Principal component analysis in limited-angle chromotomography

J. M. Mooney[a], A. K. Brodzik[b], and M. An[c]

[a]Rome Laboratory, RL/EROI, Hanscom AFB, MA, 01731

[b]Scientific Software, 50-113 Cambridge Rd., Woburn, MA 01801

[c]Solid State Scientific Co., 27-2 Wright Rd., Hollis, NH 03049

ABSTRACT

Chromotomographic spectral imaging techniques offer high spatial resolution, moderate spectral resolution and high optical throughput. However, the performance of chromotomographic systems has historically been limited by the artifacts introduced by a cone of missing information. The recent successful application of principal component analysis to spectral imagery indicates that spectral imagery is inherently redundant. We have developed an iterative technique for filling in the missing cone that relies on this redundancy. We demonstrate the effectiveness of our approach on measured data, and compare the results to those obtained with a scanned slit configuration.

Keywords: hyperspectral imaging, limited-angle tomography, principal component analysis.

1. INTRODUCTION

Our spectral imager is similar to previously reported chromotomographic devices,[1–5] all of which image a dispersed view of the entire scene over all of the spectral bands, then computationally reconstruct the chromatic image. In chromotomography, the color imagery is considered to be pseudo-three dimensional. The 3D spectral image is obtained using tomographic imaging techniques similar to those used in medical imaging. While chromotomography is computationally demanding, it has the advantage that a slit field stop is not required. We anticipate the spectral imager performance will reflect the larger optical throughput afforded by a larger field stop. The actual performance and potential multiplex advantage, hinge on the effectiveness of the reconstruction algorithm.

Our technique differs from other chromotomographic approaches in in that we utilize a direct vision prism to disperse the light, and we have designed a reconstruction method that is well suited to our hardware configuration.[6,7] Figure 1 is a schematic representation of our approach. It consists of a telescope, a direct vision prism and a camera. A direct vision prism consists of two prisms that are arranged such that one wavelength passes through undeviated while other wavelengths are deviated along a line. The direct vision prism is mounted in a bearing so that it can be rotated around the optical axis. As the prism is rotated the spectral features in the image will trace out circles where the radius of the circles are wavelength dependent. To obtain a spectral image, a series of frames of camera data are obtained with the prism in different angular orientations, then a computer is used to extract the spectral image from the measured data.

2. THE CONE OF MISSING INFORMATION

The origin of the missing cone is best understood in the context of the Radon transform combined with the central slice theorem.[8] From the central slice theorem, we know that the 2D Fourier transform of a projection is equal to a slice through the origin of the 3D Fourier transform of the spectral image. By rotating the prism, we obtain different slices through the origin, each for a different projection direction, but with the same angle, θ, between the projection plane and the optical axis (θ is defined relative to the normal to the plane). In the limit of an infinite number of

SPIE Vol. 3118 • 0277-786X/97/$10.00

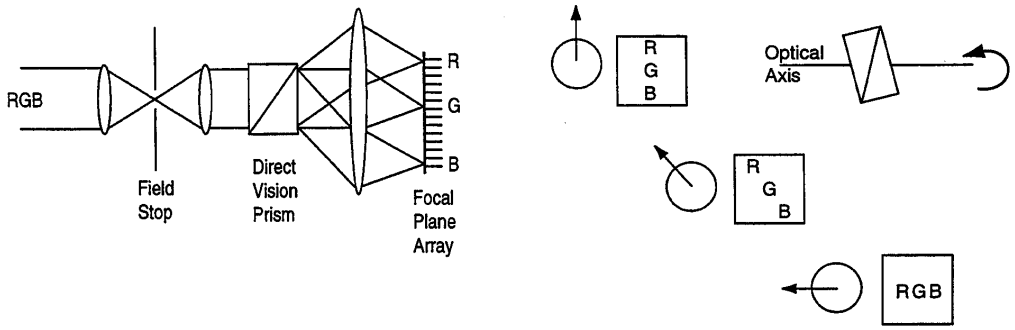

Figure 1. A schematic representation of the spectral imager and its operation. On the left the direct vision prism is shown spreading red, green and blue light across the FPA. On the right the effect of the prism rotation on the overlapping red, green and blue images is illustrated. The circle and arrow represent the prism orientation.

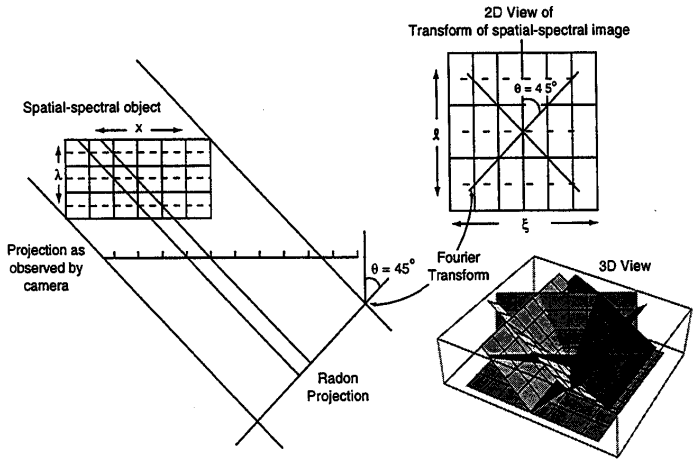

Figure 2. The correspondence of our recorded data to Radon projections. The projection angle follows from the definition that adjacent spectral bands are displaced by one pixel on the FPA and the assumption that the object volume elements are equilateral. The spectral image can be reconstructed in the frequency domain by using the central slice theorem (shown in both 2- and 3-dimensions).

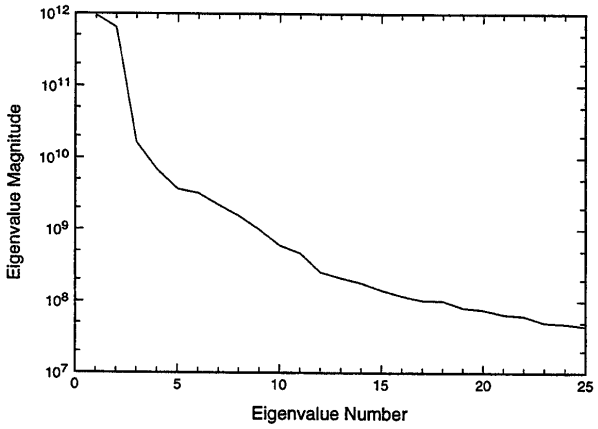

Figure 3. The eigenvalues of the covariance matrix for AVIRIS data taken of Jasper Ridge.

different projections, the projections will define a cone, where the half-angle of the cone vertex is θ. Since none of the projections provide any information about the image inside this cone, it is referred to as a 'missing cone'.

Figure 2 illustrates that the image through the direct vision prism is analogous to a projection through the pseudo-volume defined by two spatial and one spectral dimension. From Fig. 2 we can also identify the cone (triangle in the 2D view) of missing information. To estimate the impact of the missing cone on our imaging system, we calculate the fractional volume of the missing cone, which is determined by the projection angle θ. Our interpretation of the projection angle is governed by: 1) our definition of adjacent spectral bands as displaced by one pixel on the focal plane, and 2) our assumption that the volume object elements (voxels) are cubic. The definition is easily justified since it is consistent with scanned slit configurations. The assumption is completely arbitrary, but has the benefit of symmetric simplicity and leads to a missing cone half-angle of 45 degrees.

3. CHARACTERISTICS OF HYPERSPECTRAL IMAGERY

Spectral imagery is a three dimensional representation of the spatial-spectral characteristics of a scene. In general, it is possible that every spatially resolved pixel could consist of a unique spectral signature. However, it is far more likely that large regions of the scene will have similar spectra. For example, a forest might contain several different types of trees, shrubbery, and soil, but all of the trees of a given type will have similar spectra.

The reduced set of spectra that are required to represent a particular spectral image are the eigenvectors of the covariance matrix of that spectral image. Figure 3 illustrates that the first few eigenvalues are dominant. The eigenvectors that correspond to the largest eigenvalues are the basis spectra required to define a sub-space on which the spectral imagery can be projected without a significant loss of information.

The observation that an incomplete spectral basis can be used to describe spectral image is central to the technique of principal component analysis (PCA).[9] PCA is a technique for analyzing data of very high dimensionality; one first decomposes the data using singular value decomposition (SVD), then examines projections of the data onto planes suggested by the SVD. PCA has emerged as a promising technique for analyzing hyperspectral data.[10]

If only a sub-set of the basis spectra are required to effectively analyze the hyperspectral imagery, the data is redundant. We intend to use the redundancy of the hyperspectral imagery to fill in the missing cone. Our approach is

Other author information:
J.M.M.: Email: jon@max.rl.plh.af.mil; Telephone: 617-377-2876; Fax: 617-377-4814
A.K.B.: Email: brodzik@max.rl.plh.af.mil; Telephone: 617-377-4046; Fax: 617-373-4814
M.A.: Email: myoung@tiac.net; Telephone: 617-254-9149; Fax: 617-377-4814

to take advantage of the fact that the data redundancy forces the data values inside the missing cone to be consistent with those values outside. Any variation from the correct missing cone values will introduce additional spectra into the reconstruction.

4. METHOD FOR FILLING IN THE MISSING CONE

Existing methods for reconstructing chromotomographic data assume that the data values inside the missing cone are zero. This assumption is rational when we consider that both a magnitude and phase are required, and that the phase is uniformly distributed over 2π. Our approach is to approximate the missing cone values by based on the leading eigenspectra of the previous reconstruction and iterate. Since the missing cone is only 6.5% of the total volume, most of the data is unchanged as we iterate. After the final iteration, the updated values should be reasonable estimates of the actual values, and the impact of the missing cone on the reconstruction should be minimized.

The existence of a missing cone implies an infinite number of spectral image reconstructions are consistent with the measured data. Some of these reconstructions differ only slightly from the ideal reconstruction while others are completely unsatisfactory. We can eliminate the unsatisfactory solutions by requiring that the reconstruction meet an additional constraint. Our problem becomes one of selecting an additional constraint, and finding an algorithm that is consistent with both the measured data and the additional constraint.

The constraint we use is based on the assumption that incorrect data values inside the missing cone are likely to generate spectra that fall outside the sub-set of eigenspectra established by the true spectral image. Therefore, we impose the constraint the reconstruction should minimize the number of eigenspectra required to describe the spectral reconstruction.

If all of the scene information fell inside the missing cone, then our algorithm would fail. However, the field stop edges introduce high spatial frequency information into the measured data which tends to keep the data from being confined inside the cone. While our analysis indicates that the algorithm is very robust, it is important to emphasize that the performance of our approach is potentially data dependent.

The algorithm for using the data redundancy to fill in the missing cone is iterative and is similar to those in the literature.[11,12] A flowchart of the algorithm is shown in Fig. 4. After forming the initial reconstruction, it is transformed into the frequency domain. In the initial reconstruction the missing cone values are filled with zeros, in subsequent iterations updated estimates of the missing cone values are used. The reconstruction is then transformed back to the spatial domain and it is decomposed into its eigenspectra. A reduced representation of the reconstruction is obtained by projecting the data onto the major eigenspectra. The reduced representation is transformed into the frequency domain, and used to estimate the missing cone values in the following iteration. The process is repeated until the eigenvalues of the reconstructions stabilize.

5. ALGORITHM DEMONSTRATION

To demonstrate the algorithm performance, we imaged a scene that contained objects with different spectral signatures. A controlled laboratory scene enabled us to investigate the effect of the missing cone on spectral reconstruction and spectral discrimination. We note in passing that the standard technique of imaging the exit slit of a monochromator to calibrate the spectral imager will give optimistic results, since a slit object has a lot of high spatial-frequency energy that falls outside the missing cone. Our demonstration scene consists of extended objects, where a large fraction of the scene energy falls inside the missing cone, creating more of a challenge for our reconstruction algorithm.

A visible image of the scene is shown in Fig. 5. It consists of a differential source and a reflective plate. The differential source consists of two plates one in front of the other. The front plate has a four-bar pattern of holes so that the back plate can be seen through the holes. The back plate is held at a temperature elevated above that of the front plate so that in the infrared the source appears as four bright bars.

The reflective plate was illuminated with an incandescent lamp that was controlled by a variable autotransformer. The lamp voltage was adjusted so that in the broad-band infrared image, the signal from the reflective plate matched that of the differential source. A broad-band infrared image of the scene is shown in Fig. 6.

Figure 7 is an image of the scene through the prism. An artifact of the dataset that will figure prominently in the following analysis is a ghost image apparent in the projection images as illustrated in Fig. 8. The ghost image

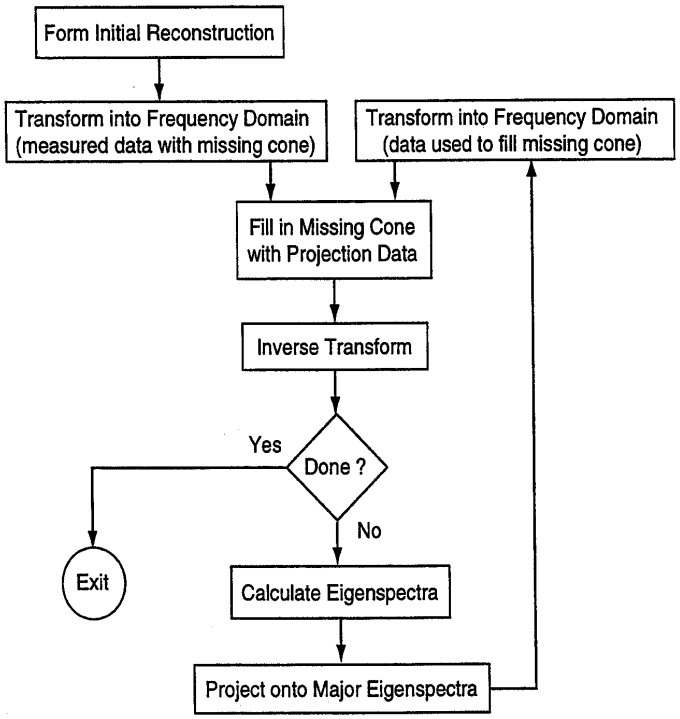

Figure 4. Flowchart of the algorithm used to fill in the missing cone.

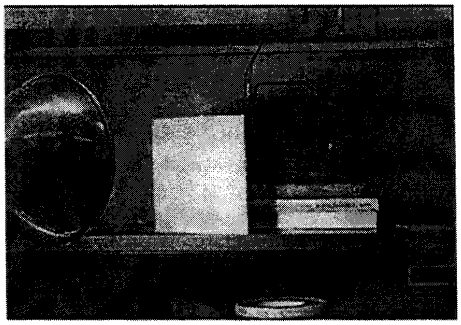

Figure 5. Visible image of the scene used for the demonstration. The reflective plate and differential source are illuminated by an incandescent lamp. Here the lamp intensity is much greater than when the experiment was performed so that the reflected light can be seen.

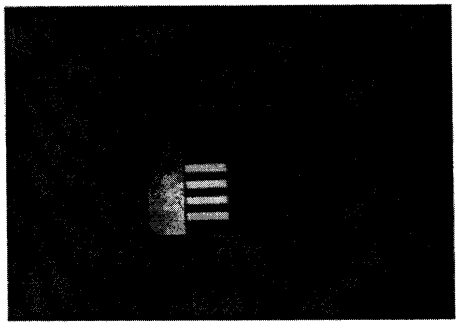

Figure 6. Broad-band infrared image of the scene used for the demonstration. The field stop limits the field of view to the central 100×100 pixels.

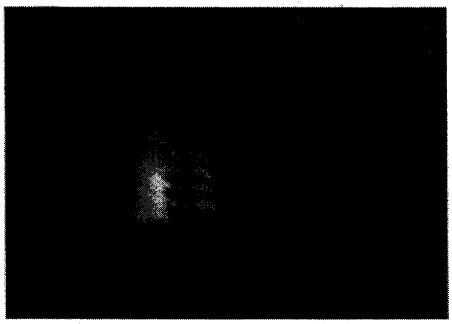

Figure 7. One of the 80 projections. Here the blurring is due to the prism dispersion.

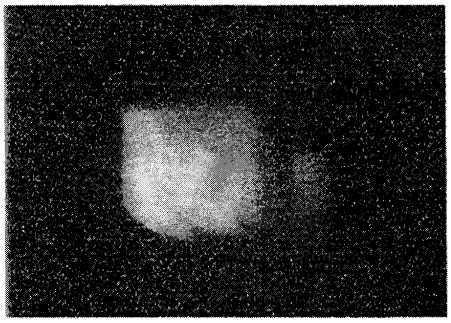

Figure 8. The same projection in Fig. 7 shown with increased gain. The ghost image of the plate is visible to the right of center.

is introduced by a pair of reflections in the optical path. Since the ghost image is inverted, it will appear to the reconstruction algorithm as if the image spectral distribution were reflected about the undeviated wavelength.

From the scene, we expect three major spectral signatures (the background, the hot bars and the reflective plate) and hope to use the differences between the spectral profiles of the plate and the bars to distinguish between them. The ghost image will introduce two additional spectral signatures, the inverted-bar and -plate signatures, both of which will have significantly reduced intensity. Since our imagery consists of five spectrally significant features, four eigenspectra are required to represent the data.

Figure 9 illustrates the spatial vectors of the SVD (spatial coefficients of the eigenspectra) for the initial reconstruction and after 9, 19, and 29 iterations. The vectors that correspond to the largest eigenvalues are across the top, followed by those corresponding to the second, third, fourth and fifth largest eigenvalues in order from top to bottom. In the left column the missing cone artifacts are observed as dark or bright halos surrounding the bar and plate features. As the number of iterations increases (moving from left to right in Fig. 9), the artifacts are pushed down to the less significant vectors.

Analysis of Fig. 9 reveals that after the initial inverse at least the first three terms are required to spectrally distinguish the bars from the plate. After 9 iterations, only the first two terms are required to distinguish the bars and plate. After 19 iterations, the reflection of the bars can be spectrally distinguished from that of the plate, and after 29 iterations, missing cone artifacts are difficult to distinguish in the fifth term. A comparison of the intensity of the ghost image to the true image, and a comparison of the eigenvalues of the various principal components indicates that the intensity of the missing cone artifacts is well below 0.3%. This number provides an upper bound on the spectral cross-talk, since the missing cone artifacts are its dominate contributing factor.

6. PERFORMANCE CHARACTERIZATION

Characterizing the performance of chromotomographic imaging devices is complicated by the presence of the missing cone. Conventional metrics such as spectral- and spatial-resolution do not reflect the presence of artifacts. (It appears that a three dimensional analog of MTF is required.) Even though characterization of artifacts is an unresolved issue, we can provide an initial evaluation of the merits of chromotomography. In particular we investigate the gain in signal-to-noise-ratio afforded by the area field stop in chromotomography versus the slit in a scanned-slit system.

To evaluate the performance of a scanned-slit system we replaced the area field stop in our system with a slit, and oriented the prism dispersion to be perpendicular to the long dimension of the slit. The slit width was adjusted so that the spectral resolution matched the chromotomographic configuration. A spectral signature from the slit configuration was obtained by sampling across one row of the focal plane array. Since the optical system, detector and electronics are the same, this comparison should reflect performance differences of spectral imaging techniques rather than instrumental variations.

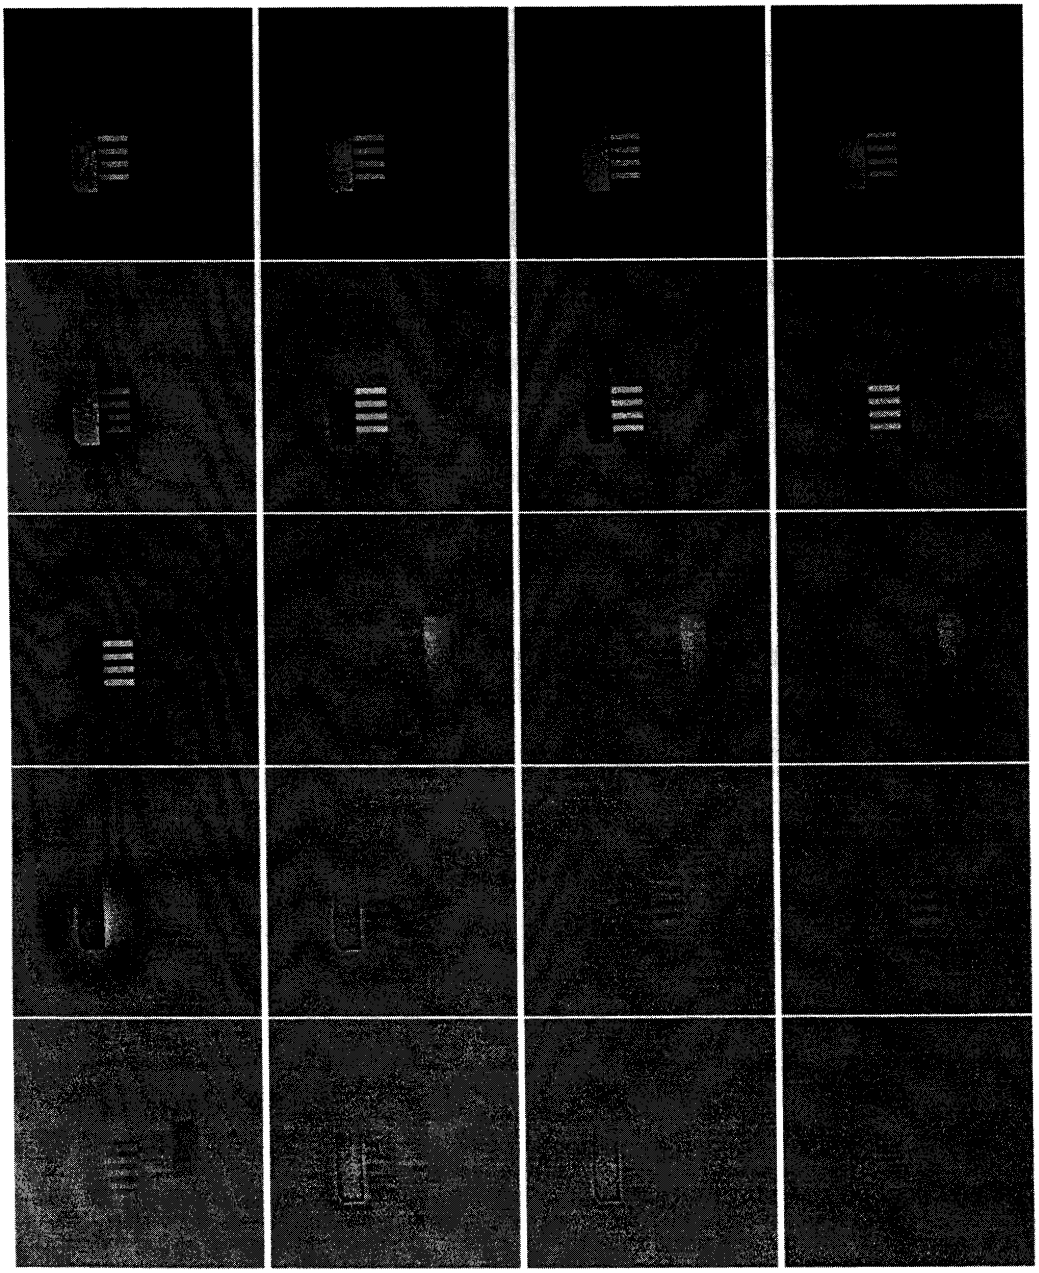

Figure 9. The vectors corresponding to the five largest singular values after 0 (left column), 9 (middle left), 19 (middle right), and 29 (right column)iterations

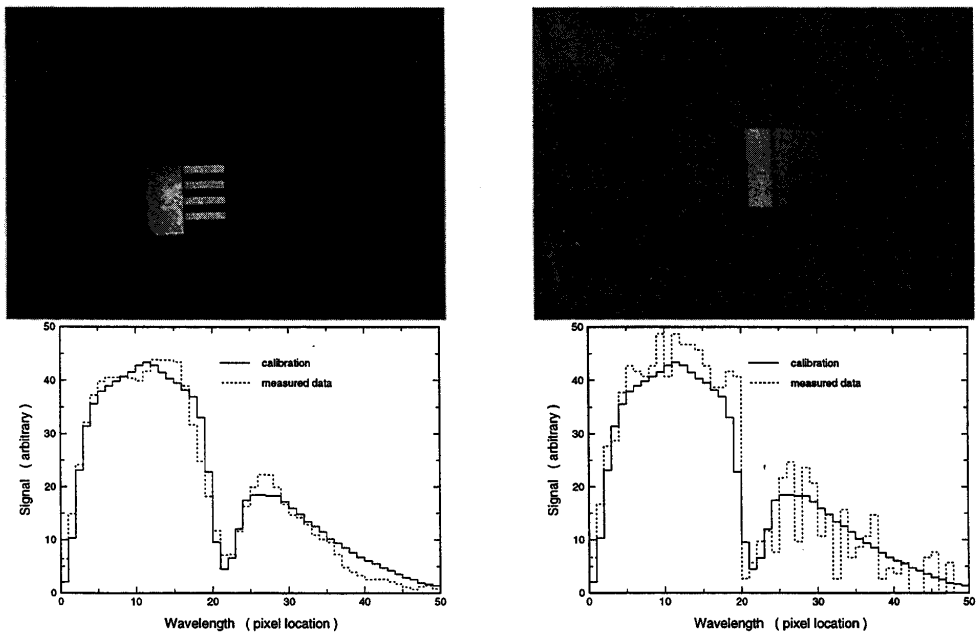

Figure 10. Comparison of chromotomography (left) and scanned slit (right) spectral imaging techniques. The slit data was obtained using a uniform source where the temperature was set to the same level as the bars on the left. The chromotomographic spectral profile was obtained for one pixel on a bar. The scanned-slit profile was obtained from one row of pixels.

In Fig. 10 results of the two spectral imaging approaches are compared. The calibration data was obtained by averaging 50 lines of data in 64 slit images, for a total of 3,200 spectral profiles. A comparison of the chromotomographic results to the calibration supports the claim that the slit width was adjusted to obtain comparable spectral resolution. Comparison of the noise in the measurements reveals that from pixel to pixel the noise is uncorrelated for the slit measurements, while it is correlated in the chromotomographic measurement. The correlation is introduced by the processing used to invert the system transfer function, and is a consequence of the missing cone.

If our techniques for filling in the missing cone functioned perfectly and the data were shot noise limited, we would expect the chromotomographic approach to out-perform the scanned-slit approach by the square-root of the ratio of the field stop area to the slit area, in this case by a factor of 7. The factor measured from the data in Fig. 10 is 2.8.

7. DISCUSSION

We have used our method for filling in the missing cone on several data sets with favorable results. We have also simulated the effectiveness of the iterative technique on AVIRIS data by introducing the missing cone into the data, then iterating to fill in the cone; again with favorable results. We intend to gather more data and further evaluate the performance and reliability of this approach.

Despite the success to date, additional algorithm development is needed. We have not yet established the requirements for convergence, the uniqueness of the reconstruction, or the number of eigenspectra to use in forming the missing cone estimate.

Our experience indicates that convergence is usually reached in less than 50 iterations. (A relaxation parameter can be introduced to speed up convergence at the risk of introducing oscillations.) It is an open question whether this technique can be used to fill the missing cone in other limited-angle tomographic imaging techniques.

8. CONCLUSIONS

We have demonstrated that the missing cone in chromotomography can be substantially filled by using an iterative algorithm combined with the constraint that the reconstruction have a minimal number of significant eigenspectra. This technique for evaluating the missing data enables us to approach the theoretical performance promised by chromotomographic imaging techniques, and promises a high efficiency alternative to conventional spectral imaging techniques.

REFERENCES

1. G. G. Levin and G. N. Vishnyakov, "On the possibilities of chronotomography of high speed processes," *Optics Communications* **56**, pp. 231–234, December (1985).
2. T. Okamoto, A. Takahashi, and I. Yamaguchi, "Simultaneous acquisition of spectral and spatial intensity distribution," *Applied Spectroscopy* **47**, pp. 1198–1202, August (1993).
3. Y. Betremieux, T. A. Cook, D. M. Cotton, and S. Chakrabarti, "Spinr: two-dimensional spectral imaging through tomographic reconstruction," *Optical Engineering* **32**, pp. 3133–3138, December (1993).
4. M. Descour and E. Dereniak, "Computed-tomography imaging spectrometer: experimental calibration and reconstruction results," *Applied Optics* **34**, pp. 4817–4826, August (1995).
5. P. A. Bernhardt, "Direct reconstruction methods for hyperspectral imaging with rotational spectrotomography," *J. Opt. Soc. Am. A* **12**, pp. 1884–1901, September (1995).
6. J. M. Mooney, "Spectral imaging via computed tomography," in *Proc. 1994 Meeting of the Infrared Information Symposia specialty group on passive sensors*, vol. 1, pp. 203–215, Defense Technical Information Center, (Alexandria, VA), March (1994).
7. J. M. Mooney, "Angularly multiplexed spectral imager," in *Imaging Spectrometry I*, M. R. Descour, J. M. Mooney, D. L. Perry, and L. Illing, eds., *Proc. SPIE* **2480**, pp. 65–77, (1995).
8. H. H. Barrett and W. Swindell, *Radiological Imaging*, vol. 2, pp. 375 - 446. Academic Press, New York, (1981).
9. I. T. Jolliffe, *Principal Component Analysis*, Springer-Verlag, New York, (1986).
10. J. W. Boardman, "Analysis, understanding and visualization of hyperspectral data as a convex set in n-space," in *Imaging Spectrometry I*, M. R. Descour, J. M. Mooney, D. L. Perry, and L. Illing, eds., *Proc. SPIE* **2480**, pp. 14–22, (1995).
11. A. Papoulis, "A new algorithm in spectral analysis and band-limited extrapolation," *IEEE Trans. Circuits and Sys.* **CAS-22**, pp. 735–742, September (1975).
12. K. C. Tam and V. Perez-Mendez, "Tomographical imaging with limited-angle input," *J. Opt. Soc. Am.* **71**, pp. 582–592, May (1981).