

Forward-and-Backward Diffusion Processes for Adaptive Image Enhancement and Denoising

Guy Gilboa, Nir Sochen, and Yehoshua Y. Zeevi

Abstract—Signal and image enhancement is considered in the context of a new type of diffusion process that simultaneously enhances, sharpens, and denoises images. The nonlinear diffusion coefficient is locally adjusted according to image features such as edges, textures, and moments. As such, it can switch the diffusion process from a forward to a backward (inverse) mode according to a given set of criteria. This results in a forward-and-backward (FAB) adaptive diffusion process that enhances features while locally denoising smoother segments of the signal or image. The proposed method, using the FAB process, is applied in a super-resolution scheme.

The FAB method is further generalized for color processing via the Beltrami flow, by adaptively modifying the structure tensor that controls the nonlinear diffusion process. The proposed structure tensor is neither positive definite nor negative, and switches between these states according to image features. This results in a forward-and-backward diffusion flow where different regions of the image are either forward or backward diffused according to the local geometry within a neighborhood.

Index Terms—Adaptive denoising, anisotropic diffusion, Beltrami flow, color processing, image enhancement, inverse diffusion, scale-space.

I. INTRODUCTION

IMAGE denoising, enhancement, and sharpening are important operations in the general fields of image processing and computer vision. The success of many applications, such as robotics, medical imaging, and quality control depends in many cases on the results of these operations. Since images cannot be described as stationary processes, it is useful to consider locally adaptive filters. These filters are efficiently modeled as solutions of partial differential equations (PDE).

The scale-space approach and other PDE techniques have been extensively applied over the last decade in signal and image processing. As Witkin [39] pointed out, the diffusion process (or heat equation), widely used in this context, is equivalent to a smoothing process with a Gaussian kernel. A major drawback of the linear scale-space framework is its uniform filtering of local signal features and noise. This problem was addressed by

Manuscript received May 29, 2001; revised March 24, 2002. This work was supported by the Ollendorf Minerva Center, the Fund for the Promotion of Research at the Technion, the Israeli Ministry of Science, and the Technion V.P.R. Fund. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Nasser Kehtarnavaz.

G. Gilboa and Y. Y. Zeevi are with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel (e-mail: gilboa@tx.technion.ac.il; zeevi@ee.technion.ac.il).

N. Sochen is with the Department of Applied Mathematics, University of Tel-Aviv, Tel-Aviv 69978, Israel (e-mail: sochen@math.tau.ac.il).

Publisher Item Identifier 10.1109/TIP.2002.800883.

Perona and Malik (P-M) [20], who proposed a nonlinear diffusion process, where diffusion can take place with a variable diffusion in order to control the smoothing effect.

The diffusion coefficient in the P-M process was chosen to be a decreasing function of the gradient of the signal. This operation selectively lowpass filters regions that do not contain large gradients (singularities such as step jumps or edges and thin lines in the case of images). Results obtained with the P-M process paved the way for a variety of PDE-based methods that were applied to various problems in low-level vision (see [31] and references cited therein).

Some drawbacks and limitations of the original model have been mentioned in the literature [2], [17], [38]. Catta *et al.* [2] proved the ill-posedness of the diffusion equation, imposed by using the P-M diffusion coefficient, and proposed a regularized version wherein the coefficient is a function of a smoothed gradient. Weickert *et al.* [37] investigated the stability of the P-M equation by spatial discretization, and proposed [23] a generalized regularization formula in the continuous domain.

The aim of this study is to further extend the nonlinear PDE-based filtering methods, and to apply them in signal and image enhancement and sharpening. We focus on enhancing and sharpening blurry signals, while still allowing some additive noise to interfere with the process. We minimize the effect of amplification of noise—an inherent byproduct of signal sharpening—by combining backward and forward diffusion processes. We then generalize the analysis of [14] by the introduction of a generalized local adaptive criterion for the forward-and-backward diffusion in sharpening and denoising of color images.

II. FORWARD-AND-BACKWARD DIFFUSION PROCESSES

A. Introduction: Enhancement by Diffusion

Most of the PDE-based studies have been devoted to denoising of images, attempting to preserve the edges. Both forward linear and nonlinear diffusion processes converge (as $t \rightarrow \infty$) to a trivial constant solution (i.e., the mean value of the signal, assuming Neumann boundary conditions). To preserve singularities, previous studies relied primarily on slower diffusion in the vicinity of singularities. Such is the P-M nonlinear diffusion equation of the form

$$I_t = \nabla \cdot (c(|\nabla I|)\nabla I) \quad (1)$$

where $c(\cdot) > 0$ is a decreasing function of the gradient, $I(t = 0) = I_0$, Neumann boundary conditions.

According to the “Minimum–Maximum” principle, no new local minima or maxima should be created at any time in the

one-dimensional (1-D) case, in order not to produce new artifacts in the diffused signal. Moreover, the value of the global minimum (maximum) along the evolution of the signal in time is bounded by that of the initial data in any dimension, and is a nondecreasing (nonincreasing) function. These conditions were obeyed by the P-M and most other nonlinear diffusion processes that were subsequently introduced in image processing. This guaranteed the stability of the PDE, avoiding explosion of the nonlinear diffusion process.

In signal enhancement, sharpening, and restoration, we do not want to restrict ourselves to the global minimum and maximum of the initial signal. On the contrary, we would like the points of extrema to be emphasized and “stretched” (if they indeed represent singularities and are not generated by the noise). Therefore, a different approach should be adopted. If we consider the gray-level value at a pixel to be analogous to the amount of particles, each having one unit of “mass,” stacked at the pixel, then in order to emphasize large gradients, we would like to move mass from the lower part of a “slope” upwards. This process can be viewed as moving back in time along the scale space, or reversing the diffusion process and applying it backward. Mathematically, this can be simply accomplished by changing the sign of the diffusion coefficient

$$I_t = \nabla \cdot J = \nabla \cdot (-c(\cdot)\nabla I), \quad c(\cdot) > 0 \quad (2)$$

where J is the flux.

Note that this is different from what was defined as “inverse diffusion” in previous studies (e.g., [2] and [21]). There, near a point x_0 where the derivative of the flux was negative, the process was defined as inverse diffusion, because the diffusion equation could be written in that neighborhood as

$$I_t(x) = -dI(x)_{xx}, \quad d > 0, \quad x \in [x_0 - \Delta x, x_0 + \Delta x]. \quad (3)$$

We refer to such processes as having local *implicit* inverse diffusion properties. Although it has the form of an inverse diffusion process, it is weaker since it does not have the important inverse diffusion property of moving signal or image “particles” upward along the slope of the gradient, at points of extrema. With positive diffusion coefficient c , this could never happen, and therefore, the minimum–maximum principle is preserved, for instance. Thus, signal sharpening requires further modification of the diffusion process. Specifically, to deblur and enhance singularities, *explicit* inverse diffusion with negative diffusion coefficient must be incorporated into the process.

The question is, can one simply use a linear inverse diffusion? Clearly, the linear inverse diffusion is a highly unstable process. As was mentioned earlier, the linear forward diffusion is analogous to convolution with a Gaussian kernel. Hence, the linear backward (inverse) diffusion is analogous to a Gaussian deconvolution, where the noise amplification explodes with frequency. Application of such a deconvolution process results in oscillations (ringings) that grow with time until they reach the limiting minimum and maximum saturation values and the original signal is completely lost (see Figs. 1 and 2 for 1-D and two-dimensional (2-D) examples of these phenomena).

Three major problems associated with the linear backward diffusion process must be addressed: the explosive instability,




Fig. 1. Linear inverse diffusion in 1-D.



Fig. 2. Two-dimensional linear inverse diffusion. From top left (clockwise): original image, blurred, inversely diffused at times 0.5, and 1, respectively.

noise amplification, and oscillations. To avoid the effect of an explosive instability, one can diminish the value of the inverse diffusion coefficient at high gradients. In this way, after the singularity exceeds a certain gradient it does not continue to affect the process any longer. One can also terminate the diffusion process after a limited time, before reaching saturation. In order not to amplify noise, which after some presmoothing, can be regarded as having mainly medium to low gradients, it is desirable to diminish the inverse diffusion process at low gradients.

To minimize the effect of oscillations, they should be suppressed the minute they are introduced. For this reason, we combine a forward diffusion force that smoothes low gradients. This force also smoothes some of the original noise that contaminates the signal from the beginning. Unfortunately, low gradients which are not due to noise, like those that are characteristic of certain textures in images, are also affected and smoothed out by this force.

The conclusion derived from this intuitive analysis is that we basically need two opposing forces of diffusion, acting simultaneously on the signal: one is a backward force (at medium gradients, where singularities are expected), and the other is a forward one, used for suppressing oscillations and reducing noise. To benefit from both, we combine them into one backward-and-forward diffusion force with a diffusion coefficient (which is a function of the gradient’s magnitude) that assumes both positive and negative values.

B. Diffusion Coefficient

Consider the following formula of the diffusion coefficient in the form of:

$$c_1(s) = \frac{1}{1 + (s/k_f)^n} - \frac{\alpha}{1 + ((s - k_b)/w)^{2m}}. \quad (4)$$

Our original formulation (used in [10], [11], and [26]) was

$$c_2(s) = \begin{cases} 1 - (s/k_f)^n, & 0 \leq s \leq k_f \\ \alpha [((s - k_b)/w)^{2m} - 1], & k_b - w \leq s \leq k_b + w \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

and its smoothed version

$$c_\sigma(s) = c_2(s) * g_\sigma(s) \quad (6)$$

where $*$ denotes convolution and g_σ is a Gaussian of standard deviation σ . The proposed coefficient (4) is of similar nature, and its parameters play similar roles. This formulation, though, appears to have the following better properties: Processes with smoother diffusion coefficients are better immuned against noise (as indicated by our computational experiments); (4) is easier to implement (no thresholding is needed) and its mathematical analysis is afforded. In our implementation, the exponent parameters (n, m) were chosen to be (4, 2) for c_1 and (4, 1) for c_2 , and $k_f < k_b - w$ for both. The P-M diffusion coefficient, in comparison is

$$c_{P-M}(s) = \frac{1}{1 + (s/k)^2}. \quad (7)$$

Plots of the coefficients and respective fluxes of (4), (5), and (7), are shown in Figs. 3–5, respectively.

A diffusion process defined by c such as in (4)–(6), or by another process of their type, switches adaptively between forward and backward diffusion processes. Therefore, we refer to it as a *forward-and-backward* (FAB) diffusion process.

The coefficient c_2 has to be continuous and differentiable. In the discrete domain, (5) could suffice (although it is only piecewise differentiable). Equations (4) and (6) can fit the general continuous case. Other formulas with similar nature may also be proposed.

Compared with the P-M equation (7), where an “edge threshold” k is the sole parameter, we now have a parameter for the forward force k_f , two parameters for the backward force (we defined them by the center k_b and width w), and the relations between the strength of the backward and forward forces (a ratio we denoted by α). We therefore discuss some rules for determining these parameters (for all three coefficient




Fig. 3. Coefficient c_1 and the corresponding flux, plotted as a function of the gradient magnitude.




Fig. 4. Coefficient c_2 and the corresponding flux, plotted as a function of the gradient magnitude.




Fig. 5. c_{P-M} and the corresponding flux as a function of the gradient magnitude.

equations). The parameter k_f —is essentially the limit of gradients to be smoothed out, and is similar in nature to the role of the k parameter of the P-M diffusion equation. The parameters k_b and w define the range of backward diffusion,

and should be determined by the values of gradients that one wishes to emphasize. In the proposed formula, the range is symmetric, and we restrain the width of the backward diffusion to avoid overlapping the forward diffusion.

One way to determine these parameters in the discrete case, without having any prior information, is by calculating the mean absolute gradient (MAG). For instance, $[k_f, k_b, w] = [2, 4, 1] * MAG$. Local adjustment of the parameters, can be done by calculating the MAG value in a window. The parameters $[k_f(x, y), k_b(x, y), w(x, y)]$ vary gradually along the signal, and enhancement is accomplished by inducing different thresholds in different locations. This is indeed required in cases of natural signals or images because of their nonstationary structure. Usually, a minimal value of forward diffusion should be kept, so that large smooth areas do not become noisy. An example implementing the local parameter adjustment is depicted by the parrot image (Fig. 10). In the deer image (Fig. 11) we adjusted the parameters according to the gradient magnitude of the initial image convolved by a Gaussian (instead of the MAG) obtaining similar results.

The parameter α determines the ratio between the backward and forward diffusion. If the backward diffusion force is too dominant, the stabilizing forward force is not strong enough to avoid oscillations. One can avoid the developing of new singularities over smooth areas in the 1-D case by bounding the maximum flux permissible in the backward diffusion to be less than the maximum of the forward one. (A proof is given in Section II-F.) Formally, we say

$$\max_{s < k_f} \{s \cdot c(s)\} > \max_{k_b - w < s < k_b + w} \{s \cdot c(s)\}. \quad (8)$$

In the case of our proposed coefficients, simple bounds for α satisfying the inequality, are obtained. (These hold for most choices of positive integer exponent parameter combinations n, m .) In the case of c_1 , we have

$$\alpha \leq k_f/2(k_b + w), \quad \text{for any } 0 < w < k_b - k_f \quad (9)$$

and for c_2 and c_σ

$$\alpha \leq k_f/2k_b, \quad \text{for any } 0 < w < k_b - k_f. \quad (10)$$

In practical applications, the above bounds can usually be increased and even doubled in value without experiencing major instabilities.

There are a few ways to regularize this PDE-based approach. Given an *a priori* information on the smallest scale of interest, one can smooth smaller scales in a noisy signal by preprocessing. As we enhance the signal afterwards, the smoothing process does not affect the end result that much. This enables operation in a much noisier environment. Another possibility is to convolve the gradient used for calculating the diffusion coefficient with a Gaussian, following the regularization method proposed by Catta *et al.* [2]. Using relatively smooth diffusion coefficients (low exponent parameters m, n) also reduces the sensitivity of the process to noise.

C. Comparison With Shock Filters

In this section, we will clarify the differences between our scheme and the PDE-based enhancement process of Alvarez and Mazorra (A-M) [1]. The latter procedure couples a diffusion term with the shock filter of Osher and Rudin (O-R) [19], yielding an equation of the form

$$I_t = -\text{sign}(G_\sigma * I_{\eta\eta})|\nabla I| + cI_{\xi\xi} \quad (11)$$

where c is a positive constant, η is the direction of the gradient (∇I), and ξ is the direction perpendicular to the gradient. The first term on the right side creates solutions approaching piecewise constant regions separated by shocks at the zero-crossings of the smoothed second derivative of η . The second term is an anisotropic diffusion along the level-set lines.

The main differences between this scheme and ours are: the scheme of (A-M) is limited to the minimum and maximum gray-level values of the original image; it is therefore more stable but cannot achieve real contrast enhancement. Their scheme is not adaptive with respect to the gradient's magnitude, it is targeted for enhancement everywhere, including smooth regions. Equation (11) contains a diffusion term along the level-sets (a curvature-flow process); therefore, small features are sometime eliminated and rounding of objects occurs. Edges are determined by the zero-crossings of the smoothed second derivative along the gradient direction, whereas we use the P-M model of gradient dependent diffusivity. The first method may well serve edge detection purposes, but may not necessarily be useful as a means of controlling an enhancement process: indicating edges by the zero-crossing is a binary decision process (and not a fuzzy one) and may be therefore less immuned against noise. The Gaussian convolution may help to some extent but many false edges will still be enhanced (as the gradient magnitude is not considered). Applying a very wide Gaussian to increase robustness will dislocate edges and wipe out many image details. In the following section, we compare, by means of an example, the scheme of (A-M) with ours.

D. Examples

We used the explicit Euler scheme with a forward difference scheme for the time derivative, and the central difference scheme with a 3×3 kernel for the spatial derivatives. Examples of signal and image restoration using the selective inverse diffusion are shown in the following.

A blurred and noisy step signal [Fig. 6(b)] was processed, assuming the availability of some prior information regarding the expected range of noise power and the approximate size of the original step. Enhancement of the step and denoising the rest of the signal are clearly depicted in Fig. 6(e). The second example illustrates simultaneous denoising and enhancement of a blurred multistep signal contaminated by uniform noise (Fig. 7).

The FAB is also effective in the enhancement of images, as is illustrated in Figs. 8, 10, and 11. In Fig. 8, we compare our results with those obtained by the application of the A-M process [(11)]. The A-M process indeed enhances the objects and forms clear edges, but small details are lost (the teddy bear's shirt patterns, for example). It consequently appears as though the image has lost its natural appearance. Also, relatively smooth regions




Fig. 6. FAB processing of a single-step noisy signal (from top to bottom): blurred step; signal contaminated by white Gaussian noise (SNR = 7 dB); results of FAB diffusion process after 1, 3, and 30 time steps, respectively. $[k_f, k_b, w] = [1/6, 1, 1/3]$.




Fig. 7. FAB processing of multiple-step noisy signal (from top to bottom): original signal (with both positive and negative discontinuities); blurred signal contaminated by white uniform noise (SNR = 8 dB); results of FAB diffusion process obtained after 2, 4, and 16 time steps, respectively. $[k_f, k_b, w] = [0.1, 0.8, 0.2]$.

(like the table) are erroneously enhanced, creating artificial textures. A plot of a horizontal line from the middle of the image (Fig. 9) highlights in some more detail the different behavior of the two enhancement processes. In Figs. 8 and 11, a fidelity term $\lambda(I_0 - I)$ was added to the evolution equation ($\lambda = 0.05$).

In the second example (Fig. 10), we implemented the automatic adjustment of local parameters, explained previously. This was needed to enhance gradients at different locations differently. For example, the gradients near the parrot's beak are very large, whereas those around the eye are much smaller. Enhancing both regions without blurring important details on one hand, and maintaining stability on the other hand, required a completely different set of parameters (k_f, k_b, w). In the process of implementing automatic adjustment, the parameters vary gradually with image location, and the enhancement appears to be natural. In Fig. 11, we compare the different results of constant versus spatially varying parameters, where



Fig. 8. Comparison between regularized shock filter (A-M) and FAB. Top: original, middle: shock filter ($\sigma = 1, c = 0.5$), and bottom: FAB process ($[k_f, k_b, w] = [2, 20, 10]$).




Fig. 9. Plot of gray level values obtained along one line of Fig. 8: top: original, middle: shock, and bottom: FAB.




Fig. 10. FAB diffusion process applied to parrot image, with local parameter adjustment using the MAG measure: Top left: original image, right: blurred image, and bottom (from left): diffusion process after time steps 1 and 8, respectively.

by using the latter method, the deer, as well as the trees behind them, are enhanced. We should comment that the process is not suited too well for the handling of textures, as seen in the spots created at the textured ground.

E. Feature-Based Diffusion Coefficient

We further extend and generalize the nonlinear PDE-based filtering method, and apply it as a combined feature-based enhancement and denoising mechanism. In order to avoid smoothing out important features of the image such as textures, we should ideally have a local feature detector that will slow down or even reverse the diffusion process in the vicinity of important features.



Fig. 11. FAB diffusion process applied to the deer image. From top: original image, result of processing with constant parameters $k_f = 2$, $k_b = 50$, $w = 10$; magnitude of smoothed gradient of original image $T(x, y) = |\nabla I_0 * G_\sigma|$, ($\sigma = 3$); result of processing with spatially varying parameters $k_f(x, y) = 0.1T$, $k_b(x, y) = 6T$, $w(x, y) = 2T$.

We minimize the amount of noise induced by the processing, which is inherently a byproduct of signal enhancement, by our generalized forward-and-backward diffusion processes. Moreover, important features are not filtered out by the forward diffusion process, enabling a complementary image processing mechanism to enhance them at a later stage, whenever it is necessary.

We propose the following general feature enhancing and denoising mechanism: let

$$c(x, y) = c(f_1, f_2, \dots, f_n)$$

where the local feature estimators $f_i \equiv f_i(x, y)$ $i = 1, \dots, n$ can be selected from a broad range of choices introduced in the fields of image processing and computer vision, e.g., edge detectors (already introduced implicitly under the gradient criterion), noise estimators, texture, scale, orientation, curvature, local power-spectrum, moments estimators, etc. The logic dictating the value of the diffusion coefficient c should be as follows: Forward diffuse features that should be filtered out because they are corrupted by noise and are of no importance to the image understanding or appearance. Backward diffuse features that should be enhanced. Avoid diffusion where either diffusion processes (forward or backward) would distort important features.

In cases where there is some *a priori* knowledge of the type of images to be processed, the diffusion process may be much better controlled.

To illustrate this feature-dependent diffusion, consider, for example, an urban scene, primarily comprised of buildings. In this case, one would like to preserve most vertical and horizontal lines and edges, significant wall textures and additional dominant edges at all orientations. To incorporate these requirements into our diffusion process, let us define by the symbols $e_e(x, y)$, $e_t(x, y)$, $e_{vl}(x, y)$, $e_{hl}(x, y)$ the local estimators that stand for edges, wall textures, vertical-lines, and horizontal-lines, respectively. An appropriate diffusion coefficient for the process is, in this case, given by

$$c(e_e, e_t, e_{vl}, e_{hl}) = \frac{1}{1 + w_e e_e + w_t e_t + w_{vl} e_{vl} + w_{hl} e_{hl}} \quad (12)$$

where w_x denotes the relative weight required to balance the desired effect of each estimator. In this simplified example, it is clear that the diffusion process will slow down considerably whenever the output of at least one of the weighted estimators is much larger than 1 ($w_x e_x \gg 1$, $x \in \{e, t, vl, vh\}$). In other areas of the image, a stronger forward diffusion will reduce the noise.

F. Stability of Smooth Regions in 1-D

Problem definitions:

- The flux J is defined as follows:

$$J \equiv c(\cdot)I_x(x, t).$$

(Note that flux in physical problems is usually defined with an opposite sign.) We assume a diffusion coefficient of the type $c = c(|I_x|)$, leading to the flux properties

$$J = J(I_x(x, t))$$

and the antisymmetry relation $J(s) = -J(-s)$.

- The nonlinear diffusion equation, with its initial and boundary conditions, is

$$\begin{aligned} I_t &= J_x = \frac{\partial}{\partial x} (c(|I_x|)I_x) \\ 0 \leq x \leq 1, \quad I(x, t=0) &= I_0(x); \\ I_x(0, t) = I_x(1, t) &= 0. \end{aligned}$$




Fig. 12. Nonmonotonic flux of a forward diffusion process and its critical points M and r .

Lemma 1: If $J(x_0, t_0)$ is a local maximum (minimum) in the spatial domain, then $I_{xt}(x_0, t_0) < 0$ ($I_{xt}(x_0, t_0) > 0$).

Proof:

$$\begin{aligned} I_t &= J_x \\ I_{xt} &= J_{xx}. \end{aligned}$$

If $J(x_0, t_0)$ is a local maximum, then $J_{xx}(x_0, t_0) < 0$. If $J(x_0, t_0)$ is a local minimum, then $J_{xx}(x_0, t_0) > 0$. ■

In Theorem 1, we regard the simpler case of positive diffusion coefficient ($c(\cdot) > 0$) with nonmonotonic flux. We prove that once a gradient gets into the smooth band $I_x \in [-r, r]$, where r is the point of maximum flux, it remains trapped there.

The maximum of the flux (see Fig. 12) is defined by

$$M \equiv \max_{I_x} \{J(I_x)\}; \quad r = \arg \max_{I_x} \{J(I_x)\}.$$

Theorem 1 (Smooth Band “Trap”): If $J(r) \equiv M = \max_{I_x} \{J(I_x)\}$, $c(s) > 0 \forall s$, and $|I_x(x_0, t_0)| < r$, then $|I_x(x_0, t)| < r$ for any $t > t_0$.

Proof: Let us assume that at a time $t_2 > t_0$ we would have $I_x(x_0, t_2) > r$. From the continuity of the gradient in time, it follows that there should be a certain time t_1 , $t_0 < t_1 < t_2$, such that

$$I_x(x_0, t_1) = r, \quad (I_{xt}(x_0, t_1)) > 0.$$

However, since $J(x_0, t_1) = M$ must be a local maximum, it follows from *Lemma 1* that $(I_{xt}(x_0, t_1)) < 0$, which contradicts our assumption.

Similarly, the assumption that at time t_2 we would have $I_x(x_0, t_2) < -r$ cannot hold. ■

Theorem 2 is the version of *Theorem 1*, adapted to our proposed FAB coefficient, having both positive and negative values of c .

The points of extrema of flux, in a FAB diffusion process, are defined as follows (see Fig. 13):

$$\begin{aligned} M_f &\equiv \max_{I_x > 0} \{J(I_x)\}; & \{r_f: J(I_x = r_f) = M_f\}, \\ M_b &\equiv -\min_{I_x > 0} \{J(I_x)\}; & \{r_b: J(I_x = r_b) = -M_b\}. \end{aligned}$$




Fig. 13. Flux and critical points of the FAB process.

This theorem states that, in the 1-D case, a point x_0 with an initial gradient magnitude below r_f will not assume a gradient magnitude larger than r_f (i.e., will stay “smooth”) through the entire forward-and-backward diffusion process, provided the forward maximum flux, M_f , is larger than the backward one, M_b .

Theorem 2 (Stability of Smooth Regions): If $M_f > M_b$, then, for every x_0 for which $|I_x(x_0, 0)| < r_f$, the derivative stays bounded at all times, i.e., $|I_x(x_0, t)| < r_f$ for any $t > 0$.

Proof: Follows directly from Theorem 1, letting $M = M_f$, $r = r_f$ and $t_0 = 0$. The fact that c can also have negative values does not affect the proof. ■

In the case where $M_b > M_f$, r_f is not guaranteed to be a local maximum.

G. Analysis of Theorem 2 in the Discrete Case

As the equations are solved numerically, we must first see if the theorem holds also in the discrete case. The main property, that we relied on in proving our theorem, is the continuity of the gradient in time, which applies only in the continuous domain. We therefore have to analyze the implications of the discrete case. Starting with the original diffusion equation

$$I_t = J_x$$

we replace the first temporal derivative by the forward difference, with a time step of Δt

$$\begin{aligned} \frac{I(x, t + \Delta t) - I(x, t)}{\Delta t} &= J_x(x, t), \\ I(x, t + \Delta t) &= I(x, t) + \Delta t J_x(x, t), \\ I_x(x, t + \Delta t) &= I_x(x, t) + \Delta t J_{xx}(x, t) \end{aligned}$$

[for brevity, we use $J_x(x, t)$ instead of $J_x(I_x(x, t))$]. Assuming I_x is in the positive “smooth band,” that is $0 < I_x(x, t) < r_f$, according to our theorem, it cannot get out of this band in the next time step, hence the following condition must be satisfied:

$$I_x(x, t + \Delta t) = I_x(x, t) + \Delta t J_{xx}(x, t) < r_f \quad (13)$$

where we regard only the case of positive I_x without loss of generality. Replacing the second spatial derivative by the central difference with a step $\Delta x \equiv h$, and using the Euler method, the condition changes to

$$I_x(x, t) + \Delta t \frac{J(x+h, t) - 2J(x, t) + J(x-h, t)}{h^2} < r_f.$$

Assigning $x = x_0$, $I_x(x_0, t) = r_0$, $J(x_0, t) = J_0 = r_0 c(r_0)$, and using the flux bound $J \leq M_f$ it is sufficient to prove that

$$r_0 + \Delta t \frac{2(M_f - J_0)}{h^2} < r_f$$

and finally

$$\frac{\Delta t}{h^2} < \frac{1}{2} \frac{r_f - r_0}{M_f - J_0}. \quad (14)$$

In order to maintain numerical stability in any such scheme, the known CFL bound [6] must be obeyed, i.e., (in the 1-D case)

$$\frac{\Delta t}{h^2} \leq \frac{1}{2}$$

and therefore we must only ensure that

$$\frac{r_f - r_0}{M_f - J_0} > 1.$$

If $\partial J(I_x)/\partial I_x$ is monotonically decreasing in the range $I_x \in (0, r_f)$ (this condition is satisfied by our proposed coefficients and by other proposed gradient-based nonlinear schemes), it is clear that

$$\frac{r_f}{M_f} < \frac{r_f - r_0}{M_f - J_0}.$$

Substituting $M_f = r_f c(r_f)$, our proof amounts to showing the simple relation of

$$\frac{1}{c(r_f)} > 1$$

and since $c(s) < 1$ for any $s > 0$, we may conclude that the theorem holds for the discrete case.

Otherwise, if $\partial J(I_x)/\partial I_x$ is not monotonically decreasing, (14) provides a bound for the time step.

III. SUPER-RESOLUTION BY THE FAB PROCESS

The FAB diffusion process is useful in applications requiring simultaneous enhancement and smoothing. We present a simple super-resolution (SR) scheme, incorporating two main subsystems: an interpolator and an enhancer–denoiser, as shown in Fig. 14.

A. Some Background: What is SR?

By SR, we refer to the process of artificially increasing the resolution of an image, using side information about the structure of any specific subset of images or of natural images in general. The processed image should not only have more pixels, but more importantly, be characterized by a wider band than that of the original image.

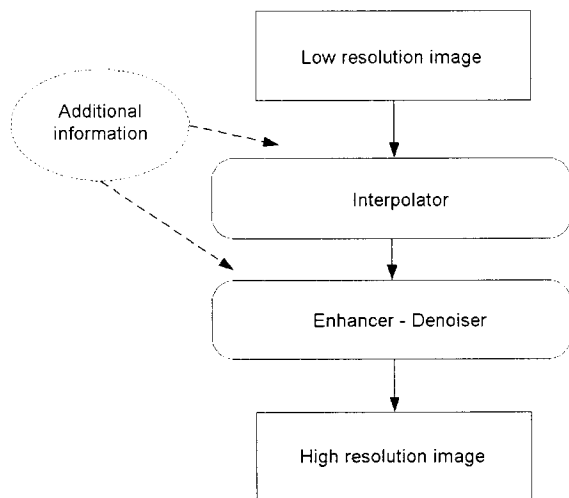


Fig. 14. Super-resolution processor.

Most applications of SR use several images obtained from the same scene or object, taken from slightly different angles or locations. After proper registration, a higher resolution image can be obtained from the low-resolution images by exploiting the combined information available at the different sets of sampling points. Examples of such SR procedures can be found in studies conducted at NASA on satellite images [4], by Schultz and Stevenson [25] processing a series of movie frames, and by [8].

B. Proposed Scheme: Single Image SR

We elaborate an approach suitable for SR based on a single image, similarly to [32]. Instead of using a sequence of video frames or multiple exposure, we exploit the properties common to a wide range of natural images. Obviously, there are cases where only one image is available, and one would still like to enhance the resolution.

Our assumption is that images can be segmented into regions falling into one of the following three categories: smooth areas, edges, and textured regions. At this point, we simplify our model and consider only images that are not endowed by significant textural attributes, that is, they can be approximated by piecewise-smooth segments separated by edges.

The proposed scheme receives a low-resolution image as an input, with possibly some prior information about the structure of the scene. The processing is executed in two steps. First, the image is interpolated to the new desired size. In our implementation, we used cubic B-spline interpolation, but other methods may also be used. The first step provides good results over smooth areas, but edges are smeared. The interpolated images often depict ringing effects, with low spatial oscillations. The purpose of the second processing step is to enhance the edges and denoise the interpolation byproducts. This is accomplished by using the FAB diffusion process. In our implementation, the parameters k_f , k_b , w were locally adjusted according to the mean gradient criterion.

C. Resolution Enhancement—An Example

Consider a narrow-band system, such as a cellular phone, that permits the communication of only low-resolution images at a reasonable rate. We wish to enhance the resolution of an image at the receiving end of the communication channel in such a way that it will appear as though a high-resolution image was transmitted over a wideband channel.

We downsample an input of a high-resolution image by four in each dimension and send the low-resolution “blocky” image (that is, 1/16 of the original size). At the receiving end, we apply the proposed SR process: the image is up-sampled and enlarged back to its original size. The FAB process is then applied. The end result [Fig. 15(d)] looks more like the original image [Fig. 15(a)] than the low-resolution image [Fig. 15(b)].

A considerable improvement can be gained by transmitting some side information in addition to the image itself. Such side information may include suitable parameters of the FAB process, specification of segments where enhancement should be avoided or emphasized, etc. Whenever the original high-resolution image is available at the transmitting end of the channel, one can find much more easily the optimal parameters suitable for the task.

In the previous example (Fig. 15), we assumed that additional information was available, specifying where enhancement should be avoided. Such image segments are typically blurry and fuzzy in the first place, clouds for instance. In Fig. 15(e), we show the result of avoiding enhancement of most of the sky (above a certain horizontal line in the image). This results in fuzzy clouds, whereas the mountains below are crisp and sharp. To compare with, in the global enhancement [Fig. 15(d)], the clouds are also sharpened and lose their natural appearance.

We should emphasize here that this resolution enhancement process does not come to replace ordinary image compression. It can be used as an additional tool that improves the overall performance in terms of bandwidth of the final image that is displayed. Indeed, the image of Fig. 15(d) [or Fig. 15(e)] is of a wider band than that of the transmitted one [Fig. 15(b)].

IV. COLOR PROCESSING

A. Beltrami Framework

The original study of Sochen *et al.* [27] unifies several approaches by means of the Beltrami framework, and offers new definitions and solutions for various image processing tasks. According to the extended Beltrami framework, images, visual objects, and their characteristics of interest such as derivatives, orientations, texture, disparity in stereo vision, optical flow, and more, are described as embedded manifolds. The embedded manifold is equipped with a Riemannian structure, i.e., a metric that encodes the geometry of the manifold. Nonlinear operations are acting on these objects according to the proper local geometry. Iterative processes are considered in this context as evolution of manifolds. The latter is a consequence of the action of a nonlinear diffusion process or another type of a nonlinear PDE. No global (time-wise) kernels can be associated with these nonlinear PDEs. Short time kernels for these processes were derived recently in [29] and [40].



(a)



(b)



(c)



(d)



(e)

Fig. 15. Application of the SR process: (a) original high-resolution image, (b) low-resolution blocky input image with 1/16 of the original pixels, (c) the image shown in (b) upsampled and interpolated by a cubic B-spline, (d) image obtained after FAB processing, and (e) FAB processing with additional side information, avoiding enhancement of most of the sky area.

We follow the studies presented by Perona and Malik [20], Sochen *et al.* [27], Kimmel *et al.* [14], Sochen and Zeevi [28], [30], and Weickert [33], and show how one can design a structure tensor that controls the nonlinear diffusion process starting from the induced metric that is given in the Beltrami framework. The proposed structure tensor is nondefinite positive, or negative, and switches between these states according to image features. This results in a *FAB diffusion flow*, and different regions of the image are either forward or backward diffused, according to the local geometry within a neighborhood. The adaptive property of the process, that finds its expression in the local decision on the *direction* of the diffusion and on its *strength*, is the main novelty of this section.

B. Geometric Measure on Embedded Maps

1) *Images as Riemannian Manifolds*: According to the geometric approach to image representation, images are considered to be 2-D Riemannian surfaces embedded in higher dimensional spatial-feature Riemannian manifolds [13]–[15], [27], [28], [30]. Let u^1, u^2 be the local coordinates on the image surface, and let $X^i, i = 1, 2, \dots, m$ be the coordinates of the embedding space; then, the embedding map is given by

$$(X^1(u^1, u^2), X^2(u^1, u^2), \dots, X^m(u^1, u^2)). \quad (15)$$

Riemannian manifolds are endowed with a bilinear positive-definite symmetric tensor which constitutes a *metric*. Let $(\Sigma, (g_{\mu\nu}))$ denote the image manifold and its metric, and $(M, (h_{ij}))$ denote the spatial-feature manifold and its corresponding metric. Then, according to Polyakov, the map $\mathbf{X}: \Sigma \rightarrow M$ has the following weight [22]:

$$E[X^i, g_{\mu\nu}, h_{ij}] = \int d^2u \sqrt{g} g^{\mu\nu} (\partial_\mu X^i) (\partial_\nu X^j) h_{ij}(\mathbf{X}) \quad (16)$$

where the range of indices is $\mu, \nu = 1, 2$, and $i, j = 1, \dots, m = \dim M$, and where we use the Einstein summation convention: identical subscript and superscript indices are summed over. We denote by g the determinant of $(g_{\mu\nu})$ and by $(g^{\mu\nu})$ the inverse of $(g_{\mu\nu})$. The measure $d^2u \sqrt{g}$ is an area element of the image manifold, and $g^{\mu\nu} (\partial_\mu X^i) (\partial_\nu X^j) h_{ij}(\mathbf{X})$ is a generalization of the L_2 norm for gradients, from Euclidean spaces to manifolds. The last two expressions do not depend on the choice of local coordinates.

A gradient descent evolution along the feature coordinates is derived from the respective Euler–Lagrange equations, multiplied by a strictly positive function and a positive definite matrix, in order to gain reparameterization invariance

$$X_t^i \equiv \frac{\partial X^i}{\partial t} = -\frac{1}{2\sqrt{g}} h^{il} \frac{\delta E}{\delta X^l}. \quad (17)$$

Given a Euclidean embedding space, with Cartesian coordinate system, the variational derivative of E with respect to the coordinates of the embedding space is given by

$$-\frac{1}{2\sqrt{g}} h^{il} \frac{\delta E}{\delta X^l} = \Delta_g X^i = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu X^i) \quad (18)$$

where Δ_g , referred to as the *Beltrami operator* [16], is a generalization of the Laplacian from flat surfaces to manifolds.

Assuming an isometric embedding, the image manifold metric can be deduced from the mapping \vec{X} and the embedding space's metric h_{ij}

$$g_{\mu\nu} = h_{ij} \partial_\mu X^i \partial_\nu X^j. \quad (19)$$

It is called the *induced metric*.

2) *Metric as a Structure Tensor*: There have been a few studies using anisotropic diffusion processes. Cottet and Germain [5] used a smoothed version of the image to orient the diffusion, while Weickert [34], [35] smoothed also the structure tensor $\nabla I \nabla I^T$ and then manipulated its eigenvalues to steer the smoothing orientation. Elimination of one eigenvalue from a structure tensor, first proposed as a color tensor in [7], was used in [24], in which case the tensors are not necessarily positive definite. In [33] and [36], the eigenvalues are manipulated to result in a positive definite tensor. (See also [3], where the diffusion is in the direction perpendicular to the maximal gradient of the three color channels; a direction that is different from that of [24].)

We will follow and generalize the analysis elaborated by Kimmel *et al.* in [14]. For completeness, we reiterate some of the relations developed in that study. Let us first show that the direction of the diffusion can be deduced from the smoothed metric coefficients $g_{\mu\nu}$ and may thus be included within the Beltrami framework under the right choice of directional diffusion coefficients.

The induced metric ($g_{\mu\nu}$) is a symmetric positive definite matrix that captures the geometry of the image surface. Let λ_1 and λ_2 be the large and the small eigenvalues of ($g_{\mu\nu}$), respectively. Since ($g_{\mu\nu}$) is a symmetric positive matrix, its corresponding eigenvectors u_1 and u_2 can be chosen orthonormal. Let $U \equiv (u_1 | u_2)$, $\Lambda \equiv \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$, and therefore

$$(g_{\mu\nu}) = U \Lambda U^T. \quad (20)$$

Let us define

$$(g^{\mu\nu}) \equiv (g_{\mu\nu})^{-1} = U \Lambda^{-1} U^T = U \begin{pmatrix} 1/\lambda_1 & 0 \\ 0 & 1/\lambda_2 \end{pmatrix} U^T, \quad (21)$$

and

$$g \equiv \det(g_{\mu\nu}) = \lambda_1 \lambda_2. \quad (22)$$

Our proposed enhancement procedure controls the above determined eigenvalues adaptively, so that only meaningful edges are enhanced, whereas smooth areas are denoised.

C. Adaptive Structure Tensor

1) *Controlling the Eigenvalues*: From the previous derivation of the induced metric $g_{\mu\nu}$, it follows that the larger eigenvalue λ_1 corresponds to the eigenvector in the gradient direction [in the three-dimensional (3-D) Euclidean case: (I_x, I_y)]. The smaller eigenvalue λ_2 corresponds to the eigenvector perpendicular to the gradient direction [in the 3-D Euclidean case: $(-I_y, I_x)$]. The eigenvectors are equal for both $g_{\mu\nu}$ and its inverse $g^{\mu\nu}$, whereas the eigenvalues have reciprocal values. We can use the eigenvalues as a means to control the Beltrami flow process. For convenience, let us define $\lambda^1 \equiv 1/\lambda_1$. As the

first eigenvalue of $g^{\mu\nu}$ (that is λ^1) increases, so does the diffusion force in the gradient direction. Thus, by changing this eigenvalue, we can reduce, eliminate, or even reverse the diffusion process in the gradient direction. Similarly, changing $\lambda^2 \equiv 1/\lambda_2$ controls the diffusion in the level-set direction.

What is the best strategy to control the diffusion process via adjustment of the relevant parameters? The following requirements may be considered as guidelines:

- the enhancement should essentially be with relevance to the important features, while smooth segments should not be enhanced;
- the contradictory processes of enhancement and noise reduction by smoothing (filtering) should coexist;
- the process should be as stable as possible, though restoration and enhancement processes are inherently unstable.

Let us define $\alpha_1(s)$ as a new adaptive eigenvalue to be considered instead of the original λ^1 . We propose an eigenvalue that is a function of the determinant of the smoothed metric. The formulation of the new eigenvalue is the same as the FAB diffusion coefficient, that is

$$\alpha_1(s) = c(s) \quad (23)$$

where $c(s)$ is defined by (4) and s , here, is chosen to be a function of the smoothed metric: $s = \sqrt{\det(g_{\mu\nu} * G_\rho)}$.

2) *Algorithm for Color Image Enhancement*: To implement the flow $\mathbf{I}_t = \Delta_{\hat{g}} \mathbf{I}$ for color image enhancement, we modify and generalize the algorithm of [14] as follows.

- 1) Compute the metric $g_{\mu\nu}$. For the N channel case (for conventional color mapping $N = 3$) we have

$$g_{\mu\nu} = \delta_{\mu\nu} + \sum_{k=1}^N I_\mu^k I_\nu^k. \quad (24)$$

- 2) Diffuse the $g_{\mu\nu}$ coefficients by convolving them with a Gaussian of variance ρ , thereby

$$\tilde{g}_{\mu\nu} = G_\rho * g_{\mu\nu}. \quad (25)$$

- 3) Compute the inverse smoothed metric $\tilde{g}^{\mu\nu}$. Change the eigenvalues of the inverse metric $\lambda^1, \lambda^2, (\lambda^1 < \lambda^2)$, of $\tilde{g}^{\mu\nu}$ to $\alpha_1(s), \alpha_2$, respectively. The new second eigenvalue should be in the range $0 < \alpha_2 \leq 1$, preferably minimal ($\alpha_2 \ll 1$) when the image is not noisy. This yields a new inverse structure tensor $\hat{g}^{\mu\nu}$ that is given by

$$(\hat{g}^{\mu\nu}) = \tilde{U} \begin{pmatrix} \alpha_1(s), & 0 \\ 0, & \alpha_2 \end{pmatrix} \tilde{U}^T = \tilde{U} \hat{\Lambda} \tilde{U}^T. \quad (26)$$

- 4) Calculate the determinant of the new structure tensor. Note that \hat{g} can now have negative values.
- 5) Evolve the k th channel via the Beltrami flow

$$I_t^k = \Delta_{\hat{g}} I^k \equiv \frac{1}{\sqrt{\hat{g}}} \partial_\mu \left(\sqrt{\hat{g}} \hat{g}^{\mu\nu} \partial_\nu I^k \right). \quad (27)$$

Remark: In this flow, we do not get imaginary values, though we have the term $\sqrt{\hat{g}}$, since in cases of negative \hat{g} the constant imaginary term $i \equiv \sqrt{-1}$ will be canceled.

3) *Comparison to Previous Studies*: There are two important differences between our scheme and that of Kimmel *et al.* [14, last section]. These concern the possible choice of eigenvalues that control the process. As will be illustrated by examples, our choice of eigenvalues may substantially improve the sharpening of natural images.

Let us first return to some analysis of the eigenvalues. For the sake of simplicity, we analyze the eigenvalues in the context of the structure tensor of the smoothed image $I * G_\rho$ (instead of the smoothed structure tensor). We examine (27) for a single channel ($N = 1$), where α_1 and α_2 are arbitrary eigenvalues of $\hat{g}^{\mu\nu}$. For the degenerate case of $\rho = 0$, we get $I_t = \alpha_1 \Delta I$. It can be shown that for any ρ other than zero, we can write the equation as

$$I_t = \alpha_1 I_{\tilde{\eta}\tilde{\eta}} + \alpha_2 I_{\tilde{\xi}\tilde{\xi}} + E \quad (28)$$

where

$$I_{\tilde{\eta}\tilde{\eta}} \doteq \frac{I_{xx}\tilde{I}_x^2 + 2I_{xy}\tilde{I}_x\tilde{I}_y + I_{yy}\tilde{I}_y^2}{\tilde{I}_x^2 + \tilde{I}_y^2}$$

$$I_{\tilde{\xi}\tilde{\xi}} \doteq \frac{I_{xx}\tilde{I}_y^2 - 2I_{xy}\tilde{I}_x\tilde{I}_y + I_{yy}\tilde{I}_x^2}{\tilde{I}_x^2 + \tilde{I}_y^2}$$

$$E \doteq O(\tilde{I}_{xx} + \tilde{I}_{xy} + \tilde{I}_{yy})$$

where a tilde above any expression indicates that it has been convolved with a Gaussian of standard deviation ρ (e.g., $\tilde{I}_x \equiv I_x * G_\rho$). For large enough ρ , we can assume the relation $I_{xx} \gg \tilde{I}_{xx}$, and similarly for the rest of the second derivatives. This relation holds especially at regions of high frequencies—typically those image regions containing edges, textures, or noise, as high frequencies decay exponentially by the Gaussian convolution. Thus, E of (28) is small in comparison to the other terms. The end result is that we get an anisotropic diffusion process with a diffusion coefficient α_1 in the direction of the smoothed gradient $(\tilde{I}_x, \tilde{I}_y) = \nabla I * G_\rho$ and with a different diffusion coefficient, α_2 , in the perpendicular direction.

This analysis holds for Weickert's coherence enhancing diffusion, Kimmel *et al.*'s scheme, and our modification (α_1, α_2 should not necessarily be constants). It may be concluded at this point that ρ controls the measure of directionality of the processes as follows:

- for small ρ , a close to isotropic diffusion takes place, controlled by α_1 ;
- for large ρ , a strong anisotropic diffusion occurs and it is being controlled by α_1 in the smoothed gradient direction, and by α_2 in the smoothed level-set direction;
- for constant $\alpha_1 > 0, \alpha_2 > 0$, the process shifts (according to ρ) from linear forward diffusion to strong coherence enhancing diffusion;
- for constant $\alpha_1 < 0, \alpha_2 > 0$, the process shifts from linear backward (inverse) diffusion to a Gabor-type process [9], where both processes are unstable.

The relation of anisotropic diffusion, using a tensor diffusivity, to Gabor's idea was mentioned previously in a few studies (such as [14] and [18]), yet, to the best of our knowledge, the relation of (28) has not been stated before.

From a numerical viewpoint, the CFL condition in explicit schemes for any ρ (in 2-D) is $\Delta t \leq h^2/4 \max(\alpha_1, \alpha_2)$. Therefore, it is a good practice to limit both eigenvalues to be smaller than one.

Let us now return to the differences of our scheme to that of [14]. In the latter study, the focus is on the first eigenvalue and on its manipulation. Therefore, the following constraint is proposed: $|\hat{g}| = |\det(\hat{g}_{\mu\nu})| = |\lambda_1\lambda_2| = 1$. Choosing $\alpha_1 = \alpha$, one gets $\alpha_2 = \alpha^{-1}$. When $|\alpha| \ll 1$, the process is completely dominated by the second eigenvalue (e.g., for $\alpha_1 = \pm 0.01$ the diffusion force in the level set direction is 10^4 times stronger, and the sign of α_1 practically does not affect the process). Smoothing along the level-set lines can be effective in images endowed with orientational structure, such as those characteristic of fingerprints. In general, though, smoothing along level set curves has the effect of turning nondirectional textures (and noise!) into artificial zebra-type stripes, which is a drawback for a general sharpening process. The examples of the Mandrill and Buttons images in [14] clearly depict this effect.

Another important property of the present study that is missing in the one reported in [14] is the adaptive characteristic of the first eigenvalue, implemented in the present study according to the FAB principles. Whereas in [14] the enhancement is being implemented everywhere, even in smooth regions, the process proposed in the present study does it selectively. In our scheme, the enhancement is directed to locations of edges and some dominant textures. Global enhancement with a constant negative α_1 causes a considerable noise amplification and the creation of artificial edges at smooth regions. Examples of these phenomena can be found in the experimental results of the subsequent Section IV-D.

Regarding stability, for $\rho \rightarrow 0$ the scheme of [14] behaves like 2-D inverse diffusion, whereas our scheme behaves like a 2-D FAB process (which is much more stable). For large ρ , in regions where the gradient direction stays constant, the scheme proposed in [14] behaves like 1-D inverse diffusion along this direction, whereas our scheme behaves like a 1-D FAB process.

D. Experimental Results

We applied three Beltrami-type processes to the Iguana color image: the original scheme of [14]; a modified version of [14], where the second eigenvalue is small; and our Beltrami-FAB process. The results presented in Fig. 16 show that in the first process, smoothing along the edges is very dominant, creating snake-like features at places of nonorientational textures (like the sand). The second process (using a small value of α_2), creates strong sharpening effects but amplifies noise at smooth regions (like the sea), as is clearly depicted in the enlargement (Fig. 17). Our Beltrami-FAB process seem to behave well in this relatively complex natural image.

In Fig. 18, we show the effects of enhancement on a compressed image. The Tulip image was highly compressed according to the JPEG standard. A known byproduct of JPEG compression is the blocking effects created at smooth regions. Indeed, the original and modified schemes of [14] enhanced the 8×8 block boundaries whereas our scheme smoothed them out.

(The color images are available at <http://visl.technion.ac.il/ip-fab>.)



Fig. 16. Iguana image processed by three Beltrami-type processes. From top: original, scheme of [14] ($\alpha_1 = 0.3, \alpha_2 = \alpha_1^{-1}$), modified [14] with small α_2 ($\alpha_1 = 0.3, \alpha_2 = 0.01$), and Beltrami-FAB process ($\alpha_1 = \alpha_1(s), [k_f, k_b, w, \alpha] = [10, 2000, 1000, 0.5], \alpha_2 = 0.01$). All processes ran 13 iterations, $dt = 0.1, \rho = 2$.

V. DISCUSSION AND CONCLUSIONS

Sharpening and denoising are contradictory requirements in image enhancement. We show how they can be reconciled by a local decision mechanism that controls the orientation, type,



Fig. 17. Enlargement of a segment of the iguana's head, with the sea at the background. From left: original; image processed by a modification of [14] with small α_2 ; and by the Beltrami-FAB. Note that smooth regions like the sea are not becoming noisy due to processing by our scheme.




Fig. 18. Segment of the compressed tulip image processed by three Beltrami-type schemes. Top (from left): original; result of processing by the scheme of [14] with $\alpha_1 = 0.5, \alpha_2 = \alpha_1^{-1}$; processed by a modified [14] with small α_2 ($\alpha_1 = 0.5, \alpha_2 = 0.1$); and processed by the Beltrami-FAB process with $\alpha_1 = \alpha_1(s), [k_f, k_b, w, \alpha] = [30, 300, 200, 0.5], \alpha_2 = 0.1$. All processes ran ten iterations, $dt = 0.1, \rho = 1$. Note that the JPEG blocking artifacts are not enhanced by the Beltrami-FAB process.

and extent of the diffusion process. The combined FAB diffusion process offers practical advantages over previously proposed studies in enhancement of image quality.

One of the important aspects of any attempt to implement a truly backward diffusion process in image processing (i.e., a process where the diffusion coefficient becomes negative) is the inherent instability. Since the physical diffusion and heat propagation occur only as a forward process, the mathematical model that well represents the physics becomes ill-posed when the diffusion coefficient changes its sign. As is well known, stability is not well-defined in ill-posed problems. It is therefore important to take a note of the fact that stability is afforded over certain regimes in the case of the FAB diffusion. We have proven stability for small gradient bands in the 1-D case, and verified the feasibility of our approach on a variety of signals and images. Intuitively, the stability in the backward process is afforded by its limitation to small areas of very few pixels, surrounded by larger areas of many more pixels, where the forward diffusion provides a “safety belt” that avoids explosion. Indeed, since the majority of pixels in natural images are characterized by low gradients and mainly singular edges give rise to the reversal of

the diffusion coefficient sign, stability is achieved. This argument does not hold any longer when the FAB diffusion process encounters a highly textured or extremely noisy image.

Yet another related facet of the present study is the generalization of the framework of the Beltrami flow for *adaptive processing* of color images. This is accomplished by replacing the eigenvalues of the color image metric by an adaptive coefficient that locally controls the orientation and extent of the diffusion. The decision of where and how to adapt the coefficient is based on the edge's direction and strength, defined by the eigenvectors and determinant of the smoothed image metric, respectively. FAB diffusion process takes place in the direction of the gradient and forward diffusion takes place in the perpendicular direction.

Examples illustrate that this approach works, and that sharpening and denoising can be combined together in the enhancement of gray-level and color images.

REFERENCES

- [1] L. Alvarez and L. Mazorra, "Signal and image restoration using shock filters and anisotropic diffusion," *SIAM J. Numer. Anal.*, vol. 31, no. 2, pp. 590–605, 1994.
- [2] F. Catte, P. L. Lions, J. M. Morel, and T. Coll, "Image selective smoothing and edge detection by nonlinear diffusion," *SIAM J. Numer. Anal.*, vol. 29, no. 1, pp. 182–193, 1992.
- [3] A. Chambolle, "Partial differential equations and image processing," in *Proc. IEEE ICIP*, vol. 1, 1994, pp. 16–20.
- [4] P. Cheeseman, B. Kanefsky, R. Kraft, J. Stutz, and R. Hanson, "Super-resolved surface reconstruction from multiple images," in *Maximum Entropy and Bayesian Methods*, G. R. Heidbreder, Ed. Norwell, MA: Kluwer, 1996, pp. 293–308.
- [5] G. H. Cottet and L. Germain, "Image processing through reaction combined with nonlinear diffusion," *Math. Comput.*, vol. 61, pp. 659–673, 1993.
- [6] R. Courant, K. O. Friedrichs, and H. Lewy, "On the partial difference equations of mathematical physics," *IBM J.*, vol. 11, pp. 215–235, 1967.
- [7] S. Di Zenzo, "A note on the gradient of a multi image," *Comput. Vis., Graph., Image Process.*, vol. 33, pp. 116–125, 1986.
- [8] M. Elad and A. Feuer, "Super-resolution restoration of continuous image sequence—Adaptive filtering approach," *IEEE Trans. Image Processing*, vol. 8, pp. 387–395, Mar. 1999.
- [9] D. Gabor, "Information theory in electron microscopy," *Lab. Invest.*, vol. 14, no. 6, pp. 801–807, 1965.
- [10] G. Gilboa, Y. Y. Zeevi, and N. Sochen, "Anisotropic selective inverse diffusion for signal enhancement in the presence of noise," in *Proc. IEEE ICASSP-2000*, vol. 1, Istanbul, Turkey, 2000, pp. 221–224.
- [11] —, "Signal and image enhancement by a generalized forward-and-backward adaptive diffusion process," in *Proc. EUSIPCO-2000*, Tampere, Finland, 2000.
- [12] —, "Resolution enhancement by forward-and-backward nonlinear diffusion processes," in *Proc. Nonlinear Signal and Image Processing*, Baltimore, MD, June 2001.
- [13] R. Kimmel, N. Sochen, and R. Malladi, "On the geometry of texture," Berkeley Labs., Univ. California, Berkeley, Rep. LBNL-39640, UC-405, Nov. 1996.
- [14] R. Kimmel, R. Malladi, and N. Sochen, "Images as embedding maps and minimal surfaces: Movies, color, texture, and volumetric medical images," *Int. J. Comput. Vis.*, vol. 39, no. 2, pp. 111–129, Sept. 2000.
- [15] R. Kimmel, N. Sochen, and R. Malladi, "From high energy physics to low level vision," in *First International Conference on Scale-Space Theory in Computer Vision*. New York: Springer-Verlag, 1997, Lecture Notes In Computer Science: 1252, pp. 236–247.
- [16] E. Kreyszing, *Differential Geometry*. New York: Dover, 1991.
- [17] X. Li and T. Chen, "Nonlinear diffusion with multiple edginess thresholds," *Pattern Recognit.*, vol. 27, no. 8, pp. 1029–1037, 1994.
- [18] M. Lindenbaum, M. Fischer, and A. Bruckstein, "On Gabor's contribution to image enhancement," *Pattern Recognit.*, vol. 27, no. 1, pp. 1–8, 1994.
- [19] S. J. Osher and L. I. Rudin, "Feature-oriented image enhancement using shock filters," in *SIAM J. Numer. Anal.*, vol. 27, 1990, pp. 919–940.
- [20] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, pp. 629–639, July 1990.
- [21] I. Pollak, A. S. Willsky, and H. Krim, "Scale space analysis by stabilized inverse diffusion equations," in *First International Conference on Scale-Space Theory in Computer Vision*. New York: Springer-Verlag, 1997, Lecture Notes In Computer Science: 1252, pp. 200–211.
- [22] A. M. Polyakov, "Quantum geometry of bosonic strings," *Phys. Lett.*, vol. 103B, pp. 207–210, 1981.
- [23] E. Radmoser, O. Scherzer, and J. Weickert, "Scale-space properties of nonstationary iterative regularization methods," *J. Vis. Commun. Image Represent.*, to be published.
- [24] G. Sapiro and D. L. Ringach, "Anisotropic diffusion of multivalued images with applications to color filtering," *IEEE Trans. Image Processing*, vol. 5, pp. 1582–1586, Nov. 1996.
- [25] R. R. Schultz and R. L. Stevenson, "Extraction of high-resolution frames from video sequences," *IEEE Trans. Image Processing*, vol. 5, pp. 996–1011, June 1996.
- [26] N. Sochen, G. Gilboa, and Y. Y. Zeevi, "Color image enhancement by a forward-and-backward adaptive Beltrami flow," in *AFPAC-2000, LNCS 1888*, G. Sommer and Y. Y. Zeevi, Eds. New York: Springer-Verlag, 2000, pp. 319–328.
- [27] N. Sochen, R. Kimmel, and R. Malladi, "A general framework for low level vision," *IEEE Trans. Image Processing*, vol. 7, pp. 310–318, Mar. 1998.
- [28] N. Sochen and Y. Y. Zeevi, "Images as manifolds embedded in a spatial-feature non-Euclidean space," Technion—Israel Inst. Technol., Haifa, CCIT Rep. 1181, Nov. 1998.
- [29] N. Sochen, "Stochastic processes in vision I: From Langevin to Beltrami," Technion—Israel Inst. Technol., Haifa, CCIT Rep. 245, June 1999.
- [30] N. Sochen and Y. Y. Zeevi, "Representation of colored images by manifolds embedded in higher dimensional non-Euclidean space," in *Proc. IEEE ICIP'98*, Chicago, IL, 1998.
- [31] B. M. ter Haar Romeny, Ed., *Geometry Driven Diffusion in Computer Vision*. Norwell, MA: Kluwer, 1994.
- [32] I. Vitsnudel, R. Ginosar, and Y. Y. Zeevi, "Neural network aided design for image processing," in *Proc. SPIE Symp. Visual Communication and Image Processing*, vol. 1606, pp. 1086–1091.
- [33] J. Weickert, "Coherence-enhancing diffusion of color images," *Image Vis. Comput.*, vol. 17, pp. 199–210, 1999.
- [34] J. Weickert, "Multiscale texture enhancement," in *Computer Analysis of Images and Patterns; Lecture Notes in Computer Science*. New York: Springer, 1995, vol. 970, pp. 230–237.
- [35] J. Weickert, "Anisotropic diffusion in image processing," Ph.D. dissertation, Kaiserslautern Univ., Germany, Nov. 1995.
- [36] J. Weickert, "Scale-space properties of nonlinear diffusion filtering with diffusion tensor," Lab. Technomathematics, Univ. Kaiserslautern, Germany, Tech. Rep. 110, 1994.
- [37] J. Weickert and B. Benhamouda, "A semidiscrete nonlinear scale-space theory and its relation to the Perona–Malik paradox," in *Advances in Computer Vision*, F. Solina, Ed. New York: Springer, 1997, pp. 1–10.
- [38] R. T. Whitaker and S. M. Pizer, "A multi-scale approach to non uniform diffusion," in *CVGIP: Image Understand.*, vol. 57, 1993, pp. 99–110.
- [39] A. P. Witkin, "Scale space filtering," in *Proc. Int. Joint Conf. Artificial Intelligence*, 1983, pp. 1019–1023.
- [40] N. Sochen, "Stochastic processes in vision I: From Langevin to Beltrami," in *Proc. Int. Conf. Computer Vision*, Vancouver, BC, Canada, July 2001, pp. 288–293.



Guy Gilboa received the B.Sc. degree in electrical engineering from Ben-Gurion University, Beer Sheva, Israel, in 1997. He is currently pursuing the Ph.D. degree in electrical engineering at the Technion—Israel Institute of Technology, Haifa.

He was with Intel Development Center, Haifa, for three years, working in the design of processors. His main research interests are related to variational and PDE-based processes applied to enhancement, sharpening, and denoising of images. He is also interested in computer vision algorithms and their applications to video surveillance and target recognition.



Nir Sochen received the B.Sc degree in physics in 1986 and the M.Sc. degree in theoretical physics 1988, both from the University of Tel-Aviv, Tel-Aviv, Israel. He received the Ph.D. degree in theoretical physics 1992 from the Université de Paris-Sud, Paris, France, while conducting his research in the Service de Physique Théorique, Centre d'Etude Nucleaire, Saclay, France.

He continued with one year of research in the Ecole Normale Supérieure, Paris, on the Haute Etude Scientifique Fellowship, and a three-year NSF fellowship in the Physics Department, University of California, Berkeley. It is there that his interest shifted from quantum field theories and integrable models, related to high-energy physics and string theory, to computer vision and image processing. He spent one year with the Physics Department, University of Tel-Aviv, and two years with the Faculty of Electrical Engineering, Technion—Israel Institute of Technology, Haifa. Since 1999, he has been a Senior Lecturer with the Department of Applied Mathematics, University of Tel-Aviv. He is also a member of the Ollendorf Center at the Technion. His main research interests are the applications of differential geometry and statistical physics in image processing and computational vision.



Yehoshua Y. (Josh) Zeevi received the Ph.D. degree from the University of California, Berkeley.

He is the Barbara and Norman Seiden Professor of Computer Science with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa. He is the Founder of the Jacobs Center for Communication and Information Technologies (CCIT) and the Ollendorff Minerva Center, and served as the Head of those centers. He was also the Dean of the Faculty of Electrical Engineering (1994-1999). He was a Visiting Scientist at Lawrence Berkeley Laboratory, University of California, Berkeley. He was a Vinton Hayes Fellow at Harvard University, Cambridge, MA, and has been a regular visitor there. He was also a Visiting Professor at the Massachusetts Institute of Technology, Cambridge, and Rutgers University, New Brunswick, NJ, and a Senior Visiting Scientist at the NTT Research Center, Yokosuka, Japan. He is presently a Visiting Professor at Columbia University, New York. He is the co-inventor of many patents and the author of over 200 publications related to vision and image sciences. He is the Editor-in-Chief of the *Journal of Visual Communication and Image Representation* and the editor of three books. He is one of the founders of i Sight, Inc., a company that developed digital video cameras that mimic the eye, and UltraGuide, Ltd., a medical technology company that develops guidance systems under ultrasound imaging.

Dr. Zeevi is a Fellow of SPIE and the Rodin Academy.