ANOMALY DETECTION WITH SPARSE UNMIXING AND GAUSSIAN MIXTURE MODELING OF HYPERSPECTRAL IMAGES

Acar Erdinç, Selim Aksoy

Department of Computer Engineering
Bilkent University
Bilkent, 06800, Ankara, Turkey
{acar.erdinc,saksoy}@cs.bilkent.edu.tr

ABSTRACT

We propose an anomaly detection method that uses Gaussian mixture models for characterizing the scene background in hyperspectral images. First, the full spectrum is divided into several contiguous band groups for dimensionality reduction as well as for exploiting the peculiarities of different parts of the spectrum. Then, sparse spectral unmixing is performed for identifying significant endmembers in the scene, and hierarchical clustering in the abundance space is used for identifying pixel groups that contain these endmembers. Next, these pixel groups are used for initializing individual Gaussian mixture models that are estimated separately for each spectral band group. Finally, the Gaussian mixture models for all groups are fused for obtaining the final anomaly map for the scene. Comparative experiments showed that the proposed method performed better than two other density-based anomaly detectors, especially for small false positive rates, on an airborne hyperspectral data set.

Index Terms— Anomaly detection, spectral unmixing, Gaussian mixture model, hyperspectral imaging

1. INTRODUCTION

High spectral resolution in hyperspectral images provides the capability to discriminate the physical characteristics of different materials and enables their identification in remotely sensed scenes. The scenario where the problem of interest is the detection of small rare objects or materials that have different spectral characteristics compared to their surroundings is called anomaly detection. Anomaly detection has been a popular problem in remote sensing where the common approach is to first model the image background and then to use a detector that quantifies the difference of a particular pixel from this background as the confidence of an anomaly existing at that pixel. Numerous approaches for both background characterization and detector functions have been proposed in the literature [1]. In this paper, our focus is on the modeling of complex scenes that contain an unknown number of background materials in an anomaly detection scenario.

The commonly used RX anomaly detector and many of its various extensions do not necessarily perform well in complex scenes due to the heterogeneity of the image background and the limitations of the single multivariate Gaussian in modeling the multiple land cover classes existing in that background. A popular alternative has been to use a Gaussian mixture model (GMM) based on the assumption that the background consists of multiple classes each of which can be modeled using a separate Gaussian component in the mixture [2]. However, GMM-based background modeling has two important issues: there is often no prior information regarding the number of components, and GMM estimation using all spectral bands may suffer from the curse of dimensionality due to increasing model complexity with increasing number of components. The former issue can be solved by an empirical search procedure when the number of different background materials is unknown in an unsupervised setting. The latter issue often necessitates a dimensionality reduction step, such as principal components analysis (PCA), prior to GMM estimation. However, there is no unanimously accepted procedure regarding which principal components should be used. There is also the potential problem that the resulting transformation that tries to maximize the data variance may not retain important discriminative information embedded in the original spectral bands.

In this paper, we use a GMM-based background model for anomaly detection. First, we perform spectral partitioning to reduce the data dimensionality while preserving the original physical characteristics of the input data. After dividing the data into several groups of contiguous spectral bands (Section 2), we perform spectral unmixing to identify the number of dominant endmembers and to automatically determine the number of components in the background model (Section 3). Lastly, we fuse the individual GMMs estimated for each spectral band group separately to obtain the final probability of having an anomaly at each pixel (Section 4). We present experimental results using an airborne hyperspectral image that contains several materials placed on a natural landscape background for an anomaly detection scenario (Section 5).

Conclusions are given at the end (Section 6).

2. SPECTRAL PARTITIONING

As mentioned above, high spectral dimensionality limits the applicability of flexible models for hyperspectral data analysis due to potential problems such as overfitting or convergence issues in the estimation process. However, popular dimensionality reduction techniques such as PCA may not be able to identify the intrinsic dimensionality of the data, and may lose information that may be useful for the subsequent analysis. Band selection methods provide an alternative by selecting an appropriate subset of original bands [3]. However, identifying how many bands to select and quantifying the appropriateness of a subset can be difficult in an unsupervised setting. Thus, internal characteristics of the data such as variance or correlation remain as a limited set of alternatives in such scenarios.

In this paper, we used an empirical division of the spectral bands into several contiguous groups. In particular, the experiments in Section 5 used 12 band groups with equal number of bands. The number of groups was determined based on the resulting dimensionality of each group so that the subsequent estimation steps can be performed without any singularity problems. We are planning to evaluate different band selection methods in future work.

3. SPECTRAL UNMIXING

After the band groups are identified, the next problem is to estimate a background GMM for each group. The main problem at this step is to identify the number of components in the mixture and to find a proper initialization of the estimation procedure so that an effective background model can be obtained. We use spectral unmixing to identify the number of dominant endmembers in the scene so that it can be related to the number of background components and can be used as input to the estimation procedure.

Linear unmixing is the most widely used method for identifying endmembers in the data. In the linear model

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}, \tag{1}$$

$\mathbf{y} \in \mathbb{R}^L$ is the observation vector for a pixel, $\mathbf{x} \in \mathbb{R}^M$ is the fractional abundance vector at that pixel for the spectral library $\mathbf{A} \in \mathbb{R}^{L\times M}$, and $\mathbf{n} \in \mathbb{R}^L$ is the error where L is the number of spectral bands and M is the number of signatures in the library. Sparse unmixing has recently been proposed for solving the linear spectral unmixing problem in (1) while not requiring any prior knowledge of the number of endmembers [4]. The unmixing problem can be formulated as an l_1-norm regularized least squares regression problem subject to an additional non-negativity constraint as

$$\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \lambda\|\mathbf{x}\|_1 \quad \text{subject to } \mathbf{x} \geq \mathbf{0} \tag{2}$$

where λ is the regularization parameter. This formulation enforces sparsity in the solution vector $\mathbf{x}$, and implicitly tries to identify the most dominant endmembers in the given observation vector $\mathbf{y}$. We also expect that sparse unmixing within each spectral band group will provide a better exploitation of the spectral characterization of the data within that group as opposed to unmixing in the full range of spectral bands due to potential mismatches between the signatures in the spectral library and the spectral signatures in the scene because of the differences in the conditions under which the data are acquired.

For a given spectral band group and a spectral library, (2) is solved for each pixel, and the resulting abundance vectors are recorded. As mentioned above, the purpose of spectral unmixing in this work is to identify the dominant components which are then used for modeling the background. However, the components that appear in a hyperspectral image are not necessarily pure due to the spatial resolution of the imagery and the size of physical materials in the scene. Thus, the so-called background components may consist of pixel groups that contain either pure pixels or pixels appearing to be a mixture of different endmembers. In such a scenario, applying a predetermined threshold to the fractional abundances to obtain the number of dominant background components proves to be insufficient, because it will only return those endmembers with a certain level of purity and higher, but not the endmembers that appear together at sub-pixel level. Thus, the desired components should be acquired by exploiting the correlation of endmembers at sub-pixel level throughout the scene.

In this paper, we use hierarchical clustering in the fractional abundance space to identify groups of pixels based on their spectral content. The previous unmixing step is expected to identify only the dominant signatures in the scene so that the clustering step can operate in an intrinsically lower dimensional space. We use the average linkage agglomerative hierarchical clustering algorithm with the Euclidean distance as the dissimilarity measure for abundance vectors. The average linkage criterion is used because we want all pixels that are selected as belonging to the same cluster to have similar abundance values. The number of clusters is automatically determined from the dendrogram resulting from the clustering [5].

4. GAUSSIAN MIXTURE BACKGROUND MODEL

The next step consists of the estimation of a GMM background model for each spectral band group. The pixel clusters identified for each spectral band group in the previous step are used to automatically initialize the Expectation-Maximization-based GMM estimation procedure for that group. The number of Gaussian components is obtained from the number of clusters, and the initial mean vector and the covariance matrix for each component are computed from the pixels within the corresponding cluster.

The final anomaly index at each pixel is obtained by fusing the resulting GMMs from individual spectral band groups. After negating the conditional background probabilities and normalizing them to the $[0, 1]$ range, fusion is performed by the "max" rule that means that a pixel that is identified as an anomaly with a high-probability by at least one of the spectral band groups is considered an anomaly after fusion.

5. EXPERIMENTS

5.1. Data set

The proposed methodology is evaluated using an airborne hyperspectral image. An anomaly detection scenario was prepared by placing different materials such as fabrics, carpets, wooden sheets, floor tiles, etc., on a natural landscape background in Turkey. The scene also has two cars that can be considered as anomalies. The image that was acquired with a visible and near infrared camera from an altitude of 500 m contains 1500×885 pixels and 182 bands covering the spectral range from 399 nm to 978 nm with a spatial resolution of 16 cm. Data collection and atmospheric correction were performed by HAVELSAN, Inc. The ground truth was prepared by manual delineation of 19 objects in the scene. The sizes of the objects vary between 4 and 900 pixels where 14 of them are smaller than 180 pixels. The spectral library used in the experiments contains signatures acquired with a spectrometer in the same scene as well as in other scenes under different atmospheric conditions. The RGB image and the corresponding ground truth mask are shown in Figures 1(a) and 1(b), respectively.

5.2. Results

In the experiments, we used 12 spectral band groups where 10 of them had 15 bands and the remaining 2 had 16 bands. The groups corresponded to different parts of the visible and near infrared spectrum. The regularization parameter for spectral unmixing was empirically set as 0.1. We used the RX anomaly detector that is built by using the full spectrum and another detector named RX-GMM that replaces the Gaussian in the RX framework by a GMM that is built by using 8 Gaussian components estimated using 10 PCA components that retain 98% of the variance in the full data as competing methods. These methods are compared to the proposed anomaly detector named Unmix-GMM in Figure 1. We also used the opening by reconstruction operator with a disk structuring element with a radius of 2 pixels for post-processing the anomaly map by eliminating isolated false positives for each method. The ROC curves were obtained by varying the detection threshold. The overall area-under-the-curve (AUC) measures were obtained as 0.8228, 0.9472, 0.9746 for the RX, RX-GMM, and Unmix-GMM detectors, respectively, after post-processing. Most of the false alarms were individual trees, large stones, and parts of the road in the scene. The results showed that the proposed method performed better than the other two, especially for small false positive rates after post-processing.

6. CONCLUSIONS

We described a method for anomaly detection that combined spectral unmixing and Gaussian mixture models. First, the full spectrum was divided into contiguous groups of spectral bands for dimensionality reduction. Then, sparse spectral unmixing was performed for identifying significant endmembers in the scene, and hierarchical clustering in the abundance space was used for identifying pixel groups that contained these endmembers. Next, these pixel groups were used for initializing Gaussian mixture models that were estimated separately for each spectral band group. Finally, the Gaussian mixture models for all groups were fused for obtaining the anomaly map for the scene. Comparative experiments showed that the proposed method performed better than two other density-based anomaly detectors on an airborne hyperspectral data set.

7. ACKNOWLEDGMENTS

The authors thank HAVELSAN, Inc. for collecting and preparing the data set.

8. REFERENCES

[1] S. Matteoli, M. Diani, and J. Theiler, "An overview of background modeling for detection of targets and anomalies in hyperspectral remotely sensed imagery," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 7, no. 6, pp. 2317–2336, June 2014.

[2] D. W. J. Stein, S. G. Beaven, L. E. Hoff, E. M. Winter, A. P. Schaum, and A. D. Stocker, "Anomaly detection from hyperspectral imagery," *IEEE Signal Processing Magazine*, vol. 19, no. 1, pp. 58–69, January 2002.

[3] Y. Liu, J. Li, A. Plaza, J. Bioucas-Dias, A. Cuartero, and P. G. Rodriguez, "Spectral partitioning for hyperspectral remote sensing image classification," in *IEEE International Geoscience and Remote Sensing Symposium*, Quebec, Canada, July 13–18, 2014, pp. 3434–3437.

[4] M.-D. Iordache, J. M. Bioucas-Dias, and A. Plaza, "Sparse unmixing of hyperspectral data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 6, pp. 2014–2039, June 2011.

[5] R. Mojena, "Hierarchical grouping methods and stopping rules: An evaluation," *The Computer Journal*, vol. 20, no. 4, pp. 359–363, 1977.

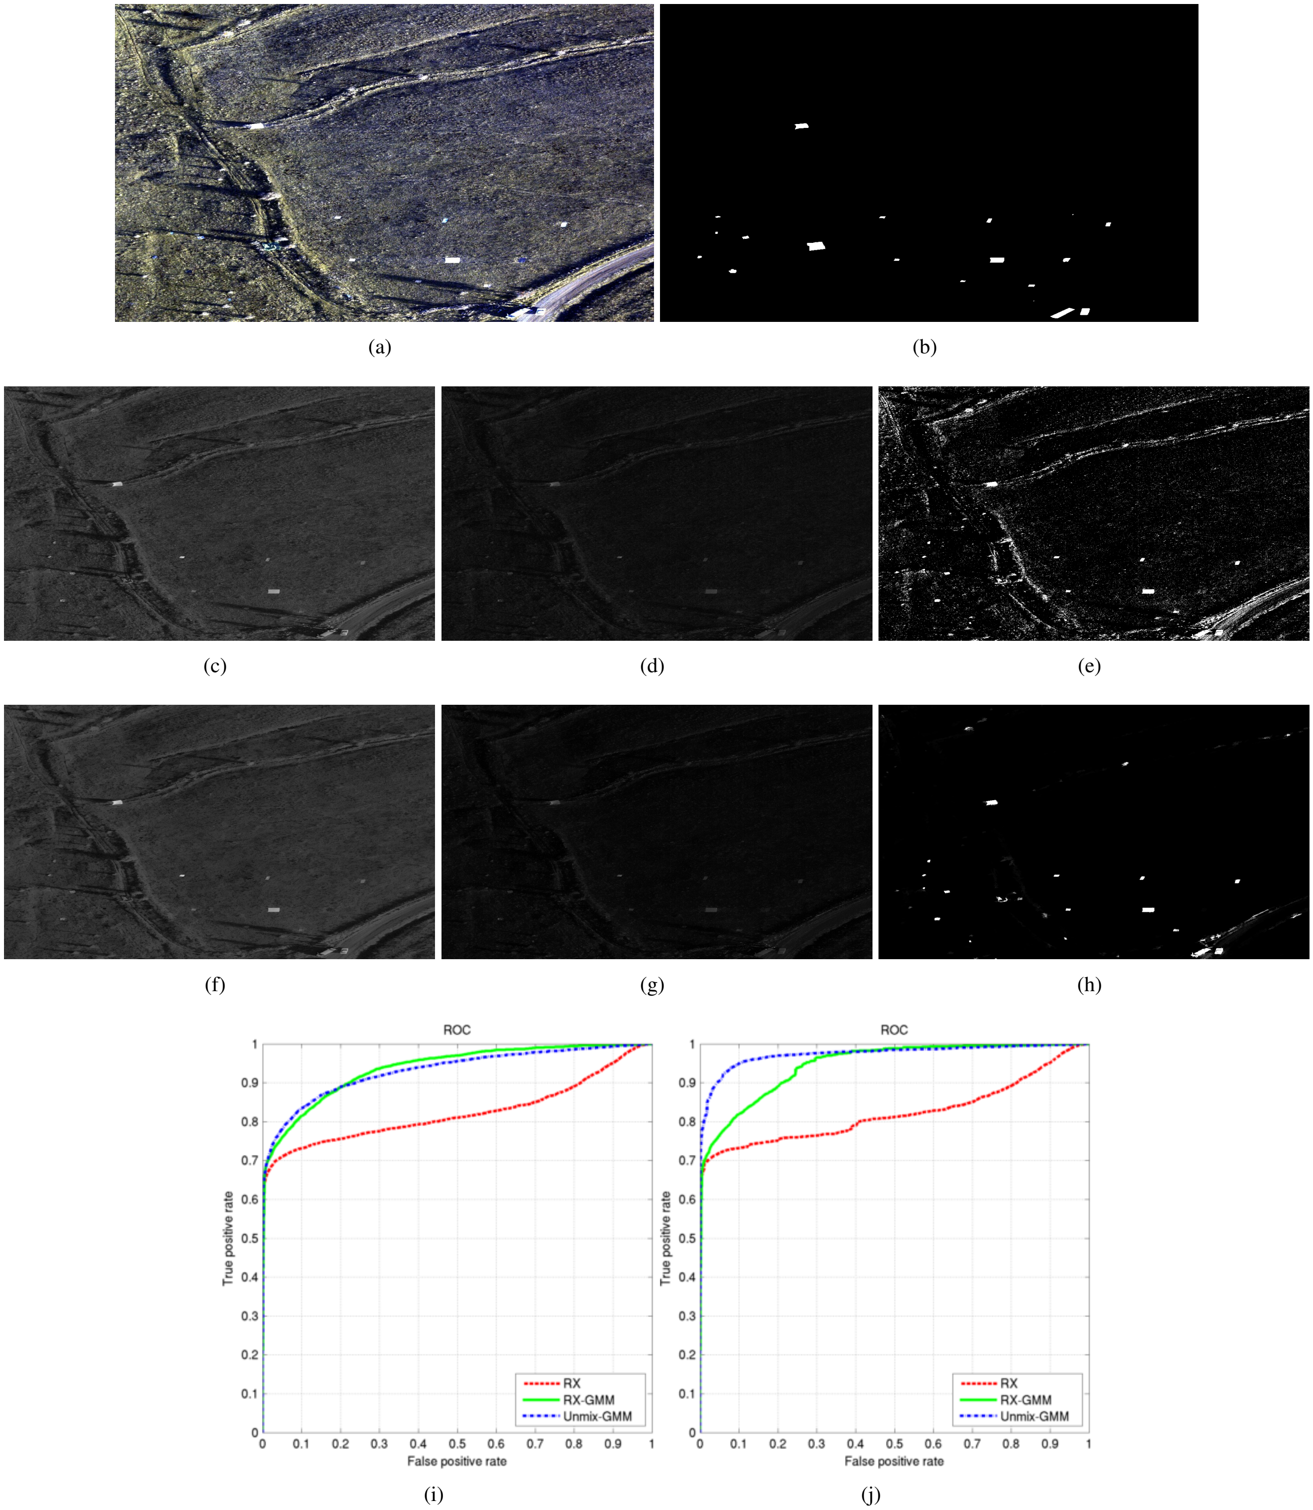

Fig. 1. Anomaly detection results. (a) RGB image. (b) Binary mask of 19 small objects. (c) Log-probability anomaly map for the RX detector before post-processing. (d) Log-probability anomaly map for the RX-GMM detector before post-processing. (e) Anomaly map for the proposed Unmix-GMM detector before post-processing. (f) Log-probability anomaly map for the RX detector after post-processing. (g) Log-probability anomaly map for the RX-GMM detector after post-processing. (h) Anomaly map for the proposed Unmix-GMM detector after post-processing. (i) ROC curve for the results before post-processing. (j) ROC curve for the results after post-processing.