

Unmixing Spectral Data for Space Objects using Low-Rank Non-Negative Matrix Factorization

V. Paul Pauca* Robert J. Plemmons[†] Maile Giffin[‡] Kris M. Hamada[§]

November 24, 2003

Abstract

This study involves spectral unmixing methodology for sensor data compression, retrieval, analysis, and classification of space objects. The methodology is based upon blind source separation by encoding the data using low-rank non-negative matrix factorization algorithms to preserve natural data non-negativity and thus avoid subtractive basis vector and encoding interactions present in techniques such as principal component analysis. Some existing non-negative matrix factorization techniques are reviewed and some new ones are proposed. Numerical experiments are reported using spectral data for non-imaging identification and classification of orbiting space objects.

Keywords: data mining, space object identification and classification, multispectral sensing, non-negative matrix factorization, blind source separation, unsupervised spectral unmixing.

1 Introduction

Non-negative matrix factorization (NMF) has recently been shown to be a very useful technique in approximating high dimensional data where the data are comprised of non-negative components. In a seminal paper published in *Nature* [13], Lee and Seung proposed the idea of using NMF techniques to find a set of basis functions to represent image data where the basis functions enable the identification and classification of intrinsic “parts” that make up the object

*Department of Computer Science, Wake Forest University, Winston-Salem, NC 27109. His research was supported in part by the Air Force Office of Scientific Research under grant FA49620-03-1-0215, and by the Army Research Office under grant DAAD19-00-1-0540.

[†]Departments of Computer Science and Mathematics, Wake Forest University, Winston-Salem, NC 27109. His research was supported in part by the Air Force Office of Scientific Research under grant F49620-02-1-0107, and by the Army Research Office under grant DAAD19-00-1-0540. email: plemmons@wfu.edu

[‡]Oceanit Laboratories, 590 Lipoa Parkway Kihei, Maui, HI 96753.

[§]Boeing LTS, Inc., 535 Lipoa Parkway, Ste. 200 Kihei, Maui, HI 96753.

being imaged by multiple observations. They showed that NMF facilitates the analysis and classification of data from image or sensor articulation databases made up of images showing a composite object in many articulations, poses, or observation views. They also found NMF to be a useful tool in text data mining (See also Pauca, Shahnaz, Berry and Plemmons [18]). In the past few years, several papers have discussed NMF techniques and successful applications to various databases where the data values are non-negative, e.g., [5, 6, 7, 8, 9, 14, 15, 21].

More generally, matrix factorization techniques in data mining fall under the category of vector space methods. Very often databases of interest lead to a very high dimensional matrix representation. Low-rank factorizations not only enable the user to work with reduced dimensional models, they also often facilitate more efficient statistical classification, clustering and organization of data, and lead to faster searches and queries for patterns or trends, e.g., Berry, Drmač, and Jessup [2].

NMF is a vector space method to obtain a representation of data using non-negativity constraints. These constraints can lead to a parts-based representation because they allow only additive, not subtractive, combinations of the original data. This is in contrast to techniques for finding a reduced dimensional representation based on singular value decomposition-type methods such as principal component analysis (PCA) [11]. One major problem with PCA is that the basis vectors have both positive and negative components, and the data are represented as linear combinations of these vectors with positive and negative coefficients. In many applications, however, the negative components contradict physical realities. For example, pixels in a gray scale image have non-negative intensities, as do the spectral radiance data obtained by remote sensing. It is the second application that leads to the topic we are concerned with in this paper: space object identification and classification from spectral reflectance data. Here, we employ unsupervised spectral unmixing techniques for recovery of constituent spectra.

2 Non-Negative Matrix Factorization

Given an initial database expressed as an $n \times m$ matrix X , where each column is an n -dimensional non-negative vector of the original database (m vectors), the standard NMF problem is to find two new reduced-dimensional matrices W and H , in order to approximate the original matrix X by the product WH in terms of some metric. Each column of W contains a *basis vector* while each column of H contains the *weights* needed to approximate the corresponding column in X using the basis from W . The dimensions of matrices W and H are $n \times r$ and $r \times m$, respectively. Usually, the number of columns in the new (basis) matrix W is chosen so that $r \ll m$, and in particular, $r(n + m) < nm$. Here, the choice of r is generally application dependent, and may also depend upon the characteristics of the particular database within the application.

The usual approach to the NMF problem is to approximate X by computing a pair W and H to minimize the Frobenius norm of the difference $X - WH$. Mathematically, the problem can be stated as follows: Let $X \in R^{n \times m}$ be a data matrix of non-negative entries. Let $W \in R^{n \times r}$ and $H \in R^{r \times m}$ for some positive integer $r < m$. The objective is then to solve the optimization problem

$$\min_{W,H} \|X - WH\|_F^2, \tag{1}$$

subject to $W_{ij} \geq 0$ and $H_{ij} \geq 0$ for each i and j .

Of course the matrices W and H are generally not unique. Conditions resulting in uniqueness in the special case of equality, $X = WH$, have been recently studied by Donoho and Stodden [5], using cone theoretic techniques (See also Chapter 1 in Berman and Plemmons [1]). Although, X is generally not a symmetric matrix, Cooper and Foote [4] have proposed summarizing video using non-negative similarity matrix factorizations, where the data similarity matrix is symmetric (See also a study of the symmetric NMF case by Catral, Han, Neumann and Plemmons [3]).

Algorithms designed to approximate X by solving the minimization problem (1) generally begin by initial estimates of the matrices W and H , followed by alternating iterations to improve these estimates. Next, some existing non-negative matrix factorization techniques are reviewed and some new ones are described.

2.1 Multiplicative Method.

A non-negative matrix factorization algorithm of Lee and Seung [13] is based on multiplicative update rules of W and H . We call this scheme the *multiplicative method*, and denote it by **MM**. A formal statement of the method is given next.

Algorithm for MM

1. Initialize W and H with non-negative values, and scale the columns of W to unit norm.
2. Iterate for each c, j , and i until convergence or stop:

$$(a) \quad H_{cj} \leftarrow H_{cj} \frac{(W^T X)_{cj}}{(W^T W H)_{cj} + eps}$$

$$(b) \quad W_{ic} \leftarrow W_{ic} \frac{(X H^T)_{ic}}{(W H H^T)_{ic} + eps}$$

- (c) Scale the columns of W to unit norm.

Clearly the approximations W and H remain non-negative during the updates. It is generally best to update W and H “simultaneously”, instead of updating each matrix fully before the other. In this case, after updating a row of H , we update the corresponding column of W . Matlab performs well with these computations. In the implementation, a small positive quantity, say the square root of the machine precision, should be added to the denominators in the approximations of W and H at each iteration step. We use a parameter $eps = 10^{-9}$ in our Matlab code for this purpose.

It is often important to normalize the columns of X in a pre-processing step, and in the algorithm to normalize the columns of the basis matrix W at each iteration. In this case we are optimizing on a unit hypersphere, as the column vectors of W are effectively mapped to the surface of a hypersphere by the repeated normalization.

The computational complexity of Algorithm MM can be shown to be $O(rnm)$ operations per iteration. However, the initial factorization needs only be computed in its entirety once.

If data are added to the database then it can either be added directly to the basis matrix W along with a minor modification of H , or else if k is fixed, then further iterations can be applied starting with the current W and H as initial approximations.

Lee and Seung [14] proved that under the MM update rules the distance $\|X - WH\|_F^2$ is monotonically non-increasing. In addition it is invariant if and only if W and H are at a stationary point of the objective function (1). From the viewpoint of nonlinear optimization, the algorithm can be classified as a diagonally-scaled gradient descent method [6].

Several other methods have been proposed for non-negative matrix factorization. (See [15] for a recent survey). In particular, Lee and Seung [13] have also provided an additive algorithm. Both the multiplicative and additive algorithms are related to expectation-maximization approaches used in image processing computations such as image restoration, e.g., [19]. Hoyer [8] has suggested a novel non-negative sparse coding scheme based on ideas from the study of neural networks, and the scheme has been applied to the decomposition of databases into independent feature subspaces by Hyvärinen and Hoyer [9]. Other methods based in part on orthogonal factorizations have been considered by Liu and Yi [15]. Studies and comparisons of various algorithms for non-negative factorizations have been given by D. Guillamet et al [6, 7] and by Liu and Yi [15]. Each approach has advantages and disadvantages, but in general the algorithms examined in these papers have been found to be effective on certain applications, and so the method of choice is often application dependent.

2.2 Non-Negative Factorization with Point Count Regularization.

In this section we present a new NMF algorithm based on a sparsity requirement for H . Hoyer's method [8] has the important feature of enforcing a statistical sparsity for the weight matrix H , thus enhancing the parts-based representation of the data in W . The method requires choice of a regularization, or penalty, parameter for use in an update scheme for H . In the next section we propose a new NMF approach which is similar to the neural network-based scheme of Hoyer, but where sparsity in H is enforced by a new point count regularization scheme by Mu, Plemmons and Santago [17]. We propose a least squares approach to solve

$$\min_{W,H} \|X - WH\|_F^2,$$

which uses a point count regularization scheme [17] in the computations that penalizes the number of nonzero entries in the weight matrix H , often leading to a basis representation in W that better represents parts or features on the database information in X . Again, the approach relies on neural network concepts and forms an extension of the work of Hoyer [8]. The algorithm we propose uses a gradient descent-type scheme for iterating with respect to W , and a point count regularized multiplicative-type scheme for iterating with respect to H . We denote the algorithm by **SC-PC** for non-negative matrix factorization with *sparse coding by point count regularization*.

Algorithm for SC-PC

1. Initialize W and H with non-negative values, and scale the columns of W to unit norm.

2. Iterate until convergence or stop:

- (a) $W \leftarrow W - \mu(WH - X)H^T$
- (b) Set any negative values of W to zero.
- (c) Rescale the columns of W to unit norm.
- (d) $H \leftarrow H.*(W^T X./(W^T W H + \lambda./(H + \gamma).^2)$

Matlab notation $.*$, $./$, and $.^2$, denoting component wise multiplication, division and exponentiation, respectively, are used in the update formula for H in Step 2(d).

2.3 A Hybrid Method.

In this section, we propose a hybrid algorithm for NMF that combines some of the better features of the methods discussed in Sections 2.1 and 2.2. First, we adopt the multiplicative algorithm approach for computing an approximation to the basis matrix W at each iterative step. This computation is essentially a matrix version of the gradient descent optimization scheme. Secondly, we compute the weight matrix H using a constrained least squares (CLS) model as the metric. The purpose is to penalize non-smoothness and non-sparsity in H . This CLS model is related to the least squares Tikhonov regularization technique commonly used in image restoration [19]. Our algorithm, which we denote **GD-CLS** for *gradient descent with constrained least squares*, is given next.

Algorithm for GD-CLS

1. Initialize W and H with non-negative values, and scale the columns of W to unit norm.
2. Iterate until convergence or stop:

- (a) $W_{ic} \leftarrow W_{ic} \frac{(X H^T)_{ic}}{(W H H^T)_{ic} + eps}$, for each c and i
- (b) Rescale the columns of W to unit norm.
- (c) Solve the constrained least squares problem:

$$\min_{H_j} \{ \|X_j - W H_j\|_2^2 + \lambda \|H_j\|_2^2 \},$$

where the subscript j denotes the j^{th} column, for $j = 1, \dots, m$. Any negative values in H_j are set to zero.

As done in Algorithm MM, we use a small positive parameter eps to avoid dividing by zero or very small numbers and enhance stability in the computations for W in Step 2(a). A small positive number λ is involved in computing the weight matrix H . In particular, λ is a regularization parameter that is used to balance the reduction of the metric

$$\|X_j - W H_j\|_2^2$$

with enforcement of smoothness and sparsity in H . Our numerical approach for solving the constrained least squares problem in Step 2(c) for the columns H_j of H makes use of an algorithm similar to one that we employed in [19] for regularized least squares image restoration. Application of the gradient descent with constrained least squares algorithm GD-CLS for NMF methodology in mining spectrographic data is considered next.

3 Non-Imaging Space Object Identification

In this section we introduce the target application area of this paper: mining scientific data for non-imaging identification and classification of space objects. Determining the status of space based assets is of national concern. Ground based telescopes have been used for years to image space based objects to determine their function and status. The sophistication of imaging techniques, such as adaptive optics and post processing algorithms, have made it possible to obtain highly resolved images of satellites and other space based assets. Ultimately the telescope and observed wavelength limit the achievable spatial resolution of any object. Satellites orbiting beyond 1000 kilometers generally cannot be resolved with enough spatial resolution to determine their status, and geosynchronous satellites which orbit much further out, 40000 kilometers, appear much like a point source to ground based telescopes.

It is therefore necessary to collect alternative information that is not spatially dependent. Spectral information using spectrometers interfaced to ground based telescopes have proven to be a useful tool in collecting information on unresolvable satellites. Complexity arises because spacecraft are composed of many different materials each with their unique spectral signature, and observations of in-orbit objects represent the additive optical contributions of each material component. Spectral traces are collected for multiple orientations of a satellite, exposing different material surfaces to illumination. Spectral traces of space based materials measured in a laboratory provide a reference for identifying absorption features inherent in a particular material. Depending on the orientation of the satellite, different proportions of each material type will be illuminated, such as solar panels or spacecraft body. By comparing the observational data with a database of material spectra it seems possible to determine the various materials that compose the satellite. This information is useful in identifying a satellite and determining its function.

Important research has already been achieved in the field of image analysis using spectral information. Multispectral sensing has mostly succeeded at classifying whole pixels although analysis of the constituents that make up a pixel are limited by the low number of wavelengths. However, hyperspectral sensing has demonstrated the capability of performing spectral unmixing on individual pixels by collecting data in numerous spectral bands. An example of hyperspectral sensing that has been used for many years is satellites that image the earth. A series of images of the same scene are obtained at different wavelengths, which provides a spectral signature for each image pixel. Spectral unmixing algorithms estimate the materials that contribute to each pixel in the series of images. One can use either linear or non-linear models to describe the observed spectral signature reflected off the materials. The inversion problem involves mapping material signatures to the observed signature to determine what materials are present and in

what quantity on the object.

NMF is a very promising technique in this respect due to the fact that the non-negativity constraint should yield results that more closely match material spectra than unconstrained methods such as PCA, which can produce unphysical negative values in its basis vectors. In particular, Wild, Curry and Dougherty [21] have recently shown that NMF techniques are very useful in hyperspectral sensing applications where satellites collect spectral radiance data for characterization of the Earth’s surface and atmosphere.

4 Numerical Experiments with Spectral Data

In this section, we present preliminary results regarding the performance of non-negative matrix factorization algorithms, described in Section 2, in the identification of constituent material spectra from spectral traces of non-imaging space objects. This important application was described in the previous section. We employed a large database of spectral data obtained using a spectrometer located in Maui, HI. Next, we describe the particular characteristics of the spectral dataset under consideration and show preliminary computational results.

4.1 Spectral Reflectance Data Characteristics

Spectroscopic data is promising for non-imaging SOI, and several material types have already been successfully observed in spectra of in-orbit spacecraft, notably white-paint and aluminum. Complexity arises because spacecraft are composed of many different materials each with their unique spectral signature, and observations of in-orbit objects represent the additive optical contributions of each material component.

The data utilized for this project was collected from the Spica spectrometer located on the rear blanchard of the 1.6 meter telescope at the Air Force Maui Optical and Supercomputing Site (AMOS) atop Haleakala, Maui. The instrument is designed around a commercially available Acton Spectrometer with a selectable three-position grating turret. In its current configuration the sensor provides 3-4 angstrom resolution spectra in two different modes, a blue mode, which spans the color regime from 4000 angstroms to 7000 angstroms, and a red mode, which collects 6000 to 9000 angstrom spectra. The spectrometer collects data through a 14 arcsecond diameter aperture, which allows all of the light from an object to enter the instrument and therefore provides absolute photometric information at the expense of slightly limiting the spectral resolution. The spectrometer CCD is a Princeton instruments 1100×330 liquid nitrogen cooled array with 25-micron pixels. The sensor and 1.6 meter telescope are shown in Figure 1.

Spica has supported several non-imaging SOI projects since 1999 including the Space Object Identification in Living Color (SILC) project and the NASA/AMOS Spectral Study (NASS). Spica continues to collect data on space objects and is building a sizable database of spectra of space objects. Currently, the database contains data on GEO and LEO satellites and a sizable number of Rocket bodies in both wavebands. It has been noted that the SOI application could require many observations at varying lighting and observing angles in order to provide the




Figure 1: The 1.6m diameter telescope at the AMOS facility is home to the Spica spectrometer on the rear blanchard.

statistical variation that could allow the implementation of methods such as neural networks in a generalized way.

In view of this, the Spica spectrometer and AMOS are dedicated to collecting observations of many different satellites at a wide variety of view geometries in order to better support SOI research projects. Spica utilizes the Image Reduction and Analysis Facility (IRAF) suite of software developed at Kitt Peak National Observatory in combination with Matlab tools designed on-site to reduce the data into final form. The IRAF reduction routines remove atmospheric effects and calculate instrumental sensitivity through the comparison of standard stars with observations, the results are calibrated intensity spectra in units of $\text{ergs}/\text{cm}^2/\text{sec}/\text{angstrom}$. Matlab tools have been developed to remove solar features by dividing out solar-type stars, which are also collected through the instrument on the night of observation. The end result of this process is a spectra that represents the color reflectance of the material features of the spacecraft.

Recent work has focused on scaling the solar analog stars to the intensity of the sun before division, which would lead to a value more closely representing the reflectance of the object. Figure 2, from a NASA report by Jorgensen et al [12], shows a Spica spectra of a white painted rocket body matched with a laboratory spectra of white paint. Figure 3 shows laboratory spectra of mylar, white paint, aluminum and solar cell material.

For this paper, an initial dataset of 1196 different spectral traces pertaining to 22 different space objects was considered. The data was collected over a period of several months. Each trace is stored in an individual data file. The file contains information regarding maximum and minimum intensity values, intensity units (normalized or $\text{ergs}/\text{cm}^2/\text{sec}/\text{angstrom}$), wavelength units (angstroms), as well as wavelength-intensity pairs (λ_i, r_i) . The values $\lambda_i = \lambda_0 + i\Delta$ are discrete values of wavelength and r_i are the corresponding intensity values. The step size Δ is a value between 3 and 4 (angstroms) and λ_0 varies per trace. Each trace contains about 1100 wavelength-intensity pairs. All traces in the dataset were corrected for cosmic rays and had background subtractions and the effects of atmospheric absorption removed. A second

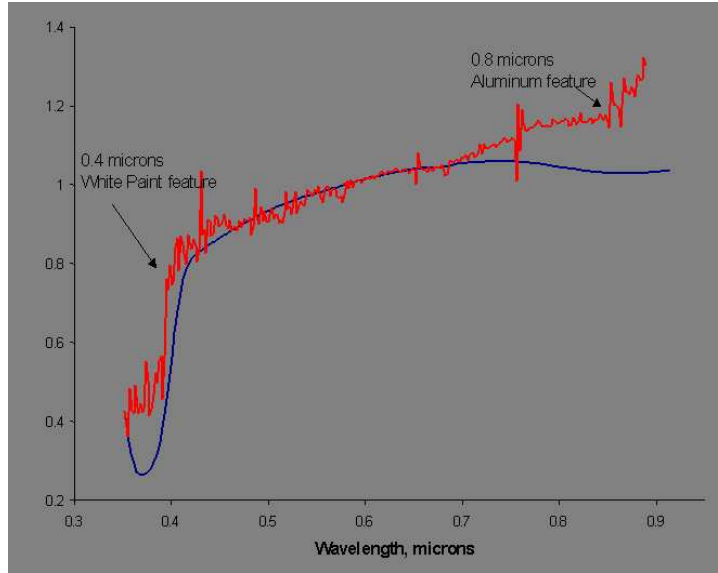


Figure 2: A match of a white painted rocket body to a laboratory white paint spectra.




Figure 3: Laboratory spectral signatures of mylar (left), white paint (center-left), aluminum (center-right), and solar cell (right) material.

dataset was obtained by post-processing the initial dataset to remove solar features as previously described. A final dataset containing 2392 spectral traces was obtained by composition of the first two datasets in order to maximize statistical variability.

4.2 Computational Results

For numerical computation purposes, a data matrix X was formed using the final dataset previously described. Because Δ and λ_0 varied per trace, all intensity values r_i were mapped to a finer wavelength grid of length 5732, allowing tabulation of intensities (columns of X) for roughly the same values of wavelength (rows of X). The dimensions of the resulting matrix X were then 5732-by-2392.

The GD-CLS algorithm described in Section 2 was used for the numerical computation of the non-negative factors W and H in the approximation of X . The number of columns k of W (or equivalently number of rows in H) was experimentally set to $k = 30$. This value also represents

an upper bound in the rank of the approximation matrix WH . The control parameter λ was set to 10^{-5} , reflecting our decision on how to balance sparsity in H with closeness of the low-rank non-negative factorization WH to X . Let $W^{(i)}$ and $H^{(i)}$ denote the approximations of W and H , respectively, at each step of the algorithm for $i = 1, 2, \dots$. At each iteration, the following cost function value was computed:

$$c_i = \|X - W^{(i)}H^{(i)}\|_F + \lambda\|H^{(i)}\|. \quad (2)$$

The iterations were terminated when $|c_i - c_{i+1}|/c_i < \tau$ for $\tau = 10^{-4}$.

The algorithm took 100 iterations to terminate. The final value of c_i was $c_{100} = 170.2901$. Better approximations to X can obviously be obtained by choosing a larger k . However a value of $k \ll n = 1192$ was sought in order to significantly reduce the dimensionality of the associated query problem. Moreover, we note that the strength of non-negative matrix factorization algorithms, such as the GD-CLS algorithm, is in producing a *parts-based* matrix W , that is, a set of basis vectors exhibiting intrinsic constituent features in the spectral data. Figure 4 provides selected columns of W showing spectral traces that are strongly correlated with mylar, white paint, and solar cell materials.




Figure 4: Selected columns of W showing possible intrinsic material spectra: mylar (left), white paint (center), and solar cell (right).

Since the columns of W are used as basis vectors in the approximation of each spectral trace, i.e. $X_j \approx WH_j$, the value of each coefficient H_{ij} represents the contribution of basis vector W_i for $i = 1 \dots 30$ in the approximation. Figure 5 shows a bar plot of the coefficients H_{ij} for each basis vector W_i , corresponding to the three spectral traces X_1, X_2, X_3 in X . The basis vectors W_1, W_2, \dots, W_{30} are associated with the entries in the x-axis. Instead of numbers $1, 2, \dots, 30$, we use labels M, P, S for mylar, white paint, and solar cell, respectively. The label U indicates a basis vector corresponding to a material that was not determined to be either mylar, aluminum, white paint, or solar cell material. The spectral traces X_1, X_2 and X_3 correspond to three different observations of the same space object, a satellite known as JCSAT 1, and are illustrated in Figure 6.

The coefficients H_{ij} associated with basis vectors corresponding to the same material type results in so-called fractional abundances. Figure 7 shows a bar plot of fractional abundances




Figure 5: Relative fractional abundances obtained using the basis representation with $k = 30$ vectors for materials found in spectral traces X_1, X_2 and X_3 . These traces correspond to an earth-orbiting satellite known as JCSAT 1. The letters M, P, S, and U stand for materials mylar, white paint, solar cell, and unknown/undetermined, respectively.




Figure 6: Three different observations X_1, X_2 and X_3 of the same space object (JCSAT 1).

for mylar, white paint, and solar cell material. Coefficients for basis vectors not determined to be one of the three types are shown in the fourth column. These relative fractional abundances have not yet been verified against known data. However, our preliminary results illustrate the potential for use of non-negative matrix factorization algorithms in solving difficult and important problems in non-imaging space object identification.

5 Concluding Remarks and Future Work

In this paper we have studied novel spectral unmixing techniques for clustering, retrieval, analysis, and classification of and spectral sensor data for orbiting space objects. The methodology is based upon blind source separation by encoding the data using low-rank non-negative matrix factorization algorithms to preserve natural data non-negativity and thus avoid subtractive basis vector and encoding interactions present in techniques such as principal component analysis. We have employ unsupervised spectral unmixing techniques for recovery of constituent spectra, where the both the basis matrix W and mixing matrix X are determined using numerical opti-




Figure 7: Relative fractional abundances for mylar, white paint, solar cell, and undetermined material.

mization techniques, and prior information such as non-negativity. This approach is in contrast to the recent work of Luu, Matson, Snodgrass, Giffin, Hamada and Lambert, which involves supervised spectral unmixing where the basis matrix W is assumed known and fixed for a given space object database [16].

We have illustrated by tests on actual space object spectral observations obtained from the Spica spectrometer that NMF techniques can be successfully used to identify object parts, i.e. specific spectral signatures as basis vectors for the spectral database, as well as fractional abundances of materials present in space objects, such as white paint, aluminum, mylar and solar cell material. Our preliminary results point to NMF as a very promising approach for this application. We are continuing to study and quantify the effectiveness of blind source separation by encoding the data using low-rank non-negative matrix factorization for non-imaging space object identification, and related applications. Simulation studies will be performed using spectral datasets of objects with known constituent materials, such as the Galaxy V satellite simulation dataset used in [16]. Future work may also include exploiting the temporal dependences of the spectral data to obtain basic shape information of space objects.

Other application areas of unsupervised spectral unmixing with NMF techniques of the type applied in this paper include: hyperspectral remote sensing for spectral unmixing of materials on the Earth’s surface from ARVIS data, Raman spectroscopy utilizing laser irradiation to provide information about biological tissue, chemical shift imaging, nuclear magnetic resonance spectroscopy, text mining, face recognition, speech recognition, and extracting summarizing excerpts from audio and video (see, e.g, [3, 4, 6, 9, 10, 13, 19, 20, 21]).

References

- [1] A. Berman and R. Plemmons. *Non-Negative Matrices in the Mathematical Sciences*, SIAM Press Classics Series, Philadelphia, 1994.

- [2] M. Berry, Z. Drmac, and E. Jessup. “Matrices, Vector Spaces, and Information Retrieval”, *SIAM Review*, Vol. 41, pp. 335-362, 1999.
- [3] M. Catral, L. Han, M. Neumann and R. Plemmons. “Reduced Rank Non-Negative Factorization for Symmetric Non-Negative Matrices”, preprint, 2003, to appear in *Linear Algebra and Applications*. See <http://www.wfu.edu/~plemmons>.
- [4] M. Cooper and J. Foote, “Summarizing Video using Non-Negative Similarity Matrix Factorization”, *Proc. IEEE Workshop on Multimedia Signal Processing* St. Thomas, US Virgin Islands, 2002.
- [5] D. Donoho and V. Stodden. “When does Non-Negative Matrix Factorization Give a Correct Decomposition into Parts?”, preprint, Department of Statistics, Stanford University, 2003.
- [6] D. Guillamet, B. Schiele and J. Vitria. “Analyzing Non-Negative Matrix Factorization for Image Classification”, *16th International Conference on Pattern Recognition (ICPR'02)*, Vol. 2, Quebec City, Canada, 2002.
- [7] D. Guillamet and J. Vitria. “Determining a Suitable Metric when Using Non-Negative Matrix Factorization”, *16th International Conference on Pattern Recognition (ICPR'02)*, Vol. 2, Quebec City, QC, Canada, 2002.
- [8] P. Hoyer. “Non-Negative Sparse Coding”, *Neural Networks for Signal Processing XII (Proc. IEEE Workshop on Neural Networks for Signal Processing)*, Martigny, Switzerland, 2002.
- [9] A. Hyvrinen and P. Hoyer. “Emergence of Phase and Shift Invariant Features by Decomposition of Natural Images into Independent Feature Subspaces”, *Neural Computation*, Vol. 12, pp. 1705-1720, 2000.
- [10] N. Keshava and J. Mustard. “Spectral Unmixing”, *IEEE Signal Processing Magazine*, pp. 44-57, January 2002.
- [11] I. Jolliffe. *Principle Component Analysis*, 2nd Ed., Springer Series in Statistics, Springer-Verlag, New York, 2002.
- [12] K. Jorgensen et al. “Squigley Lines and Why They are Important to You”, *NASA Lincoln Space Control Conference*, Lexington, MA, March 2002.
- [13] D. Lee and H. Seung. “Learning the Parts of Objects by Non-Negative Matrix Factorization”, *Nature*, Vol. 401, pp. 788-791, 1999.
- [14] D. Lee and H. Seung. “Algorithms for Non-Negative Matrix Factorization”, *Advances in Neural Processing*, 2000.
- [15] W. Liu and J. Yi. “Existing and New Algorithms for Non-negative Matrix Factorization”, preprint, Computer Sciences Dept., UT Austin, 2003.

- [16] K. Luu, C. Matson, J. Snodgrass, M. Giffin, K. Hamada and J. Lambert. “Object Characterization from Spectral Data”, *Proc. AMOS Technical Conference*, Maui, HI, 2003.
- [17] Z. Mu, R. Plemmons and P. Santago. “Iterative Ultrasonic Signal and Image Deconvolution for Estimating the Complex Medium Response”, preprint, to appear in *IEEE Transactions on Ultrasonics and Frequency Control*, 2003. See <http://www.wfu.edu/~plemmons>.
- [18] P. Pauca, F. Shahnaz, M. Berry and R. Plemmons, “Text Mining using non-negative Matrix Factorizations, preprint, 2003. See <http://www.wfu.edu/~plemmons>.
- [19] S. Prasad, T. Torgersen, P. Pauca, R. Plemmons, and J. van der Gracht. “Restoring Images with Space Variant Blur via Pupil Phase Engineering”, *Optics in Info. Systems, Special Issue on Comp. Imaging*, SPIE Int. Tech. Group Newsletter, Vol. 14, No. 2, pp. 4-5, 2003. See <http://www.wfu.edu/~plemmons>.
- [20] P. Sajda, S. Du, T. Brown, L. Parra and R. Stoyanova. “Recovery of constituent spectra in 3D chemical shift imaging using non-negative matrix factorization”, in *Proc. 4th International Symposium on Independent Component Analysis and Blind Source Separation*, pp. 71-76, 2003.
- [21] S. Wild, J. Curry and A. Dougherty. “Motivating Non-Negative Matrix Factorizations”, *Proceedings of the Eighth SIAM Conference on Applied Linear Algebra*, Williamsburg, VA, July 15-19, 2003. See <http://www.siam.org/meetings/la03/proceedings/>.