Computer Science Department of The University of Auckland
CITR at Tamaki Campus (http://www.citr.auckland.ac.nz)

CITR-TR-68 September 2000

Interactions between Number Theory and Image Analysis

Yen Chen[1] and Radim Sara[2]

Abstract

In this project, we have investigated the fusion of surface data obtained by two different surface recovery methods. In particular, we have fused the depth maps obtained by shape from contours and local surface orientation maps obtained by photometric stereo. It has been found that the surface obtained by fusing orientation and depth data is able to yield more precision when compared with the surfaces obtained by either type of data alone.

---

[1] The University of Auckland, Tamaki Campus, Centre for Image Technology and Robotics, Computer Vision Unit, Auckland, New Zealand
[2] Centre for Machine Perception, Faculty of Engineering, Czech Technical University, Czech Republic

Integration of photometric stereo and shape from occluding contours by fusing orientation and depth data

Yen Chen[1] and Radim Šára[2]

[1]Centre for Image Technology and Robotics,
Tamaki Campus, University of Auckland,
New Zealand
yen@citr.auckland.ac.nz
[2]Centre for Machine Perception
Faculty of Electrical Engineering, Czech Technical University,
Czech Republic
sara@cmp.felk.cvut.cz

1 Abstract

In this project, we have investigated the fusion of surface data obtained by two different surface recovery methods. In particular, we have fused the depth maps obtained by shape from contours and local surface orientation maps obtained by photometric stereo. It has been found that the surface obtained by fusing orientation and depth data is able to yield more precision when compared with the surfaces obtained by either type of data alone.

Keywords: shape recovery, fusion, modelling, reconstruction

2 Introduction

Current surface recovery methods have their respective advantages and drawbacks. For example, while it is possible to obtain accurate measurements using structured lighting, the process can be extremely time consuming. Photometric stereo offers fast and dense recovery of the local surface orientations, but the depth values that are calculated by the integration [2, 4] of recovered normals may be inaccurate with respect to the true depth values.

To construct a new shape recovery method that can be more robust, efficient and versatile than existing methods, we fuse the data obtained by shape recovery methods that have complementary characteristics. From previous work [1], we have decided to construct a new shape recovery method by the fusion of depth and orientation data, which are respectively obtained by shape from occluding contours and photometric stereo method. These two methods have been chosen on the basis that shape from contours is able to provide reliable measurements, but unable to recover surface concavities that are occluded from the camera by the contours of the object. Conversely, photometric stereo provides dense orientation information over the surface, but the depth measurements obtained by integration of the surface orientations are relatively scaled to the actual depth values. Therefore, the integration of these two methods may be able to yield results with higher precision than either one of the methods is able to achieve.

We have approached the project in the following steps. Firstly, we generate the synthetic surfaces and simulate the orientation and depth data as would be obtained by the shape recovery methods. From the simulated data, we calculate the weighting functions for the different types of data. The fusing of orientation and depth data is then performed according to the weights for the orientation and depth data. Finally, we compare the surfaces obtained by fusion of data, as well as the surfaces obtained by photometric stereo and shape from contours, to evaluate the performance of the fusion method. The block diagram in Figure 1 shows the steps involved in the integration of photometric stereo and shape from contours by fusing orientation and depth data.

There are two objectives for this project. The first objective is to determine whether fusion of the two kinds of data is able to provide more precision than either one of these methods. This will be achieved by comparing

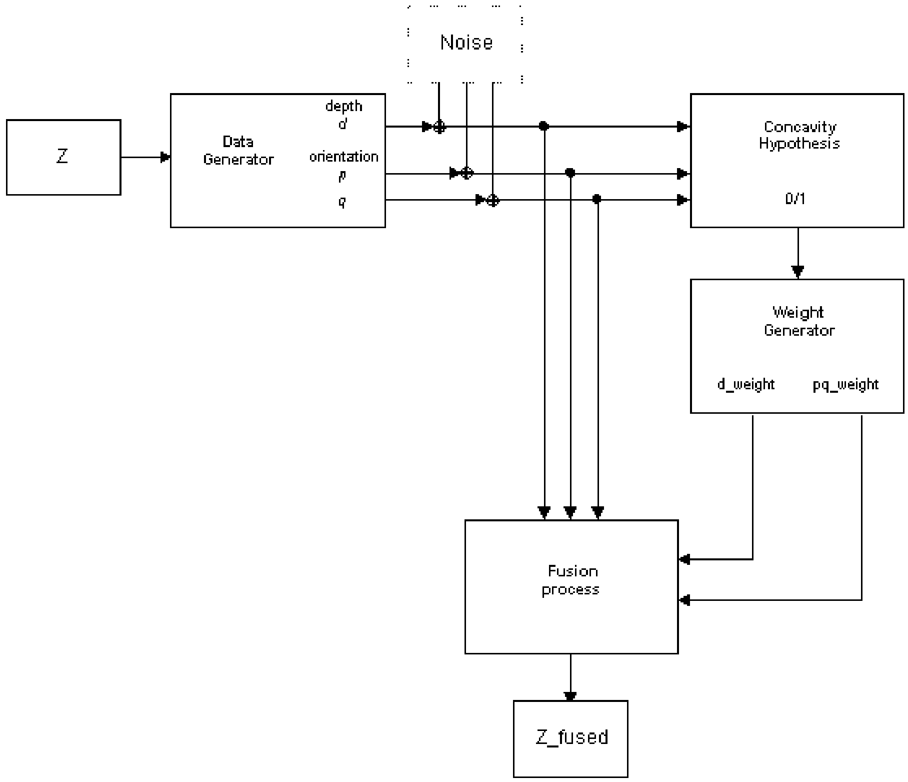

Figure 1: Block diagram for integration of shape recovery methods using fusion of depth and orientation data.

the recovered surfaces with with the true surfaces. The second objective is to determine whether there are any observable artefacts in the region where different data are fused. The surface recovered by fusion of data will be examined to achieve this objective.

3 Generation of surfaces

We have used MatLab for the generation of the surfaces. Each of the generated surfaces has a concave region which is occluded by a neighbouring region on the surface. Such surface concavity will not be recovered by shape from contours, yet the surface orientations within the concavity can be recovered by photometric stereo.

Two types of surfaces were used in this project. The first surface is generated by addition of Gaussian functions, such that the surface is C^2 continuous. It is an example of most general surfaces. The second surface, z_2, is a polyhedral surface generated by intersecting five planar surfaces. It is an example of a simple surface, such as the surface from a man-made object. Figure 2 shows the surfaces z_1 and z_2.

We have used both the continuous and polyhedral surfaces to evaluate the performance of the fusing method in different cases.

4 Simulation of orientation and depth data

The surface orientations as obtained by photometric stereo are provided by the partial derivatives of the surfaces.

The surface normals are calculated by approximating the derivatives of the original surfaces z_1 and z_2 with respect to the horizontal and vertical directions. For each surface, z_i, the local derivatives in the x and y directions are given by

$$\frac{\Delta z_i(x,y)}{\Delta x} = \frac{z_i(x+1,y) - z_i(x,y)}{\Delta x} \tag{1}$$

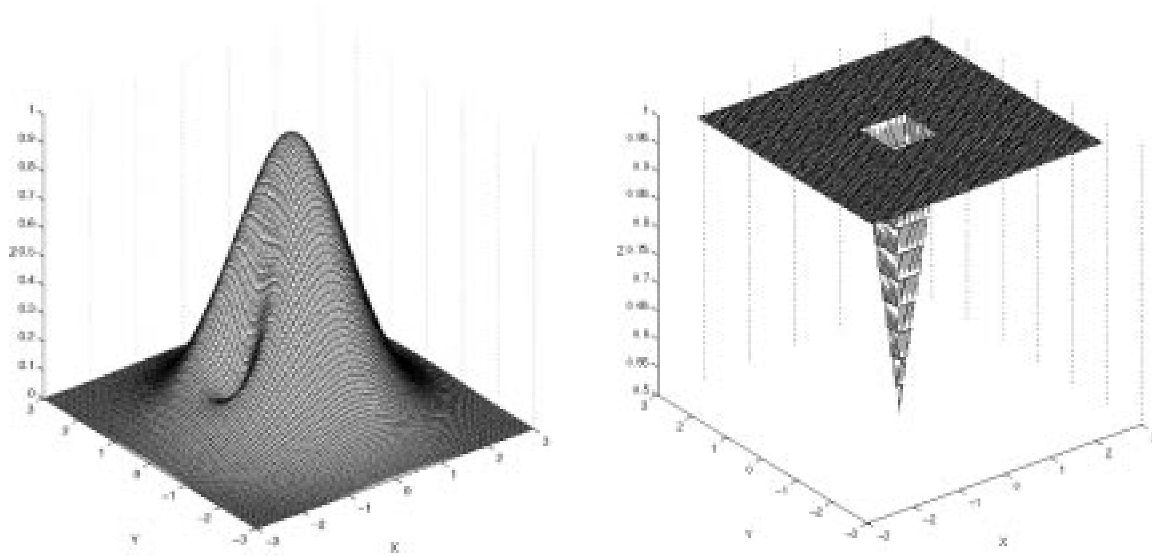

(a) (b)

Figure 2: Surface z_1 and surface z_2

$$\frac{\Delta z_i(x,y)}{\Delta y} = \frac{z_i(x,y+1) - z_i(x,y)}{\Delta y} \tag{2}$$

For our surfaces, Δx and Δy are both 0.05 units.

The orientation data are further simulated by adding random noise, N, to the derivatives. The resultant orientation data in the x and y directions are respectively given by

$$p_i(x,y) = \frac{\Delta z_i(x,y)}{\Delta x} + \alpha \cdot N(x,y)\,, \tag{3}$$

$$q_i(x,y) = \frac{\Delta z_i(x,y)}{\Delta y} + \alpha \cdot N(x,y)\,. \tag{4}$$

The simulated photometric stereo results are shown as needle maps for surfaces z_1 and z_2 are shown in Figure 3.

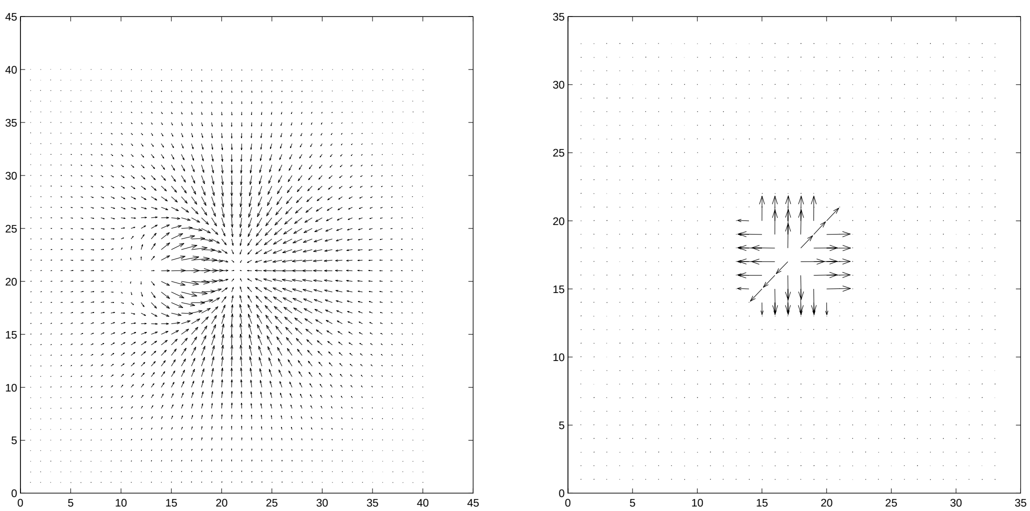

(a) (b)

Figure 3: Needle maps for z_1 and z_2

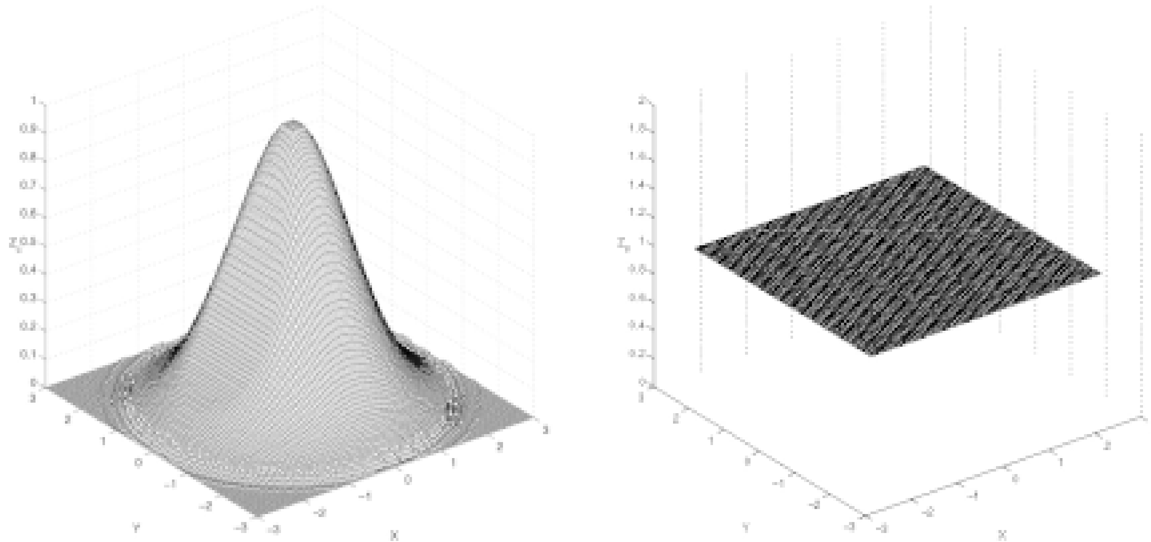

(a) (b)

Figure 4: Depth maps for z_1 and z_2

The amplitude of the original surfaces can be used directly as the depth values obtained by shape from contours, except for concaved regions on surfaces. The reason is that shape from contours is unable to recover concavities that are occluded from the viewing direction by other regions of the surface. Therefore, to simulate the depth data as obtained by shape from contours, we need to take into account the method's inability to recover surface concavities.

The surfaces are discretised into a given number of layers. For each layer, a convex polygon for the surface contour is drawn. The layers are then piled up again and the convex polygons are joined to form the surface, z_{conv}. Random noise N is added to further simulate the depth data obtained by shape from contours, given by

$$z_d(x, y) = z_{conv}(x, y) + \alpha \cdot N(x, y). \tag{5}$$

The obtained surface z_d has a developable patch covering the concavity on the original surface, as would be expected when shape from contours is used to recover the surface. The simulated surfaces are given in Figure 4.

5 Calculation of weights for fusing data

In this step, we determine the compatibility of the given surface data. The contribution of different data towards the final result is based on the computed data compatibility. This is a crucial step in the fusion of data, since suitable choice of weights enables the fusion of the complementing data to be performed in such a way that the surface given by fusion yields higher accuracy than surfaces given by either method alone.

There are two possible ways to compare the input data; either by using the depth values, or the surface orientations. We have decided to use surface orientations for the comparison of input data since it is easier to obtain. In this approach, partial derivatives are used to provide the surface normals for the surface recovered by shape from contours.

A cross section from surface z_1 is shown in Figure 5. The surface normals obtained by photometric stereo are represented by the surface normals on the solid line. The surface recovered by shape from contours is represented by the dotted line, and the surface normals calculated from the surface recovered by shape from contours are plotted on the dotted line.

From the figure, it can be seen that the orientations of surface normals along column a differ significantly, whereas the orientations of the surface normals are quite similar along column b. The difference of surface normals in column a is caused by the occurrence of a concavity. One way of calculating the angle between the two orientation vectors is to use the dot product of the surface normals. Therefore, the dot products of the normals can be used to indicate the discrepancies between these two types of data. The discrepancies caused by noise also need to be taken into account when determining the compatibility of the data. The compatibility function between the

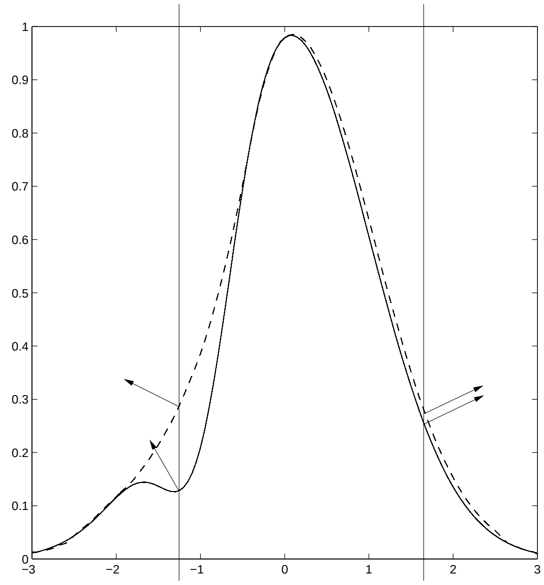

Figure 5: Cross section and surface normals of original surface and surface obtained by shape from contours.

data is a binary function given by

$$c(x,y) = \begin{cases} 1 & : \quad n_{psm} \cdot n_{sfc} < t\,, \\ 0 & : \quad otherwise \end{cases} \tag{6}$$

where t is the thresholding value that has been chosen such that the data discrepancies caused by noise can be avoided. The vectors n_{psm} and n_{sfc} respectively represent the surface normals obtained by photometric stereo and shape from contours.

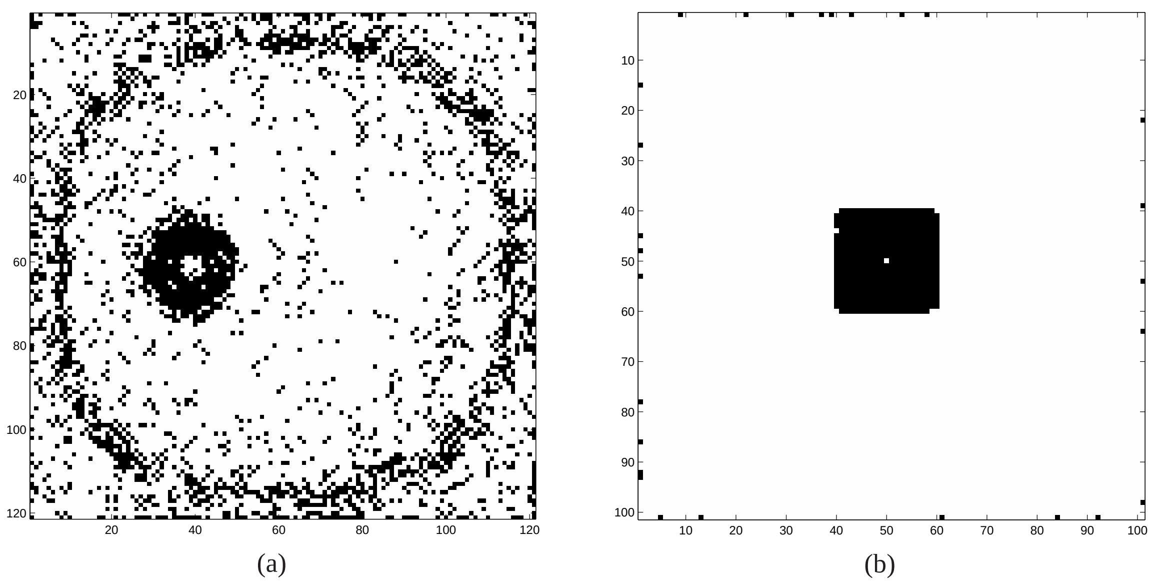

Figure 6: Calculated weights for z_1 and z_2

In Figure 6, white pixels represent regions where the data are compatible, and black pixels represent regions where the data are incompatible. From the figures, it can be seen that the general shape of the concavities have been detected for both surfaces. However, note that the surface normals from different data do agree in some positions within the concavity. This is due to the fact that for some points, the surface orientations within the concavity do agree with the surface orientations for the developable patch. Therefore, the compatibility function will need to be improved further to enable more accurate detection of concavity regions.

The data weighting functions are computed based on the compatibility of the data. In this work, the weighting function for the depth data,

$$w_d(x,y) = \begin{cases} 1 & : \quad c(x,y) = 1 \\ 0 & : \quad c(x,y) = 0 \end{cases} \tag{7}$$

is the same as the compatibility function, which is 1 when the data are compatible, and 0 when the data are incompatible.

The weighting function for the orientation data,

$$w_{pq}(x,y) = \begin{cases} 1 & : \quad c(x,y) = 0 \\ 0 & : \quad c(x,y) = 1 \end{cases} \tag{8}$$

has values of 0 in regions where the data are compatible and 1 where the data are incompatible.

The weighting functions are thus selected such that the unknown surface will generally be recovered from the depth values obtained by shape from contours, except for the concave regions, where the surface will be recovered according to the surface orientations obtained by photometric stereo.

6 Fusion of data

The fusion algorithm is implemented according to the work by D. Terzopoulos [3, 6]. It reconstructs the unknown C^2 continuous function z by combining the depth and orientation data with respect to certain weighting functions. The input parameters are d, the depth data, p and q, the horizontal and vertical orientation data, as well as w_d and w_{pq}, the weighting functions for depth and orientation data. In this work, the depth data, d, are the noise corrupted depth values obtained using shape from contours. The orientation data, p and q, are the noise corrupted orientation data obtained by photometric stereo.

The fusion algorithm minimises the error function

$$E(z) = E_d(z,d,p,q) + E_m(z) + E_t(z)\,, \tag{9}$$

where

$$E_d(z,d,p,q) = w_d \cdot S_d \sum_u \sum_v |z-d|^2 + w_{pq} \cdot S_{pq} \sum_u \sum_v |\frac{\delta z}{\delta u} - p|^2 + |\frac{\delta z}{\delta v} - q|^2 \tag{10}$$

$$E_m(z) = S_m \int\int (\frac{\delta z}{\delta u})^2 + (\frac{\delta z}{\delta v})^2 du dv \tag{11}$$

$$E_t(z) = \int\int [\frac{\delta^2 z}{\delta u^2}\frac{\delta^2 z}{\delta v^2} - (\frac{\delta^2 z}{\delta u \delta v})^2] du dv\,. \tag{12}$$

In the above formula, E_d is the data error function, which specifies how well the unknown surface z conforms to the input depth and orientation data. The second term, E_m, represents the membrane function. It specifies that the surface z should be continuous. The third term, E_t, represents the thin-plate function, which specifies that the surface z should have small changes in its curvature, that is, the surface should be smooth, with no sudden peaks. The values S_d, S_{pq}, and S_m are the coefficients of the depth error, orientation error and the continuity constraint, respectively. These values can be assigned by the designer of the algorithm. In our case, the coefficients have all been set to 1.

The conjugate gradient descend method is used to iteratively minimise the error function, such that the result is obtained when the value of the error function is at its minimum. The error function is also convex, so that the algorithm will always converge to the global minimum.

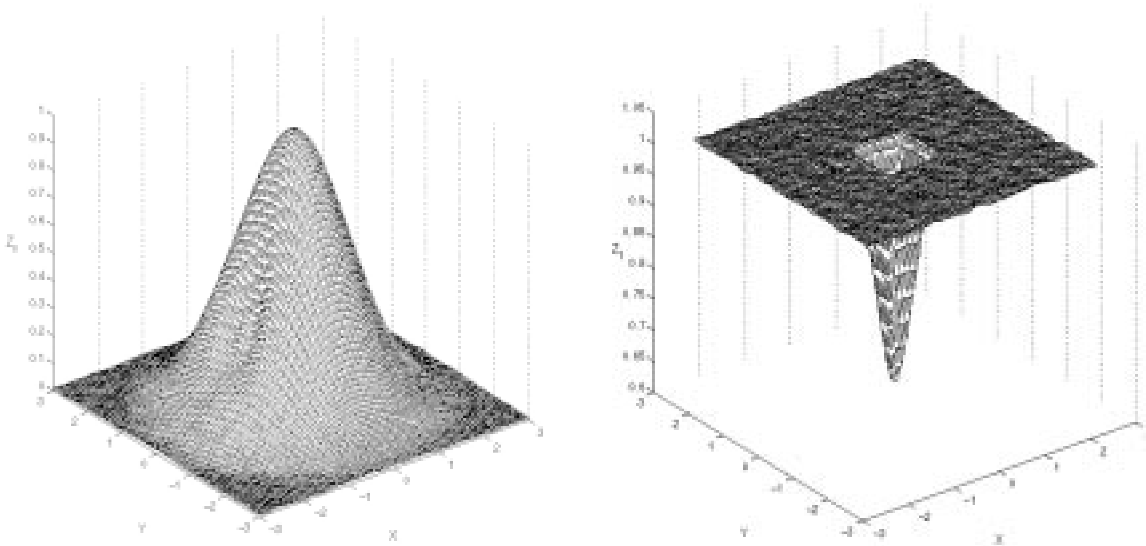

(a) (b)

Figure 7: Fused surfaces for z_1 and z_2

7 Results

After applying the fusion algorithm to the two different types of data, we obtained the surfaces as shown in Figure 7.

Figure 7(a) shows the surface obtained by fusion of orientation and depth data for surface z_1. From the figure, we can see that the shape and amplitude of the surface recovered by fusion, are quite similar to the original surface. Also notice that there are no visible artefacts along the concavity boundary in the resultant surface.

The surface recovered by fusion of data for surface z_2 is shown in Figure 7(b). The recovered surface also retains the general shape of the original surface, but the amplitude of the recovered surface differs slightly from the original surface. There are also no artefacts on the fused surface along the concavity boundary. In fact, the recovered surface is a continuous surface, whereas the original surface has discontinuous folds along the concavity boundary and within the concave region. The reason is that the fusion method assumes that the unknown function z has to be C^2 continuous, hence the sharp folds of the original surface are not recovered.

Figure 8(a) shows cross sections of the original surface z_{true} and the surfaces recovered using different methods. Figure 8(b) shows the absolute errors between the fused surface and the true surface for the same cross section.

In the cross section shown, we can see that the surface z_{psm}, recovered by photometric stereo, conforms to the original surface quite well in the concave region. Whereas the surface z_{sfc}, recovered by shape from contours, has good conformity to the original surface in most regions, except for the concavity. The surface recovered by the fusion of data, z_{fused}, lies between the surfaces recovered by photometric stereo and shape from contours. Like the surface recovered by shape from contours, the surface recovered by fusion conforms to the original surface quite well in most regions. But unlike z_{sfc}, z_{fused} has a concavity in the concaved region, even though the concavity is less emphasised when compared to the original surface. The insufficient depth of the recovered concavity is a major source of error in the surface recovered from fusion. The shallower concavity is partly due to the fact that the weighting functions are binary, thus the contribution from each type of data is either 0 or 1. The smoothness constraint of the recovered surface is another cause for the insufficiently recovered concavity. Nevertheless, the depth of the recovered concavity can be adjusted by altering the the weighting functions, such that more weight is placed on the orientation data during the fusion process.

8 Evaluation

In this section, the recovered surfaces are compared with the ground truth to evaluate the accuracy of each method.

The error function is given by $e = z - \tilde{z}$, where z is the surface obtained by fusion of original, un-corrupted

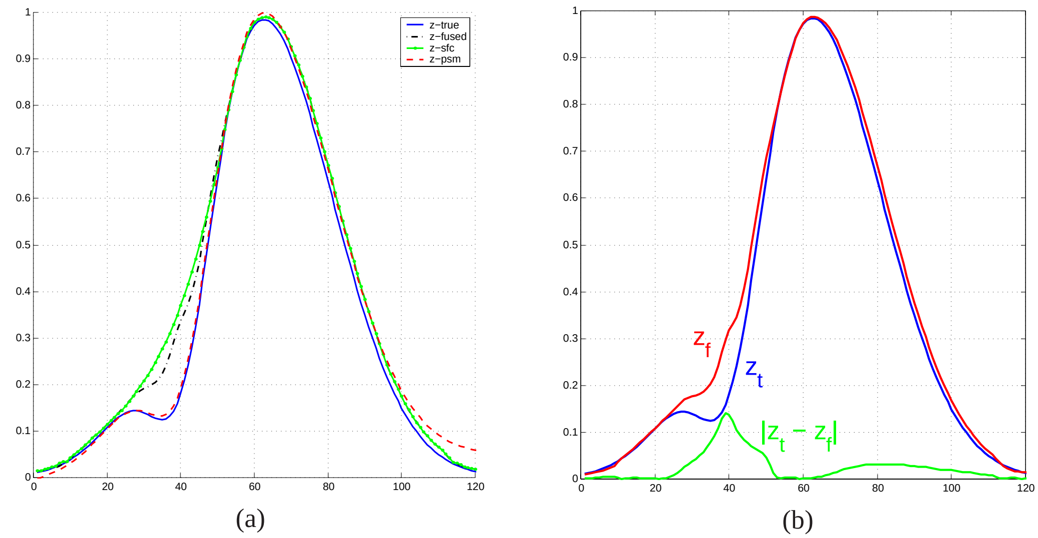

Figure 8: Cross sections of (a) resultant surfaces and (b) error

data, and $\tilde{z}$ is the surface recovered from simulated data, which are corrupted by noise.

For surface z_1,

(%)	Fused	SFC	PSM
$min\,\|e\|$	0.0005	0.0018	0.0000
$max\,\|e\|$	16.5979	19.3352	7.8308
$mean\,\|e\|$	1.5120	1.5617	2.5173
RMS	1.5630	1.8562	1.7958

For surface z_2,

(%)	Fused	SFC	PSM
$min\,\|e\|$	1.1752	1.1769	0.0070
$max\,\|e\|$	24.9002	101.5993	15.4209
$mean\,\|e\|$	2.3244	3.5532	5.3134
RMS	1.1657	8.5130	3.6062

It can be seen from the tables that the surfaces recovered by fusing of data are more accurate than the surfaces recovered by either shape from contours or photometric stereo alone.

9 Discussion

In this work, we have tested the fusion method, as well as photometric stereo and shape from contours on two different types of surfaces. The first surface is a C^2 continuous surface constructed from Gaussian functions. The second surface is a polyhedral surface constructed from planar patches. The C^2 continuous surface is an example of general surfaces, and the polyhedral surface is an example of simple surfaces.

The orientation and depth data are obtained by simulating the results as would be obtained by photometric stereo and shape from contours. In simulating the orientation data recovered by photometric stereo, errors are introduced because the differences in surface height between neighbouring positions are used to approximate the partial derivatives for surface orientations. In the simulation of depth data, the discretisation of surface heights also lead to the introduction of errors into the depth values. Apart from the errors introduced by the discretisation of data, random noises were further added to the simulated data.

Once the orientation and depth data are obtained, they are used to calculate the data compatibility function.

The main purpose of the compatibility function is to determine the concavity region on the surface. In this work, the compatibility of the different data is based on the dot product of the surface normals calculated from photometric stereo and shape from contours. However, the compatibility function calculated by this approach is unable to provide all of the points that lie within the concavity region, since at some positions within the concavity, the surface normals obtained by photometric stereo agree with the normals obtained by shape from contours. One way to improve the detection of concavities is to take the compatibilities of neighbouring positions into consideration. For example, if the neighbours of a point are all incompatible, yet the point itself is compatible, then it is highly possible that the point lies within the concavity. The compatibility function can be given as the confidence with which different types of data agree and has continuous values rather than binary. In which case, the confidence can be indicated by the difference between two types of data, as well as the compatibility of neighbouring positions.

The weighting functions are calculated according to the data compatibility function. The data weighting functions should be constructed such that the unknown surface will generally be recovered from the depth data provided by shape from contours, since shape from contours provides reliable dimensions of the object. However, in concavities, the unknown surface needs to be recovered according to the surface orientation data, because photometric stereo is able to provide accurate orientation data within the concavities. In this work, we have used binary weighting functions for the purpose of evaluation. But such inflexibility of weighting values has caused the recovered concavities to be shallower than the original concavities. One approach to improve the result is to increase the orientation data weighting function in the concavity region. However, more generally, the weighting function can be made to vary continuously with respect to the data compatibility. For example, the depth weighting function may vary proportionally to the data compatibility, such that the weight increases as the confidence in data compatibility increases. The orientation data weighting function may vary inversely with the data compatibility, and with higher values in the concavity regions to emphasis the recovered concavities.

The orientation and depth data, as well as the respective weights are given to the fusion algorithm to recover the unknown surface. The fusion algorithm has assumed that the unknown surface is C^2 continuous, which might not always be the case, as can be seen from the smoothing of the folds in the recovered surface for the simple surface z_2. The coefficients of the different constraining error functions have all been set to 1 in this work. These values may need adjustment to lessen the effect of the continuity or smoothness constraints.

It has been seen from the result that shape from contours is able to provide accurate depth recovery of the surface except for the concavity regions. The maximum error occurs in the concavity regions since the concavity cannot be recovered from the contours alone. On the other hand, the photometric stereo method is able to recover the surface concavities with accuracy, but the surface recovered from the orientations alone is more erroneous overall, as indicated by the mean errors. The surface recovered by fusing depth and orientation data is more accurate than the surfaces recovered using either depth or orientation data alone. The surface obtained by fusion also has no visible artefacts along the concavity boundary, where the transition from one type of data to the other occurs.

Our next task is to apply the fusion method on data obtained using shape from contours and photometric stereo, rather than simulated data. The errors contained in the recovered orientation and depth data cause difficulties in the fusion of real data. Therefore, the functions involved in data fusion will need to be more robust and resistant to errors. The data compatibility function requires improvement to provide reliable concavity detection in spite of the noisy data. The weighting functions may need to take the variance of the input data into consideration to avoid emphasising erroneous data. The error functions used in the fusion algorithm also need to be modified to handle real data. Furthermore, if the variances of the errors are known, they can be incorporated into the fusion process to compensate for the errors in the fused result.

10 Conclusion

This work is a preliminary step towards the integration of photometric stereo and shape from occluding contours. In this work, we have performed fusion on simulated surface data and acquired more accurate surface recovery for different types of surfaces. The data being fused are the orientation and depth data obtained by simulating the photometric stereo method and the shape from contours, respectively.

Our first object is to determine if surface obtained by fusing orientation and depth data is more accurate than surface obtained by either type of data alone. This has been achieved by quantitatively comparing the surfaces recovered from simulated data with the surface obtained from original data. It has been found that the surfaces recovered by fusion is more accurate than the surfaces recovered by either photometric stereo or shape from contours.

The second object is to see if there are any observable artefacts along the concavity boundary where different types of data are fused. A sharp transition of different types of data occur along the concavity boundary, since the surface orientations obtained by photometric stereo is used to complement the inability to recover concavities in shape from contours. By examining the surface recovered by fusion, it has been found that there are no observable artefacts along the concavity boundary.

Finally, we discussed possible modifications for the procedures involved in fusing orientation and depth data, such that the fusion method may be applied on real data.

References

[1] C.Y. Chen, K.Schlüns and R. Klette, *Comparative 3D shape recovery for archaeological artefacts*, Proc. of IASTED International Conference on Applied Modelling and Simulation, pp577-582, Sept. 1999,

[2] R.Frankot and R.Chellappa, *A method for enforcing integrability in shape from shading algorithms*, IEEE Trans. on Pattern Analysis and Machine Intelligence, 10-4 pp.439-451, 1988.

[3] W. Grimson, *A computational theory of visual surface interpolation*, Phil. Trans. R. Lond., B298 395-427, 1982.

[4] R.Klette, K.Schlüns, *Height data from gradient fields*, Proc. SPIE, 2908, pp.204–215, 1996.

[5] J. Stoker, *Differential Geometry*, New York, pp. 74-94, 1989.

[6] D. Terzopoulos, *Multi-level reconstruction of visual surfaces*, Massachusetts institute of technology artificial intelligence laboratory, April 1982.