

On the Use of Machine Vision Techniques to Detect Human Settlements in Satellite Images

*Chandrika Kamath, Sailes K. Sengupta, Douglas Poland,
and John A. H. Fetterman*

This article was submitted to
Image Processing: Algorithms and Systems II, SPIE Electronic
Imaging, Santa Clara, CA, January 22, 2003.

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

January 10, 2003

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

This report has been reproduced
directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (423) 576-8401
<http://apollo.osti.gov/bridge/>

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161
<http://www.ntis.gov/>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

On the use of Machine Vision Techniques to Detect Human Settlements in Satellite Images

Chandrika Kamath, Sailes K. Sengupta, Douglas Poland, and John A. H. Futterman

Lawrence Livermore National Laboratory
7000 East Avenue, Livermore, CA 94551, U.S.A

ABSTRACT

The automated production of maps of human settlement from recent satellite images is essential to studies of urbanization, population movement, and the like. The spectral and spatial resolution of such imagery is often high enough to successfully apply computer vision techniques. However, vast amounts of data have to be processed quickly. In this paper, we propose an approach that processes the data in several different stages. At each stage, using features appropriate to that stage, we identify the portion of the data likely to contain information relevant to the identification of human settlements. This data is used as input to the next stage of processing. Since the size of the data has reduced, we can now use more complex features in this next stage. These features can be more representative of human settlements, and also more time consuming to extract from the image data. Such a hierarchical approach enables us to process large amounts of data in a reasonable time, while maintaining the accuracy of human settlement identification. We illustrate our multi-stage approach using IKONOS 4-band and panchromatic images, and compare it with the straight-forward processing of the entire image.

Keywords: Satellite images, computer vision, human settlements, multi-stage hierarchical approach

1. INTRODUCTION

Maps of human settlement produced from satellite imagery are a key part of any application in which urban land cover, land use, or boundary is a consideration. The resolution of current satellite imagery is often fine enough that even small villages can be accurately identified. For example, the IKONOS imagery¹ that we consider in this paper is available at 4 meter Ground Sample Distance (GSD) as 4-band multi-spectral data as well as 1 meter GSD panchromatic data. However, at this fine resolution, the volumes of data are such that semi-automated techniques must be used to sift through the data.

In this paper, we describe a multi-stage approach to the semi-automated production of maps of human settlements from satellite data. We are interested in exploiting both the 4-band multi-spectral and the panchromatic imagery from IKONOS as they provide complementary information at different resolutions and wavelengths. In order to mitigate the resulting increase in the size of the data, we process the data in several stages. In the initial stages, we apply low-level image processing to the multi-spectral imagery to identify regions that are unlikely to contain human settlements. In the later stages, focusing on only those regions that are likely to contain human settlements, we use the higher resolution panchromatic imagery of the same scene to detect edges, corners, and other human-made features in the data. This refines these regions, reducing the false positives.

The use of a multi-stage approach to process the data is not new. For example, a four stage approach has been used for land-cover/land-use classification of urban areas.² These stages include feature extraction, feature coding, feature selection, and classification. Starting with 6-channel data from Landsat TM and 1-channel from ERS-1 SAR, the data is first spatially resampled to 25 meter/pixel. From this data, various statistical, textual, and Gabor features are extracted. Next, these features are coded using Self-Organizing Maps.³ As the number of features is still quite large, Classification And Regression Trees (CART)⁴ are used to select a subset of the features. Finally, these features are input to multi-layer perceptrons and the k-nearest neighbor classifier to classify the data into one of 33 classes. A similar multi-stage approach has also been used with Landsat TM

Further author information: (Send correspondence to C.K.)
C.K.: E-mail: kamath2@llnl.gov

and SPOT Pan data at 10-m pixel size to detect buildings in urban areas.⁵ Using texture analysis, the authors show that by incorporating spatial information with spectral information, they can obtain a higher accuracy for unsupervised classification of the data.

These earlier efforts focused on techniques that were appropriate to the resolution of the imagery sources available and the classification problem being addressed. In this paper, we illustrate our multi-stage approach using 4 channel (near-IR, red, green, and blue), 4m GSD multi-spectral data and 1 channel, 1m GSD panchromatic data from the IKONOS imagery. The paper is organized as follows: Section 2 describes the different stages in our approach. Section 3 evaluates our approach using sample images. The different options used are evaluated for their effectiveness, robustness, and computational efficiency. Finally, in Section 4, we summarize our work and discuss future directions.

2. TECHNICAL APPROACH

In order to extract human settlements in IKONOS multi-spectral and panchromatic satellite imagery, we use a multi-stage approach with the following stages:

- Using supervised or unsupervised classification, we first segment the multi-spectral image into k classes.
- Next, using simple statistics based on the distribution of these classes within a small window (or tile) in the image, we identify the tiles in the image that are likely to contain human settlements.
- Focusing on these tiles, we apply an edge and corner detector to the panchromatic imagery of the same scene.
- Tiles which have a sufficient number of edges and corners are identified as likely to contain human settlements.

We next describe each of these stages in further detail using a sample image.

2.1. Analysis of the multi-spectral imagery

Figure 1, panel (a), shows a 400×400 pixel, 4-band IKONOS multi-spectral image from the Nebraska region of North America. It includes both areas with human settlements (i.e. houses and other buildings) as well as open areas. Based on a visual analysis of the image, 6 separate classes were identified, corresponding to lush vegetation, not-so-lush vegetation, tarred roads and parking lots, concrete roads and rooftops, very bright surfaces such as rooftops, and dirt (with sparse vegetation). Of course, if we had selected an image from a different region of the world, or even the same image in a different season, the number and type of classes would likely be quite different.

Once the classes have been identified, a training set is generated with sufficient number of pixels from each class. In our example, we selected approximately 800 pixels from each of the classes, resulting in a training set of 4860 pixels. The training set was created by first selecting sample regions from each of the classes in the image and then extracting the features for the pixels in these regions. As we are working with multi-spectral data, the features are just the intensity values for the near-IR, red, green, and blue bands. The sample regions from which these pixels were selected are highlighted in Figure 1, panel (b).

Next, to check that the training set was a good representation of each class, we generated parallel plots for the pixels in each of the six classes. Parallel plots are often used in the visualization of high-dimensional data. Instead of the traditional coordinate system where the axes are perpendicular to each other, in parallel plots the axes are parallel to each other. As a result, more than three variables or dimensions can be visualized easily. Figure 2 displays the parallel plots for the initial set of pixels that were chosen as the training set for the image in Figure 1, panel (a). Each parallel plot has on the x-axis the five features or variables (near-IR, red, blue, green, and the class) and on the y-axis the corresponding values of the variable for each pixel. Thus a pixel is represented by the line segments that connect the values of the five variables. Note that for most of the classes, the parallel plots lie within a narrow band, indicating that the class is well defined. Also note that classes 2, 4,



Figure 1. (a) A 400×400 pixel multi-spectral image from Nebraska. (b) The regions in the image from which pixels were chosen for the 6 classes. Satellite images by Space Imaging.

5, and 6 have several outliers, that is, pixels that appear not to fit in the class. This is the result of the way in which the regions were selected to form a training set. While care was taken to select regions with pixels belonging to predominantly one class, this was not always possible, resulting in the outliers in the parallel plots. These outliers were removed from the training set before classification, resulting in the parallel plots in Figure 3 for classes 2 and 6.

Note that there is very little overlap in the parallel plots for the six classes. This indicates that the classes are well separated in feature space, and that a classifier with this training set will likely have a low cross-validation error.

In the approach described in this paper, we use supervised techniques for classifying the pixels in the multi-spectral image. Our experiences with unsupervised techniques are summarized in a companion paper.⁶

Once the training set was chosen, it was used to build a decision tree model. Decision trees^{4,7,8} belong to the category of classification algorithms wherein the algorithm learns a function that maps a data item into one of several pre-defined classes. The development of these algorithms typically has two phases. In the training phase, the algorithm is “trained” by presenting it with a set of examples with known classification. In the test phase, the model created in the training phase is tested to determine how well it classifies known examples. If the results meet expected accuracy, the model is put into operation to classify examples with unknown classification.

A decision tree is a structure composed of leaves and decision nodes. Data is introduced at a single point (the root) and proceeds toward one of the terminal nodes (or leaves), which represent classes. Each split in the tree is a decision node that specifies some test to be carried out on a feature (or a combination of features), with a branch and sub-tree for each possible outcome of the test. The decision at each node of the tree is made to reveal the structure in the data. Decision trees tend to be relatively simple to implement, yield results that can be interpreted, and have built-in dimensional reduction.

For the work in this paper, we used the decision tree software from the Sapphire data mining project.⁹ For the classification of the pixels in the multi-spectral image, we used axis parallel splits at each node of the decision tree. In such splits, the decision at each node of the tree is based only on a single feature. The splitting criterion used at each node of the tree was the Gini splitting criterion.⁴ This criterion is based on finding the split that most reduces the node impurity, where the impurity is defined as follows:

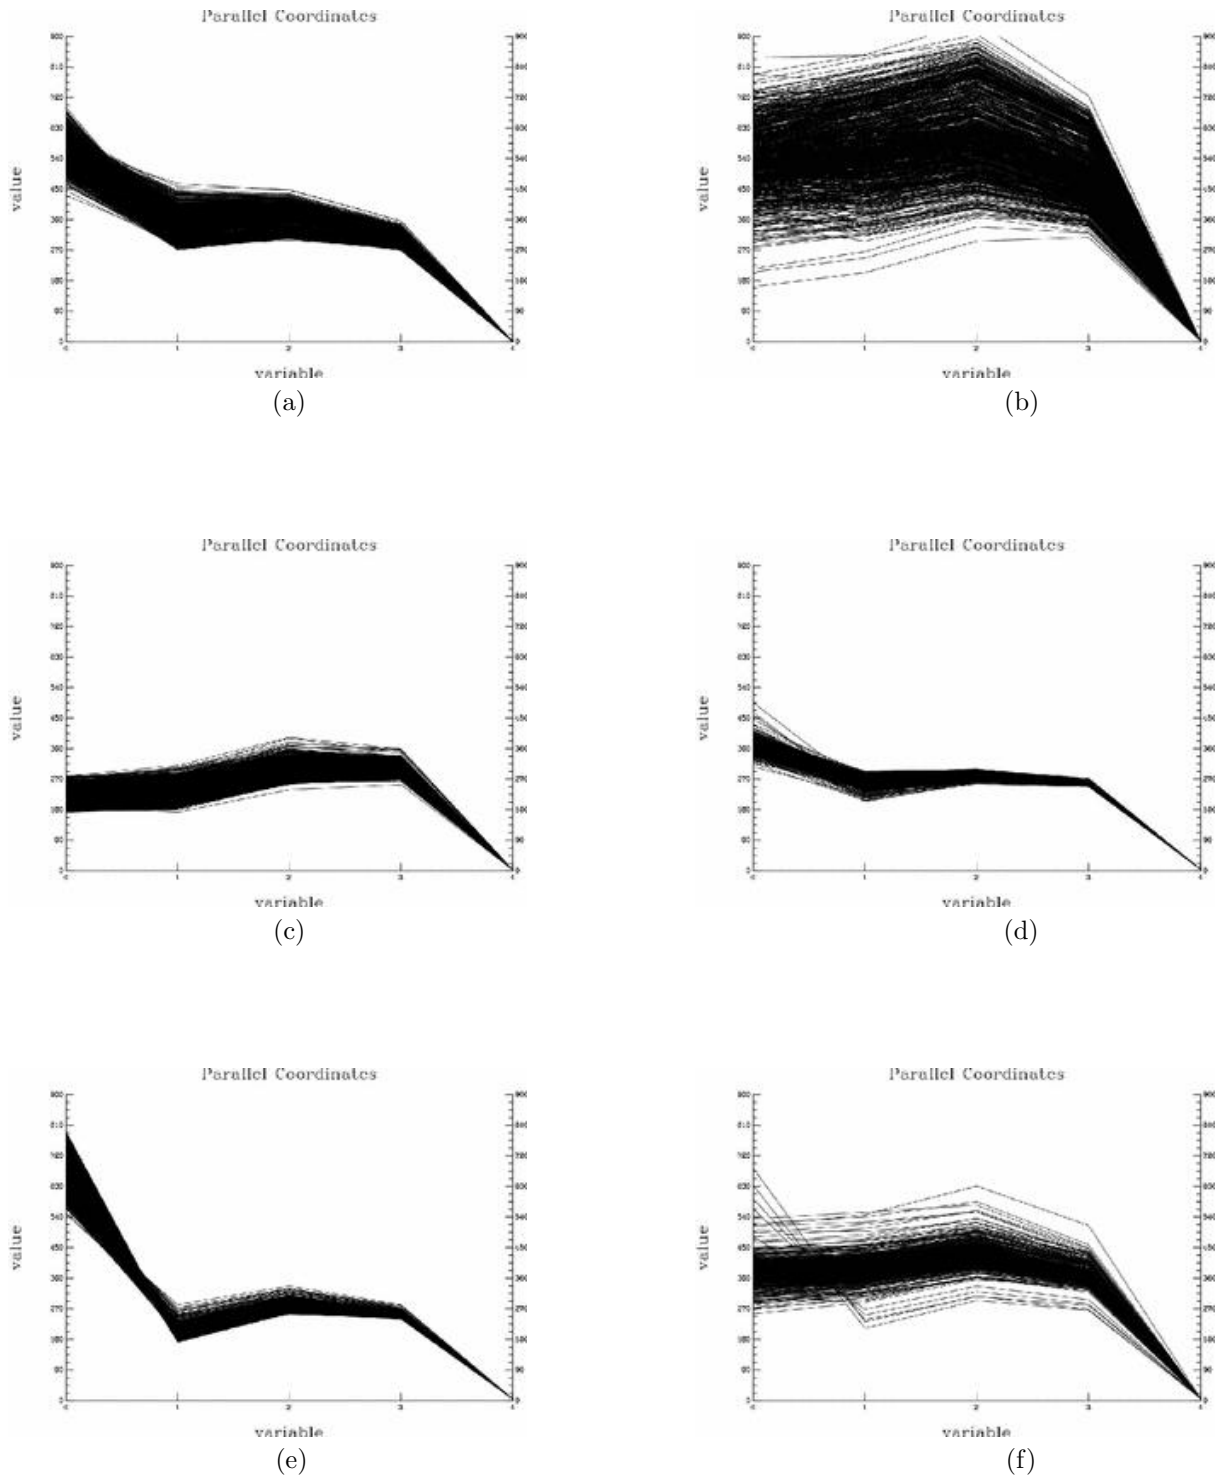


Figure 2. Parallel plots indicating the distribution of the features for the six classes before removal of outliers. The x-axis shows the five features (intensity values for the near-IR, red, green, and blue bands, and the class). The y-axis is the value of the corresponding feature. (a) Class 1 (b) Class 2 (c) Class 3 (d) Class 4 (e) Class 5 (f) Class 6




Figure 3. Parallel plots indicating the distribution of the features for the classes 2 and 6 after removal of outliers. The x-axis shows the five features (intensity values for the near-IR, red, green, and blue bands, and the class). The y-axis is the value of the corresponding feature. (a) Class 2 (b) Class 6

$$L_{Gini} = 1.0 - \sum_{i=1}^k (L_i/|T_L|)^2$$

$$R_{Gini} = 1.0 - \sum_{i=1}^k (R_i/|T_R|)^2$$

$$\text{Impurity} = (|T_L| * L_{Gini} + |T_R| * R_{Gini})/n$$

where $|T_L|$ and $|T_R|$ are the number of examples, and L_{Gini} and R_{Gini} are the Gini indices on the left and right side of the split, respectively. L_j and R_j are the number of instances of class j on the left and the right, respectively, and n is the total number of examples at a node. We used this splitting criterion as our previous experiences indicated that it worked well in practice.

The tree created using the entire training set is shown in Figure 4. The features B1, B2, B3, and B4 are the intensity values for near-IR, red, green, and blue bands respectively. The decision tree output lists the feature selected at each node, as well as the value it is compared against. The number after the colon indicates that the node in question is a leaf node, and the number is the class assigned to the leaf. At each leaf node, the numbers (a/b) indicate the (total number of samples/samples of the class not assigned to leaf node).

No pruning was applied to this tree to reduce the generalization error. An average of ten runs of ten-fold cross-validation with the training set resulted in an error rate of less than 1%. In each run, we divided the training set randomly into ten equal part, trained using nine of the parts, and tested on the tenth part, cycling through all the parts in turn.

Once the decision tree was created using 4860 pixels from the regions in Figure 1, panel (b), it was used to classify all the 160,000 pixels in the image in Figure 1, panel (a). This resulted in the image in Figure 5, panel (a), where each of the six classes is indicated by a different color. For comparison, the result of the k-means algorithm (with $k=6$) for unsupervised classification of the same image is given in Figure 5, panel (b). A comparison of the Figures 1 panel (a), and the two panels in Figure 5 indicates that the decision tree is able to generalize quite well in classifying all the pixels in the full image.


Once the first stage of classification of the pixels in the multi-spectral image is completed, we use it to identify “windows” or “tiles” in the image likely to indicate human settlements. The approach used for this was motivated by the observation that areas in the image with settlements tended to contain a mix of pixels of different classes with a random spatial distribution. Therefore, by dividing the image into non-overlapping tiles,

```


B2 < 487.5:
: B1 < 279.5: class 3 (811/4)
: B1 >= 279.5:
:   B1 < 453.5:
:   : B3 < 339:
:   : : B1 < 426: class 4 (819/0)
:   : : B1 >= 426: class 1 (2/0)
:   : B3 >= 339: class 6 (827/0)
:   B1 >= 453.5:
:   B2 < 268.5: class 5 (809/0)
:   B2 >= 268.5: class 1 (775/0)
B2 >= 487.5: class 2 (813/0)

```

Figure 4. A decision tree created using the entire training set identified in Figure 1, panel (b)



(a)



(b)

Figure 5. (a) The multi-spectral image from Nebraska, with the pixels classified using the decision tree. Each class is represented by a different color. (b) The same with the pixels classified using the k-means clustering algorithm, with $k=6$. Original images by Space Imaging

each of 10×10 pixels, we could use some simple statistics to determine if the tile was likely to contain human settlements or not. The focus in this stage was to remove areas with a homogeneous distribution of spectral classes from further consideration. The remaining areas were then processed further to determine if they contain man-made structures or not.

We tried the following simple methods to identify tiles with a mix of pixels:

- **Method 1:** In each tile, we count the number of pixels which have their four neighbors of the same class as the pixel in the center. If the percentage of such pixels in a tile exceeds a certain threshold, then we identify the tile as “unmixed” or unlikely to contain human settlements. It is also possible to include additional constraints to ensure that sufficient number of pixels from several different classes are included in each “mixed” tile.
- **Method 2:** This method uses the class-level co-occurrence matrix (CLCM) for each tile in the horizontal and vertical directions. The CLCM is analogous to the gray level co-occurrence matrix¹⁰ with the class labels being used instead of the gray levels. The sum of the off-diagonal elements of the 6×6 CLCM can be used to determine if a tile is mixed or not. If the percentage of off-diagonal elements is greater than a threshold, the tile is considered to be mixed. Note that we do not need to explicitly calculate the CLCM. We only need to count the number of times the class of a pixel is different from the class of the pixel to its right (for the horizontal CLCM) or below it (for the vertical CLCM).

- **Method 3:**

This is a generalization of methods 1 and 2. Here we use various statistics (features) derived from the CLCM as features to be used in an unsupervised classification mode. Depending on the set of statistics used, it has the potential of discriminating not only mixed and spatially homogeneous tiles but also different types of mixed tiles. For this paper, we use only two features, the angular second moment (ASM) and entropy (ENT),¹¹ defined for a 10×10 tile in a manner similar to features with the same names in the context of gray level co-occurrence matrix for a single band. We have used these two features as inputs to the Isodata¹² clustering algorithm to cluster all 10×10 non-overlapping tiles obtained as sub-images and covering the full image. We have specified the lower and upper bounds on the number of clusters to be 2 and 4, respectively. As an option, we have also used a threshold of 2.5 times the standard deviation for the cluster spread about its mean. The clustering procedure led to two clusters under each option. In a tile with mixed pixels one would expect the ASM to be low and ENT to be high. This observation allowed us to identify the two clusters representing the tiles with mixed and “unmixed” pixels respectively.

Note that the approach for identifying the mixed tiles, as well as the use of a threshold, enables us to tolerate some misclassification error in the first stage of our multi-stage process.

Once the mixed tiles have been identified, a mask is created that masks out all the unmixed tiles as shown in Figure 6, panel (a). This particular mask was generated using Method 2, with a threshold of 40%. The tiles that have been masked out are those which had the percentage of off-diagonal elements in the CLCM less than 40%. The unmasked areas in the image are the ones likely to contain human settlements. It is these areas that are considered for further processing using the panchromatic imagery.

2.2. Analysis of the panchromatic imagery

After the multi-spectral image has been processed using pixel-level classification, and the “mixed” tiles identified as the ones likely to contain human settlements, the scene is processed further using the panchromatic 1m GSD data. First, the mask that was generated using the multi-spectral data is scaled by a factor of 4 and transferred to the panchromatic image, as shown in Figure 6, panel (b). Next, we use the SUSAN¹³ edge and corner detector to identify the edges and corners of potentially man-made structures.

The output of SUSAN is illustrated using a sub-image from the panchromatic scene, Figure 7, panel (a). Panel (b) is the sub-image with the masked tiles from Figure 6, panel (b). The unmasked tiles in this sub-image include both areas with buildings and areas without buildings. Panels (c) and (d) are the outputs from the




Figure 6. (a) The multi-spectral image from Nebraska, with the “unmixed” tiles masked out. The area remaining is considered likely to have human settlements. (b) The panchromatic image of the same scene from Nebraska, with the “unmixed” tiles masked out. Original satellite images by Space Imaging.

SUSAN edge and corner detectors, respectively, after smoothing the image in panel (a). Next, the edge and corner images are again tiled using tiles with 40×40 pixels. Tiles that contain more than 40 edge pixels and more than 4 corner pixels are identified as containing built structures. Any other suitable threshold and tile size could have also been used. The tiles that do not contain built structures have been masked out in Panel (e). Finally, this mask is combined using a logical OR operation with the mask in panel (b) to give the final output in panel (e).

Once the masked image has been obtained from combining both the mask from the multi-spectral and the panchromatic imagery, it can be refined further to combine connected components and remove isolated small regions. The boundaries of the remaining regions can then be exported to a Geographic Information System for use in the production of maps.

3. EXPERIMENTAL RESULTS

In this section, we comment on our experiences in applying the multi-stage approach to IKONOS multi-spectral and panchromatic images.

3.1. Robustness experiments

To understand the robustness of our technique, we explored its applicability in different regions of the world. We are interested in understanding how much of the analysis conducted in one region can be directly applied to another region. For example, can we use the classification of the pixels in the multi-spectral image of one region to build a model that can accurately classify the pixels in another region?

To investigate this idea, we took the decision tree (Figure 4) created using the scene from Nebraska, and applied it to a scene from the Northern Mexico region (Figure 8) (a). The classes resulting from this supervised classification are shown in Figure 8 (b), along with the results from unsupervised classification on the same image in Figure 8 (c) using k-means with 6 classes. Though the two scenes from Nebraska and Northern Mexico are from two different regions of the world, they are similar enough that a decision tree model built to classify pixels in the multi-spectral image of one can be used successfully to classify the pixels in the multi-spectral image of the other. If the images were quite different, it is unlikely that such an experiment would have been equally




Figure 7. SUSAN edge and corner detector applied to a sub-image from Figure 6, panel (b). (a) Original sub-image. (b) Sub-image with the mask from Figure 6, panel (b). (c) Edges using SUSAN with option `(-e -t 40)`, after smoothing. (d) Corners using SUSAN with option `(-c -t 40)`, after smoothing. (e) The original sub-image with masked tiles indicating tiles with edge and corner frequency less than the threshold. (f) The original sub-image with both the masked tiles from panel (e) and the masked tiles from panel (b). Original satellite images by Space Imaging.

successful. This indicates that it might be possible to build models that would be tuned to different regions of the world, with possible sub-models to account for seasonal and other variations.

3.2. Computational accuracy

Our initial experiments indicate that our approach exploiting both the multi-spectral and the panchromatic images is likely to be more accurate than using either type of image by itself. For example, just using the panchromatic imagery, along with the edges and corners, may result in false positives in areas of the scene where there are roads and cars without any human settlements nearby. Or, a “mixed” tile in multi-spectral imagery may be mixed for reasons other than human settlements. We are currently conducting additional experiments to better understand how the two types of imagery help us to mitigate false positives.

Note also that the early stages of the multi-stage process do not have to be very accurate. Even if a pixel is mis-classified, such errors do not have a drastic effect on the final result. This is because we are working at the tile level, interested in regions that are not composed of isolated tiles, and processing the data in several stages.

3.3. Computational efficiency

A key part of the hierarchical approach is the savings in computational time. While the SUSAN code is quite efficient in smoothing a small (e.g., a 400×400 pixel image), and finding edges and corners in it, the cost of this processing can still be prohibitively expensive when millions of such images are being processed. By applying




Figure 8. The multi-spectral image from Northern Mexico. (a) Original image (b) The image pixels classified using the decision tree in Figure 4. (c) The image pixels classified using k-means clustering, with 6 classes. Original satellite image by Space Imaging.

SUSAN to only the mixed tiles as shown in Figure 6 (b), we can expect to reduce processing substantially. Even though the original scene from Nebraska appears to contain large areas of human settlements, the unmasked area in Figure 6 (b) is only 46 % of the original image. Thus, if the edge and corner detector were written to operate only on the unmasked regions, it would require approximately 50 % less computation time.

3.4. Experiments to identify mixed tiles

To understand the accuracy of the different methods in identifying mixed tiles, we compared the output from these methods with the ground truth generated using the panchromatic image of the Nebraska scene. This ground truth identified tiles as “mixed”, that is, requiring further processing, if the corresponding area in the panchromatic image had built-up structures. Note that the tiles are generated using the multi-spectral data, while the ground truth is based on the panchromatic imagery. For each method, we considered the number of true positives and negatives, as well as false positives and negatives. In our problem, a false positive results in additional computation in the later stages. However, a false negative implies a tile which should be processed further is being rejected. Therefore, a method with low false negatives is preferred.

Our early experiences with the different methods indicate that each method has its strengths and weaknesses. In our particular example, Method 1 with a threshold of 50 gave good results for the true positives and the false negatives. In contrast, Method 3 applied with the threshold option had fewer false positives and detected more true negatives. In future work, we plan to study the accuracy of combinations of these methods as well as their performance on images from other regions.

4. SUMMARY

In this paper, we have introduced a hierarchical approach to the automated production of maps of human settlement from the IKONOS multi-spectral and panchromatic imagery. A multi-stage approach allowed us to exploit both types of imagery to reduce false positives without excessive increase in the computational cost. For this paper, we prototyped our algorithms in ENVI/IDL.¹⁴ The work on parallel plots and decision trees was done using Sapphire software.⁹ Our future work includes understanding what types of false positives are reduced by each type of imagery, experimenting with robustness of the approach to geo-cultural, seasonal, illumination, and look-angle variations, as well as additional studies on the parameters used in each stage.

ACKNOWLEDGMENTS

UCRL-JC-150218: This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

A color version of this paper is available at the Sapphire project web-site.⁹

REFERENCES

1. "Space Imaging, Inc.." <http://www.spaceimaging.com>.
2. J. Heikkonen and A. Varfis, "Land cover/land use classification of urban areas: a remote sensing approach," *International Journal of Pattern Recognition and Artificial Intelligence* **12**(4), pp. 475–489, 1998.
3. T. Kohonen, *Self-Organizing Maps (3rd ed.)*, Springer, New York, 2001.
4. L. Breiman, J. Friedman, R. A. Olshen, and C. Stone, *Classification and Regression Trees*, CRC Press, Boca Raton, Florida, 1984.
5. Y. Zhang, "Optimization of building detection in satellite images by combining multispectral classification and texture filtering," *ISPRS Journal Photogrammetry and Remote Sensing* **54**(1), pp. 50–60, 1999.
6. S. K. Sengupta, C. Kamath, D. Poland, and J. A. H. Futterman, "Detecting human settlements in satellite images," in *Proceedings, Optical Engineering at the Lawrence Livermore National Laboratory, SPIE Lasers and Applications in Science and Engineering, San Jose, January 2003*.
7. J. Quinlan, "Induction of decision trees," *Machine Learning* **1**, pp. 81–106, 1986.
8. J. Quinlan, *C4.5: Programs for Machine Learning*, Morgan Kaufman, San Mateo, California, 1993.
9. "Sapphire: Large-Scale Data Mining and Pattern Recognition." <http://www.llnl.gov/casc/sapphire>.
10. R. M. Haralick, "Statistical image texture analysis," in *Handbook of Pattern Recognition and Image Analysis*, T. Y. Young and K.-S. Fu, eds., Academic Press, 1986.
11. R. Welch, S. Sengupta, and D. Chen, "Cloud field classification based upon high spatial resolution textural features part I: Gray-level co-occurrence matrix approach," *Journal of Geophysical Research* **93**, pp. 12663–12681, 1988.
12. G. Ball and D. Hall, "A novel method for data analysis and pattern classification," tech. rep., SRI, Menlo Park, 1965.
13. S. Smith and J. Brady, "SUSAN - a new approach to low-level image processing," *International Journal of Computer Vision* **23**(1), pp. 45–78, 1997.
14. "Research Systems Inc.." <http://www.rsinc.com>.

University of California
Lawrence Livermore National Laboratory
Technical Information Department
Livermore, CA 94551

