[14] Invariant Surface Reconstruction using Weak Continuity Constraints

Andrew Blake and Andrew Zisserman

Department of Computer Science
University of Edinburgh, Edinburgh EH9 3JZ, Scotland

Abstract

We consider surface reconstruction schemes in which the discontinuities are included explicitly by means of weak continuity constraints. It is essential that such a scheme be viewpoint invariant. We explain how this can be achieved. Two schemes are described: an invariant membrane and an invariant plate, where the regularising terms are respectively surface area and surface curvature. Results are presented of detection of discontinuities in range data, both simulated and from a laser range finder.

1 Introduction

Much work has been carried out on the problem of reconstructing surfaces from three dimensional data. Grimson[11] discusses the interpolation of smooth surfaces. He shows how surface interpolation can be done by minimising a suitably defined energy functional. The interpolating surface that results is biharmonic and under most conditions is defined uniquely. Terzopoulos[12] extended this work by using finite elements and generalising the interpolating surface to be a mixture of "membrane" and "thin plate". The computation involved relaxation which is widely used for minimisation problems in computer vision[14], largely because of its inherent parallelism.

However, there are two fundamental problems with these optimisation approaches. Firstly, the schemes do not label discontinuities *explicitly*. They can only be located by examining the gradient, after the surface has been fitted. Discontinuities can be included explicitly if the optimisation formulation uses *weak continuity constraints*. A weak constraint is one that is usually obeyed but may be broken on occasion - when there are pressing reasons to do so. This is discussed more fully in[2,6].

The second problem is that such schemes are not *invariant*. The reconstructed scheme will "wobble" as the viewpoint varies[4]. This is because instead of using the invariant quantities of surface area and surface curvature[1], the schemes use approximations to these quantities which are not invariant to change of viewpoint. Consequently the energy of the minimum energy surface changes with viewpoint. The advantage of using approximations is that energy can be minimised by solving locally coupled linear equations, and that there is a unique minimum in each frame.

In this paper we discuss two reconstruction schemes for surface approximation which incorporate weak continuity constraints. For a weak constraint a penalty is charged each time a constraint is broken. The penalty is weighed against certain "other costs". If breaking a constraint leads (somehow) to a total saving in "other costs" that exceeds the penalty, then breaking that constraint is deemed worthwhile. The energy of the "other costs" must be invariant otherwise the existence and position of discontinuities will depend on viewpoint.

Invariant schemes incorporating weak continuity constraints are harder to minimise than the earlier schemes because

- The cost function is non - convex so that naive descent tends to stick at local minima.
- They do not give rise to linear equations.

The first scheme we consider is an invariant membrane with weak (C^0) continuity constraints. Here the "other costs" mentioned above consist of a "closeness of fit" term and a term involving surface area. Minimising the energy is a trade off between fitting the data closely, reducing the surface area, and including discontinuities.

In the second scheme the "other costs" are a measure of closeness of fit and a term involving surface curvature. There are two types of discontinuity - a discontinuity in surface depth and a discontinuity in surface orientation.

Apart from reconstructing the surface an important use of the schemes is in localising discontinuities. Local edge detection operators (such as the Canny[9]) can have poor localisation in noisy data (where a large support must be used to improve the signal to noise ratio) . The problem is accentuated if the underlying step in the data is not symmetrical (for example a step with a finite gradient on one side). However, the localisation of discontinuities using weak continuity constraints is very good even under these circumstances[5].

In our numerical implementation we have used the Graduated Non-Convexity (GNC) Algorithm[3,6] as an approximate method for obtaining the global minimum. This allows the discretised energy to be minimised by local iterative methods. The GNC algorithm has been applied to dense range data obtained from a CSG body modeller and also from a laser range finder.

2 Invariant Membrane

The non-invariant schemes are

In 1D - a weak string:

$$E = \int \{(u-d)^2 + \lambda^2 (u')^2\}\, dx + P \qquad (1)$$

In 2D - a membrane:

$$E = \int \{(u-d)^2 + \lambda^2 (\nabla u)^2\}\, dx\, dy + P \qquad (2)$$

In each case there are three terms

1. A measure of faithfulness to data (the spring term).
2. A regularising term which depends on the gradient of the function.
3. A penalty term. In 1D this adds a penalty of α for each step discontinuity. In 2D the penalty is α multiplied by the length of the discontinuity.

The manner in which the parameters λ and α influence the behaviour is found by comparing the minimum energy solutions both with and without discontinuities for certain simple data[6]. The weak membrane will adopt the lowest energy configuration. The main conclusions are that

- The parameter λ is a characteristic length.
- The ratio $h_0 = \sqrt{2\alpha/\lambda}$ is a contrast threshold, determining the minimum contrast for detection of an isolated step edge.
- The ratio $g_l = h_0/2\lambda$ is a limit on gradient, above which spurious discontinuities may be marked. If the gradient exceeds g_l one or more discontinuities may appear in the fitted function.

In the first place we consider the modifications needed to make the weak string invariant, and describe how this affects the contrast threshold and gradient limit. This is then generalised to the (2D) invariant weak membrane.

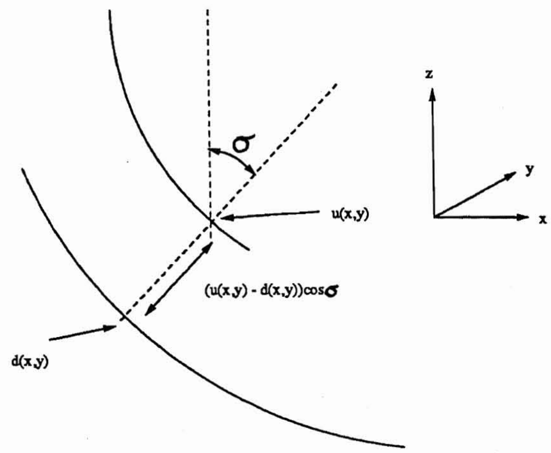

Figure 1: Invariant distance measurement: under the assumption that the surfaces u, d are roughly parallel, a fair estimate of perpendicular distance between them at (x, y) is $|u - d| \cos\sigma$.

2.1 Invariant weak string

The invariant version of (1) is

$$E = \int \{(u-d)^2 \cos^2\sigma + \lambda^2\}\, ds + P \qquad (3)$$

where

$$ds = \sqrt{1 + (u')^2} dx = \sec\sigma\, dx. \qquad (4)$$

There are two changes from (1)

1. The spring term is multiplied by $\cos^2\sigma$ (where $\tan\sigma = u'$).
2. The *arc length* is used instead of $(u')^2$ in the regularising term.

Ideally instead of $(u - d)$, which is tied to a particular reference frame, the perpendicular distance between $u(x)$ and $d(x)$ (which is invariant) should be used. The $\cos^2\sigma\, ds$ term in (3) is an imperfect attempt to do this. (this is illustrated in figure 1). The inclusion of the $\cos^2\sigma\, ds$ term reduces the contribution of the spring term especially when $\cos\sigma$ is small. This occurs when the data has high gradient. Consequently the reconstructed function will be further away from the data in high gradient regions, especially if the data is noisy.

Such a correction is appropriate under the assumption that "noise" in the system derives from the surface, from surface texture for example. However this is inappropriate if the noise derives from the sensor itself, because the noise then "lives" in the viewer (sensor) frame. Instead,

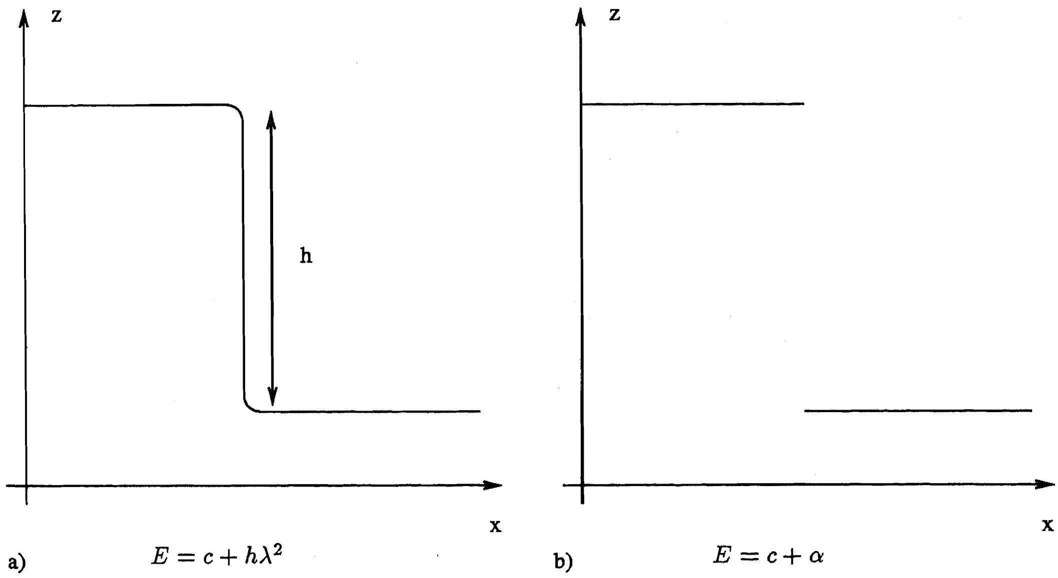

Figure 2: An isolated step (a) of height h. There are are two possible minima for the invariant weak string - a continuous function where $u(x)$ is close to $d(x)$ as in (a); or a step discontinuity (b).

the energy in that case is

$$E = \int \left\{ (u-d)^2 + \lambda^2 \sqrt{1+(u')^2} \right\} dx + P \quad (5)$$

The regularising term in (1) is the most obviously non-invariant part of the energy since it depends on gradient, which is clearly viewpoint dependent. The gradient can be small in one frame and (in theory) unbounded in another (for example at extremal boundaries - see later). Using the length as a regulariser, as in (3) and (5), entirely removes these problems.

Just as in the non-invariant case, there is a contrast threshold for the system. Its value is found by comparing the two possible minimum energies for data consisting of an isolated step of height h (see figure 2). If there is no discontinuity then to a first approximation (assuming $\lambda \ll h$) the energy is due solely to the length of the data

$$E = \lambda^2 (C + h) \quad (6)$$

where h is the length of the vertical part of the step and C is the length of the horizontal part. If there is a discontinuity at the step then the minimum energy is

$$E = \lambda^2 C + \alpha \quad (7)$$

Comparing the energies (6) and (7), the lowest energy solution will have a discontinuity if $h > \alpha/\lambda^2$.

Note there is some remission of the gradient limit g_l. This is simply because, when the surface gradient $g = u' \gg 1$, the length integrand above $\propto |g|$, rather than g^2 as for the non-invariant string. This is of most benefit near extremal boundaries where g becomes very large. This can be seen by considering a cylinder of radius r. From any viewpoint there will be an extremal boundary, and the gradient there will be extremely high (in theory unbounded). In the non-invariant string the gradient limit will certainly be exceeded and spurious discontinuities will be marked. However, in the invariant case the regularising term will be bounded by $\lambda^2 \pi r$ and so the problem is avoided. This point is clearly illustrated in the results of the range data segmentation (figure 4).

2.2 2D a weak membrane

In 2D the invariant energy is

$$E = \int \{(u-d)^2 \cos^2 \sigma + \lambda^2\} \, dS \; + P \quad (8)$$

where

$$dS = \sec \sigma \, dx \, dy = \sqrt{1 + u_x^2 + u_y^2} \, dx \, dy$$

and

$$P = \alpha \times \text{(actual length of discontinuities)},$$

Each of the terms has been made invariant

1. The spring term has been corrected as in the 1D case. σ is the the surface slant.
2. The regularising term depends on the surface area.
3. The penalty involves the actual length of the discontinuity - not simply the projected length.

Note that as for the non-invariant membrane the energy is invariant to rotations in the xy plane[7].

Again it is the regularising term which has the most severe effect if it is not invariant. The arguments in 1D concerning the contrast threshold and remission from the gradient limit are equally applicable here.

The penalty term could be made invariant by incorporating slant and tilt dependent compensation as was done by Brady and Yuille[8], for perimeter measurement under back projection. However in many circumstances the change in length of a discontinuity will be small when the viewpoint varies (Consider, for example, the projection of a circle changing to an ellipse under rotation).

3 Invariant plate

In this section the invariant forms of the plate (and, in 1D, the rod) are described. For the rod the non-invariant energy is given by

$$E = \int \{(u-d)^2 + \mu^4 (u'')^2\} dx + P \tag{9}$$

This differs from the energy of a weak elastic string (1) in including the *second* derivative of u rather than the first. The energy E is "second order". Furthermore, P incorporates a penalty β for each crease (discontinuity in du/dx) and a penalty α for each step (discontinuity in u).

The 2D version comes in two varieties[11]

Quadratic variation:

$$E = \int \{(u-d)^2 + \mu^4 (u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2)\}\, dx\, dy\ + P \tag{10}$$

Square laplacian:

$$E = \int \{(u-d)^2 + \mu^4 (u_{xx} + u_{yy})^2\}\, dx\, dy\ + P. \tag{11}$$

In fact any linear combination of these two is a feasible plate energy.

The invariant version for the rod is

$$E = \int \{(u-d)^2 \cos^2\sigma + \mu^4 (\kappa)^2\}\, ds\ + P \tag{12}$$

where κ is the curvature

$$\kappa = \frac{u''}{(1+(u')^2)^{3/2}} \tag{13}$$

The invariant equivalent of the square laplacian (11) is

$$E = \int \{(u-d)^2 \cos^2\sigma + \mu^4 (\kappa_1+\kappa_2)^2\}\, dS\ + P \tag{14}$$

where κ_1, κ_2 are principal curvatures. The squared sum of curvatures

$$(\kappa_1+\kappa_2)^2 = \frac{(Au_{xx} - 2Bu_{xy} + Cu_{yy})^2}{D^3}, \tag{15}$$

where

$$A = 1 + u_y^2\ B = u_x u_y\ C = 1 + u_x^2 \text{ and } D = \sec^2\sigma,$$

from[10]. Note that A, B, C, D are all functions of 1st derivatives of u only. If sum of squared curvatures is to be used (the invariant form of quadratic variation) then

$$\kappa_1^2 + \kappa_2^2 = (\kappa_1+\kappa_2)^2 - 2\kappa_1\kappa_2, \tag{16}$$

where, from[10] the gaussian curvature

$$\kappa_1\kappa_2 = \frac{u_{xx}u_{yy} - u_{xy}^2}{D^2}.$$

The effect of the D in the denominator is to reduce the influence of the regularising term. This has most effect when the surface gradients are large (for example noisy regions). In such regions the degree of smoothing in the viewer frame will be considerably reduced.

It can be shown[4] that, expressing the energy E as

$$E = \int \overline{E}(u_x, u_y, u_{xx}, u_{xy}, u_{yy})\, dx\, dy\ + P$$

$\overline{E}$ is a non-convex function of u_x, ... So even with a fixed set of discontinuities, it is not known whether there is an optimal u to be found[13].

To overcome this problem we suggest the following approximate method. First, estimates for $u_x(x,y), u_y(x,y)$ are obtained by fitting an invariant weak membrane (8) to the rangefinder data. This yields estimates of u_x, u_y which can be inserted as constants into $\overline{E}$ which is then convex with respect to u_{xx}, u_{xy}, u_{yy}.

Another possibility is to use gradient information directly in an invariant scheme for detecting creases. There are then two passes for detecting both crease and step discontinuities. For example in 1D the invariant weak string with a small λ (3) is used to detect steps and provide some noise rejection. The gradient estimates $u'(x)$ have been regularised and thus are well behaved (bounded). These are used for the data in a scheme with energy

$$E = \int \left\{ (u'-d')^2 \cos^2\sigma + \mu^2 \left[\frac{u''}{(1+(d')^2)^{3/2}}\right]^2 \right\} ds + P \tag{17}$$

This has the additional benefit numerically that both schemes are first order (rather than second) so that the convergence of the iteration schemes is typically improved by at least an order of magnitude[6].

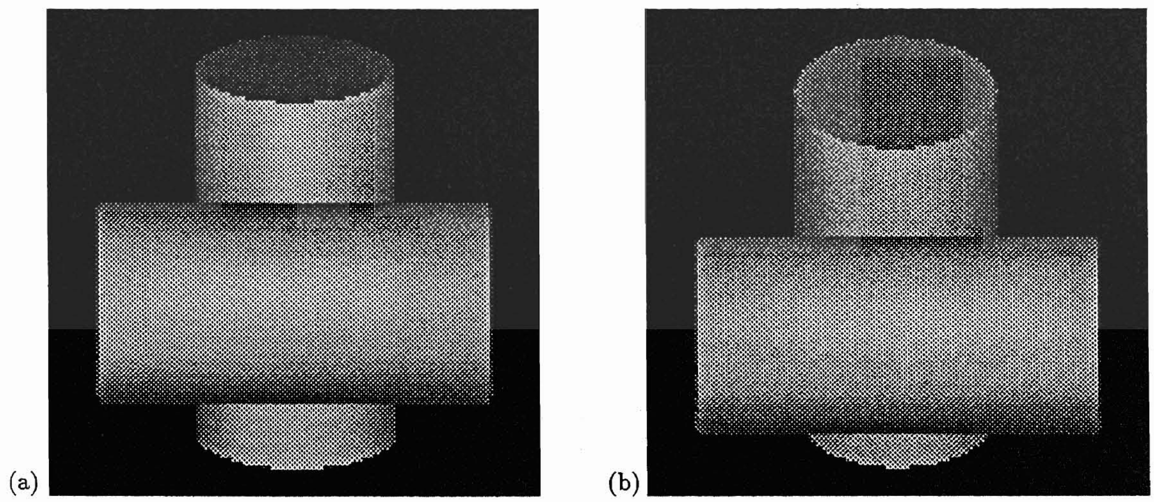

Figure 3: CSG image of two cylinders - (b) is (a) rotated by 15 degrees.

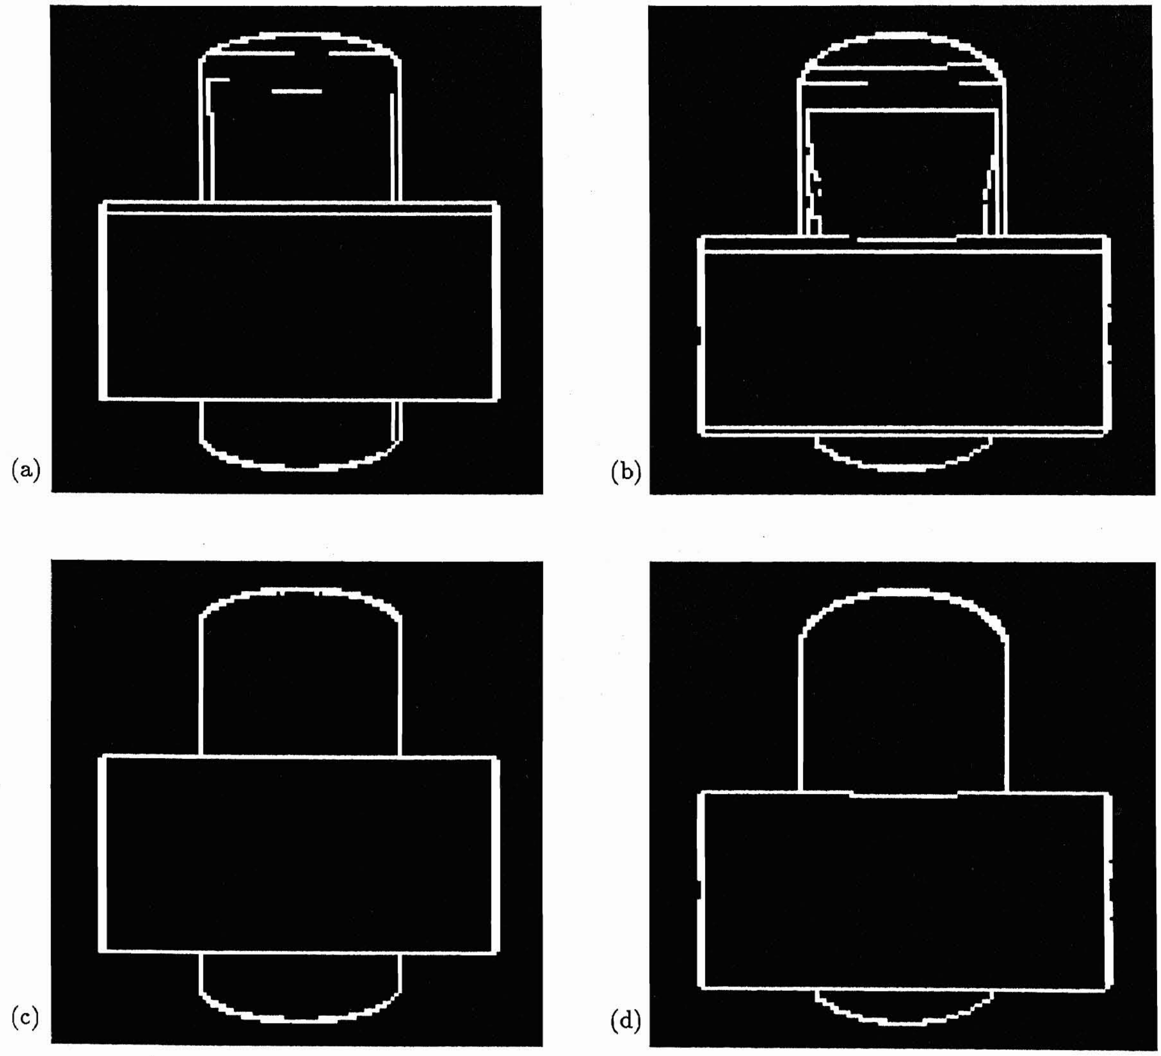

Figure 4: Edges detected in the depth data for the CSG cylinders (figure 3) . (a) and (b) are from a non-invariant membrane the double edges are where the gradient threshold is exceeded at the extremal boundaries. (c) and (d) are the edges detected using an invariant membrane.

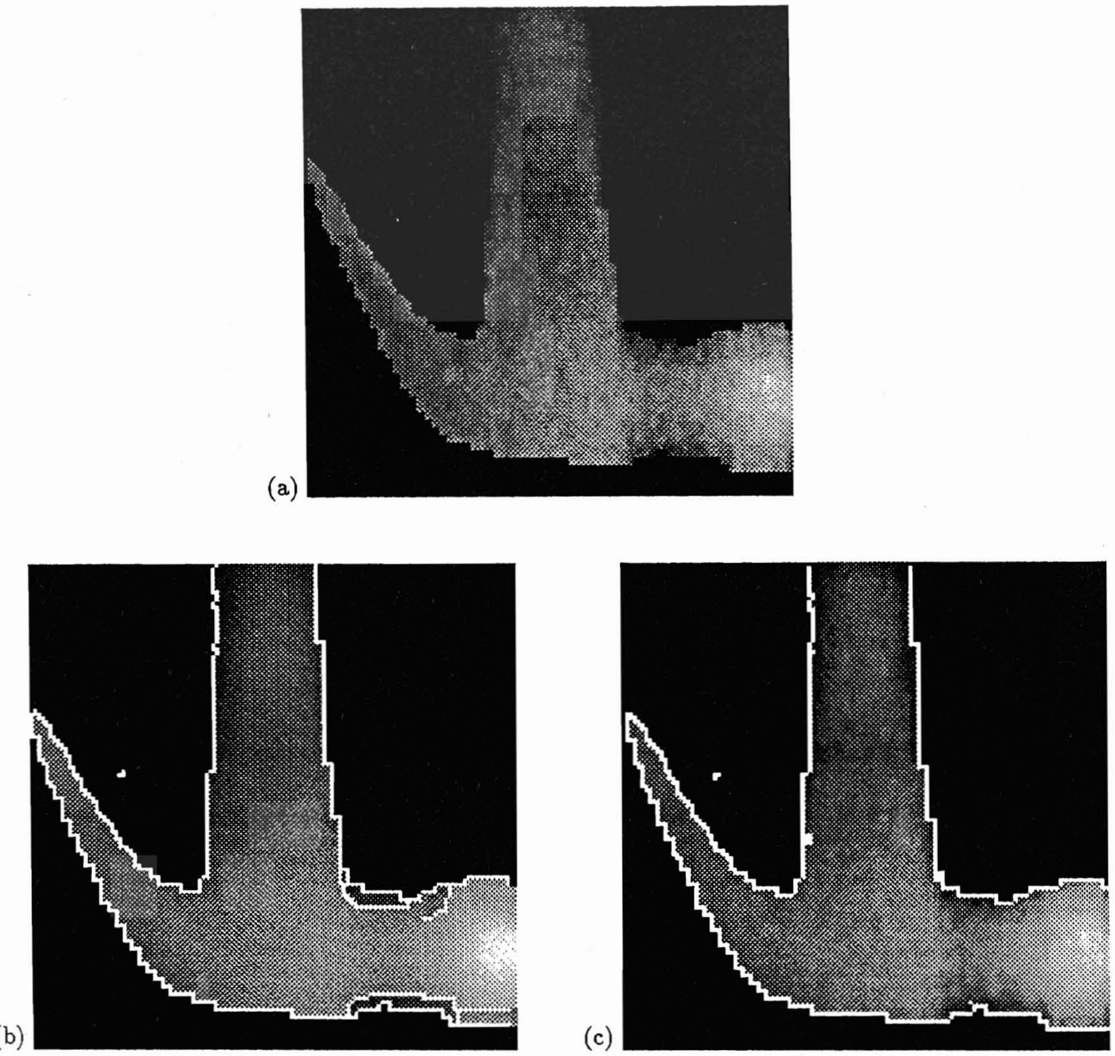

Figure 5: (a) data from a laser range finder for a hammer; (b) its edges using a non-invariant membrane; (c) its edges using an invariant membrane.

4 Results : segmentation of range data

Figure 3a shows pictures of two cylinders generated by a body modeller. Depth data, generated for this model, was segmented using both non-invariant and invariant membranes (figure 4). In both cases the localisation of the step discontinuities, both occluding and extremal, is good (even though the steps are not symmetric). At the extremal boundaries the non-invariant membrane incorrectly marks double edges (because the gradient threshold is exceeded) - but this is cured in the invariant case. Figure 3b is the model in 3a rotated by 15 degrees.

Finally figure 5 shows the laser range data for a hammer and its segmentations using the non-invariant and invariant membranes. Again the invariant membrane locates the extremal boundaries accurately despite the presence of considerable noise in some regions of the data.

The invariant scheme for the rod (17) has also been implemented. Both steps and creases are located accurately. Although the creases are very sensitive to noise. The invariant plate, using the approximate method described above, is currently being implemented.

Acknowledgments

This work was supported by SERC grant GR/D 1439.6 and by the University of Edinburgh. The Royal Society of London's IBM Research Fellowship supported A. Blake. We are very grateful to Bernard Buxton and John Canny for helpful comments, to IBM UK for the CSG modeller WINSOM, and to Gavin Brelstaff for the depth data. Thanks to Professor O. Faugeras and his group at INRIA, Versailles for laser rangefinder data.

References

[1] Besl, P. J. and Jain, R. C. (1986). Invariant surface characteristics for 3D object recognition in range data. *Comp. Vision, Graphics and Image Processing, 33, 33-80.*

[2] Blake, A. (1983). The least disturbance principle and weak constraints. *Pattern Recognition Letters, 1, 393-399.*

[3] Blake, A. (1983). *Parallel computation in low-level vision.* Ph. D. Thesis, University of Edinburgh.

[4] Blake, A. (1984). Reconstructing a visible surface. *Proc AAAI conf. 1984*, 23-26.

[5] Blake, A., Zisserman,A. and Papoulias,A.V. (1986). Weak continuity constraints generate uniform scale-space descriptions of plane curves. *Proc ECAI, Brighton, 1986.* Also reprinted in this volume.

[6] Blake A. and Zisserman A. *Visual Reconstruction.* MIT Press, 1987.

[7] Brady,M. and Horn,B.K.P. (1983). Rotationally symmetric operators for surface interpolation. *Comp. Vis. Graph. Image Proc.*, 22, 70-94.

[8] Brady,J.M and Yuille,A. (1984). An extremum principle for shape from contour. *IEEE trans. PAMI,* 6, 3, 288-301.

[9] Canny,J.F. (1983). *Finding edges and lines in images.* S.M. thesis, MIT, Cambridge, USA.

[10] do Carmo,M.P. (1976). *Differential geometry of curves and surfaces.* Prentice-Hall, Englewood Cliffs, New Jersey.

[11] Grimson,W.E.L. (1981). *From images to surfaces.* MIT Press, Cambridge, USA.

[12] Terzopoulos,D. (1983). Multilevel computational processes for visual surface reconstruction, *Computer Vision Graphics and Image Processing,* 24, 52-96.

[13] Troutman,J.L. (1983). *Variational calculus with elementary convexity.* Springer-Verlag, New York.

[14] Ullman,S. (1979). Relaxed and constrained optimisation by local processes. *Computer Graphics and Image Processing,* 10, 115-125.