

Spatial Information in Multiresolution Histograms

E. Hadjidemetriou, M. D. Grossberg, and S. K. Nayar
Computer Science, Columbia University, New York, NY 10027
{stathis, mdog, nayar}@cs.columbia.edu

Abstract

Intensity histograms have been used extensively for recognition and for retrieval of images and video from visual databases. Intensity histograms of images at different individual resolutions have also been used for indexing. They suffer, however, from the inability to encode spatial image information. Spatial information can be incorporated into histograms simply by taking histograms of an image at multiple resolutions together to form a multiresolution histogram. Multiresolution histograms can also be computed, stored, and matched efficiently. In this work we analyze and quantify the relation and sensitivity of the multiresolution histogram to spatial image information as well as to properties of shapes and textures in an image. We verify the analytical results experimentally. We demonstrate the ability of multiresolution histograms to discriminate between images, as well as their robustness to noise.

1. Introduction

Histograms have been used widely in object recognition [9, 29, 30] and are an important tool for the retrieval of images and video from visual databases [1, 22, 34]. They have many merits. For example, they are efficient, and robust to local area preserving transformations [12]. They suffer, however, from the inability to encode spatial image information.

Several investigators have tried to incorporate spatial information into histograms by defining local rather than global intensity histograms. Examples of this approach include error diffusion, the scale-imprecision space [11], and the locally orderless images [19]. In the context of indexing, intensity histograms have been combined with other spatial features that include edges, corners, textures, and regions [21, 22]. Histograms have also been combined with additional pixel statistics to form multidimensional histograms [24]. For example, co-occurrence matrices [13], or correlograms [15].

In addition, histograms of image features obtained by applying derivative filters to an image have been suggested [26]. Also, individual histograms of low image resolutions have been used to expedite retrieval [20] as well as histograms of multiple image resolutions for texture synthesis [5, 14]. Further, a representation that also employs filtering, called flexible histograms, has been applied both to object recognition [4] and texture synthesis [5].




Figure 1. The first row shows the image pyramid and the second row the multiresolution histogram.

The representation we analyze in this work, namely the multiresolution histogram, implicitly includes spatial image information. A *multiresolution histogram is the set of intensity histograms of an image at multiple resolutions*. Note that it is an image representation since multiresolution is applied to the image. It is different from histogram representations where multiresolution is applied directly to the histogram [8]. It is a simple representation that can be computed, stored, and matched efficiently. In figure 1 we show an example of a multiresolution histogram. The first row shows the image pyramid and the second row the multiresolution histogram. In addition to the initial image histogram, the multiresolution histogram also captures spatial image information through the histograms of lower resolutions. It is, however, invariant to rotations and translations.

What spatial information is embedded in a multiresolution histogram, and how sensitive is the multiresolution histogram to this information? To analyze these questions we quantify the dependence and sensitivity of the multiresolution histogram to spatial image information as well as to parameters of shapes and texels in an image. The shape and texel parameters include their size, their elongation, their border complexity, and their placement pattern. For the purpose of the analysis we exploit the equivalence between the histogram and the generalized image entropies [27], and extend this equivalence to multiresolution histograms.

Since multiresolution histograms incorporate spatial information, they are able to discriminate between different images even if the images have identical histograms. We verify the performance of multiresolution histograms experimentally. The distance between two multiresolution histograms is the sum of the distances between the histograms of all im-

age resolutions. We use a database of synthetic images and a database of natural textures in the experiments. We find that multiresolution histograms are very robust to noise and that they can discriminate between images with different spatial patterns alone without the help of any other image features.

2. Approach

We use the Gaussian image multiresolution [18, 33]. Each image resolution gives a different histogram. We analyze how the rate of change of the histogram with image resolution depends on spatial image information as well as on properties of image shapes and textures. To achieve this we employ a transformation of the histogram into the spectrum of generalized image entropies. This spectrum results by generalizing the Shannon entropy into a parametric family. We extend this transformation to multiresolution histograms.

We use the transformed spectrum because the rates at which generalized entropies change with image resolution are given by functionals of the image. These functionals, namely the generalized Fisher information measures, have the desirable property that they depend on the image gradient. Thus, they capture spatial image variation and parameters of image shapes and textures. Therefore, the generalized Fisher information measures link the rate at which the histogram changes with image resolution to parameters of image shapes and textures.

3. Multiresolution Histogram and Fisher Information Measures

We assume that image \mathcal{L} has a continuous domain D with spatial coordinates $\mathbf{x} = (x, y)$. The resolution of the image is decreased linearly with σ by convolving the image with a Gaussian kernel $\mathcal{G}(\sigma)$ [18, 33]. The convolution gives the image $\mathcal{L} * \mathcal{G}(\sigma)$ with histogram $\mathbf{h}(\mathcal{L} * \mathcal{G}(\sigma))$. The rate of change of the image histogram with respect to image resolution is given by $\frac{d}{d\sigma} \mathbf{h}(\mathcal{L} * \mathcal{G}(\sigma))$.

The histogram of an image is linearly related to the Tsallis generalized entropies of an image [32]. More precisely, the Tsallis generalized entropies of an image \mathcal{L} with a unit L_1 norm and M graylevels are given by:

$$S_q = \int_D \frac{\mathcal{L}(\mathbf{x}) - \mathcal{L}^q(\mathbf{x})}{q-1} d^2x = \sum_{j=0}^{M-1} \left(\frac{v_j - v_j^q}{q-1} \right) h_j \quad (1)$$

where q is a continuous parameter, and h_j is the histogram density of grayscale value v_j . In the limit $q \rightarrow 1$ the Tsallis generalized entropies reduce to the Shannon entropy.

Equation (1) can be used to express a vector of M different Tsallis entropies $\mathbf{S} = [S_{q_0} S_{q_1} S_{q_2} \dots S_{q_{M-1}}]^T$ as a linear function of the histogram $\mathbf{h} = [h_0 h_1 h_2 \dots h_{M-1}]^T$ to give

$$\mathbf{S}(\mathcal{L}) = \mathbf{T} \mathbf{h}(\mathcal{L}) \quad (2)$$

where \mathbf{T} is an $M \times M$ transformation. The derivation of equation (2) is almost identical to the derivation provided by

Sparring and Weickert in [27], who also show that the transformation \mathbf{T} is invertible.

The rate at which generalized image entropies change with respect to image resolution, $\frac{d}{d\sigma} S_q(\mathcal{L} * \mathcal{G}(\sigma))$, are given by the generalized Fisher information measures J_q [3, 25, 28]. More precisely,

$$\frac{d}{d\sigma} S_q(\mathcal{L} * \mathcal{G}(\sigma)) = J_q(\mathcal{L}) = \int_D \left| \frac{\nabla \mathcal{L}(\mathbf{x})}{\mathcal{L}(\mathbf{x})} \right|^2 \mathcal{L}^q(\mathbf{x}) d^2x. \quad (3)$$

The “sharpness” or spatial variation at a pixel is given by $|\nabla \mathcal{L}(\mathbf{x})/\mathcal{L}(\mathbf{x})|^2$. Hence, the average sharpness as can be seen from equation (3) is J_1 , namely the Fisher information [28]. In general, functional J_q measures a weighted average “sharpness” of an image.

To relate J_q to the rate at which the histogram changes with image resolution we differentiate equation (1) with respect to resolution and substitute the first equality in equation (3) to get:

$$J_q = \sum_{j=0}^{M-1} \left(\frac{v_j - v_j^q}{q-1} \right) \frac{dh_j}{d\sigma} \quad (4)$$

Therefore, J_q is a linear function of the rate at which the histogram bins change with resolution.

The proportionality factors, shown in equation (4) within the parentheses, are exponential functions of q . The proportionality factors of J_q weigh heavier the rate of change of the histogram bins corresponding to large intensity values (respectively small intensity values), for $q > 1$ (respectively $q < 1$) [25, 32]. The proportionality factors of the Fisher information J_1 weigh approximately equally all histogram bins. Hence, J_1 approximates the average rate of change of histogram bins. In this work we primarily analyze the Fisher information.




Figure 2. Diagram of the approach. The derivative of the histogram with respect to image resolution is transformed for analytical purposes to the Fisher information measures.

The rate of change of vector $\mathbf{S} = [S_{q_0} S_{q_1} S_{q_2} \dots S_{q_{M-1}}]^T$ with respect to image resolution is given by the $M \times 1$ vector $\mathbf{J} = [J_{q_0} J_{q_1} J_{q_2} \dots J_{q_{M-1}}]^T$. Vector \mathbf{J} can be expressed using equation (4) as a linear function of the rate at which histogram bins change. The same result is obtained by differentiating equation (2) with

respect to resolution σ to get:

$$\mathbf{J}(\mathcal{L}) = \mathbf{T} \frac{d}{d\sigma} \mathbf{h}(\mathcal{L} * \mathcal{G}(\sigma)). \quad (5)$$

As shown for the case of equation (2), transformation \mathbf{T} is invertible [27]. Therefore, the rate at which the histogram changes can be expressed as a linear function of \mathbf{J} . In turn, all J_q depend on spatial image variation. Both J_q and the multiresolution histogram, however, are preserved by rotations and translations of the image.

Figure 2 shows an overview of the approach. The connection from J_q to parameters of shapes and texels will be presented in the following two sections. We will mostly work with the Fisher information J_1 .

4. Multiresolution Histograms of Shapes

A multiresolution histogram depends on spatial image variation, which can be modeled as shape in many images. We analyze this dependence using J_q . For discrete images with N pixels J_q is a convex function in an N dimensional space. The single minimum of J_q is achieved for a radially symmetric Gaussian image [7, 25, 32]. As an image departs from Gaussian, the value of J_q increases. Several classes of transformations can lead the image along a path away from the Gaussian in this N dimensional space. We will examine the value of J_q and its sensitivity to several such classes of shape transformations and warps. Namely, elongations, warps of the boundary from rounded to pinched, and changes of the diffuseness of the intensities across the border from smooth to abrupt. As mentioned earlier, translations and rotations of an image do not affect J_q .

We first examine the effect of stretching transformations, which can elongate radially symmetric shapes. The histogram of an elongated shape changes faster with resolution than that of a radially symmetric one. We examine the linear diagonal transformations $x \rightarrow x\sqrt{\rho}$, $y \rightarrow y/\sqrt{\rho}$, where ρ is any positive number we call elongation. The relation of J_q to elongation ρ is quantified for the classes of shapes shown in figure 3. The first column of figure 3 shows the shapes to be analyzed. For these shapes ρ coincides with the ratio of the parameters along the axes, and k is the product of the parameters along the axes. For example, for the Gaussian $\rho = \sigma_x/\sigma_y$ and $k = \sigma_x\sigma_y$. In the second column we can see J_q as a function of ρ and k . The expressions for J_q are computed using the integral in equation (3).

In the second column of figure 3 we plot J_q as a function of ρ for a fixed value of k . For all images the minimum value of J_q is achieved for radially symmetric or regular shapes, whose ratio of axial parameters is one $\rho = 1$. Further, J_q increases strictly and non-linearly with elongation ρ . The elongation of the shapes does not affect their histogram since area and k remain invariant [12].

The value of J_q for shapes also depends on how rounded or complex the border is. The value of J_q is larger when

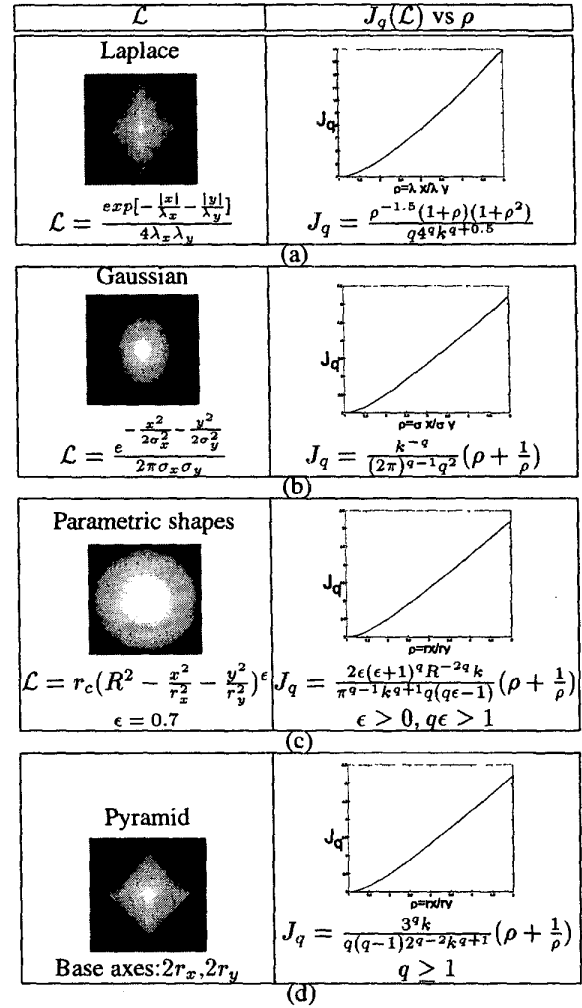


Figure 3. The values of J_q for several images \mathcal{L} as a function of the elongation ρ , where ρ is the ratio of the parameters along the axes. The first column shows the images \mathcal{L} . The second column shows the plots of J_q as a function of the elongation of the shapes. The product of the axial parameters k is kept constant. The minimum value of J_q is obtained for radially symmetric or regular shapes for which $\rho = 1$, and increases strictly with elongation.

the border is complex. This is examined for the superquadric shapes given by $\mathcal{L} = (R^\epsilon - x^\epsilon - y^\epsilon)^{0.15}$. The intensity within these shapes is almost flat because of the value of the exponent, 0.15. The size of the shapes remains fixed by adjusting the value of R . The family of the shapes is shown in figures 4 (a)–(e) which is: a pinched diamond for $\epsilon < 1$, a diamond for $\epsilon = 1$, a circle for $\epsilon = 2$, and tends to a square for $\epsilon > 2$. We plot the values of J_1 , J_2 , and J_3 as functions of ϵ in figures 4 (f), (g), and (h), respectively. They are computed directly from the image by discretizing the integral of equation (3).

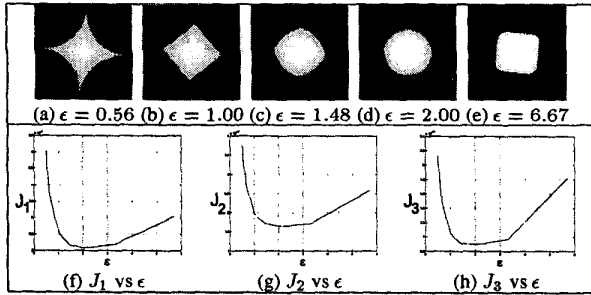


Figure 4. The values and the minimum of J_1 , J_2 , and J_3 for shapes with various borders. We use the superquadric shapes $\mathcal{L} = (R^\epsilon - x^\epsilon - y^\epsilon)^{0.15}$. The family of shapes with increasing ϵ consists of a pinched diamond, a diamond, a circle, and finally a square. In (f), (g), and (h) we can see the plots of J_1 , J_2 , and J_3 , respectively, as a function of ϵ . The minimum is attained for a circle shown in (d) for which $\epsilon = 2$. The values of J_1 , J_2 , and J_3 increase rapidly as the shape progresses from circular to pinched by decreasing ϵ .

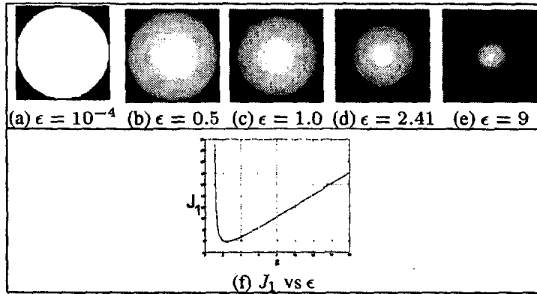


Figure 5. The Fisher information J_1 , and the minimum of the Fisher information for a shape of varying diffuseness. We use the elliptic shapes \mathcal{L}_ϵ shown in row (c) of figure 3. In (a) to (e) we can see the diffuseness progression as ϵ increases, namely, step transition, hemispherical, paraboloidal, resembling a Gaussian, and finally tending to an impulse. In (f) we can see the Fisher information given in row (c) of figure 3 as a function of ϵ . The minimum corresponds to the Gaussian resembling image shown in (d). The Fisher information increases rapidly as the diffuseness progresses to a step transition by decreasing ϵ .

The result of the numerical computations is that J_1 , J_2 , and J_3 are all minimized for a circle as shown in figure 4 (d) for which $\epsilon = 2$. Furthermore, they increase rapidly as the shape progresses from a circle to a pinched diamond by decreasing ϵ .

The value of J_q also varies depending on the diffuseness of the intensities within a shape. Shapes with smooth intensity changes across their boundaries minimize the rate of change of the histogram with resolution. The multiresolution histogram of shapes with abrupt intensity changes across their boundary, however, change faster with resolution. This

dependence is examined for the elliptical shapes given in row (c) of figure 3. We fix the size of the shapes by fixing R and vary the diffuseness by varying exponent ϵ . The progression of the diffuseness with increasing ϵ is shown in figures 5 (a)–(e) which is: a step transition for $\epsilon = 0$, a hemispherical for $\epsilon = 0.5$, a paraboloidal for $\epsilon = 1.0$, resembling a Gaussian for $\epsilon > 1.0$, and tending to an impulse as ϵ increases further.

The Fisher information for the elliptical shapes shown in row (c) of figure 3 is plotted as a function of diffuseness ϵ in figure 5 (f). Its minimum for all elongations is attained for an image with diffuseness similar to that of a Gaussian as shown in figure 5 (d) for which $\epsilon = 2.41$. As the elliptical shape progresses from a Gaussian to a step boundary by decreasing ϵ the Fisher information increases rapidly.

From the different examples we can see that J_q is minimized for shapes with rounded borders and increases with elongation. It increases even faster with warps that create spikes along the border. Linear transformations of J_q give the rate at which histogram bins change with resolution. Therefore, the results about the values of J_q also hold for the rate of change of histogram bins. Moreover, J_1 is minimized for smooth intensity transitions across the boundary, and increases for abrupt intensity transition across the shape boundary. The value of J_1 approximates the average rate of change of the histogram bins with resolution.

5. Multiresolution Histograms of Textures

We analyze the dependence of the multiresolution histogram on texture parameters using J_q . Several investigators have observed experimentally that the increase in entropy depends on the coarseness of the initial image [23, 31], the type of textures [27, 31], and the size of image regions [16]. In this work we quantify the dependence of J_q on the size of the texels, the extend at which neighboring texels overlap, and the placement pattern of the texels that can be regular or random. The models used for the texels are the shape models discussed in the previous section.

First, we examine the dependence of J_q on the number of texels within a texture of fixed size. We expect J_q to increase with the number of texels in a fixed area, since smaller texels are also “*sharper*”. Consider a simple texture model of $p \times p$ identical texels. Such textures result from tiling a texel $r = p^2$ times. To preserve the size of the texture, the texels are also contracted by a transformation A with determinant $\det A = 1/p^2$.

In the appendix we show that tiling a texel r times multiplies J_q by the factor r^{-q+1} and contracting a texel by transformation A multiplies J_q by the factor $(\det A)^{-q}$. Hence, the overall factor multiplying the value of J_q for the shape to give the value of J_q for the texture is $r^{-q+1} \cdot (\det A)^{-q} = p^2$. Such textures have the same histogram for all p since $(\det A)r = 1$, which preserves area [12].

Figures 6 (a) and (c) show two shapes and the textures formed by contracting and tiling them. Figures 6 (b) and (d)

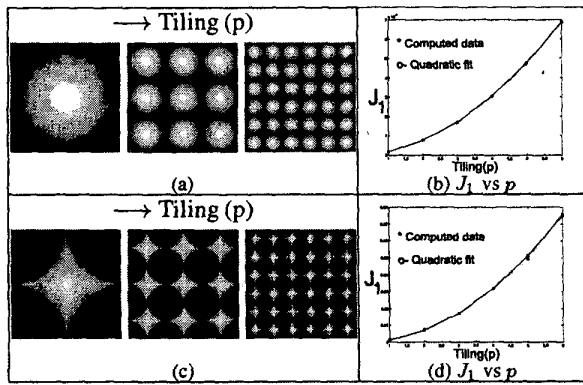


Figure 6. The Fisher information as a function of the tiling parameter p of the textures. In (a) and (c) we can see two shapes and the textures resulting by minifying and tiling them. In (b) and (d) we can see the Fisher information as a function of tiling parameter p for (a) and (c), respectively. Each plot shows the data obtained directly from the images as well as the quadratic fit. The two almost coincide as expected.

show their Fisher information, respectively, as a function of p . Each of the plots shows the Fisher information computed directly from the image as well as the quadratic fit. The fits almost coincide with the data as expected.

In the texture model shown in figure 6 adjacent texels were separated. We now examine cases where neighboring texels overlap. In particular, we examine how the Fisher information changes as the overlap among adjacent texels increases. We expect such textures to be less "sharp" and have a smaller Fisher information. Analytically, Gaussian filtering monotonically increases the size of texels and decreases their Fisher information [2, 17]. This is verified for textures consisting of mixtures of Gaussians of linearly increasing standard deviation shown in figure 7 (a). Figure 7 (b) shows that their Fisher information monotonically decreases with texel width, which in this case is the standard deviation of the Gaussians. Figure 7 (c) shows another texture with texels of linearly increasing width. Their Fisher information shown in figure 7 (d) also monotonically decreases with texel width.

In the previous texture models the texel placement was regular. We continue with textures whose texel placement is random. Randomness, on average, monotonically increases entropy and monotonically decreases Fisher information [2, 7, 17]. In figures 7 (e) and (g) we can see two textures with increasingly larger Gaussian noise perturbing the texels from their regular positions. In figures 7 (f) and (h) we can see the Fisher information as a function of the standard deviation of the Gaussian noise in pixel placement. In both cases, the Fisher information is monotonically decreasing with randomness.

Overall, the values of J_q for a texture of fixed size increase linearly with the number of texels. Hence, the rate at

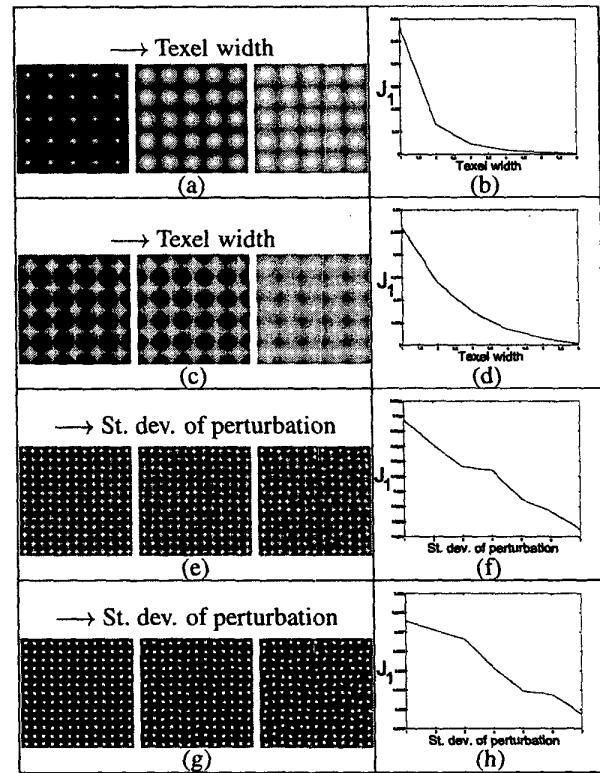


Figure 7. The Fisher information as a function of the width of overlapping texels and the randomness in the placement of texels. In (a) and (c) we can see two textures with texels of increasingly larger width. The plots in (b) and (d) show the Fisher information as a function of the texel width, respectively. In both cases the Fisher information decreases monotonically. In (e) and (g) we can see two textures with increasingly larger randomness in the placement of the texels. In (f) and (h), respectively, we can see plots of the Fisher information as a function of the standard deviation of the Gaussian noise in pixel placement. In both cases the Fisher information decreases monotonically.

which histogram bins change with resolution also increases with the number of texels. The value of J_1 , however, decreases with texel overlap and randomness in the placement of texels. Therefore, overlap and randomness decrease the average rate at which histogram bins change with resolution.

6 Experimental Verification

In the previous sections we have shown that the multiresolution histogram encodes spatial image variation. We now test directly its ability to discriminate between synthetic images with identical histograms and between natural textures that were histogram equalized. Note that J_q were only used to show the dependence and sensitivity of the multiresolution histogram on spatial image variation and are not computed




Figure 8. Several samples from the database of synthetic images. All the images in this database have the same histogram. Some of the images consist of texels that are dots, geometric shapes, and shapes of several other mathematical functions. In some cases the placement of the texels was regular and in others random.

for matching.

We construct a multiresolution image using four levels of the Burt-Adelson pyramid [6]. Each level is obtained by filtering with a Gaussian kernel and subsampling. Then, we compute the histogram of each of the four pyramid levels. The distance between two multiresolution histograms is the sum of the four individual L_1 distances between pairs of histograms corresponding to the same pyramid levels. The cost of computing the pyramid, the histograms, and the distance is of the order $O(N\sqrt{F})$, where N is the number of pixels, and F is the support of the separable Gaussian filter.

We first test the multiresolution histogram with a database of synthetic images. Many of the images in the database consist of texels and texel placements that were analyzed in the previous sections. More precisely, the texels include dots, circles, squares, several types of superquadrics, several types of triangles, and mathematical functions. Moreover, in some images the placement of the texels was regular and in others random. The database has 108 images of size 320×320 and 256 graylevels. Some images from the database are shown in figure 8. All the images have the same histogram, consist-




Figure 9. Matching of synthetic images using multiresolution histograms. The test images are shown in the columns marked Test Image and are corrupted with Gaussian noise of standard deviation 15 graylevels. The best match and second best match are shown in consecutive columns.

ing of 40% pixels with graylevel 25, and 60% of pixels with graylevel 230.

The best and second best matches for 16 test images corrupted with noise of standard deviation 15 graylevels are shown in figure 9. The rate of correct matches as a function of noise in the test images is shown in figure 12 (a). We can see that the multiresolution histogram is very robust to noise.

The multiresolution histogram was also tested with a database of natural textures that is a subset of the CURET database [10]. It consists of 305 images of size 100×100 and of 256 graylevels. The database contains images of 61 physical textures with five instances of each physical texture. The 5 instances differ in the illumination and viewing conditions. We equalized the histograms of all images in the database so that matching is based exclusively on differences in spatial image information rather than simply on differences in intensity histograms. Some equalized images are shown in figure 10. The best and second best matches for 12 equalized test images corrupted with noise of standard deviation 15 graylevels are shown in figure 11.

The matching performance of multiresolution histograms computed from histogram equalized images was compared to the performance of image intensity histograms computed




Figure 10. A few samples from the database of natural textures. It consists of 305 images of size 100×100 . It contains images of 61 physical textures, with five instances of each physical texture for five different illumination and viewing conditions. All the images are histogram equalized.

from non-equalized images. The comparison was performed using the database of natural textures for noise of standard deviation 2–37 graylevels. The results are shown in figure 12 (b). Clearly, the performance of multiresolution histograms is far superior to that of image intensity histograms.

In figures 9 and 11 we can see that images with similar shapes as well as similar texel shapes and placements were correctly matched. The multiresolution histogram performed better under noise on natural textures than on synthetic images since natural textures have a larger contrast and “sharpness”.

7. Discussion

We investigated the spatial image information embedded in a multiresolution histogram as well as the sensitivity of the multiresolution histograms to this information. We derived that the rate at which histogram bins change with image resolution are linear transforms of the values of J_q . In turn, the values of J_q are measures of image “sharpness” or variation that can encode shape and texel parameters. That is, we used J_q to quantify the relation and sensitivity of the multiresolution histogram to shape and texel parameters.

We also experimentally verified the properties of the multiresolution histogram. We found that they are able to discriminate between different images alone without the help of other features or filters. Moreover they are efficient, rotationally and translationally invariant, and robust to noise.

Appendix

Property: A uniform scaling transformation A multiplies J_q by $(\det A)^{-q}$.

Proof: As a result of transformation A and renormalization \mathcal{L} becomes $\mathcal{L}'(\mathbf{x}) = \frac{\mathcal{L}(A\mathbf{x})}{\det A}$. The value of J_q of image




Figure 11. Matching of natural textures using the multiresolution histogram. The test images are shown in the columns marked Test Image and are corrupted with Gaussian noise of standard deviation 15 graylevels. The best and second best matches are shown in consecutive columns.




Figure 12. Matching rate of multiresolution histograms under noise. The standard deviation of the noise varied from 2 to 37 graylevels. The plot in (a) is for the database of synthetic images. The plot in (b) is for the database of natural textures. In (b) we show the matching rate both of image histograms and multiresolution histograms. The matching rate of multiresolution histograms is very high in both (a) and (b). Moreover, in (b) the rate is significantly higher than that of intensity histograms.

$\mathcal{L}'(\mathbf{x})$ according to equation (3) is given by:

$$J_q \left(\frac{\mathcal{L}(A\mathbf{x})}{\det A} \right) = \int \left| \nabla \left(\frac{\mathcal{L}(A\mathbf{x})}{\det A} \right) \right|^2 \left(\frac{\mathcal{L}(A\mathbf{x})}{\det A} \right)^{q-2} d^2x$$

$$\stackrel{1}{=} (\det A)^{-q} \int |\nabla \mathcal{L}(\mathbf{x})|^2 \mathcal{L}^{q-2}(\mathbf{x}) d^2x$$

For isometries and uniform scalings $|\nabla \mathcal{L}(A\mathbf{x})|^2 =$

$\frac{1}{\det A} |\nabla \mathcal{L}(\mathbf{x})|^2$ which gives step 1 and the multiplicative factor. \square

Property: Tiling an image r times multiplies J_q by r^{-q+1} .

Proof: After tiling an image, \mathcal{L} , r times it is necessary to renormalize it to unit L_1 norm by dividing it by r to get \mathcal{L}' . Equation (3) for image \mathcal{L}' gives:

$$\begin{aligned} J_q(\mathcal{L}'(\mathbf{x})) &= r \int \left| \nabla \left(\frac{\mathcal{L}(\mathbf{x})}{r} \right) \right|^2 \left(\frac{\mathcal{L}(\mathbf{x})}{r} \right)^{q-2} d^2 \mathbf{x} \\ &= r^{-q+1} \int |\nabla \mathcal{L}(\mathbf{x})|^2 \mathcal{L}^{q-2}(\mathbf{x}) d^2 \mathbf{x} \end{aligned}$$

which gives the multiplicative factor. \square

References

- [1] J. Bach, C. Fuler, A. Gupta, A. Hampapur, B. Horowitz, R. Humphrey, R. Jain, and C. Shu. The Virage image search engine: An open framework for image management. In *Proc. of SPIE Conference on Storage and Retrieval for Image and Video Databases IV*, volume 2670, pages 76–87, March 1996.
- [2] A. Barron. Entropy and the central limit theorem. *The Annals of probability*, 14(1):336–342, 1986.
- [3] N. Blackman. The convolution inequality for entropy powers. *IEEE Transactions on Information theory*, 11:267–271, 1965.
- [4] J. Bonet, P. Viola, and J. Fisher. Flexible histograms: A multiresolution target discrimination model. In *Proc. of SPIE*, volume 3370, 1998.
- [5] J. D. Bonet. Multiresolution sampling procedure for analysis and synthesis of texture images. In *Proc. of the Annual Conference of Computer Graphics*, pages 361–368. SIGGRAPH, 1997.
- [6] P. Burt and E. Adelson. The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31(4):532–540, April 1983.
- [7] E. Carlen and A. Soffet. Entropy production by block variable summation and central limit theorems. *Communications in Mathematical Physics*, 140:339–371, 1991.
- [8] M. Carlotto. Histogram analysis using a scale–space approach. *PAMI*, 9(1):121–129, Jan. 1987.
- [9] S. Cohen and L. Guibas. The earth mover’s distance under transformation sets. In *Proc. of the 7th ICCV*, volume 2, pages 1076–1083, Kerkyra, Greece, Sep. 1999.
- [10] K. Dana, B. van Ginneken, S. Nayar, and J. Koenderink. Reflectance and texture of real-world surfaces. In *Proc. of CVPR*, pages 151–157, 1997.
- [11] L. Griffin. Scale–imprecision space. *Image and Vision Computing*, 15:369–398, 1997.
- [12] E. Hadjidemetriou, M. Grossberg, and S. Nayar. Histogram preserving image transformations. In *Proc. of CVPR*, volume 1, pages 410–416, 2000.
- [13] R. Haralick. Statistical and structural approaches to texture. *Proceedings of the IEEE*, 67:786–804, 1979.
- [14] D. Heeger and J. Bergen. Pyramid–based texture analysis/synthesis. In *Proc. of the Annual Conference of Computer Graphics*, pages 229–234. SIGGRAPH, 1995.
- [15] J. Huang, S. Kumar, M. Mitra, W. Zhu, and R. Zabih. Image indexing using color correlograms. In *Proc. of CVPR*, pages 762–768, 1997.
- [16] M. Jagersand. Saliency maps and attention selection in scale and spatial coordinates: An information theoretic approach. In *Proc. of the 5th ICCV*, pages 195–202, June 1995.
- [17] O. Johnson. Entropy inequalities and the central limit theorem. *Stochastic processes and their applications*, 88:291–304, 2000.
- [18] J. Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.
- [19] J. Koenderink and A. J. V. Doorn. The structure of locally orderless images. *IJCV*, 31(2–3):159–168, 1999.
- [20] J. Lee and B. W. Dickinson. Multiresolution video indexing for subband coded video databases. In *Proc. of SPIE Conference for Storage and Retrieval for Image and Video Databases II*, volume 2185, pages 162–173, San Jose, Feb. 1994.
- [21] B. Mel. Combining color shape and texture histogramming in a neurally inspired approach to visual object recognition. *Neural Computation*, 9:777–804, 1997.
- [22] W. Niblack. The QBIC project: Querying images by content using color, texture, and shape. In *Proc. of SPIE Conference on Storage and Retrieval for Image and Video Databases*, volume 1908, pages 173–187, April 1993.
- [23] A. Oomes and P. Snoeren. Structural information in scale–space. In *Proc. of Copenhagen Workshop on Gaussian Scale–space Theory*, volume 96/19, pages 48–57, Copenhagen, Denmark, May 1996. DIKU.
- [24] G. Pass and R. Zabih. Comparing images using joint histograms. *Multimedia Systems*, 7(3):234–240, May 1999.
- [25] A. Plastino, A. Plastino, and H. Miller. Tsallis nonextensive thermostatics and Fisher’s information measure. *Physica A*, 235:577–588, 1997.
- [26] B. Schiele and J. Crowley. Recognition without correspondence using multidimensional receptive field histograms. *IJCV*, 36(1):31–50, 2000.
- [27] J. Sporring and J. Weickert. Information measures in scale–spaces. *IEEE Transactions on Information Theory*, 45(3):1051–1058, April 1999.
- [28] A. Stam. Some inequalities satisfied by the quantities of information of Fisher and Shannon. *Information and Control*, 2:101–112, 1959.
- [29] M. Stricker and M. Orengo. Similarity of color images. In *In SPIE Conference on Storage and Retrieval for Image and Video Databases III*, volume 2420, pages 381–392, Feb. 1995.
- [30] M. Swain and D. Ballard. Color indexing. *IJCV*, 7(1):11–32, 1991.
- [31] M. Tanaka, T. Watanabe, and T. Mishima. Tsallis entropy in scale–spaces. In *Proc. of the SPIE Conference on Vision Geometry VIII*, volume 3811, pages 273–281, Denver, Colorado, July 1999.
- [32] C. Tsallis. Nonextensive statistics: Theoretical, experimental and computational evidences and connections. *Brazilian Journal of Physics*, 29(1), March 1999.
- [33] A. Witkin. Scale–space filtering. In *Proc. of IJCAI*, pages 1019–1022, Aug. 1983.
- [34] H. Zhang, C. Low, W. Smoliar, and J. Wu. Video parsing, retrieval and browsing: An integrated and content–based solution. *ACM Multimedia*, pages 15–24, 1995.