



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

K-Means Re-Clustering Algorithmic Options with Quantifiable Performance Comparisons

A. W. Meyer, D. W. Paglieroni, C. Astaneh

December 17, 2002

The International Society for Optical Engineering, Photonics
West, San Jose, CA
January 25-31, 2003

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.

K-Means Re-Clustering – algorithmic options with quantifiable performance comparisons

Alan Meyer, Senior Member, IEEE, David Paglieroni, Senior Member, IEEE, and
Cyrus Astaneh

Lawrence Livermore National Laboratory
P.O.Box 808, Livermore, CA, 94550

ABSTRACT –This paper presents various architectural options for implementing a K–Means Re–Clustering algorithm suitable for unsupervised segmentation of hyperspectral images. Performance metrics are developed based upon quantitative comparisons of convergence rates and segmentation quality. A methodology for making these comparisons is developed and used to establish K values that produce the best segmentations with minimal processing requirements. Convergence rates depend on the initial choice of cluster centers. Consequently, this same methodology may be used to evaluate the effectiveness of different initialization techniques.

1. INTRODUCTION

Image segmentation transforms pixel-level information from raw images to a higher level of abstraction in which related pixels are grouped into disjoint spatial regions. Image segmentation is a critical early step in a number of important applications and problem domains, including image understanding, automatic target cueing (automatically separating objects of interest from complex backgrounds), land use classification, etc.

The K-Means algorithm is a well-established unsupervised method for segmenting pixelated images^{1,2}. This approach does not require a priori knowledge of the pixel spectral classifications, but instead will attempt to discern these classifications over successive iterations: An effective but potentially time consuming process.

Image segmentation quality is influenced by three factors: the number of the image spectral bands, the number of spectral classes (K), and the method for initializing the spectral classes. This paper reports ON the effects of these three factors by measuring the effect on image segmentation quality and computational complexity.

2. BACKGROUND

Image segmentation is the process of extracting regions by dividing an image into disjoint sets of pixels that belong together. The input is a 2D or 3D image of pixel values, and the output is a 2D image in which each pixel is labeled with the integer-valued ID of the region to which it was assigned. This 2D image is known as a region map.

Several classes of segmentation algorithms have been developed over the last several decades, including region growers, pixel classifiers, deformable model-based methods, and morphological methods. Region growers assign pixels to regions by searching local neighborhoods centered on a pixel that is already assigned to a region: a seed pixel^{3,4}. Methods based on deformable models deform curves towards object boundaries or edge features in the image. In their final deformed state, these curves define boundaries between segmented regions^{5,6,7,8,9}. Morphological methods developed for single-band images use low-level image operators based on structuring elements to erode and dilate regions^{10,11}.

Pixel classifiers, originally intended for multi-band images, assign individual pixels to specific classes based on their spectral properties. Pixel classifiers are further identified as supervised, unsupervised, or a

hybrid such as the stochastic expectation maximization (*SEM*) algorithm that has both supervised and unsupervised elements¹².

The K-Means algorithm is a well-established method for segmenting images². This study emphasizes a variant of the K-Means algorithm that first assigns each pixel to a spectral group, or class, and then partitions these groupings by their spatial affiliations. This second partitioning is referred to as spatial re-clustering, and thus the algorithm is known as K-Means Re-clustering or simply KMR. Taxonomically, the KMR algorithm is an unsupervised algorithm. In theory, unsupervised algorithms should automatically determine appropriate spectral classes. Typically, these classes evolve iteratively until certain criteria are met.

K-Means iterations can be time consuming. One way to reduce the number of iterations is to select the class centers in a deliberate way (i.e., using what will be referred to in this paper as *informed initialization*), as opposed to randomly (the traditional approach). Another way is to relax the K-Means convergence criterion by increasing the *convergence percentage* (i.e., the percentage of pixels that are allowed to change classes on the last iteration). Figure 1 depicts the KMR algorithm identifying two approaches to initializing the spectral classes. Section 3 will develop the traditional KMR algorithm using random initialization of the spectral classes, and Section 4 will describe the informed initialization approach.

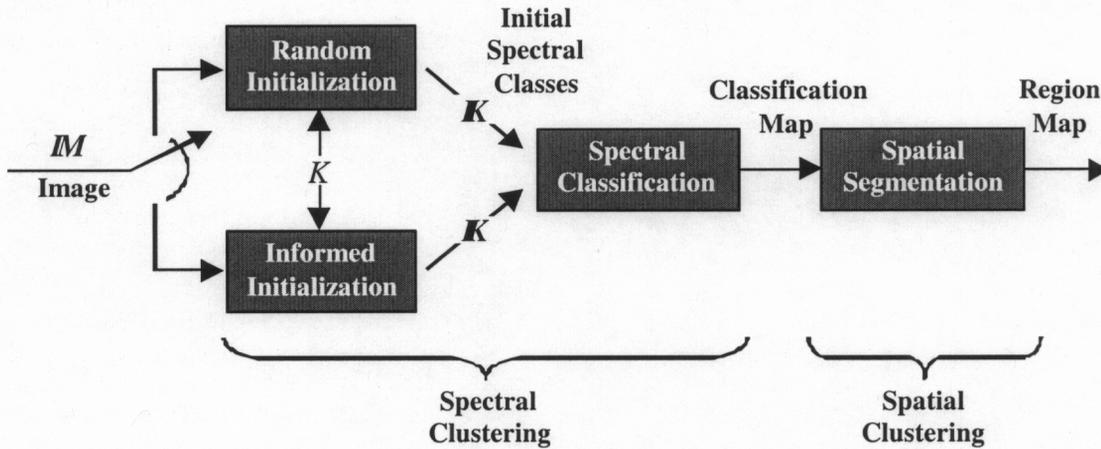


Figure 1. Block diagram of the K-Means Re-clustering algorithm

3. THE K-MEANS RE-CLUSTERING ALGORITHM

The KMR algorithm performs segmentation by grouping each pixel in the image based upon spectral and spatial affiliation. Conceptually, this requires the sequential execution of three distinct processing phases: initialization, spectral clustering and spatial re-clustering. KMR accepts as inputs the number of spectral classes and an image to be segmented. The image, \mathbf{M} , is defined as a structured set of pixels \mathbf{p} such that

$$\mathbf{M} := \{ \mathbf{p}_{i,j} \in \mathbf{R}^n : i, j, n \in \mathbf{N} \}$$

$$i := \text{row dimension}$$

$$j := \text{column dimension}$$

$$n := \text{spectral dimension}$$
(1)

The spectral classes are uniquely determined by a representative spectrum. The set of these spectra or cluster centers is given by the set \mathbf{K} and enumerated by the variable K . Thus:

$$\begin{aligned}
\mathbf{K} &:= \{ \mathbf{c} \in \mathbf{R}^n \} \\
K &= \text{length}(\mathbf{K}) \\
k &\in \mathbf{N}: 1 \leq k \leq K
\end{aligned} \tag{2}$$

3.1. Random initialization of the spectral classes

Traditionally, spectral classes are initialized by with the spectra from K pixels selected at random. Random initialization is one of the simplest K-Means class initialization strategies.

$$\begin{aligned}
\mathbf{c}_k &= \mathbf{p}_{\mu, \nu} \\
\mu &:= \{ \mu \text{ randomly selected from } [1, N] \mid \mu, N \in \mathbf{N} \} \\
\nu &:= \{ \nu \text{ randomly selected from } [1, N] \mid \nu, N \in \mathbf{N} \}
\end{aligned} \tag{3}$$

Where N is the number of rows of pixels in the image and M is the number of columns of pixels in the image.

3.2. Implementation of Spectral Clustering

The input to the spectral clustering algorithm is the set \mathbf{K} of cluster centers and the image \mathbf{M} . The output is a classification map in which each pixel is assigned to a spectral class. Spectral clustering is an iterative algorithm. Upon each iteration, every pixel is assigned to a spectral class, statistical properties for each class computed, and a new “representative” spectrum for each class selected as a function of the mean of all the assigned member pixel spectra. This iterative process is continued until a convergence criterion is satisfied.

Pixel assignment, with every iteration, is made to the class for which the two norm is minimized. Thus the assignment of a pixel, $\mathbf{p}_{i,j}$, to a spectral class, \mathbf{C}_k , with representative spectra, \mathbf{c}_k , is given by $s_{i,j}$, such that:

$$s_{i,j} = k \mid \min_k (\varepsilon_k) \tag{4}$$

$$\varepsilon_k = \|\mathbf{p}_{i,j} - \mathbf{c}_k\|_2 \tag{5}$$

Convergence is based upon two possible criteria: No more than a designated percentage (typically 5% or less) of the pixels switch classes between iterations, or a pre-specified maximum number of allowed iterations (typically 50) has occurred. A maximum number of allowed iterations are needed because the KMR algorithm is not guaranteed to converge.

Once convergence is achieved, each spectral class is pruned to constrain its variance to an integer multiple of the sample standard deviation, σ_k , of its members. Thus:

$$\begin{aligned}
\sigma_k^2 &= \text{var}(P_k) \\
P_k &:= \{ p_{a,b} \in \mathbf{C}_k \}
\end{aligned} \tag{6}$$

The rejected members are defined as

$$\begin{aligned}
\|\mathbf{p}_{a,b} - \mathbf{c}_k\|_2 &> n\sigma_k \\
n &\in \mathbf{N} \text{ and typically } n = 3
\end{aligned} \tag{7}$$

Those pixels no longer assigned to a spectral class are re-assigned to a reject ($K + 1$) class and a final iteration of the spectral clustering algorithm is then performed on the non-reject pixels.

3.3. Merging or splitting of spectral classes

Unsupervised algorithms generally allow K to vary by splitting and merging spectral classes. This option is disabled in KMR to facilitate the objective of assessing segmentation quality as a function of several fixed values of K . However, since images usually contain anomalous pixels or “transition” pixels (such as edges along borders between regions formed as combinations of different spectral classes, KMR does need a provision to handle pixels that belong to none of the K spectral classes. KMR handles anomalous pixels by creating a separate *reject class* (class $K+1$) after the last iteration.

3.4. Implementation of spatial re-clustering

For image segmentation, the ultimate goal is to partition an image into regions that can be readily identified. After spectral clustering, KMR implements a spatial clustering using classical region growing techniques⁴. Spatial re-clustering results in groupings of spatially connected pixels that belong to the same spectral cluster. Spatial re-clustering proceeds by searching the local neighborhood centered on a pixel selected at random. All neighboring pixels belonging to the same spectral class are assigned to the same region. Once a complete region has been grown, a new seed is selected from the remaining ungrouped pixels and its region is grown. This process is repeated until all of the pixels are accounted for

3.5. Problems with random initialization

This implementation of KMR requires $O(BKI)$ operations per pixel, where B is the number of spectral bands, K is the specified number of spectral classes, and I is the number of spectral clustering iterations. The computational cost of KMR thus varies more or less linearly with the number of spectral bands, the number of spectral classes and the number of spectral clustering iterations. The number of iterations can be reduced by relaxing the convergence criterion (i.e., by increasing the percentage of pixels allowed to switch spectral classes on the convergence iteration). However, segmentation quality can be detrimentally affected by relaxing the convergence criterion too much. Alternatively, the number of iterations may go way up with little improvement in segmentation quality if the convergence criterion is too strict.

Random initializations can give rise to unnecessarily slow convergence rates, particularly if the randomly chosen spectral class centers are unreasonable. Also, random initialization can create situations in which the K-Means algorithm cannot converge to a reasonable spectral class assignment, especially if the randomly chosen spectral class centers are not sufficiently disparate, or are completely unrealistic.

4. INFORMED INITIALIZATION ALGORITHM

Informed initialization is the opposite of random initialization in the sense that the spectral class centers are chosen in a completely non-random and deliberate way. Informed initialization addresses the shortcomings of random initialization that were just described.

First, the energy in each pixel is estimated by averaging all of its spectral samples (this assumes that the spectral samples are all non-negative). By this measure, the band-averaged image is a 2D image of pixel energies:

$$a_{i,j} = \frac{1}{N} \sum_n p_{i,j}(n) \tag{8}$$

$N \equiv$ number of spectral bands

The band-averaged values are then quantized to K discrete levels. The spectral class center for class k is then taken to be the pixel spectrum formed as the average of all pixel spectra that have a quantized band-averaged value of k .

$$c_k = \frac{1}{M} \sum_{\mu} \sum_{\nu} p_{\mu,\nu}$$

$$\mu \in \mathbf{X} \mid \mathbf{X} := \{i \text{ corresponding to quantization level } k\}$$

$$\nu \in \mathbf{Y} \mid \mathbf{Y} := \{j \text{ corresponding to quantization level } k\}$$

$$M \equiv \text{total number of pixels in quantization level } k$$
(9)

5. RESULTS OF THIS STUDY

Figure 2 plots the number of iterations to convergence as a function of convergence percentage for both informed and random initialization. The random initializations require the largest number of iterations for low convergence percentages. There is often a critical convergence percentage below which the number of convergence iterations increases sharply and above which the number of convergence iterations remains relatively small. For a convergence criteria greater than or equal to 2%, the informed initialization will result in less computational burden than a random initialization. It is often possible to reduce KMR computational complexity without significantly altering the region maps by choosing a convergence percentage as large as 3%. In addition to driving class initialization towards spectrally divergent classes informed initialization sometimes results in fewer KMR iterations than random initialization with a corresponding improvement in efficiency.

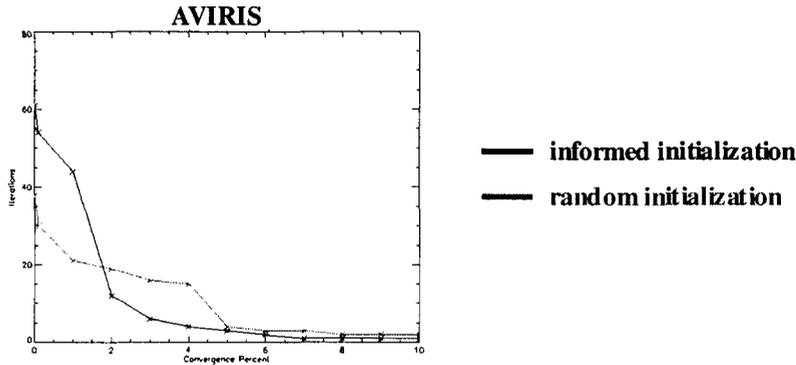


Figure 2. Number of KMR iterations as a function of convergence percentage

For a fixed convergence percentage, the number of KMR iterations varies as a function of K . Figure 3 shows the effect that K has on the number of iterations for fixed 2% convergence criteria. Informed initialization often, but not always, results in fewer KMR iterations than random initialization. The correlation between K and the number of iterations is not necessarily strong, and appears to be scene-dependent. Also, the number of KMR iterations does not appear to be strongly correlated with spectral resolution. Thus it is possible to conclude that for a modest value of K , there is little difference between the performance of either random initialization or informed initialization.

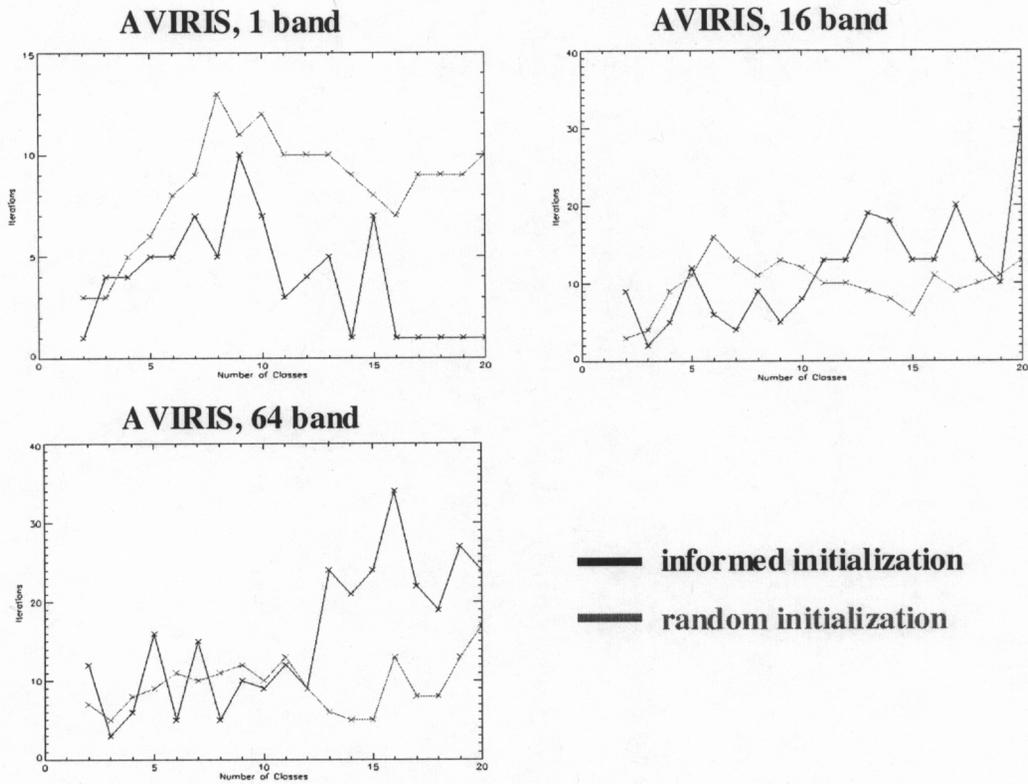


Figure 3. Number of KMR iterations as a function of the number of classes for a fixed 2% convergence criterion.

Figure 4 shows the profound effect that K has on KMR segmentation. Even though the images are busy, small K values ($K=2,3,4$) tend to produce subjectively better segmentations than larger K values, which tend to over-segment the images. Also, the informed and random initializations tend to produce comparable results.

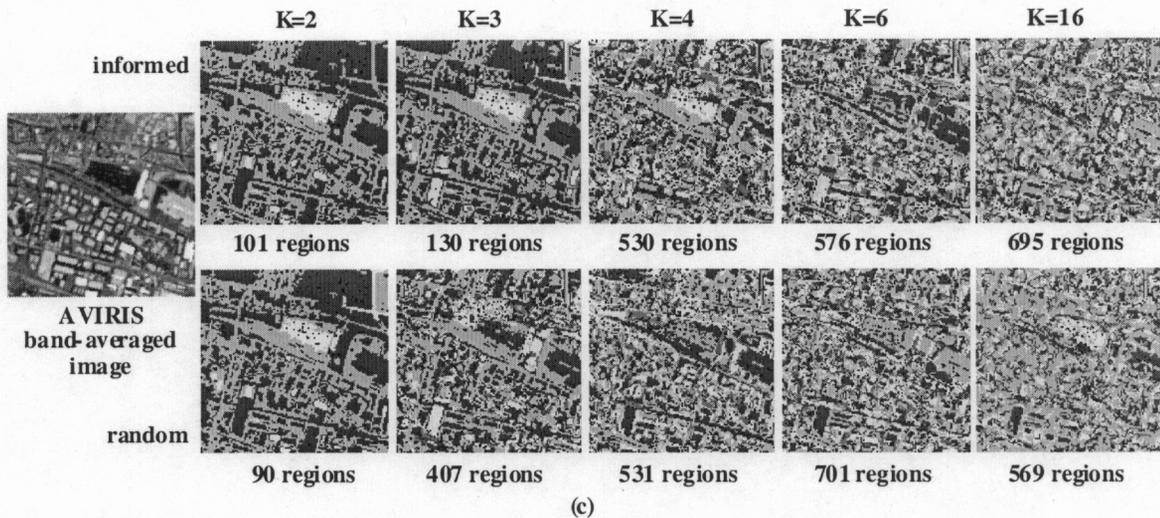


Figure 4. Comparing informed and random initialization for a KMR segmentation at a 3% convergence criteria and 16 band spectral resolution

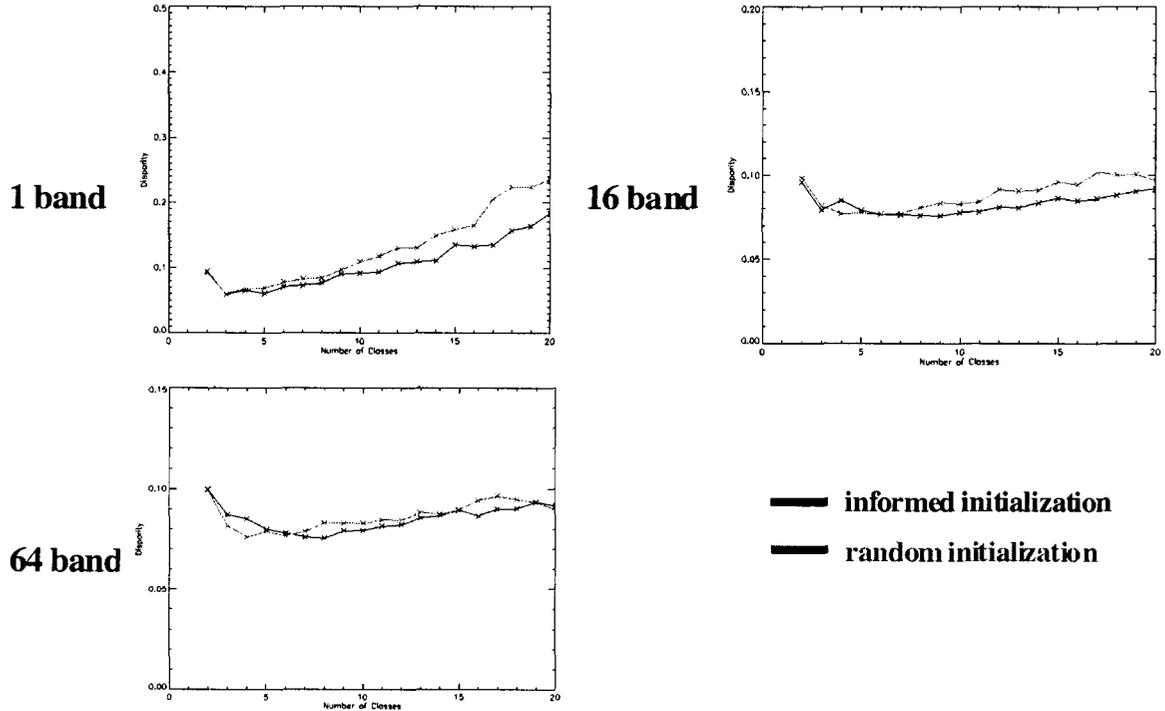


Figure 5. Segmentation quality vs. K for random and informed initializations at different spectral resolutions.

In figure 5 a disparity metric is introduced that quantifies the difference between a human derived edge map, E , and the KMR region map R . The boundary map, B , for this region is given as

$$B(i, j) = \begin{cases} 1 & R(i, j) \neq R(i+1, j) \text{ or } R(i, j+1) \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

Let n_E is the number of edge pixels in E , and n_B is the number of boundary pixels in B , then disparity, Δ_{BE} , is given by

$$\Delta_{BE} = \begin{cases} 0 & n_B = n_E = 0 \\ 1 & n_B \text{ or } n_E = 0, \text{ but not both} \\ (n_{BE} + n_{EB}) / (n_B + n_E) & n_B, n_E \neq 0 \end{cases} \quad (11)$$

Where n_{BE} is the number of boundary pixels in B not associated with an edge in E , and n_{EB} is the number of edge pixels in E not associated with a boundary in B (two pixels are said to be associated if they are close to one another – typically within 2 pixels). Disparity provides a quantitative measure of segmentation quality. It is clear from figure 5 that accuracy of KMR does not necessarily increase with spectral resolution. Also, KMR tends to perform better with smaller values of K .

6. CONCLUSIONS

Image segmentation quality and computational complexity are key performance metrics for an image segmentation algorithm. For K-Means type algorithms, these metrics are, in turn, influenced by three factors, the number of the image spectral bands, the number of spectral classes (K), and the method for initializing the spectral classes.

The number of the spectral bands often has little effect on the quality of the resulting image segmentation. However, the computational complexity of KMR is directly proportional to the number of

spectral bands. In many cases, it may thus be advisable to spectrally blur the hyperspectral image down to only a few spectral bands prior to performing KMR segmentation.

Larger numbers of spectral classes (e.g., K values beyond 5 or 10) often cause images to be over-segmented. At the same time, the computational complexity of KMR is directly proportional to K . The results from this study indicate that small values of K ($K \in [2, 3, 4, 5]$) are often most effective in segmenting images.

Finally, informed initialization addresses serious shortcomings of random initialization by reducing the likelihood of encountering unnecessarily slow K-Means convergence rates and finding unrealistic initializations for spectral class centers.

7. ACKNOWLEDGEMENTS

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

8. REFERENCES

- [1] J. MacQueen, "Some Methods for Classification and Analysis of Multivariate Data", Proc. 5th Berkeley Symp. On Probability and Statistics, University of California Press, Berkeley, 1967.
- [2] G. H. Ball and D. J. Hall, "Isodata, an Iterative Method of Multivariate Analysis and Pattern Classification", Proc. IFIPS Congress, 1965.
- [3] J. L. Muerle and D. C. Allen, "Experimental Evaluation of Techniques for Automatic Segmentation of Objects in a Complex Scene", in Pictorial Pattern Recognition (G. C. Cheng et al., Eds.), Thompson, Washington, 1968.
- [4] S. W. Zucker, "Region Growing: Childhood and Adolescence", CGIP, vol.5, no.3, September 1976, pp.382-399.
- [5] C. Xu, D. L. Pham, and J. L. Prince, "Image Segmentation Using Deformable Models", in Handbook of Medical Imaging, Volume 2, M. Sonka and J. M. Fitzpatrick (eds.), pp. 129-174, 2000.
- [6] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active Contour Models", Int'l Journal of Computer Vision, Volume 1, No. 4, pp. 321-331, 1987.
- [7] T. McInerney and D. Terzopoulos, "Topologically Adaptable Snakes", in 5th International Conference on Computer Vision, pp. 840-845, 1995.
- [8] J. A. Sethian, Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science, Cambridge University Press, 2-nd Edition, 1999.
- [9] S. R. Gunn and M. S. Nixon, "Improving Snake Performance via a Dual Active Contour", in International Conference on Computer Analysis of Images and Patterns, pp. 600-605, 1995.

- [10] P. Soille, *Morphological Image Analysis: Principles and Applications*, Springer Verlag, 1999.
- [11] J. Goutsias and S. Batman, "Morphological Methods for Biomedical Image Analysis", in *Handbook of Medical Imaging*, Volume 2, M. Sonka and J. M. Fitzpatrick (eds.), pp. 175-272, 2000.
- [12] Pascale Masson and Wojciech Pieczynski, "SEM Algorithm and Unsupervised Statistical Segmentation of Satellite Images", *IEEE Trans. Geosci. Remote Sens.*, vol.31, no.3, May 1993, pp.618-633.