Experimental Determination of Visual Color and Texture Statistics for Image Segmentation

Junqing Chen[a] and Thrasyvoulos N. Pappas[b]

[a]Unilever Research, Edgewater, NJ
[b]Dept. of Electrical and Computer Engineering, Northwestern University, USA

ABSTRACT

We consider the problem of segmenting images of natural scenes based on color and texture. A recently proposed algorithm combines knowledge of human perception with an understanding of signal characteristics in order to segment natural scenes into perceptually/semantically uniform regions. We conduct subjective tests to determine key parameters of this algorithm, which include thresholds for texture classification and feature similarity, as well as the window size for texture estimation. The goal of the tests is to relate human perception of isolated (context-free) texture patches to image statistics obtained by the segmentation procedure. The texture patches correspond to homogeneous texture and color distributions and were carefully selected to cover the entire parameter space. The parameter estimation is based on fitting statistical models to the texture data. Experimental results demonstrate that this perceptual tuning of the algorithm leads to significant improvements in segmentation performance.

Keywords: Adaptive clustering algorithm, spatially adaptive dominant colors, local median energy, content-based image retrieval (CBIR), perceptual models, natural image statistics, feature extraction, statistical modeling, optimal color composition distance, steerable filter decomposition

1. INTRODUCTION

We consider the problem of segmenting images of natural scenes based on color and texture. A recently proposed algorithm[1–5] combines knowledge of human perception with an understanding of signal characteristics in order to segment natural scenes into perceptually/semantically uniform regions. Segmentation of images of natural scenes is particularly difficult because, unlike artificial images that are composed of more or less pure textures, the texture characteristics are not uniform due to effects of lighting, perspective, scale changes, etc. To account for such characteristics, the algorithm is based on spatially adaptive color and texture features, and incorporates perceptual knowledge both in the feature extraction techniques and the design of the segmentation procedure.

In this paper, we discuss subjective tests to determine key parameters of this algorithm. Such parameters include thresholds for texture classification and feature similarity, as well as the window size for texture estimation. The purpose of the tests is to relate human perception of isolated (context-free) patches of natural textures to image statistics of those patches, and in particular, to the statistics of the color and texture features on which the segmentation algorithm is based. Thus, our goal is to link the statistics of natural textures to human perception. This is in contrast to other methods for locating texture within an image (*e.g.*, Refs. 6–8), which are more *ad hoc.*

The idea of linking human perception to the statistics of natural images is well established. One approach to studying the fundamental properties of human visual perception is to consider the natural environment in which it has evolved.[9] The concept that the statistics of natural stimuli must have influenced the development of the human visual system was originally introduced by Field,[10,11] and has led to extensive measurements of the spatial and temporal characteristics of natural scenes,[12–14] as well as the statistics of natural illuminants and reflective surfaces.[15]

In our subjective experiments, we use texture patches that correspond to homogeneous texture and color distributions. The parameter estimation is based on fitting statistical models to the texture data. A preliminary version of this work was presented in Ref. 16. Since then, we have collected additional texture data, to make sure that it covers the entire parameter

For further information, please send correspondence to T.N.P.: E-mail: pappas@ece.northwestsern.edu; address: ECE Dept., Northwestern University, 2145 Sheridan Rd, Evanston, IL 60208

space, and rerun the subjective experiments. This paper presents a comprehensive review of the experimental procedure, and also, points out a number of ways to expand and improve it. Experimental results demonstrate that the perceptual tuning of the algorithm leads to significant improvements in segmentation performance.

The paper is organized as follows. In Section 2, we present a brief overview of the segmentation algorithm. In Section 3, we discuss the subjective experiments. An analysis of the experimental results can be found in Section 4. In Section 5, we conclude with a comparison of algorithm performance before and after the perceptual tuning.

2. SEGMENTATION ALGORITHM OVERVIEW

In this section, we overview the segmentation algorithm presented in Refs. 1–3, 5. The algorithm is based on spatially adaptive color and texture features. As illustrated in Fig. 1, two types of features are developed, one describes the local color composition, and the other the spatial characteristics of the grayscale component of the texture. These features are first developed independently, and then combined to obtain an overall segmentation.

The color features describe the color composition in terms of the dominant colors and associated percentages in the vicinity of each pixel. They are based on the estimation of the spatially adaptive dominant colors, which on one hand, reflects the fact that the human visual system cannot simultaneously perceive a large number of colors, and on the other, the fact that image colors are spatially varying. The spatially adaptive dominant colors are obtained using the adaptive clustering algorithm (ACA) for segmentation.[17] The color feature representation is as follows:

$$f_c(x,y,N_{x,y}) = \{(c_i(x,y,N_{x,y}), p_i(x,y,N_{x,y})), \ \ i = 1,\ldots,M, p_i(x,y,N_{x,y}) \in [0,1]\} \tag{1}$$

where each of the dominant colors, $c_i(x,y,N_{x,y})$, is a three dimensional vector in *Lab* space and $p_i(x,y,N_{x,y})$ is the corresponding percentage. $N_{x,y}$ denotes the neighborhood around the pixel at location (x,y) and M is the total number of colors in the neighborhood. A typical value is $M = 4$. Finally, a perceptual metric (OCCD)[18] is used to determine the similarity of two color feature vectors.

The spatial texture features describe the spatial characteristics of the grayscale component of the texture, and are based on a multiscale frequency decomposition such as the steerable pyramid[19] or the Gabor transform.[20] We use the local median energy of the subband coefficients as a simple but effective characterization of spatial texture. Median operators tend to respond to texture within uniform regions and suppress responses associated with transitions between regions. The texture features consist of a classification of each pixel into one of the following categories: *smooth, horizontal, vertical, +45°, -45°*, and *complex.*

Let $s_0(x,y)$, $s_1(x,y)$, $s_2(x,y)$, and $s_3(x,y)$ represent the subband coefficient at location (x,y) that corresponds to the horizontal, diagonal with positive slope, vertical, and diagonal with negative slope directions, respectively. We will use $s_{\max}(x,y)$ to denote the maximum absolute value of the four coefficients, and $s_i(x,y)$ to denote the subband index that corresponds to that maximum. A pixel (x,y) is classified as smooth if the median of $s_{\max}(x',y')$ over a neighborhood of (x,y) is below a threshold T_0. The size of the neighborhood is $W \times W$ pixels. In Refs. 3, 5, the threshold T_0 was determined using a 2-level K-means over the image. In the next section, we will see how this threshold can be determined by subjective tests. If the pixel is not smooth, then it is further classified as follows. We compute the percentage for each value (orientation) of the index $s_i(x',y')$ in the neighborhood of (x,y). If the maximum of the percentages is higher than a threshold T_1 (*e.g.,* 36%) and the difference between the first and second maxima is greater than a threshold T_2, (*e.g.,* 15%), then there is a dominant orientation in the window and the pixel is classified accordingly. Otherwise, the pixel is classified as complex. Note that the first threshold ensures the existence of a dominant orientation and the second ensures its uniqueness. Again, these thresholds can be determined by subjective tests. In Refs. 4, 21 we showed that, while the proposed approach depends on the structure of the frequency decomposition, it is relatively independent of the detailed filter characteristics.

In Refs. 3, 5, the size W of the neighborhood for the calculation of the texture statistics described in the previous paragraph was based on experience. Accurate border localization and adaptation to local texture characteristics dictate a small window size W, while accurate texture estimation dictates a large W. As we will see in the next section, for each scale, one should select the smallest window size that captures the texture characteristics. This can be based on subjective tests.

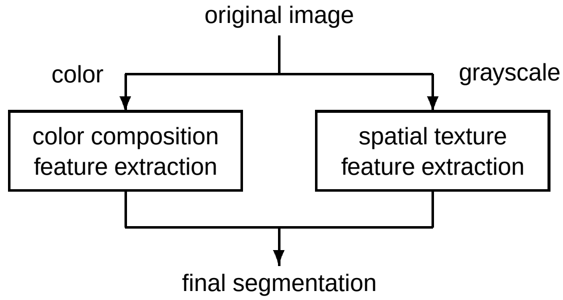

Figure 1. Schematic of segmentation algorithm

The segmentation algorithm then combines the color and spatial texture features to obtain segments of uniform texture within two steps. The first step relies on a multigrid region growing algorithm to obtain a crude segmentation. The segmentation is crude due to the fact that the estimation of the spatial and color texture features requires a finite window. This step requires a similarity threshold for classifying color patches in the same texture category. This is a critical parameter that can also be based on subjective experiments; in Refs. 3, 5, it was set based on our experience. The second step uses an elaborate border refinement procedure, which extends the idea of the ACA[17] to color texture, and progressively relies on the color segmentation to obtain accurate and precise border localization.

3. SUBJECTIVE EXPERIMENTS

Several key parameters of the segmentation algorithm described in the previous section can be determined by subjective tests. The first such parameter is the threshold T_0 for the smooth/nonsmooth classification. In Refs. 1–5 we used the K-means algorithm to determine this threshold. This relies solely on individual image statistics. It is more appropriate, however, to base this decision on how humans perceive textures and the corresponding texture statistics. Two additional parameters (T_1 and T_2) are necessary in order to further classify the nonsmooth areas into horizontal, vertical, +45°, -45°, and complex categories. Another critical parameter is the window size that is used for the determination of the texture features. As we discussed in the previous section, in order to allow accurate border localization and adaptation to local texture characteristics, it is important to keep this parameter as small as possible. On the other hand, the window size should be big in order to obtain accurate estimates of the texture characteristics. Thus, it is necessary to select the smallest window size that captures the texture characteristics at a given scale. This can also be obtained through subjective tests. Finally, another important parameter is the threshold for the color composition feature similarity.

The goal of the experiments described in Refs. 22–24 was to (a) determine perceptual categories for photographic images; (b) derive semantic names for these categories; and (c) identify low-level features that describe each category. These experiments dealt with the image as a whole. In contrast, our experiments isolate small patches of images corresponding to homogeneous texture and color distributions. It is important that such patches be considered out of context, just as the algorithms do not make use of any context information. Moreover, the subjects must not have seen any of the images from which the color-texture patches were obtained, and thus, the judgment should be solely based on what they see on the screen. Of course, this does not guarantee that the observers cannot derive any context from the scenes themselves (*e.g.*, familiar objects). Our goal is to determine what information such small image patches (devoid of context) convey to human observers and to relate those to the image statistics. Thus, our experiments are aimed directly at the problem of identifying perceptually important low-level features that can be used for image segmentation and analysis.

Experimental setup

The setup for our subjective experiments has been implemented in JAVA and has been published on the web.* The subjects were people with normal or corrected vision and normal color vision. The viewing distance was about two feet from the

*http://dunav.ece.northwestern.edu/FeatureTest/

Figure 2. Examples of Color Texture Patterns used

computer display, *i.e.*, a typical viewing distance for a person sitting in front of a computer. The subjects were advised not to move their head too close to the display. There were no other restrictions on viewing conditions. The test images were displayed on a CRT display with resolution 1280×1024 pixels; the color resolution was 24 bits. Throughout the experiment, the images were displayed against a neutral gray background. However, the subjects were allowed to adjust the background color for the best/clearest view of the color textures.

The stimuli consisted of 277 uniform color texture segments. Those were obtained from 37 photo CD images of 35mm KodakChrome slides, as well as images captured with a 5 Mpixel Nikon Coolpix 5400 digital camera with uncompressed TIFF settings at 2592×1944 pixel resolution. The photo CD images were available at four or five scales, while the digital camera pictures were downsampled to obtain additional scales. Thus, in contrast to many of the images used to test the algorithm,[5] the experimental stimuli consisted of high-quality uncompressed images. Several examples are shown in Fig. 2. The only constraint on the subject matter was that all the stimuli were images of "natural textures," *i.e.*, no synthetic, medical, microscopic, etc. images were allowed. We used a square shaped window throughout our experiments.

The experimental stimuli were carefully selected to cover the entire parameter space, namely, the local median of the maximum absolute value $s_{\max}(x,y)$ of the subband coefficients of the steerable filter decomposition (in short, local median energy). In particular, we ensured that enough samples fell into the range of 10 to 30 in the local median energy space, where the threshold for determining the texture orientations was expected to fall. This was particularly important for the final results. Overall, we had to used a larger number of texture patches than that used in the experiment reported in Ref. 16, increasing them from from 150 to 277. We also eliminated patches that were not classified as uniform texture by the users in the earlier experiment. The number of subjects that participated in the study was 20. Their ages ranged from 22 to 50 and included both experts and non-experts in image processing.

In the future, we plan to run additional experiments with more subjects using the web-based interface. For this, caution must be taken in controlling the remote display of the test images. In particular, it is important to control the visual angle (cycles per pixel) of the displayed images. This can be accomplished with a combination of image size (by interpolation) and viewing distance adjustments. The necessary adjustments can be determined by displaying a test bar that the user can measure. As the human visual system has strong color adaptation, and in agreement with previous studies that report that texture perception is not too sensitive to color variations, it is expected that differences in the color renditions by remote displays will have little impact on the subjective experiments. Hence, we do not plan any elaborate color control of the display. We now describe the subjective experiments in more detail.

Experiment 1: Texture classification

In this experiment, the subjects were asked to classify a color texture pattern into one of the following three categories:

- *Smooth:* Images with uniform or slowly varying intensity that contain no objects or sharp boundaries.
- *Texture:* Images of approximately uniform texture patterns. Since natural textures are often statistically nonuniform, slowly varying texture patterns should be included in this category.
- *Other:* Neither smooth nor texture, such as images with multiple objects or regions.

Note that the last category is important in order to filter out segments that do not contain perceptually uniform textures. Eventually, however, all nonuniform segments can be eliminated from the set of test patterns in order to speed up the experimental procedure.

The subjects were also asked to further classify the "texture" images into one of the following categories based on the perceived dominant orientations:

- *Horizontal*
- *Vertical*
- *+45^o*
- *-45^o*
- *Complex*

As we pointed out above, in all of these experiments, the "no context" requirement is crucial in order to approximate as closely as possible the context that the segmentation algorithm faces.

The size of texture window is an important parameter of this experiment. The window must be large enough for a human observer to perceive any texture. On the other hand, it must be kept small in order to avoid significant changes in the spatially varying texture characteristics. The window size that we used for the subjective experiments was 23×23 pixels. As we saw above, any windows that may contain region boundaries will be filtered out by the experimental procedure.

Experiment 1a: Minimum window size

Humans perceive texture at different scales. At each scale a minimal window size is required in order to identify a texture. This is true for both human perception and computer-based texture recognition. In Experiment 1, we used a fixed window size for all scales. At that window size, several texture scales can be perceived. However, by displaying several texture scales, we can find the minimum scale that can be perceived at that window size. Conversely, since the minimum window size at which a texture can be perceived is inversely proportional to the scale, this experiment can be used to determine the minimum window size. However, at the writing of this paper we did not have enough data to make a reliable determination of the minimum window size.

Experiment 2: Texture similarity

The goal of this experiment is to establish a threshold for the similarity of the color composition texture features. In the subjective test, two color texture segments were displayed side by side, and the subjects were asked to provide a similarity score for the displayed texture patterns. The options were:

- *Same texture*
- *Very similar*
- *Similar*
- *Somewhat similar*
- *Totally different*

No definition of similarity was given. The test included segments from the same texture and segments that the subject classified into the same category in Experiment 1. It is highly unlikely, of course, that textures belonging to different categories will be classified as anything but "totally different." This was critical in reducing the overall length of the test.

As we will see in the discussion of the experimental results below, we found that the middle category was used for image pairs about which they were uncertain. So, the following set of names would more appropriate for this experiment:

- *Same texture*
- *Similar*
- *Uncertain* (Alternatively, this category may be dropped completely in order to force the user to select one of the other categories.)
- *Somewhat similar*
- *Dissimilar*

4. ANALYSIS OF EXPERIMENTAL RESULTS

We now analyze the results of the experiments and use them to tune the segmentation algorithm.

Smooth vs. nonsmooth classification

A number of researchers have attempted to answer the question of whether a texture exists in an image. Cross and Jain[25] use a Markov random field based model to differentiate texture regions from regions that contain only white noise. Dinstein *et al.*[6] classify a pixel as "texture" by simply thresholding the difference between the maximum and minimum gray values within its neighborhood. Karu *et al.*[7] locate a texture of given coarseness based on density of the local grayscale extrema; they assume that the scale at which the texture needs to be extracted is known. Palmer and Petrou[8] locate the textured region boundaries in remotely sensed images by using the concept of "free angle" together with mathematical morphology. In contrast, we base the smooth vs. nonsmooth classification on models of the statistics of natural images extracted from data collected in Experiment 1.

As the subject judgments varied, the image category was determined using a "majority wins" rule, *i.e.,* the category that receives more than half of the votes is chosen. Note that the majority can be defined in a stricter sense (higher than 50%). Once the texture category was determined, we analyzed the image segments in the smooth and nonsmooth categories. We obtained the steerable pyramid decomposition of each image and calculated the median energy of $s_{\max}$ over the image segment that was displayed to the subjects. We then collected the median energy values for each of the textures that were classified as smooth and nonsmooth, and tried different distributions in order to find the best fit. We found that the Log Normal model is the best in terms of accuracy and simplicity. Both the Kolmogorov-Smirnov test and the Chi-square test indicated that the difference between the empirical and theoretical cumulative distributions is not significant at the significance level of $\alpha = 0.05$. The models we obtained for the smooth and nonsmooth classes are $LogN(0.73, 1.20)$ and $LogN(4.27, 1.23)$, respectively, where the first parameter denotes the mean and the second the standard deviation of the distribution. Figure 3 shows the fitted Log Normal distribution for the smooth and nonsmooth classes. When the two classes are equiprobable, the threshold below which a pixel is classified as smooth is 12.11, which is the point where the two models intersect. The threshold is a function of the means and standard deviations of the two distributions and the probability of occurrence of each class.

The smooth/nonsmooth determination can now be based on the threshold provided by the above subjective experiment. As we saw above, this threshold should also depend on the probability of occurrence of each class. This probability could be determined for each image using the following iterative scheme. First, an initial classification is obtained assuming equal probability for the smooth and nonsmooth classes. The probability of the smooth class is then recalculated based on current classification and the threshold updated. The thresholding and probability updating procedure is then repeated until convergence. Note that in this case there is no need for a cluster validation test, in contrast to the K-means classification procedure used in Refs. 1–3. We also avoid other biases that can arise in the K-means clustering procedure, for example, when one cluster is much stronger than the other. Overall, based on experimentation with hundreds of images, the use of models of natural images statistics, provides a more accurate and robust classification.

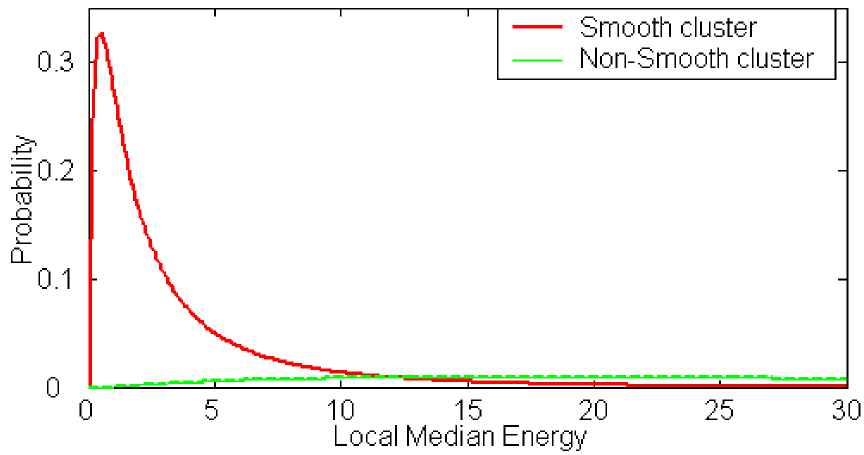

Figure 3. Distribution of smooth and nonsmooth classes

Texture orientation determination

As we saw Section 2, the determination of the texture orientation of the nonsmooth regions is based on two threshold T_1 and T_2; these thresholds ensure the existence and uniqueness of a dominant orientation. To obtain these thresholds, we collected all the images that were classified as having one dominant orientation in Experiment 1. We then calculated the histogram of maximal indices over the image, and computed the values of T_1 and T_2. We then found the smallest value of T_1 and T_2 over all subjects and all images, and used those as the thresholds. The values we obtained based on the available data were $T_1 = 42\%$ and $T_2 = 10\%$.

Color feature similarity threshold

To obtain the color feature similarity threshold, we calculated the average OCCD color feature distance of all image pairs in each similarity category, for each subject. In Fig. 4, each star represents the average OCCD color feature distance over all image pairs classified into the similarity category by one subject. Recall that no definition of similarity was provided to the subjects. Post-test interviews with the subjects indicate that they classified image pairs into the "totally different" and "same texture" categories when they were fairly confident about such classifications in terms of both color and spatial texture. We also found that the subjects typically used the "similar" category for image pairs about whose similarity they were most uncertain. When two images were similar in texture and different in color, they were likely to be placed in the "somewhat similar" category. Finally, the "very similar" category was used for textures that had similar color and spatial texture, but were not perceived as samples of the "same" texture. Overall, the subjects seemed to be quite conservative in concluding that image pairs belong to the "same texture".

Based on the above observations, we combined the data from the "same texture" and "very similar" categories into one group and the data from the "totally different" and "somewhat similar" categories into another. The data in the "similar" category were discarded. We then fitted distributions to the two groups using procedures similar to those that were used in the smooth/nonsmooth classification. The fitted distributions were $LogN(0.486, 0.243)$ and $N(6.65, 6.53)$ for the two clusters, where $N(\mu, \sigma)$ represents the Normal distribution with mean μ and standard deviation σ. Assuming that the two clusters are equally likely, the threshold then becomes 2.78.

5. PERCEPTUALLY TUNED SEGMENTATION

Based on the experimental results of the previous sections, we can now obtain a perceptually tuned version of the color-texture segmentation algorithm. The segmentation results before and after tuning are compared in Figs. 5 and 6, respectively. The image resolutions from left to right are 250×214, 140×199, 149×180, and 162×190 pixels. It is apparent that

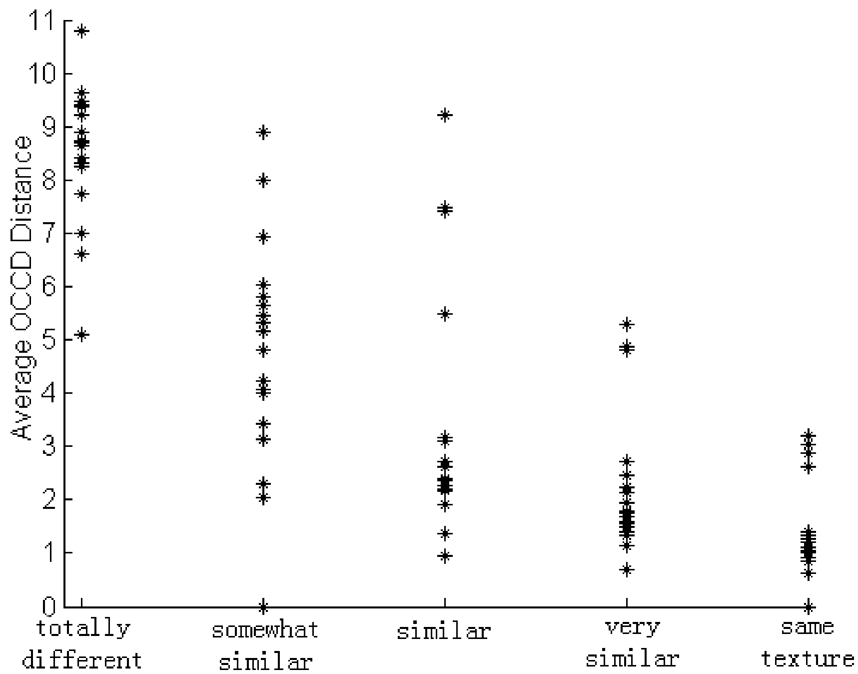

Figure 4. Scatter plot of average OCCD distance for all subjects.

perceptual tuning results in considerable improvement in image segmentation. For example, note that in the first image the edge between the sky and the mountains is more accurate in the perceptually-tuned segmentation. The perceptually-tuned results also appear to result in fewer, better defined segments, especially in the last image on the right.

As we collect data from more subjects, more accurate statistical models can be obtained, which in turn can lead to further improvements in performance.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No.CCR-0209006. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation (NSF).

REFERENCES

1. J. Chen, T. N. Pappas, A. Mojsilovic, and B. E. Rogowitz, "Adaptive image segmentation based on color and texture," in *Proc. Int. Conf. Image Processing (ICIP-02)*, **2**, pp. 789–792, (Rochester, NY), Sept. 2002.
2. J. Chen, T. N. Pappas, A. Mojsilovic, and B. E. Rogowitz, "Perceptual color and texture features for segmentation," in *Human Vision and Electronic Imaging VIII*, B. E. Rogowitz and T. N. Pappas, eds., **Proc. SPIE Vol. 5007**, pp. 340–351, (Santa Clara, CA), Jan. 2003.
3. J. Chen, T. N. Pappas, A. Mojsilovic, and B. E. Rogowitz, "Image segmentation by spatially adaptive color and texture features," in *Proc. Int. Conf. Image Processing (ICIP-03)*, **1**, pp. 1005–1008, (Barcelona, Spain), Sept. 2003.
4. J. Chen, T. N. Pappas, A. Mojsilovic, and B. E. Rogowitz, "Perceptually-tuned multiscale color-texture segmentation," in *Proc. Int. Conf. Image Processing (ICIP-04)*, pp. 921–924, (Singapore), Oct. 2004.
5. J. Chen, T. N. Pappas, A. Mojsilovic, and B. E. Rogowitz, "Adaptive perceptual color-texture image segmentation," *IEEE Trans. Image Processing* . To appear.

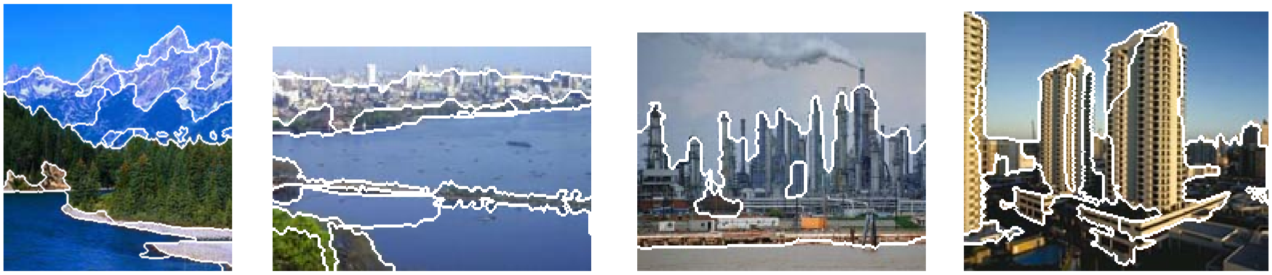

Figure 5. Image Segmentation without perceptual tuning.

Figure 6. Image Segmentation with perceptual tuning.

6. I. Dinstein, A. C. Fong, L. M. Ni, and K. Y. Wong, “Fast discrimination between homogeneous and textured regions,” in *Proc. Asian Conf. Comput. Vis.*, **84**, pp. 361–363, (Osaka, Japan), 1993.
7. K. Karu, A. K. Jain, and R. M. Bolle, “Is there any texture in the image?,” *Pattern Recognition* **29**(9), pp. 1437–1446, 1996.
8. P. L. Palmer and M. Petrou, “Locating boundaries of textured regions,” *IEEE Trans. Geoscience and Remote Sensing* **35**, pp. 1367–1371, Sept. 1997.
9. B. E. Rogowitz, T. N. Pappas, and J. P. Allebach, “Human vision and electronic imaging,” *Journal of Electronic Imaging* **10**, pp. 10–19, Jan. 2001.
10. D. J. Field, “Relations between the statistics of natural images and the response properties of cortical cells,” *Journal of the Optical Society of America A* **4**, pp. 2379–2394, 1987.
11. D. Field, “What the statistics of natural images tell us about visual coding,” in *Human Vision, Visual Proc., and Digital Display*, B. E. Rogowitz, ed., **Proc. SPIE Vol. 1077**, pp. 269–276, (Los Angeles, CA), Jan. 18–20 1989.
12. D. Field, B. Olshausen, and N. Brady, “Wavelets, blur and the sources of variability in the amplitude spectra of natural scenes,” in *Human Vision and Electronic Imaging*, B. E. Rogowitz and J. P. Allebach, eds., **Proc. SPIE Vol. 2657**, pp. 108–119, (San Jose, CA), Jan. 29–Feb. 1 1996.
13. B. Olshausen and D. Field, “Learning efficient linear codes for natural images: the roles of sparseness, overcompleteness, and statistical independence,” in *Human Vision and Electronic Imaging*, B. E. Rogowitz and J. P. Allebach, eds., **Proc. SPIE Vol. 2657**, pp. 132–138, (San Jose, CA), Jan. 29–Feb. 1 1996.
14. D. J. Field, “Wavelets, vision and the statistics of natural scenes,” in *Wavelets: The key to intermittent information?*, B. W. Silverman and J. C. Vassilicos, eds., ch. 8, pp. 147–164, Oxford University Press, 1st ed., 2000.
15. L. Maloney, “Photoreceptor spectral sensitivities and color correction,” in *Perceiving, Measuring, and Using Color*, M. H. Brill, ed., **Proc. SPIE Vol. 1250**, pp. 103–110, (Santa Clara, CA), Feb. 15–16 1990.

16. J. Chen, T. N. Pappas, A. Mojsilovic, and B. E. Rogowitz, “Perceptual tuning of low-level color and texture features for image segmentation,” in *Asilomar Conf. on Signals, Systems, and Computers*, (Pacific Grove, CA), Nov. 2003. To appear in 2004 proceedings.
17. T. N. Pappas, “An adaptive clustering algorithm for image segmentation,” *IEEE Trans. Signal Processing* **SP-40**, pp. 901–914, Apr. 1992.
18. A. Mojsilović, J. Hu, and E. Soljanin, “Extraction of perceptually important colors and similarity measurement for image matching, retrieval, and analysis,” *IEEE Trans. Image Processing* **11**, pp. 1238–1248, Nov. 2002.
19. E. P. Simoncelli and W. T. Freeman, “The steerable pyramid: A flexible architecture for multi-scale derivative computation,” in *Proc. ICIP-95, vol. III*, pp. 444–447, (Washington, DC), Oct. 1995.
20. J. G. Daugman and D. M. Kammen, “Pure orientation filtering: A scale invariant image-processing tool for perception research and data compression,” *Behavior Research Methods, Instruments, and Computers* **18**(6), pp. 559–564, 1986.
21. J. Chen, *Perceptually-Based Color and Texture Features for Image Segmentation and Retrieval*. PhD thesis, Northwestern Univ., Evanston, IL, Dec. 2003.
22. B. E. Rogowitz, T. Frese, J. R. Smith, C. A. Bouman, and E. Kalin, “Perceptual image similarity experiments,” in *Human Vision and Electronic Imaging III*, B. E. Rogowitz and T. N. Pappas, eds., **Proc. SPIE, Vol. 3299**, pp. 576–590, (San Jose, CA), Jan. 1998.
23. A. Mojsilovic and B. Rogowitz, “A psychophysical approach to modeling image semantics,” in *Human Vision and Electronic Imaging VI*, B. E. Rogowitz and T. N. Pappas, eds., **Proc. SPIE Vol. 4299**, pp. 470–477, (San Jose, CA), Jan. 2001.
24. A. Mojsilovic and B. Rogowitz, “Capturing image semantics with low-level descriptors,” in *Proc. Int. Conf. Image Processing (ICIP-01)*, pp. 18–21, (Thessaloniki, Greece), Oct. 2001.
25. G. Cross and A. K. Jain, “Measures of homogeneity in textures,” in *Proc. CVPR Conf.*, pp. 211–216, (Washington DC), 1983.