Iterative Denoising for Adaptive Sensors

David J. Marchette* Carey E. Priebe† Youngser Park† Damianos Karakos†

Abstract

We describe a methodology for the adaptive selection of features in high dimensional data, and apply it to adaptive sensing of hyperspectral imagery. The methodology constructs a tree defining the spectra to collect, conditional on the results of previous collection. This allows the sensor to collect a small subset of the available spectrum, focused specifically on the current environment. An example of detecting chemical beakers in hyperspectral imagery is used to illustrate the approach.

Keywords: Pattern recognition; image processing; hyperspectral imagery; adaptive sensor, tree-based clustering.

1 Introduction

One of the consequences of the accelerating pace of sensor design is that sensors are becoming more and more capable of collecting data in vast ranges of the electromagnetic spectrum, while the processing of the data lags behind. This can be a result of communication constraints, data processing constraints, or the curse of dimensionality. In this paper we consider a simple adaptive sensor, capable of collecting imagery in any of many spectral bands. One can think of this as a type of hyperspectral sensor. It is impractical, due to time and bandwidth constraints, to collect images in all bands (in our case, the dimensionality of the hyperspectral sensor is 266); instead one wishes to select just those bands which contain information sufficient to solve the problem at hand. Thus we want a methodology to select the bands to collect, and to adjust the collection based on previously collected information.

It is important to note that we do not propose trying to solve the problem at hand (for example, pixel classification) with a single collection; rather we propose a methodology similar to the "20-questions" or coarse-to-fine methodology of Blanchard and Geman [2005]. We first collect a few spectral bands, sufficient to determine where a target of interest may be in the image. Then, for those regions of potential targets, we select a new sensor setting to collect (potentially different) spectral bands. This results in a methodology in which the processing is focused on the areas of potential targets, and the sensor adapts to collect the bands relevant to the observed data.

This work is a follow-on to Priebe et al. [2004a], Priebe et al. [2004b], and Priebe et al. [2004c], where it was shown that under certain conditions, the optimal strategy is a tree: an initial collection is used to select the bands (features) for a subsequent collection, which in turn drives the next collection, and so on. In our case, we will assume that the sensor can collect a linear combination of the bands, and that this linear combination is to some degree selectable. Similar work is reported in Priebe et al. [2005], where the problem investigated is the detection of a car.

*Naval Surface Warfare Center, Code B10, Dahlgren, VA 22448. dmarchette@gmail.com
†The Johns Hopkins University, Center for Imaging Science, Baltimore, MD, 21218. {cep,youngser,damianos}@jhu.edu

2 Hyperspectral Data

The data in this study consists of a hyperspectral image taken of beakers containing two chemicals (ether and acetone). Figure 1 shows a representative image from the data cube. The two beakers can be seen as dark rectangles in the lower left corner of the image. Part of a vehicle is visible in the right hand side. Figure 2 depicts the first two principal components for the data cube. Note that the beakers (red (ether) and green (acetone) in this figure) are not at all separated from the background or each other in this projection. However, there is obvious clustering in the data, indicating that it might be of value to consider local projections.

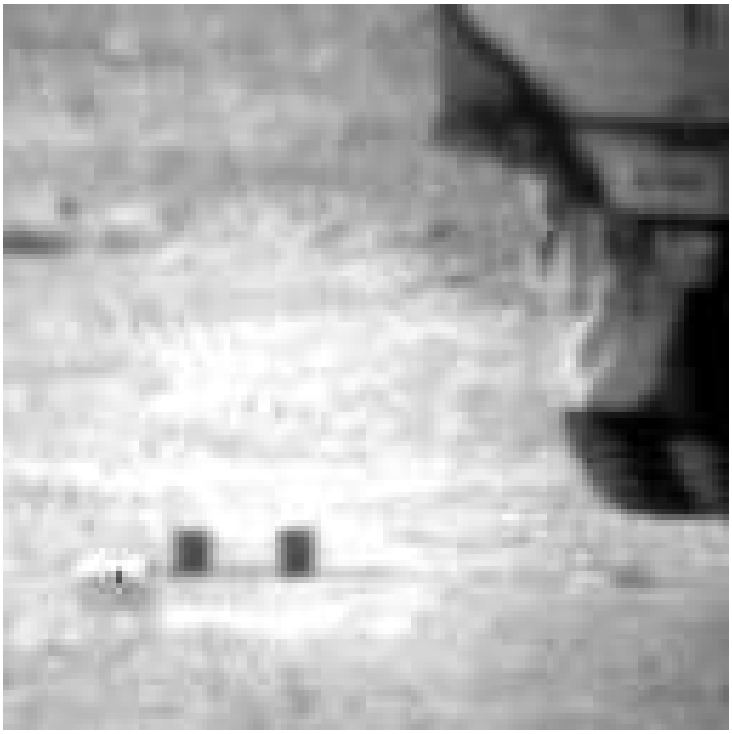

Figure 1: One band of the image showing the two beakers (small dark rectangles). The image is 128x128 pixels, rendered in an 8-bit gray scale. The dark mass on the right is part of a vehicle.

In addition to the full data cube, the sensor is capable of producing a linear combination of spectral bands. In fact, the sensor actually returns a linear combination (Hadamard frames, see Priebe et al. [2005] for more information) which is subsequently preprocessed into the data cube. Thus, we will also consider the case where we operate on the sensor response before processing into the data cube. In this manner, we can implement a sensor with the desired properties; rather than collecting the entire data cube, we can collect just a few features at a time. The purpose of this paper is to demonstrate that this can lead to good results.

3 Results

We now construct the iterative denoising tree for the hyperspectral data cube. We assume that for each time step the sensor can collect and transmit two images, corresponding to linear combinations of the bands. Thus, the projected data will be two dimensional. The dimensionality was chosen arbitrarily, simply to make the data displays easier. In a real system this would also be chosen to optimize performance (and need not be the same at each node of the tree).

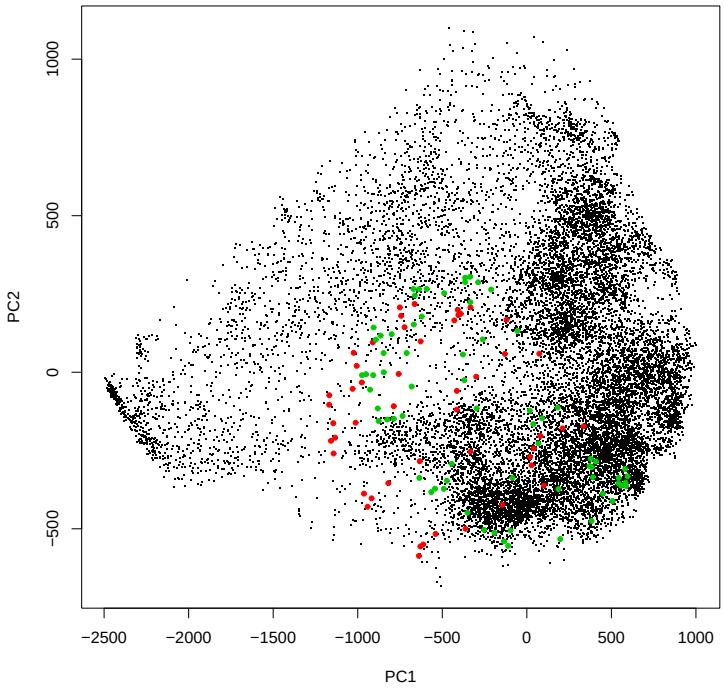

Figure 2: The first two principal components of all the data in the 128x128x266 data cube. The points from the two beakers are colored green (acetone) and red (ether).

The selection of the projection, for any given set of pixels, is problematic. What we want is a projection that either provides good classification or good clustering, and then to proceed according to which criterion is met. This is daunting due to the combinatorics of selecting two projections from all possible pairs. Instead, a suboptimal method is used, and for the purposes of illustration, we will use principal components, conscious of the fact that this is not likely to be optimal for either classification or clustering, but that empirical studies have shown that it is often useful.

We illustrate the procedure using hyperspectral data taken of beakers containing two chemicals, within a complex background (Figures 1, 2). Pixels corresponding to the two chemicals are color coded in the scatter plots of the projections.

The iterative denoising tree proceeds as follows. Starting with the full data set we project to the first two principal components (see Figure 2). The clustering is performed using model based clustering (*mclust*, see Fraley and Raftery [2002]). For the purposes of illustration, *mclust* was constrained to find no more than six clusters, and it chose to use a model of maximal size. The clusters are illustrated in Figure 3 (left) as different colors plotted on the image (the colors are arbitrary), and in Figure 4 both using the image to indicate (in red) where the pixels for the individual clusters lie, and below in a scatterplot of the two dimensional data in the cluster. Note that the projections shown below the images in Figure 4 have been recomputed on the subsetted data. This results in a new projection (hence a new sensor collection in an adaptive sensor) that is better suited for the data represented by this cluster.

Within each cluster (branch of the tree) the process repeats. We have illustrated this for only the cluster containing the majority of the observations from the two beakers (the third cluster from the left). The pixels in the other clusters are primarily edge pixels, which could be expected to be problematic for classification in any event. The *mclust* algorithm selects three clusters for this group, two of which separate the chemical

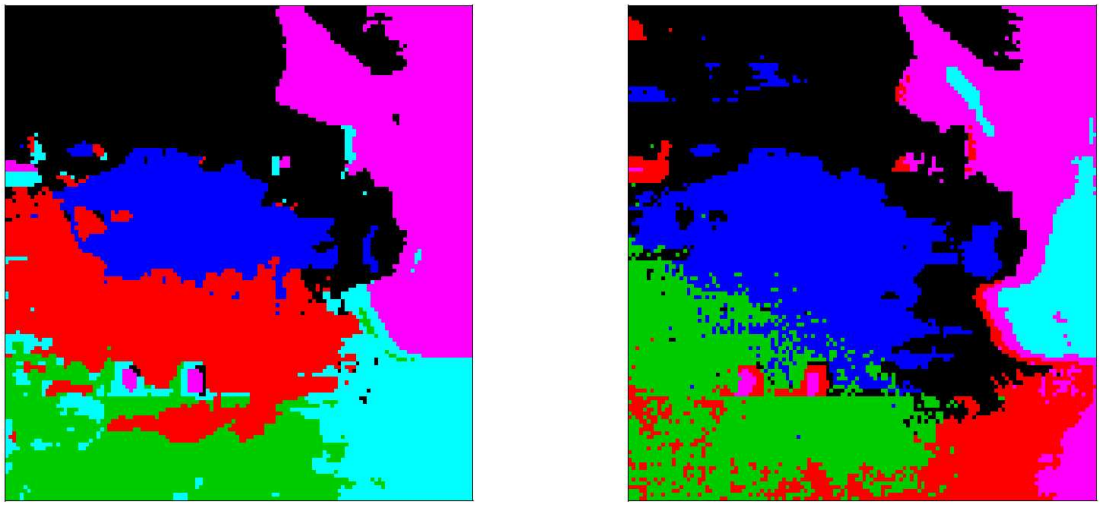

Figure 3: The initial clusters found by *mclust* for the data cube (left) and Hadamard frames (right). In each case six clusters were chosen, and the different clusters are color coded (the coloring is arbitrary within each image).

observations quite well. The middle cluster shows some overlap, and we repeat the procedure on these data. In this case *mclust* selects two clusters, one containing two pixels from the beakers, which is not shown in this figure. As before, each projection is recomputed on the data in the cluster.

The three leaves containing most of the chemical pixels are shown enlarged in Figure 4. As can be seen, the chemical pixels separate quite well in these projections. Thus, the iterative denoising tree has selected projections which do an excellent job of classification. We have not used the class labels at all in this procedure, except to select the portions of the tree to show in the figure. It is at the leaves that classification is performed. Note that one might well continue the tree down from these leaves for the purpose of classifying the pixels according to chemical; however there sparsity of data suggests this is probably not worth pursuing on this data set.

The sensor of interest here cannot collect arbitrary projections, but only those associated with Hadamard frames. We illustrate how the work might be applied in this case in Figure 5. Note that unlike principal components, there is no natural ordering of the Hadamard frames. As a result, one must either undertake the computationally expensive task of searching for good patterns, or use some sub-optimal approximation. The idea we pursue here is to use the principal components to guide this search. We select the two Hadamard frames so that the projected data best matches the principal components projection. We do this sequentially, first finding the best match to the first principal component, then finding the best match to the second. At first, one might doubt that this approach has much hope of success, but as can be seen in Figure 5 it works quite well on these data.

The first level clustering, again to six clusters, is depicted in Figure 3 (right). This is quite similar to that produced from the principal components. Also, the groupings are not unreasonable, given what is known about the objects within the image. Thus, we have some reason for optimism. The iterative denoising tree is quite simple in this case. It stops after just one iteration (we are only considering the cluster containing the majority of beaker pixels; other branches do in fact continue to split). After a single clustering the beaker pixels are separated out quite well from the background. This is not at all the case in the original projection, and so the clustering and re-projection is essential to this result.

Further processing is required to separate the different chemicals. Unfortunately, in this initial proof-

of-concept data collection, there are not enough examples of the beakers to support any further processing. However, in principle, the procedure continues by further clustering on the leaves, then choosing projections within the clusters as above. At the leaves in Figures 4 or 5, one could cluster, or one could do a beaker/background classification. In these data this would result in the same thing, at least as far as the beaker pixels are concerned. The projections on the leaves would then be chosen not by principal components or clustering criteria, but rather by classification criteria, such as Fisher's linear discriminant, or other feature extraction methods (Duda et al. [2000]).

Clearly, a better method for constructing the tree would be to use the Hadamard frames directly, and select from these to optimize a clustering criterion rather than fitting them to the principal components. This was done in Priebe et al. [2005] on a different problem than considered here. One purpose of this study was to show that matching to the principal components could be used as a solution to the combinatorial problem of selecting the best pair (or k-tuple) of Hadamard frames at each stage.

4 Summary

Iterative denoising can be thought of as a method for searching through complex data, selecting different views of the data at different stages of the processing. It can also be thought of as a method for selecting the parameters (wavelengths to sense) for an adaptive sensor. In many situations one cannot collect all the features that the sensor is capable of, nor would one want to. One solution is to design adaptive sensors, where the information collected can be adjusted according to the environment or the context. Based on the information at hand, the sensor is tasked to collect a particular set of measurements.

The key to the iterative denoising approach is the observation that classification should not be attempted globally. Different regions of feature space require different processing. Furthermore, it is often a bad idea to focus exclusively on classification, even if the ultimate task is classification. The context of the environment presented to the sensor must be taken into account. This is particularly important if there are many possible classes. In the example presented in this paper, it is important to first separate the background (dirt) from potential beakers (man made objects, if you will). Then, one must separate out beakers from the other objects in the image. Only then can one hope to classify the different chemicals within the beakers.

The problem of constructing a training set for an adaptive sensor is a difficult one. We have approached this by using a well-known approach (principal components analysis (PCA)) to the processed data cube, then attempting to match the principal components selected to the raw bands sensed. This is sub-optimal in at least two ways: PCA itself is sub-optimal for this problem; further, it is unlikely that the principal components of the processed data cube are close to the raw sensor output. With this said, the results are extremely promising. Finally, note that we did not use the actual classification of the pixels, except in selecting which branches to expand for the figures. Thus, only at the leaves of the tree is classification performed. We obtain a tree which clusters the pixels, as well as one in which classification can be performed at the leaves. These observations suggest that iterative denoising is a useful tool for analysis of data from adaptive sensors.

Using a tree to encode the sensing decisions allows fast decisions for those cases where there are no targets in view. This allows the sensor to spend the majority of its scarce resources on those regions in which previous senses have indicated that there is a high likelihood of target. This allows the trade-off of false rejection rate with bandwidth and processing constraints, and allows for the design of sensors that can adapt to a changing environment without resulting in unmanageable data volume.

Acknowledgments

Work supported in part by the Applied and Computational Mathematics Program (ACMP) of the Defense Advanced Research Projects Agency (DARPA).

References

Gilles Blanchard and Donald Geman. Hierarchical testing designs for pattern recognition. *Annals of Statistics*, 33:1155–1202, 2005.

Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification*. Wiley, New York, 2000.

C. Fraley and A. E. Raftery. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, 97:611–631, 2002.

Carey E. Priebe, David J. Marchette, and Dennis M. Healy. Integrated sensing and processing decision trees. *IEEE PAMI*, 26:699–708, 2004a.

Carey E. Priebe, David J. Marchette, and Dennis M. Healy. Integrated sensing and processing for statistical pattern recognition. In D. Rockmore and D. M. Healy, editors, *Modern Signal Processing*, pages 223–246. Cambridge University Press, 2004b.

Carey E. Priebe, David J. Marchette, Youngser Park, and Robert R. Muise. An application of integrated sensing and processing decision trees for target detection and localization on digital mirror array imagery. *Applied Optics*, to appear, 2005.

Carey E. Priebe, David J. Marchette, Youngser Park, Edward J. Wegman, Jeffrey L. Solka, Diego A. Socolinsky, Damianos Karakos, Ken W. Church, Roland Guglielmi, Ronald R. Coifman, Dekang Lin, Dennis M. Healy, Marc Q. Jacobs, and Anna Tsao. Iterative denoising for cross corpus discovery. In *Proceedings of COMPSTAT 2004*, pages 381–392, 2004c.

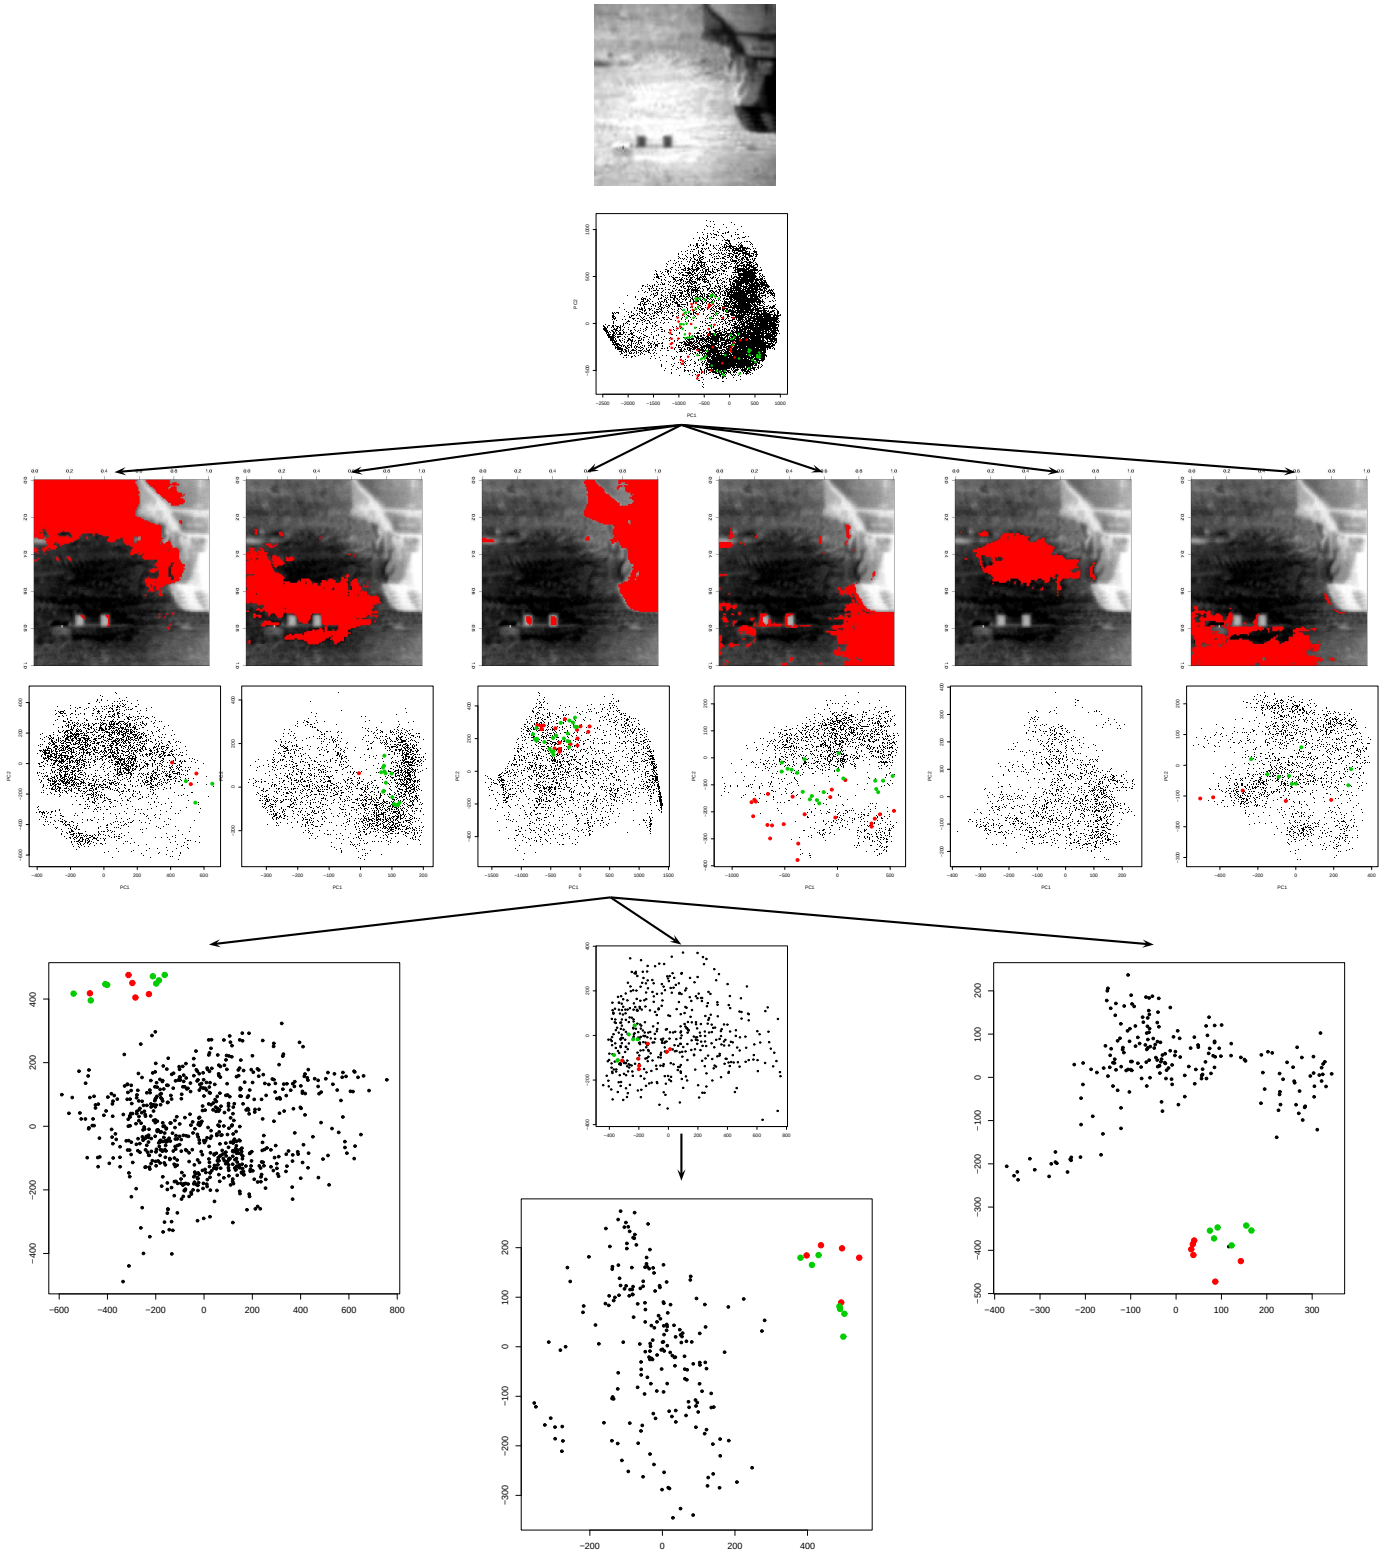

Figure 4: The iterative denoising tree on the data cube. The top (root) node shows the original projection. Next are the six clusters chosen by *mclust*. The clusters are indicated in red in the images, with the corresponding data shown below. In all cases, the projection is recomputed on these subsets. The third cluster from the left contains the largest subset of the chemical observations, and the resulting clusters are shown below. The overlap in the middle cluster suggests a further clustering, resulting in two clusters, only one of which contains significant numbers of chemical observations. This is shown in the final leaf.

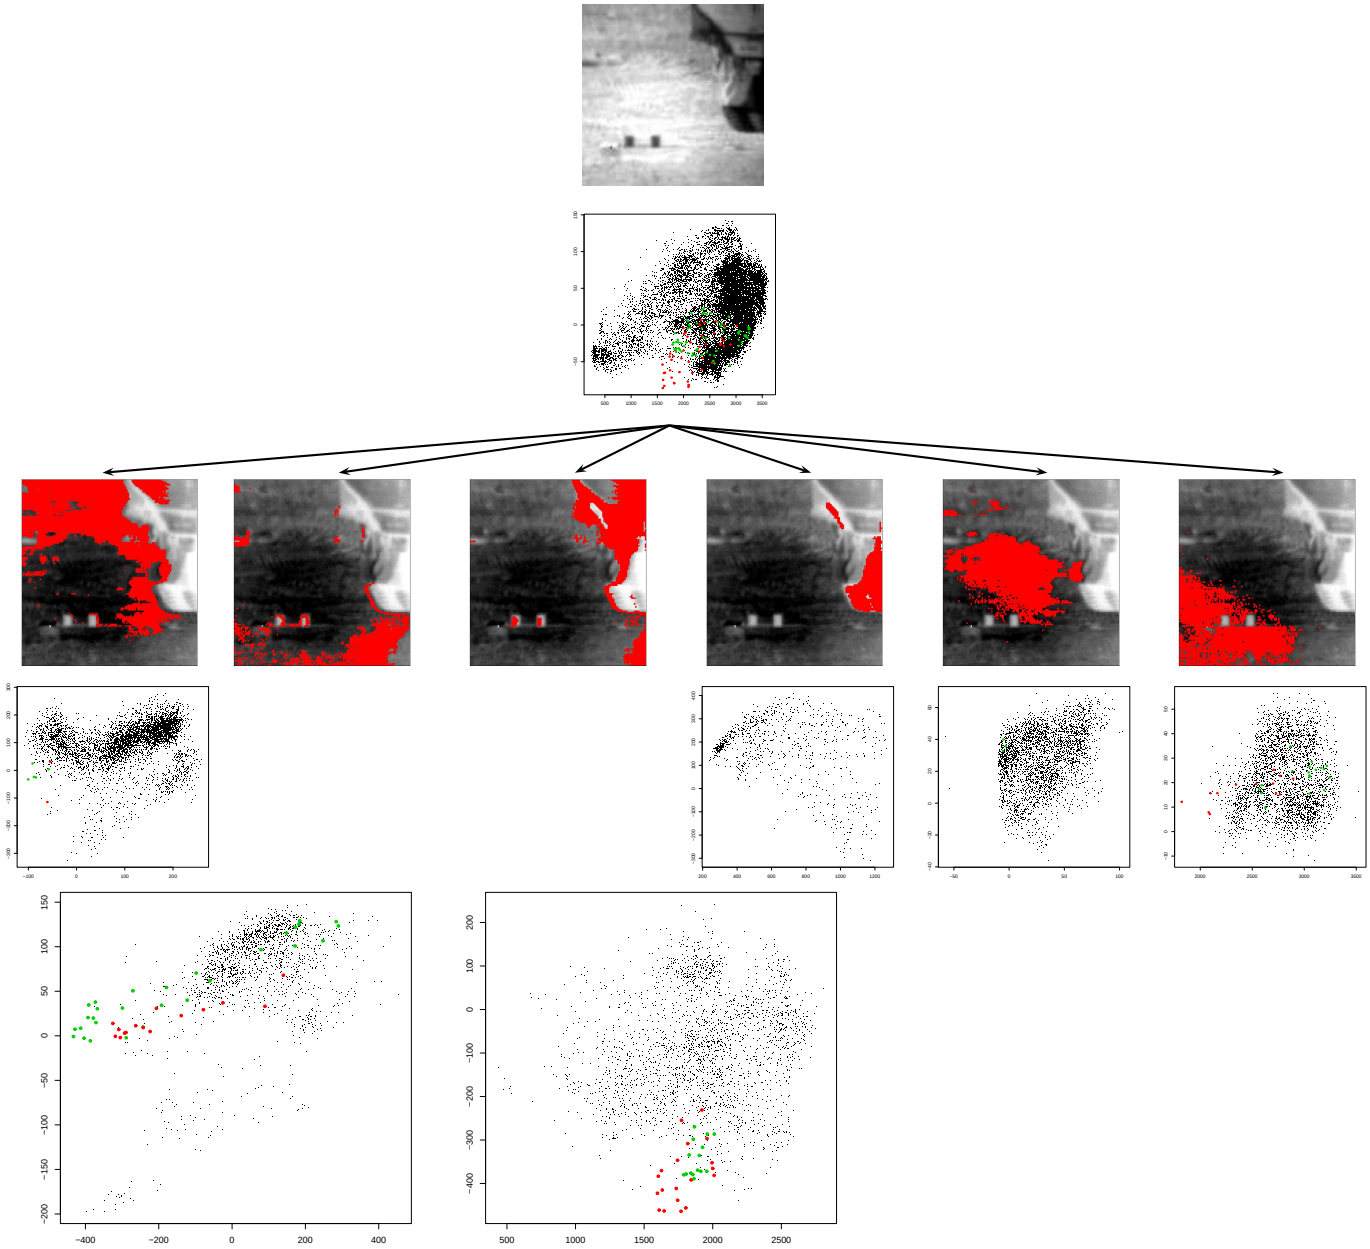

Figure 5: The iterative denoising tree on the Hadamard frames. The Hadamard frames are chosen at each stage to be close to the principal components projection. The top (root) node shows the original projection. Next are the six clusters chosen by *mclust*. The clusters are indicated in red in the images, with the corresponding data shown below. In all cases, the projection is recomputed on these subsets. The second and third clusters from the left contains the largest subset of the chemical observations (these are shown at a larger scale to better see the class separation). The first and sixth clusters also contain a number of pixels labeled as chemicals; however, as can be seen by the images above, these are primarily edge pixels. The separation of the chemicals and the background is quite good after just one iteration of the denoising tree, although a slightly better separation can be obtained for the observations in the these clusters by iterating one more step.