# Applications of Locally Orderless Images

Bram van Ginneken and Bart M. ter Haar Romeny

Image Sciences Institute, Utrecht University, The Netherlands.
E-mail: bram@isi.uu.nl, URL: http://www.isi.uu.nl/

**Abstract.** In a recent work [1], Koenderink and van Doorn consider a family of three intertwined scale-spaces coined the *locally orderless image* (LOI). The LOI represents the image, observed at inner scale $\sigma$, as a local histogram with bin-width $\beta$, at each location, with a Gaussian-shaped region of interest of extent $\alpha$. LOIs form a natural and elegant extension of scale-space theory, show causal consistency and enable the smooth transition between pixels, histograms and isophotes. The aim of this work is to demonstrate the wide applicability and versatility of LOIs. We consider a range of image processing tasks, including variations of adaptive histogram equalization, several methods for noise and scratch removal, texture rendering, classification and segmentation.

## 1 Introduction

Histograms are ubiquitous in image processing. They embody the notion that for many tasks, it is not the spatial order but the intensity distribution within a region of interest that contains the required information. One can argue that even at a single location the intensity has an uncertainty, and should therefore be described by a probability distribution: physical plausibility requires non-zero imprecision. This led Griffin [2] to propose a *scale-imprecision space* with spatial scale parameter $\sigma$ and an intensity, or tonal scale $\beta$, which can be identified with the familiar *bin-width* of histograms.

Koenderink and Van Doorn [1] extended this concept to *locally orderless images* (LOIs), an image representation with three scale parameters in which there is no local but only a global topology defined. LOIs are *local histograms*, constructed according to scale-space principles, viz. without violating the causality principle. As such, one can apply to LOIs the whole machinery of techniques that has been developed in the context of scale-space research.

In this paper, we aim to demonstrate that LOIs are a versatile and flexible framework for image processing applications. The reader may conceive this article as a broad feasibility study. Due to space limitations, we cannot give thorough evaluations for each application presented. Obviously, local histograms are in common use, and the notion to consider histograms at different scales (soft binning) isn't new either. Yet we believe that the use of a consistent mathematical framework in which all scale parameters are made explicit can aid the design of effective algorithms by reusing existing scale-space concepts. Additional insight may be gained by taking into account the behavior of LOIs over scale.

M. Nielsen et al. (Eds.): Scale-Space'99, LNCS 1682, pp. 10–21, 1999.

## 2 Locally orderless images

We first briefly review locally orderless images [1] by considering the scale parameters involved in the calculation of a histogram:

- the inner scale $\sigma$ with which the image is observed;
- the outer scale, or extent, or scope $\alpha$ that parameterizes the size of the field of view over which the histogram is calculated;
- the scale at which the histogram is observed, tonal scale, or bin-width $\beta$.

The locally orderless images $H(\mathbf{x_0}, i; \sigma, \alpha, \beta)$ are defined as the family of histograms, i.e. a function of the intensity $i$, with bin-width $\beta$ of the image observed at scale $\sigma$ calculated over a field of view centered around $\mathbf{x_0}$ with extent $\alpha$. The unique way to decrease resolution without creating spurious resolution is by convolution with Gaussian kernels [3] [4]. Therefore Gaussian kernels are used for $\sigma$, $\alpha$ and $\beta$. We summarize this with a recipe for calculating LOIs:

1. Choose an inner scale $\sigma$ and blur the image $L(x; \sigma)$ using the diffusion

$$\Delta_{(\mathbf{x})} L(\mathbf{x}; \sigma) = \frac{\partial L(\mathbf{x}; \sigma)}{\partial \frac{\sigma^2}{2}}. \tag{1}$$

2. Choose a number of (equally spaced) bins of intensity levels $i$ and calculate the "soft isophote images", representing the "stuff" in each bin through the Gaussian gray-scale transformation

$$R(\mathbf{x}, i; \sigma, \beta) = \exp(-\frac{(L(\mathbf{x}; \sigma) - i)^2}{2\beta^2}) \tag{2}$$

3. Choose a scope $\alpha$ for a Gaussian aperture, normalized to unit amplitude

$$A(\mathbf{x}; \mathbf{x_0}, \alpha) = \exp \frac{-(\mathbf{x} - \mathbf{x_0})(\mathbf{x} - \mathbf{x_0})}{2\alpha^2} \tag{3}$$

and compute the locally orderless image through convolution

$$H(\mathbf{x_0}, i; \sigma, \alpha, \beta) = \frac{A(\mathbf{x}; \mathbf{x_0}, \alpha)}{2\pi\alpha^2} * R(\mathbf{x}, i; \sigma, \beta). \tag{4}$$

Note that $H(\mathbf{x_0}, i; \sigma, \beta, \alpha)$, is a stack of isophote images, and therefore has a dimensionality 1 higher than that of the input image.

The term *locally orderless* image refers to the fact that we have at each location the probability distribution at our disposal, which is a mere orderless set; the spatial structure within the field of view $\alpha$ centered at $\mathbf{x}$ has been obliterated. This is the key point: instead of a (scalar) intensity, we associate a probability distribution with each spatial location, parameterized by $\sigma$,$\alpha$,$\beta$. Since a distribution contains more information then the intensity alone, we may hope to be able to use this information in various image processing tasks.

The LOI contains several conventional concepts. The original image and its scale-space $L(\mathbf{x}; \sigma)$ that can be recovered by integrating $iH(\mathbf{x_0}, i; \sigma, \alpha, \beta)$ over $i$.

The "conventional" histogram is obtained by letting $\alpha \to \infty$. The construction also includes families of isophote images, which for $\beta > 0$ are named *soft isophote images* by Koenderink. And maybe even more important, by tuning the scale parameters the LOI can fill intermediate stages between the image, its histogram and its isophotes. This can be useful in practice. The framework generalizes trivially to nD images or color images, if a color metric is selected.

## 3 Median and maximum mode evolution

If we replace the histogram at each location with its mean, we obtain the input image $L(\mathbf{x}; \sigma)$ blurred with a kernel with width $\alpha$. This holds independently of $\beta$, since blurring a histogram does not alter its mean. If, however, we replace the histogram with its median or its maximum mode (intensity with highest probability), we obtain a diffusion with scale parameter $\alpha$ that is reminiscent of some non-linear diffusion schemes. The tonal scale $\beta$ works as a tuning parameter that determines the amount of non-linearity. For $\beta \to \infty$, the median and the maximum mode are equal to the mean, so the diffusion is linear. Griffin [2] has studied the evolution of the median, and the stable mode (defined as the mode surviving as $\beta$ increases), which is usually equal to the maximum mode. He always sets $\sigma = \alpha$. This ensures that for $\alpha \to \infty$ the image attains its mean everywhere, as in linear diffusion. With only a few soft isophote level images in the LOI, maximum mode diffusion also performs some sort of quantizing, and one obtains piecewise homogenous patches with user-selectable values. This can be useful, e.g. in coding for data-compression and knowledge driven enhancements.

## 4 Switching modes in bi-modal histograms

Instead of replacing each pixel with a feature of its local histogram, such as the median or the maximum mode, we can perform more sophisticated processing if we take the structure of the local histograms into account. If this histogram is bi-modal, this indicates the presence of multiple "objects" in the neighborhood of that location. Noest and Koenderink [5] have suggested to deal with partial occlusion in this way.

**Fig. 1.** Left: Text hidden in a sinusoidal background, dimensions $230 \times 111$, intensities in the range $[0, 1]$. Middle: bi-modal locations in an LOI of $\sigma = 0$, $\beta = 0.15$ and $\alpha = 1.5$. Right: bi-modal locations have been replaced with the high mode. Text is removed and the background restored.

Consider locations with bi-modal histograms. We let locations "switch mode", i.e. if they are in the high mode (that is, their original value is on the right of the minimum mode in between the two modes), we replace their value with the low mode (or vice versa, depending on the desired effect). The idea is to replace a bright/dark object with the most likely value that the darker/brighter object that surrounds it has, namely the low/high maximum mode. Note that this a two-step process: the detection of bi-modal locations is a segmentation step, and replacing pixels fills in a value from its surroundings, using statistical information from only those pixels that belong to the object to be filled in.

This scheme allows for a scale-selection procedure. For fixed $\sigma, \beta, \alpha$, there may be locations with more than two modes in their local distribution. This indicates that it is worthwhile to decrease $\alpha$, focusing on a smaller neighborhood, until just two modes remain. Thus we use a locally adaptive $\alpha$, ensuring that the replaced pixel value comes from information from locations "as close as possible" to the pixel to be replaced.

We have applied this scheme successfully for the removal of text on a complicated background (Figure 1), the detection of dense objects in chest radiographs, and noise removal. Figure 2 shows how shot noise can be detected and replaced with a probable value, obtained from the local histogram. The restoration is near perfect. Figure 3 shows three consecutive frames from an old movie with severe deteriorations. To avoid to find bi-modal locations due to movement between frames, we considered two LOIs, one in which the frame to be restored was the first and one in which it was the last image. Only locations that were bi-modal in both cases were taken in consideration. Although most artifact are removed, there is ample room for improvements. One can verify from the middle

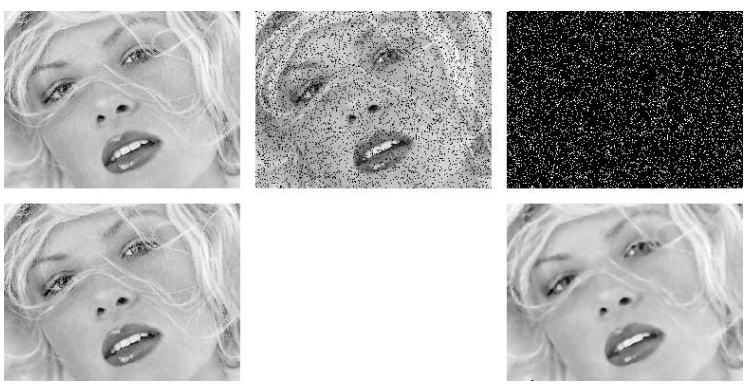

**Fig. 2.** (top-left) Original image, 249 × 188 pixels, intensities scaled to $[0, 1]$; (top-middle) Original image with shot noise. This is the input image for the restoration procedure; (top-right) Locations in (top-middle) with bi-modal histograms and pixels in the lowest mode using $\sigma = 0, \beta = 0.04, \alpha = 0.5$. (bottom-left) Restoration using mode-switching for bi-modal locations gives excellent results; (bottom-right) Restoration using using 5x5 median filter. This removes most shot noise, but blurs the image.

**Fig. 3.** The top row shows three consecutive frames (337 × 271 pixels, intensities scaled to $[0, 1]$) from a movie with severe local degradations, especially in the first frame shown. LOIs were calculated with $\sigma = 0$, $\beta = 0.1$, and $\alpha = 2.0$. The second row shows the detected artifact locations for each frame. The bottom row shows the restored frames, using histogram mode switching.

column in Fig. 3 that the hand which makes a rapid movement has been partly removed. Distinguishing such movements from deteriorations is in general a very complicated task, that would probably require a detailed analysis of the optic flow between frames.

## 5 Histogram transformations

Any generic histogram can be transformed into any other histogram by a non-linear, monotonic gray-level transformation. To see this, consider an input histogram $h_1(i)$ and its cumulative histogram $\int_{-\infty}^{i} h_1(i')di' = H_1(i)$ and the desired output histogram $h_2(i)$ and $H_2(i)$. If we replace every $i$ with the $i'$ for which $H_1(i) = H_2(i')$ we have transformed the cumulative histogram $H_1$ into $H_2$ and thus also $h_1$ into $h_2$. Since cumulative histograms are monotonically increasing, the mapping is monotonically increasing as well.

An example is histogram equalization. When displaying an image with a uniform histogram (within a certain range), all available gray levels or colors will be used in equal amounts and thus "perceptual contrast" is maximal. The idea to use local histograms (that is, selecting a proper $\alpha$ for the LOI) for equalization, to obtain optimal contrast over each region in the image stems from the 1970s

original $\alpha = 8; \beta = 0.4$ $\alpha = 8; \beta = 0.2$ $\alpha = 4; \beta = 0.2$

**Fig. 4.** A normal PA chest radiograph of 512 by 512 pixels with intensities in the range [0,1]. (a) Original image, in which details in lung regions and mediastinum are not well visible due to the large dynamic range of gray levels. (b)-(d) Adaptive histogram equalization (AHE) based on the LOI with $\sigma = 0$ and 3 combinations of $\alpha$ and $\beta$.

[6] and is called *adaptive histogram equalization* (AHE). However, it was noted that these operations blow up noise in homogeneous regions. Pizer et al. [7] proposed to *clip* histograms, viz. for each bin with more pixels than a certain threshold, truncate the number of pixels and redistribute these uniformly over all other bins. It can be seen that this ad hoc technique amounts to the same

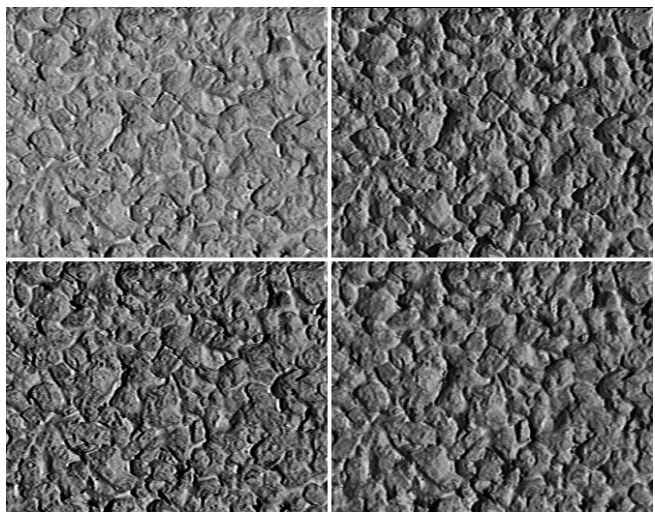

**Fig. 5.** Top-left is an image (332 by 259 pixels) of rough concrete viewed frontally, illuminated from 22°. Top-right: the same material illuminated from 45°. Bottom-left shows the top-left image with its histogram mapped to the top-right image to approximate the change in texture. Bottom-right shows the result of local histogram transformation, with $\alpha = 2$. The approximation is especially improved in areas that show up white in the images on the left. These areas are often partly shadowed with illumination from 45°, and using a local histogram may correctly "predict" such transitions.

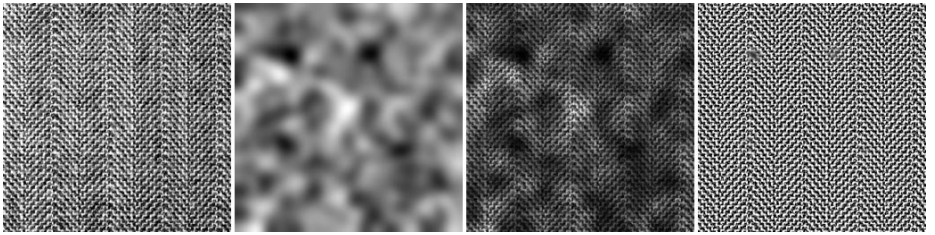

(a) (b) (c) (d)

**Fig. 6.** (a) A texture from the Brodatz set [11], resolution $256^2$, intensities in the range $[0, 1]$. (b) Blurred Gaussian noise, scaled to range from $[0, 1]$. (c) Multiplication of (a) and (b). (d) Reconstruction of (a) from (c) from the LOI with $\sigma = 0, \beta = 0.1, \alpha = 2$ and computing for each point a mapping to the local histogram at the randomly chosen location (80,80).

effect as increasing $\beta$ in the LOI; notably, for $\beta \to \infty$, AHE has no effect. Thus we see that the two scale parameters $\alpha$ and $\beta$ determine the size of structures that are enhanced and the amount of enhancement, respectively. Figure 4 shows a practical example of such a continuously tuned AHE for a medical modality (thorax X-ray) with a wide latitude of intensities.

An alternative to histogram equalization is to increase the standard deviation of the histogram by a constant factor, which can be done by a linear gray level transformation, or variations on such schemes [8]. Again, the LOI provides us with an elegant framework in which the scale parameters that determine the results of such operations are made explicit.

Another application of histogram transformation is to approximate changes in texture due to different viewing and illumination directions [9]. In general, the textural appearance of many common real-world materials is a complex function of the light field and viewing position. In computer graphics it is common practice, however, to simply apply a projective transformation to a texture patch in order to account for a change in viewing direction and to adjust the mean brightness using a bi-directional reflection distribution function (BRDF), often assumed to be simply Lambertian. In [9] it is shown that this gives poor results for many materials, and that histogram transformations often produce far more realistic results. A logical next step is to consider *local* histogram transformations. An example is shown in Figure 5, using a texture of rough concrete taken from the CURET database [10]. Instead of using one mapping function for all pixel intensities, the mapping is now based on the pixel intensity and the intensities in its surroundings. Physical considerations make clear that this approach does make sense: bright pixels which have dark pixels due to shadowing in their neighborhood are more likely to become shadowed for more oblique illumination than those that are in the center of a bright region.

Finally, histogram transformations can be applied to restore images that have been corrupted by some noise process, but for which the local histogram properties are known or can be estimated from the corrupted image. Such cases are

encountered frequently in practice. Many image acquisition systems contain artifacts that are hard to correct with calibration schemes. One example in medical image processing is the inhomogeneity of the magnetic field of an MR scanner or of the sensitivity of MR surface coils, leading to low frequency gradients over the image. A generated example is shown in Figure 6 where we multiplied a texture image with Gaussian noise. By randomly choosing a point in the corrupted image and computing the mapping that transforms each local histogram to the local histogram at that particular location we obtain the restored image in Figure 6(d). Apart from it being low frequency, the LOI method does not make any assumption about the noise process and works for multiplicative, additive, and other kinds of noise processes.

## 6 Texture classification and discrimination

LOIs can be used to set up a framework for texture classification. The histogram is one of the simplest texture descriptions; the spatial structure has been completely disregarded and only the probability distribution remains. This implies that any feature derived from LOIs is rotationally invariant. There are several ways possible to extend LOIs:

**Locally orderless derivatives**
Instead of using $L(\mathbf{x};\sigma)$ as input for the calculation of LOIs, we can use $L_n^\theta(\mathbf{x};\sigma)$, which denotes the $n$th order spatial derivative of the image at scale $\sigma$ in the direction $\theta$. These images can be calculated for any $\theta$ from a fixed set of basis filters in several ways, for a discussion see [12], [13]. For $n = 0$, these locally orderless derivatives (LODs) reduce to the LOIs. Alternatively, one could choose another family of filters instead of directional derivatives of Gaussians, such as differences of offset Gaussians [14], [15], or Gabor filters [16].

**Directional locally orderless images**
Another way to introduce orientation sensitivity in LOIs is to use anisotropic Gaussians as local regions of interest. This would extend the construction with an orientation $0 < \theta < \pi$, and an anisotropy factor.

**Cooccurrence matrices**
Haralick [17],[18] introduced cooccurrence matrices, which are joint probability densities for locations at a prescribed distance and orientation. Texture features can be computed from these matrices. It is straightforward to modify the LOIs into a construction equivalent to cooccurrence matrices. It leads to joint probability functions as a function of location.

Results from psychophysics suggest that if two textures are to be pre-attentive discriminable by human observers, they must have different spatial average $\int\int_{T_1} R(x,y)$ and $\int\int_{T_2} R(x,y)$ of some locally computed neural response $R$ [14]. We use this as a starting point and compute features derived from LODs, av-

eraged over texture patches. Averaging will give identical results for any $\alpha$ if we use linear operations on LODs to compute local features. Thus we include non-linear operations on the local histograms as well. An obvious choice is to use higher-order moments.

Which combinations of scales $\sigma, \alpha, \beta$ are interesting to select? First of all, we should have $\sigma < \alpha$, otherwise local histograms are peaked distributions and higher order moments of these distributions are fully predictable. Furthermore, in practice $\sigma$ and $\beta$ will often be mutually exclusive. Haralick [18] defined texture as a collection of typical elements, called tonal primitives or textons, put together according to placement rules. As the scope $\alpha$ is much larger than the spatial size of the textons, the local histograms will not change much anymore. Therefore it does not make sense to consider more than one LOI with $\alpha$ much larger than the texton size. Using $\alpha = \infty$ is the obvious choice for this large scope histogram. Secondly, if we vary $\sigma$ at values below the texton size, we study the spatial structure of the textons. For $\sigma$ much larger than the texton size, we are investigating the characteristics of the placement rules.

We performed an experiment using texture patches from 16 large texture images from the USC-SIPI database available at `http://sipi.usc.edu`, 11 of which originated from the Brodatz collection [11]. From each texture, 16 nonoverlapping regions were cropped and subsampled to a resolution of $128 \times 128$. Intensity values of each patch were normalized to zero mean and unit variance. Figure 7 shows one patch for each texture class.

We classified with the nearest-neigbor rule and the leave-one-out method. A small set of 9 features was already able to classify 255 out of all 256 textures correctly. This set consisted of 3 input images, $L_0(\mathbf{x}; \sigma = 0)$ (used in feature 1-3), $L_1^{0^\circ}(\mathbf{x}; \sigma = 1)$ (used in feature 4-6), and $L_0^{90^\circ}(\mathbf{x}; \sigma = 1)$ (used in feature 7-9) for which we calculated the averaged second moment (viz. the local standard deviation) for $\beta = 0.1$ and $\alpha = 1, 2, \infty$.

To gain more insight into the discriminative power of each of the calculated features separately, we performed the classification for any possible combination of 1, 2 or 3 out of the 9 features. The best and worst results are given in Table 1. It is interesting to see that there is no common feature in the best single set, the best 2 features and the 3 best ones, which indicates that all features contain discriminant power. Since we use only 2nd moments, features are invariant to gray-level inversion. This can be solved by adding higher moments, which was apparently unnecessary for the test set we considered.

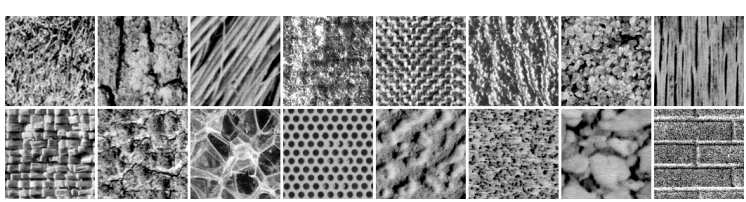

**Fig. 7.** The 16 different textures used in a texture classification experiment.

| features set | best single | worst single | best 2 | worst 2 | best 3 | worst 3 | full set |
|---|---|---|---|---|---|---|---|
| features used | 7 | 3 | 4,9 | 2,3 | 1,5,8 | 3,8,9 | all |
| result | 47.6% | 12.9% | 91.4% | 41.0% | 99.2% | 71.5% | 99.6% |

**Table 1.** Classification results for various combinations of features.

## 7 Texture segmentation based on local histograms

Many "general" (semi-)automatic segmentation schemes are based on the notion that points in spatial proximity with similar intensity values are likely to belong to the same object. Such methods have problems with textured areas, because the intensity values may show wild local variations. A solution is to *locally* compute texture features and replace pixel values with these features, assuming that pixels that belong to the same texture region will now have a similar value. The framework of LOIs is ideally suited to be used for the computation of such local features. One could use LODs, or another extension of LOIs put forward in the previous section. Shi and Malik [15] have applied their normalized cut segmentation scheme to texture segmentation in this way, using local histograms and the correlation between them as a metric.

Here we present an adapted version of a seeded region growing (SRG), that is popular in medical image processing. For $\alpha \to 0$, our adaptation reduces to a scheme very similar to the original SRG. This is directly due to the fact that LOIs *contain* the original image.

SRG segments an image starting from seed regions. A list is maintained of pixels connected to one of the regions, sorted according to some metric. This metric is originally defined as the squared difference between the pxiel intensity and the mean intensity of the region. The pixel at the top of the list is added to

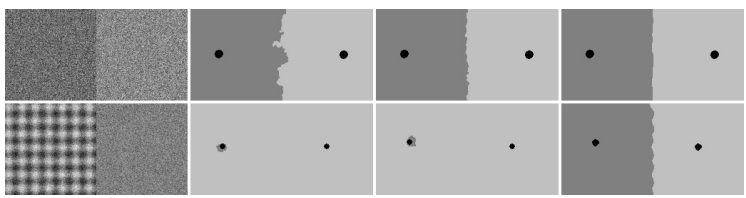

**Fig. 8.** Top row, from left to right: A 256 × 128 test image composed of two homogenous regions with intensity 0 and 1 and Gaussian noise with zero mean and unit variance. An LOI with $\sigma = 0$ and $\beta = 0.2$ and $\alpha = 0, 1, 4$, respectively is used for seeded region growing from the two seeds shown in white. Since the mean of the two region is different, regular seeded region growing ($\alpha = 0$) works well. Bottom row: same procedure for a partly textured image; the left half was filled with $\sin(x/3) + \sin(y/3)$, the right half was set to zero, and Gaussian noise with zero mean and $\sigma = 0.5$ was added. Regular seeded region growing now fails, but if $\alpha$ is large enough, the segmentation is correct.

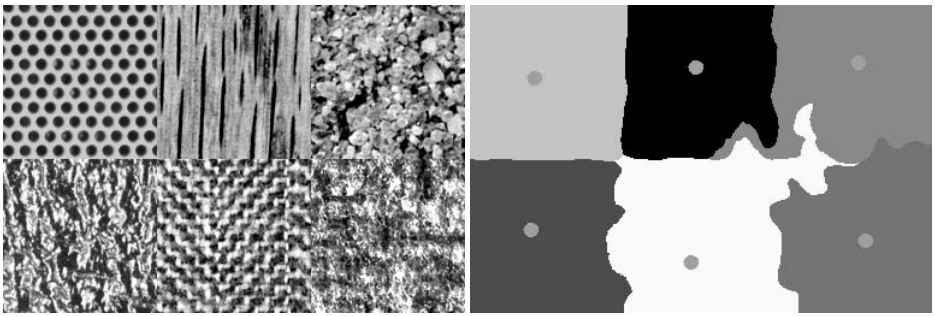

**Fig. 9.** Left: A test image composed of 6 texture patches of pixel size 128 × 128 each. Intensity values in each patch are normalized to zero mean and unit variance. Right: The result of segmentation with seeded region growing based on a LOI with $\sigma = 0$, $\beta = 0.2$ and $\alpha = 8$. The circles are the seeds.

the region it is connected to, and the neighbors of the added pixel are added to the list. This procedure is repeated until all pixels are assigned to a region.

We propose to compute a metric based on the local histograms of a pixel and a region. We subtract the histograms and take the sum of the absolute values of what is left in the bins. For $\alpha \to 0$ this reduces to a scheme similar to the original scheme, except that one considers for the region the global mode instead of the mean (most likely pixel value instead of the mean pixel value). Figures 8 to 10 illustrate the use of seeded region growing based on local histograms.

**Fig. 10.** Top row, left: Wildlife scene with leopard, size 329 × 253 pixels, intensities scaled between $[0, 1]$; Bottom row, left: A locally ($\sigma = 8$) normalized version of the input image; Middle and right: Segmentation by SRG based upon LOI with $\sigma = 0$, $\beta = 0.05$ and $\alpha = 0, 4$, respectively. Note how well the textured area is segmented in the lower right image.

## 8 Concluding remarks

In the applications presented, we have used many aspects of LOIs. They are a natural extension of techniques that usually use pixels, e.g. seeded region growing. They extend techniques that use "conventional" histograms with an extra degree of freedom, e.g. histogram transformation techniques. Other applications exploit the behavior of LOIs over scale to obtain non-linear diffusions, for scale selection in noise removal, and to derive texture features. We conclude that LOIs are image representations of great practical value.

**Acknowledgments**
This work is supported by the IOP Image Processing funded by the Dutch Ministry of Economic Affairs. We thank Peter van Roosmalen for providing the movie sequence.

## References

1. J.J. Koenderink and A.J. van Doorn. The structure of locally orderless images. *IJCV*, 31(2/3):159–168, 1999.
2. L.D. Griffin. Scale-imprecision space. *Image and Vision Comp.*, 15:369–398, 1997.
3. J.J. Koenderink. The structure of images. *Biol. Cybern.*, 50:363–370, 1984.
4. J. Weickert, S. Ishikawa, and A. Imiya. On the history of Gaussian scale-space axiomatics. In *Gaussian Scale-Space Theory*, pp. 45–59. Kluwer, Dordrecht, 1997.
5. A.J. Noest and J.J. Koenderink. Visual coherence despite transparency or partial occlusion. *Perception*, 19:384, 1990.
6. R.A. Hummel. Image enhancement by histogram transformation. *Comp. Graph. and Im. Proc.*, 6:184–195, 1977.
7. S. Pizer, E. Amburn, J. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. ter Haar Romeny, J. Zimmerman, and K. Zuiderveld. Adaptive histogram equalization and its variations. *Comp. Vis., Graph. and Im. Proc.*, 39:355–368, 1987.
8. D.C. Chan and W.R. Wu. Image contrast enhancement based on a histogram transformation of local standard deviation. *IEEE TMI*, 17(4):518–531, 1998.
9. B. van Ginneken, J.J. Koenderink, and K.J. Dana. Texture histograms as a function of illumination and viewing direction. *IJCV*, 31(2/3):169–184, 1999.
10. K.J. Dana, B. van Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. on Graphics*, 18(1):1–34, 1999.
11. P. Brodatz. *Textures*. Dover, New York, 1966.
12. W.T. Freeman and E.H. Adelson. The design and use of steerable filters. *IEEE PAMI*, 13(9):891–906, 1991.
13. P. Perona. Deformable kernels for early vision. *IEEE PAMI*, 17(5):488–499, 1995.
14. J. Malik and P. Perona. Preattentive texture discrimination with early vision mechanisms. *JOSA-A*, 7(5):923–932, 1990.
15. J. Shi and J. Malik. Self inducing relational distance and its application to image segmentation. In *ECCV*, 1998.
16. A.C. Bovik, M. Clark, and W.S. Geisler. Multichannel texture analysis using localized spatial filters. *IEEE PAMI*, 12(1):55–73, 1990.
17. R.M. Haralick, K. Shanmugam, and I. Dinstein. Textural features for image classification. *IEEE PAMI*, 3:610–621, 1973.
18. R.M. Haralick. Statistical and structural approaches to texture. *Proc. of the IEEE*, 67(5):786–804, 1979.