

Resolution Selection Using Generalized Entropies of Multiresolution Histograms

Efstathios Hadjidemetriou, Michael D. Grossberg, and Shree K. Nayar

Computer Science, Columbia University, New York, NY 10027, USA
{stathis, mdog, nayar}@cs.columbia.edu

Abstract. The performances of many image analysis tasks depend on the image resolution at which they are applied. Traditionally, resolution selection methods rely on spatial derivatives of image intensities. Differential measurements, however, are sensitive to noise and are local. They cannot characterize patterns, such as textures, which are defined over extensive image regions. In this work, we present a novel tool for resolution selection that considers sufficiently large image regions and is robust to noise. It is based on the generalized entropies of the histograms of an image at multiple resolutions. We first examine, in general, the variation of histogram entropies with image resolution. Then, we examine the sensitivity of this variation for shapes and textures in an image. Finally, we discuss the significance of resolutions of maximum histogram entropy. It is shown that computing features at these resolutions increases the discriminability between images. It is also shown that maximum histogram entropy values can be used to improve optical flow estimates for block based algorithms in image sequences with a changing zoom factor.

1 Introduction

The performances of many image analysis and interpretation algorithms depend on the image resolution at which they are applied. Therefore, the selection of the appropriate resolution is a critical preprocessing step. In this work we suggest the use of the generalized entropies of multiresolution histograms of an image for resolution selection. We first compute the multiresolution of an image, where resolution decreases with the standard deviation σ of a Gaussian filter [1,2]. Then, we transform the images of the various resolutions into their histograms. Finally, we compute the Tsallis generalized entropies of the histograms [3] for certain orders q . *We call the plot of the Tsallis histogram entropy of order q as a function of image resolution σ the entropy-resolution plot of order q .*

Histogram entropies have several properties which enable their use for resolution selection. One such property is that their values are directly related to the significance of a resolution. A high resolution busy image with many high count histogram bins has large histogram entropy. In the limit of low resolution, an image has uniform intensities, such as that shown in figure 1(g), and zero histogram entropy. Moreover, histogram entropies are inherently non-monotonic with resolution with one or more maxima. The local maxima correspond to significant resolutions.

Generalized histogram entropies are defined over the whole image. This is desirable when the detection of significant resolutions of an image requires the consideration of a certain image extent. This is obviously the case for textures which are commonly defined over extensive parts of images. Generalized histogram entropies can also detect multiple significant resolutions; for example, for textures with two levels of texel aggregation. Finally, they are robust to noise.

We first examine the entropy–resolution plots of images containing shapes and the sensitivity of these plots to the shape boundary. We then relate the entropy–resolution plots of shapes to those of textures. We also discuss the dependence of the entropy–resolution plots on their order q .

It is shown that computing image features at the resolution of maximum histogram entropy increases the discriminability between images. It is also shown that the maximum entropy values can be used to improve the performance of block based optical flow estimation for image sequences with a changing zoom factor.

2 Previous Work

The need for resolution selection was first realized, and exclusively used, in the context of edge detection. For example, Marr and Hildreth used multiple image resolutions to detect an arbitrary edge [4]. Later, several authors suggested the selection of image resolutions using the derivatives of pixels normalized by resolution [5,6]. To reduce noise, Lindeberg retained edges only if they fell along chains of connected pixels [6]. Elder and Zucker [7], as well as Marimont and Rubner [8], computed the edge magnitude at a number of resolutions and selected the lowest resolution that exceeded the sensor noise. Jeong and Kim [9] used pixel differential measurements to formulate a global regularizing function to select resolution.

Pixel differential measurements have more recently been used to select resolutions in problems other than edge detection. Lindeberg has used them to determine the characteristic length of objects [6], and to track features in an image sequence [10]. Another application has been to select resolutions that are appropriate to compute image features [11,12]. Finally, they have been used directly in image indexing systems [13,14].

It is not obvious, however, that differential measurements are the most appropriate technique for general resolution selection problems. Differential measurements have several limitations. In general, they are monotonically decreasing with image resolution. They can only be made non-monotonic by using normalized differential expressions. Both the expression and its normalization, however, are selected heuristically.

Pixel differential measurements are very sensitive to noise and are local. They cannot characterize textures, for example, which can cover extensive parts of an image. For large shapes, resolution selection requires two steps. First, the

detection of the boundary pixels, and then the connection of the boundary pixels. Finally, differential measurements give only one significant resolution [6,7,8].

To alleviate these problems several researchers have suggested resolution selection using image entropy, which is a global image function. Further, image entropy is not dependent upon edge measurements, and is robust to noise [15, 16,17]. Sporring and Weickert [17] have considered resolutions as significant if they correspond to local maxima of the rate of entropy increase. Jagersand computed the characteristic length of objects based on resolutions which correspond to a local maximum of the rate of increase of mutual image entropy [15].

It has been shown, however, that image entropy is monotonic with resolution [18,19]. Similarly, image mutual entropy and generalized entropies are monotonic with resolution [18,20]. Moreover, the rates of change of both image entropy and mutual entropy with resolution have also been shown to be monotonic [21,22]. The monotonicity of image entropies limits their ability to select resolutions. Hence, in this work we use exclusively histogram entropies, which are inherently non-monotonic with resolution. They are also based on extensive image regions, they are robust to noise, and can detect multiple significant resolutions.

A similar application of entropies of histograms has been as an auto-focus criterion in the problem of depth from defocus [23], where it has been assumed that the histogram entropy is a monotonic function of focus. Another similar technique has been to compute the change of histogram entropy over an image, of a specific resolution, by varying the extend of the image region considered [24, 25,26]. This technique has been used for image segmentation.

In an early work Wong and Vogel [27] computed the Shannon entropy of histograms at multiple image resolutions for 2–3 images. Moreover, Roldan et al [28] derived some mathematical results about the Shannon histogram entropy of a critically subsampled pyramid constructed with a box filter. They showed that the histogram entropy normalized by the pyramid level is monotonically decreasing [28]. They also showed that an image in the form of a Gibbs distribution has maximum entropy at all pyramid levels [28].

3 Tsallis Entropies of Histograms

The Shannon entropy measures the average information of all the densities in the histogram. In some cases, however, it is desirable to use an information measure whose value is influenced selectively by some of the intensities. For example, ignore high bin count intensities of a uniform background, or increase sensitivity to high bin count significant intensities. To achieve this we use nonlinear generalized entropies. We choose the Tsallis generalized entropies because they simplify the analytical part of this work.

The Tsallis generalized entropy of order q over histogram \mathbf{h} with unit L_1 norm is given by [3,29]:

$$S_q(\mathbf{h}) = \sum_{j=0}^{m-1} \frac{h_j - h_j^q}{q - 1} \quad (1)$$




Fig. 1. The image in (a) is binary. Filtering the image in (a) moderately gives the image in (c). Filtering the image in (a) significantly gives the image in (e). Eventually filtering the image in (a) gives the image in (g) of uniform intensity equal to the average intensity of the image in (a). Next to each image is its histogram.

where m is the number of intensity levels, and h_j is the histogram density of intensity j . In the limit $q \rightarrow 1$ the Tsallis generalized entropies reduce to the Shannon entropy.

The minimum value of all Tsallis entropies is zero and occurs for histograms with non-zero count in a single intensity bin. That is, for uniform intensity images such as that shown in figure 1 (g). All Tsallis entropies obtain their maximum value for histograms where all intensities are equally frequent.

The sensitivity of the entropies for histograms which do not have minimum or maximum entropy values depend not only on the histogram, but also on the order q of the entropy. The order q appears as exponent in the numerator of the entropy expression given in equation (1). Hence, the entropies of $q < 1$ have a large value for histograms with many low bin count intensities and a low value for histograms with many high bin count intensities [3,20]. Conversely, entropies of $q \geq 1$ have a large value for histograms with many frequent intensities and a small value for histograms with many low bin count intensities [3,20].

An illustrative example are histograms which consist of a central main lobe together with sidelobes. In histograms with a narrow central main lobe and wide sidelobes the side lobes contribute many low bin count intensities. Thus, entropies with $q < 1$ have a large value. As the width of the main lobe increases, the number of frequent intensities increases with it. Consequently, the orders q of the entropies which attain large values also increase. For very large values of q large entropy values are attained only for extreme histograms with a very wide main lobe. All other histograms have a small entropy values. Similarly, for very small value of q large entropy values are attained only for extreme histograms with a very narrow main lobe. Again, all other histograms have a small entropy values. Therefore, in this work, we examine entropies in a range of q close to unity, $-0.5 \leq q \leq 2.5$.

The number of independent generalized entropies of the histogram of a discrete image is finite. It is shown in appendix A.1 that this number can be as large

as the number of image pixels. In this work, we only use five different histogram entropies of orders $q = (-0.5, 0.2, 1.0, 1.3, 2.3)$.

4 Entropy–Resolution Plots of Shapes

A binary image of a shape has a histogram which consists of two impulses like those shown in figure 1 (b). Filtering the image with a Gaussian changes the binary image into grayscale by smoothing the intensity step along the shape boundary. In the histogram, Gaussian image filtering changes the impulses into wider distributions with long sidelobes like those shown in figure 1 (d). Each impulse in the histogram of the original binary image gives rise to a different distribution in the histogram. The low bin count intensities between the initial impulses correspond to the intensities of regions of steep intensity changes at the border between the shape and the background. This is particularly true for images whose initial histograms consist of distant impulses such as binary and halftone images.

When an image is filtered extensively the widths of the distributions in the histogram increase. Further, the peaks of the distributions move towards the mean image intensity. At some resolution the border between the shape and the background disappears and the two are joined. At that resolution, the individual distributions in the histogram meet to form a single distribution, like that shown in figure 1(f). Beyond that resolution the intensities of the image become uniform, the histogram has a count at a single bin such as that shown in figure 1 (h).

Limited image filtering simply increases the width of the main lobes of the histogram distributions. Therefore, it increases the order q of the histogram entropies which attain large values. Filtering, eventually, contracts the histogram towards the mean intensity and finally turns it into an impulse. Hence, the value of histogram entropy eventually decreases and finally becomes zero.

The images in figures 2 (a) and (g) have the same histogram. Some entropy–resolution plots of the two images are shown next to them in order of increasing q . The plots verify that the resolution of maximum entropy increases as a function of q . Note that the maxima of the entropy–resolution plots for which $q \geq 1$ occur at resolutions beyond those shown in the entropy–resolution plots of figure 2.

The resolutions at which the entropy–resolution plots attain large values also depend on the boundary of the shape in the original image. In appendix A.2 we examine the rates at which the Tsallis entropies of the histogram of image \mathcal{L} change with Gaussian filtering. We show that they are linearly proportional to Fisher information measures of m different orders q . The generalized Fisher information measures $J_q(\mathcal{L})$ are given by [20,21,22]:

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

where $\mathcal{L}(\mathbf{x})$ is the intensity value of image pixel \mathbf{x} . The Fisher information measures are nonlinearly weighted averages of pixel sharpness, which is defined as




Fig. 2. The binary images of two shapes are shown in (a) and (g). Next to each image is a sequence of entropy–resolution plots of orders $q = (-0.5, 0.2, 1.0, 1.3, 2.3)$ left to right. The maxima in the entropy–resolution plots with $q < 1$ of the convoluted shape in (g) occur at finer resolutions than those of the entropy–resolution plots of the same orders of the image in (a).

$|\nabla\mathcal{L}(\mathbf{x})/\mathcal{L}(\mathbf{x})|^2$ [22]. The generalized Fisher information measures increase non-linearly as the shapes become more eccentric or convoluted [30].

The maxima of the entropy–resolution plots in figures 2(h) and (i) occur at finer resolutions than the maxima of the plots in figures 2(b) and (c), respectively. This is because, compared to the shape in figure 2 (a), the boundary of the shape in figure 2 (g) is more convoluted with a higher rate of change of histogram entropy.

In summary, the entropy–resolution plots of binary shapes are initially increasing, they reach a maximum, and eventually decrease to zero. The resolution at which a plot reaches its maximum depends on its order q and on the boundary of the shape.

5 Entropy–Resolution Plots of Textures

In this section we consider a simple texture model where the texture is constructed by contracting shapes such as those discussed in the previous section to form texels. Subsequently the texels are tiled $r = p^2$ times to form a regular $p \times p$ texture. To preserve the size of the texture the texels are also contracted by a uniform transformation A whose determinant is given by $\det A = 1/r$. Since, $(\det A)r = 1$, textures for all p have the same histogram [31].

In appendix A.2 we show that the rate of change in the entropy–resolution plots of a shape multiplied by p^2 gives the rate of change in the entropy–resolution plots of the corresponding regular texture. That is, in going from a shape to the corresponding regular texture the horizontal resolution coordinate of the entropy–resolution plot is scaled by $1/p$. Thus, the resolution of maximum entropy is scaled by the same factor. As discussed in section 4 the resolution of maximum entropy for shapes is attained at the resolution at which the border between the shape and the background disappears. For texture images, the resolution of maximum entropy corresponds to the resolution at which

different texels come in contact. Note that the discussion above assumes circular boundary conditions.




Fig. 3. The images in (a) and (g) are two synthetic textures. Next to each image is a sequence of some of its entropy–resolution plots of orders $q = (-0.5, 0.2, 1.0, 1.3, 2.3)$ left to right. The texture in (a) can be contracted to give part of the texture in (g). Similarly, the horizontal σ axis of each entropy–resolution plot of the image in (a) can be contracted to give part of the corresponding entropy–resolution plot of the image in (g).

The textures shown in figure 3(a) and figure 3(g) are obtained by minifying the same shape, shown in figure 2(a). The width of the texels in figure 3(g) is 70% of those in figure 3(a). Both textures have the same histogram entropies. The entropy–resolution plots are shown next to the textures. The horizontal coordinate σ of the entropy–resolution plots of the image in figure 3(a), shown in figure 3 (b–f), can indeed be contracted by 0.7 to give part of the entropy–resolution plots of the image in figure 3(g), shown in figures 3 (h–l). Moreover, the resolution of maximum entropy of the image in figure 3(a) can be multiplied by 0.7 to give the resolution of maximum entropy of the image in figure 3(g).

Above we examined texture models and entropy–resolution plots of regular textures. We now examine entropy–resolution plots of textures with random texel placement. Randomness monotonically increases image entropies and decreases the generalized Fisher informations [20], which are linearly related to the rate at which the histogram entropies change with resolution. Therefore, randomness in the placement of texels decreases the rate at which the histogram entropies change [20] and shifts their maxima to lower resolutions.

In figure 4(a) we show a regular texture and in figure 4(c) we show a random texture. Both textures are synthetic with identical texels and histograms. The maximum of the entropy–resolution plot of $q = 1.3$ of the regular texture shown in figure 4(b) occurs at a lower σ than the maximum of the entropy–resolution plot of the same order of the random texture in figure 4(d). The random textures consist of aggregates of texels. The distance between aggregate texels is larger than the distance between individual texels regularly placed. Thus, the maximum for randomly placed textures in figure 4 is shifted to a higher σ value.




Fig. 4. The textures in (a) and (c) are synthetic with identical texels and histograms. The texel placement in the texture in (a) is regular, whereas the texel placement in the texture in (c) is random. Next to each image is its entropy–resolution plot of order $q = 1.3$. The maximum in the entropy–resolution plot of the regular texture occurs at a finer resolution than the maximum in the entropy–resolution plot of the random texture.

In summary, the resolution of maximum entropy depends on the entropy order, on the shape of the texels, on the size of the texels, and on the distance between them. Moreover, it depends on their placement pattern.




Fig. 5. The images, shown in (a) and (e), have identical foreground but differ in the area of the background. Next to each image are three of its entropy–resolution plots of orders $q = (-0.5, 1.0, 2.3)$ left to right. The entropy–resolution plot of order $q = -0.5$ is not sensitive to a change in the area of the background. The entropy–resolution plots of order $q \geq 1$, however, are significantly affected by the increase in the background area.

6 Significance of Entropy–Resolution Plots of $q \neq 1$

The entropy–resolution plots of order $q < 1$ are significant in cases where it is desirable to diminish the effect of the variation of the background. Figures 5 (a) and (e) show two images with identical foregrounds [32], but backgrounds of different size. In figure 5(a) the background is much smaller than in figure 5(e).




Fig. 6. In (a), (d), (g), and (j) are 4 images with fine structure. Next to each image are two of its entropy–resolution plots of orders $q = (0.2, 2.3)$ left to right. The images in (a) and (d) contain texels of more than one shape and size. The images in (g) and (j) contain two levels of structure. This is revealed by the two maxima in the entropy–resolution plots of order $q = 2.3$.

Some of the entropy–resolution plots are shown next to the images. These plots show that increasing the area of the background shifts the maxima of all entropy–resolution plots to higher σ values. This is because the filtering required to reach equilibrium, that is an image of uniform intensity, is more extensive.

The shift of the entropy–resolution plots of $q < 1$ is much smaller than that of the entropy–resolution plots of $q \geq 1$. This is because entropies with $q < 1$ are more sensitive to low bin count intensities and less sensitive to frequent intensities, which in this case correspond to the background [3]. On the other hand, entropy–resolution plots of $q > 1$ are significantly affected by the frequent image intensities of the background.

The entropy–resolution plots of order $q > 1$ are significant in cases where it is desirable to increase sensitivity to high bin count significant intensities of an image and decrease the sensitivity to low bin count ones. This can be used to decrease noise sensitivity. Also, in some histograms there are no low bin count intensities. Instead, all intensities are frequent. In such a case generalized entropies of $q \leq 1$ are not sensitive to histogram changes with image resolution, and it is necessary to use entropies of order $q > 1$.

A histogram with many high count bins can arise as a binary image containing multiple different shapes is filtered. The two impulses in the histogram of such a binary image meet more than once. In consecutive meetings after the first, however, the histogram has many high count bins. Therefore, only entropies of order $q > 1$ are sensitive to maxima corresponding to consecutive meetings.

In figures 6 (a), (d), (g), and (j) we can see textures that have texels of more than one shape and size. Next to each image we show 2 of its entropy–resolution plots of orders $q = (0.2, 2.3)$. Each of the images in figures 6 (a) and (d) has texels of 2 different sizes. This causes their entropy–resolution plots in figures 6 (c), and (e–f) to have 2 maxima. The image in figure 6 (g) shows a checkerboard pattern with stripes superimposed upon it. In the entropy–resolution plots in

figures 6 (h–i) the first maximum corresponds to the size of the stripes. The second maximum corresponds to the size of the squares.

The image in figure 6 (j) shows a board with electronic components. The entropy–resolution plot in figure 6(l) has 2 maxima. The first maximum corresponds to the size of the components of the board. The second maximum corresponds to the size of the entire board. As we can see in figure 6, the multiple maxima in the entropy–resolution plots are more obvious for entropies of order $q > 1$. Also, entropy–resolution plots of orders $q > 1$ are less sensitive to noise.

In summary, entropy–resolution plots of order $q < 1$ are useful for shapes embedded in background, whereas plots of order $q > 1$ are useful for textures with multiple significant resolutions.

7 Applications

We demonstrate two different applications of the entropy–resolution plots. We first show that the resolutions at which the entropy–resolution plots are maximized can also maximize the discriminability between images in a database. We then use the entropy–resolution plot to adapt the block size in block based optical flow estimation for a zoom–in sequence. We show that motion estimation is improved.

7.1 Discrimination between Images

The discriminability of an image feature can be improved by appropriately selecting the image resolution at which it is computed. This is particularly true for image features which are non–monotonic with resolution. We demonstrate this for two features. The first is a vector which consists of the energy $E(C)$ and the Shannon entropy $S_1(C)$ of the cooccurrence matrix C [33]. The cooccurrence matrix is computed over a 3×3 neighborhood. The energy is given by $E(C) = \sum_{i,j=0}^{m-1} c_{ij}^2$, where c_{ij} is the element of the cooccurrence matrix at the i^{th} row and j^{th} column. The Shannon entropy is given by $S_1(C) = \sum_{i,j=0}^{m-1} c_{ij} \log c_{ij}$.

The second feature are the parameters of a Gauss Markov random field (GMRF) [34]. This model assumes that the intensity at each pixel is a linear combination of the intensities in the surrounding 5×5 window. Thus, each pixel gives an equation for 24 unknowns. Over the entire image this leads to an overconstraint system of linear equations that we solved with least squares.

The discriminability of the feature vectors can be measured using a database of images consisting of multiple disjoint classes [35]. Consider a database with k classes of p_j images in class j . The discriminability D can be modeled by the ratio of the between class scatter V_b to the within class scatter V_w , that is $D = \frac{V_b}{V_w}$ [35]. The within–class scatter is given by: $V_w = \sum_{j=1}^k \sum_{i=1}^{p_j} \|f_i - \mu_j\|_2^2$, where μ_j is the mean of class j . The between–class scatter is given by $V_b = \sum_{i=1}^k p_i \|\mu_i - \mu_{tot}\|_2^2$, where μ_{tot} is the mean of the entire database. We compute discriminability D as a function of image resolution σ for 2 databases of images.




Fig. 7. The images in (a) and (c) show fingerprints of two different persons. The images in (b) and (d) show the entropy–resolution plots of order $q = 1.3$ for the fingerprints in (a) and (c), respectively. The plot in (e) shows the discriminability–resolution plot of the cooccurrence matrix features for the fingerprint database. The plot in (f) shows the discriminability–resolution plot of the GMRF features for the fingerprint database. For both features the resolution of maximum discriminability coincides with that of the maxima of the entropy–resolution plots.




Fig. 8. The images in (a) and (c) show two textures. The images in (b) and (d), respectively, show their entropy–resolution plots of order $q = 1.3$. The plot in (e) shows the discriminability–resolution plot of the cooccurrence matrix features for the texture database. The plot in (f) shows the discriminability–resolution plot of the GMRF features for the texture database. For both features the resolutions of maximum discriminability is approximately the same as the resolutions of the maxima of the entropy–resolution plots.

The first database consists of images of fingerprints of 4 different persons [36]. The fingerprints of each person form a different class. Thus, there are $k = 4$ different classes. Each class consists of 8 images of fingerprints from different fingers of the same person. The intensity resolution of all images is 8 bits and their size is 364×256 pixels. Figures 7(a) and (c) show two images from two different classes. Next to each image, in figures 7(b) and (d), respectively, is its entropy–resolution plot of order $q = 1.3$. The order of the entropy–resolution plots is chosen to be greater than one since the structure of the images is fine. Note that the resolution at which the entropy is maximum is the effective distance between the ridges of the fingerprints.

The plot in figure 7(e) shows the discriminability of the cooccurrence matrix features between the fingerprint classes as a function of resolution, namely, the discrimination–resolution plot. The plot in figure 7(f) shows the discrimination–resolution plot for the GMRF features. The range of the horizontal coordinate σ of all plots in figure 7 is the same. The discriminability in figures 7(e–f) is maximum at the resolution at which the entropy–resolution plots are maximized.

The second database consists of natural textures. It is a subset of the CURET database [37]. It consists of images from $k = 61$ classes of textures. Each class consists of 5 instances of a physical texture. The 5 instances differ in the illumination and viewing directions. To account for these differences all images were histogram equalized. In total there are 305 images. The intensity resolution of all images is 8 bits and their size is 100×100 pixels.

In figure 8(a) and (c) we see images of two different textures from the database. Next to each image, in figures 8(b) and (d), respectively, is its entropy–resolution plot of order $q = 1.3$. The plot in figure 8(e) shows the discrimination–resolution plot of the cooccurrence matrix features. The plot in figure 8(f) shows the discrimination–resolution plot of the GMRF features. The range of the horizontal coordinate σ of all the plots in figure 8 is identical. The discriminability in figures 8(e–f) is maximized at approximately the resolution at which the entropy–resolution plots of the images achieve their maxima.

The plot in figure 8(e) has two maxima. This is because the images in the texture database have different resolutions of maximum histogram entropy. The peak at the finer resolution is higher because most textures are fine. Note that the discriminability in the plots in figures 7 (e–f) and figure 8(e) increases for large σ . This is because at those resolutions all images in a class have uniform intensity and are similar. Thus, the between–class scatter becomes very small.

7.2 Optical Flow Estimation

In optical flow estimation using block matching a very important parameter is the block size. Clearly, the appropriate block length b must be related to the size of the objects or the coarseness of the textures in an image sequence. Thus, for a zoom–in sequence the block size should vary.

We adapt the block size in a zoom–in sequence using the entropy–resolution plots of order $q = 0.2$. The order of the entropy–resolution plots used is less than one since the sequence zooms into the central part of the image which is surrounded by background. We use the maximum value of the entropy, $S_{0.2}^{max}$, which corresponds to the dominant image structure. The length of the block b_i used for the i^{th} image is given by the block length used for image $(i - 1)$ multiplied by the factor by which the maximum entropy changes. That is, $b_i = b_{i-1}(S_{0.2}^{i,max}/S_{0.2}^{i-1,max})$.

Figures 9(a) and (d) show 2 frames from two different synthetic zoom–in sequences where the optical axis coincides with the geometric center of the image. The size of the images is 281×224 pixels and 301×200 pixels, respectively. In each image in figures 9(b) and (e) we can see superimposed the vector field estimated from the zoom–in sequence. In figure 9(c) and (f) we plot the motion estimation error as a function of the block length. For the adaptive algorithm the horizontal axis is the block length used for the first image in the sequence. The error measure is proportional to the negative of the cosine of the angle by which the motion vectors deviate from the correct motion direction.

The dotted lines in the plots in figures 9(c) and (f) show the motion estimation error for constant block size. The solid lines show the motion estimation




Fig. 9. The images in (a) and (d) are from 2 different zoom-in sequences. The images in (b) and (e) show images of the zoom-in sequences with motion vectors superimposed upon them. The plots in (c) and (f) are the motion estimation error as a function of the block size. The error of the constant block size algorithm is the dotted line and the error of the adaptive block size algorithm is the solid line. For both sequences the minimum error of the adaptively selected block size algorithm is smaller than the minimum error of the constant block size algorithm.

error for adaptively selected block size. In both plots the minimum error of the adaptively selected block size method is smaller than the minimum error of the fixed block size method.

8 Summary

We have suggested and examined the use of multiresolution histogram entropies for resolution selection. Multiresolution histogram entropies are appropriate for this objective because they are non-monotonic with resolution and robust to noise. Moreover, they represent sufficiently large image regions, and can detect multiple significant resolutions.

We examined the plot of the Tsallis entropy of order q of the histogram as a function of Gaussian image resolution, namely, the entropy-resolution plot of order q . We examined the entropy-resolution plots of images of shapes, regular textures, and random textures. Moreover, we discussed the fact that for $q < 1$ the entropy-resolution plots emphasize the foreground and for $q > 1$ the entropy-resolution plots are more sensitive to multiple significant resolutions.

We showed that the discriminability between images is larger at the resolution at which the histogram entropies are maximized. We also showed that the entropy–resolution plots can be used to adapt the block size in an optical flow algorithm and improve its performance.

A Appendix

A.1 Number of Independent Histogram Entropies

Property: The number of independent entropies of the histogram of an image can be at most equal to the number of the pixels in the image.

Proof: The histogram can be considered to be a distribution. It has been shown that the entropies of a distribution can be transformed into the histogram of the distribution [26]. Therefore, the entropies of the histogram can be transformed into the histogram of the histogram. The dimensionality of the histogram of the histogram can be as large as the number of image pixels. Therefore, the number of independent histogram entropies can be at most equal to the number of image pixels. □

A.2 Histograms Entropies and Fisher Information Measures

Lemma: The rate of change of the histogram with respect to Gaussian image filtering is linearly proportional to m different generalized Fisher information measures of the image.

Proof: The histogram \mathbf{h} is equivalent to any complete set of Lyapunov exponents [17,26]. In particular, a histogram of m graylevels can be transformed linearly to a vector of m distinct Tsallis entropies. The transformation is given by $\mathbf{S}(\mathcal{L}) = R \mathbf{h}$, where R is an $m \times m$ matrix. Appropriate selection of the entropy orders makes R invertible [17]. Therefore, we have $\mathbf{h} = R^{-1}\mathbf{S}(\mathcal{L}) = T \mathbf{S}(\mathcal{L})$, where T is also an $m \times m$ matrix.

The rate at which the histogram changes with respect to image resolution is given by $\frac{d\mathbf{h}}{d\sigma} = T \frac{d\mathbf{S}(\mathcal{L})}{d\sigma}$. In turn, the rate at which the Tsallis entropies of the image change with respect to Gaussian filtering is proportional to the generalized Fisher information measures of the same order J_q given by equation (2) [20,21, 22]. Therefore, we obtain $\frac{d\mathbf{h}}{d\sigma} = T \mathbf{J}(\mathcal{L})$, where $\mathbf{J}(\mathcal{L}) = (J_1(\mathcal{L}) J_2(\mathcal{L}) \dots J_m(\mathcal{L}))^T$. In particular, for histogram density of intensity j we have $\frac{dh_j}{d\sigma} = \sum_{l=0}^{m-1} t_{jl} J_l(\mathcal{L})$, where t_{jl} are elements of matrix T . □

Property: The rate of change of Tsallis entropies of the histogram, $S_q(\mathbf{h})$, with respect to Gaussian filtering of the image is linearly proportional to generalized Fisher information measures of the image of m different orders q .

Proof: The rate of change of the Tsallis entropies of the histogram with respect to image resolution, σ , is obtained by differentiating equation (1) with respect to σ to obtain $\frac{dS_q(\mathbf{h})}{d\sigma} = \frac{1}{q-1} \sum_{j=0}^{m-1} \left((1 - qh_j^{q-1}) \frac{dh_j}{d\sigma} \right)$. By substituting the lemma above in this relation we obtain: $\frac{dS_q(\mathbf{h})}{d\sigma} = \frac{1}{q-1} \sum_{j,l=0}^{m-1} \left((1 - qh_j^{q-1}) t_{jl} J_l(\mathcal{L}) \right)$. □

A.3 Effect of Texel Repetition on the Rate of Change of the Generalized Entropies of the Histogram

Lemma 1: Consider regular textures constructed from shapes as described in section 5. The relation between the generalized Fisher information of a shape, $J_q(\mathcal{L}_s)$, to that of the corresponding regular $p \times p$ texture, $J_q(\mathcal{L}_t)$, is: $J_q(\mathcal{L}_t) = p^2 J_q(\mathcal{L}_s)$.

The proof of Lemma 1 is given by Hadjidemetriou et al [30].

Property: The rate of change of the entropy of the histogram of the texture is p^2 times the rate of change of the entropy of the histogram of the shape.

Proof: The property derived in appendix A.2 for textures becomes $\frac{dS_q(\mathbf{h}_t)}{d\sigma} = \frac{1}{q-1} \sum_{j,l=0}^{m-1} \left((1 - qh_j^{q-1}) t_{jl} J_l(\mathcal{L}_t) \right)$. The corresponding relation holds for shapes $\frac{dS_q(\mathbf{h}_t)}{d\sigma}$. Substituting $J_q(\mathcal{L}_t) = p^2 J_q(\mathcal{L}_s)$ in the relation for textures we obtain $\frac{dS_q(\mathbf{h}_t)}{d\sigma} = p^2 \frac{dS_q(\mathbf{h}_s)}{d\sigma}$. \square

References

1. Koenderink, J.: The structure of images. *Biological Cybernetics* **50** (1984) 363–370
2. Witkin, A.: Scale–space filtering. In: *Proc. of IJCAI*. (1983) 1019–1022
3. Tsallis, C.: Nonextensive statistics: Theoretical, experimental and computational evidences and connections. *Brazilian Journal of Physics* **29** (1999)
4. Marr, D., Hildreth, E.: Theory of edge detection. *Proc. Royal society of London B* **207** (1980) 187–217
5. Korn, A.: Toward a symbolic representation of intensity changes in images. *IEEE Trans. on PAMI* **10** (1988) 610–625
6. Lindeberg, T.: Edge detection and ridge detection with automatic scale selection. *IJCV* **30** (1998) 117–154
7. Elder, J., Zucker, S.: Local scale control for edge detection and blur estimation. *IEEE Trans. on PAMI* **20** (1998) 699–716
8. Marimont, D., Rubner, Y.: A probabilistic framework for edge detection and scale selection. In: *Proc. of ICCV*. Volume 1. (1998) 207–214
9. Jeong, H., Kim, C.: Adaptive discrimination of filter scales for edge detection. *IEEE Trans. of PAMI* **14** (1992) 579–585
10. Bretzner, L., Lindberg, T.: Feature tracking with automatic selection of spatial scales. *CVIU* **71** (1998) 385–392
11. Almansa, A., Lindberg, T.: Fingerprint enhancement by shape adaptation of scale–space operators with automatic scale selection. *IEEE Trans. on IP* **9** (2000) 2027–2042
12. Wiltschi, K., Lindberg, T., Pinz, A.: Classification of carbide distributions using scale selection and directional distributions. In: *Proc. of ICIP*. Volume 3. (1997) 122–125
13. Chomat, O., Verdieri, V., Hall, D., Crowley, J.: Local scale selection for Gaussian based description techniques. In: *Proc. ECCV*. Volume 1., Dublin (2000) 117–133
14. Crowley, J., A.C.Parker: A representation for shape based on peaks and ridges in the difference of low-pass transform. *IEEE Trans. on PAMI* **6** (1984) 156–170
15. Jagersand, M.: Saliency maps and attention selection in scale and spatial coordinates: An information theoretic approach. In: *Proc. of ICCV*. (1995) 195–202

16. Oomes, A., Snoeren, P.: Structural information in scale-space. In: Proc. of Workshop on Gaussian Scale-space Theory. Volume 96/19., Copenhagen, DIKU (1996) 48–57
17. Sporring, J., Weickert, J.: Information measures in scale-spaces. *IEEE Trans. on IT* **45** (1999) 1051–1058
18. Barron, A.: Entropy and the central limit theorem. *The Annals of probability* **14** (1986) 336–342
19. Carlen, E., Soffet, A.: Entropy production by block variable summation and central limit theorems. *Communications in Mathematical Physics* **140** (1991) 339–371
20. Plastino, A., Plastino, A., Miller, H.: Tsallis nonextensive thermostatics and Fisher’s information measure. *Physica A* **235** (1997) 577–588
21. Blackman, N.: The convolution inequality for entropy powers. *IEEE Trans. on IT* **11** (1965) 267–271
22. Stam, A.: Some inequalities satisfied by the quantities of information of Fisher and Shannon. *Information and Control* **2** (1959) 101–112
23. Thum, C.: Measurement of the entropy of an image with application to image focusing. *Optica Acta* **31** (1984) 203–211
24. Kadir, T., Brady, M.: Saliency, scale and image description. *IJCV* **45** (2001) 83–105
25. Koenderink, J., Doorn, A.: The structure of locally orderless images. *IJCV* **31** (1999) 159–168
26. Sporring, J., Colios, C., Trahanias, P.: Generalized scale-space. In: Proc. of ICIP. Volume 1. (2000) 920–923
27. Wong, A., Vogel, M.: Resolution-dependent information measures for image analysis. *IEEE Trans. on Systems Man and Cybernetics* **SMC-7** (1977) 49–61
28. Roldan, R., Molina, J., Aroza, J.: Multiresolution-information analysis for images. *Signal Processing* **24** (1991) 77–91
29. Tanaka, M., Watanabe, T., Mishima, T.: Tsallis entropy in scale-spaces. In: Proc. of the SPIE Conference on Vision Geometry VIII. Volume 3811. (1999) 273–281
30. Hadjidemetriou, E., Grossberg, M., Nayar, S.: Spatial information in multiresolution histograms. In: Proc. of CVPR. Volume 1. (2001) 702–709
31. Hadjidemetriou, E., Grossberg, M., Nayar, S.: Histogram preserving image transformations. In: Proc. of CVPR. (2000) I:410–416
32. Robotics Institute of CMU: (1998) CIL image database.
33. Haralick, R.: Statistical and structural approaches to texture. *Proceedings of the IEEE* **67** (1979) 786–804
34. Mao, J., Jain, A.: Texture classification and segmentation using multiresolution simultaneous autoregressive models. *Pattern Recognition* **25** (1992) 173–188
35. Fugunaka, K.: Introduction to Statistical Pattern Recognition. Academic Press, New York (1990)
36. Michigan State University: (2000) FVC database.
37. Dana, K., Nayar, S., van Ginneken, B., Koenderink, J.: Reflectance and texture of real-world surfaces. In: Proc. of CVPR. (1997) 151–157