

Significantly Different Textures: A Computational Model of Pre-attentive Texture Segmentation

Ruth Rosenholtz

Xerox PARC, 3333 Coyote Hill Rd., Palo Alto, CA 94304, rruth@parc.xerox.com

Abstract. Recent human vision research [1] suggests modelling pre-attentive texture segmentation by taking a set of feature samples from a local region on each side of a hypothesized edge, and then performing standard statistical tests to determine if the two samples differ significantly in their mean or variance. If the difference is significant at a specified level of confidence, a human observer will tend to pre-attentively see a texture edge at that location. I present an algorithm based upon these results, with a well specified decision stage and intuitive, easily fit parameters. Previous models of pre-attentive texture segmentation have poorly specified decision stages, more unknown free parameters, and in some cases incorrectly model human performance. The algorithm uses heuristics for guessing the orientation of a texture edge at a given location, thus improving computational efficiency by performing the statistical tests at only one orientation for each spatial location.

1 Pre-attentive Texture Segmentation

Pre-attentive texture segmentation refers to the phenomenon in human vision in which two regions of texture quickly (i.e. in less than 250 ms), and effortlessly segregate. Observers may perceive a boundary or edge between the two regions.

In computer vision, we would like to find semantically meaningful boundaries between different textures. One way of estimating these boundaries is to find boundaries that would be found by a human observer. The boundaries thus defined should be sufficient for most computer vision applications. Whether a human observer can distinguish two textures depends upon whether the discrimination is pre-attentive or attentive. The experimental literature tells us far more about pre-attentive segmentation than attentive discrimination.

Researchers have suggested both feature- and filter-based models of pre-attentive texture segmentation. Many of the feature-based models have been statistical in nature. Julesz [2] suggested that pre-attentive segmentation is determined by differences in the 2nd-order statistics of the texture, or differences in the 1st-order statistics of "textons" such as line terminators and corners [3]. Beck, Prazdny, & Rosenfeld [4] suggested that texture segmentation is based upon differences in the first-order statistics of stimulus features such as orientation, size, and contrast. However, these theories do not indicate how these differences might be quantified, or what properties of the statistics might be used. Furthermore, such models have not typically been implemented such that they could be tested on actual images.

Filter-based models [e.g. 5, 6, 7, 8] have suggested that texture segmentation is determined by the responses of spatial-frequency channels, where the channels contain

both linear filtering mechanisms and various non-linearities. Malik & Perona's model [7] provides the most developed example of this type of model. It involves linear bandpass filtering, followed by half-wave rectification, non-linear inhibition and excitation among channels and among neighboring spatial locations, filtering with large-scale Gaussian first derivative filters, and a decision based upon the maximum response from the final filtering stage.

These models often contain many unknown parameters. What weights specify the inhibition and excitation among the different filter responses? What scale should one use for the Gaussian 1st- derivative filters? Perhaps most importantly, such models often contain an arbitrary, unspecified, threshold for determining the existence of a perceived edge between two textures. Many of these filter-based models are notoriously vague about the final decision stage. Furthermore, such models don't give us much insight into which textures will segment, since the comparison carried out by the model is often obscured by the details of the filtering, non-linearities, and image-based decision stage. What is the *meaning* of the texture gradient computed in the penultimate stage of the Malik & Perona model?

This paper describes a working texture segmentation algorithm that mimics human pre-attentive texture segmentation. Section 2 reviews recent human vision experiments [1], which were aimed at studying what first-order statistics determine texture segmentation. These results suggest that modelling pre-attentive texture segmentation by standard statistical tests for a difference in mean and standard deviation of various features such as orientation and contrast. Section 3 reviews previous models of pre-attentive texture segmentation in light of these results, and discusses the relationship to other edge detection and image segmentation algorithms. Section 4, presents a biologically plausible, filter-based algorithm based upon the experimental results in Section 2 and those of Kingdom & Keeble [9]. Section 5 presents results of this algorithm on artificial and natural images.

2 Recent Experimental Results in Pre-attentive Texture Segmentation

In [1], I studied segmentation of orientation-defined textures such as those shown in Figure 1. In each of the three experiments, observers viewed each texture pair for 250 ms, and the task was to indicate whether the boundary between the two textures fell to the left or right of the center of the display.

If, as Beck et al [4] suggested, texture segmentation of orientation-defined textures is based upon differences in the 1st-order statistics of orientation, to what 1st-order statistics does this refer? If the difference in mean orientation is the crucial quantity, two textures should segment if this difference lies above a certain threshold, independent of other properties of the orientation distributions. A more plausible possibility is that the determining factor is the *significance* of the difference in mean orientations. The significance of the difference takes into account the variability of the textures, so that two homogeneous textures with means differing by 30 degrees may segment, while two heterogeneous textures with the same difference in mean may not. Perhaps observers can also segment two textures that differ only in their variability. Other parameters of the distribution might also be relevant, such as the

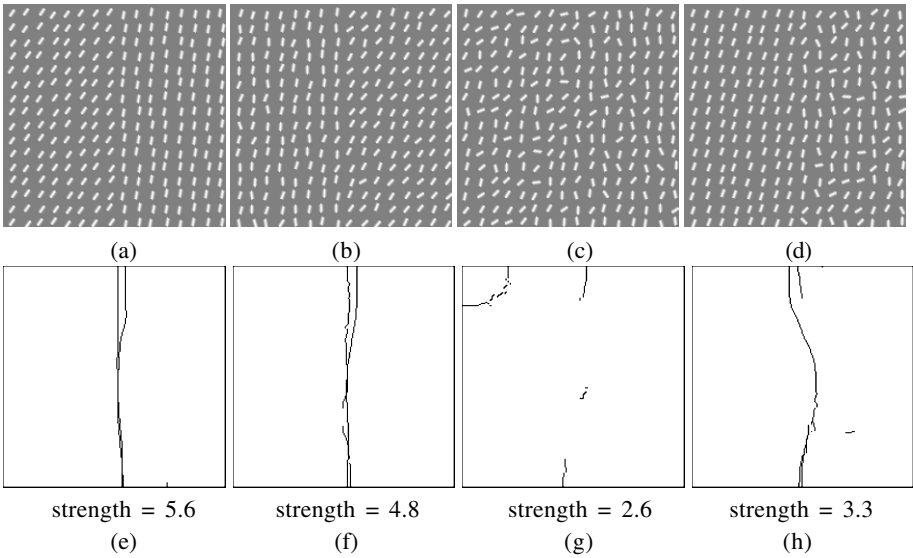


Figure 1: Three orientation-defined textures (a, b, c, d) and the significant edges found by our algorithm (e, f, g, h). The “strength” gives the average value of the test statistic at all locations for which the edge was significant. For (e, f, g), compare this with 2.04, the 82% threshold for a difference in mean orientation. For (h), compare the strength with 2.5, the 82% threshold for a difference in orientation variability.

skew or kurtosis. Alternatively, observers might be able to segment two textures given any sufficiently large difference in their first-order statistics.

The first experiment asked observers to segment two textures that differed only in their mean orientation. Each texture had orientations drawn from a *wrapped normal* distribution [10]. The experiment determined the threshold difference in mean orientation, at which observers can correctly localize the texture boundary 82% of the time, for 4 different values of the orientation standard deviation. Figure 2a shows the results. Clearly observers can segment two textures differing only in their mean orientation. Furthermore, the difference in mean required to perform the segmentation task depends upon the standard deviation.

The second experiment determined whether or not observers could pre-attentively segment textures that differed only in the variance of their orientation distributions. For two possible baseline standard deviations, the experiment measured the threshold increment in standard deviation at which observers could correctly localize the texture boundary 82% of the time. Observers could segment textures differing only in their variance, and Figure 2b shows the thresholds found. The difference in variance required depends upon the baseline standard deviation.

The third experiment tested segmentation of a unimodal wrapped-normal distribution from a discrete, bimodal distribution with the same mean orientation and variance. This experiment measured percent correct performance, for 4 possible spacings of the modes of the bimodal distribution. The results are shown in Figure 2c. All observers

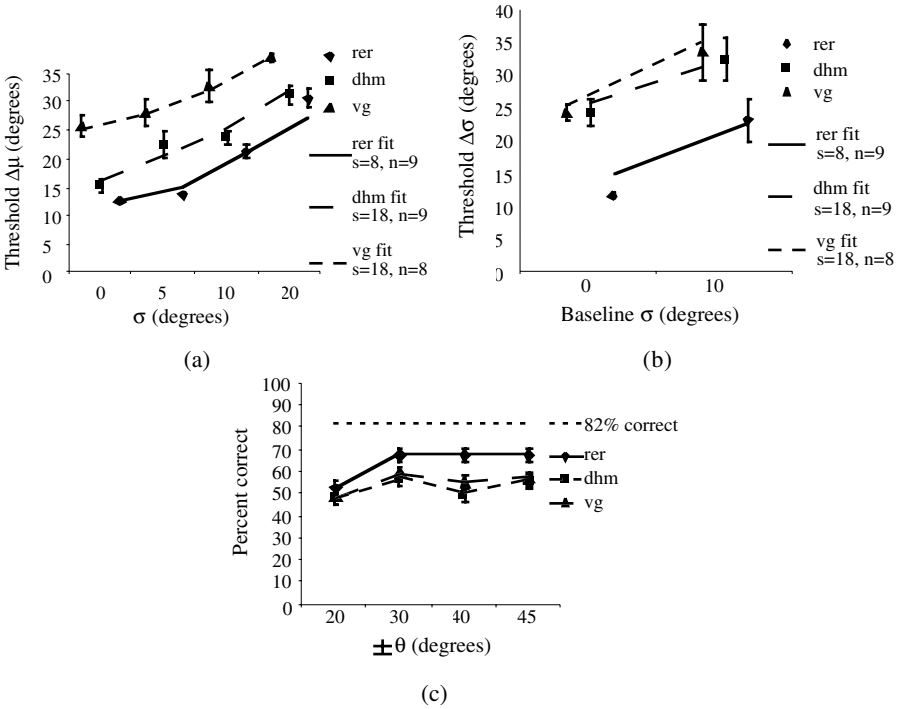


Figure 2: Results of psychophysical experiments and modelling. Symbols plus error bars indicate data; curves indicate the fit of the model to that data. (See text.) In (a) and (b) data for the 3 observers and best fit curves have been shifted horizontally to facilitate viewing.

performed well below the 82% correct required in Experiments 1 and 2, with only RER (the author) performing significantly above chance. These results do not rule out the possibility that observers may segment textures differing in statistics other than their mean and variance. However, the inability to segment textures that differ so greatly suggests that observers do not make use of the full first-order statistics in performing this task.

These results suggest a model of pre-attentive texture segmentation of orientation-defined textures. Figure 3 depicts the stages of this model. The observer first extracts noisy estimates of orientation, with the internal noise distributed according to a




Figure 3: Diagram of our computational model of texture segmentation.

wrapped normal distribution, with some standard deviation s . The observer then collects n orientation estimates from each side of a hypothesized edge. If the two sets of samples differ significantly, at the $\alpha=0.82$ confidence level, in either their mean orientation or variance, then the observer sees a boundary. As implemented, this model uses the standard Watson-Williams test for a significant difference in mean orientation, and the Watson-Williams and Mardia test for a significant difference in orientation variance [10]. Section 4 discusses these tests in more detail.

This model, then, has two free parameters: the internal noise standard deviation, s , and the number of samples, n . The curves Figure 2a-b show the fit of this model to the experimental data, and the legend indicates the best-fit values of these parameters. For any given observer, this model provides a good fit of the thresholds for segmentation both of textures differing only in mean orientation, and textures differing only in orientation variance. Furthermore, the fit to the data of each of the three observers yields roughly the same value for the number of samples, n , suggesting that the size of the integration region may not vary much across observers. In modelling experimental results on homogeneous textures such as those used in these experiments, the n samples may be taken from anywhere in the texture. However, previous work suggests that the visual system performs local texture segmentation computations, with samples taken from regions adjacent to the hypothesized edge [see, e.g., 8, 11]. The fit to the data of the two less experienced subjects, DHM and VG, yields roughly the same value for the internal noise parameter, s . The fit to the data of experienced observer RER yields a lower value for this parameter, consistent with evidence that learning, for low-level perceptual tasks such as this one, may be mediated by a reduction of internal noise [12].

It appears, then, that for these orientation-defined textures, a good model for segmentation extracts noisy orientation estimates, then tests for a significant difference between the distributions of estimates using standard parametric statistical tests. For the purposes of this paper, I assume that a similar model describes segmentation based upon features besides orientation, such as contrast, color, etc. As of yet there is little experimental evidence addressing this point. Section 4 presents a texture segmentation algorithm based upon this model, using orientation and contrast as the texture features.

3 Previous Work: Texture Segmentation Models and Algorithms

To summarize the experimental results discussed in the previous section, a texture segmentation algorithm should be able to segment textures differing in mean orientation or in orientation variability. The segmentation difficulty should increase with increasing variability, in a way described by standard statistical tests. The segmentation algorithm should not find an edge when two textures have the same mean and variance, yet one is bimodally distributed and the other is unimodal.

Most filter-based models of pre-attentive texture segmentation will find an edge when the two textures differ in their mean orientation. Though these algorithms do not explicitly perform a statistical test for the difference in mean orientation, many of them will replicate the findings that edges become weaker as the variance increases [e.g. 7, 8]. It remains to be seen whether performance will degrade in a way that

matches the experimental findings. If it does not, then these algorithms require more than one threshold -- the threshold would need to vary with orientation variance. Furthermore, my suggested model has two intuitive parameters, s and n , easily determined by a fit to the data. To the extent that previous models have *implicitly* performed a sort of statistical test, it may prove difficult to fit such models to the data. Space concerns prohibit our displaying results from Bergen/Landy [8] and Malik/Perona [7]. However, in our experience, as the variability of the texture increases, both algorithms find a large number of spurious edges, and Bergen/Landy may generate spurious edges at a finer scale than the true edge. Both models fail to find an edge when the textures differ only in variance.

A number of standard clustering techniques will cluster in such a way as to maximize the difference in mean relative to the variability of each cluster, or some other standard statistical measure [13, 14, 15]. However, clustering techniques require that the user specify the number of clusters. My edge detection algorithm will test for edges without knowledge of the number of edges.

Ruzon & Tomasi [16] perform color edge detection using the Earth Mover's Distance (EMD) as a measure of the difference between two distributions of color samples. This work uses a more complicated statistical test for an edge, like the algorithm presented here -- this paper does not analyze the relationship between their EMD and standard statistical tests. Voorhees & Poggio [17] also use a more complicated distance measure between two textures, extracting blobs from the textures, then using a non-parametric statistical test to compare the textures. However, not all textures lend themselves to easy blobs extraction. Furthermore, the experimental results discussed in Section 2 show that the non-parametric test in [17] will find edges where human observers do not pre-attentively perceive them [1], a criticism that also holds for the χ^2 measure of [14].

At each location, Ruzon & Tomasi test for an edge in a number of different directions. This requires a great deal of computation. Section 4.5 presents heuristics for testing for an edge at only one orientation.

Elder [18] and Marimont & Rubner [19], working in standard, luminance-based edge detection, perform a statistical test for whether an edge exists at a given location, but they use a global measure of the variability, and thus will yield unexpected results when the variance changes over the image.

Some of the closest previous work comes from Fesharki & Hellestrand [20], who perform edge detection by using a Student's t-test to test for a significant difference in mean luminance. Similarly, Weber & Malik [21] test for motion boundaries, by using standard parametric statistical tests.

4 The Texture Segmentation Algorithm

4.1 Statistical Tests for a Significant Difference in Mean or Spread

My texture segmentation algorithm declares the presence of an edge if certain basic texture features differ significantly in their mean or spread. For the purposes of this paper, the algorithm extracts the features of orientation and contrast. This subsection describes the statistical tests used by the algorithm for these two kinds of features.

Testing for a Difference in Mean Orientation. Given two samples of orientation estimates, the Watson-Williams test [10] indicates whether the mean orientations of the two samples differ significantly. Assume the two independent random samples of the same size, n , and denote them $\{\phi_i\}_{i=1}^n$ and $\{\psi_i\}_{i=1}^n$. The Watson-Williams test assumes two samples drawn from *von Mises* distributions with the same *concentration parameter*, κ . The von Mises distribution for directional data has many of the nice features of the normal distribution for linear data. It is quite similar to the wrapped normal distribution used in Section 2, and for many purposes they may be considered equivalent [10]. The higher the concentration parameter, the more concentrated the distribution about the mean orientation – i.e. the lower the spread of the distribution. The Watson-Williams test is ideally for κ larger than 2 -- roughly equivalent to an angular standard deviation of less than 22 degrees.

We first compute the components of the *mean resultant vectors*:

$$\bar{C}_1 = \frac{1}{n} \sum_i \cos(2\phi_i), \quad \bar{S}_1 = \frac{1}{n} \sum_i \sin(2\phi_i), \quad \bar{C}_2 = \frac{1}{n} \sum_i \cos(2\psi_i), \quad \bar{S}_2 = \frac{1}{n} \sum_i \sin(2\psi_i)$$

From this, we compute the length of the j th mean resultant vector:

$$\bar{R}_j = \sqrt{\bar{C}_j^2 + \bar{S}_j^2}$$

We also compute the length of the resultant vector for the combined data sample:

$$\bar{R} = \sqrt{\bar{C}^2 + \bar{S}^2}, \text{ where } \bar{C} = (\bar{C}_1 + \bar{C}_2)/2 \text{ and } \bar{S} = (\bar{S}_1 + \bar{S}_2)/2$$

The test statistic is

$$F_\mu = 2 \left(1 + \frac{3}{8\tilde{\kappa}} \right) (n-1) \frac{\bar{R}_1 + \bar{R}_2 - 2\bar{R}}{2 - \bar{R}_1 - \bar{R}_2} \quad (1)$$

where $\tilde{\kappa}$ is the value of the concentration parameter estimated from the two samples.

Under the null hypothesis of no edge, F_μ is approximately distributed according to an F distribution, tabulated in standard statistics books. If F_μ is larger than the tabulated value for significance level α and degrees of freedom $(1, 2n-2)$, the difference in mean orientation is significant at level α . Throughout this paper, $\alpha=0.82$.

Testing for a Difference in Orientation Variability. The Watson & Williams test, modified by Mardia [see 10], tests for a difference in concentration parameter. Assuming two samples, both of size n , drawn from von Mises distributions, with the mean resultant length of the combined sample greater than 0.70 (equivalently, an angular deviation of less than 22 degrees), the test statistic is:

$$F_\kappa = \frac{(1 - \bar{R}_1)}{(1 - \bar{R}_2)} \quad (2)$$

or $1/F_\kappa$, whichever is >1 . Under the null hypothesis that the two samples have the same concentration, F_κ is approximately distributed according to an F distribution. To test for an edge of a given significance, α , one compares F_κ with the value in the $\alpha/2$ F distribution table, with degrees of freedom $(n-1, n-1)$. If the statistic is larger than

the value from the table, the difference in concentration parameter is statistically significant.

Testing for a Difference in Mean Contrast. The standard Student's t -test tests for a significant difference in mean for linear variables such as contrast. Assuming two samples, both of size n , drawn from normal distributions with unknown but equal variance, the test statistic is

$$t = (\bar{x}_1 - \bar{x}_2)\sqrt{n} / \sqrt{s_1^2 + s_2^2} \quad (3)$$

where \bar{x}_j and s_j are the j th sample mean and standard deviation, respectively. Under the null hypothesis that the two samples are drawn from the same distribution, this statistic is distributed according to a Student's t distribution, tabulated in standard statistics books. If the absolute value of t is larger than $t_{\alpha/2}$ from these tables, the difference in mean is significant at significance level α .

Testing for a Difference in Contrast Variance. To test for a significant difference in variance, for a linear variable such as contrast, one uses a standard F test. Again assuming two samples, each of size n , drawn from normal distributions, the test statistic is

$$F_\sigma = \frac{s_1^2}{s_2^2} \quad (4)$$

or $1/F_\sigma$, whichever is greater than 1. Under the null hypothesis that the two samples are drawn from distributions with the same variance, this statistic is once again distributed according to an F distribution. For a given significance, α , the difference in variance is significant if F_σ is larger than the value in the statistics table for this significance, with degrees of freedom $(n-1, n-1)$.

General Comments. Note that the two tests for an orientation edge may also be used for any texture feature that is circularly distributed, i.e. for which a value k is equivalent to a value $k \bmod m$, for some m . The two tests for a contrast edge may be used for any one-dimensional, linearly distributed texture feature. Similar tests exist for multi-dimensional data.

Before going on to describe the rest of the texture segmentation algorithm, consider what processing occurs in each of these four statistical tests. In each case, computing the statistic involves first taking the average of one or more functions of the feature estimates. For example, the orientation tests require first computing the mean resultant vectors, by taking the average of the cosine and sine of twice the orientation. In all 4 cases, the final test statistic involves some function of these averages. This suggests that an algorithm that computes these test statistics will involve first (1) computing a set of feature images, then (2) integrating some function(s) of these images over some integration region, and finally (3) computing the test statistics and testing for the presence of an edge. These stages correspond to those of the model depicted in Figure 3. Additional details of these stages are given below.

4.2 Extract Texture Features

For the purposes of this paper, my algorithm extracts the features of orientation and contrast. The use of contrast as a feature deserves some discussion. Several researchers [6, 7] have argued that segmentation of texture patterns such as those studied by Julesz [2] is determined by the "contrast" features extracted by center-surround filters such as Difference-of-Gaussian (DOG) filters. Others [3, 22] have suggested that one must extract more complicated features such as junctions and line endings. This remains an unresolved issue, but for many examples there is little difference between these two approaches. In addition, my algorithm assumes that segmentation of contrast-defined textures is determined by the significance of the differences between the mean and standard deviation of the contrast, as with orientation. This has yet to be demonstrated experimentally. Finally, the fit of the model to the orientation data from [1] yields estimates of the various model parameters, and the same would need to be done for other texture features, if one wishes to mimic human performance. My philosophy here is to extrapolate the orientation segmentation model to contrast textures, test it on various images, and look forward to additional psychophysics to resolve some of the above issues.

The algorithm extracts the features of orientation and contrast using biologically plausible spatial-frequency channel operations. For orientation, this amounts to using steerable filtering [23] to extract estimates of orientation. Steerable filtering produces two images at each scale, representing $\cos(2\theta)$ and $\sin(2\theta)$, where θ is the orientation, locally. These are precisely the functions of orientation one needs in order to compute the mean resultant vectors required by the statistical tests in Eqns. 1 and 2. For contrast, we follow [7], filtering with first- and second-derivative DOG filters, followed by half-wave rectification, and a local max operation. The local max is intended to give a single, phase-invariant measure of local contrast.

The algorithm extracts orientation and contrast at a number of different scales. For orientation, we use the oriented Gaussian pyramid from [8]. Thus the scales for orientation filtering differ by a factor of two. For contrast, we use the more densely sampled scales of [7].

The algorithm processes each scale and feature independently. As necessary, it adds noise to the feature estimates, so as to match the internal noise parameter in fits to human data. The orientation estimation procedure contains inherent noise equivalent to the internal noise for the two more naïve observers in Figure 2, so the algorithm requires no added noise. For contrast, we have no fit to human data, and have experimented with a number of possible amounts of added noise. In the examples shown here, use an added noise with variance $\sigma^2=150$.

4.3 Collect Samples of Feature Estimates and Compute the Relevant Statistics

Next, the algorithm hypothesizes that an edge, with a particular orientation, exists at a particular location in the image. It then tests this hypothesis by running the various statistical tests described above.

As mentioned in the discussion of these statistical tests, the first step in calculating each of the statistics is to compute averages of various functions of the feature values.

Therefore, the step in the model in which the observer collects n feature samples from each side of a hypothesized edge is equivalent to integrating some functions of the feature values over a local integration region. The algorithm integrates using a Gaussian window.

Previous work by Kingdom & Keeble [9] suggests that, for orientation, the size of the integration region is a constant multiple of the width of the oriented lines that make up the texture. Alternatively, one may think of the region size as being a constant multiple of the support of the filters used to extract orientation estimates. For finer features, the region is proportionally smaller than for coarser features. Once again, the algorithm described here assumes that the same principle holds for features other than orientation.

The fit to the experimental data in Figure 2 indicated that human observers seemed to collect orientation estimates from $n=9$ elements on each side of the hypothesized edge. Based upon the width and spacing of the line elements used in the experimental displays, this implies that human observers collect samples from a circular region of diameter approximately 5 times the support of their oriented filters. The algorithm uses an integration region for contrast features of a similar size relative to the center-surround filters used to extract contrast.

4.4 Combine Results into a Single Edge Map

The results of the previous steps give us an image for each scale, feature type, and statistical test (mean and spread). Each image indicates, for each location, the presence or absence of a significant edge of the given type. The value of the statistic gives an indication of the strength of the edge. One could stop at this stage, or there are various things one could do to clean up these edge maps and combine them into a single map. "False" variance edges tend to occur next to edges due to a difference in mean, since a region which spans the edge will have a higher variance than neighboring regions which include only one texture. Therefore, the algorithm inhibits variance edges near significant difference-in-mean edges. It combines edge images across scales using a straightforward "or" operation. If an edge is significant at any scale, it is significant, regardless of what happens at any other scale. Finally, texture edges tend to be more poorly localized than luminance edges, due to their statistical nature. Texture edges, in the examples in this paper, are typically statistically significant in a band about 2 texels wide. One can "thin" edges by finding the maximum of the test statistic in the direction of the edge, as done for the results presented here.

4.5 How to Hypothesize an Edge

Edge detection methods that involve taking the gradient of the mean of some feature value [e.g. 24, 7, and many others] have the advantage that they are "steerable" [23]. This means that such methods can take only a horizontal and a vertical derivative, and infer the direction in which the derivative – and thus the edge strength – is maximized. Edge detection methods that involve more complex comparisons of the features on each side of the edge [e.g. 16] typically require that one check each possible orientation for the presence of an edge. This process is time consuming. Thus in this




Figure 4: Relationship between mean resultant vectors and heuristics for guessing edge direction (see text).

subsection I ask whether one can make a good guess as to the most likely orientation of an edge at each location, and test only that edge orientation.

This subsection presents a set of heuristics for guessing the edge orientation, one for each statistical test s . For each s , one would like to test only one orientation for an edge at each location, yet find edges at the same locations as if one had tested all possible orientations. Therefore, for each s , one desires a steerable function, f_s , of the feature values, such that *if there exists an orientation such that test s indicates a significant edge, test s will indicate a significant edge at the orientation that maximizes f_s* . If these conditions hold, one may easily find the orientation, θ_{edge} , that maximizes f_s , and be assured that if test s would find a significant edge at that location, it will find one at orientation θ_{edge} .

For each test, s , I present a function f_s that is the Gaussian 1st-derivative of some function of the feature values. The maximizing orientation of such f_s is given by the direction of the gradient. The heuristics presented produced the desired behaviour over 99% of the time in Monte Carlo simulations. Thus these heuristics provide a reliable way of avoiding the computational complexity of testing each possible edge orientation at each spatial location.

Heuristic for a Mean-Orientation Edge. As mentioned above, the orientation estimation stage returns two images for each scale, representing $\cos(2\theta)$ and $\sin(2\theta)$, where θ is the local orientation estimate at that scale. First, rotate all orientation estimates by angle β , such that the mean resultant vector for the combined sample has an orientation of 0° . This generates the image:

$$\sin(2\theta - \beta) = \sin(2\theta)\cos(\beta) - \cos(2\theta)\sin(\beta)$$

Next, take the Gaussian 1st-derivative of this image in the horizontal and vertical directions. The direction of the gradient provides a guess for the edge direction. The intuition follows:

Figure 4a shows example resultant vectors, rotated to the canonical position, for the case in which the two distributions have the same concentration but different mean orientations. The Gaussian 1st-derivative computes essentially the difference between the y-components of the two resultant vectors. The length of the resultant vector gives a measure of the concentration of the distribution. The orientation of the resultant vectors indicates an estimate of the mean orientations of the underlying distributions. For given lengths of the two mean resultant vectors, \bar{R}_1 and \bar{R}_2 , the

vertical distance between their tips will be larger the larger the difference in angle between them, and for a given difference in angle, the vertical distance will be smaller the shorter their lengths. The same is true of the test statistic, F_μ , thus the intuition that this heuristic should give us a reasonable guess for the direction of the edge.

Heuristic for an Orientation-Variability Edge. Again rotate the orientation estimates such that the resultant vector for the combined sample has an orientation of 0° . This time, compute the image:

$$\cos(2\theta - \beta) = \cos(2\theta)\cos(\beta) + \sin(2\theta)\sin(\beta)$$

Again, take Gaussian 1st-derivatives of this image. The direction of the gradient provides the guess for the edge direction. The intuition follows:

Figure 4b shows example resultant vectors for the case in which the underlying distributions differ only in their concentration parameter. The horizontal distance between the tips of these resultant vectors gives a measure of the difference between the two concentration parameters. The test statistic (see Eqn. 2) takes a ratio of the resultant vector lengths, as opposed to a difference, but the difference serves well when it comes to guessing the direction of the edge.

Heuristic for a Mean-Contrast Edge. Here we steer the Gaussian 1st-derivative of contrast. The direction of the gradient, which gives the direction with the largest change in mean contrast, provides the guess for the edge direction.

In this case, one can make a stronger statement: The direction of the gradient of mean contrast indicates the direction in which the test statistic reaches its maximum. Thus the direction of the gradient always provides the best guess for the edge direction. Recall that the test statistic for a difference in mean contrast is:

$$t = (\bar{x}_1 - \bar{x}_2)\sqrt{n} / \sqrt{s_1^2 + s_2^2}$$

The gradient indicates the direction in which $(\bar{x}_1 - \bar{x}_2)$ is maximized. The only way that t could reach a maximum in a different direction is if that direction reduced the denominator more than the numerator. It is a simple matter of algebra to show that an edge direction with smaller $(\bar{x}_1 - \bar{x}_2)$ in fact also yields a larger $\sqrt{s_1^2 + s_2^2}$, thus reducing the test statistic. Due to space concerns, we do not reproduce the proof here. The intuition follows: Suppose that in the direction of the gradient, $\bar{x}_1 > \bar{x}_2$. When the direction of the hypothesized edge changes, m samples from the first set transfer to the second, and vice versa. Since the difference between the two means decreases, the change in direction adds, on average, smaller elements to the 1st set, and takes away larger elements. But such a manipulation will increase the sum of the variances of the sets, since it adds small elements to the set with the larger mean, and vice versa.

Heuristic for a Contrast-Variance Edge. First calculate the mean, μ , of both sample sets. Then compute the image

$$(x_i - \mu)^2$$

Again, steer the Gaussian 1st-derivative of this function, and the direction of its maximum gives us our guess for the edge direction.

This heuristic estimates the variance of each sample set, and uses the difference

between these variances as a measure of the strength of the edge in a given direction. The test statistic takes a ratio of the two variance estimates, as opposed to a difference, but the difference serves well when it comes to guessing the direction of the edge, and is easily steered.

5 Results and Discussion

All of the examples in this paper used the same version of the algorithm, with parameters set as described above. Figure 1 shows the results on four images much like those used in the experiments described in Section 2. The first three texture pairs differ in mean orientation by 25° . Low variability textures, as in Figure 1a, allow observers to localize the boundary well over 82% of the time. In Figure 1b, a larger orientation variance makes the difference in mean just above the 82% correct threshold. In Figure 1c, the increased variance makes the difference in mean well below threshold – observers would have great difficulty localizing this edge.

Figures 1e-1g show the results of our algorithm on these three texture pairs. The results will often show a number of edges laid on top of each other, when edges are found at a number of different scales or by more than one of the 4 statistical tests. The algorithm correctly finds a strong boundary in the first image, a weaker one in the second image, and essentially no edge in the third.

Figure 1d shows an image in which the two textures differ only in the spread of their orientations. This difference is above the thresholds found in Experiment 2, and it should be possible to see the boundary. Our algorithm finds the boundary, as shown in Figure 1h.

Many texture segmentation experiments and theories have revolved around images like those in Figure 5. Malik & Perona [7] tested their algorithm on such images, and used it to predict the segmentability of these texture pairs. Their predictions agreed with the experimental data of Gurnsey & Browse [25]. (Malik & Perona also compared the results of their algorithm with those of Krose [26]. However, Krose studied a visual search task, in which observers looked for an “odd man out.” It is inappropriate to compare those results with results of segmentation of images such as those in Figure 5, for reasons given in [25].) The results of my algorithm, shown in




Figure 5: Four texture pairs from [7], with the edges found by our algorithm. The strength gives the average value of the mean-contrast test statistic for each edge, and should be compared with the threshold value, 1.42. Only mean-contrast edges were significant.




Figure 6: Patchwork of natural textures (a), along with the contrast (b) and orientation (c) edges found by our algorithm.




Figure 7: Results on a natural image.

Figure 5, agree with those of [25].

Figure 6a shows an array of real textures, stitched together. Note that the true boundaries between the textures are *not* straight. The algorithm does quite well at extracting contrast edges (Figure 6b), often closely following the boundaries. It performs less well near corners, where assumptions about the number of textures present in a region are violated. Few of the textures in this image have strong orientation, providing a challenge to finding orientation edges, yet we still see reasonable results, as shown in Figure 6c. The algorithm predicts no pre-attentive segmentation for texture pairs with no edge shown between them. Figure 7 shows a highly oriented natural image, for which the algorithm extracts strong orientation edges.

This paper has presented an algorithm for mimicking pre-attentive texture segmentation. The algorithm bridges the gap between earlier statistical and filter-based models. It has essentially two intuitive key parameters, the internal noise in the feature estimates and the integration scale, and obtains values of those parameters from fits to experimental data in [1, 9]. Performing tests for an edge more complicated than a gradient has traditionally led to increased computational cost due to a need to test for an edge at multiple orientations. The heuristics presented here essentially allow the algorithm's statistical tests to be "steered."

References

- 1 R. Rosenholtz. What statistics determine segmentation of orientation-defined textures? *Perception (Suppl.)*, 26:111, 1997.
- 2 B. Julesz. Experiments in the visual perception of texture. *Sci. Amer.*, 232:34-43, 1975.
- 3 B. Julesz. Texton gradients: the texton theory revisited. *Biol. Cybern.*, 54:245-51, 1986.
- 4 J. Beck, K. Prazdny, & A. Rosenfeld. A theory of textural segmentation. In *Human & Machine Vision*, Beck, Hope, & Rosenfeld, eds. (New York: Academic Press), 1983.
- 5 M. R. Turner. Texture discrimination by Gabor functions. *Biol. Cybern.*, 55:71-82, 1986.
- 6 J. Bergen & E. Adelson. Early vision and texture perception. *Nature*, 333:363-4, 1988.
- 7 J. Malik & P. Perona. Preattentive texture discrimination with early vision mechanisms. *J. Opt. Soc. Am. A.*, 7(5):923-32, 1990.
- 8 J. R. Bergen & M. S. Landy. Computational modeling of visual texture segregation. In *Computational models of visual perception*, Landy & Movshon, eds. (Cambridge, MA: MIT Press), 1991.
- 9 F. A. A. Kingdom & D. R. T. Keeble. On the mechanism for scale invariance in orientation-defined textures. *Vision Research*, 39:1477-89, 1999.
- 10 E. Batschelet. *Circular statistics in biology*. (London: Academic Press), 1981.
- 11 H. C. Nothdurft. Sensitivity for structure gradient in texture discrimination tasks. *Vision Research*, 25:1957-68, 1985.
- 12 B. A. Doshier & Z.-L. Lu. Mechanisms of perceptual learning. *Investigative Ophthalmology & Visual Science (Suppl.)*, 39(4):912, 1998.
- 13 R. Duda & P. Hart. *Pattern classification and scene analysis*. (New York: Wiley), 1973.
- 14 J. Puzicha, T. Hoffman, J. Buhmann. Non-parametric similarity measures for unsupervised texture segmentation and image retrieval. *Proc. CVPR*, pp. 267-72, Puerto Rico, June, 1997.
- 15 J. Shi & J. Malik. Self-inducing relational distance and its application to image segmentation. *Proc. ECCV*, pp. 528-43, Freiburg, Germany, June, 1998.
- 16 M. A. Ruzon & C. Tomasi. Color edge detection with the compass operator. *Proc. CVPR*, pp. 160-6, Fort Collins, CO, June, 1999.
- 17 H. Voorhees & T. Poggio. Computing texture boundaries from images. *Nature*, 333:364-7, 1988.
- 18 J. H. Elder & S. W. Zucker. Local scale control for edge detection and blue estimation. *Proc. ECCV*, pp. 57-69, 1996.
- 19 D. H. Marimont & Y. Rubner. A probabilistic framework for edge detection and scale selection. *Proc. ICCV*, pp. 207-14, Bombay, India, January, 1998.
- 20 M. N. Fesharki & G. R. Hellestrand. A new edge detection algorithm based on a statistical approach. *Proc. ISSIPNN*, 1:21-4, 1994.
- 21 J. Weber & J. Malik. Scene partitioning via statistic-based region growing. Tech. report UCB/CSD-94-817, Comp. Sci. Div. (EECS), Univ. of California, Berkeley, 1994.
- 22 E. Barth, C. Zetsche, F. Giulianini, & I. Rentschler. Intrinsic 2D features as textons. *J. Opt. Soc. Am. A.*, 15(7):1723-32, 1998.
- 23 W. T. Freeman & E. H. Adelson. The design and use of steerable filters. *IEEE PAMI*, 13(9):891-906, 1995.
- 24 R. M. Haralick. Digital step edges from zero crossings of the 2nd directional derivative. *IEEE PAMI*, 6(1):58-68, 1984.
- 25 R. Gurnsey & R. Browse. Micropattern properties and presentation conditions influencing visual texture discrimination. *Percept. Psychophys.*, 41:239-52, 1987.
- 26 B. Krose. Local structure analyzers as determinants of preattentive pattern discrimination. *Biol. Cybern.*, 55:289-98, 1987.