Convex projections algorithm for restoration of limited-angle chromotomographic images

Andrzej K. Brodzik

Scientific Software, Woburn, Massachusetts 01801

Jonathan M. Mooney

U.S. Air Force Research Laboratory, Hanscom Air Force Base, Massachusetts 01731

Received May 27, 1998; accepted August 5, 1988; revised manuscript received Spetember 21, 1998

We present a new algorithm for image restoration in limited-angle chromotomography. The algorithm is a generalization of the technique considered previously by the authors, based on a hybrid of a direct method of inversion and the iterative method of projections onto convex sets. The generalization is achieved by introducing a new object domain constraint. This constraint takes advantage of hyperspectral data redundancy and is realized by truncating the singular-value decomposition of the spatial–chromatic image matrix. As previously, the transform domain constraint is defined in terms of nonzero singular values of the system transfer function matrix. The new algorithm delivers high image fidelity, converges rapidly, and is easy to implement. Results of experiments on real data are included. © 1999 Optical Society of America
[S0740-3232(99)00802-9]

OCIS codes: 000.3860, 040.3060, 100.3020, 100.6950, 120.6200.

1. INTRODUCTION

Chromotomography provides an estimate of image intensity as a function of position and wavelength. The estimate is based on evaluation of the two-dimensional tomographic projections of the three-dimensional object related to the image through the x-ray transform.[1] The objective of chromotomographic image restoration is to recover the complete three-dimensional spatial–chromatic scene from chromatically sampled two-dimensional projections.

The restoration is obstructed by the so-called limited-angle problem, generic to many computed-tomography applications, or by the fact that, as a result of certain physical limitations of the measuring instrument, not enough data can be collected.[2–5] Effects of this can be seen in the Fourier domain by appealing to the projection-slice theorem.[6] The theorem states that the two-dimensional Fourier transform of the tomographic projection is equal to the three-dimensional Fourier transform of the image evaluated on a plane through the origin in a direction perpendicular to the projection beam. If a complete set of tomographic projections measured over a full range of angles is not acquired (as is the case in applications such as synthetic aperture radar or medical computerized tomography[7]), then an angular region of a Fourier transform of the image is not sampled.

A similar manifestation of the same phenomenon occurs in chromotomography. Although a full range of projections can in this case be obtained, since the projection beam rotates at a fixed *acute* angle with the chromatic axis, the projection plane sweeps out only part of the three-dimensional space, and the union of Fourier transforms of all tomographic projections contains only partial information about the three-dimensional Fourier transform of the object. This is reflected in a singularity of the system transfer function (STF) matrix, which relates the tomographic projections with the object, thus obstructing computation of the hyperspectral image by a direct method of inversion. Use of the pseudoinverse obtained through singular-value decomposition (SVD) of the STF matrix offers the minimum-norm least-squares solution; this estimate, however, suffers from poor feature resolution and contains severe artifacts.

In an alternative approach, which bypasses the computation of an inverse of a singular matrix, recovery of the unknown object is achieved by utilizing *a priori* information about the image. Typically, an initial (often arbitrary) guess about the unknown object is made, which is then subjected to a sequence of corrections, forcing it to satisfy a number of desirable characteristics. This sequence of corrections is applied repetitively until convergence occurs. An example of this approach is the Gerchberg–Papoulis procedure,[8,9] where knowledge of the image spatial boundary and part of its Fourier transform is used to recover the spectral image. The basis for this technique is that a spatially limited object has an analytic Fourier transform and therefore in principle can be uniquely determined from any finite interval of its spectrum. The Gerchberg–Papoulis procedure is an iterative algorithm, which approaches the solution by alternating between the object domain and the Fourier domain. The known parts of the spectrum and the known boundary of the image are imposed on the iterated solution as constraints.

The Gerchberg–Papoulis technique is generalized by the method of projections onto convex sets (POCS), which allows other types of *a priori* knowledge about the object to be incorporated into the algorithm. As a result, an increased convergence rate and an improved performance are achieved, at only a modest increase in the computa-

tional cost for implementing the additional constraint operators. As with the Gerchberg–Papoulis technique, the discrete Fourier transform still remains the most expensive part of the iteration. This advantageous performance/complexity trade-off is of particular importance in image processing applications, where data dimensionality prohibits computationally complex approaches.

The main challenge of the POCS method is to identify constraint operators that can be easily implemented and that lead to a rapid convergence. This is particularly important in the processing of high-dimensionality images, such as hyperspectral data cubes. Traditional constraint operators involving nonnegativity, magnitude bounds, or finite support yield slow convergence and unsatisfactory performance, since hyperspectral image estimates do not differ significantly from the image characteristics demanded by the constraints.[10]

In this work we propose a new object domain constraint. This constraint is based on the SVD technique, which is used to determine the dominant structure of the data and to construct reduced-dimensionality approximations by projecting the data onto a subspace that is consistent with image spectral characteristics. Results of experiments demonstrate that the new constraint yields rapid convergence and leads to restoration of a significant portion of the missing information. The algorithm also provides a reasonable way of monitoring the progress of the iteration by means of the singular-value spectrum. The costly two-dimensional Fourier transform is transferred to the precomputational stage, and the computation required by the iteration is reduced to matrix–vector multiplication. Since the new algorithm relies on data redundancy, a characteristic of many applications, it is anticipated that the SVD–POCS algorithm can be effectively applied to other data restoration problems, such as ectomography.[11]

Although SVD has been known for approximately 100 years, recently there has been a revival of interest in this technique in the signal processing community.[12–14] SVD has been used in such diverse applications as data dimensionality reduction[15]; separation of signals of interest, or signal from noise in medical applications[16]; in radar,[17] in communications[18]; and in classification of end members in hyperspectral data[19–21] (to give just a few examples). This renewed interest is largely due to the rapid advance in computing technology, amplified by the increasing sophistication of numerical algorithms, which has made calculation of the SVD of medium-size matrices practical. Since application of SVD to signal processing problems is still in its infancy, many exciting new results are to be expected in the future. The new iterative technique proposed in this paper relies on SVD, using its various aspects at three different stages of the algorithm: computation of an initial estimate of the unknown image (robust matrix inversion), computation of the object domain constraint (subspace identification), and computation of the transform domain constraint (compaction of information).

The paper is organized as follows: we present an algebraic formulation of the reconstruction problem (Section 2), review the POCS method (Section 3), introduce the new SVD–POCS algorithm (Section 4), and report results of experiments (Section 5).

2. PROBLEM

An imaging spectrometer reconstructs a three-dimensional spatial–chromatic scene from a sequence of two-dimensional images. The reconstruction can be accomplished in several ways, depending on whether multiplexing of information is performed in either the spatial or the chromatic domain or jointly. Different multiplexing schemes imply different trade-offs in terms of efficiency, flexibility, and complexity of the spectrometer,[22] but the numerical formulation of the reconstruction problem remains essentially the same.

In Ref. 22 Mooney has proposed a new computed-tomography image spectrometry technique (chromotomography). In his approach the multiplexing is accomplished by a rotating prism (Fig. 1). As the prism rotates, each chromatic slice of the object cube follows a circular path, with the radius of the path determined by the prism dispersion. A sequence of spatial tomographic projections $g(\bar{x}, \phi)$ is thus obtained, each tomographic projection being an integral of the three-dimensional spatial–chromatic object cube $f(\bar{x}, \lambda)$ in the chromatic variable λ:

$$g(\bar{x}, \phi) = \int_{-\infty}^{+\infty} f(\bar{x} - k(\lambda - \lambda_0)\bar{p}_\phi, \lambda)\mathrm{d}\lambda, \qquad (1)$$

where $\bar{x} = (x_1, x_2)$, $\bar{p}_\phi = (\cos\phi, \sin\phi)$, $0 \leq \phi < 2\pi$, λ_0 is the center wavelength, and k is a spectrometer constant determined by the sensor focal length and prism dispersion (Fig. 2).[22] This can be recognized as a three-dimensional x-ray transform of f (Fig. 3). Taking the two-dimensional Fourier transform of Eq. (1) in the spatial variable $\bar{x}$, we have

$$\mathbf{g}(\bar{\xi}, \phi) = \int_{-\infty}^{+\infty} \exp[-2\pi i(k\bar{p}_\phi \cdot \bar{\xi})(\lambda - \lambda_0)]\mathbf{f}(\bar{\xi}, \lambda)\mathrm{d}\lambda, \qquad (2)$$

where $\mathbf{f}(\bar{\xi}, \lambda)$ and $\mathbf{g}(\bar{\xi}, \phi)$ are the two-dimensional Fourier transforms of $f(\bar{x}, \lambda)$ and $g(\bar{x}, \phi)$, respectively, in $\bar{x}$ and $\bar{\xi} = (\xi_1, \xi_2)$ is the frequency variable. For a fixed ϕ, the right-hand side of Eq. (2) can be recognized as the one-dimensional Fourier transform of $\mathbf{f}$ in the chromatic

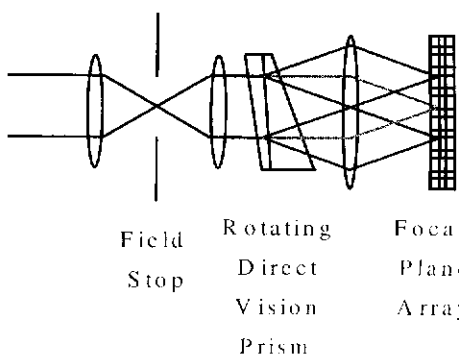

Fig. 1. Schematic representation of the chromotomographic imager. The direct vision prism is shown spreading red, green, and blue light across the focal-plane array (FPA).

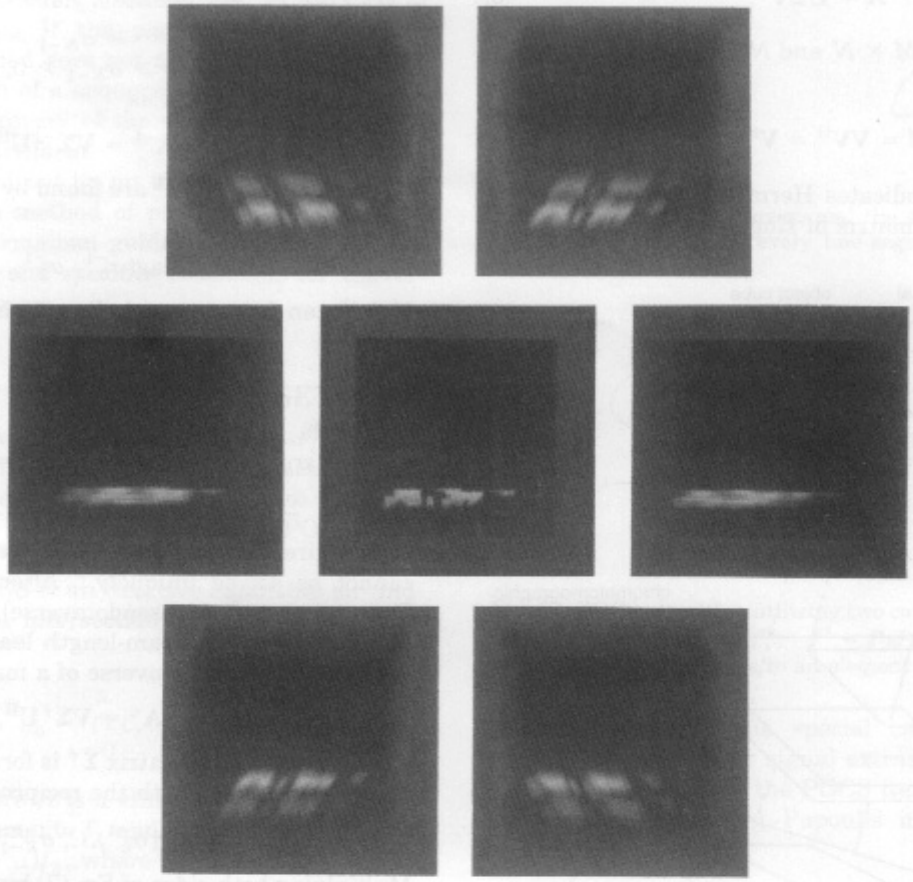

Fig. 2. Chromotomographic projections of the hyperspectral data cube of Hanscom Air Force Base, Mass., registered by FPA for prism rotation angles $\phi = 0, \pi/3, 2\pi/3, \pi, 4\pi/3$, and $5\pi/3$. The projections are ordered from middle left, clockwise. The center image corresponds to the sum of all monochromatic slices, as seen by the spectrometer.

variable λ evaluated on a plane $\{(\bar{\xi}, \zeta)\colon \zeta_\phi = \langle k\bar{p}_\phi, \bar{\xi}\rangle\}$ through the origin of the three-dimensional frequency space[4]:

$$\mathbf{g}(\bar{\xi}, \zeta_\phi) = \mathcal{F}_\lambda\{\mathbf{f}(\bar{\xi}, \lambda - \lambda_0)\}. \tag{3}$$

The union of all chromatic Fourier transforms (3) for $0 \leqslant \phi < 2\pi$ is an exterior of the cone $\{(\bar{\xi}, \zeta)\colon |\zeta| \geqslant |\langle k\bar{p}_\phi, \bar{\xi}\rangle|, 0 \leqslant \phi < 2\pi\}$ in the three-dimensional Fourier space (Fig. 4).

The objective of chromotomography is to find the object f, given the measurements g. Since part of the three-dimensional spectrum of g is unknown, there is insufficient information to recover f directly. This phenomenon is known in tomographic image restoration as the limited-angle problem. It is often interpreted in the context of the projection-slice theorem, which states that the two-dimensional Fourier transform of each line integral projection is a plane through the three-dimensional Fourier transform of the object. If the axis of rotation of the prism were perpendicular to the projection direction, the plane in the Fourier space would rotate around a line contained in the plane, sweeping the whole space. However, since the axis of rotation forms an acute angle, say α, with the projection direction, the plane rotates about a line not contained in the plane, ignoring a region of a double cone of an angle $\pi/2 - \alpha$ (Fig. 3).

Consider a version of Eq. (2), sampled at discrete chromatic bands and discrete angles[23]:

$$\mathbf{g}_m(\bar{\xi}) = \sum_{n=0}^{N-1} \exp[-2\pi i\langle \bar{p}_m, \bar{\xi}\rangle(n - n_0)]\mathbf{f}_n(\bar{\xi}), \tag{4}$$

where $\bar{p}_m = (\cos(2\pi m/M), \sin(2\pi m/M))$, $0 \leqslant m < M$, $M \geqslant N$, $n = k\lambda$, and $n_0 = k\lambda_0$; so

$$\begin{bmatrix} \mathbf{g}_0(\bar{\xi}) \\ \mathbf{g}_1(\bar{\xi}) \\ \cdot \\ \cdot \\ \cdot \\ \mathbf{g}_{M-1}(\bar{\xi}) \end{bmatrix} = \mathbf{A}(\bar{\xi}) \begin{bmatrix} \mathbf{f}_0(\bar{\xi}) \\ \mathbf{f}_1(\bar{\xi}) \\ \cdot \\ \cdot \\ \cdot \\ \mathbf{f}_{N-1}(\bar{\xi}) \end{bmatrix}, \tag{5}$$

where the $\mathbf{A}(\bar{\xi})$ is an $M \times N$ matrix with elements

$$\mathbf{A}_{m,n}(\bar{\xi}) = \exp[-2\pi i\langle \bar{p}_m, \bar{\xi}\rangle(n - n_0)]. \tag{6}$$

For brevity we write Eq. (5) as

$$\mathbf{g} = \mathbf{A}\mathbf{f}. \tag{7}$$

The existence and the uniqueness of the solution of Eq. (7) depend on the rank of $\mathbf{A}$, which is equal to the number of independent rows of $\mathbf{A}$. It is clear from Eq. (6) that $\mathbf{A}$ is ill conditioned for many values of $\bar{\xi}$. A convenient tool for evaluating the rank of a matrix is SVD. The SVD of a matrix $\mathbf{A}$ is defined as[24]

$$\mathbf{A} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^{\mathrm{H}}, \tag{8}$$

where $\mathbf{U}$ and $\mathbf{V}$ are $M \times N$ and $N \times N$ matrices, respectively, such that

$$\mathbf{U}^{\mathrm{H}}\mathbf{U} = \mathbf{V}\mathbf{V}^{\mathrm{H}} = \mathbf{V}^{\mathrm{H}}\mathbf{V} = \mathbf{I},$$

the superscript H indicates Hermitian adjoint, and $\boldsymbol{\Sigma}$ is an $N \times N$ diagonal matrix of singular values,

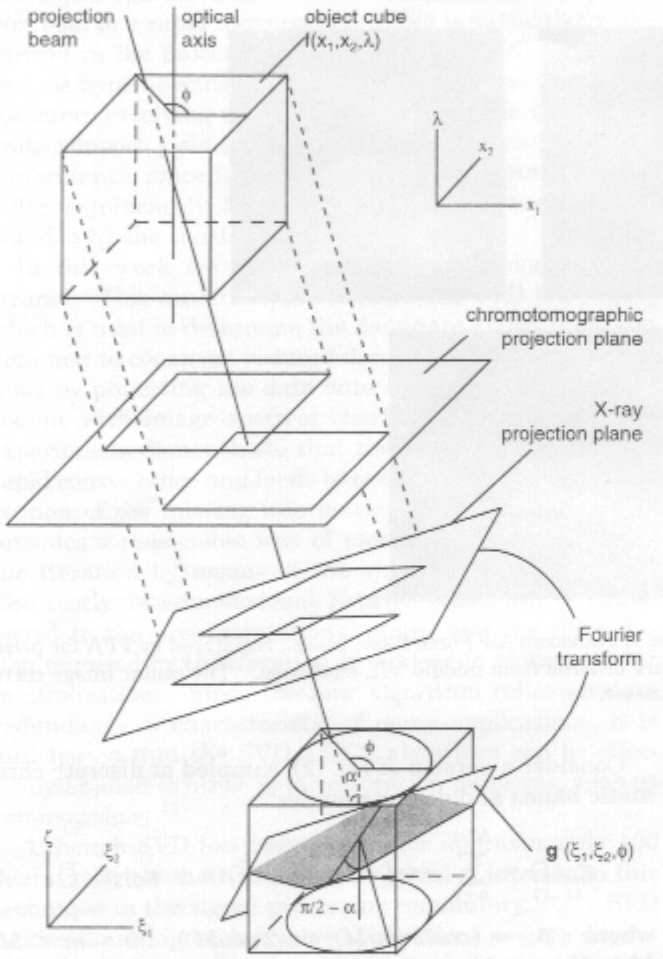

Fig. 3. Geometry of chromotomographic data collection and its relation to the x-ray transform.

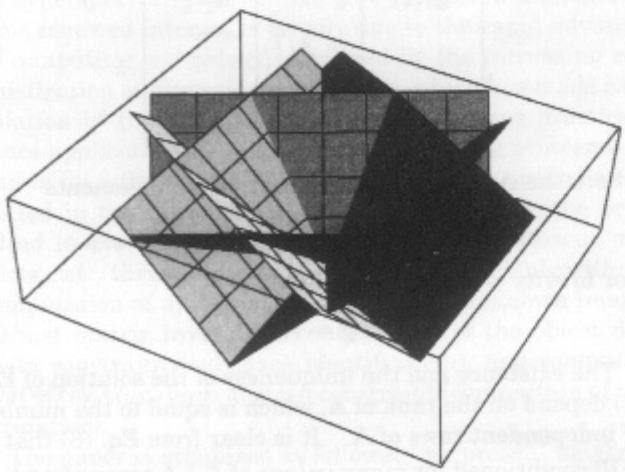

Fig. 4. Two-dimensional Fourier transform planes corresponding to the six chromotomographic projections shown in Fig. 2.

$$\boldsymbol{\Sigma} = \mathrm{diag}(\sigma_0, \sigma_1, \ldots, \sigma_{N-1}),$$

such that $\sigma_0 \geqslant \sigma_1 \geqslant \cdots \geqslant \sigma_{N-1} \geqslant 0$. If $\mathbf{A}$ is nonsingular, i.e., $\sigma_0 \geqslant \sigma_1 \geqslant \cdots \geqslant \sigma_{N-1} > 0$, then a matrix inverse to $\mathbf{A}$ can be computed as

$$\mathbf{A}^{-1} = \mathbf{V}\boldsymbol{\Sigma}^{-1}\mathbf{U}^{\mathrm{H}},$$

where elements of $\boldsymbol{\Sigma}^{-1}$ are found by inverting elements of $\boldsymbol{\Sigma}$:

$$\boldsymbol{\Sigma}^{-1} = \mathrm{diag}(\sigma_0^{-1}, \sigma_1^{-1}, \ldots, \sigma_{N-1}^{-1}).$$

If $\mathbf{A}^{-1}$ can be computed, Eq. (7) has a unique solution given by

$$\mathbf{f} = \mathbf{A}^{-1}\mathbf{g}.$$

If $\mathbf{A}$ is singular, i.e., there is $K < N$ such that $\sigma_0 \geqslant \cdots \geqslant \sigma_{K-1} > \sigma_K = \cdots = \sigma_{N-1} = 0$, so that

$$\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_K = \mathrm{diag}(\sigma_0, \ldots, \sigma_{K-1}, 0, \ldots, 0), \tag{9}$$

then a direct inverse $\mathbf{A}^{-1}$ cannot be obtained and Eq. (7) cannot be solved uniquely. Alternatively, the Moore–Penrose inverse (a pseudoinverse) $\mathbf{A}^{+}$ (Ref. 25) can be used to find a minimum-length least-squares solution of Eq. (7). The pseudoinverse of a matrix $\mathbf{A}$ is defined as

$$\mathbf{A}^{+} = \mathbf{V}\boldsymbol{\Sigma}^{+}\mathbf{U}^{\mathrm{H}}, \tag{10}$$

where the diagonal matrix $\boldsymbol{\Sigma}^{+}$ is formed by replacing nonzero elements of $\boldsymbol{\Sigma}$ with the reciprocal values

$$\boldsymbol{\Sigma}^{+} = \mathrm{diag}(\sigma_0^{-1}, \ldots, \sigma_{K-1}^{-1}, 0, \ldots, 0). \tag{11}$$

Multiplying both sides of Eq. (7) by $\mathbf{A}^{+}$ yields the pseudosolution

$$\mathbf{f}^{+} = \mathbf{A}^{+}\mathbf{g}. \tag{12}$$

In practice, the recorded data $\mathbf{g}$ are contaminated by noise, i.e.,

$$\mathbf{g} = \mathbf{A}\mathbf{f} + \mathbf{n}. \tag{13}$$

In effect, small nonzero singular values of $\mathbf{A}$ result in instabilities. This can be seen by considering

$$\mathbf{A}^{+}\mathbf{g} = \mathbf{A}^{+}\mathbf{A}\mathbf{f} + \mathbf{A}^{+}\mathbf{n} = \mathbf{V}\boldsymbol{\Sigma}^{+}(\boldsymbol{\Sigma}\mathbf{V}^{\mathrm{H}}\mathbf{f} + \mathbf{U}^{\mathrm{H}}\mathbf{n}).$$

If elements of $\boldsymbol{\Sigma}$ are close to zero, then elements of $\boldsymbol{\Sigma}^{+}$ become very large and the filtered noise dominates restoration. To balance the loss of spectral resolution and the noise amplification that is due to small singular values, a modified version of Eq. (11) can be used, where small singular values close to the noise variance are set to zero. Alternatively, a regularization technique can be applied, which allows for gradual transition of singular values to zero.[26]

The direct method of inversion, as implemented by Eq. (12), provides a reasonable estimate of the hyperspectral image; it does not, however, uniquely reconstruct the three-dimensional scene. This can lead to artifacts, particularly in scenes with significant content in the low-spatial-frequency/high-chromatic-frequency regime, which coincides with the null space of $\mathbf{A}$. To improve fidelity of the hyperspectral image, one needs to recover the null-space information. Clearly, this cannot be accomplished by using the direct method, and other approaches need to be tried. One such approach involves using *a priori* information about the scene, such as finite extent,

finite intensity range, energy bounds, etc., in the form of solution constraints. If the pseudosolution resulting from the direct method does not meet these constraints, repetitive application of a sequence of constraints to the estimate leads to recovery of the null-space information and to reduction of artifacts.

This idea is formalized by an image restoration technique known as the method of projections onto convex sets (POCS). Its formalism guides selection of eligible image property sets and specifies conditions for convergence of the iteration.

3. PROJECTIONS ONTO CONVEX SETS

The POCS method was introduced by Bregman[27] and Gubin *et al.*[28] and popularized by Youla and Webb,[29] Levi and Stark,[30,31] and Sezan and Stark,[32] who also applied it to image restoration. A detailed overview of POCS is the paper by Combettes.[33]

The method of POCS is an iterative algorithm for finding an image f' in the intersection of a given sequence of R closed convex sets:

$$\mathcal{C}_0 = \bigcap_{r=1}^{R} \mathcal{C}_r .$$

A subset $\mathcal{C}$ of $\mathcal{H}$, where $\mathcal{H}$ is a Hilbert space, is convex if for any two of its elements f_1 and f_2 it contains the element $f = \mu f_1 + (1 - \mu) f_2$, where $0 \leq \mu \leq 1$ (Fig. 5). A subset $\mathcal{C}$ of $\mathcal{H}$ is closed if the limit element of any sequence of elements in $\mathcal{C}$ is contained in $\mathcal{C}$. Associated with each closed (not necessarily convex) set $\mathcal{C}_r$ is a projection operator $P_r : \mathcal{H} \rightarrow \mathcal{C}_r$, such that

$$\|f - P_r f\| = \min \|f - h\| \qquad \text{over all } h \in \mathcal{C}_r ,$$

so the nearest element to f in $\mathcal{C}_r$ is $P_r f$. If $\mathcal{C}_r$ is convex, then $P_r f$ is unique. Given (in general, nonlinear) projection operators P_r associated with closed convex sets $\mathcal{C}_r$, a sequence of images $\{f^k\}$ is generated by the recursive relation

$$f^{k+1} = P_R P_{R-1} \cdots P_1 f^k . \tag{14}$$

The sequence $\{f_k\}$ converges to f' in $\mathcal{C}_0$ (Fig. 6); i.e., for every $f \in \mathcal{H}$,

$$\lim_{k \rightarrow \infty} \langle f^k, f \rangle = \langle f', f \rangle .$$

Relation (14) can be stated more generally:

$$f^{k+1} = T_R T_{R-1} \cdots T_1 f^k , \tag{15}$$

where $T_r = I + \lambda_r (P_r - I)$, $0 < \lambda_r < 2$, and I is the identity operator. The λ_r's are relaxation parameters and can be used to accelerate the rate of convergence of the algorithm.

Typical convex sets include sets of images restricted by spatial extent (optical field stop) ($\mathcal{C}_{SL}$), the known part of the spectrum ($\mathcal{C}_{SP}$), band limitedness ($\mathcal{C}_{BL}$), the known part of the image ($\mathcal{C}_{IP}$), nonnegativity, amplitude bound (intensity range), l^2 energy, etc. In the special case when only the two sets $\mathcal{C}_{IP}$ and $\mathcal{C}_{BL}$ are used, the POCS iteration reduces to its famous special case, the Gerchberg–Papoulis algorithm for signal extrapolation. When only $\mathcal{C}_{SL}$ and $\mathcal{C}_{SP}$ are used, the POCS iteration takes the dual form of the Gerchberg–Papoulis iteration, i.e., that of spectral extrapolation.

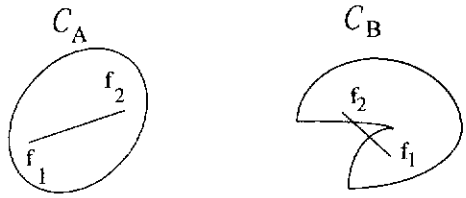

Fig. 5. The set $\mathcal{C}_A$ is convex, and the set $\mathcal{C}_B$ is not convex. A convex set must contain every line segment with end points in the set.

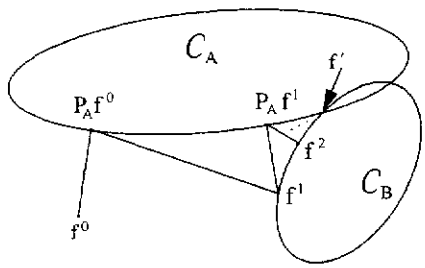

Fig. 6. POCS algorithm utilizing two convex sets $\mathcal{C}_A$ and $\mathcal{C}_B$. A sequence of images f^0, $f^1 = P_B P_A f^0$, $f^2 = P_B P_A f^1, \ldots$, $f^k = P_B P_A f^{k-1}$ converges to an element f' in $\mathcal{C}_A \cap \mathcal{C}_B$.

4. ALGORITHM

The Gerchberg–Papoulis algorithm for spectral extrapolation is an example of the POCS method utilizing two constraint sets: the set of images with a known part of its Fourier transform and the set of compactly supported images. Many other algorithms utilizing different constraint sets can be formulated by using the framework of POCS, depending on the side information available. This freedom to match an algorithm to an application is indeed one of the greatest advantages of a signal processing approach based on POCS. However, finding useful convex property sets also poses a challenge. The constraints need to describe physical properties of objects with a high degree of accuracy. At least some of these properties should not be shared by the initial estimate. The constraints have to be computationally efficient. Finally, convergence has to be reached in a small number of iterations. The role of the algorithm designer is to identify and implement signal property sets that best fulfill these requirements. Traditional property sets, such as finite spatial extent, amplitude bounds, positivity, or spectral limits, usually yield slow convergence, partly because the image properties associated with these sets often do not differ significantly from properties of the image estimate. In this paper we introduce a POCS algorithm, based on two new constraint sets, that delivers promising results.

The first constraint set is determined by the data collection system, which can be modeled by means of Eq. (7) or (13). In either case the SVD of **A** naturally leads to partition of the image into two components:

$$\mathbf{f} = \mathbf{f}^+ + \mathbf{f}_N. \tag{16}$$

The pseudoinverse $\mathbf{f}^+$ is the known component of the image to be recovered, corresponding to the nonzero singular values of $\mathbf{A}$. $\mathbf{f}_N$ is the unknown image component, corresponding to the null space of $\mathbf{A}$. We will use decomposition (16) to form a constraint set of images with a known part equal to $\mathbf{f}^+$. Since, as will be seen in Subsection 4.A, the decomposition takes place in a space spanned by right singular vectors of the STF $\mathbf{A}$, we will call this constraint a transform domain constraint.

The second, object domain constraint set is determined directly by the hyperspectral data. It is well known that hyperspectral images are highly redundant in both the spatial and chromatic variables. The redundant information can be compacted by applying SVD to the data organized in a matrix form. The compacted information can be extracted from the data matrix and applied as an estimate for the unknown image component $\mathbf{f}_N$. Since SVD is the optimal transformation for information redundancy reduction, the constraint leads to an efficient and rapidly converging algorithm.

The following subsections describe both constraint sets in detail.

A. Transform Domain Constraint

As it was seen in Section 2, singularity of the STF matrix $\mathbf{A}$ leads to parameterization of the solution space by the null space of $\mathbf{A}$. Use of the pseudoinverse $\mathbf{A}^+$ yields a unique solution by discarding the null-space component. Since the pseudoinverse identifies the known data, it can be used to form a constraint set. This constraint set is similar to the set of images with a known spectral part, C_{SP}, utilized in the Gerchberg–Papoulis iteration. Indeed, we can write an iteration of the form

$$\mathbf{f}^{k+1} = P_{\mathbf{A}}\mathbf{f} + \bar{P}_{\mathbf{A}}P^f\mathbf{f}^k, \tag{17}$$

where $P_{\mathbf{A}} = \mathbf{A}^+\mathbf{A}$ and $\bar{P}_{\mathbf{A}} = I - \mathbf{A}^+\mathbf{A}$ are the range and null-space projection operators of $\mathbf{A}$, respectively, and P^f is a projection operator associated with the object domain constraint. In short, Eq. (17) can be written as

$$\mathbf{f}^{k+1} = \mathbf{f}^+ + \bar{P}_{\mathbf{A}}P^f\mathbf{f}^k \tag{18}$$

$$= P^{\mathbf{A}}P^f\mathbf{f}^k, \tag{19}$$

where $P^{\mathbf{A}}(\mathbf{f}') = \mathbf{f}^+ + \bar{P}_{\mathbf{A}}(\mathbf{f}')$ is a projection of $\mathbf{f}' = P^f\mathbf{f}^k$ onto $C_{\mathbf{A}}$, a set of signals with a known component $\mathbf{f}^+ = P_{\mathbf{A}}\mathbf{f}$. The known part $\mathbf{f}^+$ is used in the iteration as the initial estimate $\mathbf{f}^0$. The unknown component $\mathbf{f}^k$ is iteratively refined by alternatively applying projectors P^f and $\bar{P}_{\mathbf{A}}$. Since

$$\mathbf{A}^+\mathbf{A} = \mathbf{V}\Sigma_K^+\Sigma_K\mathbf{V}^{\mathrm{H}} = \mathbf{V}I_K\mathbf{V}^{\mathrm{H}}, \tag{20}$$

where Σ_K is defined by Eq. (9) and I_K is a rank K identify matrix, the iteration (17) can be written as

$$\mathbf{f}^{k+1} = \mathbf{V}I_K\mathbf{V}^{\mathrm{H}}\mathbf{f} + (I - \mathbf{V}I_K\mathbf{V}^{\mathrm{H}})P^f\mathbf{f}^k. \tag{21}$$

Equation (21) allows us to interpret the constraint projector $P^{\mathbf{A}}$ as a linear filtering operator in a transform space defined by the $\mathbf{V}$ matrix. To see that, premultiply Eq. (7) by $\Sigma^+\mathbf{U}^{\mathrm{H}}$, so that

$$\Sigma^+\mathbf{U}^{\mathrm{H}}\mathbf{g} = I_K\mathbf{V}^{\mathrm{H}}\mathbf{f}, \tag{22}$$

and set $\mathbf{y} = \Sigma^+\mathbf{U}^{\mathrm{H}}\mathbf{g}$ and $\mathbf{x} = \mathbf{V}^{\mathrm{H}}\mathbf{f}$. Equation (7) can now be expressed as a simple filtering operation:

$$\mathbf{y} = I_K\mathbf{x}. \tag{23}$$

The unitary matrix $\mathbf{V}^{\mathrm{H}}$ forms a set of orthogonal vectors, similar to the sinusoidal Fourier basis functions. Multiplication of $\mathbf{f}$ by $\mathbf{V}^{\mathrm{H}}$ unitarily transforms the vector $\mathbf{f}$ into a new vector $\mathbf{x} = \mathbf{V}^{\mathrm{H}}\mathbf{f}$. Premultiplication of $\mathbf{x}$ by a rank K identity matrix I_K in Eq. (23) can therefore be viewed as projecting $\mathbf{f}$ onto a K-dimensional space in the $\mathbf{V}$-transform domain, and the pseudoinverse

$$\mathbf{f}^+ = \mathbf{V}\mathbf{y} = \mathbf{V}I_K\mathbf{V}^{\mathrm{H}}\mathbf{f} \tag{24}$$

can be viewed as a $\mathbf{V}$-domain low-pass filtered version of $\mathbf{f}$.

To simplify the discussion in this section, we have considered the noise-free data collection model (7). The more realistic data+noise model (13) does not affect the derivation or the interpretation of the $\mathbf{V}$-transform constraint in a significant way. Computation of the pseudoinverse, as mentioned in Section 2, has to be modified by setting small singular values to zero to filter out signal components excessively contaminated by noise. Transition from the noise-free model (7) to the data+noise model (13) and modification of the pseudoinverse have two important consequences. First, discarding small singular values decreases the rank of $\mathbf{A}$. Second, the transition from model (7) to model (13) implies that the pseudosolution $\mathbf{f}^+$, depending on the choice of K, will be affected by noise. The iterative algorithm will not attempt to reduce degradation of $\mathbf{f}^+$, once K is selected.

B. Object Domain Constraint

It is possible in principle for every spatially resolved pixel to possess a unique chromatic signature. In practice, however, large regions of hyperspectral scenes, such as mixed forests or crop fields, have similar chromatic profiles. Figure 7, which shows a three-dimensional rendering of a hyperspectral data cube of Moffett Field, Calif., measured with the airborne visible/infrared imaging spectrometer instrument, demonstrates this case. The

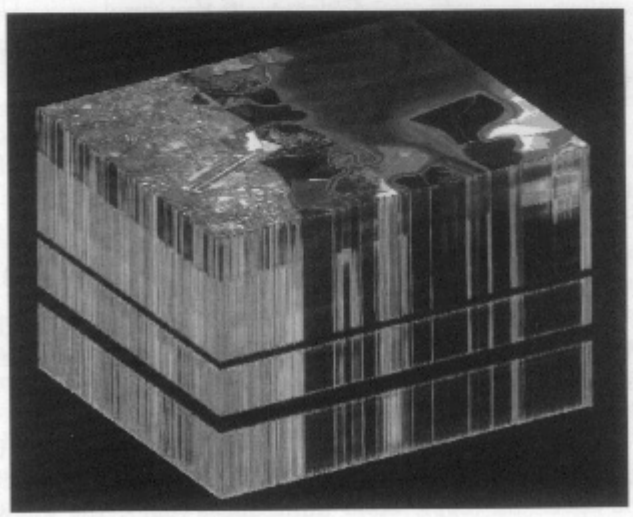

Fig. 7. Hyperspectral image of Moffett Field, Calif., taken with the AVIRIS instrument (courtesy of the Jet Propulsion Laboratory, Pasadena, Calif.).

strong axial alignment of the image features in the vertical faces of the data cube indicates a high degree of correlation (or information redundancy) among the chromatic slices. The existence of a significant degree of correlation in hyperspectral images has been form erly recognized and used to identify chromatic end members, perform chromatic unmixing, and compress information.[19,21,34] Here we propose to take advantage of redundancy to construct the object domain constraint, which is the cornerstone of our iterative algorithm for restoration of chromotomographic data.

Since data are redundant, part of the Fourier transform can uniquely represent the hyperspectral image. In chromotomography the conical shape of the missing data region implies that some information exists for all spatial frequencies and that some information exists for all chromatic frequencies. Using the redundancy present in known parts of all horizontal planes (spatial information) and vertical lines (chromatic information) throughout the three-dimensional Fourier cube, we can form an estimate of the unknown image part by forcing the missing data values to be consistent with the known region of the image.

The redundancy of a hyperspectral image can be assessed by computing the SVD of the image organized as a two-dimensional spatial–chromatic matrix. Consider a zero-mean data matrix

$$F = (f_0 - \bar{f}_0, f_1 - \bar{f}_1, \ldots, f_{N-1} - \bar{f}_{N-1})^{\mathrm{T}}, \tag{25}$$

formed by taking as its rows

$$f_n = f_n(x' = x_1 + X_1 x_2),$$
$$0 \le x_1 < X_1, \qquad 0 \le x_2 < X_2, \tag{26}$$

the lexicographically ordered monochromatic slices f_n of the object cube estimate, where X_1 and X_2 are the hyperspectral image spatial dimensions and $\bar{f}_0$, $\bar{f}_1$, and $\bar{f}_{N-1}$ are the mean values of vectors f_0, f_1, and f_{N-1}. Application of SVD to F, i.e.,

$$F = U\Sigma V^{\mathrm{T}} = \sum_{n=0}^{N-1} \sigma_n \bar{u}_n \bar{v}_n^{\mathrm{T}}, \tag{27}$$

produces a new set of triples of singular values σ_n, spatial right singular vectors $\bar{v}_n$ (eigenimages), and chromatic left singular vectors $\bar{u}_n$ (eigenchroma), forming the weighted outer product sum of Eq. (27) (to distinguish the SVD of the data matrix F from the SVD of the system STF matrix $\mathbf{A}$, we use italic letters U, Σ, and V for the former and boldface letters $\mathbf{U}$, $\mathbf{\Sigma}$, and $\mathbf{V}$ for the latter). The eigenimages and the eigenchroma are ordered in terms of decreasing singular values or, equivalently, decreasing information.

Usually, a few singular values of the SVD dominate the singular-value spectrum (Fig. 8). These singular values correspond to outer products having the richest information content. Since the outer products associated with the lower-order singular values represent noise and artifacts, an image can be fully represented by the dominant few outer products: Specifying *a priori* the number of outer products that represent a hyperspectral image is similar in spirit to constraining a two-dimensional image spatially, as in the Gerchberg–Papoulis procedure. The main factor that differentiates the SVD–POCS algorithm object domain constraint from the Gerchberg–Papoulis algorithm finite-spatial-extent constraint is that, in the former, information is constrained in an indirect way by selection of the image feature subspace spanned by the dominant singular vectors of the data matrix, whereas in the latter the constraint is realized by selection of the image valid spatial indexes. An approach based on the former constraint is both efficient and intuitively pleasing, since selection of the outer products has an immediate and profound effect on the data and affects the entire information content of the spatial–chromatic image rather than a single aspect of it, spatial boundaries, whose importance is often uncertain.

Define projector $P^f = UI_L U^{\mathrm{T}}$, where U is the eigenchroma matrix of the pseudoinverse $F = F^0$ [as in Eq. (27)], I_L is a rank L identity matrix, and L is the constrained number of eigenimage–eigenchroma pairs of F. The image in Eq. (27) can then be decomposed as

$$F = F_L + F_L', \tag{28}$$

where

$$F_L = P^f F. \tag{29}$$

F_L is a projection of F onto the feature space U_L, represents the compacted image information, and is used as an estimate for the missing part of the spatial–chromatic data. $F_L' = F - F_L$ is the orthogonal complement of F_L, which is a projection of F onto the noise space $U - U_L$, represents noise and artifacts, and is discarded.

Optimal selection of the feature space dimension L is of great importance, since it bears on fidelity of the restoration and convergence speed of the algorithm. The problem of estimating an optimal model dimension is encountered in many applications, particularly in detection and communications. One of the earliest attempts to solve it can be found in Shannon,[35] who tried to establish a mathematical connection between the complexity of the source and its distortion. In principle, the problem can be reduced to estimating a low-rank approximation to the data matrix corrupted by noise. If noise variance is known, the feature space dimension that optimally trades model distortion for noise reduction can be expressed in terms of signal-to-noise ratio.[36] In practice, however, noise variance can be difficult to obtain. Moreover, in data resto-

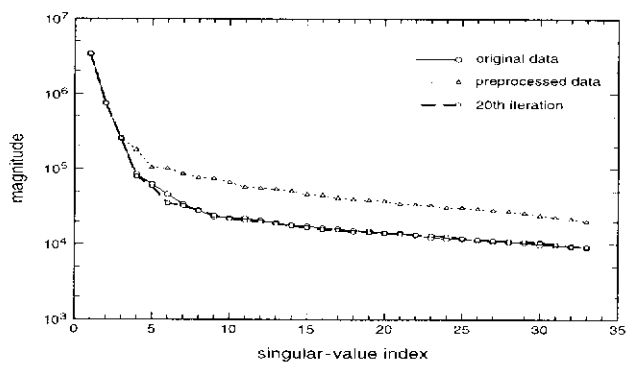

Fig. 8. Singular-value spectrum of Jasper Ridge, Calif.

ration, the model includes artifacts, as well as signal and noise. Ideally, the SVD would separate signal and artifacts (or signal and interference). However, this is rarely the case; singular vectors usually contain information related to all three components of the model.

A more favorable situation occurs when signal-to-artifact ratio (and signal-to-noise ratio) is sufficiently high; in effect, the dominant singular vectors contain mostly signal information. One intuitive approach is then to identify an abrupt change in the singular-value spectrum and use it as a demarcation point for subspace decomposition. A still better strategy for an iterative algorithm is to start with a more economical approximation of the feature space and then gradually increase its dimension as iteration progresses and signal-to-artifact ratio increases. Such a strategy allows one to build into the restoration algorithm a hierarchical structure, where data restoration is applied to one singular vector at a time, progressing downward from more relevant to less relevant singular vectors. The disadvantage of this approach is that with each relaxation of the model dimension L, the feature space U_L (onto which the data are projected) changes, which makes predicting convergence properties of the algorithm difficult.

We plan to investigate this approach further, establish theoretical convergence conditions for the algorithm, and derive a quantitative rule for optimal subspace selection. Initial investigations have shown that this strategy can be fruitful. However, our present results indicate that even suboptimal heuristic selection of a fixed feature space dimension L can result in a satisfactory performance/convergence trade-off for the SVD–POCS algorithm.

C. Algorithm

The iterative algorithm is of the form

$$\mathbf{f}^{k+1} = P^{\mathbf{A}}P^f\mathbf{f}^k = \mathbf{f}^0 + \bar{P}_{\mathbf{A}}P^f\mathbf{f}^k. \tag{30}$$

The object domain projector P^f, which reflects properties of the hyperspectral data,[23] and the V-transform domain projector $P^{\mathbf{A}}$, which reflects properties of the data collection system, are applied alternately to the estimated image. The projector P^f limits the spatial–chromatic extent of the image by truncating its SVD expansion. The projector $P^{\mathbf{A}}$ constrains the image to have a spectral segment $\mathbf{f}^0$ equal to the pseudosolution $\mathbf{f}^+$. Convergence of the iteration is ensured by convexity of the constraint sets. The stop criterion for the iteration is given by a minimum allowed increment in singular values of the iterated matrix F^k, e.g., 0.1% of the magnitude of the largest singular value.

The SVD–POCS algorithm requires computation of the constraint operators P^f and $P^{\mathbf{A}}$ in the preparation stage and computation of the product $P^{\mathbf{A}}P^f\mathbf{f}^k$ in the iteration stage.

The constraint operator P^f can be found by application of SVD directly to F, but since the data matrix is $p \times N$, where $p \gg N$, this method is highly inefficient. Alternatively, the constraint operator P^f can be found either by solving the $N \times N$ eigenvalue problem of the covariance matrix

$$R_{FF} = FF^{\mathrm{T}} = U\Lambda U^{\mathrm{T}},$$

where $\Lambda = \Sigma^2$ is a diagonal matrix of eigenvalues and U is the eigenvector matrix of FF^{T}, or, equivalently, by solving the eigenvalue problem of the covariance matrix $R_{\mathbf{FF}}$, where $\mathbf{F}$ is a two-dimensional matrix formed from $\mathbf{f}$ (the two-dimensional Fourier transform of f) by organization analogous to Eqs. (25) and (26). Since the eigenchroma matrix U is identical for both F and $\mathbf{F}$, the projection operator P^f can be applied directly to $\mathbf{F}$, thereby allowing the image to stay in the spatial Fourier space for the entire iteration. The observation that $UI_LU^{\mathrm{T}} = U_LU_L^{\mathrm{T}}$ further facilitates the computation by reducing the $N \times N$ matrix–vector multiplication to a $2L$ set of N-point inner product computations, where typically $L \ll N$. Since for $L < N/2$ the computation of $U_L(U_L^{\mathrm{T}}\mathbf{F})$ is preferable to the computation of $P^f\mathbf{F}$, U_L rather than P^f is stored in the preparation stage.

The constraint operator $P^{\mathbf{A}}$ requires computation of the known spectral part $\mathbf{f}^0$, as well as the operator $\bar{P}_{\mathbf{A}}$, which is used to project the estimated spectral image onto the null space of $\mathbf{A}$. Computing $\mathbf{f}^0$ requires the SVD of $\mathbf{A}$, which is obtained off line, and the multiplication of $\mathbf{g}$ by the pseudoinverse of $\mathbf{A}$. To find the projection matrix $\bar{P}_{\mathbf{A}}$, one can use the identity $I_N - \mathbf{V}I_K\mathbf{V}^H = \mathbf{V}_{N-K}\mathbf{V}_{N-K}^H$, where $\mathbf{V}_{N-K}$ is formed by replacing the first K columns of $\mathbf{V}$ with zero column vectors. Since $N - K \approx N/4 < N/2$ (on average across all spatial frequencies), computation of $\mathbf{V}_{N-K}(\mathbf{V}_{N-K}^H\mathbf{f})$ is again preferable to computation of $\bar{P}_{\mathbf{A}}\mathbf{f}$, and $\mathbf{V}_{N-K}$ rather than $\bar{P}_{\mathbf{A}}$ is stored in the preparation stage.

The algorithm proceeds as follows.

Precomputation:

- compute SVD of $\mathbf{A}$.

Computation of projectors:

- compute $\mathbf{f}^0 = \mathbf{A}^+\mathbf{g}$,
- store $\mathbf{V}_{N-K}$,

} P^A

- compute the covariance matrix $R_{\mathbf{F}^0\mathbf{F}^0}$,
- find U and Σ of F through eigendecomposition of $R_{\mathbf{F}^0\mathbf{F}^0}$.

} P^f

Iteration:

- obtain $\mathbf{f}_1^k = P^f\mathbf{f}^k$ by computing $U_L(U_L^{\mathrm{T}}\mathbf{F}^k)$,
- obtain $\mathbf{f}_2^k = \bar{P}_{\mathbf{A}}\mathbf{f}_1^k$ by computing $\mathbf{V}_{N-K}(\mathbf{V}_{N-K}^H\mathbf{f}_1^k)$,
- obtain $\mathbf{f}^{k+1} = \mathbf{f}^0 + \mathbf{f}_2^k$.

The multiplicative complexity of the iteration is

$$2pN\bar{K} + 2pNL, \tag{31}$$

where $p = X_1X_2$ is the number of pixels, N is the number of resolved wavelengths, $\bar{K}$ is the number of zero singular values of $\mathbf{A}(\bar{\xi})$ averaged across all spatial frequencies $\bar{\xi}$, and L is the dimension of the rank-reduced matrix F_L. The algorithm requires storage of the $p \times N \times \bar{K}$ matrix $\mathbf{V}_{N-K}$, the small $L \times N$ matrix U_L, and the $p \times N$ vector $\mathbf{f}^0$. The first term of expression (31) accounts for computing $\bar{P}_{\mathbf{A}}\mathbf{f}_1^k$, and the second accounts for computing $P^f\mathbf{f}^k$. The main computational and storage burden is as-

sociated with the projector P_A. For $p = 240^2$ and $N = 64$, an average execution time for a single iteration cycle on a dual Pentium Pro processor is 10 s.

5. EXPERIMENTS

We have tested the algorithm on several synthesized and real datasets. Real datasets were either chromotomographic sequences of scenes acquired with a hyperspectral camera built in our laboratory and passed through both stages of the algorithm (inverse+iterations), or sample AVIRIS data, preprocessed to mimic the loss of information occurring in the inversion stage and subjected only to the iteration stage of the algorithm.

To evaluate the performance of the algorithm on complex data, we used a sample AVIRIS image of Jasper Ridge, Calif. We selected a 240×240 segment from the original 512×614 image and chose 64 consecutive frames (starting with the 44th frame) from the sequence of 224 to fit our processing environment. We preprocessed the initial set of images, subtracting the low-spatial-frequency information by multiplying the set by

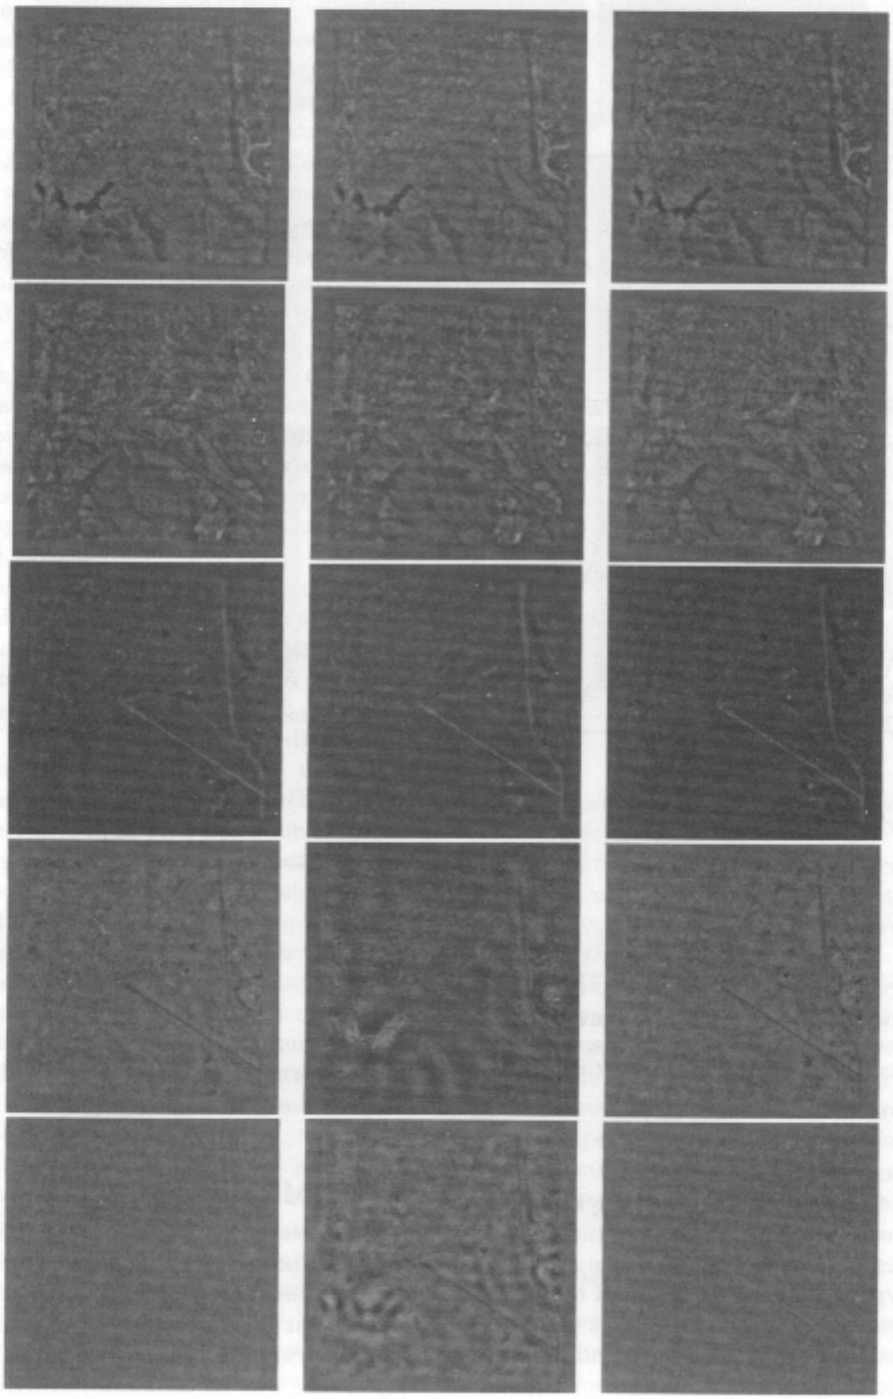

Fig. 9. First five eigenimages of Jasper Ridge: the original AVIRIS sequence (first column), the pseudoinverse reconstruction (second column), and the 20th iteration (third column).

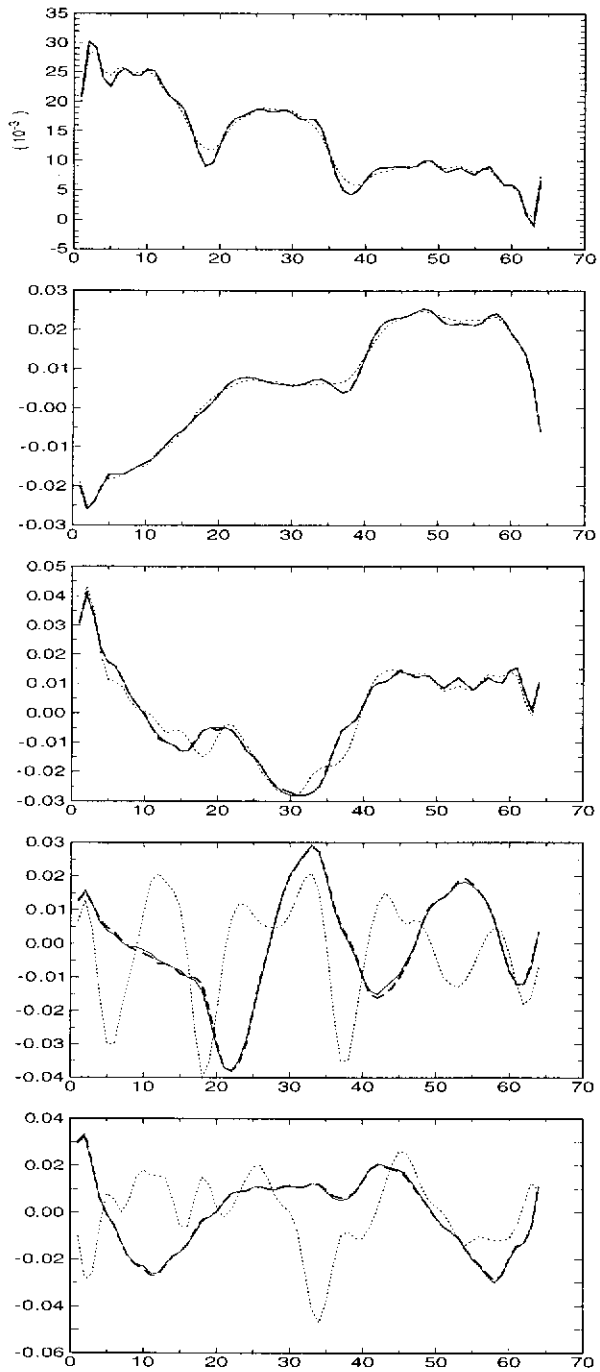

Fig. 10. First five chromatic singular vectors of Jasper Ridge: the original AVIRIS sequence (solid curves), the pseudoinverse reconstruction (dotted curves), and the 20th iteration (dashed curves).

$\mathbf{A}^H\mathbf{A}$. We then applied the iterative algorithm, using the first five eigenimages (representing most of the image energy) to form a reduced-rank data matrix estimate.

Figure 9 illustrates the progression of the first five eigenimages of the hyperspectral data through the different stages of the algorithm: the original set (first column), the preprocessed set (second column), and the 20th iterate (third column). Figure 10 shows the corresponding spectral singular vectors. The artifacts evident in the fourth and fifth eigenimages of the preprocessed data (and to a lesser degree in the third eigenimage) largely disappear in the corresponding eigenimages of the 20th iterate. Similarly, the singular-value spectrum of the iterate approaches that of the original image (Fig. 8). The rms error of the 20th iterate of the first, second, and third eigenimages yields twofold, fourfold, and fivefold decreases, respectively, as compared with the rms error of the preprocessed data. Even higher rms error reduction rates occur for the lower-order eigenimages.

The second test sequence was obtained with an infrared camera built in our laboratory.[4] The $f/4$ InSb camera utilizes a 256×256 focal-plane array (FPA), and operates in the 3–5-μm band at 60 frames per second with a 2-ms integration time. The camera collects one frame of data for each of the 80 prism orientations, uniformly spaced over 2π. We imaged a target of opportunity (a building) from the laboratory window. The data were obtained on January 12, 1998, at 9.30 a.m. To minimize the effects of FPA nonuniformities and stray reflections from the rotating prism, an additional sequence of data was recorded with the entrance pupil blocked. The blocked sequence was used to perform a one-point nonuniformity correction on a frame-by-frame basis.

Figure 2 shows examples of the dispersed images for six of the 80 orientations. The effect of atmospheric absorption by CO_2 at 4.27 μm is visible as a horizontal stripe in the vertically dispersed images. In the reconstruction (not shown), the solar radiation dominates the 3-μm side of the band, and the scene thermal emission dominates the 5-μm side. The atmospheric absorption by CO_2 in the 4.2–4.5-μm band results in dramatic reduction of contrast in the spatial image and in strong attenuation of the 4.2–4.5-μm band in the chromatic profiles. The chromatic end members of the scene are the thermal spectrum of the background, the thermal spectrum of the hot segments of the scene, the solar spectrum reflected from the bright parts of the building, and the spectrum of material differences of the brick and painted metal portions of the building.

Figure 11 illustrates the top six eigenimages of the scene. The upper half of the figure shows eigenimages of the pseudoinverted data. All six eigenimages contain reconstruction artifacts in the form of high-contrast halos around the building. These halos are characteristic of the high-pass filtering and the noise amplification introduced by the pseudoinverse. The lower half of Fig. 11 illustrates the effect of the iterative algorithm on the artifacts; the halos have been eliminated. Incidentally, the eigendecomposition demonstrates the importance of FPA nonuniformities, which appear in the top-right corner of the third eigenimage in the lower half.

6. SUMMARY

The objective of this work was improving the quality of the chromotomographic image restoration. The objective was accomplished by introducing a novel object domain constraint based on the inherent redundancy of the hyperspectral data and on the information-compacting ability of singular-value decomposition. As a result, a highly efficient and effective image restoration algorithm for processing chromotomographic images was obtained. It was

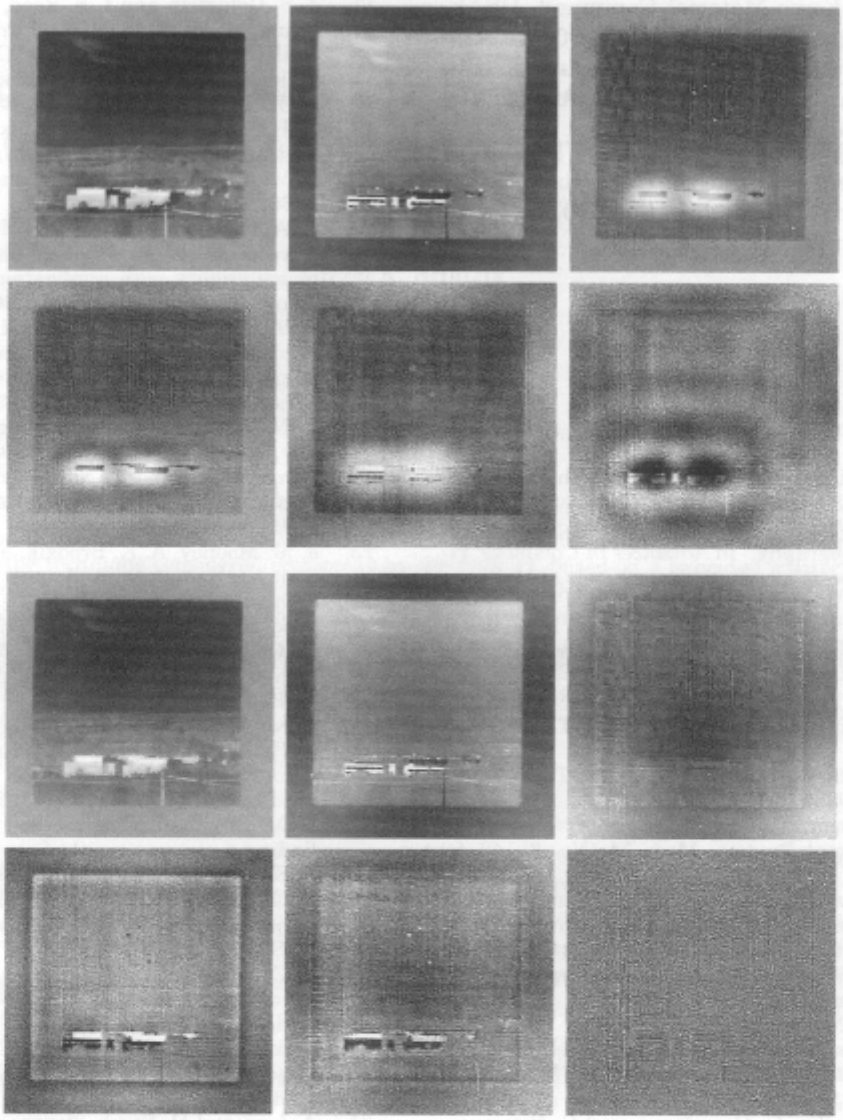

Fig. 11. First six eigenimages of Hanscom: the pseudoinverse (upper half) and the 20th iteration (lower half).

demonstrated that the algorithm is able to suppress artifacts and noise characteristic of the pseudosolution and to improve resolution of the distinct features present in the scene. The algorithm allows one to monitor the progress of iteration in both a quantitative (singular values) and a qualitative (eigenimages) manner. The computational cost of the iteration was highly reduced by transferring the two-dimensional discrete Fourier transform calculations to the precomputational stage. All three aspects—performance, tractability, and efficiency—indicate that the new constraints are far superior to the standard constraints in processing of chromotomographic data.

Most of the gain in the performance of the algorithm was achieved within the first ten iterations. Although, at the present stage of development, processing requires 2–3 min of CPU time, a tenfold processing time improvement can be expected through algorithm tuning alone. Further CPU time reduction can be accomplished by implementation of the algorithm in a parallel processing environment, which should allow near-real-time processing. An additional benefit of the algorithm, intrinsic to its iterative nature, is that some estimate of the object is available at *all* times, successive iterations providing only higher-fidelity updates.

The main issues to be investigated in the future are the design of an easily computable criterion for determining the degree of hyperspectral model reduction and the validation of the assumption that high-order eigenchroma are shared by different spatial frequency regions of the hyperspectral image. The reliability of the model reduction criterion and eigenchroma error estimates are going to impact the performance of the algorithm. Since both the scene chromatic content and the relative importance of various spectral components can be highly dependent on the application, characterization of a wide range of hyperspectral scenes needs to be performed.

ACKNOWLEDGMENTS

The authors thank Jerry Silverman, Virgil Vickers, and Bill Ewing for critical reading of the manuscript, Myoung An and Virgil Vickers for software support, and Linda Bouthillette for the artwork. This effort was sponsored by the U.S. Air Force Office of Scientific Research and

Rome Laboratory, Air Force Material Command, USAF, under cooperative agreement F30602-95-2-0034. The U.S. government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon.

Send all correspondence to Andrzej K. Brodzik, Scientific Software, Woburn, Massachusetts 01801; e-mail, brodzik@max.rl.plh.af.mil.

REFERENCES

1. D. C. Solmon, "The x-ray transform," J. Math. Anal. Appl. **56**, 61–83 (1976).
2. H. H. Barrett, "Limited-angle tomography for the nineties," J. Nucl. Med. **31**, 1689–1692 (1990).
3. D. A. Hayner and W. K. Jenkins, "The missing cone problem in computer tomography," in *Advances in Computer Vision and Image Processing*, T. S. Huang, ed. (JAI Press, London, 1984), Vol. 1, pp. 83–144.
4. J. M. Mooney, V. E. Vickers, M. An, and A. K. Brodzik, "A high-throughput hyperspectral infrared camera," J. Opt. Soc. Am. A **14**, 2951–2961 (1997).
5. K. C. Tam and V. Perez-Mendez, "Tomographic imaging with limited-angle input," J. Opt. Soc. Am. **71**, 582–592 (1981).
6. F. Natterer, *The Mathematics of Computerized Tomography* (Wiley, Stuttgart, 1986).
7. H. Hiriyannaiai, "X-ray computed tomography for medical imaging," IEEE Signal Process. Mag. **14** (No. 2), 42–59 (1997).
8. A. Papoulis, "A new algorithm in spectral analysis and band-limited extrapolation," IEEE Trans. Circuits Syst. **CAS-22**, 735–742 (1975).
9. R. W. Gerchberg, "Super-resolution through error energy reduction," Opt. Acta **21**, 709–720 (1974).
10. A. K. Brodzik, J. M. Mooney, and M. An, "Image restoration by convex projections: application to image spectrometry," in *Imaging Spectrometry*, M. R. Descour and J. M. Mooney, eds., Proc. SPIE **2819**, 231–242 (1996).
11. H. Knutsson, P. Edholm, G. Grandlund, and C. Petersson, "Ectomography—a new radiographic reconstruction method—I. Theory and error estimates," IEEE Trans. Biomed. Eng. **BME-27**, 640–648 (1980).
12. F. Deprettere, ed., *SVD and Signal Processing: Algorithms, Applications and Architectures* (Elsevier, Amsterdam, 1988).
13. R. Vaccaro, ed., *SVD and Signal Processing II: Algorithms, Analysis and Applications* (Elsevier, Amsterdam, 1991).
14. H. C. Andrews and C. L. Patterson, "Singular value decomposition and digital image processing," IEEE Trans. Acoust., Speech, Signal Process. **ASSP-24**, 26–53 (1976).
15. K. Konstantinides, B. Natarajan, and G. S. Yovanof, "Noise estimation and filtering using block-based singular value decomposition," IEEE Trans. Image Process. **6** (No. 3), 479–483 (1997).
16. P. K. Sadasivan and D. N. Dutt, "SVD based technique for noise reduction in electroencephalographic signals," Signal Process. **55**, 179–189 (1996).
17. J. S. Goldstein and I. S. Reed, "Reduced-rank adaptive filtering," IEEE Trans. Signal Process. **45**, 492–496 (1997).
18. S. Heidari and C. L. Nikias, "Co-channel interference mitigation in the time-scale domain: the CIMTS algorithm," IEEE Trans. Signal Process. **44**, 2151–2162 (1996).
19. N. H. Endsley, "Spectral unmixing algorithms based on statistical models," in *Imaging Spectrometry*, M. R. Descour, J. M. Mooney, D. L. Perry, and L. Illing, eds., Proc. SPIE **2480**, 23–36 (1995).
20. A. A. Green, M. Berman, P. Switzer, and M. D. Craig, "A transformation for ordering multispectral data in terms of image quality with implications for noise removal," IEEE Trans. Geosci. Remote Sens. **26** (No. 1), 65–74 (1988).
21. J. B. Lee, A. S. Woodyatt, and M. Berman, "Enhancement of high spectral resolution remote-sensing data by a noise-adjusted principal component transform," IEEE Trans. Geosci. Remote Sens. **28** (No. 3), 295–304 (1990).
22. J. M. Mooney, "Spectral imaging via computed tomography," in *Proceedings of the 1994 Meeting of the Infrared Information Symposia Specialty Group on Passive Sensors* (Defense Technical Information Center, Alexandria, Va., 1994), pp. 203–215.
23. J. M. Mooney, A. K. Brodzik, and M. An, "Principal component analysis in limited angle chromotomography," in *Imaging Spectrometry*, M. R. Descour and S. S. Shen, eds., Proc. SPIE **3118**, 170–178 (1997).
24. G. Golub and C. Van Loan, *Matrix Computations* (Johns Hopkins U. Press, Baltimore, Md., 1993).
25. A. Ben-Israel and T. N. E. Greville, *Generalized Inverses: Theory and Applications* (Wiley, New York, 1980).
26. A. Sano, "Optimally regularized inverse of singular value decomposition and application to signal extrapolation," Signal Process. **30**, 163–176 (1993).
27. L. M. Bregman, "The method of successive projections for finding a common point of convex sets," Dokl. Akad. Nauk SSSR **162** (No. 3), 487–490 (1965).
28. L. G. Gubin, B. T. Polyak, and E. V. Raik, "The method of projections for finding a common point of convex sets," USSR Comput. Math. Math. Phys. **7** (No. 6), 1–24 (1967).
29. D. C. Youla and H. Webb, "Image restoration by the method of convex projections: Part 1—theory," IEEE Trans. Med. Imaging **MI-1** (No. 2), 81–94 (1982).
30. A. Levi and H. Stark, "Signal restoration from phase by projections onto convex sets," J. Opt. Soc. Am. **73**, 810–822 (1983).
31. A. Levi and H. Stark, "Image restoration by the method of generalized projections with application to restoration from magnitude," J. Opt. Soc. Am. A **1**, 932–943 (1984).
32. M. I. Sezan and H. Stark, "Image restoration by the method of convex projections: Part 2—applications and numerical results," IEEE Trans. Med. Imaging **MI-1** (No. 2), 95–101 (1982).
33. P. L. Combettes, "The foundations of set theoretic estimation," Proc. IEEE **81** (No. 2), 182–208 (1993).
34. P. J. Ready and P. A. Wintz, "Information extraction, SNR improvement, and data compression in multispectral imagery," IEEE Trans. Commun. **COM-21**, 1123–1130 (1973).
35. C. E. Shannon, "Coding theorems for discrete source with a fidelity criterion," in *Institute of Radio Engineers National Convention Record* (Institute of Radio Engineers, New York, 1959), Part 4, pp. 142–163.
36. L. L. Scharf, "The SVD and reduced-rank signal processing," in *SVD and Signal Processing II: Algorithms, Analysis and Applications*, R. Vaccaro, ed. (Elsevier, Amsterdam, 1991), pp. 3–31.