Local Mode Filtering

J. van de Weijer R. van den Boomgaard

Intelligent Sensory Information Systems
Faculty of Science, University of Amsterdam
Kruislaan 403, 1098 SJ Amsterdam, The Netherlands
{joostw, rein}@science.uva.nl

Abstract

Linear filters have two major drawbacks. First, edges in the image are smoothed with increasing filter size. Second, by extending the filters to multi-channel data, correlation between the channels is lost. Only a few researchers have explored the possibilities of mode filtering to overcome these problems. In this article mode filtering will be motivated from both a local histogram with tonal scale and a robust statistics point of view. The tonal scale is proved to be equal to the scale of the error norm function within the robust statistics framework.

Instead of the more commonly studied global mode, our focus is on the local mode. It preserves edges and details and is easily extensible to multi-channel data. A generalization of the spatial Gaussian filtering to a spatial and tonal Gaussian filter is used to iterate to the local mode. Results on color images include successful noise attenuation while preserving edges and detail by local mode filtering.

1. Introduction

Linear filtering is a widely used framework for many image processing tasks. Apart from the desired noise reduction and scale selection, this technique has two major drawbacks: 1. details and edges are smoothed, and 2. due to the absence of a natural basis for vector ordering extension to multi-channel data is done by applying the operation on the channels separately. Hence the correlation between the channels is neglected. Especially for these reasons several nonlinear filters have been proposed in literature.

The distribution of the pixel values in a neighborhood, i.e. a local histogram, contains significant information about the zero order local image structure. All maxima in a distribution are called modes. The global mode is the highest maximum of the distribution and hence is the most occurring value in the neighborhood. In this paper the importance of the local mode is demonstrated. This is the mode found by incorporating the pixel value in the center of the neighborhood as a-priori knowledge to the iteration procedure. The local mode has the property of noise reduction without the two drawbacks mentioned before.

A solid framework for working with local histograms is the imprecision space proposed by Griffin [4]. The imprecision space is determined by two scales. Next to a spatial scale a tonal scale is introduced to incorporate the imprecision of the measurement. Griffin uses the imprecision space to study the evolution of the median and the stable mode to find perceptual edges.

Koenderink et al. [6] extend the imprecision space with a third scale, the inner scale. This is the scale at which the image is observed. The combination of the three scales is called locally orderless images. This is because a global topology, but not a local topology is defined. Van Ginneken et al. [11] have considered a number of applications on grey scale images. In [10] the framework is extended to multi-channel images and results of global mode filtering for color images are given.

The field of robust statistics studies the sensitivity of statistical methods to deviations of the data from the underlying models [5]. It has been used in many computer vision problems to reduce the influence of outliers (e.g. [2]). Black et al. [3] prove that anisotropic diffusion [7] can be related to robust statistics.

In this paper the connection between the histogram spaces of Griffin [4] and Koenderink [6] and robust statistics is made. The tonal scale is proved to be equal to the scale of the error norm function in robust statistics theory. The local mode will be motivated from both frameworks.

The classical Gaussian filter is extended with a tonal weight. Not only do pixels further away in spatial space, $x - x_0$, have less influence, but also pixels further away in sensor space, $f(x) - f(x_0)$. This easy to implement filter is shown to iterate to the local mode of the

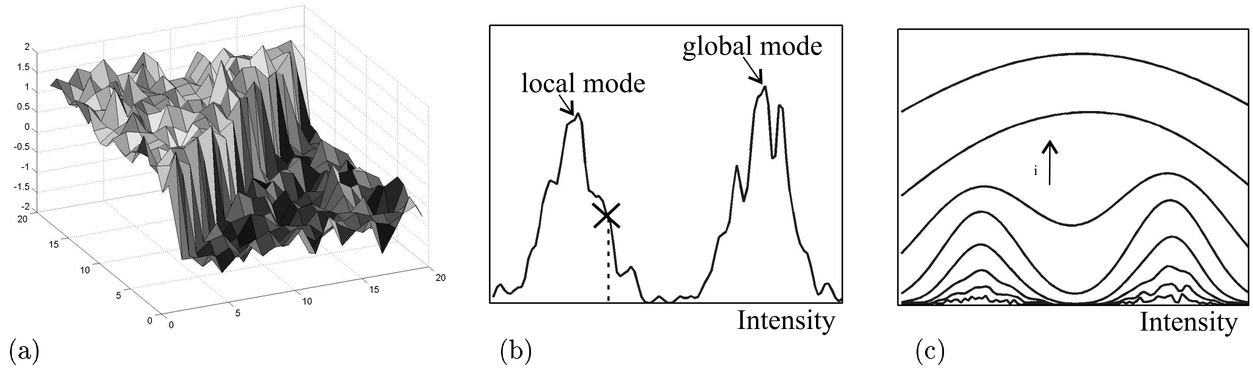

Figure 1: a) edge image. b) histogram of figure 1a. c) smoothed versions of histogram

pixel when keeping the neighborhood data constant. If used iteratively on the image with small spatial extend it behaves like anisotropic diffusion [8].

This paper is organized as follows. In section 2 a histogram framework and a robust statistical framework will be used as a foundation for the local mode filter. In section 3 a generalized version of the spatial Gaussian filter is proposed which is shown to iterate toward the local mode. In section 4 results of local mode filtering are shown. Section 5 finishes with concluding remarks.

2 Local Mode filtering

In figure 1b the histogram of the edge signal of figure 1a is depicted. The two peaks concur with the two areas in figure 1a. All maxima in the distribution are called modes. We make a distinction between two modes:

- the *global mode* which is the highest mode of the distribution.
- the *local mode*, which is related to a starting point of iteration. In figure 1b the local mode of the point indicated by a cross is given. Starting an iterative search for a maximum from this cross results in the indicated local mode. The concept of both global and local mode is easily extensible to multi-channel data.

In the following section the local and global mode are motivated from a local histogram point of view. Subsequently the same is done within the framework of robust statistics.

2.1 Local histogram space

In this section the *local histogram space* (LHS) is introduced. It was originally introduced by Griffin [4]. The LHS is a combination of the spatial space and the sensor space. In fact a local histogram is constructed for every position in the image.

Given an image $f : R^n \rightarrow R^m$ where n is the spatial dimension and m the dimension of the sensor space. The LHS at spatial-scale and sensor-scale zero is defined as

$$H\left(x, i, \sigma_x = 0, \sigma_i = 0\right) = \delta\left(i - f(x)\right) \tag{1}$$

with δ the Dirac delta function. The dimensionsionality of the LHS is $n + m$.

The influence of the neighboring values on the local histogram is obtained by convolution along the spatial plane, indicated by $\otimes_x$, of the LHS at scale zero with a Gaussian

$$H(x, i, \sigma_x, 0) = H\left(x, i, 0, 0\right) \otimes_x G(x, \sigma_x), \tag{2}$$

where the Gaussian of dimensionality n is defined as

$$G\left(x, \sigma\right) = \frac{1}{\left(2\pi\right)^{\frac{n}{2}} \sigma^n} e^{-\frac{x^T \cdot x}{2\sigma^2}}. \tag{3}$$

At this point the LHS consists of a Gaussian weighted local histogram at every position with the size of the neighborhood determined by the *spatial scale* σ_x.

Finally, the *tonal scale* of the histogram is tuned. This is the scale at which the histogram is observed. This is achieved by convolution with a m dimensional Gaussian without normalization

$$H(x, i, \sigma_x, \sigma_i) = H\left(x, i, \sigma_x, 0\right) \otimes_i \left(\left(2\pi\right)^{\frac{m}{2}} \sigma^m G\left(i, \sigma_i\right)\right) \tag{4}$$

The reason for convolution with a Gaussian without normalization becomes clear in the next subsection where the connection between this kernel and the error norm is made. Clearly, it has no influence on the location of the modes.

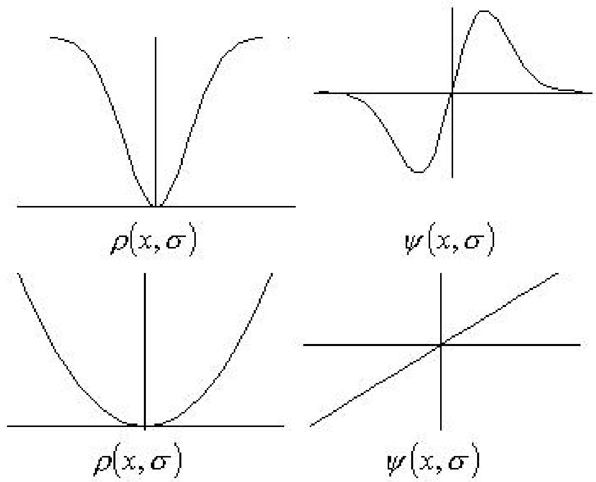

Figure 2: top: the error norm of eq. 8 and its influence function. bottom: least squares error norm and influence function

The LHS is dependent on the two scale parameters: the spatial scale σ_x, and the tonal scale σ_i. The spatial scale σ_x regulates the size of the neighborhood of the local histogram. With a spatial scale of zero, only the center pixel itself is represented in the local histogram. In the case $\sigma_x = \infty$ the LHS at every position is equal to the histogram of the image observed at scale σ_i .

The tonal scale σ_i is the scale at which the local histogram is smoothed. In figure 1c the evolution of the histogram is depicted for different sensor space scales. For $\sigma_i = 0$ the traditional mode is found, i.e. the most common value of a set. For $\sigma_i = \infty$ both the local and the global mode are equal to applying a classical Gaussian filter.

For small details in a larger background the local histogram will be dominated by two peaks. A large peak for the background and a smaller peak for the detail. Applying a global mode filter to such images erases the detail. In this case local mode filtering is preferable. The center pixel of the neighborhood will be taken as a starting point of iteration for the local mode and the averaged color of the detail will be found instead of the background color. A consequence of this approach is that local mode filtering will consider speckle noise also as a detail, and preserve it.

2.2 Robust statistics and mode filtering

Robust statistics is used to reduce the influence of outliers [5]. Given a set of data measurements $d = \{d_0, d_1, ..., d_s\}$, the parameters, $a = \{a_0, a_1, ..., a_s\}$ of the model $u(s;a)$ are found by minimizing the *residual error*

$$\varepsilon = \sum_{s \in S} \rho \left(d_s - u\left(s;a\right), \sigma \right) \tag{5}$$

where ρ is an *error norm* and σ a scale parameter. Depending on the desired robustness for outliers different functions for ρ can be chosen (see figure 2). For normally distributed noise the optimal ρ is the least squares estimate. The influence function ψ, which is the derivative of ρ is also depicted in figure 2. The influence function describes the effects of an measurement of value x on the estimate. The sensitivity of the least-squares method to outliers clearly follows from the influence function.

Besl et al. [2] used robust statistics for the design of filters which are robust to outliers. Outliers are produced by both noise other than normally distributed, and by pixels which belong to another set.

The zeroth order estimate, $u = i$, is found by minimizing

$$\varepsilon\left(x_0, i, \sigma_x, \sigma_i\right) = \int_x \rho\left(f\left(x\right) - i, \sigma_i\right) s\left(x - x_0, \sigma_i\right) dx \tag{6}$$

with respect to i. We added a function s to incorporate spatial dependence of data. If a Gaussian function for s is used the residual error can be rewritten as an integral over all the values of the sensor-space by using the LHS defined in section 2.

$$\varepsilon\left(x_0, i, \sigma_x, \sigma_i\right) = \int_j H(x_0, j, \sigma_x, 0) \rho\left(j - i, \sigma_i\right) dj, \tag{7}$$

which can be rewritten as a convolution of the local histogram with the ρ kernel.

$$\varepsilon\left(x_0, i, \sigma_x, \sigma_i\right) = H(x_0, i, \sigma_x, 0) \otimes_i \rho\left(i, \sigma_i\right). \tag{8}$$

The correspondence between this equation and eq. 4 is apparent. If the following function for ρ is chosen

$$\rho\left(x, \sigma\right) = 1 - e^{-\frac{x^2}{2\sigma^2}}, \tag{9}$$

the following holds

$$\varepsilon(x, i, \sigma_x, \sigma_i) = 1 - H(x, i, \sigma_x, \sigma_i). \tag{10}$$

The error norm of eq. 9 and its influence function are shown in figure 2. From the influence function the robustness to outliers is clearly visible in the fact that for large $|x|$ the influence function goes to zero.

Since the residual error is equal to 1 minus the local histogram, also the residual error of the zero order estimate of figure 1a as a function of σ_i is depicted in

Figure 3: a) color image (300x300) b) degraded with uncorrelated Gaussian noise of $\sigma = 20$ c) 1 iteration of tonal Gaussian with $\sigma_x = 3$ and $\sigma_i = 40$ d) 1 iteration of tonal Gaussian with $\sigma_x = 3$ and $\sigma_i = \infty$. This is equal to classical spatial Gaussian filter .

figure 1c (, this can be seen by holding the figure upside down). And hence *the global mode, which is the maximum of the smoothed histogram is equal to the zeroth order robust estimate, which is the minimum of the residual error.* The tonal scale σ_i at which the histogram is observed can also be interpreted as a scale of robustness. With increasing tonal scale the robustness decreases. Hence, *whereas the spatial scale has a statistical interpretation as the size of the measurement set, the tonal scale determines the robustness of the measurement to outliers.*

The local mode uses the value of the center pixel as a-priori knowledge to chose between the different minima of the residual error ε. The value of the center pixel is the starting point of iteration to the local minimum.

3 Local mode finding by spatial and tonal filtering

In this section the concept of tonal filtering is introduced. An intuitive introduction of the filter was given in [9]. Here we prove that iteratively applying this filter results in the local mode operation.

3.1 Tonal and spatial Gaussian filtering

Spatial Gaussian filtering of an image f leads to an image g:

$$g(x_0) = \int_{R^n} f(x)\, G(x - x_0, \sigma_x)\, dx \qquad (11)$$

Because the Gaussian kernel G is normalized there is no explicit need for normalization of the weighted average performed in the above Gaussian convolution. However, we can write:

$$g(x_0) = \frac{\int_{R^n} f(x)\, G(x - x_0, \sigma_x)\, dx}{\int_{R^n} G(x - x_0, \sigma_x)\, dx} \qquad (12)$$

to make the normalization explicit.

As mentioned before, a well known disadvantage when using the Gaussian convolution to eliminate additive noise, is that image structure is also smoothed.

Consider the example of an image showing a steep edge (figure 1a) . The Gaussian filter not only suppresses the noise but also smoothes the edge. The reason for this is clear, we are mixing the grey values from two constant regions in the image, leading to the (weighted) mean of both grey values. The weights being dependent on the area within the Gaussian aperture of the two areas.

We propose a straightforward generalization of the Gaussian filter where we not only use a spatial scale σ_x but also a tonal scale σ_i. The spatial scale σ_x is used to weigh the contribution of a value $f(x)$ depending on its distance from the central point x_0. The *tonal scale* σ_i weighs the contribution of $f(x)$ in the local averaging according to its *tonal distance*, $f(x) - f(x_0)$, from the tonal value in the central point;

$$g(x_0) = \frac{\int_{R^n} f(x)\, G(x - x_0, \sigma_x)\, G(|f(x) - f(x_0)|, \sigma_i)\, dx}{\int_{R^n} G(x - x_0, \sigma_x)\, G(|f(x) - f(x_0)|, \sigma_i)\, dx} \qquad (13)$$

This filter is easily implemented, and shows some remarkable properties in practical applications. In figure 3 examples are shown of tonal (and spatial) Gaussian filtering. We may observe that:

- Edges are preserved better then in classical (spatial) Gaussian filtering.

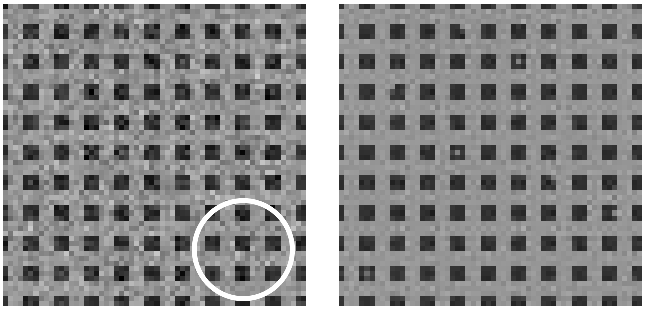

(a) (b)

Figure 4: a) grayscale signal, with grayvalue difference of 100 between valleys and ridges polluted with Gaussian Noise $\sigma = 15$. b) local mode filtered result with $\sigma_x = 5$ and $\sigma_i = 20$. The spatial filter size with $radius = 2\sigma_x$ is superimposed in a).

- For $\sigma_i \to \infty$ tonal Gaussian filtering is equal to spatial Gaussian filtering.
- Small details are preserved. The spatial scale is of little influence on this property.

3.2 Local mode finding

Here we show that iteratively applying the spatial and tonal Gaussian will lead to the local mode. For the local mode of the residual error ε with ρ as given in eq. 8 the following holds,

$$\frac{\partial \varepsilon}{\partial i} = \int_x G\left(x - x_0, \sigma_x\right) \frac{\partial}{\partial i}\left(1 - e^{-\frac{(f(x)-i)^2}{2\sigma_i^2}}\right) dx = 0 \quad (14)$$

This can be written as

$$i = \frac{\int_x f(x)\, G\left(x - x_0, \sigma_x\right) e^{-\frac{(f(x)-i)}{2\sigma_i^2}} dx}{\int_x G\left(x - x_0, \sigma_x\right) e^{-\frac{(f(x)-i)}{2\sigma_i^2}} dx} \quad (15)$$

which is of the form $i = F(i)$ and can be solved by functional iteration, i.e. $i_n = F\left(i_{n-1}\right)$. Hence f_n will approach the local mode for increasing n,

$$f_n\left(x_0, \sigma_x, \sigma_i\right) = \frac{\int f(x) G(x-x_0,\sigma_x) G(f_0(x)-f_{n-1}(x_0),\sigma_i) dx}{\int G(x-x_0,\sigma_x) G(f_0(x)-f_{n-1}(x_0),\sigma_i) dx} \quad (16)$$

This equation is an iterative version of the proposed tonal and spatial Gaussian in eq. 13.

Once more, applying eq.13 to the data and replacing the starting point of iteration with every step by the outcome of the previous step, results in the local mode.

3.3 Anisotropic Diffusion

In this subsection we shortly discuss the link between local mode filtering and anisotropic diffusion.

To prevent edges from blurring Perona and Malik [7] adapted the diffusion equation to

$$\frac{\partial f\left(x, y, t\right)}{\partial t} = div\left[g\left(\nabla I\right) \nabla I\right] \quad (17)$$

where t is the direction of iteration, g the edge-stopping function, which reduces diffusion across edges.

In [3] Black proves that iteratively minimizing the residual error ε is equal to anisotropic diffusion. The fact that the local minimum is desired and found, instead of the global minimum, is not mentioned.

In the previous subsection we showed that iterating towards the local mode is equal to applying the spatial and tonal Gaussian. Hence, diffusion can be written as follows,

$$f_n\left(x_0, \sigma_x, \sigma_i\right) = \frac{\int f(x) G(x-x_0,\sigma_x) G(f_{n-1}(x)-f_{n-1}(x_0),\sigma_i) dx}{\int G(x-x_0,\sigma_x) G(f_{n-1}(x)-f_{n-1}(x_0),\sigma_i) dx} \quad (18)$$

The difference with eq. 16 is that instead of keeping the neighborhood data constant, the neighborhood data is updated with every iteration step.

3.4 Vector valued images

Extension of mode filtering to color images is straightforward. For color images, $f = \{R, G, B\}$, the LHS is five dimensional space. The local histogram now consists of 3D Gaussians positioned at the R, G, B values of the neighborhood. The modes can again be found by iteratively applying eq. 13 with the Euclidean distance in sensor space as the tonal distance

$$|f(x) - f(x_0)| = \sqrt{(R(x)-R(x_0))^2+(G(x)-G(x_0))^2+(B(x)-B(x_0))^2} \quad (19)$$

4 Results

In figure 4 the local mode is applied to a synthetical gray-scale image degraded with Gaussian noise. The image is a 2D version of a block function for which local mode filtering is especially suitable. The size of the mode filter is superimposed on the image and is several periods of the block signal. The local histogram at this scale contains two modes. One for the ridges and one for the basins. Since local mode filtering uses the centerpixel of the neighborhood as a starting point of iteration almost everywhere the correct mode was found. Applying a classical Gaussian filter or a median filter to this pattern has disastrous effects.

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

Figure 5: a) color document image (150x150) b) corrupted with uncorrelated Gaussian noise $\sigma = 20$ c) local mode filter with $\sigma_x = 5$ and $\sigma_i = 16$ d) vector median filter (3x3)(see [1])

In figure 5b a color document image degraded by Gaussian noise is depicted. The result of noise attenuation without blurring the details is depicted in figure 5c. At the applied tonal scale the characters are far enough separated in sensorspace from the background to initiate a local mode. Consequently they are preserved. Increasing the tonal scale will eventually induce a blurring effect. A vector median filter of size 3×3 already blurs the characters, see figure 5d.

Note that a consequence of the preservation of detail is that also other sorts of noise, like e.g. speckle noise, are preserved.

5 Concluding Remarks

In this paper we have connected a framework of histogram based computations to the theory of robust statistics. Both are used as a theoretical foundation for local mode filtering. This filter reduces normally distributed noise without blurring edges and details and is easily extended to multi-channel images.

For the implementation of the local mode a generalized version of the Gaussian filter is proposed. Next to the spatial distance also the distance in sensor space is used to weight the values in the neighborhood. This filter preserves edges and details and is easily extensible to multi-channel data. If used iteratively on an image it is shown to approach the local mode.

Future research will focus on extending the framework to higher order local image structure. "Clever choices" to iterate to a local minimum in the residual error for these models are still subject of research.

Application of a Gaussian filter in the sensor domain was motivated from the histogram space point of view. From the robust statistics field several other functions are proposed, with different robustness to outliers [5]. Adaptation of the theory to incorporate these functions is expected diserable for some applications.

References

[1] J. Astola, P. Haavisto, and Y. Neuvo. Vector median filters. *IEEE Proceedings*, 78(4):678–689, April 1990.

[2] P.J. Besl, J.B. Birch, and L.T. Watson. Robust window operators. *Machine Vision and Applications*, 2:179–191, 1989.

[3] M. Black, G. Sapiro, D. Marimont, and D. Heeger. Robust anisotropic diffusion. *IEEE Trans. Image Processing*, 7(3):421–432, March 1998.

[4] L.D. Griffin. Scale-imprecision space. *Image and Vision Computing*, 15:369–398, 1997.

[5] P.J. Huber. *Robust Statistics*. "probability and mathematical statistics". Wiley, 1981.

[6] J.J. Koenderink and A.J. van Doorn. The structure of locally orderless images. *International Journal of Computer Vision*, 31(2/3):159–168, 1999.

[7] P.Perona and J. Malik. Scale space and edge detection using anisotropic diffusion. *IEEE trans. on pattern analysis and machine intelligence*, 12(7):629–639, 1990.

[8] G. Sapiro and D. Ringach. Anisotropic diffusion of multivalued images with applications to color filtering. *IEEE Trans. Image Processing*, 5(11):1582–1586, Oct 1996.

[9] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. of the Sixth International Conference on Computer Vision*, Bombay, India, January 1998.

[10] J. van de Weijer and Th. Gevers. Color mode filtering. In *Proc. Int. Conference on Image Processing*, Thessaloniki, Greece, Oct. 2001, accepted.

[11] B. van Ginneken and B.M. ter Haar Romeny. Applications of locally orderless images. In *Proceedings Scale-Space 99*, pages 10–22, 1999.