Shapelets Correlated with Surface Normals Produce Surfaces

Peter Kovesi
School of Computer Science & Software Engineering
The University of Western Australia
Crawley, Western Australia 6009
pk@csse.uwa.edu.au

Abstract

This paper addresses the problem of deducing the surface shape of an object given just the surface normals. Many shape measurement algorithms such as shape from shading and shape from texture only return the surface normals of an object, often with an ambiguity of π in the surface tilt. The surface shape has to be inferred from these normals, typically via some integration process. However, reconstruction through the integration of surface gradients is sensitive to noise and the choice of integration paths across the surface. In addition, existing techniques cannot accommodate ambiguities in tilt. This paper presents a new approach to the reconstruction of surfaces from surface normals using basis functions, referred to here as shapelets. *The surface gradients of the shapelets are correlated with the gradients of the surface and the correlations summed to form the reconstruction. This results in a simple reconstruction process that is very robust to noise. Where there is an ambiguity of π in the surface tilt, reconstructions of reduced quality are still possible up to a positive/negative shape ambiguity. Intriguingly, some form of reconstruction is also possible using just slant information.*

1. Introduction

There are a number computer vision techniques that attempt to determine the surface normals within a scene. These include shape from shading and its extension, photometric stereo [4, 11, 12, 17], shape from texture [1, 2, 28, 18, 8, 7], and slant and tilt from differential invariants [15, 9]. Often the surface normals obtained from these methods are ill-conditioned and subject to noise.

This paper addresses the problem of deducing the surface shape of an object given just the surface normals. The traditional approach to reconstructing a surface from its surface normals is via integration. The difficulty with this approach is that it can be very sensitive to noise and there is the problem of finding appropriate regularization techniques to impose the requirement that the surface gradient be integrable [3, 14, 30, 31, 22, 23].

Terzopoulos [29] proposes a variational formulation of the reconstruction process. A concern with this approach is that in forming an energy function that seeks to minimize deviation from depth and orientation constraints, and also minimize surface discontinuity as measured by some thin plate spline, one produces an expression with gross dimensional inconsistencies. The penalty function that measures deviation from depth and orientation constraints involves the addition of distance error squared terms to gradient error squared terms. The surface stability function involves terms where the square of the second derivative of the surface is added to the square of the first derivative. Finally the penalty and surface stability functions are added to produce the energy function. A number of parameters have to be supplied which have to attempt to act as unit conversion factors for these mixed quantities. Another difficulty is that the iterative energy minimization solution process is slow to converge. To improve convergence speed Terzopoulos proposes a multiresolution solution process. Solutions at coarse scales allow constraints to propagate across the surface more rapidly. Coarse scale solutions are then used as a starting point for iterative solution at a finer scale. The dimensional inconsistency of the energy function means that the behaviour of the iterative solution process can vary considerably at different scales, further complicating the setting of parameters. Most schemes involving regularization terms suffer from these difficulties.

Frankot and Chellappa [10] introduced a very simple and powerful way of ensuring integrability of a surface as part of their shape from shading algorithm. The surface gradients are projected onto a set of integrable basis functions, and the surface is reconstructed from these. They use the Fourier basis functions. This is a quick one-step algorithm that is highly robust to noise. Indeed, it possibly remains one of the most noise tolerant algorithms to date.

The natural extension of this approach is to use localized wavelet basis functions rather than Fourier ones. Hsieh et al. [13] employ this in their shape from shading algorithm. Karaçali and Snyder [5, 6] employ a reconstruction approach based on constructing an orthonormal set of gradient fields that span a feasible subspace of the gradient space using wavelets. The measured gradient field is projected onto the feasible subspace to produce a surface with gradient field closest to the measured gradients. They adapt wavelet denoising techniques to very successfully reduce the influence of noise on the reconstruction. A drawback of this approach is that discontinuities in the surface have to be deduced and treated specially within the context of the feasible gradient subspace.

The work presented in this paper follows the approach of projecting gradients onto basis functions. Where this work differs is that a redundant set of non-orthogonal basis functions of finite support are used, and the correlation with the basis functions is formulated with respect to slant and tilt, rather than in terms of the gradient of the surface with respect to x and y. The advantage of this is that the correlation measures can be modified to allow for ambiguities in tilt of π, and even allow for the case where no tilt data is available. Where there is ambiguity in the tilt data this new approach allows potential surface shape solutions to be hypothesized for subsequent verification or rejection.

Many forms of basis functions have been used for image decomposition and reconstruction with various forms of wavelets being popular in recent years [21]. When the decomposition is done in terms of a continuous wavelet transform, say using Gabor wavelets, the correlation results correspond to band-passed versions of the image; these can be summed directly to reconstruct the image. It is an adaptation of this latter approach, applied to surface normals, that is employed in this paper for the reconstruction of surfaces.

2. Reconstruction from Shapelets

The aim of shape from shading, photometric stereo and shape from texture is to produce a range image of the scene. Clearly the range image can be represented as the sum of a set of basis functions, however we have to determine the set of the basis functions that make up the range image solely from the gradient data. This proves not to be too problematic if we make the following observation.

A correlation performed between the gradients of a signal and the gradients of a basis function can provide information equivalent to direct correlation between the signal and basis function (up to a signal offset) because differentiation is linear. If the values in a range image are doubled ('doubling' its shape) so too will the surface gradients. Thus, if we correlate the surface gradient information with the gradients of a bank of shapelet basis functions we can reconstruct the surface shape, up to an offset, by simply summing the correlation results. The summing of the basis correlations automatically imposes a continuity constraint and performs an implicit integration of the surface from its gradients.

2.1. Choice of Shapelet

Potentially there are many basis functions that could be used as shapelets[1]. In designing a bank of shapelet filters we start by noting that correlating the gradient of one shapelet filter with the gradient of the signal will correspond to extracting a band of frequencies from the signal gradient. The need to reconstruct a surface from correlations between surface normals and shapelet gradients imposes some constraints on the shapelet function. These are:

- The gradient of the shaplet function must satisfy the admissibility condition of zero mean.
- The shapelet must have minimal ambiguity of shape with respect to its gradient.
- The shapelet function must allow preservation of phase information in the signal.
- The bank of shapelet filters should, ideally, provide uniform coverage of the signal spectrum so that it is faithfully reconstructed.

To achieve a shapelet with minimal ambiguity of shape with respect to its gradient it should be simple, and ideally take the form of a single, symmetric peak so that (in 1D) the gradient function will have a single positive peak and a single negative peak. This also ensures that the admissibility condition is satisfied. With the assumption that the shapelet is non-negative and symmetric, say a Gaussian, the gradient function will be odd-symmetric, resulting in a frequency domain transfer function that is complex and odd-symmetric. Changing the scale of the shapelet will have the effect of changing the separation of the peaks of the gradient function producing a transfer function that responds to a different band of frequencies. Thus, if several scales of shapelets are used one can achieve fairly complete coverage of the spectrum.

An important constraint is that the transfer function of the shapelet gradient should not corrupt phase information in the gradient of the signal [25]. This implies that the transfer function of the shapelet gradient should be non-negative for positive frequencies and non-positive for negative frequencies. Achieving this with a simple shape of

[1] The term 'shapelet' was first coined by Refregier for his orthonormal basis set consisting of weighted Hermite polynomials [26]. Here we use to term to describe any basis function of finite support used for representing shape.

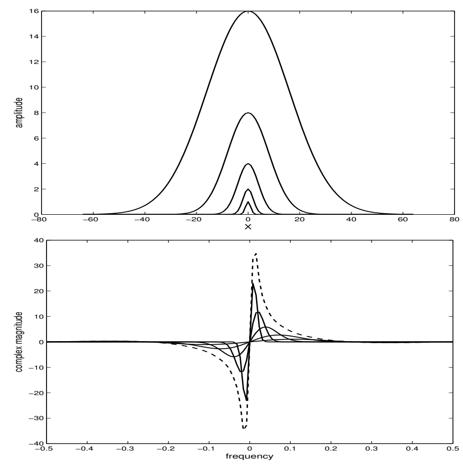

Figure 1. A shapelet bank formed by five Gaussians with height proportional to scale (top), and the corresponding transfer functions of their gradients (bottom). The dashed line shows the sum of the transfer functions, ideally this should form a curve that is inversely proportional to frequency.

finite support limits the shapelet function to a Gaussian, or near Gaussian, shape.

Finally we must address coverage of the signal spectrum achieved by the bank of shapelets. As mentioned above, correlating the gradient of one shapelet filter with the gradient of the signal will correspond to extracting a band of frequencies from the signal gradient. The sum of all the transfer functions of the shapelet gradients will represent the total energy extracted from the signal gradient spectrum. The signal gradient spectrum differs from the original signal spectrum in that its Fourier components are phase shifted by $\pi/2$ and the amplitudes are scaled by the frequency. Accordingly the sum of the shapelet gradient transfer functions should attempt to form a curve that is inversely proportional to frequency to counteract this scaling of the gradient spectrum. Ideally the net result is that we obtain uniform coverage of the original signal spectrum. This ensures all frequency components of the signal are represented equally and we obtain a faithful reconstruction, up to a scale factor.

For 1D signals a suitable shapelet basis set can be constructed by a set of Gaussian functions with the scaling between successive filters being a factor of two and the heights of the Gaussian functions being proportional to scale. Figure 1 shows such a shapelet bank formed by five Gaussians.

In this situation we clearly do not have an orthogonal basis set. This is not of great concern if we ensure we have a strongly redundant, or overcomplete, basis set. Redundancy in the basis set is achieved by using a continuous wavelet/shapelet decomposition and by using a small geometric scaling factor between successive scales of basis function shapes. Within a strongly redundant non-orthogonal basis set one can take each basis function and find another that is approximately orthogonal to it. Thus we can consider a redundant non-orthogonal basis set as approximating the sum of several *orthogonal* basis sets. This means that we can achieve accurate reconstruction of a function up to a scale factor. If necessary this scale factor can be determined and the scale of the reconstruction corrected. Overcomplete basis sets have been advocated by Simoncelli et al. [27], and Olshausen and Field [24]. They have the advantage that small changes in local signal features result in smooth transitions in the basis coefficients, overcompleteness also provide robustness in the presence of noise.

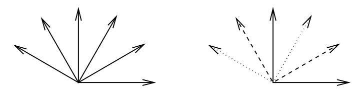

Figure 2. A redundant, or overcomplete, non-orthogonal basis set, shown on the left, can be considered to be the sum of several orthogonal basis sets, as seen on the right. Thus accurate reconstruction, up to a scale factor, can be obtained.

For the reconstruction of a 1D signal from its gradients we calculate the gradient correlation at each shapelet scale using

$$C_i = \nabla_f \star \nabla_{si} \ , \tag{1}$$

where ∇_f and ∇_{si} denote the gradients of the surface and shapelet at scale i respectively, and $\star$ denotes correlation. These correlations are then summed to form the reconstruction, R. Note that the reconstruction will be scaled by the degree of redundancy of the shapelet basis set.

$$R = \sum_i C_i \ . \tag{2}$$

3. Shapelet reconstruction in 2D

Two-dimensional fields of surface normals are typically calculated and/or specified in terms of slant and tilt. The following discussion assumes we have orthographic knowledge of the slant and tilt values over the surface. It is desirable that correlation between the gradients of the surface and shapelets be defined in terms of slant and tilt separately and then combined because often the calculation of these two quantities is obtained by very different means, and they

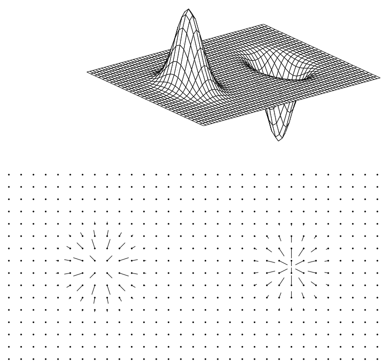

Figure 3. A surface with a peak and a trough along with its corresponding needle diagram showing the different tilt values.

can be subject to differing degrees of uncertainty. For example, tilt often has an ambiguity of π.

The procedure that has been adopted is as follows: Correlation of the surface and shapelet slants is done in terms of the gradient *magnitude* which is given by the tangent of the slant, σ

$$|\nabla| = \tan(\sigma) \ . \tag{3}$$

The gradient correlation is then formed as

$$C_{\nabla i} = |\nabla_f| \star |\nabla_{si}| \ . \tag{4}$$

With no tilt information this correlation matches positive and negative gradients equally because we only have access to the gradient magnitude. Thus a mound in the surface will correlate equally with a similarly shaped depression in the surface. To be able to make a distinction between positive and negative shapes tilt information must be used.

If at some point the surface and shapelet gradient magnitudes match and the tilt directions also match then the component of the shapelet at this point must be positive. If the tilts of the surface and shapelet are in opposite directions then the shapelet component must be negative. If the tilts are orthogonal then there is no correlation, positive or negative, between the surface and shapelet. Thus the gradient correlation must be multiplied by a tilt correlation measure that varies between 1, when tilts are aligned, to -1 when the tilts are in opposite directions. An obvious measure that satisfies this requirement is the cosine of the tilt angle difference. To form the tilt correlation between a shapelet at scale i and the surface we sum the cosine of the tilt differences between points on the surface and shapelet, and use the standard trigonometric difference equation to overcome any angle wraparound problems at the origin. Thus

$$C_{\tau i} \quad = \quad \cos(\tau_f) \star \cos(\tau_{si}) + \sin(\tau_f) \star \sin(\tau_{si}) \ . \tag{5}$$

where τ_f and τ_{si} denote the tilts of the surface and shapelet at scale i respectively.

The overall correlation measure between surface and shapelet at scale i is obtained by the point-wise product of the gradient and tilt correlations

$$\begin{aligned} \mathbf{C}_i &= C_{\nabla i}.C_{\tau i} \\ &= |\nabla_f| \star |\nabla_{si}|.[\cos(\tau_f) \star \cos(\tau_{si}) \\ &\quad + \sin(\tau_f) \star \sin(\tau_{si})] \\ &= [\ |\nabla_f|.\cos(\tau_f)\ ] \star [\ |\nabla_{si}|.\cos(\tau_{si})\ ] \\ &\quad +[\ |\nabla_f|.\sin(\tau_f)\ ] \star [\ |\nabla_{si}|.\sin(\tau_{si})\ ] \ , \end{aligned} \tag{6}$$

where . denotes point-wise multiplication. This process is performed over multiple shapelet scales and the results summed to form the reconstruction

$$R = \sum_i \mathbf{C}_i \ . \tag{7}$$

4. Results

The following results were obtained using a bank of Gaussian shapelets with a geometric scaling factor of 2 between successive shapelets, the standard deviation of the smallest shapelet was 1. Note that for 2D reconstructions the heights of the Gaussian shapelets should be constant, rather than being proportional to scale as with the 1D case. A synthetic test surface was constructed and slant and tilt values extracted. These were then supplied to the algorithm to see how well the original surface was regenerated. Figure 4 shows how, starting with the smallest shapelet scales and working up, the different scales combine to produce the reconstruction. Figure 5 shows the results of applying the shapelet reconstruction algorithm to real data. Here the surface slant and tilt values of the pear were obtained via Loh's shape from texture algorithm [19, 20].

4.1. Reconstruction with the presence of noise

To illustrate the robustness of the reconstruction Gaussian noise was added to the slant and tilt values of the test surface with standard deviations of 0.3 radians. Noisy slant values were clamped to the range 0 to $\pi/2 - \epsilon$, where ϵ was set to 0.05 radians, to exclude impossible values. With this level of noise reconstruction via direct integration based approaches fail completely. The shapelet reconstruction results using 6 scales are shown in Figure 6. For comparison the surface reconstructions obtained via Frankot and Chellappa's algorithm [10] and Terzopoulos' multigrid method [29] are also shown.

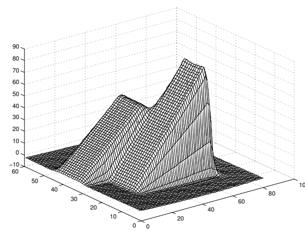

Original surface.

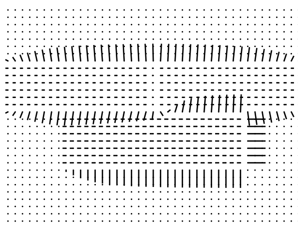

Surface normal needle plot.

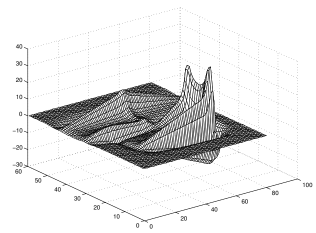

Reconstruction over 2 scales.

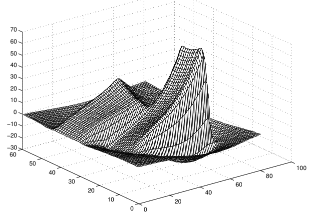

Reconstruction over 4 scales.

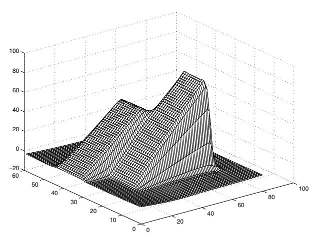

Reconstruction over 6 scales.

Figure 4. Reconstructions of a test surface from slant and tilt data using 2, 4 and 6 shapelet scales. These show the effect of progressively adding lower frequency surface gradient information.

These results demonstrate that the shapelet reconstruction algorithm shares the strong robustness to noise that the Frankot and Chellappa algorithm has. Typically the reconstruction will be slightly smoother than the Frankot and Chellappa reconstruction because the smallest shapelet scale still excludes some of the high frequency components of the noisy surface gradient data. For Terzopoulos' multigrid method 200 iterations, applied at each of 3 scales of analysis, were used. At the locations of the discontinuities in the surface the multigrid thin plate spline parameters were manually preset to allow orientation discontinuities, this improved the reconstruction but increased the influence of noise at these points. Note that at the lower scale grid representations one cannot place these discontinuities in their ideal locations. The resulting errors in the low scale reconstructions propagate to the finer scales. With appropriate parameters the multigrid approach can be tolerant to noise but it has difficulty propagating constraints across the surface when only surface orientation data is supplied.

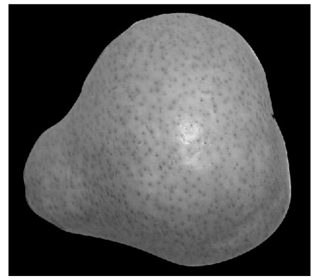

Segmented pear image.

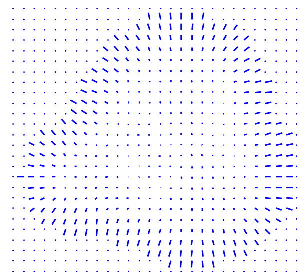

Normals from texture

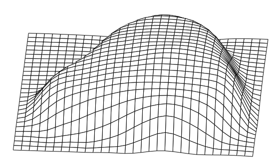

Reconstructed surface.

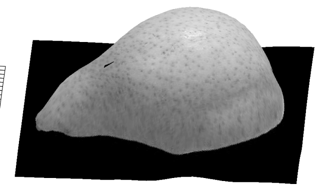

Texture mapped reconstruction.

Figure 5. Reconstruction of a pear using slant and tilt deduced from texture.

It should be noted that the results presented here are for noise levels that are much larger than are typically presented in the literature. For example Wei and Klette [31] present results for Gaussian noise with a standard deviation of only 0.01, similarly Noakes and Kozera [22] provide results for Gaussian noise of 0.04 standard deviation. However, a recent exception to this is the work by Karaçali and Snyder [6] where they present a good quality reconstruction of a surface where the signal-to-noise level, expressed in terms of gradient space noise, was 0dB. The target surface that they were reconstructing was continuous with a maximum gradient of approximately 2.5.

4.2. Reconstruction with tilt ambiguity of π

The previous section has demonstrated that the shapelet reconstruction approach provides reconstructions that are highly robust to noise. However, the key advantage of the approach is that it also allows reconstructions to be considered where there is an ambiguity of π in the tilt data, or where there is no tilt data at all. Measurement of surface tilt via shape from shading or texture will typically have an ambiguity of π. With such an ambiguity one is unable to determine whether a shape is positive or negative. This causes considerable problems for any reconstruction algorithm that requires as input gradient information with respect to x and y. However, some form of surface reconstruction is possible via shapelets.

To account for the ambiguity of π the tilt correlation

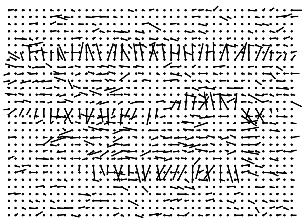

Surface normal needle plot.

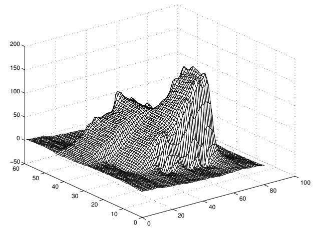

Shapelet reconstruction.

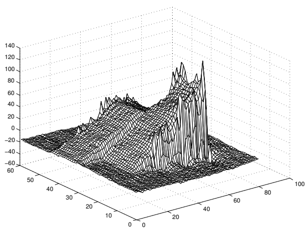

Frankot Chellappa reconstruction.

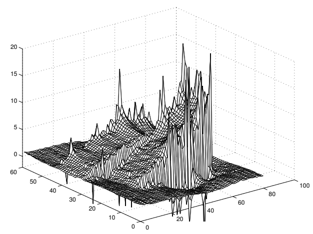

Multigrid reconstruction.

Figure 6. Surface reconstructions from slant and tilt data with additive Gaussian noise having standard deviation of 0.3 radians.

measure is modified so that we work with the squared cosine of the tilt differences. Using the trigonometric double angle formula equation 5 is modified to be

$$C_{\tau i} = \frac{\cos(2\tau_f) \star \cos(2\tau_{si}) + \sin(2\tau_f) \star \sin(2\tau_{si}) + 1}{2} . \tag{8}$$

This provides a measure that varies from 1, when tilts are aligned, down to 0 when they are orthogonal, and then back to 1 when they are in opposite directions. The slant correlation remains as before.

The reconstructions obtained on the peak and trough, and ramps test surfaces are shown in Figure 7. Note that 'negative' shapes are reconstructed in the 'positive' direction as this is the implicit assumption of the reconstruction process. With the tilt correlation modified to allow for tilt ambiguity both the slant and tilt correlation measures are always non-negative. Thus the surface is constructed purely by the addition of the non-negative basis functions. This limits the shapes that can be reconstructed. The quality of the reconstructed shapes is also degraded and some artifacts are produced. Note the additional lobes that have been produced at the sharp discontinuities at the edges of the ramps surface. Here the shapelet function would normally have a strong negative correlation with the surface normal of the shapelet being in the opposite direction to the surface normal of the surface. However, the modification to the tilt correlation to allow ambiguity of π results in a correlation that is positive at these points.

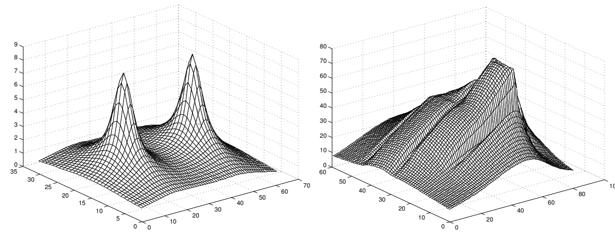

Figure 7. Reconstructions where there is a tilt ambiguity of π. The peak and trough surface is shown left, and the ramps test surface, right.

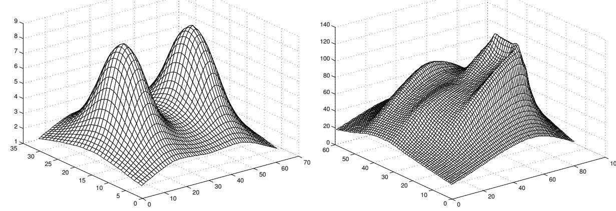

Figure 8. Reconstructions of the two test surfaces using slant data alone.

4.3. Reconstruction from slant or tilt alone

Interestingly, some form of reconstruction is achieved if one only uses the gradient magnitude correlation and ignores the tilt, this is shown for the two test surfaces in Figure 8. Note that the slant correlation is always non-negative so again we only reconstruct 'positive' surfaces. Despite this the reconstructions based on slant alone are surprisingly good. Reconstruction using only the tilt correlation is weaker than using slant alone. Even though tilt provides information about positive and negative shape it does not provide knowledge of the magnitude and reconstructions using tilt information alone are typically very poor.

4.4. Shape from occluding contours

The boundary of an object provides a cue to the basic shape of an object not only through the position of the boundary but by the fact that the normal to the surface at an occluding contour is at 90 degrees to the viewing direction. Occluding contours were marked manually on an image of the Mona Lisa. Along these contours the slant was set to $\pi/2$ and tilt was assigned to be perpendicular to the contour with the appropriate direction. At all other points in the scene the slant was assumed 0. This gradient field is far from integrable (at line terminations and T junctions)

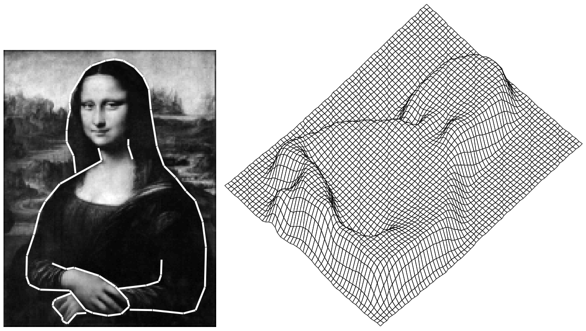

Figure 9. Shape reconstruction of the Mona Lisa from occluding contours.

but the shapelet reconstruction approach has no difficulties with this. This very crude approximation of what the surface normals are for the whole scene is enough to construct a simple $2\frac{1}{2}$D representation of the scene.

5. Conclusions

Reconstructions of surfaces from their surface normals via shapelets provides a new reconstruction method that is simple to implement and is highly robust to noise. The use of basis functions implicitly imposes a continuity constraint in the reconstruction, yet at the same time allows sharp transitions to be represented in the surface via the finer scales of shapelets. There is no need to infer the locations of discontinuities in the surface and/or apply special conditions at these points. The correlation process treats slant and tilt separately and makes the different roles of slant and tilt explicit in the reconstruction process. This permits data with tilt ambiguity to be considered and allows slant and tilt data to be considered in isolation. These options would be impossible to consider using a reconstruction process based on integration of gradients. It is envisaged that the flexibility and robustness of the shapelet reconstruction approach will create opportunities for new shape from texture and shape from shading algorithms. For example, where tilt information is uncertain one may be able to adopt an iterative approach. First, a surface would be hypothesized using slant information and, if available, any estimates of tilt. The surface would then be verified by comparing an estimate of the appearance of the reconstructed surface against the original image, and on the basis of this the tilt estimates would be updated. This would be repeated until convergence occurred.

MATLAB code is available for those wishing to replicate the results presented here [16].

Acknowledgment

The author would like to thank Angeline Loh for providing the shape from texture results.

References

[1] J. Aloimonos. Visual shape computation. *Proceedings of the IEEE*, 76(8):899–916, August 1988.

[2] D. Blostein and N. Ahuja. Shape from texture: Integrating texture-element extraction and surface estimation. *PAMI*, 11(12):1233–1251, 1989.

[3] M. J. Brooks. Surface normals from closed paths. In *Int. Jt. Conf. Artificial Intelligence*, pages 98–101, 1979. Tokyo, Japan.

[4] M. J. Brooks and B. K. P. Horn. Shape and source from shading. In *Int. Jt. Conf. Artificial Intelligence*, pages 932–936, August 1985. Los Angeles, CA.

[5] B. Karaçali and W. Snyder. Partial integrability in surface reconstruction from a given gradient field. In *IEEE International Conference on Image Processing*, volume 2, pages 525–528, 2002.

[6] B. Karaçali and W. Snyder. Noise reduction in surface reconstruction from a given gradient field. *International Journal of Computer Vision*, 60(1):24–44, 2004.

[7] M. Clerc and S. Mallat. The texture gradient equation for recovering shape from texture. *PAMI*, 24(4):536–549, 2002.

[8] D. A. Forsyth. Shape from texture and integrability. In *ICCV01*, pages 447–452, 2001.

[9] E. Francois and P. Bouthemy. Derivation of qualitative information in motion analysis. *Image and Vision Computing*, 8(4):279–288, November 1990.

[10] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading. *IEEE PAMI*, 10(4):439–451, July 1988.

[11] B. K. P. Horn. *Robot Vision*. McGraw-Hill, 1986. New York.

[12] B. K. P. Horn. Height and gradient from shading. *International Journal of Computer Vision*, 5:37–75, 1990.

[13] J. Hsieh, H. Liao, M. Ko, and K. Fan. Wavelet-based shape from shading. *Graphical Models and Image Processing*, 57(4):343–362, 1995.

[14] R. Klette and K.Schluns. Height data from gradient fields. In *Proceedings of SPIE Machine Vision Applications, Architectures, and Systems Integration (2908)*, pages 204–215, 1996. Boston, USA.

[15] J. J. Koenderink and A. J. van Doorn. How an ambulant observer can construct a model of the environment from the geometrical structure of the visual inflow. *Kybernetik*, pages 224–247, 1978.

[16] P. D. Kovesi. MATLAB functions for computer vision and image analysis. School of Computer Science & Software Engineering, The University of Western Australia. Available from:
<http://www.csse.uwa.edu.au/∼pk/research/matlabfns/>.

[17] R. Kozera. An overview of the shape-from-shading problem. *Machine Graphics and Vision*, 7(1/2):291–312, 1998.

[18] K. Lee and C. Kuo. Direct shape from texture using a parametric surface model and an adaptive filtering technique. In *CVPR98*, pages 402–407, 1998.

[19] A. M. Loh and P. D. Kovesi. Estimation of surface normal of a curved surface using texture. In *Proceedings of the 7th Australian Pattern Recognition Society Conference – Digital Image Computing: Techniques and Applications*, pages 155–164, December 2003. Sydney.

[20] A. M. Loh and P. D. Kovesi. Shape from texture without estimating transforms. Technical Report UWA-CSSE-05-001, School of Computer Science & Software Engineering, The University of Western Australia, 2005. http://www.csse.uwa.edu.au/research/Technical_Reports.

[21] S. G. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.

[22] L. Noakes and R. Kozera. *2-D Leap Frog: Integrability, noise and digitization*, pages 339–351. Springer-Verlag, 2001.

[23] L. Noakes and R. Kozera. Nonlinearities and noise reduction in 3-source photometric stereo. *Journal of Mathematical Imaging and Vision*, 18(2):119–127, 2003.

[24] B. Olshausen and D. Field. Sparse coding with an overcomplete basis set: A strategy employed by V1?, 1997.

[25] A. V. Oppenheim and J. S. Lim. The importance of phase in signals. In *Proceedings of the IEEE 69*, pages 529–541, 1981.

[26] A. Refregier. Shapelets: I. A Method for Image Analysis. *Mon. Not. R. Astron. Soc.*, 338(35), 2003. http://arXiv.org/abs/astro-ph/0105178.

[27] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger. Shiftable multiscale transforms. *IEEE Transactions on Information Theory*, 38(2):587–607, March 1992.

[28] B. J. Super and A. C. Bovik. Shape from texture using local spectral moments. *PAMI*, 17(4):333–343, April 1995.

[29] D. Terzopoulos. The computation of visible-surface representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4):417–438, 1988.

[30] T. Wei and R. Klette. A wavelet-based algorithm for height from gradients. *Lecture Notes in Computer Science*, 1998:84–90, 2001.

[31] T. Wei and R. Klette. Height from gradient using surface curvature and area constraints. In *3rd Indian Conference on Computer Vision, Graphics and Image Processing*, 2002. Ahmedabad.