Band Selection for Hyperspectral Image Classification Using Mutual Information

Baofeng Guo, Steve R. Gunn, R. I. Damper *Senior Member, IEEE* and J. D. B. Nelson

***Abstract*— Spectral band selection is a fundamental problem in hyperspectral data processing. In this paper, a new band-selection method based on mutual information (MI) is proposed. MI measures the statistical dependence between two random variables and can therefore be used to evaluate the relative utility of each band to classification. A new strategy is described to estimate the mutual information using *a priori* knowledge of the scene, reducing reliance on a 'ground truth' reference map, by retaining bands with high associated MI values (subject to certain so-called 'complementary' conditions). Simulations of classification performance on 16 classes of vegetation from the AVIRIS 92AV3C dataset show the effectiveness of the method, which outperforms an MI-based method using the associated reference map, an entropy-based method and a correlation-based method. It is also competitive with the steepest ascent (SA) algorithm at much lower computational cost.**

I. Introduction

Hyperspectral sensors simultaneously measure hundreds of contiguous spectral bands with a fine spectral resolution, e.g., $0.01\,\mu$m. For instance, the AVIRIS hyperspectral sensor [1] has 224 spectral bands ranging from $0.4\,\mu$m to $2.5\,\mu$m. Such a large number of bands implies high-dimensionality data, presenting several significant challenges to image classification. The dimensionality of input space strongly affects the performance of many supervised classification methods [2]. There is likely to be redundancy between bands; and some bands may contain less discriminatory information than others. Finally, high-dimensional data impose requirements for storage space, computational load and communication bandwidth that tell against real-time applications.

It is therefore advantageous to remove bands which convey little or no discriminatory information. Many band-selection techniques have been proposed, such as search-based methods [3], [4], [5], [6], transform-based methods [7], [8], and information-based methods [6], [9]. Other techniques include a scheme trading spectral for spatial resolution [10], maximization of Spectral Angle Mapper [11], high-order moments [12], wavelet analysis [13], etc. However, there are still some challenges to apply these techniques effectively, such as high computational cost, presence of local minima problems, difficulties for real-time implementation, etc.

Manuscript received December 12, 2005; revised April 4, 2006. Baofeng Guo, Steve R. Gunn, R. I. Damper and J. D. B. Nelson are with Information: Signals, Images, Systems (ISIS) Research Group, School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, UK. This work was supported by the Data and Information Fusion (DIF) Defence Technology Centre funded by the UK Ministry of Defence and managed by General Dynamics Limited and QinetiQ.

In this paper, a new band-selection method based on a measure of mutual information (MI) is proposed. Dimensionality-reduction techniques such as principal components have been deliberately avoided so as to retain the raw hyperspectral data for purposes of registration with other source images (e.g., SAR imagery), and not to lose the original physical meaning. MI measures the statistical dependence between two random variables and so can be used to evaluate the relative utility of each band to classification. Although entropy [6], [9] and mutual information [7], [14] have obvious potential for band selection, this has not been fully exploited in the past.

First, a novel strategy is proposed to estimate the mutual information using *a priori* knowledge, to reduce the reliance of MI estimation on availability of a reference map (i.e., the 'ground truth' map, in which each pixel is correctly labeled by its class). This will extend the application scope and adaptability of the MI-based method. Second, bands are selected not only on the basis of their associated MI values, but also their correlation (using the so-called 'complementary' level measure) with their neighboring bands. The proposed method avoids iterative searching, and so provides a low computational cost solution for time-critical applications. This method is also useful when it is difficult to obtain enough training data to validate selection of the classifiers' parameters, since it relies less on given ground truth.

II. AVIRIS 92AV3C Database

The public AVIRIS 92AV3C hyperspectral database has been researched extensively. The database is illustrative of the problem of hyperspectral image analysis to determine land use. It can be downloaded from `ftp://ftp.ecn.purdue.edu/biehl/MultiSpec/`. Although the AVIRIS sensor collects nominally 224 bands of data, 4 of these contain only zeros and so are discarded, leaving 220 bands in the 92AV3C dataset. At certain frequencies, the spectral images are known to be adversely affected by atmospheric water absorption. This affects some 20 bands. Each image is of size 145×145 pixels. The datacube was collected over a test site called Indian Pine in north-western Indiana, USA [15], [16].

The database is accompanied by a reference map, indicating partial ground truth, whereby pixels are labeled as belonging to one of 16 classes of vegetation. Not all pixels are so labeled (e.g., highway, rail track, etc.), presumably because they correspond to uninteresting regions or were too difficult to label. There are 10366 pixels from $145 \times 145 = 21025$ classified as belonging to one of the 16 classes (i.e., 49.3%), with the unlabeled pixels accounting for the remaining 50.7%.

III. Band selection using information theory

Let A be a random variable taking values in the set $\mathcal{A}$ with probability distributions $p(A)$. The entropy is defined by:

$$H(A) = -\sum_{A \in \mathcal{A}} p(A) \log p(A) \tag{1}$$

Some methods (e.g., [6], [9]) use the entropy directly as a criterion for band selection. In these methods, the entropy is used to estimate the information contained in each individual band and these are then ranked in order. Bands are selected by choosing those with the top-ranking entropy values.

However, the entropy measure is not always, as often assumed, consistent with the requirement of image classification. In fact, the method is only effective when the uncertainty calculated by the entropy encodes the most discriminatory information. From equation (1), it is seen that $H(A)$ is calculated with respect to the single variable A, without reference to any objective. Thus, the amount of information measured by the entropy lacks a point of reference or benchmark, and there is no guarantee that it matches the information content useful for target classification. To improve the approach, it is logical to extend the information measure to two variables: one for the spectral image itself and the other for the target (reference) image that is directly related to the classification objective. Mutual information (MI) provides a framework to measure the similarity between two random variables, and was introduced for band selection in [14], [7].

Given two random variables A and B with marginal probability distributions $p(A)$ and $p(B)$ and joint probability distribution $p(A, B)$, the mutual information is defined as:

$$I(A, B) = \sum_{A \in \mathcal{A}, B \in \mathcal{B}} p(A, B) \log \frac{p(A, B)}{p(A) \cdot p(B)} \tag{2}$$

It follows that MI is related to entropy as:

$$I(A, B) = H(A) + H(B) - H(A, B) \tag{3}$$

Here, $H(A)$ and $H(B)$ are the entropies of A and B, respectively, and $H(A, B)$ is their joint entropy.

Treating the spectral images and the corresponding reference map as random variables, MI can be used to estimate the dependency between them. Since the reference map implicitly defines the required classification result, the MI measures the relative utility of each spectral band to the classification objective. Using equation (3), the mutual information between each of the 220 spectral images (i.e., each band) and the corresponding reference map accompanying the 92AV3C dataset was calculated. The result is shown in Figure 1 as the solid curve. In the next section, an alternative MI measurement based on an estimated reference map (shown as a dashed curve in the figure) is presented.

Figure 1 shows that the spectral bands from approximately 5 to 35, 110 to 150 and 165 to 215 have higher MI values than other bands. The MI curve also reveals clearly the effect of atmospheric water absorption, giving the lowest MI values in bands 104–108 and 150–163 at precisely those frequencies where absorption occurs.

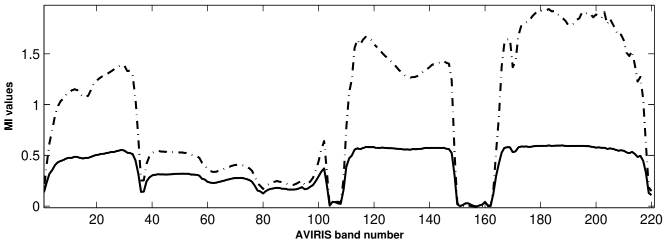

Fig. 1. Mutual information of AVIRIS bands 1 to 220 with respect to the actual reference map (solid line) and the estimated reference map as described in Section IV (dashed line).

It is instructive to compare this MI curve to selected examples of AVIRIS images shown in Figure 2(a). It can be seen that the bands most similar to the reference map of Figure 2(b) are those having higher values of MI[1]. For example, the images of the first row of Figure 2(a) (i.e., the spectral bands from 121 to 125) bear more obvious resemblance to the reference map than those in the second row (i.e., from 81 to 85). Their MI values are correspondingly much higher than those of the spectral bands in the second row. In this particular example, it is seen that the MI of a spectral band with respect to the reference map is consistent with visual observations regarding the relative importance of each spectral band to classification. Hence, it is reasonable to consider using MI to encode the relative utility of a spectral band, and to use this as the basis of band selection. Since the reference map may be prepared in different ways (e.g., according to the choice of pixels to label, the user's choice of false color, etc.), one should only use the intrinsic properties of spectral reflectance and not over-interpret the notion of similarity between the images in Figure 2(a) and the reference map in Figure 2(b).

A weaknesses of the straightforward MI method is that it relies on a reference map, which may not always be available. To improve the applicability of the method, a new band-selection scheme based on estimating a reference map is proposed, as now described.

IV. Band-selection scheme

Motivated by the discussion above, a band-selection scheme based on mutual information is considered. In this case, instead of calculating the MI based on the reference map, R, an estimated reference map, $\widehat{R}$, is used. This is assumed to be easy to obtain and to be a good estimate of R. In the proposed method, the estimated reference map is designed using available *a priori* knowledge about the spectral 'signature' of frequently-encountered materials. In the context of particular imaging applications, such knowledge is often encoded in a so-called spectral-signature library, such as the USGS Digital Spectral Libraries at `http://speclab.cr.usgs.gov`.

From the spectral-signature library, it is easy to identify an approximate wavelength range in which the signatures of in-

[1] The reader should note that the spectral reflectance values in Figure 2(a) have been transformed using a false color palette and then converted to gray-scale; the reference map of Figure 2(b) has had pixels from each of the 16 classes (plus unlabeled pixels) color-coded and also converted to gray-scale. Although it is hard to distinguish all 16 classes, these figures are adequate for gaining an impression of similarity of spectral images to the reference map.

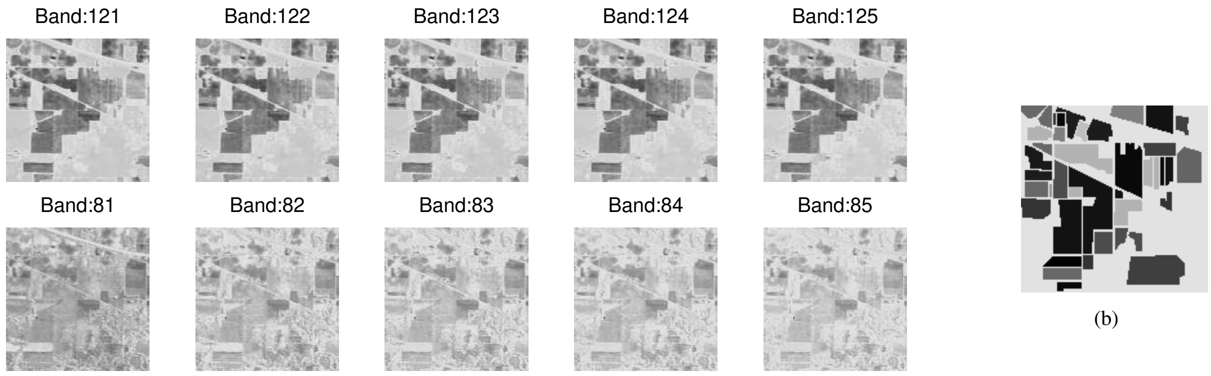

(b)

(a)

Fig. 2. (a) Examples of AVIRIS spectral images in informative bands 121–125 and less informative bands 81–85; and (b) the reference map.

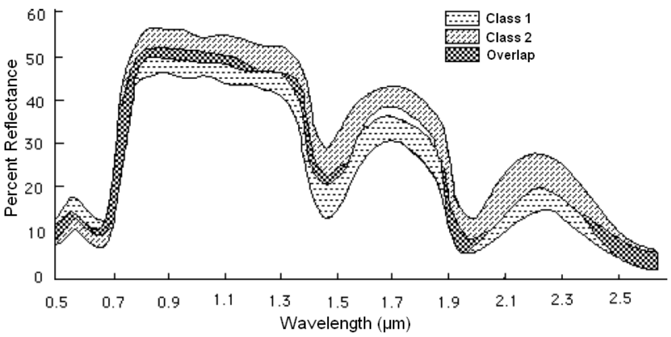

Fig. 3. Spectral responses for two classes of vegetation from [15].

terest are significantly differentiated. Such regions are referred to here as *key spectra*. For example, in the task of classifying two particular plants [15], their spectral reflectances can be obtained from laboratory measurements, shown in Figure 3. There is a frequency-dependent overlap between the two classes due to the natural variability of reflectance. Thus, the two classes are more separable in the region 1.6–1.8 μm than elsewhere. These relatively more informative bands correspond to the key spectra.

Generation of the estimated reference map from the available images exploits *a priori* knowledge, which can be achieved either through subjective evaluation or objective extraction. In the vegetation-dominant scene of AVIRIS 92AV3C, the task is to classify land use for agriculture. Domain experts can visually inspect example images (such as those illustrated in Figure 2(a)) to determine which are likely to be key spectra, useful in discriminating classes. For a multiclass problem, this can be implemented by observing (pairwise) the signatures in the library, selecting a suitable group of bands to cover the spectra with the biggest divergence, and averaging these. So, if $M_j \in \mathcal{S}$, $1 \leq j \leq J$ are images from the set of key spectra $\mathcal{S}$, then the estimated reference map $\widehat{R}$ is obtained as:

$$\widehat{R} = \frac{1}{J}\sum_{j=1}^{J} M_j \tag{4}$$

Here, bands 170 to 210 are averaged using equation (4). This was chosen as a region in which mutual information is high (see Fig. 1) and relatively constant. Other regions could have been chosen on the same basis. For a one class versus the rest problem, signatures of all the other materials can be averaged as 'background', and compared with the target signature. A more sophisticated scheme would be to generate the reference map based on some differentiable discriminating metric and greedy algorithm. The estimated reference map is unlabeled but this is immaterial as it is not used in classifier training.

To validate this estimation method, MI computed using on the estimated reference map, $I(M_i, \widehat{R})$, are compared with those obtained using the reference map accompanying the 92AV3c database, $I(M_i, R)$, $1 \leq i \leq 220$. This is shown in Figure 1, where the dashed curve depicts $I(M_i, \widehat{R})$ and the solid curve depicts $I(M_i, R)$. The overall shapes of the two curves are very similar, but the MI using the estimated reference map is consistently higher than that using the real reference map, because it is computed from means of spectral reflectances of all materials whereas the real map is computed from integer values arbitrarily coding 16 class labels. Nonetheless, the excellent agreement in overall shape means that the estimated reference map can be considered as a good alternative representation of the ideal classification objective.

With this new approach, MI-based band selection can now be extended to applications in which the reference map is difficult to obtain. Spectral bands i can be ranked according to their corresponding values of $I(M_i, \widehat{R})$ and a threshold applied to retain only the top-ranking bands. From equation (4), it can be seen that $\widehat{R}$ depends only on the data, and would automatically be updated with new data. In this way, the MI estimate is 'adaptive' to the image content as the investigated scene changes.

The above scheme works well when the threshold for selection is low. However, when fewer bands are to be selected (i.e., high threshold), this scheme may give too much preference to contiguous bands with high MI values but also high correlation between bands, indicating considerable redundancy. To avoid including redundant information, two ancillary selection parameters are introduced: rejection bandwidth B and complementary threshold η to control if the neighboring bands surrounding the top-ranked bands should be included besides considering their MI values. Here, B stands for a bandwidth centred at a selected band, indicating the bands within the bandwidth are highly correlated, and possibly redundant to the central one, and η evaluates how much complementary 'information' a neighboring band can supplement to the selected one. If significant change of 'information' has been detected for a neighboring band, this band should be included even if it may fall within the rejection bandwidth. Algorithm 1 shows the process in detail. First, define $d(n) = \Delta(MI(n) - MI(n-1))$. Let X be the number of bands to be selected. At some point in the selection process, let $\mathcal{S}$ be the set of selected bands and let $\mathcal{R}$ be the set of of remaining (unselected) bands. We initialise the process with $\mathcal{S} = \emptyset$ and $\mathcal{R} = \{1, 2, \ldots, 220\}$.

Algorithm 1 Band selection using rejection bandwidth and complementary threshold

while $|\mathcal{S}| < X$ **do**
  select band index $S = \arg\max_s MI(s)$ where $s \in \mathcal{S}$
  neighbor set $\mathcal{N} = \{n \mid n = S-(B+1), \ldots, S, \ldots, S+B\}$
  if $\max_n d(n) < \eta$ **then**
    $\mathcal{S} \leftarrow \mathcal{S} \cup S$ and $\mathcal{R} \leftarrow \mathcal{R} \setminus \mathcal{S} \setminus \mathcal{N}$
  else
    $\mathcal{S} \leftarrow \mathcal{S} \cup S$ and $\mathcal{R} \leftarrow \mathcal{R} \setminus \mathcal{S}$
  end if
end while

The particular values of B and η were determined by validation testing, and chosen as 5 and 0.5 respectively. Comparing the present approach with previous research [6], [7], [9], [14], the main differences are:

1) MI is used to measure the utility of each band.
2) The MI is calculated with respect to an estimated reference map rather than a pre-specified reference map.
3) Two parameters, B and η, are introduced to avoid selecting redundant neighboring bands.

V. Simulation results

Simulations of classification performance have been carried out to assess the proposed band-selection method on the hyperspectral dataset AVIRIS 92AV3C. Currently-popular support vector machines (SVMs) [17] were chosen as the classifiers in these simulations since previous studies of hyperspectral data classification have shown competitive performance with the best available algorithms [16], [4]. Although SVMs are used here, the proposed method is not limited to supervised algorithms. Other classification algorithms are also applicable since the MI metric is calculated directly from the data, without feedback from classifiers.

Half of the pixels from each class were randomly chosen for training, with the remaining 50% forming the test set on which performance was assessed (Table I). The class labels from the reference map accompanying the dataset were utilized for the supervised training. This was repeated 10 times to allow an estimate of the error in this sampling process. Since SVMs are inherently binary (two-class) classifiers, $C(16, 2) = 120$ one-against-one classifiers were used with subsequent majority voting to give a multi-class result. The kernel function used is an inhomogeneous polynomial with order 5. The penalty parameter C is tested between 10^{-3} and 10^{8} by a validation procedure using training data, and 10^{4} (which is the same for all classifiers) was chosen as a suitable compromise value.

TABLE I

Number of training and testing pixels in each class.

Class	Pixels in training set	Pixels in testing set
1. Alfalfa	27	27
2. Corn-notill	717	717
3. Corn-min	417	417
4. Corn	117	117
5. Grass/Pasture	249	248
6. Grass/Trees	374	373
7. Grass/pasture-mowed	13	13
8. Hay-windrowed	245	244
9. Oats	10	10
10. Soybeans-notill	484	484
11. Soybeans-min	1234	1234
12. Soybean-clean	307	307
13. Wheat	106	106
14. Woods	647	647
15. Bldg-Grass-Tree-Drives	190	190
16. Stone-steel towers	48	47

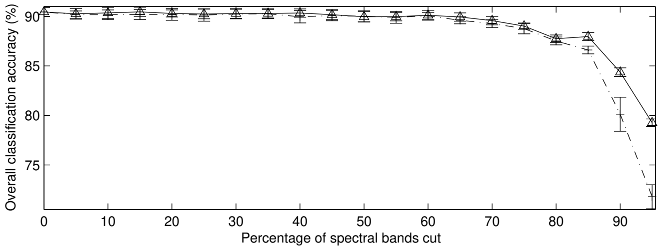

Fig. 4. Overall classification accuracy versus percentage of bands removed for band-selection scheme: solid line is the result using the estimated reference map, and dashed line is the result using the reference map supplied with the 92AV3C dataset. Here, 'overall' means all classes are treated together so that those with large numbers of pixels (see Table I) have greater effect. Error bars represent plus-and-minus one standard deviation over 10 simulation runs.

The main objective of band selection is to remove redundant spectral bands without degrading the classification accuracy too much. The simulations were designed to assess the change of classification accuracy as spectral bands are progressively removed according to Algorithm 1. Figure 4 shows the results for the two cases where the MI was calculated with respect to the reference map accompanying the 92AV3C dataset (dashed line) and where it was calculated with respect to the estimated reference map (solid line). Data points are at 5% increments corresponding to removal of 11 bands at each step.

The results depicted in Figure 4 indicate that:

1) The two curves are very similar when the percentage of bands cut is less than 80%, showing that the estimated reference map is a good approximation to the real one;
2) When this percentage is greater than 80%, the performance based on the estimated reference map is better.

The second result reveals an important issue on the bias of probability estimation. In the accompanying reference map, less than half of all pixels in the scene are labeled (see Section II), and among them 50% are used as the training set. This means that less than 25% of pixels are usable based on

TABLE II
COMPARISON OF PROPOSED METHOD WITH COMPETITOR ALGORITHMS ON THE BASIS OF OVERALL ACCURACY (%)

Bands Retained	Method			
	MI_est	SA	Entropy	Correlation
20	86.57	88.07	74.04	60.39
25	87.18	88.96	75.66	63.62
30	88.23	89.60	76.47	67.28
35	89.00	89.67	78.83	70.49
40	89.38	86.62	80.35	71.88
45	89.67	90.12	81.86	73.48
50	90.25	89.42	82.13	74.08
60	90.14	88.15	83.34	75.26
70	90.54	89.75	85.41	75.93
80	90.62	90.33	86.43	82.69

the accompanying reference map, which may lead to unreliable estimates of MI. On the other hand, all pixels are exploitable based on the estimated reference map. As more and more bands are selected, the two methods tend to give similar results because most of the essential bands are included.

As expected from Figure 1, the bands with lowest MI (which were therefore removed earliest) were those affected by atmospheric problems, e.g., from 104 to 108, 150 to 163 and 220. In previous research [16], [18], these are normally identified by the prior knowledge of experts. From a practical point of view, this approach is subjective and impractical in rapidly changing scenarios. In this sense, the proposed method provides a better way to discard these bands.

VI. COMPARISON WITH OTHER APPROACHES

To validate the new proposed method using the estimated reference map (MI_est) more fully, performance has been compared against some other representative band-selection algorithms on the 92AV3C dataset, namely steepest ascent (SA) searching [3], [4], entropy [6], [9] and the very well-known Pearson correlation coefficient. The comparison is carried out for 20 to 80 bands retained, since the SA algorithm is very computationally-expensive and for the 92AV3C dataset the classification accuracy shows little change beyond 80 selected. bands.

Table II shows the results. To account for local maxima, the SA algorithm was run six times with random initialisations, and the best result chosen for tabulation. The proposed method easily outperformed the methods based on entropy and correlation. It is also competitive with the searching-based SA algorithm, with slightly lower accuracy at band selection numbers 20 to 35. Considering that the SA algorithm involves computationally-expensive iterative search and Jeffries-Matusita distance evaluation, the proposed method provides a useful alternative with much lower computational cost, suiting it to time-critical applications. It was also noticed in implementing the SA algorithm that when the number of selected bands increased above about 40, the covariance matrices used to calculate the Bhattacharyya distance tended to become singular, producing early cessation of the search and a sub-optimal result. We have tried to mitigate this problem by adding a regularisation factor, but avoidance of this problem is another reason to favor the new method.

VII. CONCLUSION

A band-selection method for hyperspectral image analysis based on estimation of mutual information has been described. Given the relationship between MI and Bayes error, maximizing MI is analogous to the idea of maximizing separability that many other methods employ. Experiments on the AVIRIS dataset show that the proposed method outperformed or was competitive with state-of-the-art methods, but having the advantages of easy implementation and low computational cost. Future work will concentrate on improved methods for estimating the reference map from prior knowledge.

REFERENCES

[1] AVIRIS, “Airborne visible/infrared imaging spectrometer.” [Online]. Available: http://aviris.jpl.nasa.gov/.
[2] G. Hughes, “On the mean accuracy of statistical pattern recognizers,” *IEEE Transactions on Information Theory*, vol. 14, pp. 55–63, 1968.
[3] S. Serpico and L. Bruzzone, “A new search algorithm for feature selection in hyperspectral remote sensing images,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 7, pp. 1360–1367, 2001.
[4] F. Melgani and L. Bruzzone, “Classification of hyperspectral remote sensing images with support vector machines,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 8, pp. 1778–1790, 2004.
[5] G. Petrie, P. Heasler, and T. Warner, “Optimal band selection strategies for hyperspectral data sets,” in *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, Seattle, WA, 1998, pp. 1582–1584.
[6] P. Groves and P. Bajcsy, “Methodology for hyperspectral band and classification model selection,” in *IEEE Workshop on Advances in Techniques for Analysis of Remotely Sensed Data*, Greenbelt, MD, 2003, pp. 120–128.
[7] C.-I. Chang, D. Qian, T.-L. Sun, and M. Althouse, “A joint band prioritization and band-decorrelation approach to band selection for hyperspectral image classification,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 37, no. 6, pp. 2631–2641, 1999.
[8] M. Velez-Reyes and L. Jimenez, “Subset selection analysis for the reduction of hyperspectral imagery,” in *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, vol. 3, Seattle, WA, 1998, pp. 1577–1581.
[9] P. Bajcsy and P. Groves, “Methodology for hyperspectral band selection,” *Photogrammetric Engineering and Remote Sensing Journal*, vol. 70, no. 7, pp. 793–802, 2004.
[10] J. Price, “Spectral band selection for visible-near infrared remote sensing: spectral-spatial resolution tradeoffs,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 35, no. 5, pp. 1277–1285, 1997.
[11] N. Keshava, “Best bands selection for detection in hyperspectral processing,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP’01*, vol. 5, Salt lake City, UT, 2001, pp. 3149–3152.
[12] D. Qian, “Band selection and its impact on target detection and classification in hyperspectral image analysis,” in *IEEE Workshop on Advances in Techniques for Analysis of Remotely Sensed Data*, Greenbelt, MD, 2003, pp. 374–377.
[13] S. Kaewpijit, J. L. Moigne, and T. El-Ghazawi, “Automatic reduction of hyperspectral imagery using wavelet spectral analysis,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 4, pp. 863–871, 2003.
[14] C. Conesea and F. Masellia, “Selection of optimum bands from TM scenes through mutual information analysis,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 48, no. 3, pp. 2–11, 1993.
[15] D. Landgrebe, “On information extraction principles for hyperspectral data: A white paper,” Purdue University, West Lafayette, IN, Technical Report, School of Electrical and Computer Engineering, 1997. [Online]. Available: http://dynamo.ecn.purdue.edu/~landgreb/whitepaper.pdf.
[16] J. Gualtieri and R. Cromp, “Support vector machines for hyperspectral remote sensing classification,” in *Proceedings of the 27th AIPR Workshop on Advances in Computer Assisted Recognition*, Washington DC, 1998, pp. 121–132.
[17] C. Cortes and V. N. Vapnik, “Support-vector networks,” *Machine Learning*, vol. 20, no. 3, pp. 1–25, 1995.
[18] C. A. Shah, P. Watanachaturaporn, M. K. Arora, and P. K. Varshney, “Some recent results on hyperspectral image classification,” in *Proceedings of IEEE Workshop on Advances in Techniques for Analysis of Remotely Sensed Data*, Greenbelt, MD, 2003, pp. 346–353.

List of Figures