EXPLOITING SPECTRAL CORRELATIONS FOR SEGMENTATION AND SHAPE DETERMINATION FROM HYPERSPECTRAL DATACUBES OF ROTATING SATELLITES

Sudhakar Prasad
Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM 87131

Qiang Zhang
Dept. of Biostatistical Sciences, Wake Forest University, Winston-Salem, NC 27109

We present simulation-based studies of a spectral correlation approach to segment and disambiguate materials on the surface of a satellite from its hyperspectral image (HSI) data. The kinematical evolution of the segment boundaries as inferred from a time series of poses of a rotating satellite contains unique signatures – and should thus enable the extraction – of the rotational parameters and local surface morphology of the object for the purposes of SOI. As a first step in the realization of this goal, we demonstrate the uniqueness of recovery of the geometrical shape, orientation, and linear dimensions of a simple 3D object from the 2D line-of-sight projections of curves on its surface when these projections are given for two poses of the object.

1 Introduction

Hyperspectral datasets typically contain far more information than needed for a typical remote sensing task of material identification, segmentation, and composition. A typical solar-illuminated man-made satellite consists of a relatively small number of geometrical primitives and materials with unique spectral reflectance signatures. This implies a highly correlated hyperspectral data cube, which may be idealized as a low-rank sum of terms, each expressible in a spatial-spectral factorized form. The high degree of spectral correlation within each spatial segment of the object admits a simple moving-averages (MA) approach for identifying the material components and their spatial boundaries, as we have recently demonstrated [1, 2]. This approach works equally well with a less idealized form of the data cube too, one that contains arbitrary, non-uniform values within the spatial support of each material component.

While the MA approach works well in the presence of low to moderate levels of noise, an approach based on spectrally correlating HSI data that have been first coarse-grained along the spectral dimension by combining them over a number of contiguous spectral channels at a time works particularly well under more noisy conditions. We shall present the rather promising results on segment boundary detection and reconstructing material spectral traces that we have obtained from applying this approach to simulated HSI data.

An accurate determination of the segment boundaries can help extract local shape and kinematical parameters of a satellite or another man-made rigid space object. We argue that the unique relationship among the two-dimensional (2D) projections of a curve fixed on a rigidly rotating three-dimensional (3D) object contains unequivocal information about the object shape and size parameters as well as the Euler angles describing the poses assumed by the object as it rotates. We can exploit this observation for a rotating 3D space object by following the kinematical evolution of its segment boundaries. If each material component of the object has a simple geometrical shape, then we can parameterize the components in terms of simple geometrical primitive shapes, such as spheres, ellipsoids, cylinders, cones, parallelepipeds, triangles, tetrahedra, etc, drawn from a database of shapes.

We demonstrate the viability and robustness of this approach by means of a computer simulation. Given the 2D image-plane projections of a curve and a certain primitive shape, we construct a fit-to-data L2 functional to optimize for the shape parameters, Euler angles for each pose, the 3D body coordinates of points on the curve, and the rotation matrix for each pose. For this highly nonlinear problem, we use a block coordinate descent algorithm to alternately optimize for each set of parameters in an iterative way. And by searching exhaustively through this database of primitives, we can infer the specific primitive shape and its size parameters that correspond most closely to the time series of observed 2D projections of a specific boundary. In our initial studies, through forward models we simulated the 2D projections of curves on an ellipsoid and a cylinder at two different poses, from which we tried to recover the original parameters and body coordinates used in the simulation. The results have been very promising. Even with just two poses, we could infer the correct shape from similar looking projections while discriminating against other shapes. Estimated shape parameters and Euler angles are also quite close to the original.

2 An Improved Spectral Correlation Approach for Segmentation

The MA approach is based on a direct correlation of the spectral traces at the different spatial pixels in an HSI datacube. Specifically, we construct a map of coarse-grained spectral correlations obtained by averaging pixel-to-pixel (*i.e.,*, fine-grained) normalized spectral correlations (NSC) over a small averaging cell, typically between 2x2 and 4x4 pixels in size. As the cell is translated across the image, the spatially averaged NSC exhibits dips at pixels at the boundaries between materials that are spectrally dissimilar. A map of dips in the NSC is then a map of the boundary pixels, but at a reduced spatial resolution given by the linear dimension of the averaging cell. Note that the MA approach is not particularly discriminating when the noise per data-cube voxel is large enough to wash out any spectral differences between spectrally similar materials.

Here we propose a more robust approach that exploits more explicitly the sparsity of the HSI dataset for a satellite when represented in a product basis of spatial primitives and spectral traces. Specifically, we first reduce the spectral size of the datacube by integrating along the spectral dimension over a number of spectral channels at a time. Such a coarser spectral binning into broader spectral bands improves the effective SNR of the reconstituted data via noise averaging, while reducing the data size and thus improving the computational efficiency of further data processing. The reduced data still contain enough information, however, to discriminate between materials and determine their segment boundaries as well as their fractional abundances.

To perform segmentation, we evaluate the magnitude of the two-dimensional spatial gradient vector in each spectral band. These magnitudes over W bands in all constitute the elements of a W-dimensional vector at each spatial pixel.

$$\mathbf{g}^{(ij)} = (g_1^{(ij)}, \ldots, g_W^{(ij)})^T, \tag{1}$$

where $g_w^{(ij)}$ is the magnitude of the spatial gradient of the total brightness $f_w^{(ij)}$ at the (ij)th spatial pixel in the wth spectral band,

$$g_w^{(ij)} = \sqrt{[f_w^{(i+1,j)} - f_w^{(ij)}]^2 + [f_w^{(i,j+1)} - f_w^{(ij)}]^2}. \tag{2}$$

The p-norm of the gradient vector $\mathbf{g}^{(ij)}$, for a sufficiently large p, then contains a weighted measure of the change of pixel brightness in the various reconstituted spectral bands. A map of this quantity yields all segment-boundary pixels as locations where it takes values significantly larger than at neighboring pixels, signaling a change of brightness across such boundary pixels. We varied p between 2 and 4 in our algorithm.

Note that this approach is superior to the MA approach in three ways. First, it suffers from no loss of spatial resolution. Second, by coarse-graining over the spectral dimension we deal with noise-averaged data and thus have better noise performance. Finally, by making p sufficiently large we allow the more discriminating spectral bands to be more heavily weighted than the less discriminating ones in performing the segmentation, a feature entirely absent from the MA approach.

Our new approach can also be used to determine the spectral traces and fractional abundances of the different material components. To do so, we start with a cell-averaged spectral trace at each spatial pixel of the HSI data that have been spectrally reduced as described earlier. Much as in the MA approach, the averaging cell here is taken to be small enough that one is sure of its being contained fully inside a single material component, e.g., by considering the map of purely spatial data obtained by fully integrating the original HSI data over its spectral dimension and staying away from regions of large changes in the spatial map. We then compute the NSC of the cell-averaged spectral trace with the spectral traces at the various spatial pixels in the original data, and use a thresholding approach to discriminate it from the traces of other material components on the satellite. The pixels in the map of the NSC that exceed the threshold are then assigned to the material to which the initial cell-averaged spectral trace belongs. The threshold map is a map of the spatial segment of that material, and the number of pixels in that map proportional to its fractional abundance. By subsequently averaging the spectral traces spatially over these pixels generates a better noise-averaged spectral trace of that material than the initial cell-averaged spectral trace.

2.1 Results

We employed a computer simulation of the Hubble Space Telescope (HST) generated from simple geometrical primitives like rectangles, cylinders, ellipses, etc. that are juxtaposed in definite spatial positions and orientations to mimic an idealized version of the HST. To all spatial pixels of a specific material component, an appropriate spectral trace from a database of spectral traces generated from actual laboratory measurements is assigned. Different components

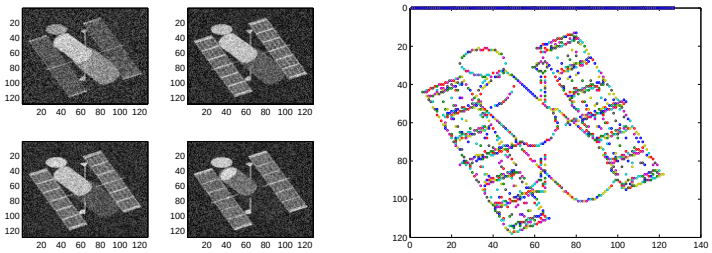

(a) (b)

Figure 1: (a) Noisy HST images in the hyperspectral channels 25, 50, 75, and 100, for the noise parameter $\eta = 0.5$. (b) Segment boundary pixels, indicated by small circles, as determined by our gradient-based multi-spectral segmentation approach.

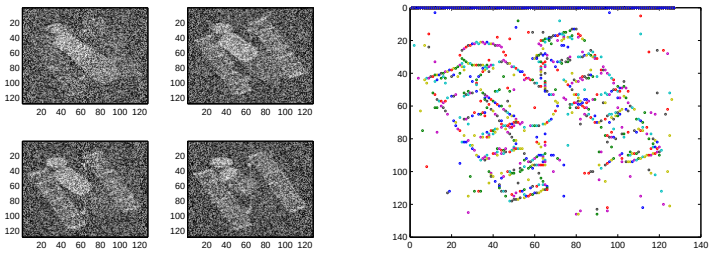

(a) (b)

Figure 2: Same as Fig. 1 except $\eta = 2.0$

of the HST mock-up have different spectral traces assigned to them. A 2D spatial projection of a particular orientation of the full 3D HST, along with the spectral traces at the various spatial pixels as described above, then generates the 3D HSI datacube that we consider our pristine dataset. Adding noise and blur to this datacube then simulates observationally more realistic HSI datasets. More details of the simulated datacube may be found in Ref. [3].

In Figs. 1 and 2, we display the results of spatial segmentation obtained by our improved approach under two different noise conditions. The additive noise is taken to be uniform over the range $\eta[0, \max(f)]$, where $\max(f)$ is the maximum brightness over all HSI voxels and η a noise strength parameter taken to be 0.5 and 2 for the two figures. The corresponding peak SNR values are $4\sqrt{3} \approx 7$ and $\sqrt{3} \approx 1.7$, while the true SNR values are considerably lower. The left panel in each figure contains four narrowband images of the noisy HST, specifically those contained in the 25th, 50th, 75th, and 100th hyperspectral channels of the noisy HSI datacube. The right panel displays the spatially segmented HST obtained by our approach, in which the boundary pixels have been identified for clarity by small circles. Note the excellent performance of our approach even under the highly noisy conditions corresponding to $\eta = 2$ in Fig. 2, as a visual comparison between the left and right panels readily indicates.

An example of two material spectral traces estimated by our approach compared to the corresponding "truth" trace is shown in Figs. 3 for the case of low and moderate noise levels, $\eta = 0.1$ and $\eta = 0.5$. Note the excellent estimate we obtain even under the moderate-noise conditions corresponding to $\eta = 0.5$. Although due to space limitation we cannot display our results on spectral trace estimates exhaustively here, it is not surprising that materials with the larger fractional abundances return the better estimates for their spectral traces in our method.

We now turn to the problem of exploiting a time series of 2D projected segment boundaries of a rotating satellite to determine both the corresponding time series of orientational parameters (Euler angles) and the fixed 3D shapes of its material segments.

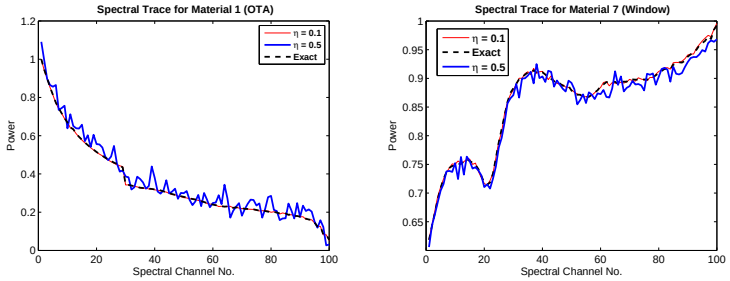

(a) (b)

Figure 3: (a) The spectral trace of the material component no. 1 estimated by our approach vs. truth, for two noise-parameter values, $\eta = 0.1$ and $\eta = 0.5$. (b) Same as (a) except the spectral trace for the material component no. 7 is plotted.

3 Relation between the Body and Space Coordinates for a Rotating Rigid Body

Consider the coordinates of a point on the rigid space object relative to two different coordinate systems, one that is fixed in the rotating body and the other that is fixed in space. Let us assume for simplicity that the translational motion of the object is frozen out by means of a perfect tracking system, leaving only the rotational motion of the object to be analyzed. We may take the origin of the two coordinate systems to be coincident at the object's center of mass (CM), taken fixed in the frozen translational frame of the object.

We shall take the Z axis of the space-fixed axes to be along the line of sight (LoS) from the observing telescope to the object's CM, so that the two-dimensional (2D) projected view of the object is simply the distribution of its (X, Y) coordinates with Z set to 0. Let (x, y, z) be the coordinates relative to a conveniently chosen set of body axes fixed in the body. The corresponding coordinates relative to the space axes, which we shall call space coordinates, are related to the above body coordinates via orthogonal matrices describing three coordinate rotations, the first by Euler angle α about the z axis, the second by Euler angle β about the resulting y axis, and the third by Euler angle γ about the resulting z axis. Specifically, we have

$$\begin{aligned} \begin{pmatrix} x \\ y \\ z \end{pmatrix} &= \mathbf{R}_x(\gamma)\mathbf{R}_y(\beta)\mathbf{R}_z(\alpha) \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} \\ &\equiv \mathbf{R}(\alpha, \beta, \gamma) \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} \end{aligned} \tag{3}$$

where $\mathbf{R}(\alpha, \beta, \gamma) \equiv \mathbf{R}_z(\alpha)\mathbf{R}_y(\beta)\mathbf{R}_z(\gamma)$ is the composite rotation matrix. The 3×3 matrices $\mathbf{R}_y$ and $\mathbf{R}_z$ take the following form in terms of the involved Euler angles:

$$\mathbf{R_z}(\alpha) = \begin{pmatrix} \cos\alpha & \sin\alpha & 0 \\ -\sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \mathbf{R_y}(\beta) = \begin{pmatrix} \cos\beta & 0 & -\sin\beta \\ 0 & 1 & 0 \\ \sin\beta & 0 & \cos\beta \end{pmatrix}. \tag{4}$$

Inverting Eq. (3) yields the space coordinates in terms of the body coordinates,

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \mathbf{R}^{-1}(\alpha, \beta, \gamma) \begin{pmatrix} x \\ y \\ z \end{pmatrix}. \tag{5}$$

The first two Euler angles specify the (azimuthal and polar angular coordinates of the) rotation axis while the third specifies the rotation angle. Each orientation, or pose, of the rigid object is determined by a specific set of the three Euler angles. A discrete time series of N object poses is then specified uniquely and fully in terms of the corresponding

discrete time series, $\{\alpha_n, \beta_n, \gamma_n | n = 1, \ldots, N\}$, of these three Euler angles. These $3N$ parameters contain complete information about the rotational kinematics of the rigid space object relative to its CM.

4 Relation of Coordinates to the Object Shape and Size Parameters

As the object rotates, its surface elements such as segment boundaries and other localized features evolve according to the local surface morphology, including its geometrical shape and size parameters. We shall henceforth make the simplifying assumption that the object is comprised of a small number of relatively simple geometrical primitives. Among such primitives we shall include parallelepipeds, ellipsoids, cones, and cylinders. Note that spheres, disks, and planar structures are limiting cases of the above shapes. Our assumption, which exploits an important aspect of the compressibility of man-made objects of primary interest here, enables us to idealize even the most complex objects in terms of possibly interpenetrating but otherwise simple geometrical shapes of different sizes and relative locations. Below we provide for four different primitive shapes how the body coordinates of points on curves on each such primitive are transformed to the space coordinates relative to the translationally frozen space axes.

4.1 Ellipsoid

The surface of an ellipsoid of semi major axes equal to a, b, and c in a coordinate system with origin at the center and coordinate axes along the principal axes of the ellipsoid is described parametrically in terms of the angular coordinates θ and ϕ by the equations

$$\xi = a\, \sin\theta\, \cos\phi; \;\; \eta = b\, \sin\theta\, \sin\phi; \;\; \zeta = c\, \cos\theta, \tag{6}$$

where ξ, η, ζ are the three Cartesian coordinates of a surface point, and the polar angle θ and azimuthal angle ϕ vary in the ranges $[0, \pi]$ and $[0, 2\pi)$, respectively. A curve on the ellipsoid may be expressed parametrically by two functional equations,

$$\theta = \theta(s), \;\; \phi = \phi(s), \tag{7}$$

that describe how the angular coordinates of the points on the curve evolve as functions of a single parameter s limited to a finite range, $s \in [s_1, s_2]$.

If the center of the ellipsoid is at location (x_0, y_0, z_0) relative to the CM of the parent object and the orientation of the elliposidal axes relative to the body axes is described by the Euler angles $(\alpha_0, \beta_0, \gamma_0)$, then we may relate the body coordinates (x, y, z) of an ellipsoidal surface point to its principal coordinates (ξ, η, ζ) by means of the transformation (5) decribed in the previous subsection as

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x_0 \\ y_0 \\ z_0 \end{pmatrix} + \mathbf{R}^{-1}(\alpha_0, \beta_0, \gamma_0) \begin{pmatrix} \xi \\ \eta \\ \zeta \end{pmatrix}. \tag{8}$$

A use of the parametric equations (6) and (7) then determines the body coordinates of points on a curve resident on the ellipsoidal surface. A final substitution of this body-coordinate column vector into Eq. (5) yields the space coordinates (X, Y) of the 2D projected image of that curve as a function of the current set of Euler parameters (α, β, γ) of the orientation of the parent object relative to the space-fixed axes.

4.2 Cylinder

The principal coordinates of a right elliptical cylinder of semi-major and semi-minor axes a and b, and length L are parametrically defined as

$$\xi = a\, \cos\phi; \;\; \eta = b\, \sin\phi; \;\; \zeta \in [0, L]. \tag{9}$$

A curve on the cylindrical surface too may be specified parametrically as

$$\phi = \phi(s), \;\; \zeta = \zeta(s). \tag{10}$$

A substitution of the form (9) into Eq. (8), with the parametric functions (10) specifying a curve of interest, then yields the requisite space coordinates of the points on the cylindrical curve as a function of the orientation of the parent object.

Similar parameterizations can also be given for other primitive shapes, including the cone and parallelepiped, but for lack of space we do not give any details in this paper.

5 Recovery of Primitive Shape Type and Parameters from Rotational Kinematics of Surface Curves

The preceding section describes how curves on surfaces project onto the observation plane of the remote camera. This description will now be used to develop the full forward model for recovery of the orientation parameters of the overall space object and the surface shape, size, location, and relative-orientation parameters for its primitives. We shall assume that the curves we work with are all unobstructed segment boundaries, or unobstructed parts thereof, that remain fixed in the object, thus excluding those boundaries that are either formed merely by the shadowing of one component by another or are exterior boundaries.

Our approach for recovering the primitive type and parameters consists of the following steps:

1. Acquire a time series of hyperspectral data sets for successive, closely spaced orientations of the space object;
2. Use the proposed spectral-correlation approach to map out all unobstructed segment boundary curves on the space object for each hyperspectral data cube in the time series;
3. Identify the time series of a specific segment boundary curve that separates two different components of the object;
4. Try to fit the time series of the segment boundary assuming that it lies on an arbitrarily chosen primitive shape out of a data base of different shapes;
5. If the fit is good, then accept the primitive shape; otherwise reject and move to the next primitive shape in the data base and try to fit;
6. Repeat the process until the chosen primitive shape is found to be consistent with the time series of the segment boundary;
7. Using the same approach as above find the second primitive shape consistent with this time series;
8. Repeat for the time series of other segment boundaries until they are exhausted.

Note that since a specific boundary lies on at least two primitive surfaces, it would seem to be impossible to assign the two correctly matched shapes unambiguously to the two surfaces, resulting in a two-fold ambiguity of shape assignment. However, since multiple boundaries and multiple primitive surfaces are typically present on a man-made object, this two-fold ambiguity is automatically resolved for the typical object when a sufficiently long time series of poses is used.

6 The inverse projection problem

As a first cut at the problem, let us restrict to a single snapshot, in which we are looking at curves on a primitive shape at the pose defined by the Euler angles (α, β, γ). Given observed projections and a primitive shape k, the inverse projection problem is defined as to minimize the following least square functional and to jointly optimize for the unknown body coordinates p_i, Euler angles (α, β, γ), and the shape-specific parameters θ_k.

$$\mathcal{F}_k = \min_{\alpha,\beta,\gamma,\theta_k} \sum_{i=1}^{N} \|P_i - A(\alpha, \beta, \gamma) p_i\|_2^2, \tag{11}$$

where $A(\alpha, \beta, \gamma) = [R_z^{-1}(\alpha) R_y^{-1}(\beta) R_x^{-1}(\gamma)]_{1:2,:}$ is the rotation matrix but including only the first two rows, $P_i = (X_i, Y_i)^T$ are observed 2D space coordinates, $p_i = (x_i, y_i, z_i)^T$ are the unknown body coordinates, θ_k is the shape parameter vector. Points on a primitive shape must satisfy a shape constraint function f_k, e.g. for a sphere, we would have $f_k(x_i, y_i, z_i) = x_i^2 + y_i^2 + z_i^2 - r^2$. Hence (11) becomes a constrainted optimization problem subject to

$$f_k(p_i) = 0. \tag{12}$$

Since the ultimate purpose is to identify a certain shape, we would have to solve (11) K times to search for the shape that produces the minimum $\mathcal{F}$.

Ambiguity can arise when different objects having the same projection, or curves on the same object having the same projection. For example, all points on a cylinder will project to a circle along the cylinder length axis. A second view might solve this ambiguity, or it might not. For example, there are two viewing directions along the cylinder length axis, both will have the same projection. And we would not be able to detect the Euler angles and the curves exactly on a sphere, because if we rotate the original curves by an arbitrary angle, their projections would be identical as the observed after we add the original Euler angle with the added ones.

Next we try to solve (11) by taking a more detailed look. To explicitly write out P_i, p_i and incorporate the shape constraints into (11), we have

$$\mathcal{F}_k = \min_{\alpha,\beta,\gamma,\theta_k} \sum_{i=1}^{N} \left\| \begin{pmatrix} X_i \\ Y_i \end{pmatrix} - A(\alpha,\beta,\gamma) \begin{pmatrix} x_i \\ y_i \\ f_k(x_i,y_i) \end{pmatrix} \right\|_2^2. \tag{13}$$

Note that only (X_i, Y_i) are observables, while the primitive shape requirements would restrict the body coordinates to two parameters (x_i, y_i), the problem is still underdetermined, with the number of unknowns at $2N + S_k + 3$, where N is the number of observed points, and S_k is the number of shape parameters, e.g. $S_k = 1$ for a sphere where the only parameter needed is the radius r. However, if we have observations from a different angle, $(\alpha_2, \beta_2, \gamma_2)$, while denoting the original angles as $(\alpha_1, \beta_1, \gamma_1)$, and assume a full registration, that is to say, body coordinates (x_i, y_i) would remain the same for two sets of observations (X_{i1}, Y_{i1}) and (X_{i2}, Y_{i2}), then we only need $\lceil [\sup_k(S_k) + 6]/2 \rceil$ number of extra points in (1) to make the problem exact. Since $N \gg \sup_k(S_k)$, we would only need two views and register only a few points between the two views. After adding in the multi-view observations, the functional in (13) becomes

$$\min_{\alpha,\beta,\gamma,\theta_k} \sum_{i=1}^{N} \sum_{j=1}^{2} \left\| \begin{pmatrix} X_{ij} \\ Y_{ij} \end{pmatrix} - R_z^{-1}(\alpha_j) R_y^{-1}(\beta_j) R_x^{-1}(\gamma_j) \begin{pmatrix} x_i \\ y_i \\ f_k(x_i,y_i) \end{pmatrix} \right\|_2^2.$$

Notice that though the unknowns $\alpha_j, \beta_j, \gamma_j$, and (x_i, y_i) join together in a multiplication fashion, thus the problem is highly nonlinear, each of them is only contained in a separated factor, hence we can use the block descent algorithm to alternatingly optimize for each of them for a given set of shape parameter values. Thus, the original problem is turned into four simpler subproblems, namely

$$\begin{pmatrix} x_i^{(p+1)} \\ y_i^{(p+1)} \end{pmatrix} = \operatorname*{argmin}_{(x_i,y_i)} \left\| \begin{pmatrix} X_{ij} \\ Y_{ij} \end{pmatrix} - R_z^{-1}(\alpha_j^{(p)}) R_y^{-1}(\beta_j^{(p)}) R_x^{-1}(\gamma_j^{(p)}) \begin{pmatrix} x_i \\ y_i \\ f_k(x_i,y_i) \end{pmatrix} \right\|_2^2, \tag{14}$$

$$\alpha_j^{(p+1)} = \operatorname*{argmin}_{\alpha_j} \left\| \begin{pmatrix} X_{ij} \\ Y_{ij} \end{pmatrix} - R_z^{-1}(\alpha_j^{(p)}) R_y^{-1}(\beta_j^{(p)}) R_x^{-1}(\gamma_j) \begin{pmatrix} x_i^{(p+1)} \\ y_i^{(p+1)} \\ f_k\left(x_i^{(p+1)}, y_i^{(p+1)}\right) \end{pmatrix} \right\|_2^2, \tag{15}$$

and two analogous ones for β_j and γ_j. Due to the nonlinearity introduced by the shape function, f_k, we must use a constrainted nonlinear algorithm such as the Trust-Region algorithm [4].

Next we optimize the shape parameters θ_k with known Euler angles at two different views $(\alpha_j, \beta_j, \gamma_j)$ and known body coordinates (x_i, y_i). The optimal solution of θ_k can be analytically derived by setting the gradient of the least square functional to zero and solve it. Lastly we can analytically solve for the three Euler angles again by setting the gradient of the least square functional to zero and solving it. It's worthwhile to note that the solutions for the Euler angles are independent of the primitive shapes and thus can be used for any type of shape.

7 Tests with simulated data

In our tests, we have simulated the two projections of curves on four primitive shapes, namely ellipsiod, cylinder, cone, and parallelepiped. Projections are computed through four forward models, one for each primitive shape. Two sets of fixed Euler angles are used as input to the forward models. Here we have bypassed the registration step, because the two sets of points are always fully registered in simulation.

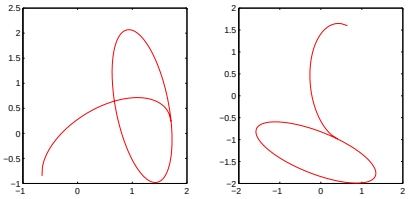

Figure 4: Simulated projections of two curves on an ellipsoid.

7.1 Ellipsoid

On an ellipsoid, with radii $(1, 2, 3)$, we sample evenly spaced 100 points on two curves, a full circle and an half circle, lying in two planes perpendicular to each other and the two sets of Euler angles are fixed at $(-\pi/6, \pi/4, -\pi/4)$ and $(\pi/6, -\pi/4, \pi/4)$. The body coordinates of sampled points along with the Euler angles are used as inputs to a forward model to generate the two observed projections shown in Figure 4. By solving the inverse projection problem through the block descent algorithm with random starting values, in Figure 5(a) we have plotted the estimated projections, i.e. $A(\hat{\alpha}_j, \hat{\beta}_j, \hat{\gamma}_j)\hat{p}_i$, in blue and the observed projections in red in every five iterations for the first 100 iterations. Clearly the two sets of curves start to converge after 50 iterations. The same inverse projection problem is solved but by using a cylinder shape, the same plots of observed and estimated curves are shown in Figure 5(b) and clearly the two sets of curves could not converge. The same problem is also repeated using the other two primitive shapes but only the ellipsoid shape provides the satisfactory answer, thus we have successfully identified the true shape in this experiment. The estimated shape parameters are shown in first column of Table 1 along with the true parameters for comparison.

7.2 Cylinder

Similar to the curves on an ellipsoid, this time on a cylinder with radii $(1, 1)$ we sample 100 evenly spaced points on two curves, a full circle and a curve representing the intersection of two same size cylinders intersecting each other perpendicularly. And the Euler angles remain the same as $(-\pi/6, \pi/4, -\pi/4)$ and $(\pi/6, -\pi/4, \pi/4)$. The inverse projection problem is again solved by first using the cylinder shape and then the ellipsoid shape. As for the ellipsoid, the two sets of curves converge nicely when using the true (cylinder) shape, but not so well when using the wrong (ellipsoid) shape. This shape mismatch is also well verified in the value of the functional $\mathcal{F}_k$, which is smaller for the cylinder than for the ellipsoid. We thus correctly identify the true shape. The estimated shape parameters are shown in second column of Table 1 along with the true parameters for comparison.

7.3 Other Shapes

We tested our algorithm for two other shapes, namely a cone and a parallelepiped. For the cone, we chose 100 evenly spaced points on the 2D projections of a regular curve, namely a circle formed by a transverse section of the cone, and a saddle curve, for the same two poses as for the ellipsoid and the cylinder. The reconstructed 2D projections converged nicely to the simulated projection data within a few tens of iterations for the cone shape, while they were mismatched for other shapes even after 100 iterations.

For the parallelepiped, we sampled 150 points on three curves, one on each of three intersecting surfaces for the same two poses as for the other shapes. The two sets of curves converge to each other within the first 20 iterations. Due to the coupling of principal edges and the Euler angles, however, it becomes very difficult to estimate both at the same time. But, in terms of identifying the true parallelepiped shape, the algorithm still does a fine job, that is to say, when using other shapes, the functional values are always much higher.

In the table below, we summarize our results for all of the shapes we simulated and tested except the parallelepiped. We note that in each case the recovery of shape parameters is excellent, while the Euler angles of the poses are determined only to within about 10% accuracy. We expect the situation to improve, however, when more poses are included. But it remains to be seen whether having more segment boundaries as on a space object containing multiple interpenetrating primitives, such as out simulated HST, and noise yield an inverse problem that will be amenable to a properly regularized version of our present algorithm.

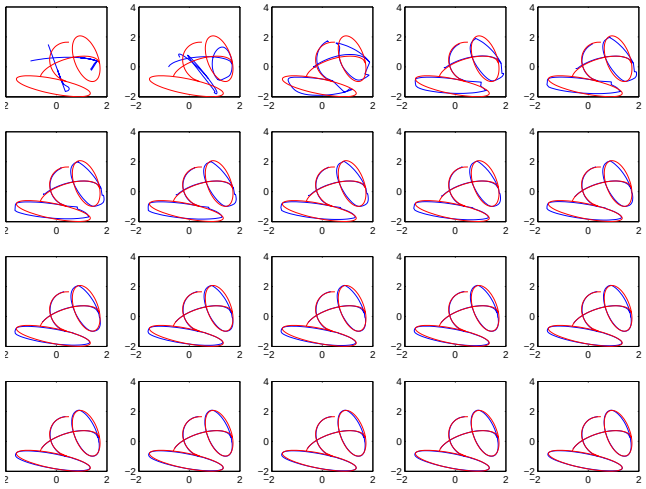

(a)

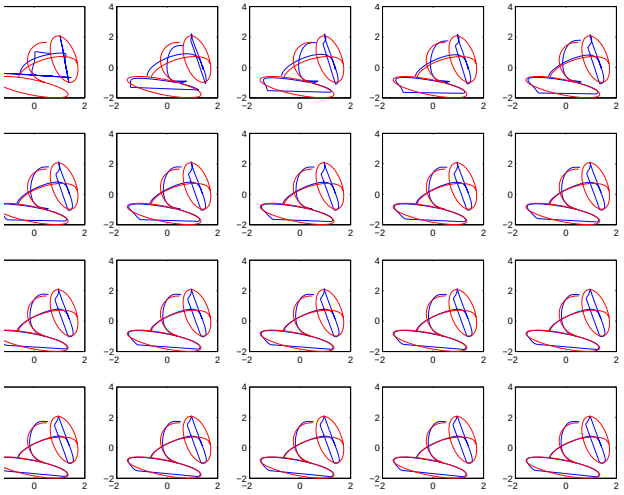

(b)

Figure 5: (a) Plots of estimated and observed curves using an ellipsoid shape in every 5 iterations. (b) Plots of estimated and observed curves using a cylinder shape in every 5 iterations.

Table 1: Estimated vs. True Parameters

	Curves on an Ellipsoid		Curves on a Cylinder		Curves on a Cone	
Parameters	True	Estimated	True	Estimated	True	Estimated
α_1	-.5236	-.4106	-.5236	-.4919	-.5236	-.4486
β_1	.7854	.7983	.7854	.8128	.7854	.8101
γ_1	-.7854	-.6559	-.7854	-.7610	-.5236	-.4677
α_2	.5236	.6375	.5236	.5541	.5236	.5980
β_2	-.7854	-.8001	-.7854	-.8136	-.7854	-.8095
γ_2	.7854	.6967	.7854	.7668	.5236	.4696
a	1	.9572	1	1.0085	1	.994
b	2	1.9891	1	.9990	1	.996
c	3	2.9274				

8 Conclusions

The vastly redundant information contained in the hyperspectral image datacubes for man-made space objects can be mined by means of a spectral-correlation analysis. This yields a simple method for performing spatial segmentation and material identification for such objects. We have shown that our approach performs quite well even in the presence of large amounts of data noise.

The kinematics of the segment-boundary projections along the line of sight as the space object rotates rigidly carry information about the local shape and geometry of the surface as well as the kinematics of its overall orientational parameters. A time series of HSI datacubes corresponding to changing object poses can then be exploited to recover its local surface shape, size, and geometry. An iterative numerical fitting approach based on rigid-body kinematics has been suggested here to achieve this goal which is likely to be important for general SOI/SSA missions.

As a first step toward validating this fitting approach, we have introduced an iterative nonlinear least-squares optimization approach to recover the shape, the dimensions, and the Euler angles of a simple primitive from a knowledge of two LoS projections of two curves on its surface for two different poses. For this nonlinear problem, the block coordinate descent algorithm introduced here works well for optimizing the functional in (11) given a specific shape, and we can correctly differentiate the true primitive shape from others using the functional values. If one uses the wrong shape, the algorithm may never converge, while it always converges with the true shape. The estimated Euler angles and the shape parameters agree reasonably well with their original values for the first three primitive shapes, but for the parallelepiped shape, due to an intrinsic ambiguity introduced by the principal edges, it remains difficult to find the exact Euler angles. More poses and curves may be needed to remove such ambiguities for the more complicated surfaces.

Acknowledgment

The work reported here was supported by the US Air Force Office of Scientific Research under award numbers FA9550-08-1-0151 and FA9550-09-1-0495. We also acknowledge helpful conversations with Dr. Douglas Hope.

References

[1] F. Li. M. Ng, R. Plemmons, Q. Zhang, and S. Prasad, "Hyperspectral image segmentation, deblurring, and spectral analysis for material identification," SPIE Proceedings, vol. 7701, to be published (2010).

[2] S. Prasad and D. Hope, "Information theoretic characterizations of compressive-sensing based space object identification," SPIE Proceedings, vol. 7800, to be published (2010).

[3] Q. Zhang, H. Wang, R. Plemmons, and P. Pauca, "Tensor Methods for Hyperspectral Data Analysis: A Space Object Material Identification Study", Journal of the Optical Society of America, A., vol. 25, No. 12, pp. 3001-3012, December, 2008.

[4] T. Coleman and Y. Li, "An interior trust region approach for nonlinear minimization subject to bounds," SIAM J. Optimiz., vol. 6, pp. 418-425 (1996).