-SNE, effectively segregating the continuous PC space into many distinct clusters and gradients.

property is inherent to t-SNE and limits its utility for imaging spectroscopy data – unless additional information is introduced.

Joint characterization provides this additional information. In the case of these data, the joint space clearly both retains useful clustering from t-SNE, and also orients those clusters using the global variance structure captured by PCA. As a result, subtle but consistent interpretable differences can be isolated within the imaging spectroscopy data which could not readily be identified from global or local characterization alone.



Figure 5: Multiscale spectral feature space of an AVIRIS hyperspectral image cube of agricultural plots in the Great Central Valley of California. The spectral feature space of PCs 1 and 2 (lower left) shows global structure related to characteristic reflectance spectra of distinct land cover types. The complementary feature space of PC1 and t-SNE1 (upper left) shows finer scale clustering related to local structure corresponding to subtle differences among similar land cover spectra. Distinct clusters (colored) of spectra generally correspond to spatially contiguous patches within individual plots (upper right) while the mean spectra of each cluster (lower right) all correspond to green vegetation, but with varying differences in amplitude and curvature related to varying depth of H2O absorption bands resulting from varying leaf water content of different crops. Pairwise differences of cluster mean reflectances (inset) illustrate wide range of varying spectral slopes indicative of differing H2O absorption strength, despite similar shapes of vegetation spectra that result in cluster overlap in PC1/PC2 feature space.

FIGURE 5 shows an example using 10 regions of interest (ROIs) from the PC+t-SNE space. These ROIs are chosen to illustrate the observed local variability in reflectance of green vegetation spectra.

All green vegetation spectra in the image cluster in one corner of the PC space (lower left scatterplot), and form distinct clusters throughout the t-SNE space (inset lower left), but form clearly separated clusters across the range of t-SNE values in the joint characterization space (upper left scatterplot). Mean reflectance spectra of each ROI are plotted in the lower right. All have a similar overall shape (associated with their proximity in PC space) with characteristic absorptions typical of green vegetation, but differ interpretably at wavelengths corresponding to known biophysical absorption features like leaf water, cellulose/lignin, and photosynthetic pigments. Differences among ROIs are easier to visualize when accentuated by differencing pairs of mean spectra (inset right). The differences among the mean spectra are also indicated by the varying amplitude sequences in each wavelength interval (inset labels). The variations in SWIR (1500 nm to 2500 nm) amplitude are primarily a result of differences in depth of the liquid water absorptions in the H<sub>2</sub>O gaps driven by differences in leaf water. It is notable that the spectral subsets of every pair of ROIs has Transformed Divergence of 2.0, indicating complete spectral separability in the original 224-dimensional feature space of the hyperspectral cube. When the pixels in each contiguous ROI from the PC+t-SNE space are plotted in their corresponding locations in geographic space (upper right image) it is immediately apparent that each ROI cluster from PC+t-SNE space corresponds to a distinct, spatially contiguous region within individual or adjacent fields in geographic space. This spatial contiguity is consistent with the local structure revealed by the PC+t-SNE space being related to crop-specific features in the reflectance spectra. These subtle intra-field variations in reflectance can also be seen on the sides of the hyperspectral cube in FIGURE 5.

As noted throughout, t-SNE is but one of a plethora of possible algorithms to capture local variance. In order to demonstrate the potential for other algorithms to contribute additional information under the joint characterization framework, we also illustrate the results of a second manifold-learning algorithm, Laplacian Eigenmaps (LE) [7]. Briefly, LE (a special case of the more general term 'Spectral Embedding') forms a data affinity matrix and computes the spectral decomposition of this matrix. Here, we use the 'nearest neighbor' affinity metric and default number of n\_neighbors (i.e., max(N/10,1)). We also note that LE is more computationally intensive than t-SNE, requiring spatial subsampling (see Appendix for additional computational details).

The resulting LE feature space for the AVIRIS data is shown in FIGURE 6. It is evident that LE and t-SNE provide complementary information about both global and local manifold structure. For these data, the LE-generated feature space topology clearly captures apexes linked by both rectilinear and curvilinear edges. LE clearly identifies the endmembers identified from FIGURE 5, providing simpler, more readily interpretable topology and endmember identification. Specifically, senescent vegetation (yellow) and wet (cyan) endmembers form clear binary mixing relations with partial vegetation (magenta), which mixes with bare dry soil (red) to converge on closed canopy vegetation (green). Additionally, the joint PC+LE feature space combines the connectivity structure information provided by LE with the physical interpretability of the PC basis vectors, clearly capturing coherent, interpretable landscape patterns. While FIGURES 1-5 illustrate that t-SNE can add useful clustering information in the context of joint characterization, FIGURE 6 clearly shows that other approaches to capturing local variance are also worth exploring. Ultimately, the optimal approaches for capturing global and local variance are likely to be data dependent, and the wise and careful analyst may find it useful to explore multiple avenues of inquiry before determining an optimal characterization for any particular case.



Figure 6: Using an additional local variance metric can provide additional insight into feature space topology. In this case, implementation of the Laplacian Eigenmaps (LE) algorithm generates a feature space with clearer apexes than PCs and less clustering than t-SNE.

# 4. DISCUSSION

From FIGURES 4, 5, and 6, it is apparent that joint characterization can be useful for analysis of imaging spectroscopy datasets, at least of agricultural landscapes. While some of the features of imaging spectroscopy data may render it particularly amenable to this type of analysis, the approach is highly general and might easily be extended to other high dimensional datasets. Potential implications and avenues for future work are discussed below.

### 4.1 Importance of Both Local Manifolds & Global Variance

The foundation of this study is the recognition that the information content of high dimensional datasets can map onto multiple scales of variance. In essence, some data are too rich to be well-described by variation measured at a single scale. In the case of imaging spectroscopy, global and local scales are frequently recognized as the "spectral continuum" and "absorption features" (respectively), and so have established physical attributes [9]. The global variance structure of Earth surface reflectance also has a well-understood physical meaning in terms of a mixing space of spectrally distinct endmember materials [8,10-14]. This may not generally be the case for all types of data.

One common approach to determining the feature space basis vectors for imaging spectroscopy data is to use low-order dimensions from PCA. However, this approach implicitly assumes that global variance corresponds to information – that is, high variance features (e.g., differences in the spectral continuum) are prioritized over low variance features (e.g., absorptions). Such an approach may not be optimal in all cases, as one of the great strengths of imaging spectroscopy is its ability to accurately capture narrow, low-amplitude features. In this context, joint characterization can be used to generate an alternate set of spectral feature space basis vectors. Specifically, the spectral continuum might be effectively captured using a global algorithm like PCA, while low amplitude differences in absorptions using an algorithm more tailored to capturing local-scale manifold structure, like t-SNE. It is possible that using both characterization approaches together might yield a joint spectral feature space which could more effectively capture important differences in both spectral continuum and absorptions. This might prove particularly valuable for exploiting low-amplitude variance structures in the context of the spectral mixture residual approach proposed by Sousa et al., 2021 [15,16].

### 4.2 Implications for the Remote Sensing Community

Imaging spectroscopy is an emerging subdiscipline in remote sensing. After decades of instrument development, and early successes like Hyperion and HICO [17,18], a new generation of satellite missions are at or nearing operational status. While multispectral satellite imagery has been collected for over 40 years, image analysis methods must continue to evolve to match the higher dimensional demands of hyperspectral data. Tools like joint characterization will be necessary for the scientific community to fully exploit the information captured by these data. Indeed, a theoretical basis exists for joint characterization to prove particularly useful for geospatial datasets. In such data, spatial autocorrelation can translate into physically meaningful variance at local scales, and the finite range of reflectance of Earth surface properties can translate into consistent, physically interpretable patterns of global variance.

Specifically, the results of this analysis suggest that, in some circumstances, joint characterization may be an effective tool for the imaging spectroscopy community by providing useful basis vectors for spectral feature spaces. For many study areas, the global variance patterns captured by PCA can be easily interpretable in the context of common, well understood classes of Earth materials [8]. Linking these interpretable global features to the local, low-variance manifolds captured by t-SNE, LE or other similar algorithms may considerably enhance the value of imaging spectroscopy by allowing for effective visual discrimination of subtle biophysical differences among spectra with similar global variance profiles (e.g., the green vegetation endmembers shown in FIGURE 5).

An additional promising area of inquiry is in the realm of spatiotemporal analysis. Temporal feature spaces are inherently higher dimensional than spectral feature spaces because they are a result of temporal progressions of physical states less constrained by the finite number of material properties and molecular structures that give rise to feature spaces of reflectance spectra [19]. It is likely that image time series may also be characterized by signals existing at multiple scales, especially given the wide range of ways that Earth surface reflectance can evolve through time (vegetation phenology, fire, flooding, land cover change, etc.). This idea has recently begun to be explored [2021], but significant avenues for future work remain.

#### 4.3 Open Questions for Future Work

Several avenues for future work exist. First, we use PCA, t-SNE, and LE in this analysis as tools for measuring global and local variance, but others exist. Which, if any are most effective? Using what criteria, and under what circumstances? Second, t-SNE and LE (as well as several other manifold learning approaches) can yield variable results. How sensitive is the result to hyperparameter choice? How well do clusters observed by joint characterization align with physical properties of the data? In what, if any, cases are clusters spurious, or have no clear physical meaning? Questions like these, and many others, will be critical to examine in detail as joint characterization is applied to a greater diversity of data.

### **5. CONCLUSIONS**

We introduce the idea of joint multiscale characterization. The underlying principle is that information in high dimensional data can exist at multiple scales of variance. Joint characterization uses global and local variance metrics together to identify data patterns which are not evident from either metric alone. We illustrate this principle using Principal Components Analysis as a global variance metric, and t-distributed Stochastic Neighbor Embedding and Laplacian Eigenmaps as local variance metrics. Examples begin with low dimensional synthetic image data and progress in complexity and dimensionality to airborne imaging spectroscopy data of an agricultural landscape. The utility of joint characterization increases as data complexity increases, culminating in clear separability of green vegetation spectra which differ at wavelengths corresponding to key biophysical traits. Joint characterization may prove particularly useful for geospatial datasets given both robust local structure (due to spatial autocorrelation) and physical meaning of global variance (due to spectral mixing), but the idea could be generalized to other data types as well.

## 6. ACKNOWLEDGEMENT

D. Sousa was funded by the taxpayers of the State of California. C. Small gratefully acknowledges the support of the endowment of the Lamont Doherty Earth Observatory of Columbia University. © 2021. All rights reserved.

### 7. SUPPLEMENTARY INFORMATION

#### 7.1 t-SNE Hyperparameter Sensitivity

FIGURE S1 illustrates the effect of varying perplexity on the trial dataset from FIGURE 2. Here, perplexity is varied across the range of suggested values (from 5 to 50, in increments of 5). The overall effect of changing perplexity in this dataset is on the dispersion of the central grayscale clusters. In general, t-SNE clusters comprised of equal-valued grayscale pixels are most compact and most separated from each other at low perplexity values. As perplexity increases, the dispersion of the clusters increases and they grow towards each other, ultimately filling the majority of the space of the central portion of the plot. Interestingly, this occurs only for the clusters corresponding to grayscale pixels, and not clusters corresponding to binary color mixtures.



Figure S1: The t-SNE hyperparameter of perplexity impacts cluster dispersion. As the t-SNE hyperparameter of perplexity increases across its recommended range of values, the central clusters corresponding to equal-valued grayscale matrix diagonals expand to fill the center of the plot. The outer clusters corresponding to binary color mixtures also expand but show considerably less sensitivity.

Other key hyperparameters (early exaggeration, learning rate, number of iterations) were varied across their suggested ranges as well, but none demonstrated any appreciable sensitivity comparable to that of perplexity. In particular, recent literature[20,21], suggests the potential for particular sensitivity to learning rate and recommends determining this hyperparameter using the following formula:



max(N / early\_exaggeration / 4, 50)

Figure S2: The t-SNE hyperparameter of learning rate has a minor impact on cluster dispersion in this trial dataset. As the t-SNE hyperparameter of learning rate increases between the minimum recommended value of 50 [20,21], the scikit-learn default of 200, and the data-adaptive recommended value of 833, central clusters corresponding to equal-valued grayscale matrix diagonals contract slightly to leave more empty space near the center of the plot. The outer clusters corresponding to binary color mixtures show considerably less sensitivity. For the trial dataset from FIGURE 2, with N = 30,000 pixels and using the default scikit-learn early\_exaggeration value of 12, this formula yields a learning rate value of:

 $\max(40,000 / 12 / 4, 50) = \max(833, 50) = 833$ 

We illustrate the effect of learning rate at this hyperparameter choice, in comparison to the scikitlearn default value of 200 as well as additional values between the lower bound of 50 and upper bound of 833, in FIGURE S2. We also note that recent work [22,23], suggests relative hyperparameter insensitivity for some biogeophysical systems, such as landscape-scale Earth observation.

### 7.2 Computational Specifics

Computation of was performed using the *Principal Component Rotation* tool in ENVI 5.6, which follows the approach of [24]. Details are given at: https://www.l3harrisgeospatial.com/docs/ principalcomponentanalysis.html. Covariance-based transforms were used in all cases. Results were compared to the scikit-learn output (*sklearn.decomposition.PCA*) and no perceptible difference was found.

Computation of t-SNE and LE was performed in Python 3.8.8 using scikit-learn v0.24.2. Computation time was found to be sensitive to perplexity but less sensitive to learning rate. On a 2 GHz Quad-Core Intel Core i5 CPU with 32 GB of 3733 MHz LPDDR4X RAM and an Intel Iris Plus Graphics 1536 MB GPU, computation time for a single t-SNE realization for FIGURE 2 with perplexity = 100 was just under 3 minutes (199 seconds for learning rate = 200; 215 seconds for learning rate = 833). Computation time decreased when the perplexity hyperparameter was lowered to its default of 30 using either the higher or lower learning rate (138 seconds at learning rate = 833; 127 seconds at learning rate = 200).

Computation of LE was found to require substantially more RAM than t-SNE. Specifically, the number of samples (i.e., pixels) was found to be a significant limiting factor in computation and inversion of the LE connectivity matrix. Computation using the full  $500 \times 500$  pixel AVIRIS example dataset was not possible using this hardware, so data were spatially subsampled (decimated) by a factor of 5 in each dimension to yield a  $100 \times 100$  pixel subset. However, once the dataset was reduced within RAM limitations, computation time was faster than t-SNE (50 seconds).

### References

- Donoho DL. High-dimensional data analysis: The Curses and Blessings of Dimensionality. AMS Math Challenges Lecture, AMS Conference on Math Challenges of the 21st Century.2000:32.
- [2] Kahn S, Kurita N, Gilmore K, Nordby M, O'ConnorP, et al. Design and Development of the 3.2 Gigapixel Camera for the Large Synoptic Survey Telescope. Presented at the Ground-based and airborne instrumentation for astronomy III, International Society for Optics and Photonics. 2010;7735:77350J.
- [3] Jenkins J, Tenenbaum P, Brosnan I G, Middleton E, Campell PK, et al. Surface Biology & Geology Pathfinder Data Analysis Pipeline. Presented at the AGU Fall Meeting Abstracts. 2019; 2019:GC54A-03.

- [4] Pearson K. LIII. On Lines and Planes of Closest Fit to Systems of Points in Space. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science. 1901;2:559-572.
- [5] Cayton L. Algorithms for Manifold Learning. Univ. of California at San Diego Tech. Rep, 2005;12:1-17.
- [6] Van der Maaten L, & Hinton G. Visualizing Data Using t-SNE. Journal of Machine Learning Research. 2008;9:2579-2605.
- [7] Belkin M, Niyogi P. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. Neural Computation. 2003;15:1373-1396.
- [8] Sousa D, Small C. Multisensor Analysis of Spectral Dimensionality and Soil Diversity in the Great Central Valley of California. Sensors. 2018;18:583.
- [9] Clark RN, RoushTL. Reflectance Spectroscopy: Quantitative Analysis Techniques for Remote Sensing Applications. Journal of Geophysical Research: Solid Earth.1984;89(B7):6329-6340.
- [10] Kauth RJ, Thomas GS. The Tasselled Cap A Graphic Description of the Spectral-Temporal Development of Agricultural Crops as Seen by LANDSAT. In Symposium on Machine Processing of Remotely Sensed Data. West Lafayette, Indiana, 4B: The Laboratory for Applications of Remote Sensing, Purdue University. 1976;159: 41-51.
- [11] Crist EP, Cicone RC. A Physically-Based Transformation of Thematic Mapper Data—The TM Tasseled Cap. IEEE Transactions on Geoscience and Remote Sensing, GE. 1984; 22:256–263.
- [12] Small C. The Landsat ETM+ Spectral Mixing Space. Remote Sensing of Environment. 2004;93:1-17.
- [13] Small C, Milesi C. Multi-Scale Standardized Spectral Mixture Models. Remote Sensing of Environment.2013);136, 442-454.
- [14] Sousa D, SmallC. Global cross-calibration of Landsat spectral mixture models. Remote Sensing of Environment. 2017;192,139-149.
- [15] https://arxiv.org/ftp/arxiv/papers/2108/2108.09545.pdf
- [16] https://doi.org/10.1002/essoar.10508388.1
- [17] Pearlman JS, BarryPS, Segal CC, Shepanski J, Beiso D, et. al. Hyperion, a Space-Based Imaging Spectrometer. IEEE Transactions on Geoscience and Remote Sensing. 2003;41:1160-1173.
- [18] Corson MR, Korwan DR, Lucke RL, Snyder WA, et. al. The Hyperspectral Imager for the Coastal Ocean (Hico) On the International Space Station. Presented at the IGARSS 2008-2008 IEEE International Geoscience and Remote Sensing Symposium, IEEE. 2008; 4: IV-101.
- [19] Small C, Milesi C. Multi-Scale Standardized Spectral Mixture Models. Remote Sensing of Environment.2013;136, 442-454.
- [20] Schaum N, Karkanias J, NeffNF, MayAP, Quake SR, et al. Single-Cell Transcriptomics of 20 Mouse Organs Creates a Tabula Muris. Nature. 2018;562,367-372.

- [21] Zeisel A, Hochgerner H, LönnerbergP, Johnsson A, Memic, F, et al. Molecular Architecture of the Mouse Nervous System. Cell. 2018;174, 999-1014.e22.
- [22] https://arxiv.org/ftp/arxiv/papers/2112/2112.01416.pdf
- [23] Small C, Sousa D. The Climatic Temporal Feature Space: Continuous and Discrete. Advances In Artificial Intelligence and Machine Learning. 2021;1:165–183.
- [24] Richards JA, Richards J. Remote Sensing Digital Image Analysis. Springer.1999;3.