

INVESTIGATION OF THE RANDOM FOREST FRAMEWORK FOR CLASSIFICATION OF HYPERSPECTRAL DATA

JiSoo Ham¹, Yangchi Chen¹, Melba M. Crawford¹, and Joydeep Ghosh²

¹Center for Space Research, 3925 W. Braker Lane, Austin, TX 78759

²Department of Electrical and Computer Engineering

The University of Texas at Austin

{jham@csr.utexas.edu, yanji@csr.utexas.edu, crawford@csr.utexas.edu, ghosh@ece.utexas.edu}

Abstract – Statistical classification of hyperspectral data is challenging because the inputs are high in dimension and represent multiple classes that are sometimes quite mixed, while the amount and quality of ground truth in the form of labeled data is typically limited. The resulting classifiers are often unstable and have poor generalization. This paper investigates two approaches based on the concept of random forests of classifiers implemented within a binary hierarchical multi-classifier system, with the goal of achieving improved generalization of the classifier in analysis of hyperspectral data, particularly when the quantity of training data is limited. A new classifier is proposed which incorporates bagging of training samples and adaptive random subspace feature selection within a Binary Hierarchical Classifier (BHC), such that the number of features that is selected at each node of the tree is dependent on the quantity of associated training data. Results are compared to a random forest implementation based on the CART framework. For both methods, classification results obtained from experiments on data acquired by the NASA airborne AVIRIS instrument over the Kennedy Space Center, Florida, and by Hyperion on the NASA EO-1 satellite over the Okavango Delta of Botswana are superior to those from the original best basis BHC algorithm and a random subspace extension of the BHC.

Keywords – binary hierarchical classifier, classification, CART, Hyperion, hyperspectral, Okavango Delta, random forests, random subspace feature selection.

I. INTRODUCTION

The increasing availability of data from hyperspectral sensors, particularly with the launch of the Hyperion instrument on the NASA EO-1 satellite, has generated tremendous interest in the remote sensing community. These instruments characterize spectral signatures with much greater detail than traditional multispectral sensors and thereby can potentially provide improved discrimination of targets [1]. However, hyperspectral data also present difficult challenges for supervised statistical classification, where labeled training data are used to estimate the parameters of the label-conditional probability density

functions [2]. The dimensionality of the data is high (~200), there are often tens of classes C , and the quantity of training data is often small. Sample statistics of training data may also not be representative of the true probability distributions of the individual class signatures, particularly for remote, inaccessible areas where training data are logistically difficult and expensive to acquire. Generalization of the resulting classifiers is often poor, thereby resulting in poor quality mapping over extended areas.

Various approaches have been investigated to mitigate the impact of *small sample sizes and high dimensionality*, which are inherently coupled issues since the adequacy of a data sample depends on the data dimensionality, among other factors [3]. For example, regularization methods try to stabilize the covariance matrix by weighting the sample covariance matrix and a pooled covariance matrix or by shrinking the sample covariance matrix toward the identity matrix [4]. While this may reduce the variance of the parameter estimates, the bias of the estimates can increase dramatically. Alternatively, the input space can be transformed into a reduced feature space via feature selection [5] or feature extraction. Although these two approaches reduce the effect of the high dimensionality problem, feature selection methods are often trapped in a local optimal feature subset, while feature extraction methods lose the interpretability of the original features. Another way of dealing with a small training set is to augment it with unlabeled data and then use semi-supervised learning techniques. These methods have been shown to enhance supervised classification [6-7]. However, convergence of the updating scheme can be problematic, and it is affected by selection of the initial training samples and by outliers.

In analysis of hyperspectral data, Lee and Landgrebe proposed methods for feature extraction based on decision boundaries that maximize separation of data in multiple two-class problems [8]. These decision boundary feature extraction (DBFE) methods are often effective for two-class problems, but do not exploit correlation between sequential bands. Jia and Richards developed the Segmented Principal Components Transformation (SPCT) whereby the original bands are grouped into subsets of highly correlated adjacent bands to which the K-L transform is applied. The most significant principal components are then selected from each subset to yield a feature vector with reduced dimension [9]. The approach treats inter-band correlation globally and does not guarantee good discrimination capability because the PCT preserves variance in the data rather than maximizing discrimination between classes. Kumar et al. investigated band combining techniques, motivated by best-basis functions, as a means of feature extraction in a pairwise classifier framework [10]. Adjacent bands are selected for merging (alt. splitting) in a bottom-up (alt. top down) fashion using the product of a correlation measure and a Fisher discriminant. Morgan et al. [11] suggested a similar correlation-based band combining approach, in conjunction with a covariance shrinkage method, for both a top-down and bottom up hierarchical classifier to ameliorate the small training data problem.

The theory and practice of classifier ensembles also provide ways of alleviating sample size and high dimensionality concerns [12]. Bagging involves bootstrapped sampling of the original data, generating a classifier specific to each sample, and then averaging the classifier outputs [13]. This method takes

advantage of data reuse, but when the training data set in the (sub-)sample is very small, the potential for improved diversity and reduced impact of outliers is offset by degradation in individual classifier performance [14]. Boosting also combines weak individual classifiers to develop an improved classifier, but by re-weighting training data to increase sensitivity to incorrectly classified training observations. While boosting can improve performance for large training samples, it is not useful for small sample problems, particularly in the presence of outliers. When the input space is large, random subspace (RS) feature selection can potentially provide improved classifier diversity, while stabilizing parameter estimates, by randomly reducing the number of inputs to each classifier in the ensemble and constructing multiple classifiers in the resulting random input space [15-16]. The method is potentially attractive for problems with redundant input features (e.g. hyperspectral data) and when outliers exist in the training data. Recently, approaches referred to as “random forests of classifiers” have been proposed. These involve developing multiple trees from randomly sampled subspaces of input features, then combining the resulting outputs via voting or a maximum *a posteriori* rule [17]. These methods typically achieve superior generalization for small training samples, but are computationally intensive.

Land cover classification problems usually involve a large number of classes, i.e. the *output space* is large. Output decomposition using binary classifiers in a multi-classifier framework has been shown to be more successful than traditional 1-of- C classifiers for many problems involving large output spaces [18]. Decomposition methods using pairwise classifiers, error correcting output codes (ECOC) [19], and binary decision trees have all been investigated in this context (see [11] for an overview). Pairwise classifiers develop a separate classifier for each pair of classes, thereby resulting in $O(C^2)$ classifiers which must be combined to determine the final class label. These methods often yield simple classifiers with excellent discrimination for specific pairs, but are generally inefficient for problems with a large number of output classes. In the ECOC, a C -class problem is decomposed into \bar{C} binary problems, whereby the original class is then encoded into a \bar{C} binary vector of a coding matrix. It has been shown that the ECOC method yields robust, stable classifiers. However, since the code matrix design is not based on the characteristics of the classes it represents, interpretability of the classifier is limited.

Binary trees, which often provide an attractive approach for decomposing large output space problems, can be constructed using a variety of splitting functions involving single or multiple features and output classes. To address the high dimensional output problem while exploiting the affinity for spectrally similar classes, Kumar et al. proposed a Binary Hierarchical Classifier (BHC) [20] to decompose a $(C > 2)$ -class problem into a binary hierarchy of $(C - 1)$ simpler 2-class problems that can be solved using a corresponding hierarchy of classifiers, each based on a simple linear discriminant. The method was extended by Morgan et al. [11] for small training samples using an adaptive best basis BHC, which exploits the class specific correlation structure between sequential bands of hyperspectral data and utilizes an adaptive regularization approach to stabilize covariance estimates. An adaptive random subspace

feature selection approach was also investigated within the BHC framework (RS-BHC) as a means of improving classifier performance when the number of training samples is extremely small [21].

In this paper, we investigate a random forest of binary classifiers as a means of increasing diversity of hierarchical classifiers. We evaluate the results obtained for trees produced by our BHC classifier and the original CART-based random forest method [17]. For the BHC, the goal is to exploit the advantages of natural class affinity while improving generalization in classification of hyperspectral data when the number of training samples is small. The CART-based approach is not directly affected by small sample size statistics and potentially provides greater diversity within the forest, but typically produces trees of enormous size if the output space is large. The paper is organized as follows: the best basis (BB-BHC), random subspace (RS-BHC), and random forest (RF-BHC) implementations of the BHC method and the CART-based framework (RF-CART) are all described in Section II; classification results using the random forest approaches obtained for data acquired by AVIRIS (Airborne Visible/Infrared Imaging Spectrometer) over the Kennedy Space Center, Florida, and EO-1 Hyperion over the Okavango Delta, Botswana, are presented in Section III and compared to those obtained from the BB-BHC, RS-BHC. Results from all the methods are evaluated, and new directions for future work are suggested in Section IV.

II. RANDOM FOREST BINARY HIERARCHICAL CLASSIFICATION METHOD

The top-down Binary Hierarchical Classifier (BHC) framework recursively decomposes a C -class problem into $C-1$ two-(meta)class problems via a deterministic simulated annealing method [20]. The root classifier tries to optimally partition the original set of classes into two disjoint meta-classes while simultaneously determining the Fisher discriminant that separates these two subsets. This procedure is recursed, i.e., the meta-class Ω_n at node n is partitioned into two meta-classes $(\Omega_{2n}, \Omega_{2n+1})$, until the original C classes are obtained at the leaves. The tree structure, as shown in Fig. 1, allows the more natural, easier discriminations to be accomplished earlier. Fewer classes are involved in the partitioning at lower levels of the BHC hierarchy. Thus, while the classification task typically becomes simpler, the number of relevant training samples also decreases. The BB-BHC ameliorates this effect by utilizing an ancestor covariance matrix while exploiting the inter-band serial correlation through an adaptive, class dependent, band aggregation process [11]. A band combining step is performed on highly correlated, spectrally adjacent bands prior to the partitioning of meta-classes, thereby reducing the number of inputs relative to the number of training data points. Bands are aggregated until a user defined ratio, R , between the number of training samples for the respective (meta)classes and input dimension is achieved. Typically, R is selected to be at least 5.

The BHC method is extended in the RS-BHC approach by utilizing the random subspace method as a post-processing stage to tree construction, with the goal of reducing the number of inputs while refining decision boundaries [22]. The BB-BHC method is used to first construct the hierarchy, then random subspace sampling is performed at each node of the tree where the criterion for R is not satisfied. For each

(meta) class m with n_m vector-valued observations, $\mathbf{X}_m = (\mathbf{X}_1, \dots, \mathbf{X}_{n_m})$, a subset of elements of $\mathbf{X}_i = (x_{i1}, \dots, x_{ik})$ with dimension $p_m = n_m/R < k$ is then randomly selected from the k -dimensional set of features. The resulting modified training set $\mathbf{X}_m^r = (\mathbf{X}_1^r, \dots, \mathbf{X}_{n_m}^r)$ consists of observation vectors, $\mathbf{X}_i^r = (x_{i1}^r, \dots, x_{ip}^r)$, where the same subset of features is selected for each element $\mathbf{X}_i^r \in \mathbf{X}^r, (i = 1, \dots, n_m)$. The number of random subspaces selected at each such node is $N_s = (k/p_m) \times F$, where the value of F is a user supplied input. A discriminant vector is constructed for each random subspace, and the N_s vectors are combined at each node of the hierarchy via majority voting. Our empirical evidence indicates that good results are typically achieved for $2 < F < 4$, which provides adequate coverage of the feature space. Improvement in classification accuracy is not significant for $F > 4$.

The random forest implementation of the BHC (RF-BHC) extends the RS-BHC by incorporating random subspace feature selection in the actual development of the tree. This is particularly advantageous as random subspace sampling is performed by the RS-BHC only at nodes where the ratio, R , is not exceeded. Thus, sub-sampling of the input features typically occurs only at lower levels of the tree, thereby limiting diversity. For moderate sized training samples, bagging can increase diversity of the multiclassifier system, so a bootstrap sample of observations is selected for each tree in the RF-BHC. At each meta-class node m , a random subspace of features of dimension $p_m = \min(p_m, N_f)$ is selected to determine the decision boundary for the classifier at that node, where N_f is a user selected input. To guarantee greater diversity, we choose $N_f \ll k$. The tree is then developed using the resulting set of features selected at each node. The process is repeated to grow a forest of identically, independently distributed random vectors associated with the individual trees.

The fundamental difference between BHC and other decision trees is that the former focuses on decomposing the output space; partitioning of the input space occurs as a consequence. Both RF-BHC and RF-CART use the random forest ensemble method to increase the diversity of each base learning module, then combine results of the individual modules (trees). While Breiman's CART-based random forest follows a typical binary divide-and-conquer hierarchical scheme, it differs from the BHC in the base learning module. The BHC uses the GAMLS [23] algorithm to split each node into meta-classes which are separated by the maximum Fisher distance. Using a sequence of binary tests, CART seeks the split that maximizes the reduction of the impurity of the parent nodes and its two child nodes as measured by the Gini index [24]. The most discriminating feature is selected to perform the split. Used in the random forest context, a random subspace of the original k features is selected at each node of the tree, and the most discriminating feature is then selected. Further, unlike the actual CART method, the RF-CART approach

does not perform pruning of nodes as pruning reduces diversity of trees in the forest. Analogous to the RF-BHC, each tree is grown using a bootstrap sample of the training set.

III. RESULTS

Hyperspectral data from two sources were analyzed in this paper:

1). *Kennedy Space Center, Florida*: The NASA AVIRIS (Airborne Visible/Infrared Imaging Spectrometer) instrument acquired data over the Kennedy Space Center (KSC), Florida, on March 23, 1996. AVIRIS acquires data in 224 bands of 10 nm width with center wavelengths from 400 - 2500 nm. The KSC data, acquired from an altitude of approximately 20 km, have a spatial resolution of 18 m. After removing water absorption and low SNR bands, 176 bands were used for the analysis. Training data were selected using land cover maps derived from color infrared photography provided by the Kennedy Space Center and Landsat Thematic Mapper (TM) imagery. The vegetation classification scheme was developed by KSC personnel in an effort to define functional types that are discernable at the spatial resolution of Landsat and these AVIRIS data. Discrimination of land cover for this environment is difficult due to the similarity of spectral signatures for certain vegetation types. For classification purposes, 13 classes representing the various land cover types that occur in this environment were defined for the site. (Table 1). Classes 4, and 6 represent mixed classes.

2) *Okavango Delta, Botswana*: The NASA EO-1 satellite acquired a sequence of data over the Okavango Delta, Botswana in 2001-2004. The Hyperion sensor on EO-1 acquires data at 30 m pixel resolution over a 7.7 km strip in 242 bands covering the 400-2500 nm portion of the spectrum in 10 nm windows. Preprocessing of the data was performed by the UT Center for Space Research to mitigate the effects of bad detectors, inter-detector miscalibration, and intermittent anomalies. Uncalibrated and noisy bands that cover water absorption features were removed, and the remaining 145 bands were included as candidate features: [10-55, 82-97, 102-119, 134-164, 187-220]. The data analyzed in this study, acquired May 31, 2001, consist of observations from 14 identified classes representing the land cover types in seasonal swamps, occasional swamps, and drier woodlands located in the distal portion of the Delta [22]. These classes were chosen to reflect the impact of flooding on vegetation in the study area. The class names and corresponding numbers of ground truth observations used in the experiments are listed in Table 2. Classes 3 and 4 are both floodplain grasses that are seasonally inundated, but differ in their hydroperiod (the amount of time inundated). Classes 9, 10, and 11 represent different mixtures of acacia woodlands, shrublands, and grasslands and are named according to the dominant class. Training data were selected manually using a combination of GPS located vegetation surveys, aerial photography from the Aquarap (2000) project, and 2.6 m resolution IKONOS multispectral imagery. The class priors for both data sets, as indicated by the labeled data, are only moderately skewed. For simplicity, we assume the class priors to be equal while developing the BHC classifier. This assumption shall be reconsidered later.

For both data sets, ten randomly sampled partitions of the training data were sub-sampled such that 75% of the original data were used for training and 25% for testing. In order to investigate the impact of the quantity of training data on classifier performance, these training data were then sub-sampled to obtain ten samples comprised of 50%, 30%, and 15% of the original training data. All classifiers were evaluated using the ten test samples composed of 25% of the original training data.

Experiments were performed using the BB-BHC, RS-BHC, RF-BHC, and RF-CART. Although authors recommend various values for the dimension of the random subspace and the number of trees in a random forest, there do not appear to have been any systematic studies of the issue to date. In the results reported here, the ratio R was set at 5, and the value of F was 4 for the RS-BHC method. In our experiments, the dimension of the random subspace was determined adaptively in the BHC, but was always selected such that the value of R was at least 5. For the RF-BHC, the value of N_f was selected to be 20. In order to have somewhat comparable inputs, 20 input features were randomly selected in the RF-CART method. 100 trees were grown for each experiment as our sensitivity studies showed that larger forests did not provide improved results for these data sets.

Results: Original Training and Test Areas

Kennedy Space Center, Florida: The true color image shown in Fig. 2a, along with the classification results obtained from the RF-BHC in Fig. 2b, shows the spatial distribution of classes and training sites over the 614 x 512 pixel study area. Average classification accuracies for test data and associated standard deviations for the 10 experiments conducted with each classifier are plotted in Fig. 3. The overall trends in accuracies relative to the quantity of training data are similar for all methods when applied to the test data set. At the 75% sampling rate, the accuracies of all methods are nearly the same, although the RF-BHC yields somewhat higher accuracies than the other methods. The results obtained by the BB-BHC and RF-CART methods are very similar over all sampling rates, with RF-CART yielding slightly lower accuracies. These methods consistently produced the lowest overall average accuracies. The RS-BHC yielded approximately the same accuracies as the BB-BHC at the 75% sampling rate, but improved relative to the BB-BHC and the RF-CART approach at lower sampling rates. For the BHC-based methods, this appears to demonstrate the value of reduced redundancy in the input space and improvements achieved by better tuning of the decision boundaries, even though the tree structure is identical to the BB-BHC and random sampling of the feature space is not required until lower levels of the tree (particularly for the higher training data fractions). Results were also obtained using the original RS-BHC and best basis aggregated data. The BB-RS-BHC consistently yielded slightly lower accuracies than the RS-BHC because of the reduced diversity of trees, but results were not statistically different and are not reported here. The overall average accuracy of the RF-BHC is consistently the highest and improves relative to other BHC methods and RF-CART as the fraction of training data is reduced.

The RF-BHC is also the most stable method over all training fractions, as measured by the standard deviation of the accuracies. The standard deviations of the accuracies of the random subspace based methods appear to benefit from the diversity of the input space. The standard deviation of the accuracies obtained by the BB-BHC increased dramatically at the 30% sampling rate because it was necessary to aggregate a large number of bands to satisfy the ratio R . The problem, which is manifested both in the tree building and in the decision boundary of the BB-BHC, is offset in the determination of the RS-BHC decision boundary. It should be noted that although the standard deviation of BB-BHC decreases at the 15% sampling rate, the associated average classification accuracy is also poor, further demonstrating it is uniformly inferior at low sampling rates. The reduced accuracy of RF-CART at low sampling rates, relative to the BHC based random forest methods, is attributed to the value of the inherent exploitation of class affinities by the BHC approaches. Further, although the standard deviation of the accuracies for the RF-CART approach is low for high sampling rates, it increases consistently as the sampling rate of the training data is reduced, likely because the discrimination capability of the single best feature within a small random sample of inputs may be quite variable. Further, the benefits of bagging the training sample occur at the higher sampling rates for both the RF-BHC and RF-CART methods.

Okavango Delta, Botswana: The RGB image in Fig. 4a and the classification results obtained by the RF-BHC in Fig. 4b show that the spatial distribution of classes is extremely complex over this 256 x 1476 pixel area. Using the same random sampling strategy as for the KSC data, results were obtained at each percentage for all four classifiers. Plots of classification accuracies at the various sampling rates are shown in Fig. 5. The overall trends in accuracies relative to the fraction of training data are similar those of the KSC AVIRIS test data. Among the classifiers, RF-BHC yielded the highest classification accuracy on all training sample fractions, and performance degraded only slightly at lower sampling rates. The standard deviations of the accuracies obtained using the RF-BHC are low and remain nearly constant over the various sampling rates. At the 30% sampling rate, the standard deviations of the accuracies yielded by the BB-BHC and RS-BHC are substantially higher. For the BB-BHC, this again appears to be due to the amount of band aggregation required to achieve the ratio R . Unlike the KSC case, the RS-BHC is apparently unable to mitigate problems associated with band aggregation during the tree construction phase for the Okavango data.

The difference in results produced by the RF-BHC and RF-CART methods was unexpected since both utilize an ensemble of 100 trees to build a stronger classifier. Previous research by Tumer and Ghosh [14] indicated that the accuracy of an ensemble method relies on the diversity of the base classifier. To investigate the performance of the individual trees, we further analyzed the performance of both random forest methods at the 75% sampling rate. The average accuracy over the set of individual trees developed by the RF-BHC is 89.2%, and the standard deviation is 1.3. The overall accuracy for the RF-BHC, which is determined by simple voting, increased by 5.7% to 94.9%. For the RF-CART method, the average accuracy obtained using 100 trees is 84.2%, with standard deviation 1.3. The ensemble of these 100 trees,

using simple voting utilized in the original Brieman random forest, yielded a 7.8% increase to 92%. For this type of classification problem, it appears that the BHC is a better base classifier than CART, although CART realizes substantial improvement when trees are combined.

Generalization to Spatially Disjoint Areas.

Traditionally, the training and test data are spatially co-located and can thus be assumed to be samples from the same distribution. In practice, however, it is also useful to estimate how a classifier will perform in areas that are somewhat different, in order to indicate how much additional data labeling and retraining is needed to make the model applicable to much larger areas. With this goal in mind, a “spatially disjoint” test set was also acquired from a geographically separate location at the Botswana site and used to evaluate the classifiers developed previously.

These spatially disjoint data have somewhat different characteristics from the training/test data, so the performance of all classifiers is reduced, as expected. Still, as with the test data, the BB-BHC yielded the lowest overall average accuracy at all sampling rates. The incremental improvement in average accuracy achieved by the random subspace method increases with reduced sampling rates, but is not statistically significant as the standard deviations of the accuracies also increase substantially with lower sampling rates. The RF-BHC implementation and the RF-CART method yielded higher accuracies for the spatially disjoint test data at all sampling rates than both the BB-BHC and the RS-BHC, thereby demonstrating the greater generalization of these approaches. Similar to results from the test data, the RF-BHC consistently produced the highest overall average accuracies for the spatially disjoint test set, indicating the value of exploiting class affinity, coupled with the increased diversity of trees achieved by forcing random sampling of the input space at all nodes. The RF-CART method also achieved good generalization, as indicated by its performance on these spatially disjoint test data, although results for the original test set were inferior to the other methods. This is attributed both to the diversity that it achieves and its reduced dependence on the training sample statistics. Similar to the test data, the performance of both the RF-BHC and RF-CART methods was further investigated for the 100 trees obtained from the Hyperion spatially disjoint test set at the 75% sampling rate. The average classification accuracy over the set of individual RF-BHC trees is 68.2%, and the standard deviation is 2.9. The ensemble random forest result using simple voting is 75.2%, an increase of 7%. For RF-CART the values are 60.8% and 2.4, respectively. The classification accuracy increased by 9.6% to 70.4% when the 100 trees were combined using simple voting.

Since RF-BHC and RF-CART use the same random forest framework, their differences lie both in the tree construction and the underlying classifier. For the remotely sensed data in this study, the BHC exploits class affinity, while the good performance of the CART-like method on the spatially disjoint test set suggests that it provides more diversity. To further investigate this issue, we calculated the entropy, a non-pairwise diversity measure [12], of trees obtained from both RF-BHC and RF-CART (Table 3). The results indicate that RF-CART method produces more diverse trees than RF-BHC at all four sampling rates.

RF-CART achieves an 11.2% increase in accuracies via the ensemble, while RF-BHC results improve only 7.7%, thereby reinforcing the idea that ensemble methods benefit more from combining diverse classifiers. Further, as the sampling rate increases, the diversity of RF-BHC trees decreases. Under the same situation, however, the diversity of the RF-CART forest remains comparatively consistent. This means that the RF-BHC inputs become more homogeneous as the number of samples increases, while RF-CART does not follow the same trend. Overall, the advantages of an ensemble approach are clear as the RS-BHC used only one tree structure rather than an ensemble of potentially different trees, which significantly reduced the generalization of its classification accuracies on the spatially disjoint test set.

The differences between the overall accuracies for the spatially disjoint test set and those for the test set are quite remarkable. As noted earlier, the test data are spatially co-located with the training data, whereas the spatially disjoint test set is not. Clearly, either the class priors or the class conditional feature distributions (or both) are substantially different, at least for some classes in the more remote area. This motivated us to further investigate class specific results. Class dependent accuracies for the Hyperion test and spatially disjoint test sets are provided in Fig. 6, and the detailed confusion matrix for the RF-BHC is contained in Table 4. Results in Table 2 indicate that the priors were indeed somewhat different. In particular, there were relatively more samples of Classes 2 and 11, and less of Classes 1 and 9. However, while false negative errors increase for Classes 2 and 11, there is no overall clear trend. For example, classification accuracies for Class 1 (water) which is spectrally quite distinct, are unaffected by the change of priors. Moreover, several class accuracies are now much lower than 80%, while others are almost unaffected. This leads us to believe that change in class-conditional distributions in certain classes that are spectrally quite similar is the main cause of the marked degradation in their classification accuracies. In particular, the overall classification accuracies of RS-BHC, RF-BHC and RF-CART methods are strongly influenced by the performance of Classes 2 and 11. Class 2, hippo grass, which grows within the river channels, has a small training sample and is spectrally similar to water as many pixels are mixed with water. Class 11, acacia grasslands, is a mixed class that is most often confused with other grasses or acacia shrubs, which is also a mixed class.

Using Fig. 6(b), we can also compare the class specific accuracies for the spatially disjoint test set for the RS-BHC, RF-BHC and RF-CART approaches. Consistent with the overall accuracies, the performance of the RF-BHC is generally better than RF-CART method at both the 75% and 15% sampling rates. Similarly, the RS-BHC yields consistently lower accuracies, particularly for Classes 2 and 11. Although higher classification accuracies were achieved for Class 2 by RF-CART than the two BHC methods, it is not statistically significant as the standard deviation of the average sample accuracy is more than 12.

In comparing the overall computational requirements of the BHC-based and RF-CART methods, there are several tradeoffs. BHC-based trees always solve C-1 binary problems. At the 75% sampling rate, the average CART decision tree for Botswana Hyperion data contained 326 nodes (std. dev.= 16). For this 14 class problem, the BHC tree had only 27 (1+13*2) nodes. For the same experiment, the CPU time

for the RF-CART method was 8m 42s, while it was 1h 4m 4s for the RF-BHC. Both experiments were performed on a 3GHz Pentium 4 CPU machine. The RF-BHC required more CPU time than the RF-CART method because GAMLSS is a deterministic simulated annealing algorithm. It should be noted that while neither algorithm was coded as an operational method, average timing results reflect their relative computational requirements.

IV. CONCLUSIONS AND FUTURE WORK

The primary purpose of the study was to investigate the performance of random feature subset selection methods in terms of generalization. The secondary goal was to investigate the performance of the methods when applied to data acquired by Hyperion data, which have low SNR. The performance of an implementation which focused on tuning decision boundaries of the BHC and that of two random forest approaches was investigated. Classification accuracies achieved by ensemble methods rely heavily on achieving diversity within the ensemble. The conflicting effects of improved SNR and reduced spectral resolution from band aggregation appear to be positively complemented by the improved diversity achieved by the RS-BHC through random sampling of the original features. We also noted that the change in classification accuracies achieved by using a forest rather than a single tree indicates that the RF-CART method actually achieves greater incremental benefit from the ensemble than the RF-BHC. Thus, the ensemble both exploits the greater diversity provided by the single feature splits and mitigates the potential impact of selecting features that are redundant or have poor discrimination capability.

A critical characteristic of the BHC is that it exploits the natural groupings of similar classes, which often occur in remotely sensed data acquired over natural landscapes. This provides a natural hierarchy which is often well handled by the simple Fisher discriminant. The random forest methods all yielded superior results for both test and spatially disjoint test data at our two study sites, with the improvement being greater for the spatially disjoint test set, thereby indicating improved generalization to extended areas. For these data, RF-BHC produced stable results over all sampling rates. Additional study is required to better characterize this issue. In this context, elimination of irrelevant and possibly redundant input features should also be considered in the RF-BHC. Other classifiers, such as the ECOC and SVM, should also be investigated within the RF-BHC framework. Overall, the RF-BHC methods appear to be quite promising in terms of generalization, but should be applied to many more data sets with different characteristics in order to better assess their overall performance. Also, much work remains to be done on determining how to improve performance on extended areas represented by the spatially disjoint data set, especially since both the class mixtures and class conditional spectral properties can change in such situations. If this problem can be solved, then one can more confidently label much larger regions than those directly described by the available labeled data. For mixed classes, the issue may be mitigated in some cases by determining relative abundances of component classes via unmixing of hyperspectral data, if

representative signatures of pure classes can be obtained [25, 26]. Approaches for representing spatially nonstationary spectral signatures may also be appropriate.

ACKNOWLEDGEMENTS

This research was supported by the NASA EO-1 Program (Grant NCC5-463), the Terrestrial Sciences Program of the Army Research Office (DAAG55-98-1-0287) and NSF (Grant IIS-0312471). We thank Amy Neuenschwander of the UT Center for Space Research for help in pre-processing the Hyperion data and interpreting the overall classification results.

REFERENCES

- [1] J.S. Pearlman, P.S. Berry, C. C. Segal, J. Shapanski, D. Beiso, and S.L. Carman, "Hyperion: a space-based imaging spectrometer," *IEEE Trans. Geosci. Rem. Sens.* **41**(6), 1160-1173, 2003.
- [2] D. Landgrebe, "Hyperspectral image data analysis as a high dimensional signal processing problem," (Invited), Special Issue of the *IEEE Signal Processing Magazine*, **19**(1): 17-28, 2002.
- [3] S.J. Raudys and A.K. Jain, "Small sample size effects in statistical pattern recognition: recommendations for practitioners," *IEEE Trans. Pattern Anal. Mach. Intell.*, **13**(3): 252-64, 1991.
- [4] S. Tadjudin and D.A. Landgrebe, "Covariance estimation with limited training samples," *IEEE Trans. Geosci. Rem. Sens.*, **37**(4): 2113-8, 1999.
- [5] S.B. Serpico and L. Bruzzone, "A new search algorithm for feature selection in hyperspectral remote sensing images," *IEEE Trans. Geosci. Rem. Sens.*, **39**(7): 1360-1367, 2001.
- [6] Q. Jackson and David Landgrebe, "An adaptive classifier design for high-dimensional data analysis with a limited training data set," *IEEE Trans. Geosci. Rem. Sens.*, **39**(12): 2664-79, 2001.
- [7] A. Blum and T. Mitchell, "Combining labeled and unlabeled data with co-training," *Proc. 11th Annual Conf. Computational Learning Theory*, 92-100, 1998.
- [8] C. Lee and D. Landgrebe, "Decision boundary feature extraction for neural networks," *IEEE Trans. Neural Networks*, **8**(1): 75-83, 1997.
- [9] X. Jia and J.A. Richards, "Segmented principal components transformation for efficient hyperspectral remote-sensing image display and classification," *IEEE Trans. Geosci. Rem. Sens.*, **37**(1): 538-42, 1999.
- [10] S. Kumar, J. Ghosh, and M.M. Crawford, "Best basis feature extraction algorithms for classification of hyperspectral data," *IEEE Trans. Geosci. Rem. Sens.* **29**(7): 1368-79, 2001.
- [11] J.T. Morgan, A. Henneguelle, M.M. Crawford, J. Ghosh, and A. Neuenschwander, "Adaptive feature spaces for land cover classification with limited ground truth," *Intl. J. Pattern Recognition and*

Artificial Intelligence, Spl. Issue on Fusion of Multiple Classifiers, J. Kittler and F. Roli (Eds.), **18**(5): 777-800, 2004.

[12] L. I. Kuncheva, *Combining Pattern Classifiers: Methods and Algorithms*. Wiley, 2004.

[13] L. Breiman, "Bagging predictors," *Machine Learning*, **24**(2): 123-40, 1996.

[14] K. Tumer and J. Ghosh, "Error correlation and error reduction in ensemble classifiers," *Connection Science*, Special Issue on Combining, **8**(3/4), 385-404, 1996.

[15] T.K. Ho, "The random subspace method for constructing decision forests," *IEEE Trans. Pattern Anal. Mach. Intell.*, **20**(8): 832-844, 1998.

[16] M. Skurichina and R.P. W. Duin, "Bagging, boosting, and the random subspace method for linear classifiers," *Int. J. Pattern Analysis and Applications*, **5**(2): 121-135, 2002.

[17] L. Breiman, "Random forests," *Machine Learning*, **45**: 5-32, 2001.

[18] J. Furnkranz, "Round robin classification," *J. Machine Learning Research*, **2**: 721-747, 2002.

[19] T. G. Dietterich and R. Bakiri, "Solving multiclass learning problems using error correcting output codes," *J. Artificial Intelligence Research*, **2**(1): 263-286, 1995.

[20] S. Kumar, J. Ghosh, and M.M. Crawford, "Hierarchical Fusion of Multiple Classifiers for Hyperspectral Data Analysis," *Int. J. Pattern Analysis and Applications*, **5**(2): 210-220, 2002.

[21] M. Crawford, J. Ham, and J. Ghosh, "Robust classifiers for hyperspectral data analysis using limited training data," presented at the *2003 Tyrrhenian International Workshop on Remote Sensing*, Elba Island, Italy, 15-18, Sept. 2003.

[22] A.L. Neuenschwander, M.M. Crawford, and S. Ringrose, "Results of the EO-1 experiment - Use of Earth Observing-1 Advanced Land Imager (ALI) data to assess the vegetational response to flooding in the Okavango Delta, Botswana," *Int. J. Remote Sens.*, accepted for publication.

[23] S. Kumar and J. Ghosh, "GAMLS: A generalized framework for associative modular learning systems," *Proc. Applications and Science of Comp. Intelli. II*, Orlando, Florida, 24-34, 1999.

[24] L. Breiman, J.H. Friedman, R.A. Olshen, C.J. Stone, *Classification and Regression Trees*, Wadsworth, Belmont, 1984.

[25] J.W. Boardman, "Geometric mixture analysis of imaging spectrometry data," *Proc. Int. Geosci. Rem. Sens. Symp.*, **4**, 2369-2371, 1994.

[26] G.P. Asner and K.B. Heidebrecht, "Spectral unmixing of vegetation, soil and dry carbon in arid regions: Comparing multi-spectral and hyperspectral observations," *Int. J. Remote Sens.*, **23**: 3939-3958, 2002.

LIST OF TABLES

Table 1. Class codes, names, and number of training samples for Kennedy Space Center AVIRIS data

Table 2. Class codes, names, and number of training and spatially disjoint test samples for Botswana Hyperion data.

Table 3. Entropy-based diversity of ensemble members observed for the spatially disjoint Botswana Hyperion data at different sampling rates.

Table 4. Confusion matrix for Hyperion spatially disjoint test set at 75% sampling rate, RF-BHC classifier.

Table 1.

	Class	No. samples
1	Scrub	761 (14.6%)
2	Willow swamp	243 (4.66%)
3	Cabbage palm hammock	256 (4.92%)
4	Cabbage palm/oak hammock	252 (4.84%)
5	Slash pine	161 (3.07%)
6	Oak/broadleaf hammock	229 (4.38%)
7	Hardwood swamp	105 (2.0%)
8	Graminoid marsh	431 (8.27%)
9	Spartina marsh	520 (9.99%)
10	Cattail marsh	404 (7.76%)
11	Salt marsh	419 (8.04%)
12	Mud flats	503 (9.66%)
13	Water	927 (17.8%)

Table 2.

	Class	No. samples, Test set	No. samples, S. D. Test set
1	Water	270 (8.31%)	126 (5.05%)
2	Hippo grass	101 (3.09%)	162 (6.5%)
3	Floodplain grasses1	251 (7.74%)	158 (6.34%)
4	Floodplain grasses2	215 (6.63%)	165 (6.62%)
5	Reeds1	269 (8.27%)	168 (6.74%)
6	Riparian	269 (8.27%)	211 (8.46%)
7	Firescar2	259 (7.98%)	176 (7.06%)
8	Island interior	203 (6.26%)	154 (6.17%)
9	Acacia woodlands	314 (9.67%)	151 (6.05%)
10	Acacia shrublands	248 (7.65%)	190 (7.62%)
11	Acacia grasslands	305 (9.38%)	358 (14.35%)
12	Short mopane	181 (5.56%)	153 (6.13%)
13	Mixed mopane	268 (8.27%)	233 (9.34%)
14	Exposed soils	95 (2.92%)	89 (3.57%)

Table 3.

Methods \ Sampling Rate		15%	30%	50%	75%
RF-BHC	Average	0.440	0.383	0.345	0.326
	Std	0.013	0.008	0.006	0.007
RF-CART	Average	0.516	0.476	0.460	0.460
	Std	0.024	0.020	0.011	0.010

Table 4.

<i>Classified\Actual (col)</i>	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	126	44	0	0	0	0	1	0	0	0	0	0	0	0
2	0	84	0	0	4	0	0	0	0	0	0	0	0	0
3	0	0	139	18	3	0	5	0	1	0	0	0	0	0
4	0	0	0	144	4	2	0	0	0	0	0	0	0	0
5	0	0	0	2	96	23	0	0	0	0	0	0	0	0
6	0	15	0	0	16	149	0	0	30	0	0	0	2	0
7	0	2	0	0	0	0	160	0	0	0	0	0	0	0
8	0	0	0	0	12	0	0	150	0	0	92	1	0	0
9	0	17	4	1	4	36	0	0	115	0	0	0	60	0
10	0	0	15	0	11	0	7	0	1	172	38	0	6	0
11	0	0	0	0	0	0	0	0	0	17	208	0	0	0
12	0	0	0	0	7	0	0	0	0	0	1	129	2	0
13	0	0	0	0	9	1	3	0	4	1	6	23	163	0
14	0	0	0	0	2	0	0	4	0	0	13	0	0	89
Total	126	162	158	165	168	211	176	154	151	190	358	153	233	89
Class accuracies	100	51.9	88	87.3	57.1	70.6	90.9	97.4	76.2	90.5	58.1	84.3	70	100

LIST OF FIGURES

- Figure 1.** BINARY HIERACHICAL(multi)-CLASSIFIER framework for solving a C-class problem.
- Figure 2 a).** AVIRIS data, (Bands 31, 21, 11) acquired over KSC, training sites overlaid
- Figure 2 b).** Classified image of KSC AVIRIS data using RF-BHC classifier.
- Figure 3.** Ave and std. dev. of classification accuracies for AVIRIS test data.
- Figure 4 a).** Hyperion data, (Bands 51, 149, 31) acquired over Okavango Delta, training sites overlaid.
- Figure 4 b).** Classified image of Hyperion data over Okavango Delta using RF-BHC classifier.
- Figure 5.** Ave. and std. dev. of classification accuracies for Hyperion test and spatially disjoint test sets.
- Figure 6 a).** Class dependent accuracies for Hyperion test set.
- Figure 6 b).** Class dependent accuracies for Hyperion spatially disjoint test set.

Figure 1.

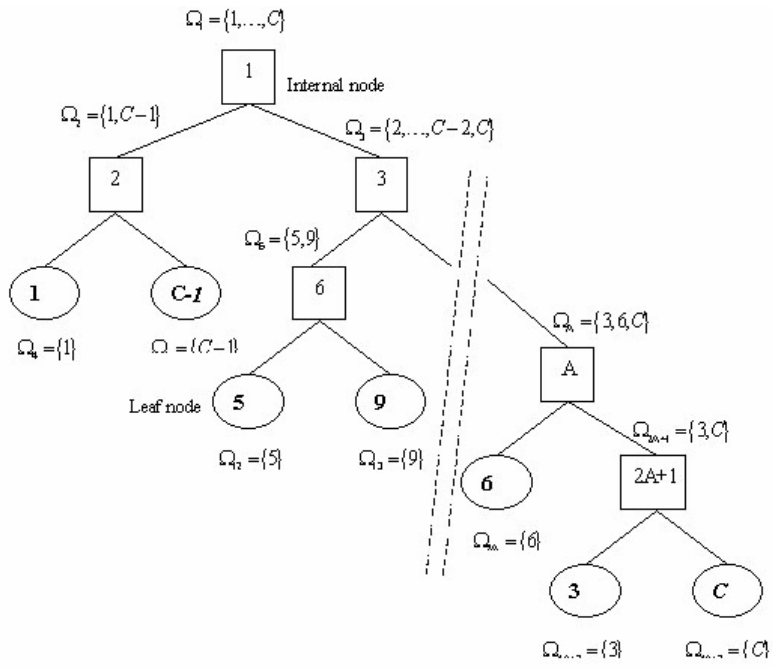


Figure 2 a).




Figure 2 b).




Figure 3.

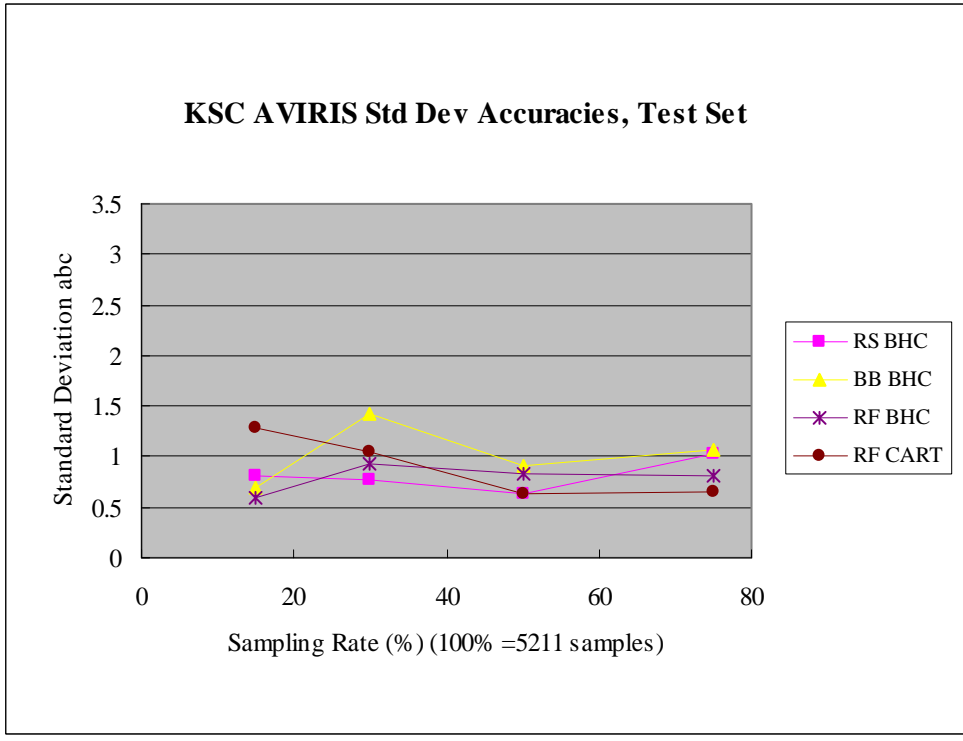



Figure 4 a).




Figure 4 b).


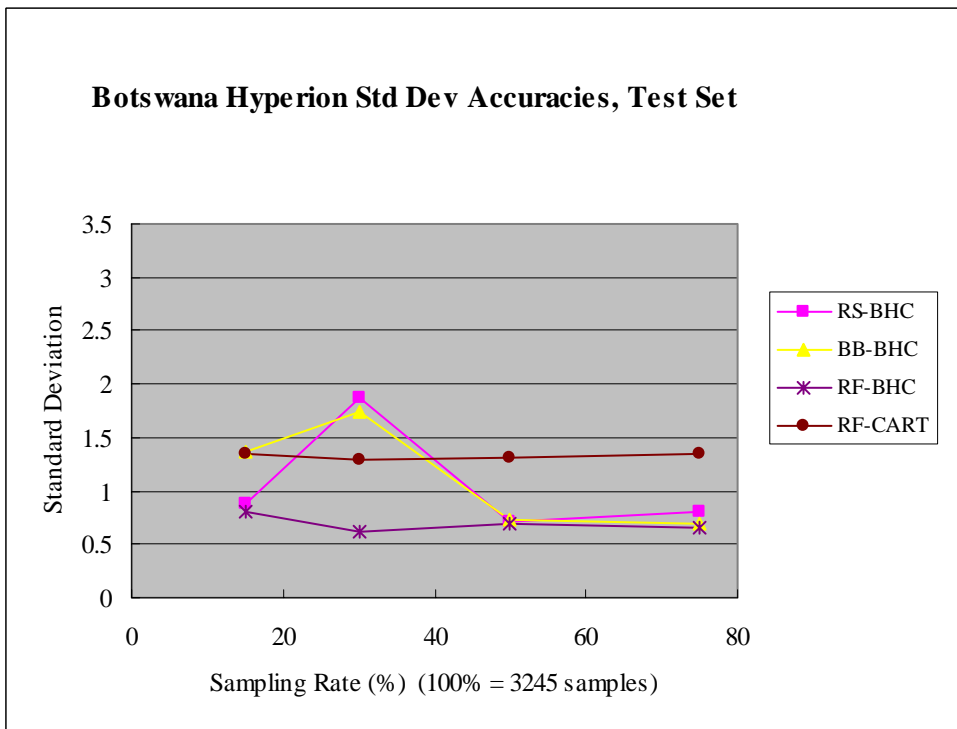
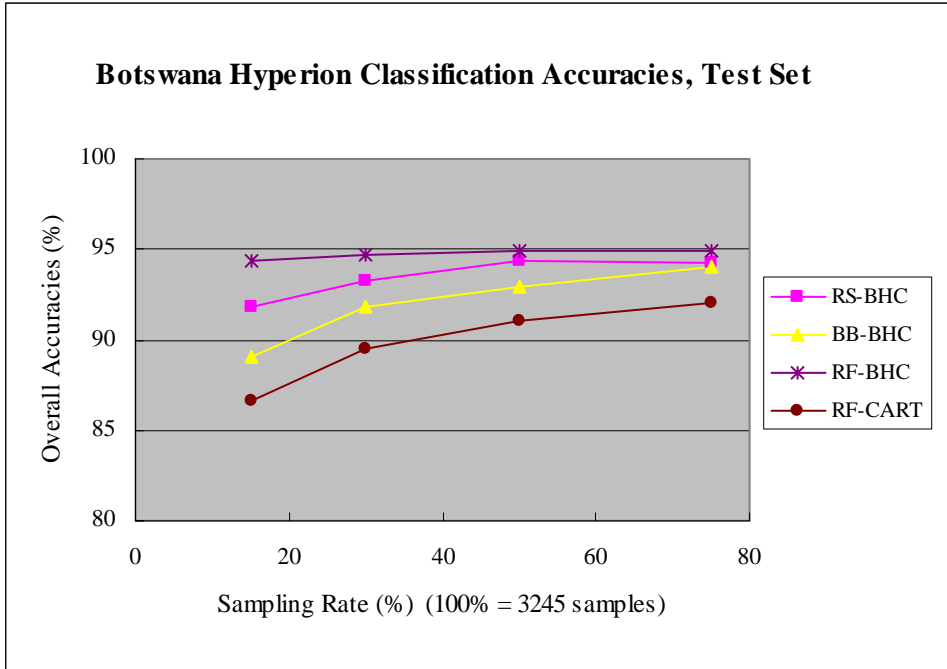



Figure 5.



Botswana Hyperion Classification Accuracies, Spatially Disjoint Test Set



Botswana Hyperion Std Dev Accuracies, Spatially Disjoint Test Set

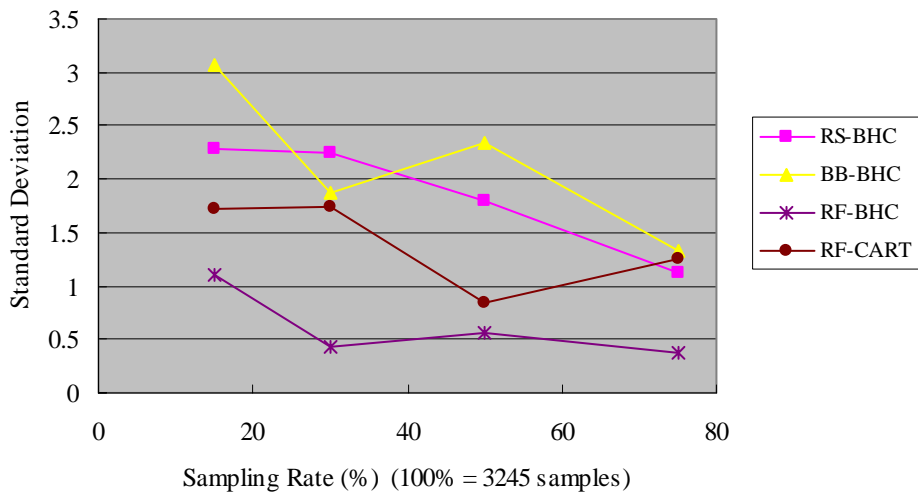


Figure 6 a).


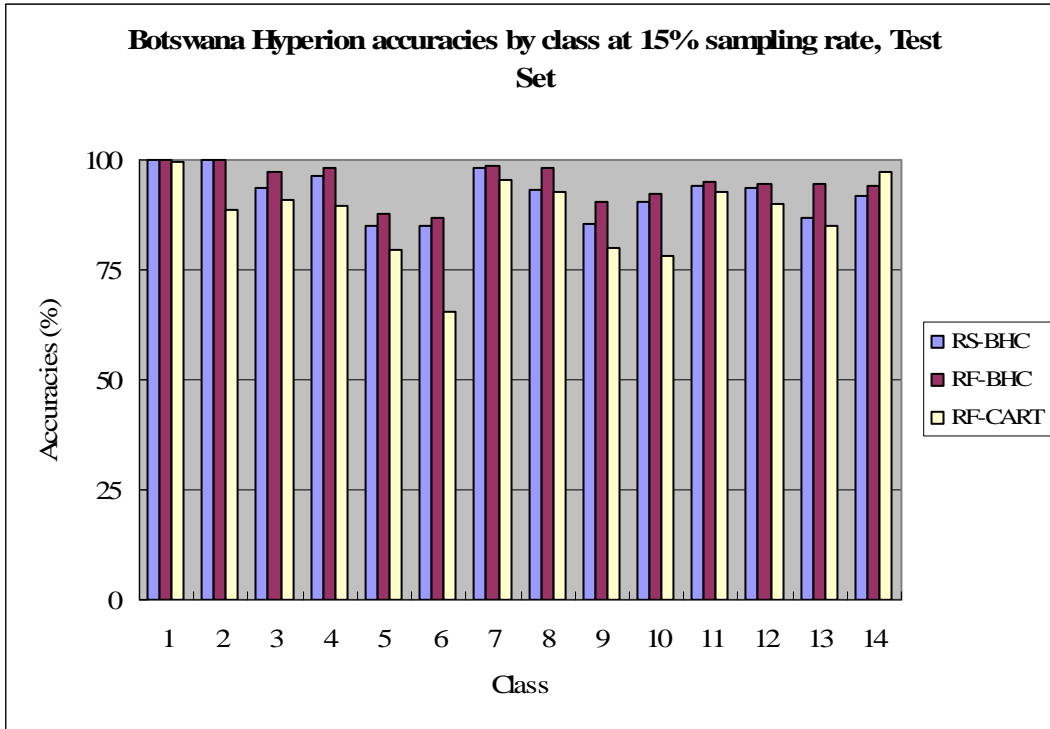


Figure 6 b).

