A Network Approach to
Reconstructions from Zero-Crossings

by

Robert Hummel [1]
Courant Institute of Mathematical Science

Robert Moniot [2]
Fordham University

---

[1] New York University
251 Mercer Street
New York, NY 10012
hummel@nyu.nyu.edu

[2] Division of Science and Mathematics
Fordham University
College at Lincoln Center
New York, NY 10023
moniot@acf8.nyu.edu

A Network Approach to
Reconstructions from Zero-crossings

Robert Hummel
Robert Moniot

Abstract

The zero-crossings of Laplacian-of-Gaussian filtered image data have been used for edge detection and segmentation of digital images. Under certain conditions, the set of all zero-crossings represented at all scales of resolution completely determines the image data, up to a multiplicative constant. Reconstruction methods are desired in order to study the applicability of these theories when the conditions do not hold, and to study the stability of the representation. We apply a method of solving ill-conditioned inverse problems, called "minimization of equation error," to this problem, in order to produce a signal with zero-crossings at all scales that are nearly equivalent to zero-crossings of natural data. Our results suggest that representation using zero-crossing data alone is quite unstable.

1. Zero-crossings

The zero-crossings of the Laplacian of a Gaussian-filtered image form closed contours that have been suggested as a representation of the edges of a digital image [1]. Even though some of the zero-crossings can occur at locations that do not correspond to physical edges in the image, the existence of a zero-crossing at some location, for a specific Gaussian, can be used as a feature signifying information about the image. Thus zero-crossings have been used for motion analysis, object extraction, texture analysis, and many other image processing tasks.

By definition, a zero-crossing is the border between a region where the data is positive and a region where the data is negative. If $f(x,y)$ denotes a gray level image defined on the two-dimensional domain $\mathbb{R}^2$, then the zero-crossings of interest are those pertaining to the data $\Delta G * f$, where G is a Gaussian,

$$G(x,y) = \frac{1}{2\pi\sigma^2} \cdot e^{-(x^2+y^2)/2\sigma^2}.$$

By varying σ, we obtain zero-crossings at different *scales*, and it is suggested that combining information about zero-crossings at multiple scales can lead to better identification of physical edges [2]. If we think of σ as a continuous parameter, then the resulting domain $(x,y,\sigma) \in \mathbb{R}^2 \times \mathbb{R}^+$ is called scale-space as introduced by Witkin [3]. We will consider zero-crossings in this domain. In fact, we will reparameterize scale-space, for convenience, using the variable $t = (1/2)\sigma^2$. Throughout this paper, "zero-crossings" will refer to the zero-crossings of the

filtered data. The filtered data will be denoted by the function $v(x,y,t)$, defined by

$$v(x,y,t) = \int_{\mathbb{R}^2} K(x-x',y-y',t)f(x',y')dx'dy',$$

where

$$K(x,y,t) = \frac{1}{4\pi t}e^{-\frac{x^2+y^2}{4t}}.$$

We use here the function K to represent a parameterized set of Gaussians that approaches the delta function as $t \to 0$, and diffuses as t increases. For any given fixed t_0, the zero-crossings of $v(x,y,t_0)$ form the contours as defined by the Marr-Hildreth zero-crossing edge detector for the scale corresponding to t_0. As t varies, the zero-crossings of $v(x,y,t)$ form surfaces that evolve and disappear as t increases.

The function $v(x,y,t)$ and its zero-crossings enjoy many special properties. For example, v is a solution to the Heat Equation:

$$\Delta v = \frac{\partial v}{\partial t}.$$

If the initial image data $f(x,y)$ is sufficiently smooth, v satisfies the initial conditions:

$$v(x,y,0) = \Delta f(x,y).$$

Using the fact that v satisfies the Heat Equation, one can prove that the zero-crossings satisfy the "evolution property" [4], which states roughly that zero-crossing surfaces cannot be created at intermediate values of t, but instead evolve and can only disappear as t increases. This property, noted by Witkin [3], is potentially useful if the image data is to be represented by a simplified description of the zero-crossings, as opposed to analytical or bit-map sampled representations of the locations of the zero-crossings. Using this property, zero-crossing surfaces form tubes that emanate from the bottom plane $\{t=0\}$, and might be described in terms of a containment tree of surfaces, approximate shapes, and locations of top points.

It is also possible to define a scale-space formulation for images $f(x,y)$ defined on a bounded domain: $(x,y) \in \mathcal{D} \subseteq \mathbb{R}^2$. This can be done by solving a boundary-value Heat Equation problem, where either Dirichlet or Neumann conditions are applied on the boundary of the resulting scale-space cylinder. More details are given in [5], and we will make use of such a formulation for our experiments later in this paper.

The information of the zero-crossings of $v(x,y,t)$ is equivalent to complete knowledge of the sign of v, i.e., $\text{sgn}(v(x,y,t))$, except for a possible global change in sign. Here, the signum function is defined by:

$$\text{sgn}(x) = \begin{cases} -1 \text{ if } x<0 \\ 0 \text{ if } x=0 \\ +1 \text{ if } x>0 \end{cases}$$

We do not specify how the zero-crossing information should be represented, and will assume that the information can be obtained in either form. In our

reconstruction method, we will make use of the data $\text{sgn}(v(x,y,t))$.

From a given image $f(x,y)$, we can thus easily form the data $v(x,y,t)$, and represent the zero-crossings of v. The zero-crossing descriptions might be simplified by making use of the evolution property. An alternate form of simplification arises if the locations of the zero-crossings are represented by a bit-map at a quantized set of levels, and if the quantization resolution of the bit-maps decrease as the scale parameter t increases. In this case, a scale-space pyramid is formed, and the locations of the zero-crossings can be recorded within the pyramid. The rationale for subsampling as the scale increases is that the data $v(x,y,t)$ becomes smoother as t increases, and the corresponding structure of the zero-crossings also tends to become smoother. A sampled version of the $v(x,y,t)$ data, with decreasing resolution as t increases, is closely related to the Laplacian pyramid data structure [6].

In this paper, we are concerned with the amount and quality of information that is retained in the zero-crossing representation. We will thus assume that little simplification of the zero-crossings of $v(x,y,t)$ data takes place. Instead, we will assume that the data is represented at the same resolution at all levels of t. Typically, the image data $f(x,y)$ is sampled on a grid of points, and similarly $v(x,y,t)$ is represented on a three-dimensional grid of points. We can thus speak of values $f(i,j)$ and data $v(i,j,k)$, where the variables take on integer values (and the scale-space parameter k is a nonnegative integer. We propose to call the data $v(i,j,k)$ the "Laplacian monolith," which seems especially appropriate when the domain of $f(i,j)$ is a rectangular array of points, as is usually the case. The zero-crossings can then be recorded either as border pixels in (i,j,k) space where v is positive, or by simply recording the $\text{sgn}(v(i,j,k))$ values at each pixel.

2. Completeness

The collection of all zero-crossings, at all scales of resolution, can, under certain restrictions, determine the initial data. Thus given the zero-crossings in scale-space of the Laplacian-of-Gaussian convolved against an initial function $f(x,y)$, one should be able in principle to reconstruct $f(x,y)$ up to a multiplicative constant, and an additive harmonic function, providing f satisfies certain conditions. The ambiguity arises because $cf+h$ will clearly have exactly the same zero-crossings as f, where c is a scalar multiplicative constant, and h is a harmonic function ($\Delta h = 0$). However, if we assume that f approaches zero at locations distant in its domain, then the ambiguity with respect to h disappears. The multiplicative ambiguity is considered unimportant.

The conditions required for completeness can be quite technical. For example, if f is a polynomial in x and y, then the zero-crossings are the real analytic varieties of the polynomial v, and it can be shown that in most cases (for example, if v is irreducible) the real analytic varieties determine the polynomial. In fact, the analytic structure of the zero-crossing surfaces is required at just two points, under most circumstances [7]. Alternatively, if f is band limited, then again in most cases f is determined by the zero-crossings of v; see [8]. A celebrated result by Logan [9] asserts that a one octave bandpass one-dimensional signal is reconstructible from its zero-crossings provided that the signal and its Hilbert transform share no common simple zeros. Although the slices of the Laplacian pyramid and Laplacian monolith do not provide precisely bandpass signals, and despite the problems in extending the

result to higher dimensions, there have nonetheless been alternative formulations of the zero-crossing data structure that have been suggested that would ensure completeness, based on this theorem, provided the Hilbert transform condition holds on a sufficient number of scanlines. In particular, Zeevi and Rotem [10] use a bandpass decomposition of an image by tiling the digital Fourier domain with rectangles of sufficiently small size, noting that a digital image that is bandpass in two dimensions has rows that are bandpass in one dimension.

The theoretical results that indicate completeness of the representation by zero-crossings are typically non-constructive. Thus, reconstruction methods are based on methods distinct from the proofs of completeness. Both the Zeevi work [10] and the Sanz reconstruction experiments [8] are based on iterative algorithms. Convergence is not guaranteed in either method, and one can at best assert that if convergence happens, then the solution will have the correct zero-crossings. Based on the theoretical results, one can then expect that if the technical conditions hold for the original imagery, then the reconstruction will be the same as the original image. Zeevi and Rotem report good results, using the discrete bandpass filtered images, and Sanz and Huang report disappointing results using zero-crossings of Laplacian-of-Gaussian filtered images. Two difficulties can lead to poor reconstructions (with any iterative method). First, it can easily happen that the iterations fail to converge. Second, instability can lead to reconstructions that are significantly different from the original image, even though the representation has been proven to be complete.

Let us amplify the instability issue. Reconstruction from zero-crossings is an example of an ill-conditioned inverse problem. A similar situation arises in deblurring images: representation of image data by a blurred version of the image is complete, and reconstruction from blurred data is theoretically possible [11]. However, in practice, there are many images that when blurred give essentially the same data as the given blurred image, so that the problem becomes one of choosing the best deblurring. Similarly, we expect there to be many images that have very nearly the same zero-crossings, at all levels, as a given image. A reconstruction method must be able to deal with inaccuracies in the representation of the zero-crossings, and thus needs to choose the appropriate reconstructed image based on assumptions about natural images.

Why is reconstruction of any interest? If the information in the zero-crossings is complete, then all further image analysis can take place using the zero-crossing information, and reconstruction becomes irrelevant. Our interest in studying methods of reconstruction is to study the stability issue mentioned above. At the worst, we will find instances of signals that have nearly the same zero-crossings as a given signal, and thus be able to assess the ambiguity in the representation by zero-crossings. If the two signals are significantly (i.e., perceptually) different, then this puts into doubt the viability of the zero-crossings as a method of representation, in the absence of other information. Of course, there may exist different methods of reconstruction. Some reconstruction methods may implicitly use assumptions about natural images in order to yield better results. In this paper, we concentrate on a method that we call "minimizing equation error." We have reason to believe that if stable reconstruction of natural images from zero-crossings is possible, then the method of minimizing equation error should work well.

3. A network approach

Our network approach to reconstructions from zero-crossings is based on the method of minimizing equation error. We previously reported this method, and demonstrated good results in applications involving image deblurring [12]. The method reconstructs the function v in scale-space, and makes use of the fact that v satisfies the Heat Equation. The data $v_{i,j,k}$ can be viewed as a multilevel grid of units communicating locally. The essence of the idea is that the units should achieve values $v_{i,j,k}$ satisfying the given zero-crossing constraints, and satisfying, to the extent possible, a discrete version of the Heat Equation. As in a network, the values are updated iteratively using information from local values to minimize a measure of error.

We first sketch the formulation of the approach in a continuous domain. Consider the initial image $f_0(x,y)$ and its scale-space Laplacian-of-Gaussian filtered data $v_0(x,y,t)$. We assume that the zero-crossings in scale space of v_0 are given, so we assume knowledge of $s_0(x,y,t) = \text{sgn}(v_0(x,y,t))$. We then pose the problem:

$$\text{Find } v \text{ minimizing } \|\Delta v - \frac{\partial v}{\partial t}\|^2,$$

$$\text{subject to } \text{sgn}(v(x,y,t)) = s_0(x,y,t).$$

Here, the norm $\|\cdot\|$ is the L^2 norm over the scale-space $\mathbb{R}^2 \times \mathbb{R}^+$. This is a different formulation from minimizing the data error, which would amount to minimizing a measure of the error in the locations of the zero-crossings.

In a discrete formulation, we assume the existence of variables at all nodes, $v_{i,j,k}$, $k \geq 0$, and formulate the Heat Equation as:

$$v_{i,j,k+1} = \frac{1}{16}[v_{i-1,j-1,k} + 2v_{i-1,j,k} + v_{i-1,j+1,k}$$

$$+ 2v_{i,j-1,k} + 4v_{i,j,k} + 2v_{i,j+1,k}$$

$$+ v_{i+1,j-1,k} + 2v_{i+1,j,k} + v_{i+1,j+1,k}].$$

We can then formulate equation error as a sum of square difference of the left and right sides of this discretized Heat Equation over all nodes, and seek to minimize the resulting quadratic functional.

In practice, things are complicated by the need to define boundary conditions, choose an appropriate constrained minimization procedure, and to ensure numerical stability and reasonable convergence rates.

We will briefly discuss some of these concerns in the case of one space dimension. In this case, we are given initial data $f(i)$, and construct the scale-space function $v_{i,k}$. This function is computed recursively from the formulas:

$$u_{i,0} = f(i),$$

$$u_{i,k+1} = \frac{1}{4}[u_{i-1,k} + 2u_{i,k} + u_{i+1,k}].$$

$$v_{i,k} = u_{i,k+1} - u_{i,k}.$$

For the border points, let us assume that $f(i)$ is defined for $-N \leq i \leq N$. Then we

replace in the above formulas, for $i = \pm N$, the formulas:

$$u_{N,k+1} = \frac{1}{4}u_{N-1,k} + \frac{3}{4}u_{N,k},$$

$$u_{-N,k+1} = \frac{1}{4}u_{-N+1,k} + \frac{3}{4}u_{-N,k},$$

and otherwise things are the same.

For our numerical experiments (Section 4), we in fact reformulate the Heat Equation as a first order differential system:

$$\nabla v = \bar{\sigma},$$

$$\nabla \cdot \bar{\sigma} = \frac{\partial v}{\partial t}.$$

When discretized, we then use the the set of variables $v_{i,k}$ and $\sigma_{i,k}$ (for the one-dimensional space variable case), and formulate the equation error using the quadratic functional:

$$E = \sum_{k \geq 0} \sum_{i} \left(\frac{v_{i,k} - v_{i-1,k}}{2} - \sigma_{i,k} \right)^2$$

$$+ \left(v_{i,k+1} - v_{i,k} - \frac{\sigma_{i+1,k} - \sigma_{i,k}}{2} \right)^2.$$

We wish to minimize E using the constraints

$$v_{i,k} \geq 0 \text{ when } s_0(i,k) = 1$$

$$v_{i,k} \leq 0 \text{ when } s_0(i,k) = -1.$$

Here s_0 is the given sign data which encodes the information about the zero-crossings. A further constraint is needed to ensure that the solution $v=0$ is not obtained. This can be done, for example, by specifying a single value for v at some node. We chose instead to specify v at the maximum level $k=T$, where the specified data is ordinarily quite smooth. One way to impose the inequality constraints is by employing an iterative algorithm, and projecting the values of $v_{i,k}$ to zero after each iteration, whenever a constraint is violated. We instead used a penalization method, adding to E a term that grows quadratically in the variable whenever a constraint is violated. By placing a large weight on the penalization term, we obtain a functional suitable for minimizing that will ensure, upon convergence, that the constraints generally hold. Specifically, we use a penalization term of the form:

$$\lambda \sum \left(|\text{sgn}(v) - s_0| \cdot v^2 \right).$$

Other details of the finite difference formulation and an alternate finite element formulation may be found in [12]. Once the functional is specified, any of many different optimization procedures may be used to perform the minimization. A simple method would be to employ gradient descent. This has the advantage of requiring only local interactions between the elements. We in fact used a conjugate gradient approach, which improves computation speed, and requires only very simple global computations.

We could imagine something similar to a neural network to implement the minimization procedure. For a gradient descent procedure, the computation is especially simple. We discretize scale-space with a network of nodes, with each node associated with some position and scale parameter. Each node in the network stores a current value, or activation level, representing the scale-space value for the filtered data at that location. Each node attempts to adjust its value so that the values in the level below average to give the value at the given scale, as prescribed by the Heat Equation, which the filtered data is known to satisfy. This adjustment is mediated by local connections. Specifically, the gradient of the quadratic equation error is a function yielding a value at each node, and can be computed at each site by a local linear sum, involving values at nodes in neighboring positions and neighboring scales. By adjusting a node's value by subtracting a portion of the component of the gradient at that site, the total equation error will be reduced. The zero-crossing information, or signum of v_0, imposes constraints on the nodes which are not violated. These constraints can be imposed by pegging the values at nodes to lie in some range. Perhaps, for example, values of nodes lying along the zero-crossings are clamped to zero. Alternatively, some nodes might be enforced to remain positive, and other nodes forced negative. In any case, the network achieves some relaxation point based on these conditions, using the minimization process, thereby constructing a reconstruction of the scale-space function v.

The network viewpoint suggests that an analog implementation of this computation would be feasible. Such an implementation could potentially be stable and much faster than a digital implementation.

How can the original function $f(x,y)$ be obtained from the reconstructed function v? One method would be to solve the Poisson equation:

$$\Delta f(x,y) = v(x,y,0).$$

The lack of boundary conditions reflects the additive harmonic function ambiguity. We can alternatively make use of the fact that

$$f(x,y) = -\int_0^\infty v(x,y,t)dt.$$

In a discrete setting, this means that the layers of the Laplacian monolith are summed. In practice, there are only a finite number of levels that have been computed in the equation error minimization procedure. Thus the data that can be reconstructed by summing these levels is

$$r_{i,j} = -\sum_{k=0}^{T} v_{i,j,k}.$$

Then the $r_{i,j}$ data can be seen to equal $f - G_{T+1}*f$, where f is the original image, and G_{T+1} is a Gaussian obtained by $T+1$ levels of digital blurring, as described above. One way to recover the f data from the r data is to form the sum

$$\sum_{k=0}^{\infty} G_{k(T+1)}*r + \mu$$

where μ is the mean value of the f data. The sum converges fairly rapidly. This was the procedure used to reconstruct an image f.

4. Results

We present in this section results of using the network approach to reconstructions from zero-crossings. We first conducted experiments using one-dimensional data obtained from a scanline of a digitized image. Figure 1 shows a plot of this data. We construct the scale-space filtered function $v_0(i,k)$ as described in the previous section, and compute the sign of the result. We used 41 layers, so that $0 \leq k \leq 40$. In Figure 2, we display the sign of the function v_0, using a dark color for negative regions, white for positive regions, and a gray color for pixels whose values lie near zero. This is the zero-crossing information that forms the representation of the data in Figure 1.

Reconstruction proceeds by minimizing equation error, using the sign of v_0 as the data shown in Figure 2. The initial estimate for v is the same as the given sign data, i.e., one in the positive regions, negative one in the negative regions, and zero within the "gray" regions. A penalization method was employed to coerce the inequality constraints. A conjugate gradient approach was used, instead of gradient descent, so as to speed computation. The computations were done on a Convex computer. Convergence was reached with a total equation error plus penalty of less than 10^{-10}. At this point, 77 points in the scale-space, out of 5120 points, violated the sign constraint, and the magnitudes of the violations were quite small. Using the data obtained in this manner, a reconstructed function is obtained from the sum of levels as described in the previous section. In Figure 3, we show this reconstructed function, and Figure 4 shows its sign of Laplacian-of-Gaussian filtered data. Clearly, the zero-crossing information is nearly identical, but the signals have some differences. On the other hand, the reconstructed function is qualitatively similar to
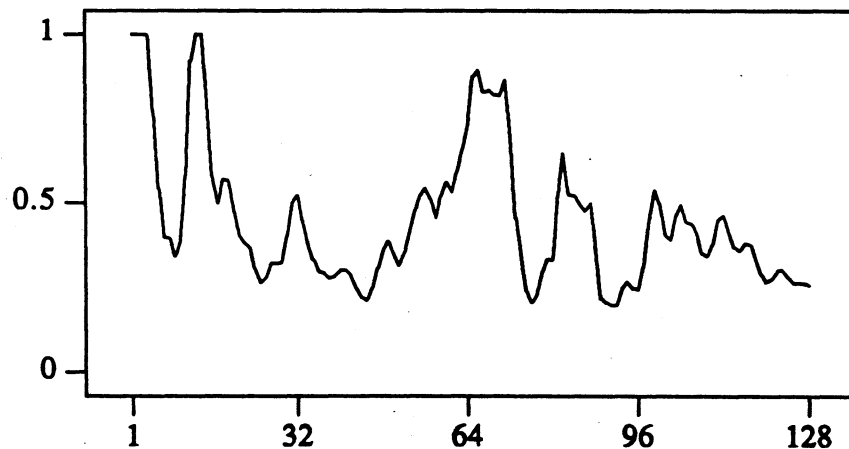

Figure 1. A scanline taken from a digitized natural image. There are 128 pixels in the scanline.
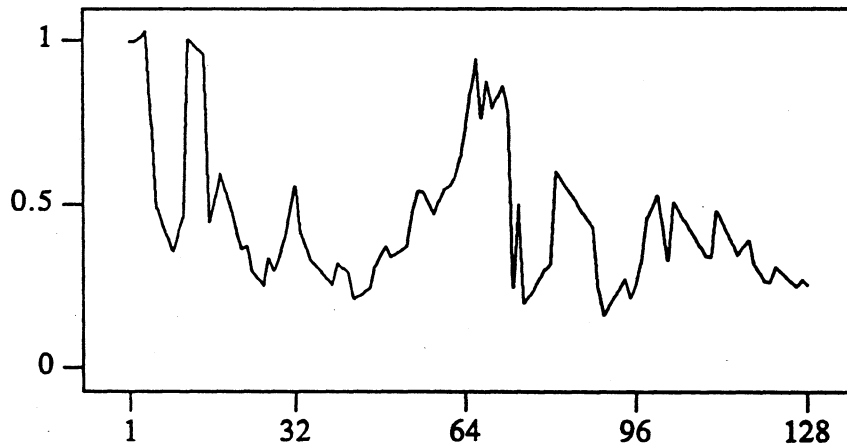
Figure..2

Figure 2. The sign of the scale-space filtered data shown in Figure 1. Positive regions are shown in white, negative in black, and values very near zero are in gray. The zero-crossing occur on the borders between black and white regions and within gray areas.

the original, so that the zero-crossing information has captured some of the essential information about the signal.

We next applied the reconstruction procedure to digitized two-dimensional images. Figure 5 shows an initial 64 by 64 image. The filtered data was computed, and the signum of the result was used as input to the reconstruction procedure.

Figure 3. A reconstruction f of a signal from the reconstructed Laplacian-of-Gaussian filtered data v. This function should be the same as Figure 1 if stable reconstruction from zero-crossings were possible.

Figure..4

Figure 4. The sign of the scale-space filtered data shown in Figure 3. The data is nearly identical to that of Figure 2, showing that the signals in Figures 1 and 3 have virtually the same zero-crossings for their Laplacian-of-Gaussian filtered data at all scales of resolution.

Twenty levels of scale were used, and the signum of levels 0, 3, 6, 9, 12, 15, and 18 are shown in the leftmost column of Figure 7. (The finest resolution level, i.e. smallest scale, is at the top.) The reconstructed v was produced, and the resulting monolith has the same sign as the original in nearly all of the points in the three-dimensional grid. Using the v function, the column sum r was computed, and a reconstructed f produced according to the method given in Section 3. This reconstructed image is shown in Figure 6. Finally, the Laplacian monolith data of the reconstructed f was produced. The sign of the result is shown in the middle column of Figure 7. The rightmost column of Figure 7 shows the points where there are discrepancies between the sign of the Laplacian monoliths of the original and reconstructed images. We see that the two images have nearly the same zero-crossings except at level 0, and yet the images look rather different. Some of the inaccuracies in the reconstruction reflect the high-frequency distortions observed in the zero-crossing errors. One could argue that since the reconstruction does not look like a natural image, there might be a better reconstruction method that obtains a closer approximation to the original image, by making assumptions about images. However, we believe that the equation error minimization procedure is a good reconstruction method, based on its excellent performance with deblurring. In any case, the existence of the two distinct images with nearly identical zero-crossings at all levels suggests a serious lack of stability with the representation.

We have also performed some experiments in which, in addition to specifying the sign of v, the gradient of v was specified along the zero-crossings. In these experiments, convergence was very rapid, and the reconstruction of the image was virtually flawless. This supports the notion that the zero-crossings together with gradient data along the zero-crossings forms a complete representation and permits stable reconstruction. However, it should be noted that as implemented, the amount of information in the latter representation exceeds the amount of data in the original image.

Figure..5

Figure 5. Original 64 by 64 digital image.

Figure..6

Figure 6. The reconstructed image, obtained from the reconstructed Laplacian monolith data, as produced by minimizing equation error using the sign of Laplacian monolith data of the original image as constraints. As shown in the next figure, this image has nearly identical zero-crossings, at all levels, as the image in Figure 5.

Figure..7a   Figure..7b   Figure..7c

Figure 7. Levels 0, 3, 6, 9, 12, 15, and 18 of the sign of the Laplacian monolith of (left column) the original image, and (middle column) the reconstructed image. The rightmost column shows the locations of points where the two signs differ.

Acknowledgements

References

[1] Marr, D. and E. Hildreth, "Theory of edge detection," *Proceedings Royal Society London (B)*, p. 187 (1980).

[2] Marr, D., *Vision*, W. H. Freeman and Company (1982).

[3] Witkin, A., "Scale space filtering," *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, p. 1019 (1983).

[4] Hummel, Robert, "Representations based on zero-crossings in scale-space," *Proceedings of the IEEE Computer Vision and Pattern Recognition Conference*, pp. 204-209 (June, 1986).

[5] Hummel, Robert, "The scale-space formulation of pyramid data structures," pp. 107-123 in *Parallel Computer Vision*, ed. Len Uhr, Academic Press, New York (1987).

[6] Burt, P. and T. Adelson, "The laplacian pyramid as a compact image code," *IEEE Trans. on Communications*, p. 532 (1983).

[7] Yuille, A. L. and T. Poggio, "Fingerprints theorems for zero crossings," *J. Optical Society of America* 2, pp. 683-692 (1985).

[8] Sanz, Jorge and Thomas Huang, "Theorem and Experiments on Image Reconstruction from Zero Crossings," Research Report RJ5460, IBM Almaden Research Center (January, 1987).

[9] Logan, B., "Information in the zero-crossings of band pass signals," *Bell Systems Technical Journal* 56, p. 510 (1977).

[10] Zeevi, Y. Y. and D. Rotem, "Image reconstruction from zero crossings," *IEEE ASSP* 34, pp. 1269-1277 (1986).

[11] Hummel, Robert A., B. Kimia, and S. Zucker, "Deblurring gaussian blur," *Computer Vision, Graphics, and Image Processing* 38, pp. 66-80 (1987).

[12] Hummel, Robert and Robert Moniot, "Solving ill-conditioned problems by minimizing equation error," *Proceedings of the First International Conference on Computer Vision*, (1987).