



Noise Reduction in Surface Reconstruction from a Given Gradient Field

B. KARAÇALI

*Department of Radiology, University of Pennsylvania, Section on Biomedical Image Analysis,
3600 Market Street Suite 380, Philadelphia, PA 19104, USA*

bilge@rad.upenn.edu

W. SNYDER

*Electrical and Computer Engineering Department, North Carolina State University,
1010 Main Campus Drive, Box 7911, Raleigh, NC 27695, USA*

wes@eos.ncsu.edu

Received November 20, 2002; Revised August 18, 2003; Accepted January 5, 2004

Abstract. We present a gradient space technique for noise reduction in surfaces reconstructed from a noisy gradient field. We first analyze the error sources in the recovered gradient field of a surface using a three-image photometric stereo method. Based on this analysis, we propose an additive noise model to describe the errors in the surface gradient estimates. We then use a vector space formulation and construct a multiscale orthonormal expansion for gradient fields. Using the sparse representation properties of this expansion, we develop techniques for reducing the gradient field noise by coefficient selection with thresholding. The simulation results indicate that the proposed technique provides significant improvement on the noise levels of both the estimated gradient fields and the reconstructed surfaces under heavy noise levels. Furthermore, the experiments using noisy photometric stereo image triplets of real range data suggest that the additive model remains viable after the nonlinear photometric stereo operation to provide accurate noise removal.

Keywords: surface reconstruction, gradient field, noise reduction, thresholding, projection, gradient space, orthonormal representation, multiscale analysis

1. Introduction

Reconstruction of a three-dimensional scene is a well studied problem in computer vision literature. Methods for reconstructing a three-dimensional surface topography from one or more brightness images are referred as “shape-from-shading” methods (Deift and Sylvester, 1981; Horn and Brooks, 1989; Hurt, 1991; Zhang et al., 1999). Shape-from-shading can be posed as the problem of finding a surface directly that would result in the observed images, or can be attempted in two steps: obtainment of a gradient field characterizing the surface normals at every pixel, and reconstruction of a surface that corresponds to this gradient field.

A common technique to infer the surface normal at a given pixel based on the observed brightness values is the “photometric stereo” method (Woodham, 1978; Ikeuchi, 1981; Hayakawa, 1994). Given three images $E_i(x, y)$, $i = 1, 2, 3$, taken from the same viewing angles, it solves the system of equations represented by

$$E_i(x, y) = R_i(p(x, y), q(x, y)), \quad i = 1, 2, 3 \quad (1)$$

for $p(x, y)$ and $q(x, y)$, for all (x, y) in the surface domain. The imaging conditions for each image are summarized by the reflectance maps $R_i(p, q)$, which determine the brightness value to be observed at the i 'th image, at pixels where the surface gradient is

equal to $[p \ q]^T$. Woodham (1978) shows that for non-degenerate $R_i(p, q)$, the surface gradient at a given pixel (x, y) may be found as the unique intersection of the three curves representing the system of Eq. (1) in the $p - q$ space. Tagare and de Figueiredo (1991) show that for a large class of non-Lambertian diffuse surfaces, three light sources are indeed sufficient for unique representation of the set of surface orientations. Kozera (1991) studies the existence and uniqueness of photometric stereo for Lambertian surfaces in two views.

Aside from the photometric stereo technique, there are other attempts to invert the reflectance map. Rajaram et al. (1995) train a feed-forward neural network over a set of brightness value triplets obtained from pictures of a sphere in order to generate a needle map, which they use to reconstruct the surface topography. Iwahori et al. (1994) construct two radial basis function neural networks to model the surface normal estimation and image generation processes. They train and use these networks over three images to estimate Phong (1975) and Lambertian (1760) reflectance model parameters. Later, the same approach is generalized to use multiple images obtained by light sources that are rotationally symmetrical around the viewing direction (Iwahori et al., 1997).

There are several different approaches to generate a surface corresponding to a given gradient field. If the field is not integrable, plain integration yields ambiguous results depending on the path of integration (Horn, 1986). One approach is to use variational calculus to minimize a cost function resulting in an integrable surface with a gradient field as close to the given field as possible (Ikeuchi and Horn, 1981; Horn and Brooks, 1986). Frankot and Chellappa (1988, 1989) show that in continuous settings, it is possible to characterize integrable gradient fields as a closed subspace in the space of all gradient fields. They first express the (band-limited) surface as a linear combination of exponential elementary surfaces as

$$z(x, y) = \sum_{f_x, f_y} \alpha(f_x, f_y) \exp(j(f_x x + f_y y)). \quad (2)$$

Differentiating both sides with respect to x and y , one obtains

$$\frac{\partial z}{\partial x}(x, y) = \sum_{f_x, f_y} j f_x \alpha(f_x, f_y) \exp(j(f_x x + f_y y)) \quad (3)$$

and

$$\frac{\partial z}{\partial y}(x, y) = \sum_{f_x, f_y} j f_y \alpha(f_x, f_y) \exp(j(f_x x + f_y y)). \quad (4)$$

It is then clear that the gradient field (p, q) corresponding to the surface z resides in a subspace of the space of all gradient fields, which is spanned by the set

$$\{(j f_x \exp(j(f_x x + f_y y)), j f_y \exp(j(f_x x + f_y y))\}.$$

Furthermore, since the only requirement on z is that it be integrable, the subspace spanned by this set of elementary gradient fields is indeed the space of all integrable gradient fields. Consequently, the integrable surface closest to a given (possibly non-integrable) gradient field can be computed by projecting the given gradient field onto the integrable subspace, and finding the corresponding surface using Fourier transform techniques. Later, other authors followed the same argument and developed similar techniques using wavelets (Hsieh et al., 1995) and Legendre polynomials (Kim and Park, 1997). Simchony et al. (1990) arrive at a similar solution by considering the Laplacian that corresponds to the non-integrable gradient field using direct Poisson methods (Buzbee et al., 1970). Noakes et al. (1999) and Noakes and Kozera (1999), formulate a least-squares approach to solve a linear system of equations to calculate the integrable gradient field closest to a given, possibly non-integrable one. The optimization, however, fails for images of size 32×32 or larger, and block processing techniques are implemented that eventually converge to the global optimum (Noakes and Kozera, 2001).

Recently, Karaçalı and Snyder (2002) formulated an adaptive approach to reconstruct surfaces that correspond to non-integrable gradient fields, such as surfaces containing edges or occluding boundaries. Using discrete approximations to surface gradients, they construct an orthonormal set of gradient fields spanning a *feasible* subspace of the gradient space, in which the gradient fields of surfaces with (adaptively detected) edges reside.

In this paper, we investigate noise reduction capabilities of the technique proposed by Karaçalı (2002) and Karaçalı and Snyder (2002). In general settings, noise reduction in photometric stereo refers to finding

a surface that produces brightness images as close to the given photometric stereo triplet as possible. Noakes and Kozerá (2003b) follow this strategy by iteratively optimizing an energy function that penalizes the deviation of the produced brightness images from the given ones. This entails solving a non-quadratic optimization problem in many variables. We follow a different approach that solves the problem in two steps. We first consider gradient fields obtained by solving the quadratic photometric stereo problem. We then employ another quadratic approach to eliminate the noise in the gradient field obtained as such.

We begin by analyzing the sparsity properties of the multi-scale gradient field representation. Next, we implement a gradient field noise removal method by means of coefficient trimming on the corresponding orthonormal gradient field expansion, using an additive noise model. Using experiments on both synthetically generated gradient fields and photometric stereo image triples from real range data, we show that the proposed method achieves high levels of noise reduction under heavy noise.

The next section is devoted to a brief overview of this method. In Section 3, we analyze the sources of noise and distortion in the gradient field obtained using three-image photometric stereo. Section 4 presents the results of the proposed noise reduction technique on synthetic and real data.

2. Overview

A surface z over a set of (uniformly sampled with unit distance) grid points $[1, 2, \dots, m] \times [1, 2, \dots, n]$ is represented as one element of the set

$$\mathcal{Z} = \{z \mid z : [1, 2, \dots, m] \times [1, 2, \dots, n] \rightarrow \mathbb{R}\}.$$

Define the inner product $\langle \cdot, \cdot \rangle_{\mathcal{Z}}$ on \mathcal{Z} by

$$\langle z_1, z_2 \rangle_{\mathcal{Z}} = \sum_x \sum_y z_1(x, y) z_2(x, y). \quad (5)$$

The discrete surface gradient may be given by the finite difference approximation to partial derivatives

$$\nabla z = [p \ q]^T$$

where p and q may be derived by an linear estimation of derivatives. For example,

$$p(x, y) = \begin{cases} z(x+1, y) - z(x, y) & \text{if } x = 1 \\ z(x, y) - z(x-1, y) & \text{if } x = m \\ (z(x+1, y) - z(x-1, y))/2 & \text{otherwise} \end{cases} \quad (6)$$

and

$$q(x, y) = \begin{cases} z(x, y+1) - z(x, y) & \text{if } y = 1 \\ z(x, y) - z(x, y-1) & \text{if } y = n \\ (z(x, y+1) - z(x, y-1))/2 & \text{otherwise} \end{cases} \quad (7)$$

The coupled pair (p, q) is then referred to as a *gradient field*. Let the inner product $\langle \cdot, \cdot \rangle_{\mathcal{P} \times \mathcal{Q}}$ on the gradient space $\mathcal{P} \times \mathcal{Q}$ be defined by

$$\langle (p, q)_1, (p, q)_2 \rangle_{\mathcal{P} \times \mathcal{Q}} = \sum_x \sum_y p_1(x, y) p_2(x, y) + q_1(x, y) q_2(x, y) \quad (8)$$

Note that the gradient field (p, q) obtained from a surface z by Eqs. (6) and (7) can be considered as the image of z under a linear operator $O : \mathcal{Z} \rightarrow \mathcal{P} \times \mathcal{Q}$ (Karaçalı, 2002).

Let $z_i, i = 1, \dots, mn$, be the Kronecker delta functions defined over the grid $[1, \dots, m] \times [1, \dots, n]$, i.e.,

$$z_i(x, y) = \begin{cases} 1 & \text{if } (x-1)n + y = i \\ 0 & \text{otherwise} \end{cases}, \quad (9)$$

$x \in [1, \dots, m], y \in [1, \dots, n]$. Clearly, the elements of the set $\{z_i\}$ span \mathcal{Z} , which can be shown to be a Hilbert space with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{Z}}$. Define the gradient fields $(p, q)_i \triangleq O(z_i)$ for $i = 1, \dots, mn$. Since $\mathcal{P} \times \mathcal{Q}$ can also be shown to be a Hilbert space with $\langle \cdot, \cdot \rangle_{\mathcal{P} \times \mathcal{Q}}$, the set of gradient fields span a $mn - 1$ -dimensional subspace S of $\mathcal{P} \times \mathcal{Q}$, referred as the *feasible subspace* (Karaçalı, 2002).

Let $(\bar{p}, \bar{q})_j, j = 1, \dots, mn - 1$, be the set of gradient fields obtained by applying the Gram-Schmidt orthonormalization procedure on the $(p, q)_i$. Now, suppose a gradient field (p, q) is given. If $(p, q) \in S$, then

the surface \hat{z} defined as

$$\hat{z} = \sum_{j=1}^k \langle (p, q), \overline{(p, q)}_j \rangle_{\mathcal{P} \times \mathcal{Q}} \bar{z}_j. \quad (10)$$

solves the equation $O(z) = (p, q)$ for z . Moreover, this solution is unique up to an equivalence class of surfaces which have the same image under the operator O . If $(p, q) \notin S$, the gradient field $\overline{(p, q)}$ given by

$$\overline{(p, q)} = \sum_{j=1}^k \langle (p, q), \overline{(p, q)}_j \rangle_{\mathcal{P} \times \mathcal{Q}} \overline{(p, q)}_j \quad (11)$$

is, by construction, the projection of (p, q) onto S , and thus, \hat{z} is a surface with gradient field closest to (p, q) in the sense of the gradient space norm induced by $\langle \cdot, \cdot \rangle_{\mathcal{P} \times \mathcal{Q}}$.

Minimum norm inversion of linear transformations by means of projections as such is a well known technique, which is also mentioned in Noakes and Kozera (2001). Instead of a direct approach as presented above, however, the authors propose solving a corresponding large system of linear equations using an iterative scheme.

Karaçali and Snyder (2002) show, in addition, that it is possible to characterize the subspace S in such a way to allow discontinuities in the reconstructed surface by modifying the operator O and the definitions of p and q around the discontinuities. Using proper adjustments to the operator O which reflects to the sets $\{(p, q)_j\}$ and $\{\bar{z}_j\}$, they achieve adequate reconstruction of surfaces that correspond to non-integrable gradient fields. In the present paper, however, we will assume that the true gradient fields are integrable, keeping in mind that the results also apply to the partially integrable case.

3. Error Sources in Photometric Stereo

The error in the reconstructed surface is an aggregation of misrepresentations and misdetections along the process of imaging, photometric stereo inversion of estimated reflectance maps, and surface reconstruction based on the obtained gradient field. Suppose the true surface is continuous except across discontinuity such as a step edge or an occluding boundary. Furthermore, suppose that each pixel on the surface is obtained under the same viewing direction but different lighting yielding non-degenerate three-image photometric

stereo. The steps of the reconstruction process and corresponding error sources can then be summarized as follows.

- *Imaging*: This step represents the acquisition of the three brightness images of the unknown surface. The main error source is the measurement error and the internal thermal noise of the imaging device. Both these errors are often aggregated in one additive Gaussian white noise component, while a more accurate model for the former is a signal-dependent Poisson distribution (Jain, 1989).

Another source of error involved in the imaging process is the orthographic projection assumption, postulating that the detected light rays from each pixel were traveling in parallel trajectories on their way to the imaging device. The true nature of image acquisition, however, is characterized by the perspective projection which manifests itself in geometric distortions (Grossberg and Nayar, 2001). As a result of these distortions, not only the general perceived geometry of the surface changes, but also the invariant viewing direction assumption becomes invalid.

- *Three-image photometric stereo*: The main component of error in the application of photometric stereo originates from the reflectance maps. In theory, reflectance maps can be approximated to have a particular closed form expression (Woodham, 1978; Horn, 1986), but real surface reflectances deviate from such approximations quite often (Horn, 1986). The reflectance maps are therefore constructed as look-up tables, specifying brightness values on a finite discrete grid in the $p - q$ plane, introducing sampling errors to the surface gradient representations.

The estimation of reflectance maps is typically done by imaging objects of known shapes, such as spheres and cylinders, whose surface reflectivity properties match the surface of interest (Jones and Taylor, 1994). This introduces additional errors directly into the values stored in the reflectance maps, through the involved imaging, gradient matching and interpolation steps of the estimation process.

Another error component involved with the use of reflectance maps is the implied assumptions. The reflectance maps characterize the dependence of the surface brightness on the surface normal, by keeping all other dependences fixed. This corresponds to globally (over the surface domain) invariant viewing direction (hence, orthographic projection), albedo

and illumination strength (in terms of distances to and incidence angles of the light sources).

- *Surface reconstruction*: The surface reconstruction technique of Karaçalı and Snyder (2002) summarized in Section 2 represents the partial derivatives of the surface by the discrete approximations in Eqs. (6) and (7). The accuracy of these approximations, however, downgrades at regions where the surface exhibits rapid changes. In fact, to say that the first partial derivative may be represented by these approximations is the same as saying the surface is of order no more than two. This suggests adding yet another non-stationary noise component to the gradient field, based on how the surface reconstruction is done.

Geometric distortions due to perspective projection can be alleviated to a degree by applying various orthorectification techniques frequently used in remote sensing (Shevlin, 1994; Zheng et al., 1997; Kiang, 1997). Surface reconstruction techniques that attempt to account for the effects of perspective projection have also been suggested (Lee and Kuo, 1994; Galo and Tozzi, 1996). In general, however, the orthographic projection assumption is sufficiently accurate when the observation distance is a few orders larger than the surface size.

3.1. Effect of Noise on Photometric Stereo

Suppose the orthographic projection reflectance map assumptions hold. An additive noise model for the photometric stereo image irradiance equation in Eq. (1) becomes

$$E_i(x, y) + E_i^n(x, y) = R_i(p(x, y), q(x, y)) + R_i^n(p(x, y), q(x, y)) \quad (12)$$

where $E_i^n(x, y)$ represents the image noise and $R_i^n(p(x, y), q(x, y))$ represents the estimation and quantization errors in the reflectance maps at pixels (x, y) and gradient $(p(x, y), q(x, y))$ respectively. Consider pixel (x, y) . Let

$$C_i = \{(\phi, \psi) \mid \tilde{E}_i(x, y) = \tilde{R}_i(\phi, \psi)\} \quad \text{for } i = 1, 2, 3, \quad (13)$$

where

$$\tilde{E}_i(x, y) = E_i(x, y) + E_i^n(x, y) \quad (14)$$

and

$$\tilde{R}_i(p(x, y), q(x, y)) = R_i(p(x, y), q(x, y)) + R_i^n(p(x, y), q(x, y)). \quad (15)$$

In ideal three-image photometric stereo, the curves C_1 , C_2 and C_3 intersect at exactly one point in the $p - q$ space (for a general class of reflectance maps (Tagare and deFigueiredo, 1991)), which is the solution of (1) and is the surface gradient at pixel (x, y) . In the general noisy case, however, this is not necessarily true. The noise terms $E_i^n(x, y)$ and $R_i^n(p(x, y), q(x, y))$ cause the curves C_1 , C_2 and C_3 to change shape and position. As a result, the system in Eq. (12) may not have a solution at all. Several illustrations of the system in Eq. (12) for various cases of noise effects are displayed in Fig. 1.

One way of processing the noisy curves that do not intersect is to introduce slack variables δ_i to the system of equations in Eq. (12)

$$\tilde{E}_i(x, y) + \delta_i = \tilde{R}_i(p(x, y), q(x, y)) \quad (16)$$

and to minimize a cost function penalizing large values of δ_i (in absolute value) such as

$$H(\delta_1, \delta_2, \delta_3) = \frac{1}{2} \sum_{i=1}^3 \delta_i^2. \quad (17)$$

Geometrically, this corresponds to enlarging or shrinking C_1 , C_2 and C_3 according to their respective reflectance maps so that they intersect at a unique point. Since the curves do not possess closed form expressions (the reflectance maps are determined empirically), this optimization cannot be carried out using functional optimization tools such as the calculus of variations (Courant and Hilbert, 1962). Various search techniques such as line search and random search can, however, be implemented in the slack parameters δ_i to obtain acceptable solutions.

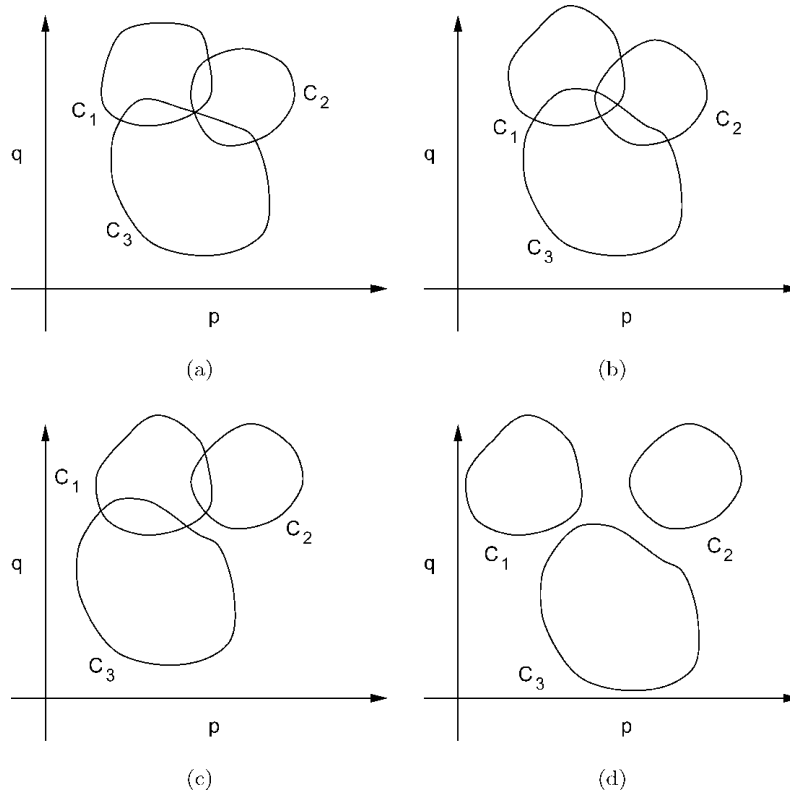


Figure 1. (a) Curves C_1 , C_2 and C_3 in ideal three-image photometric stereo, (b), (c), (d) examples of noisy cases.

4. Noise Reduction Using Sparse Gradient Field Representations

The formulation in the discussion above suggests representing the overall noise in the estimated surface gradients by an additive component such as

$$(p(x, y), q(x, y)) = (p^t(x, y), q^t(x, y)) + (p^n(x, y), q^n(x, y)) \quad (18)$$

where $(p^t(x, y), q^t(x, y))$ denotes the true surface gradient, $(p^n(x, y), q^n(x, y))$ the additive noise, and $(p(x, y), q(x, y))$ the estimated gradient at the pixel (x, y) . Thus, the gradient field becomes

$$(p, q) = (p, q)^t + (p, q)^n \quad (19)$$

with (p, q) , $(p, q)^t$ and $(p, q)^n$ denoting the observed, true and noise gradient fields respectively.

In reality, it is well known that the photometric stereo method to recover surface gradients from the brightness

image triplets is a nonlinear operation which breaks the additive and white nature of the brightness image noise (Noakes and Kozera, 2003b). In fact, after the photometric stereo operation, the gradient field noise not only changes distribution, but also becomes signal dependent: regions of the surface with larger gradients suffer more than relatively flat regions from the modified gradient noise. The illumination of the unknown surface determined by the source directions also play an important role on the characteristics of the resulting gradient noise. We adopt such an additive noise model for essentially three reasons. First, without any prior information on the surface, the signal dependent characteristics of the gradient noise cannot be inferred. White Gaussian gradient noise assumption thus becomes reasonable since it maximizes the entropy. Secondly, this particular model allows incorporating well-known noise reduction techniques from the signal recovery literature into the problem of gradient field denoising in surface reconstruction. Finally, the experiments we have conducted on photometric stereo image triplets of real range data corrupted

with heavy noise indicate that the additive model in Eq. (18) remains valid to a large degree so that significant noise reductions can be achieved by such methods.

Karaçalı and Snyder (2002) develop an orthonormal multiscale representation for the gradient fields in the feasible subspace and use it to do adaptive edge detection. The wavelet-based representation they propose is of the form

$$(p, q)^S = \sum_{j=1}^k w_j \overline{(p, q)}_j \quad (20)$$

with

$$w_j = \langle (p, q), \overline{(p, q)}_j \rangle_{\mathcal{P} \times \mathcal{Q}} \quad (21)$$

so that $(p, q)^S$ is the projection of (p, q) onto the feasible subspace S spanned by the gradient fields $\{\overline{(p, q)}_j\}$, $j = 1, \dots, k$. The gradient fields $\{\overline{(p, q)}_j\}$ have the additional property that the information content of $\overline{(p, q)}_{j_1}$ is at a strictly higher scale than $\overline{(p, q)}_{j_2}$ if $j_1 > j_2$.

Multiscale orthonormal transformations such as wavelet decomposition are known to provide sparsity in signal representations (Mallat, 1989, 1998). Sparse signal representations where a large portion of signal energy is concentrated on a small number of terms have been widely studied in noise reduction applications (Mallat, 1998; Donoho and Johnstone, 1994; Krim et al., 1995; Karaçalı, 1999; Chang et al., 2000). We follow the same philosophy and exploit the sparsity properties of the representation in Eq. (21) for noise reduction in the gradient fields (p, q) through the use of coefficient thresholding.

Thresholding on the representation coefficients is a tool commonly used in wavelet analysis for coefficient selection in signal denoising, particularly using wavelet decomposition (Donoho and Johnstone, 1994; Krim et al., 1995; Karaçalı, 1999; Chang et al., 2000). The main rationale behind coefficient trimming with respect to a (previously or adaptively determined) threshold is the uncorrelated nature of the independent and identically distributed (*iid*) additive noise in the sparse signal representation. For simplicity, suppose that the gradient field $(p, q) \in S$. Evaluating Eq. (21) on the noise

model in Eq. (19), we obtain

$$(p, q) = \sum_{j=1}^k \langle (p, q), \overline{(p, q)}_j \rangle_{\mathcal{P} \times \mathcal{Q}} \overline{(p, q)}_j \quad (22)$$

$$= \sum_{j=1}^k \langle ((p, q)^t + (p, q)^n), \overline{(p, q)}_j \rangle_{\mathcal{P} \times \mathcal{Q}} \overline{(p, q)}_j \quad (23)$$

$$= \sum_{j=1}^k \langle ((p, q)^t, \overline{(p, q)}_j) \rangle_{\mathcal{P} \times \mathcal{Q}} \overline{(p, q)}_j + \langle (p, q)^n, \overline{(p, q)}_j \rangle_{\mathcal{P} \times \mathcal{Q}} \overline{(p, q)}_j. \quad (24)$$

Defining w_j^t , and w_j^n as

$$w_j^t = \langle (\widehat{(p, q)})^t, \overline{(p, q)}_j \rangle_{\mathcal{P} \times \mathcal{Q}} \quad (25)$$

$$w_j^n = \langle (p, q)^n, \overline{(p, q)}_j \rangle_{\mathcal{P} \times \mathcal{Q}} \quad (26)$$

and noting the orthonormality of $\{\overline{(p, q)}_j\}$, Eq. (19) becomes equivalent to

$$w_j = w_j^t + w_j^n \quad (27)$$

for all $j = 1, \dots, k$. Furthermore, if the additive noise term $(p, q)^n$ is assumed to be *iid* Gaussian with zero mean and variance σ^2 , by the orthonormality of the expansion in Eq. (20), the coefficients w_j^n become *iid* with zero mean and variance σ^2 .

The thresholding estimate $\widehat{w} = [\widehat{w}_1 \widehat{w}_2 \dots \widehat{w}_k]$ of the true set of coefficients w^t with respect to a threshold T is defined as

$$\widehat{w}_j = \begin{cases} w_j & \text{if } |w_j| > T \\ 0 & \text{otherwise} \end{cases} \quad (28)$$

The risk, R , the error in the recovered gradient field $(\widehat{p, q})$, is

$$R = \|(p, q)^t - (\widehat{p, q})\|_{\mathcal{P} \times \mathcal{Q}}^2 = \sum_{j=1}^k (w_j^t - \widehat{w}_j)^2 \quad (29)$$

noting the orthonormality of $\{\overline{(p, q)}_j\}$ and that $(p, q) \in S$, where

$$(\widehat{p, q}) = \sum_{j=1}^k \widehat{w}_j \overline{(p, q)}_j. \quad (30)$$

The ideal coefficient selection for the optimal (in the mean squared error sense) estimate \hat{w}^{id} of w^t is given by

$$\hat{w}_j^{id} = \begin{cases} w_j & \text{if } |w_j^t| > \sigma \\ 0 & \text{otherwise} \end{cases} \quad (31)$$

with minimal error

$$R^{id} = \sum_{j=1}^k \min \{|w_j^t|, \sigma\}^2. \quad (32)$$

Note that this is not a thresholding method, since there may not be a threshold for which the Eq. (28) would produce \hat{w}^{id} . In fact, \hat{w}^{id} is not computable since whether $|w_j^t| > \sigma$ holds or not is not known for any j . We will, however, use \hat{w}^{id} in our simulations as a benchmark estimate to compare with the other estimated gradient fields and corresponding surface reconstructions.

The variants of basic thresholding method in Eq. (28) in noise reduction applications are numerous (Chang et al., 2000). We will, however, confine ourselves with the basic method and study two types of threshold selection schemes.

The first threshold selection approach is originally proposed by Donoho and Johnstone (1994) in the context of wavelet shrinkage, and is based on the following limiting property of a zero mean Gaussian distribution. Let N be a random variable with distribution $\mathcal{N}(0, \sigma^2)$. Let $N(k)$ denote a total of K realizations of N , $k = 1, \dots, K$. It can be shown (Donoho and Johnstone, 1994; Krim et al., 1995) that

$$\lim_{K \rightarrow \infty} \Pr \left\{ \max_{k=1, \dots, K} |N(k)| < \sigma \sqrt{2 \log_2 K} \right\} = 1 \quad (33)$$

This threshold can be computed separately for each scale $s = 1, \dots, J$ where the image to be reconstructed is of size $2^J \times 2^J$. Since $(\overline{p, q})_j$ corresponds to scale $\lfloor (\log_2 j)/2 \rfloor + 1$, and since there are $3 \cdot 2^{2(s-1)}$ coefficients at scale s , the resulting thresholding scheme is therefore given by

$$\hat{w}_j^{ws} = \begin{cases} w_j & \text{if } |w_j| > T^{ws} \\ 0 & \text{otherwise} \end{cases} \quad (34)$$

where

$$\begin{aligned} T^{ws} &= \sigma \sqrt{2 \log_2 (3 \cdot 2^{2 \lfloor (\log_2 j)/2 \rfloor})} \\ &= \sigma \sqrt{2(\log_2 3 + 2 \lfloor (\log_2 j)/2 \rfloor)} \end{aligned} \quad (35)$$

and σ is the standard deviation of w_j^n . We will refer to this particular choice of threshold as *wavelet* thresholding because of its origins in wavelet analysis (Donoho and Johnstone, 1994).

We obtain another thresholding rule using results from Extreme Value Theory (EVT) (Embrechts et al., 1997) which studies the statistical properties of extremal events. Let $P(w)$ denote the probability distribution of w_j^n , i.e.,

$$P(w) = \int_{-\infty}^w \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{w^2}{2\sigma^2}}. \quad (36)$$

Since w_j^n are *iid*, the probability $\Pr\{\max_{j=1, \dots, k}(w_j^n) \leq w\}$ is given by

$$\begin{aligned} &\Pr \left\{ \max_{j=1, \dots, k} (w_j^n) \leq w \right\} \\ &= \Pr \{w_1^n \leq w, \dots, w_k^n \leq w\} \\ &= \Pr \{w_1^n \leq w\} \dots \Pr \{w_k^n \leq w\} \\ &= \Pr \{w_1^n \leq w\}^k \\ &= (P(w))^k. \end{aligned} \quad (37)$$

Similarly,

$$\begin{aligned} &\Pr \left\{ \min_{j=1, \dots, k} (w_j^n) \leq w \right\} \\ &= \Pr \left\{ -\max_{j=1, \dots, k} (-w_j^n) \leq w \right\} \\ &= \Pr \left\{ \max_{j=1, \dots, k} (-w_j^n) \geq -w \right\} \\ &= 1 - \Pr \left\{ \max_{j=1, \dots, k} (-w_j^n) \leq -w \right\} \\ &= 1 - \Pr \{-w_1^n \leq -w\} \dots \Pr \{-w_k^n \leq -w\} \\ &= 1 - \Pr \{-w_1^n \leq -w\}^k \\ &= 1 - (P(-w))^k \\ &= 1 - (1 - P(w))^k. \end{aligned} \quad (38)$$

Given a threshold T , the probability of all noise coefficients w_j^n to be in the interval $[-T, T]$ is therefore

$$\begin{aligned} \Pr \{|w_j^n| \leq T\} &= \Pr \left\{ \max_{j=1, \dots, k} (w_j^n) \leq T \right\} \\ &\quad \cdot \Pr \left\{ \min_{j=1, \dots, k} (w_j^n) \geq -T \right\} \end{aligned}$$

$$\begin{aligned}
 &= (P(T))^k(1 - P(-T))^k \\
 &= (P(T))^{2k}. \tag{39}
 \end{aligned}$$

Let $\alpha \in [0, 1]$ be such that $\Pr\{|w_j^n| \leq T\} = 1 - \alpha$. Solving Eq. (39) for T , we obtain

$$\begin{aligned}
 1 - \alpha &= (P(T))^{2k} \\
 \sqrt[2k]{1 - \alpha} &= P(T) \\
 T &= P^{-1}(\sqrt[2k]{1 - \alpha}). \tag{40}
 \end{aligned}$$

Setting α to

$$\alpha = \frac{1}{3 \cdot 2^{2(s-1)}}, \tag{41}$$

yields the thresholding rule

$$\hat{w}_j^{\text{evt}} = \begin{cases} w_j & \text{if } |w_j| > T^{\text{evt}} \\ 0 & \text{otherwise} \end{cases} \tag{42}$$

where

$$T^{\text{evt}} = P^{-1}((1 - \alpha)^{(2^{-2\lfloor \log_2 j \rfloor / 2})/6}) \tag{43}$$

We will refer to this thresholding rule as *EVT* thresholding. The particular choice of α in Eq. (41) reflects the allowance of roughly one noise coefficient into the estimate at every scale. Note that the computation of P^{-1} can be done numerically to an arbitrary level of precision.

5. Simulation Results

5.1. Simulation Results on Synthetic Data

We have generated a synthetic surface, and computed reconstructions using the techniques discussed in the previous section on the corresponding gradient fields corrupted by additive noise at various levels. The generated surface bears the analytic form

$$\begin{aligned}
 &\frac{15}{15.5947} \left(2 - \cos\left(\frac{2\pi(x-1)}{31}\right) \right) \\
 &- \cos\left(\frac{2\pi 3(x-1)}{31}\right) \left(2 - \cos\left(\frac{2\pi(y-1)}{31}\right) \right) \\
 &- \cos\left(\frac{2\pi 3(y-1)}{31}\right)
 \end{aligned}$$




Figure 2. True surface.

where $x, y \in \{1, 2, \dots, 32\}$. The true surface and its gradient field are shown in Figs. 2 and 3. The zero-mean white Gaussian gradient space noise at signal-to-noise (SNR) levels 20, 10 and 0 dB are added to the gradient field prior to noise reduction techniques followed by reconstruction of a surface. Simulation results consist of four separate surface reconstructions:

- *plain reconstruction*: Surface reconstruction with no noise reduction; the surface that corresponds to the projection of (p, q) onto the feasible subspace S
- *ideal reconstruction*: Surface reconstruction using the coefficients w_j over which the true gradients have more power than the additive noise
- *wavelet thresholding reconstruction*: Surface reconstruction following thresholding of w_j by the rule in Eq. (34)
- *EVT thresholding reconstruction*: Surface reconstruction following thresholding of w_j by the rule in Eq. (42)

We have generated the sets $\{\overline{(p, q)}_j\}$ and $\{\bar{z}_j\}$ providing a multiscale representation of the feasible subspace from a Coiflet basis for the surface space \mathcal{Z} (Daubechies, 1988). In the simulations below, we assume the noise variance is known. In a more general setting, it can be estimated given a noisy gradient field. For the purposes of the analysis presented here, the absolute accuracy of this estimate is not relevant, since the main objective is to study the relative performance of the proposed techniques. The individual noise performances will evidently be dependent on the accuracy of the noise variance estimate.




Figure 3. Gradient field of the surface in Fig. 2: (a) p and (b) q .

The resulting surfaces for 20 dB SNR are shown in Fig. 4. Figures 5 and 6 show the reconstructions at noise levels 10 and 0 dB SNR respectively. The coefficients \hat{w}_j that produce these reconstructions are displayed in Figs. 7–9.

We quantify the performance of these reconstructions in both the gradient space and the surface space by measuring the mean squared errors and the corresponding achieved SNR levels at the reconstructions. The error analysis of the reconstruction results shown in Table 1 indicate that all methods yield significant improvements on the noise level in both the gradient field and the reconstructed surface. The wavelet and EVT thresholding schemes consistently provide over 5 dB of improvement in the gradient space, and even challenge the statistically optimal ideal coefficient selection result at 20 dB SNR case. Additionally, the EVT-based thresholding rule consistently performs better than the wavelet-based rule in all cases. The coefficient plots in Figs. 7–9 demonstrate that the EVT-based rule achieves this performance by including additional valuable coefficients to the reconstruction that are shrunk by the wavelet-based threshold. This suggests that the threshold in Eq. (35) is superfluous, and the threshold in Eq. (43) provides a more appropriate alternative in the case of additive white Gaussian noise. Moreover, while the wavelet-based threshold explicitly relies on the Gaussian noise assumption, the EVT-based threshold can be computed for different types of noise distributions such as Laplacian or Poisson, using appropriate distribution functions.

Note that the plain surface reconstruction algorithm with no additional noise processing produces an improvement of approximately 3 dB in the reconstructed surface gradient field. The projection of (p, q) from the $2mn$ -dimensional gradient space onto the $(mn - 1)$ -dimensional feasible subspace effectively reduces the noise energy by a half, producing roughly a 3 dB increase in the resulting SNR.

Table 1. Error statistics of noise reduction techniques.

| Method | Surface error | | Gradient field error | |
|-----------|---------------|----------|----------------------|----------|
| | MSE | SNR (dB) | MSE | SNR (dB) |
| 20 dB SNR | | | | |
| Plain | 0.0165 | 27.1990 | 0.0039 | 23.2124 |
| Ideal | 0.0081 | 30.3248 | 0.0013 | 27.9142 |
| Wavelet | 0.0095 | 29.6215 | 0.0022 | 25.6371 |
| EVT | 0.0095 | 29.6140 | 0.0020 | 26.1629 |
| 10 dB SNR | | | | |
| Plain | 0.1415 | 17.8751 | 0.0408 | 12.9653 |
| Ideal | 0.0260 | 25.2407 | 0.0070 | 20.5924 |
| Wavelet | 0.0693 | 20.9755 | 0.0244 | 15.2080 |
| EVT | 0.0552 | 21.9666 | 0.0174 | 16.6762 |
| 0 dB SNR | | | | |
| Plain | 1.6756 | 7.1422 | 0.4081 | 2.9665 |
| Ideal | 0.2333 | 15.7051 | 0.0521 | 11.9014 |
| Wavelet | 0.8582 | 10.0482 | 0.1651 | 6.8974 |
| EVT | 0.6699 | 11.1236 | 0.1366 | 7.7177 |

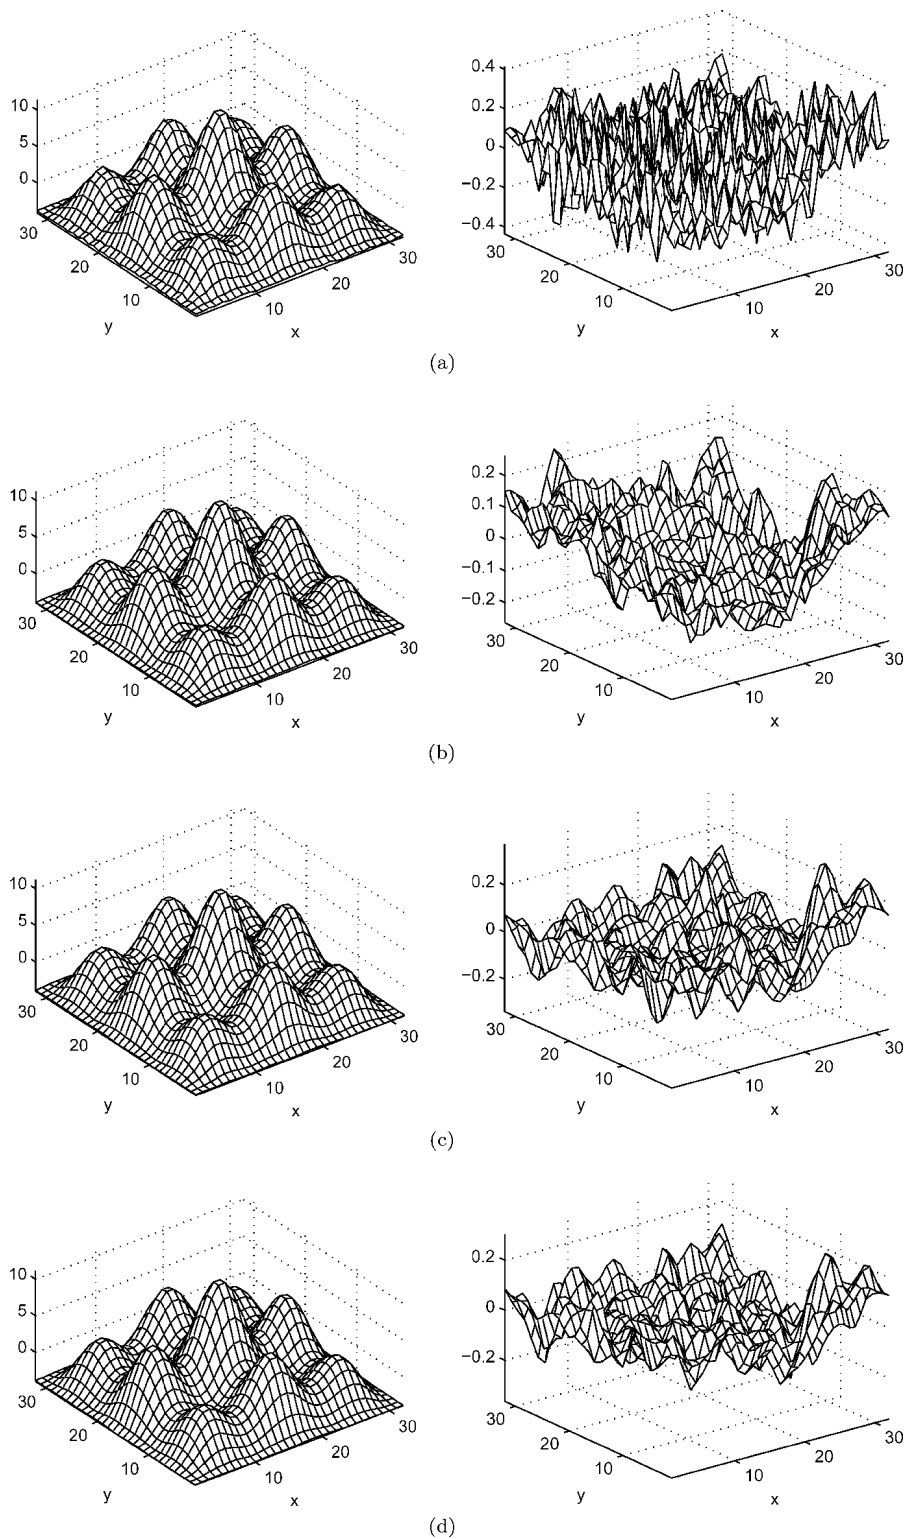


Figure 4. Surface reconstruction at 20 dB SNR, reconstructed surfaces and the error surfaces; (a) plain, (b) ideal, (c) wavelet thresholding, and (d) EVT thresholding.

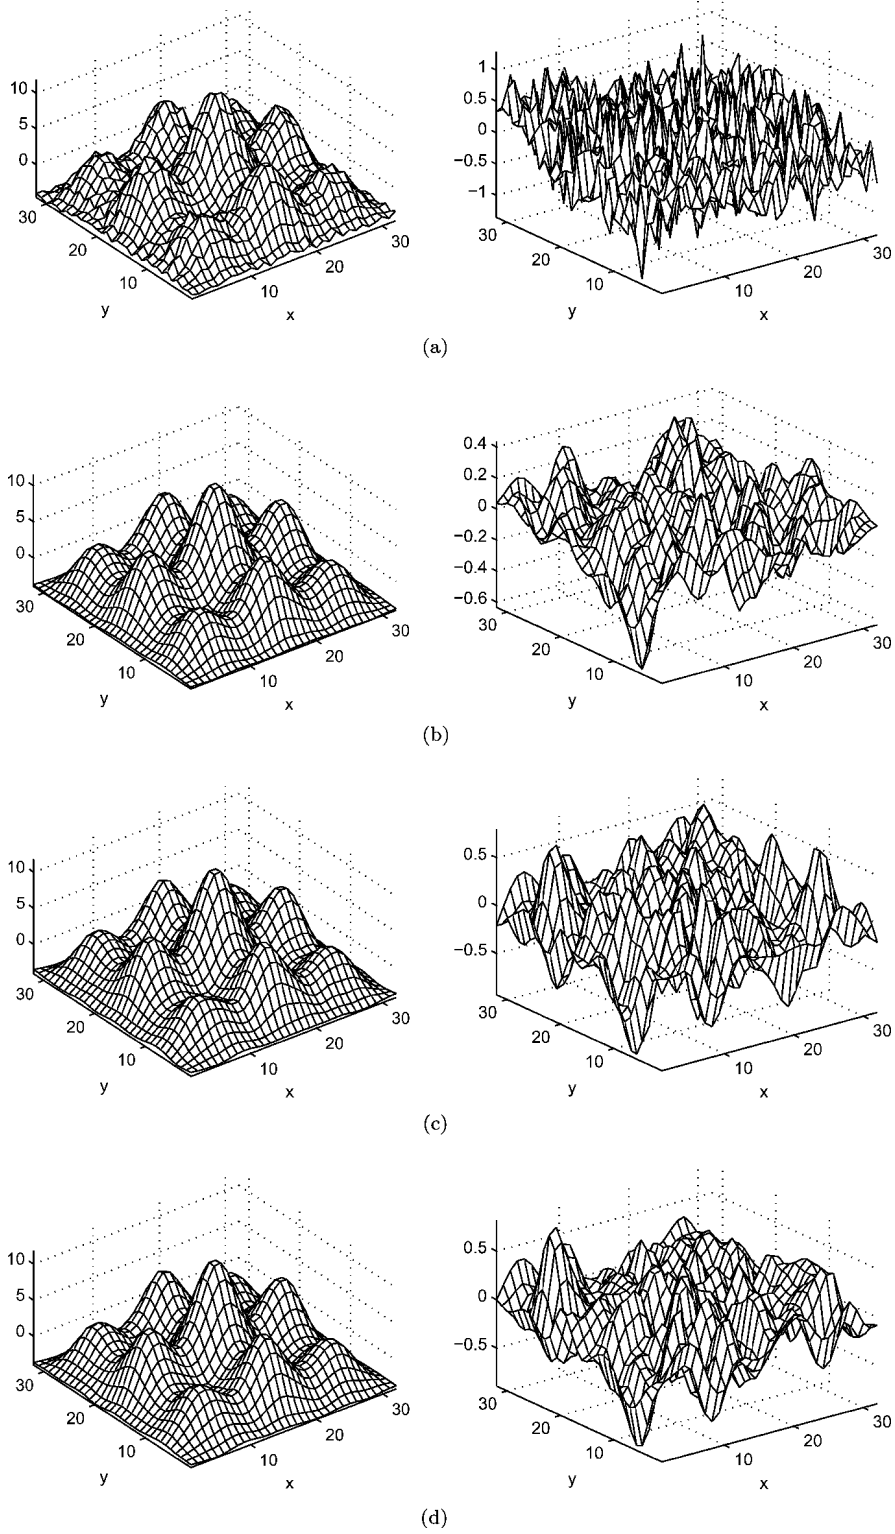


Figure 5. Surface reconstruction at 10 dB SNR, reconstructed surfaces and the error surfaces; (a) plain, (b) ideal, (c) wavelet thresholding, and (d) EVT thresholding.

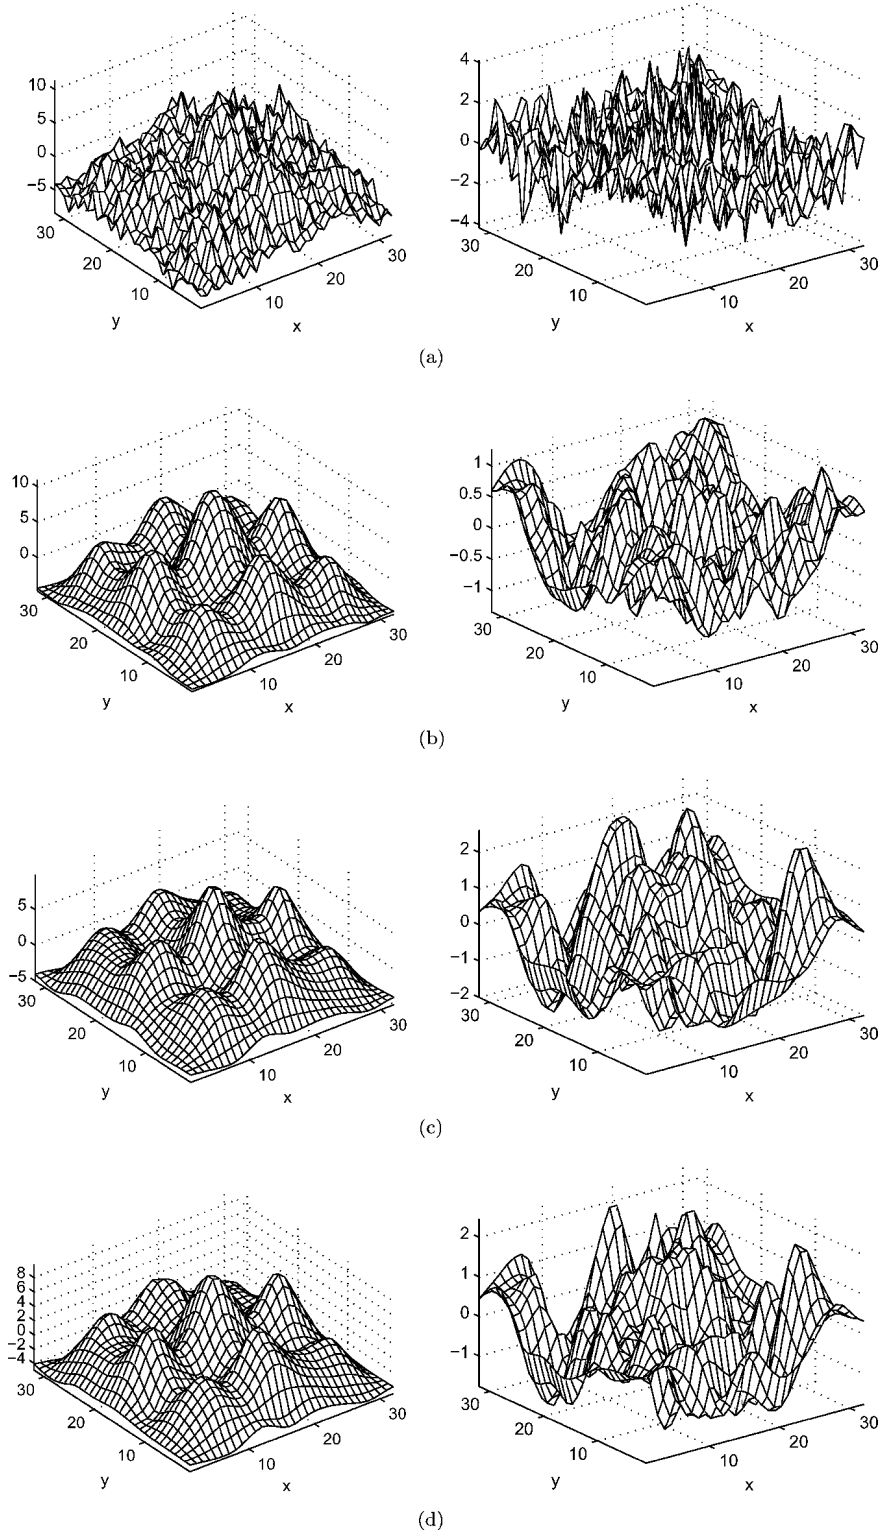


Figure 6. Surface reconstruction at 0 dB SNR, reconstructed surfaces and the error surfaces; (a) plain, (b) ideal, (c) wavelet thresholding, and (d) EVT thresholding.




Figure 7. Coefficients \hat{w}_j at 20 dB SNR, (a) plain, (b) ideal, (c) wavelet thresholding, and (d) EVT thresholding.

Another observation on the results of Table 1 is the noise levels of gradient fields and those of corresponding reconstructed surfaces. Since the surfaces are computed by virtually integrating the gradient fields projected onto the feasible subspace, the noise reduction properties of integration provides improved SNR levels in the surface space. The effects of noise reduction in the gradient space, therefore, manifest further improvements in the surface space.

5.2. Simulation Results on Real Data

We have also applied our gradient field noise reduction method on 64×64 blocks extracted from real range images obtained from the USC Range Image Database¹ (Stein and Medioni, 1992). For each range image, we have computed three brightness images assuming Lambertian reflectance illuminated by three sources with

illumination directions v_1 , v_2 , and v_3 given by

$$v_1 = \frac{1}{\sqrt{5}} \begin{bmatrix} 0 \\ 1 \\ 2 \end{bmatrix}, \quad v_2 = \frac{1}{\sqrt{5}} \begin{bmatrix} -\frac{\sqrt{3}}{2} \\ -\frac{1}{2} \\ 2 \end{bmatrix}, \quad \text{and} \\ v_3 = \frac{1}{\sqrt{5}} \begin{bmatrix} \frac{\sqrt{3}}{2} \\ -\frac{1}{2} \\ 2 \end{bmatrix}.$$

The true brightness images E_i are then given by

$$E_i(x, y) = R_i(n(x, y)) = v_i^T n(x, y), \quad i = 1, 2, 3, \quad (44)$$

where $n(x, y)$ denotes the normal of the surface at the point (x, y) . We have then added white Gaussian noise to these brightness images to obtain the noisy




Figure 8. Coefficients \hat{w}_j at 10 dB SNR, (a) plain, (b) ideal, (c) wavelet thresholding, and (d) EVT thresholding.

photometric stereo image triplets I_i at 0 dB SNR to be used in the experiments. The original range images and the noisy brightness images are shown in Fig. 10.

Given three noisy brightness images I_1 , I_2 , and I_3 , we have implemented a photometric stereo algorithm to find the surface normal at a point (x, y) by minimizing

$$H(n) = \frac{1}{2} \sum_{i=1}^3 (I_i(x, y) - v_i^T n)^2 \quad (45)$$

subject to the constraint $n^T n = 1$. Using the method of Lagrange multipliers, Eq. (45) becomes

$$L(n, \lambda) = \frac{1}{2} \sum_{i=1}^3 (I_i(x, y) - v_i^T n)^2 - \lambda(n^T n - 1). \quad (46)$$

Taking the derivative and equating to zero at the solution $n = n^*$ provides

$$\frac{\partial L}{\partial n} \Big|_{n=n^*} = \sum_{i=1}^3 (I_i(x, y) - v_i^T n^*) v_i - \lambda n^* = 0. \quad (47)$$

Solving the above equation for n^* and replacing in Eq. (46) results in a function $L(n^*, \lambda)$ having a derivative with respect to λ of form

$$\frac{\partial L}{\partial \lambda} \Big|_{n=n^*} = \frac{P(\lambda)}{Q(\lambda)} \quad (48)$$

where $P(\lambda)$ is a polynomial in λ of third order. The photometric stereo solution to the surface normal at (x, y) is then obtained by solving $P(\lambda) = 0$ for λ which is guaranteed to have at least one real root, and computing n^* accordingly.




Figure 9. Coefficients \hat{w}_j at 0 dB SNR, (a) plain, (b) ideal, (c) wavelet thresholding, and (d) EVT thresholding.

It should be noted that the estimate for the surface normal obtained by the method described above may provide unrealistic solutions under heavy noise such as normals nearly parallel to the $x - y$ plane or even normals pointing in the negative z direction. Such normals correspond to very large surface gradients which then cause wild fluctuations in the reconstructed surfaces. Examples of such instances are presented in Noakes and Kozera (2003a, 2003b). For this reason, we have implemented the photometric stereo method above with an additional step that checks for surface normal estimates making an angle larger than 80° with the z axis, and replacing them with a closest acceptable normal. This step effectively removes the unrealistic surface normal estimates for more agreeable alternatives. The ratio of corrected surface normal estimates for each case is shown in Table 2. Note that the number of surface normal estimates that require the correction described above is propor-

tional to the standard deviation of the brightness image noise.

The surface reconstruction results are shown in Fig. 11. Both the wavelet thresholding and the EVT thresholding filter most of the noise present in the plain surface reconstruction while retaining the major features of the true surfaces. More importantly, these results indicate that even though the additive, white and

Table 2. Standard deviation of the additive measurement noise in the photometric stereo brightness image triplets and the ratio of the corrected surface normal estimates.

| | Standard deviation of measurement noise | Ratio of corrected surface normal estimates |
|--------|-----------------------------------------|---------------------------------------------|
| Bach | 0.3169 | 470/4096 |
| Brahms | 0.2419 | 35/4096 |
| Chopin | 0.2576 | 45/4096 |




Figure 10. Real range images and corresponding photometric stereo triplets. In order to simulate measurement error, white Gaussian noise at 0 dB SNR is added to the calculated brightness images.

independent character of the brightness image noise is lost by the photometric stereo method when the surface normals and the gradients are computed, the noise model set forth in Eq. (18) still holds to a large extent

to allow accurate noise reduction by quadratic means. This result suggests that even under heavy noise, surface reconstruction by first solving the photometric stereo problem and then integrating the noisy gradients

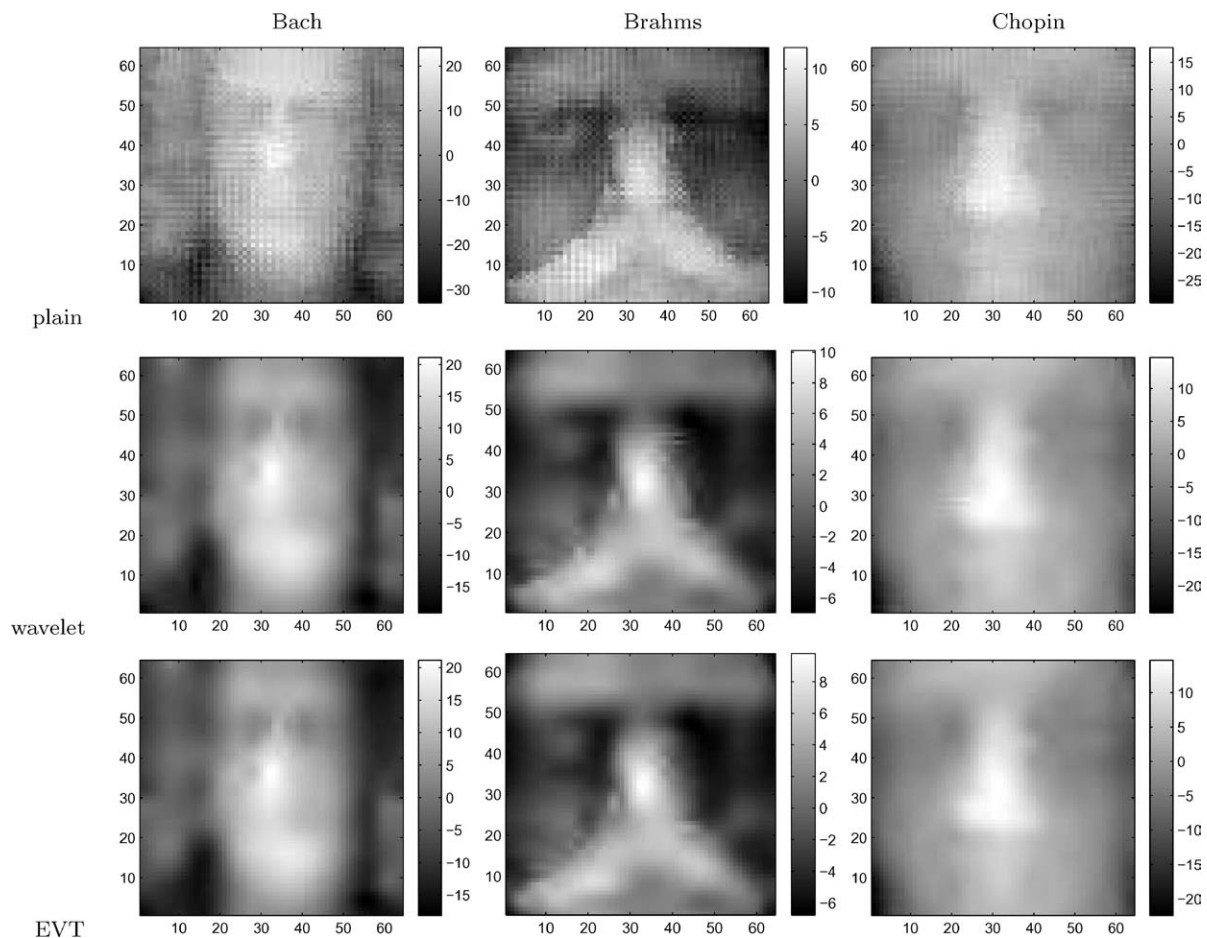


Figure 11. Surfaces reconstructed from the noisy photometric stereo triplets. Both the wavelet and EVT based thresholding strategies successfully remove most of the degenerative noise present in the plain reconstruction.

by performing noise reduction as shown here is a cost-effective and powerful alternative to solving iteratively for a surface that would otherwise correspond to the given photometric stereo image triplet as closely as possible.

6. Conclusion

We have presented an analysis of error sources affecting the accuracy of the surface gradient estimates of a three-image photometric stereo method in a noisy environment. We have discussed both the structural errors, caused by the deviations from the involved assumptions, and the measurement noise related with the imaging conditions and the estimation of reflectance maps. Based on this analysis, we have proposed an

additive noise model to represent the aggregate errors in the process of estimating surface gradients from intensity images.

We then exploited noise reduction techniques using the surface reconstruction method proposed by Karaçalı (2002) and Karaçalı and Snyder (2002). We show empirically that the multiscale gradient field representation is in fact a sparse orthonormal expansion in the space of possible gradient fields. Simulation results indicate that coefficient selection using thresholds from wavelet analysis and extreme value theory provides significant improvements on the noise levels of both the given noisy gradient fields and the surfaces reconstructed subsequently, large enough to challenge the unachievable statistically optimal reconstruction in some cases. Furthermore, the experiments on photometric stereo images of real range data corrupted by

additive noise indicate that the additive noise model remains viable for noise reduction purposes even under adverse noise conditions, even though the additive white Gaussian gradient noise model is broken by the nonlinear response of the photometric stereo method.

Note

1. USC Institute for Robotics and Intelligent Systems, Fridtjof Stein.

References

- Buzbee, B., Golub, G., and Nielson, C. 1970. On direct methods of solving poisson's equations. *SIAM Journal of Numerical Analysis*, 7:627–656.
- Chang, S.G., Bin, Y., and Vetterli, M. 2000. Adaptive wavelet thresholding for image denoising and compression. *IEEE Transactions on Image Processing*, 9(9):1532–1546.
- Courant, R. and Hilbert, D. 1962. *Methods of Mathematical Physics*. John Wiley & Sons: New York.
- Daubechies, I. 1988. Orthonormal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, 41:909–996.
- Deift, P. and Sylvester, J. 1981. Some remarks on shape from shading problems in computer vision. *Journal of Mathematical Analysis and Applications*, 34:235–248.
- Donoho, D.L. and Johnstone, I.M. 1994. Ideal denoising in an orthonormal basis chosen from a library of bases. *Comptes Rendus de l'Académie des Sciences, Série I*, 319.
- Embrechts, P., Klüppelberg, C., and Mikosch, T. 1997. Modelling extremal events for insurance and finance. In *Applications of Mathematics: Stochastic Modelling and Applied Probability*. Springer: Berlin.
- Frankot, R.T. and Chellappa, R. 1988. A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4):439–451.
- Frankot, R.T. and Chellappa, R. 1989. A method for enforcing integrability in shape from shading algorithms. In *Shape from Shading*, B.K.P. Horn and M.J. Brooks (Eds.). The MIT Press, Chapt. 5, pp. 89–122.
- Galo, M. and Tozzi, C.L. 1996. Surface reconstruction using multiple light sources and perspective projection. In *IEEE International Conference on Image Processing*, vol. 2, pp. 309–312.
- Grossberg, M.D. and Nayar, S.K. 2001. A general imaging model and a method for finding its parameters. In *IEEE Eighth International Conference on Computer Vision*, vol. 2, pp. 108–115.
- Hayakawa, H. 1994. Photometric stereo under a light-source with arbitrary motion. *Journal of Optical Society of America, Serie A*, 11(11):3079–3089.
- Horn, B.K.P. 1986. *Robot Vision*. MIT Press: Cambridge.
- Horn, B.K.P. and Brooks, M.J. 1986. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing*, 33:174–208.
- Horn, K.B.P. and Brooks, M.J. (Eds.). 1989. *Shape from Shading*. The MIT Press.
- Hsieh, J.-W., Liao, H.-Y.M., Ko, M.-T., and Fan, K.-C. 1995. Wavelet-based shape from shading. *Graphical Models and Image Processing*, 57(4):343–362.
- Hurt, N.E. 1991. Mathematical methods in shape-from-shading: A review of recent results. *Acta Applicandae Mathematicae*, 23(2):163–188.
- Ikeuchi, K. 1981. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(6):661–669.
- Ikeuchi, K. and Horn, B.K.P. 1981. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17(1–3):141–184.
- Iwahori, Y., Tanaka, H., Woodham, R.J., and Ishii, N. 1994. Photometric stereo for specular surface shape based on neural network. *IEICE Transactions on Information and Systems*, E77-D(4):498–506.
- Iwahori, Y., Woodham, R.J., Ozaki, M., Tanaka, H., and Ishii, N. 1997. Neural network based photometric stereo with a nearby rotational moving light source. *IEICE Transactions on Information and Systems*, E80-D(9):948–957.
- Jain, A.K. 1989. *Fundamentals of Digital Image Processing*. Prentice Hall.
- Jones, A. and Taylor, C. 1994. Robust shape from shading. *IVC*, 12(7):411–421.
- Karaçalı, B. and Snyder, W. 2002. Partial integrability in surface reconstruction from a given gradient field. In *IEEE International Conference on Image Processing*, vol. 2, pp. 525–528.
- Karaçalı, B. 1999. Signal reconstruction using the cross-scale correlation properties of the wavelet decomposition. Technical Report TR 99/13, North Carolina State University, Center for Advanced Computing and Communications.
- Karaçalı, B. 2002. Vector space methods in surface reconstruction from one or more images acquired from the same view with application to scanning electron microscopy images. Ph.D. thesis, North Carolina State University, Department of Electrical and Computer Engineering.
- Kiang, R.K. 1997. Rectifying airborne scanner measurements using neural networks. In *Proceedings of SPIE, Applications and Science of Artificial Neural Networks III*, vol. 3077, pp. 525–534.
- Kim, B.-H. and Park, R.-H. 1997. Shape from shading and photometric stereo using surface approximation by legendre polynomials. *Computer Vision and Image Understanding*, 66(3):255–270.
- Kozera, R. 1991. Existence and uniqueness in photometric stereo. *Applied Mathematics and Computation*, 44(2):1–103.
- Krim, H., Mallat, S., Donoho, D., and Willsky, A.S. 1995. Best basis algorithm for signal enhancement. In *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*.
- Lambert, J.H. 1760. *Photometria sive de Mensura et Gratibus Luminis, Colorum et Umbrae*. Eberhard Klett, Augsburg.
- Lee, K.M. and Kuo, C.-C.J. 1994. Shape from shading with perspective projection. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 59(2):202–212.
- Mallat, S. 1989. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693.
- Mallat, S.G. 1998. *A Wavelet Tour of Signal Processing*. Academic Press: San Diego.
- Noakes, L. and Kozera, R. 1999. A 2D leap-frog algorithm for optimal surface reconstruction. In *Proceedings of the 44th Annual*

- Meeting of the SPIE, Stream Vision Geometry VIII*, vol. 3811, pp. 317–328.
- Noakes, L. and Kozera, R. 2001. 2D Leap-Frog: Integrability, noise and digitization. In *Lecture Notes in Computer Science*, vol. 2243, pp. 352–364.
- Noakes, L. and Kozera, R. 2003a. Denoising images: Non-linear Leap-Frog for shape and light-source recovery. In *Lecture Notes in Computer Science*, T. Asano, R. Klette, and C. Ronse (Eds.), vol. 2616, pp. 419–436.
- Noakes, L. and Kozera, R. 2003b. Nonlinearities and noise reduction in 3-source photometric stereo. *Journal of Mathematical Imaging and Vision*, 18(2):119–127.
- Noakes, L., Kozera, R., and Klette, R. 1999. The lawn-mowing algorithm for noisy gradient vector fields. In *Proceedings of the 44th Annual Meeting of the SPIE, Stream Vision Geometry VIII*, vol. 3811, pp. 305–316.
- Phong, B.T. 1975. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317.
- Rajaram, K.V., Parthasarathy, G., and Faruqi, M.A. 1995. Neural network approach to photometric stereo inversion of real-world reflectance maps for extracting 3-D shapes of objects. *IEEE Transactions on Systems, Man and Cybernetics*, 25(1):1289–1300.
- Shevlin, F. 1994. Rectifying airborne scanner measurements using neural networks. In *IEEE International Conference on Image Processing*, vol. 1, pp. 1007–1011.
- Stein, F. and Medioni, G. 1992. Structural indexing: Efficient 3D object recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):125–145.
- Symchony, T., Chellappa, R., and Shao, M. 1990. Direct analytical methods for solving poisson equation in computer vision problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(5):435–446.
- Tagare, H.D. and deFigueiredo, R.J.P. 1991. A theory of photometric stereo for a class of diffuse non-Lambertian surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(2):133–152.
- Woodham, R.J. 1978. Photometric stereo: A reflectance map technique for determining surface orientations from a single view. In *Proceedings of SPIE, Image Understanding Systems & Industrial Applications*, vol. 155, pp. 136–143.
- Zhang, R., Tsai, P.-S., Cryer, J.E., and Shah, M. 1999. Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706.
- Zheng, Q., Klemas, V., Yan, X.-H., Wang, Z., and Kagleder, K. 1997. Digital orthorectification of space shuttle coastal ocean photographs. *International Journal of Remote Sensing*, 18(1):197–211.