



National Library
of Canada

Acquisitions and
Bibliographic Services Branch

395 Wellington Street
Ottawa, Ontario
K1A 0N4

Bibliothèque nationale
du Canada

Direction des acquisitions et
des services bibliographiques

395, rue Wellington
Ottawa (Ontario)
K1A 0N4

Your file Votre référence

Our file Notre référence

NOTICE

The quality of this microform is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Reproduction in full or in part of this microform is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30, and subsequent amendments.

AVIS

La qualité de cette microforme dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

La reproduction, même partielle, de cette microforme est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30, et ses amendements subséquents.

**Edge Localization in Surface
Reconstruction Using Optimal
Estimation Theory**

Shailendra Mathur

B. Eng. McGill University 1989

Department of Electrical Engineering

McGill University

Montreal

March, 1996

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of
Master of Engineering

© Shailendra Mathur, 1996



National Library
of Canada

Acquisitions and
Bibliographic Services Branch

395 Wellington Street
Ottawa, Ontario
K1A 0N4

Bibliothèque nationale
du Canada

Direction des acquisitions et
des services bibliographiques

395, rue Wellington
Ottawa (Ontario)
K1A 0N4

Your file Votre référence

Our file Notre référence

The author has granted an irrevocable non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of his/her thesis by any means and in any form or format, making this thesis available to interested persons.

L'auteur a accordé une licence irrévocable et non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de sa thèse de quelque manière et sous quelque forme que ce soit pour mettre des exemplaires de cette thèse à la disposition des personnes intéressées.

The author retains ownership of the copyright in his/her thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without his/her permission.

L'auteur conserve la propriété du droit d'auteur qui protège sa thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

ISBN 0-612-12128-3

Canada

Abstract

In this thesis the problem of localizing discontinuities while smoothing noisy data is solved for the surface reconstruction method known as Curvature Consistency. In this algorithm, noisy initial estimates of surface patches are refined according to a continuity model, using a relaxation process. The interaction between neighbouring pixels in local neighbourhoods during relaxation is shown to be equivalent to a multiple measurement fusion process, where each pixel acts as a measurement source. Using optimal estimation theory as a basis, an adaptive weighting technique is developed to estimate interpolant surface patch parameters from neighbouring pixels. By applying the weighting process iteratively within local neighbourhoods, discontinuities are localized, and a piecewise-continuous surface description is achieved. The resulting discontinuity localization algorithm is adaptive over different signal to noise ratios, robust over discontinuities of different scales and independent of user set parameters.

Résumé

Cette thèse résout le problème de localisation des discontinuités de la méthode de reconstruction de surface 'Curvature Consistency', tout en lissant les données bruitées. Dans cet algorithme, les estimés initialement bruités des méthodes de reconstruction de surfaces sont améliorées pour satisfaire un modèle de continuité en utilisant un procédé de relaxation. L'interaction entre les pixels voisins pendant le procédé de relaxation est présenté comme étant équivalent à un procédé de fusion de données où chaque pixel agit comme une source de mesure. En se basant sur une théorie de mesure optimale, une méthode de pondération adaptative est utilisée pour évaluer des informations provenant des différents pixels voisins. En appliquant cette méthode itérativement, les discontinuités sont graduellement détectées et une description de surface localement continue est réalisée. La détection des discontinuités résultante est adaptative aux différents rapports de signal à bruit, robuste pour différentes échelles de discontinuités et pratiquement libre de paramètres à définir par l'utilisateur.

Acknowledgments

A first note of thanks goes to my thesis supervisor Frank Ferrie, who has shown incredible patience and support during the development of this thesis. In addition, thanks goes to my group-mates in the computer vision lab at CIM for the stimulating and friendly environment in which we all co-existed. Special thanks go to Gilbert Soucy, who has to be the most helpful guy around at CIM. The many fruitful discussions with Gilbert, Andre Lejeune, Pierre Tremblay and Duncan Baird have contributed significantly to ideas in this thesis.

My special appreciation goes to my parents Shashi and Lakshmi Mathur and my wife Julie for the support and extra push I always needed to finish this thesis.

Table of Contents

Chapter 1	Introduction	1
1.1	Motivation	2
1.2	Overview	6
1.3	Contribution	9
1.4	Organization of Thesis	10
Chapter 2	Literature Overview	12
2.1	Adaptive Noise Smoothing Techniques	13
2.2	Visual Reconstruction in the Regularization Framework	17
2.3	Use of Optimal Estimation Theory in Visual Reconstruction	23
Chapter 3	The Curvature Consistency Algorithm	26
3.1	The Local Surface Representation	27
3.2	Initial Estimates	28
3.3	Iterative Refinement Using Curvature Consistency	29
3.4	Applying Curvature Consistency to Range Images	31
3.5	Chapter Summary	35
Chapter 4	An Optimal Estimation Problem	36
4.1	Optimal Estimation Theory	37
4.2	An Optimal Estimation Problem	38
4.3	State Space Formulation	40
4.3.1	Accommodating Sequences of Measurements	41
4.3.2	Recursively Estimating the Noise Variance	42
4.3.3	Sources With Changing Noise Properties	43
4.4	Chapter Summary	44
Chapter 5	Using Estimation Theory In Curvature Consistency	46

5.1	Control Processes at the Pixel Level in Curvature Consistency	47
5.1.1	Information Flow in Relaxation Processes	48
5.1.2	Defining the Noise Property	49
5.1.3	Optimally Combining Measurements From Neighbours	52
5.1.4	Optimal Estimation in the Curvature Consistency Formulation	55
5.1.5	Modification to the Update Equations	56
5.2	Accommodating Pixel Groupings in Discontinuous Regions	60
5.2.1	Mapping Error Variances to a New Function	62
5.2.2	Making γ Adaptive	64
5.3	Chapter Summary	68
Chapter 6	Results and Discussion	70
6.1	Qualitative Analysis	71
6.1.1	General Behaviour of the Algorithm	72
6.1.2	Localization of a Step Edge	72
6.1.3	Localization of a Roof Edge	77
6.1.4	Discontinuity Localization in Complex Images	81
6.2	Iteration Control	84
6.3	Adaptiveness of the Modified Algorithm	88
6.3.1	Scale Space Robustness	89
6.3.2	Noise Level Robustness	93
Chapter 7	Conclusions	97
Bibliography	99

List of Figures

1.1	(a) Gaussian noise added to a synthetic step edge. Distortion in the step after (b) 30 iterations of curvature consistency algorithm, (c) at convergence (100 iterations). (d) Reconstruction using the edge localization method at convergence.	4
1.2	(a) Gaussian noise added to a synthetic roof edge. Distortion in the step after (b) 20 iterations of curvature consistency algorithm, (c) at convergence (100 iterations). (d) Reconstruction using edge localization method after 20 iterations.	5
3.1	Local surface representation - the augmented Darboux frame	27
3.2	The local transport model determines how a frame at a neighbouring point Q appears when it is extrapolated over to P	30
3.3	(a) A synthetic range image. (b) Noise added to the range points (SNR 5/3). (c) The initial estimate of the darboux frames and surface patches. Darboux frames and recovered surface after (d) 3 iterations, (e) 10 iterations of the curvature consistency algorithm. (f) Energy residual plot.	33
3.4	(a) A real range image of a toy car. (b) Initial estimate of the surface patches. Recovered surface after (c) 3 iterations, (d) 5 iterations, (e) 10 iterations of the curvature consistency algorithm. (f) Plot of energy residual. . .	34
5.1	Direction of information flow into a 3×3 neighbourhood during a relaxation process	48
5.2	(a) A simulated noisy range image of a step junction. (b) Initial estimate of the surface patches. (c) Underlying surface recovered using the original curvature consistency algorithm after 40 iterations. (d) Surface reconstruction using the curvature consistency modified to incorporate variance weighting. . .	59

5.3	A 5×5 neighbourhood mask straddling a (a) step discontinuity, (b) a roof discontinuity. The mask is centered at P. The discontinuity divides the neighbourhood into two subregions A and B	62
5.4	Distribution of weights according to a Gaussian error function	63
5.5	Distribution of weights for the depth component according to a Gaussian function amongst pixel sources in a 3×3 (a) continuous neighbourhood, and (b) discontinuous neighbourhood. The 'x' symbols on the X-axis represent the error variance estimate of the different pixel sources; the '*' symbols represent the corresponding weights assigned according to the Gaussian function; the 'c' symbol on the X-axis is the mean of error variances of all the pixel sources; the '+' symbol shows the corresponding mapping of this mean onto the Gaussian trace.	66
5.6	(a) A simulated noisy range image of a step junction. Surface reconstruction using a 5×5 neighbourhood size after 40 iterations (b) using the original curvature consistency algorithm, (c) using variance weighting in the algorithm, (d) using Gaussian weighting.	67
6.1	Changes in the surface structure around discontinuities with progressive iterations of the modified curvature consistency algorithm, using a 5×5 mask, after (a) initial patch estimate, (b) 2 iterations, (c) 4 iterations, (d) 10 iterations, (e) 20 iterations, (f) 50 iterations	74
6.2	The estimation errors and the corresponding weights for the position (depth in range image) component of the Extended Darboux Frame at a pixel straddling the discontinuity using the modified curvature consistency algorithm. Using a 5×5 mask after (a) 1 iteration, (b) 4 iterations, (c) 10 iterations, (d) 20 iterations.	75

6.3	The estimation errors and the corresponding weights for the position (depth in range image) component of the Extended Darboux Frame at a pixel in a continuous region of the surface data using the modified curvature consistency algorithm. Using a 5×5 mask, after (a) 1 iteration, (b) 4 iterations, (c) 10 iterations, (d) 20 iterations.	76
6.4	Surface reconstruction of a simulated noisy range image containing a roof discontinuity. (a) The original image, (b) initial estimate of surface patches. Using a 5×5 mask, distortion of the edge after (c) 10 iterations and (d) 20 iterations of the original curvature consistency. Preservation of the edge after (e) 10 iterations and (f) 20 iterations of the modified algorithm.	78
6.5	The estimation errors and the corresponding weights for the Normal component of the Extended Darboux Frames at a pixel straddling the roof discontinuity while applying the modified curvature consistency algorithm. Using a 5×5 mask, the estimation error variance and corresponding weights after (a) 1 iteration, (b) 5 iterations, (c) 10 iterations, (d) 20 iterations.	79
6.6	Surface reconstruction of an image containing discontinuity in position, as well as curvatures and normals. (a) Original data of a discontinuous sinusoidal trace in grid form, (b) initial fit of surface patches to the image. Using a 3×3 mask, (c) surface reconstruction after 20 iterations of the original curvature consistency algorithm, (d) surface reconstruction after 20 iterations of the modified algorithm.	81
6.7	(a) Range image of an owl statue acquired with the NRCC/McGill laser range-finder. (b) Reconstructed surface after 8 iterations of the original algorithm. (c) Reconstructed surface after 8 iterations of the modified algorithm.	82
6.8	Surface reconstruction of an image of two overlapping rocks, acquired using the NRCC/McGill University laser range-finder; (a) Original range data in grid form; (b) Initial fit of surface patches to the image; (c) Surface reconstruction after 50 iterations of the original curvature consistency algorithm; (d) Surface reconstruction after 50 iterations of the modified algorithm.	83

6.9	Comparison of global error of fit when the original (solid line) and modified (dashed line) curvature consistency algorithm are applied to the image containing (a) the step discontinuity, (b) the roof discontinuity.	87
6.10	Scale space performance of the original curvature consistency algorithm on an image containing discontinuities of varying scales. (a) Original image, surface reconstruction using a (b) 3×3 mask, (c) 5×5 mask, (d) 7×7 mask, (e) 9×9 mask.	91
6.11	Scale space performance of the modified curvature consistency algorithm on an image containing discontinuities of varying scales. (a) Original image, surface reconstruction using using a (b) 3×3 mask, (c) 5×5 mask, (d) 7×7 mask, (e) 9×9 mask.	92
6.12	Noise adaptation of the modified curvature consistency algorithm. (a) Original image with noise of standard deviation 1 (outermost), 2, 3 and 4 (innermost) added to the surface regions. (b) Surface reconstruction after 25 iterations using a 5×5 mask.	94
6.13	Noise space adaptation of the modified curvature consistency algorithm at multiple scales. Image contains regions with signal to noise ratios of 10 (outermost region), 3.33, 2, and 1 (middle region). (a) Original image, surface reconstruction using a (b) 3×3 mask, (c) 5×5 mask, (d) 7×7 mask, (e) 9×9 mask.	95

This thesis presents a new way of looking at inter-pixel interaction during relaxation processes. Using this approach, a novel method is proposed for the localization of discontinuities during surface reconstruction from noisy range data. The neighbourhood-level processing during relaxation is shown to be equivalent to a multiple measurement fusion problem which is solved using optimal estimation theory. Pixels in a given neighbourhood act as information sources, combining their information in some fashion to update the state of that neighbourhood. The surface reconstruction algorithm of focus is the so called curvature consistency algorithm [12] [13] [16] [20] [55] [56]. Since smoothing of noisy surface patch estimates in this algorithm involves a variational relaxation process, the neighbourhood information processing in this stage can be formulated in multiple measurement fusion terms using the theory developed in this thesis. Optimal estimates of surface patch parameters in each neighbourhood are obtained by combining information from all the neighbours. By redefining the concept of noise in the image, the same optimal estimator is then adjusted to provide discontinuity localization while eliminating errors due to noise.

The purpose of most surface reconstruction algorithms is to provide a piecewise continuous surface description from noisy image data. The problem of removing noise while preserving edges in the image data has been studied by many researchers [8] [9] [38] [39] [42] [49] [54] [69] [70]. Most of the algorithms developed by these researchers are adaptive in the sense that noise smoothing is applied only in continuous data regions and inhibited in the presence of an edge. In the field of surface reconstruction, many researchers base their approach on the regularization framework [50]. Various methods have been proposed to incorporate discontinuity information into this framework [2] [4] [27] [41] [59] [66]. Nearly all algorithms reviewed require some user set parameters. Some of the more recent adaptive algorithms [41] [49] [54] still require a parameter which controls the level of smoothing to be applied to the image. In the

attempt to make the discontinuity localization algorithm in this thesis completely adaptive, methods have been found to set all the parameters automatically.

Optimal estimation theory has found many applications in the field of image, and more generally, signal processing [23] [71]. In this thesis, a novel use of this theory has been found in providing optimal surface patch estimates while preserving discontinuities in the curvature consistency surface reconstruction algorithm.

1.1 Motivation

Visual reconstruction is an important area of research in the field of computer vision. The goal is to reduce visual data to stable descriptions [4]. It is viewed as an ill-posed inverse mathematical problem [50]. In the absence of constraints, the existence and stability of a unique solution cannot be guaranteed.

Surface reconstruction techniques deal with a particular aspect of visual reconstruction, which is the inference of surfaces from sets of noisy image data. Constraints are needed to narrow the choices of surface functions representing the image samples. Surface reconstruction is seen as being an early vision process, after which the higher level processes act to aggregate the stable surface descriptions to produce more domain dependent knowledge. In early vision processes, the constraints need necessarily be generic and should reflect general assumptions about the physical world being imaged and the imaging modality itself.

A local surface smoothness constraint is one such generic assumption which helps in restricting the domain of possible solutions for a function which interpolates a set of image samples. However, a smoothness constraint by itself may be incomplete, if not inaccurate, in some regions of the sampled data. In particular, inaccurate representation of a sampled surface is obtained in regions containing discontinuities if the smoothness constraint is applied blindly. A smooth interpolating function does not follow sampled data in discontinuous regions. To localize discontinuities in local neighbourhoods, it is imperative to either

1. include a model of discontinuities as part of the constraints, or

2. apply the smoothness constraint selectively, based on a belief in the continuity of the sampled data.

Under-constraining is precisely the problem with the surface reconstruction algorithm referred to as the “curvature consistency” algorithm, which is the focus of this thesis. The curvature consistency framework was first introduced by Sander and Zucker for surface reconstruction in 3D voxel based images [56], and was consequently modified for range images in [14, 15]. This algorithm operates on local measurements of surface orientation and curvature to produce a stable surface description from noisy discrete surface samples. Here the idea is to iteratively refine an initial set of surface descriptors such that, at convergence, each provides a consistent representation of its local surface region with respect to its neighbours. The problem is formulated as an energy minimization where the task is to minimize a functional form that limits the variation of curvature with respect to an implied model of the local surface [56, 37]. The resulting class of algorithms is robust, view-independent, and converges quite rapidly to stable descriptions [14, 15, 37, 56]. However, as presently formulated, the curvature consistency framework does not have an explicit model of discontinuity and applies the continuity constraint even across discontinuous sample data¹. Surface discontinuities are all pervasive in real life data: step edges occur at the junction of two adjacent objects, between object and background or between occluding object; roof edges occur between adjacent faces of a single object [4]. While recovering a stable surface description, it is important to preserve the discontinuities as a precursor to object segmentation and identification.

The effect of the original curvature consistency algorithm in reconstructing a discontinuous surface is demonstrated in Figure 1.1. Figure 1.1(a) shows noisy surface patch estimates from a synthetically generated step edge with Gaussian noise added. Figures 1.1(b) and 1.1(c) show the effects of the original curvature consistency algorithm on the step edge, after 30 and 100 iterations respectively. The discontinuity distortion can be seen to get progressively worse with the number of iterations. Fig-

¹We are mainly concerned with C^0 and C^1 discontinuities.




Figure 1.1: (a) Gaussian noise added to a synthetic step edge. Distortion in the step after (b) 30 iterations of curvature consistency algorithm, (c) at convergence (100 iterations). (d) Reconstruction using the edge localization method at convergence.

Figure 1.1(d) shows the same step discontinuity recovered after 30 iterations of the discontinuity localization method developed in this thesis. The discontinuity structure can be seen to be preserved while providing the same smoothing effect in the continuous regions of the data. Figure 1.2 shows the distortion for a roof discontinuity using the original algorithm in Figure 1.2(b) and (c), and the preservation of the edge in Figure 1.2(d) using the modified algorithm.

To tackle the problem of edge distortion, a modification is required to the functional minimization procedure in the curvature consistency algorithm to adapt to local discontinuities. It is desirable to make the solution adaptive over the image so that no user set thresholds are required and images of different noise properties can be handled. It should also be able to accommodate discontinuous surface features at




Figure 1.2: (a) Gaussian noise added to a synthetic roof edge. Distortion in the step after (b) 20 iterations of curvature consistency algorithm, (c) at convergence (100 iterations). (d) Reconstruction using edge localization method after 20 iterations.

different scales. Properties like view invariance [5] [61] which are exhibited by the existing curvature consistency algorithm should be preserved.

1.2 Overview

The discontinuity localization problem can be solved either by including a model of discontinuity in the constraints or by applying the constraints selectively². Several researchers have adopted the former approach in their surface reconstruction methods, where they have incorporated models of discontinuity as part of the constraints in the functional minimization leading to a unique surface function [2] [41]. One of the terms in the minimization function is dependent on the n^{th} derivative of sample data, where n determines the continuity of the region.

In the second approach, the smoothness constraint is applied selectively. It is enforced only in those image regions where there is reason to believe that the image data comes from a continuous region of the surface. Some researchers [24] [66] control the application of the continuity constraint by assuming that discontinuity information is available beforehand and smoothing the data only in the continuous regions of the data. However, discontinuity information is not always available. This approach then has the disadvantage of requiring a priori knowledge of the image properties and not being adaptive over variable noise in the image samples. Adaptive techniques can be used to modulate the application of the continuity constraints according to the degree of discontinuity in the sample data rather than make a binary decision on the presence or absence of an edge. The degree of dis(continuity) is learnt on the fly, while the smoothing is applied. Adaptive techniques have the advantage in that image properties are estimated and consecutively refined automatically as part of the process.

Traditionally, first derivatives have been used to define the “edginess” in a region during surface reconstruction [4] [24] [41] [49] [54] [66]. The edginess control is used to determine the balance between two separate processes acting on the noisy data.

²Various implementations of the curvature consistency algorithm deal with discontinuities in ad hoc fashion, typically by not updating across steep curvature gradients.

One process forces the interpolating function to smooth over the noisy sections of the data, while the other accommodates discontinuities and closeness of fit to the data. In some of the works [4] [24] [66], an energy threshold was used to make a binary decision on the presence or absence of a discontinuity in the energy minimization process. Recent modifications to these algorithms [22] [41] have made them more adaptive to the image and do not require such a user set threshold.

In this thesis the discontinuity localization problem in the curvature consistency algorithm is solved using an adaptive method, which controls the application of the constant curvature constraint according to a current belief in the continuity of the data. Using a recursive estimation technique, the reconstruction process “learns” the continuous or discontinuous nature of the local surface and applies the smoothing constraint accordingly.

In the curvature consistency framework, the continuity in a neighbourhood is determined by taking the difference between what a neighbour predicted as the interpolant surface patch parameters, and what was actually obtained after integrating information from all the other neighbours. Conceptually, this method of determining (dis)continuity is significantly different from simply taking a first derivative.

Over several iterations, within a local neighbourhood, the current belief of continuity or discontinuity is represented by an “error” variance associated with each pixel in that neighbourhood. The error variance of a given pixel tracks the history of its accuracy in estimating the interpolant patch parameters over the iterations. In the presence of an edge, some neighbours would consistently demonstrate a worse performance in the patch estimation process than others. A dichotomy in error variances within the same neighbourhood would indicate the presence of an edge. However, no attempt is made to make a binary decision about its presence or absence, as this would involve setting a user set threshold. Rather, the error variance is used to weight the pixel’s influence in determining the interpolant patch parameters in the subsequent iterations.

Optimal estimation theory offers a method of incorporating measurements from multiple sources with different noise properties. In a major contribution of this the-

sis, it is shown that inter-pixel interaction during relaxation processes can in fact be construed as a multiple measurement fusion problem, with each pixel acting as a source of measurement. The smoothing operation in the curvature consistency algorithm is performed using variational relaxation, which is based on the paradigm of relaxation labeling [34], except, it performs iterative constraint satisfaction in the continuous domain. Relaxation algorithms propagate information across the image iteratively. At each iteration a new “measurement” is available at a pixel-source. An optimal estimation technique can then be used to combine the measurements from pixel-sources in a neighbourhood to predict the interpolant patch parameters. The optimal estimator developed in this thesis uses a function of the prediction error variance of each pixel-source to weight its contribution to the interpolant surface patch parameter update.

By accommodating discontinuities in the “noise” model attributed to each pixel source, we arrive at a method by which the prediction error variance at each pixel is now capable of storing information about the continuity of data in the neighbourhood. By using an appropriate function of the error variance, it now becomes possible to encourage intra-region smoothing rather than inter-region smoothing across discontinuities.

Though the direct application of this discontinuity localization method is in surface reconstruction, it could conceivably be applied to any general smoothing problem where edges need to be preserved. The weighting function in the resulting edge preserving smoother, is similar to the one proposed by Saint-Marc et al. in [54]. Whereas the origins of their work lie in the anisotropic diffusion work done by Perona and Malik [49], the origin of the adaptive smoother presented in this thesis lies in optimal estimation theory. To the best of the author’s knowledge, the idea of using optimal estimation *between neighbouring image pixels* for discontinuity preserving smoothing has not been addressed in existing literature. Saint-Marc et al. and Perona and Malik, leave some parameters as user set arguments, such as the one which controls the amount of smoothing to be applied. In my approach all such parameters are calculated automatically and adaptively over an image.

Note that this work addresses the problem of localizing and preserving the local structure of discontinuities adaptively, without explicitly identifying them. The decision on the presence or absence of a discontinuity is not needed a priori and could be made *after* the surface is reconstructed.

The assumption in the development of the original curvature consistency algorithm was of local C^2 continuity of the underlying surface. In this thesis, the edge preserving reconstruction algorithm will also be applied to data from C^0 and C^1 continuous surfaces. On these surfaces the continuity constraint depends only on the position and normals of the adjacent patches. The discontinuities localized on the C^0 surface is only in depth, the ones on the C^1 surface is in depth and normals, and the ones on the C^2 surfaces are in depth, normals as well as curvature components. Another assumption made is that the noise in the image is zero mean Gaussian.

In this thesis, the curvature consistency algorithm without discontinuity localization is referred to as the “original algorithm”, and the extended algorithm with discontinuity localization, as the “modified algorithm”. In the context of the curvature consistency algorithm, the discontinuities which are localized are in depth, orientation and curvature.

1.3 Contribution

In this thesis, an original and unique discontinuity localization method is described in the framework of the curvature consistency surface reconstruction algorithm. More specifically, the contributions are:

1. It is shown that inter-pixel interactions which take place in relaxation processes can be construed in estimation theoretic terms. At each iteration in the relaxation process, the problem of finding the new state of a neighbourhood using the information from the constituent pixels is in fact a problem of multiple-measurement fusion and can be solved using optimal estimation theory.
2. Using optimal estimation theory as a basis, a new method is proposed in the curvature consistency framework to estimate interpolant surface patch param-

eters from neighbouring pixels. The patch parameter estimates are optimal in the minimum mean square error (MMSE) sense.

3. While keeping within the optimal estimation framework, a change in the definition of noise provides a method to localize discontinuities while providing a smooth surface fit in continuous areas using the curvature consistency smoothing algorithm.
4. The edge preserving smoother developed in this thesis is shown to be completely adaptive to discontinuities in depth, orientation and curvatures and over varying noise levels within the same image.
5. The edge preserving algorithm is shown to be robust over scale space. Processing can be done using different neighbourhood sizes, achieving the same edge localization effect.
6. The algorithm is made completely adaptive, with no user set parameters to adjust the amount of smoothing required.
7. An analysis of the proposed algorithm is presented, with an in-depth look at the role of each parameter.
8. Experiments are presented on synthetic and real range images to validate the theories set forth in this thesis.

1.4 Organization of Thesis

Before tackling the theory behind the discontinuity localization method, a survey of existing methods is presented in Chapter 2. Previous work is reviewed in the field of adaptive noise removal algorithms, surface reconstruction in the regularization framework and associated discontinuity localization methods, and finally the use of optimal estimation theory in the field of image and signal processing.

In Chapter 3, an overview of the curvature consistency surface reconstruction algorithm is presented. Since it is this algorithm which we will augment to include

the discontinuity localization feature, it is worthwhile to review the theory behind it. The algorithm is presented in the context of reconstruction from range images. The surface representation, methods for obtaining a first estimate of the surface patches, and the method to iteratively refine these estimates are covered.

Chapter 4 provides an introduction to optimal estimation theory. The usage of a simple version of the Kalman filter is illustrated by applying it to a simple problem of multiple-measurement fusion. The example problem is carefully chosen so that it parallels the measurement fusion process between pixels in a relaxation process in an image. An optimal estimator is derived and modified to accommodate time varying noise processes.

Chapter 5 forms the core of this thesis. In this chapter it is shown that during variational relaxation in the curvature consistency algorithm, the process by which surface patch parameters are updated is analogous to the multiple measurement fusion problem described in the previous chapter. It is shown that the optimal estimator developed for the previous chapter could be applied directly in the curvature consistency framework. The concept of noise in the pixel sources in the relaxation framework is discussed in detail and expanded to include discontinuities. The edge preserving smoothing technique is then derived using optimal estimation theory and the redefinition of noise as a basis.

Results and analysis are presented in Chapter 6. The discontinuity localization algorithm is tested on a number of synthetic and real range images and its behaviour characterized. The role of the different parameters are discussed. Using several experiments, the robustness and adaptive behaviour of the algorithm is proven.

Finally, conclusions are presented in Chapter 7, in which areas of further research are described and the main ideas behind this thesis are summarized.

In this chapter, the historical context for the research presented in this thesis is established. The literature survey spans over three related research areas: adaptive noise smoothing techniques, visual reconstruction algorithms in the regularization framework and associated discontinuity localization methods, and finally the use of optimal estimation theory in image processing.

First we will review the early literature related to noise removal methods. The emphasis is on algorithms which are adaptive in nature, i.e., which adapt to the presence of the edges in the images, and do not require a priori knowledge of their positions. Another feature to look for in these algorithms is adaptation to different noise levels in the images. The application of these techniques varies from 1-D signals, 2-D gray level images, to $2\frac{1}{2}$ -D range images.

Next, more recent literature related to the application of adaptive discontinuity localization techniques to visual reconstruction methods will be reviewed. More specifically, visual reconstruction algorithms in a regularization framework will be surveyed. From a mathematical and statistical point of view, the problem of interpolating data from *surfaces* has received much attention. The goal of these algorithms is to recover smooth and piecewise continuous surface interpolant functions from noisy image sample data by minimizing some energy cost functional. Some of the more important reconstruction algorithms and the corresponding discontinuity localization methods have been reviewed.

One of the major contributions of this thesis is to show a link between optimal estimation theory and relaxation processes, such as the one which takes place in the curvature consistency algorithm. Though the application of optimal estimation theory to inter-pixel interactions is new, the use of optimal estimation theory to the general field of image processing is not. Some of the image processing algorithms which use optimal estimation frameworks for feature analysis and signal estimation

are reviewed in this chapter.

2.1 Adaptive Noise Smoothing Techniques

Edges in images contribute significantly to the process of human or artificial perception [30] [40]. Due to their importance in building a scene description, there has been a tremendous amount of research performed in extracting, analyzing and preserving edge features in images. Most of this research has centered around the detection of reliable edge information, and its use in building credible object descriptions. However, the research literature relevant to this thesis topic deals with the *preservation* rather than the explicit detection of edges.

Distortion of edges takes place most commonly in noise smoothing or image restoration algorithms. See [26] and [53] for an introduction to the subject of smoothing noisy images. A low pass filtering process to remove the higher frequency noise results in the removal of the high frequency components of the edge signal as well [26]. A similar fate is expected when linear algorithms such as a mean filter (running average) are used. Noise removal filters which adapt to varying noise and image statistics over an image are non-linear. Mastin in [42] and Chin et al. in [8] provide a review and evaluation of several such non-linear adaptive filtering algorithms.

One of the most popular edge preserving noise smoothing algorithms has been the median filter first proposed by Tukey [69]. It was initially developed for application in time series analysis. It was subsequently adapted to image noise smoothing in [21] and [51]. This filter operates on the simple concept of picking the median pixel value within a sliding window to replace the window's central pixel. Huang et al. [33] provided a much faster version of the algorithm.

Lev et al., in [39], provide one of the early adaptive methods to enhance noisy images containing discontinuities. Theirs is an iterative algorithm which uses the observation that while averaging in a neighbourhood which contains an edge, only those neighbours which are on the same side of the edge should be used in the averaging process. To establish which neighbours are on the same side of the edge, a kind of

template matching is performed to determine which of 12 configurations of edges a 3×3 neighbourhood matches. For each possible edge configuration A_i , ($1 \leq i \leq 12$), a corresponding 3×3 weight matrix D_i is assigned. Once template matching of a 3×3 image neighbourhood A has taken place over each of the 12 possible neighbourhood configurations, the final weighting for each neighbour in the window is determined by

$$D = \sum_{i=1}^{12} w_i D_i, \quad (2.1)$$

where, w_i is a scalar measure of how closely the image neighbourhood A matches the configuration A_i . A more elaborate second approach presented in the paper assigns a gray level probability density to each pixel P in the image, rather than just a gray level. At each iteration this density is adjusted by comparing it with the densities of P 's neighbours, where the contribution of a given neighbour Q depends on the confidence that Q belongs to the same region as P , as determined by the template matching process described earlier. One limitation of this algorithm is its dependence upon the existence of local cues to the underlying piecewise-constant structure. Another limitation is that the algorithm is developed only for a 3×3 window size. Ideally for noisier images, more image pixels need to be used in the averaging process to eliminate the noise, however it is not trivial and computationally expensive to extend this algorithm to larger scales.

Davis et al. in [9] also work on the same assumption that due to the presence of an edge, there is a grouping of pixels according to similarity of gray levels. However they propose a much simpler approach than [39], in which the gray level of the sliding window's central pixel is replaced by the average of its k neighbours whose gray levels are *closest* to that of the central pixel. The biggest drawback of this algorithm relative to [39] is that the parameter k now needs to be set by the user to determine how many pixels in a neighbourhood should contribute to the average. In [39] on the other hand, the neighbourhood contributions are determined automatically. This algorithm lends itself better to application with a larger window size, even though there are no results presented to this effect in the paper.

Wang et al. in [70] recommend a gradient-inverse weighted noise smoothing algorithm. Whereas the observation in the two previous techniques ([9] and [39]) was that within each subregion separated by a discontinuity, the gray levels are similar, Wang et al. work on the assumption that the *variations* of gray levels inside each continuous region are smaller than those between regions. For each 3×3 neighbourhood, they use a matrix of weighting coefficients. The coefficients are the normalized gradient inverses at each pixel. This approach is best applied to images with sharp discontinuities and with discontinuities with magnitudes much greater than the noise magnitude.

In [38] Lee proposed an adaptive method to smooth images with additive Gaussian noise. The basic assumption is that the sample mean and variance of a pixel at the center of a window are equal to the local mean and variance of all pixels within the neighbourhood. The mean and variance of the recovered image are approximated from the local means and variances from the noisy image. Each pixel in the recovered image is given by

$$\hat{x}(i, j) = \bar{x}(i, j) + \frac{Q(i, j)}{Q(i, j) - \sigma^2} [z(i, j) - \bar{z}(i, j)], \quad (2.2)$$

where, $z(i, j)$ is a noisy image pixel, $\bar{z}(i, j)$ represents the mean gray level in the neighbourhood around (i, j) , and

$$Q(i, j) = E[(z(i, j) - \bar{z}(i, j))^2] - \sigma^2. \quad (2.3)$$

The σ^2 parameter is an estimate of the additive noise variance. It is estimated by averaging the lowest 10% of the window variances computed in a flat intensity region. The user is in fact responsible for directing the algorithm to the “flat” area in the image.

Tomita et al. describe a smoothing method in [68], in which smoothing is performed only in a selected neighbourhoods around each pixel. The most homogeneous of five rectangular neighbourhoods around a given pixel is chosen and the gray level of the pixel replaced by the average gray level in the selected neighbourhood. The

homogeneity is determined by applying a gradient operator within a given neighbourhood. This algorithm tends to over-smooth complex shaped boundaries, and does not preserve edges of smaller scales. In an enhancement to the original method, Nagao et al. [45] propose using nine bar shaped regions instead of the original rectangular regions.

In one of the more important works in this field, Perona and Malik in [49] have extended the concept of scale-space recovery of edge features in images using anisotropic diffusion. They have described a very effective method of localizing edges in the image by using a spatially varying diffusion coefficient. The coefficient encourages intra-region smoothing in preference to inter-region smoothing. The anisotropic diffusion equation for image I at scale (time) t is

$$I_t = \text{Div}(c(x, y, t)\nabla I) = c(x, y, t)\Delta(I) + \nabla c \cdot \nabla I, \quad (2.4)$$

where $c(x, y, t)$ is the diffusion coefficient, Δ is the Laplacian operator ∇ is the grad operator, and Div is the divergence operator. For discontinuity localization they propose setting this coefficient to a function $g(\cdot)$ of the magnitude of the local gradient in the image

$$c(x, y, t) = g(\|\Delta I(x, y, t)\|). \quad (2.5)$$

The function $g(\cdot)$ is chosen to be a nonnegative monotonically decreasing function with $g(0) = 1$. Amongst the functions they examined were

$$g(\Delta I) = e^{-(\|\Delta I\|/K)^2}, \quad (2.6)$$

and

$$g(\Delta I) = \frac{1}{1 + (\frac{\|\Delta I\|}{K})^2}. \quad (2.7)$$

The parameter K determines how much smoothing should be applied in the image. Since this is a user set parameter, the algorithm is not completely automatic. Even though the algorithm adapts to edges, due to a fixed constant, it does not adapt to varying noise statistics over the image. There is a danger of over-smoothing discon-

tinuous regions due to an inappropriate choice of K .

In a recent simplification of Perona and Malik's work, Saint Marc et al. [54] present a more general algorithm for adaptive smoothing. They propose a filtering mechanism where neighbourhood weighting is applied using the function (2.6). They show operation of their filter on 1-D signals, 2-D intensity images as well as range images. A special adaptation of their algorithm has to be made to deal with range images. In their original formulation, the weight assigned to a given neighbour is a function of the signal difference between them. To localize discontinuities in orientation (roof edges), instead of using the difference in signal value between two pixels, the difference in the gradients at those two pixels is used. Similar to Perona and Malik's work, they also have no automatic method of setting the K parameter to determine the amount of smoothing required over the image. Though they have made good use of the K parameter by designating it as the scale space control, they present no method of determining its value according to the chosen scale.

2.2 Visual Reconstruction in the Regularization Framework

The process of reducing noisy visual data to stable descriptions is often referred to as visual reconstruction. The process of surface recovery from image data falls under the same definition. Visual reconstruction from data is an ill-posed problem, which by definition implies that the existence, uniqueness and stability of solutions cannot be guaranteed. Poggio et al. in [50] in their review paper, presented a regularization framework [67] for solving it. The idea is to make the problem well-posed by introducing a priori knowledge in the form of constraints. Several researchers have used the regularization approach to solve the visual reconstruction problem, with and without discontinuity localization. Some of that work will be reviewed here. See [5] and [61] for a detailed analysis of surface reconstruction theory and implementation in the regularization framework.

Poggio et al. set the stage for the use of regularization techniques in visual and

surface reconstruction in [50]. Surface reconstruction is akin to finding a solution to the general problem of the form

$$Az = y \quad (2.8)$$

The aim is to find z from the data y . This requires choosing a norm $||\cdot||$ and a stabilizing functional $||Pz||$. A is a linear operator, the norms are usually quadratic and P is linear. Poggio et al. present two methods of regularization:

1. Among all the z which satisfy

$$||Az - y|| \leq \epsilon, \quad (2.9)$$

where ϵ is a minimum error, find the z which minimizes

$$||Pz||^2. \quad (2.10)$$

In surface reconstruction terms, the idea then would be to first find a function which provides the closest fit to the given sample data, and then apply the smoothness constraint criterion $||Pz||^2$. The criterion could be that of C^1 or C^2 continuity.

2. Find the z that minimizes

$$||Az - y||^2 + \lambda ||Pz||^2, \quad (2.11)$$

where λ is a regularization parameter. It controls the trade-off between “closeness to the data” and “following the continuity constraint”.

Several researchers have formulated the surface reconstruction problem in this manner [2] [4] [27] [28] [41] [59] [66]. Though the particular reconstruction problem formalism may be based on physical, stochastic, or information theoretical models, the solutions are based on functional energy minimization. The surface f can be

found by minimizing a composite cost energy functional

$$E(f) = \sum_{i=1}^k |z_i - f(x_i, y_i)|^2 + \lambda(f_{xx}^2 + 2f_{xy}^2 + f_{yy}^2) dx dy \quad (2.12)$$

The first term in the energy functional measures the cost due to error between a function f and the data z . The second term is the regularizer term which measures the irregularity of the data. As presented in (2.12) and used by Grimson in [28], this term gives a measure of the C^2 continuity. The minimization yields a function which is termed as the *thin plate spline* approximation. Schumaker in [57] was one of the first to use this function in interpolation from scattered data.

By applying the C^2 smoothness criterion imposed by this regularizer globally over the whole image, discontinuous regions in the surface may not be recovered accurately. A better approach is to use a variable order regularizer [2]. By considering higher orders of derivatives the regularizer term in the cost energy functional becomes

$$E_R(f) = \sum_{n=1}^N \lambda_n \int \int_{\Omega} [f^{(n)}(x, y)]^2 dx dy, \quad (2.13)$$

where N is the highest order derivative, $f^{(n)}$ is the n^{th} order derivative of f , λ_n is the weight assigned to each individual n^{th} order component in the summation, and Ω is the region over which the interpolation takes place.

Geman and Geman in [24] introduce the idea of a “line process” into the regularization framework. The line process embeds prior knowledge of the geometry of discontinuities in the regularizer term. For the n^{th} order derivative term in (2.13) the cost energy term becomes

$$E_R^n(f) = \int \int_{\Omega} [\lambda_n \{ (1 - u_{x,y,n})(f_x^{(n)}(x, y))^2 + (1 - v_{x,y,n})(f_y^{(n)}(x, y))^2 \} + \alpha_{x,y,n}^u u_{x,y,n} + \alpha_{x,y,n}^v v_{x,y,n}] dx dy, \quad (2.14)$$

where f_x and f_y denote the derivatives in x and y directions respectively, $u_{x,y,n}$ and $v_{x,y,n}$ are the line processes in the two directions, $\alpha_{x,y,n}^u$ and $\alpha_{x,y,n}^v$ are the penalties imposed for the inclusion of the discontinuity and controls the resistance to noise

in the image. For quadratic regularization terms in standard regularization, the search space for the function which minimizes the energy functional has only one local minima [50]. However for non-quadratic terms like the one proposed by Geman and Geman, it is easy to get stuck in a local minima and never converge to the global one.

Geman and Geman use a stochastic method (simulated annealing) to optimize the non-convex energies, in which random disturbances are applied to enable the system to jump out of local minima in the energy function. The magnitude of the random disturbances are controlled by a temperature parameter. Their formulation is based on Bayesian estimation and Markov random field (MRF) models.

Blake and Zisserman [4] have proposed a deterministic procedure for optimizing the non-convex energies associated with the piecewise continuous reconstruction problems. The weak-membrane model energy minimization has been solved using their Graduated Non-convexity (GNC) algorithm. For a comparison of deterministic (relaxation) vs. a stochastic method for surface reconstruction with discontinuity localization see [3].

Drawbacks of MRF models are the computational complexity and the difficulty in estimating the parameters of the model. Geiger and Girosi in [22] have derived a deterministic approximation to the MRF models, one of which gives in a natural way, the GNC algorithm of Blake and Zisserman.

Deterministic relaxation algorithms can be inefficient. Since they are posed in variational terms or as partial differential equations, they are amenable to numerical solutions by algorithms which are local, iterative and often parallel. In [65] Terzopoulos points out that such relaxation algorithms have an inherent inefficiency at propagating constraints between widely separated processing regions in an image. Due to this reason, these algorithms converge extremely slowly. He proposes a multi-grid relaxation method to deal with this inefficiency. Since constraints travel faster across the image at coarser scales, they can be used to constrain the processing at a smaller scale [63] [64].

To incorporate discontinuities in the regularization framework, Terzopoulos [66]

introduces a controlled-continuity stabilizer in which the regularizer term contains a weighted sum of the energies associated with a membrane and a thin plate approximation over the surface. The regularizer term E^R is described as

$$E^R(f) = \frac{1}{2} \int \int_{\Omega} \rho(x, y) \{ \tau(x, y) (f_{xx}^2 + 2f_{xy}^2 + f_{yy}^2) + [1 - \tau(x, y)] (f_x^2 + f_y^2) \} dx dy. \quad (2.15)$$

The parameters $\rho(x, y)$ ($0 \leq \rho \leq 1$) and $\tau(x, y)$ ($0 \leq \tau \leq 1$) provide control over the behaviour of the surface as a thin plate or a membrane. When $\rho(x, y) = \tau(x, y) = 1$, a thin -plate approximation function is obtained and the solution is as computed by Grimson in [28]. When $\rho(x, y) = 1$ and $\tau(x, y) = 0$, a membrane approximation is obtained. When $\rho(x, y) = 0$, the surface is discontinuous and surface “tearing” occurs.

The methods described above allow the incorporation of depth and orientation discontinuities into the regularization model, however, the localization of binary discontinuities must be either pre-specified or determined at the end of the reconstruction process. If we consider the regularizer term in one dimension (weak string formulation), the regularizers which incorporate discontinuity localization can be represented by

$$E_R(f) = \sum_{n=1}^N \lambda_n \pi_n(x) \int_a^b [f^{(n)}(x)]^2 dx, \quad (2.16)$$

where $\pi_n(x)$ controls the n^{th} order discontinuity localization. In the presence of a discontinuity ($\pi_n(x) = 0$), smoothing is inhibited, and $\pi_n(x) = 1$ implies that the interpolation takes place over a continuous region and smoothing of the noise is permitted. Stan Li in [41] points out that in most of the reconstruction methods, the value of this controller is either restricted to, or forced towards 0 or 1. He proposes an adaptive regularizer which does not use a weighting term like (2.16). This regularizer solves the conflict between over-smoothing across discontinuities, and finding discontinuities. His proposed regularizer is of the form

$$E_R(f) = \sum_{n=1}^N \lambda_n \int_a^b -e^{-[f^{(n)}(x)]^2 / \tau_n} dx, \quad (2.17)$$

where $[f^{(n)}(x)]^2$ gives a measure of the n^{th} order irregularity of the data, and τ_n is a user set parameter which controls the level of smoothing. His implementation uses only the first order derivatives in the regularizer. Though this method manages to adapt the surface reconstruction algorithm to different levels of discontinuities in depth and orientation, there is no automatic method of setting the smoothing control τ .

Viewpoint invariance in the reconstruction process has been recognized as a very desirable feature. The recovery of the surface description should be independent of the choice of viewer coordinate system [4]. In [5] and [61] different surface reconstruction algorithms have been analyzed from the view invariance aspect. See [60] for one such example of viewpoint invariant recovery of surface interpolants from sparse data.

In [20] [13] [16] [12] Ferrie et al. recover the local structure of a surface described in terms of differential geometry. The reconstruction algorithm is an adaptation to range images of the work done by Sander and Zucker in [56] and [55] and is popularly referred to as the curvature consistency algorithm. The curvature consistency algorithm has been shown to be very effective in recovering the local structure of smooth surfaces obtained from C.T. data [55, 56], range images [15, 17], and shape-from-shading [14]. As originally proposed by Sander, the purpose of the algorithm was to infer trace points through which the surface of the object passes in 3D images. Coupled to that was the problem of estimating parameters for the local shape descriptors: the extended darbox frames. The extended darbox frames describe quadric patches, which provide a covering for the object surface. They capture the differential structure of the local surface region based on principal curvatures and tangent fields. The constraints are in terms of curvature variation between local neighbouring patches. The recovery of the surface parameters is view-invariant.

The surface recovery problem in the curvature consistency framework is not molded in the same regularization framework as the algorithms reviewed so far. However it is still based on variational principles. At the beginning of this section two methods were recommended for solving ill-posed problems (2.8). Most of the reconstruction algorithms presented so far have followed the second method, where the “closeness” term and the “regularizer” term are both minimized at the same time. How-

ever the method proposed by Ferrie et al. is similar to the first method of regularization, where the closest fit to the data is obtained first, and then as a consequent step, the continuity constraint is applied. It is this algorithm which will be augmented in this thesis to include discontinuity localization. The underlying curvature-consistency data smoothing algorithm will be reviewed in some detail in the following chapter.

2.3 Use of Optimal Estimation Theory in Visual Reconstruction

With close similarities between the fields of digital signal processing and estimation theory, the use of the optimal estimation theory has seen tremendous use in signal estimation, image restoration, de-blurring, multi-sensor fusion, and other fields of image processing. In [23] Gelb and in [71] Willsky have provided a comprehensive introduction to optimal estimation theory and its use in digital signal processing.

Some early research in the recursive estimation of image signals from data corrupted by additive white Gaussian noise is presented in [35]. In the field of visual reconstruction, recursive estimation has traditionally been used as a means of incorporating image data from different sources or from time-varying imagery [7]. Nearly all the papers surveyed in this field use some variation of the Kalman filter [73].

Matthies et al. [44] have used Kalman filtering for obtaining on-line estimates of depth from motion sequences. In the first stage they use correlation to produce a disparity map at each pixel and an estimate of the associated variance. In a second stage, the new disparity map is integrated with the one predicted from the previous time step. The third stage uses a regularization approach to smooth the measurement noise and to interpolate under-constrained areas. The smoothing function used is the generalized piecewise continuous spline under tension [65]. As a last stage, the disparity field in the next time step is predicted according to a camera translation. In [43], an extension to arbitrary camera motion is described. It is interesting to note that in this later work they attempt to localize discontinuities while performing the regularization step by thresholding the angle between the camera view vector and the

local surface normal to detect surface foreshortening.

As opposed to estimating depth maps, Szeliski in [62] presents a Kalman filter based method to estimate camera motion parameters from time-varying range maps. No correspondence is established between images rather, the motion estimate is found by finding the geometric transformation which makes it most likely (in a Bayesian sense) that the points come from the same surface. His method also “registers” pixels in two consecutive images, so that they may be integrated into an updated piecewise continuous surface estimate. A dense depth map is incrementally built up by interpolating and integrating sparse range data. The points are then matched to some known surface model. He uses a thin plate surface model and Terzopoulos’s method [65] to localize depth and orientation discontinuities.

Singh in [58] uses a Kalman filtering technique to compute image-flow from time-varying imagery. One very interesting aspect about his work is that he uses a neighbourhood weighting method to localize discontinuous flow fields. Neighbours which do not provide a good estimate of the central pixel’s motion vector are weighted down, while the ones which provide a good estimate are given more emphasis in the subsequent time steps of image frame integration.

A recent paper by Heel [29] has shown an equivalence between recursive estimation theory and surface reconstruction in the regularization framework from time-varying imagery. Using a thin plate model, and representing the surface by a depth map he shows that the solution of the resulting optimization problem is identical to the update procedure of a Kalman filter. He also shows the equivalence between the regularization parameter and the variance of the depth estimate at a given pixel. Given two depth samples Z_{1ij} and Z_{2ij} , to compute a depth map Z_{ij} which is closest to both values the following optimization problem needs to be solved

$$\min_Z \sum_{i,j} \lambda_1 (Z_{ij} - Z_{1ij})^2 + \lambda_2 (Z_{ij} - Z_{2ij})^2, \quad (2.18)$$

the solution to which is

$$Z_{ij} = \frac{\lambda_1 Z_{1ij} + \lambda_2 Z_{2ij}}{\lambda_1 + \lambda_2}. \quad (2.19)$$

In showing the equivalence between the Kalman filtering update equation which integrates information from one image to another, the λ 's can be replaced by the inverse of the variance estimate for that particular pixel to give

$$Z_{ij} = \frac{Z_{1ij}/p_{1ij} + Z_{2ij}/p_{2ij}}{1/p_{1ij} + 1/p_{2ij}}. \quad (2.20)$$

To incorporate discontinuities in the same model he includes the line-process terms [28] in the functional to be minimized and arrives at another corresponding Kalman filter based update procedure.

Measurements of depth and/or orientation information over surfaces are available from several imaging modalities. Several Shape-from-X {X = shading, motion, texture, stereo} algorithms are capable of providing these measurements with varying degrees of accuracy [32]. The goal of a reconstruction algorithm is to provide piecewise smooth surface descriptions from these measurements, while preserving the important features of the surface.

In Sander's and Zucker's original work [55, 56], the surface reconstruction algorithm for 3D images was tied to inferring the trace points of the surface itself. In the later adaptation to range images, and shading analysis, [37, 14], the trace of the surface is given a priori, hence the problem reduces to finding the surface coverings from the measurements. The surface reconstruction process begins with an initial set of surface descriptions computed from local estimates of the orientation or depth, e.g. [1, 11, 18, 19, 48, 47]. The curvature consistency algorithm amounts to an iterative refinement of these estimates such that curvature varies according to an expected local model of a surface. In this thesis the context of the reconstruction algorithm will be depth maps acquired with a laser range-finding system.

The surface descriptors used in a reconstruction algorithm should be able to describe the salient features in the image, and should be independent of the viewer coordinate system. In Section 3.1 one such descriptor, based on differential geometry of local surfaces is presented. The method used to obtain initial estimates of this surface descriptor from noisy range data is presented in Section 3.2. In Section 3.3 these initial estimates are refined using curvature consistency: a transport model sets up the continuity criterion between neighbouring surface patches; the updating rules describe how neighbouring patches contribute to the refinement of the surface descriptor at a range data point.




Figure 3.1: Local surface representation - the augmented Darboux frame

3.1 The Local Surface Representation

The goal of our surface reconstruction algorithm is to provide a patchwork of overlapping, interpolation functions, each of which describes the surface locally. Following the convention of Sander and Zucker, we use an augmented Darboux frame to represent the local neighbourhood of a point P on a surface S [10, 55, 56]. This representation is described as follows. Let the local neighbourhood of P be represented by a quadric patch of the form

$$w = au^2 + cv^2, \quad (3.1)$$

with origin at P and the w axis aligned with the surface normal at P , N_P , as shown in Figure 3.1. The orientation of this local frame is such that the u and v coordinate axes align with two special directions on S at P . These are the directions for which the normal curvature at P (a directional property) takes on maximum and minimum values, κ_{M_P} and $\kappa_{\mathcal{M}_P}$, and are referred to as the principal directions M_P and \mathcal{M}_P respectively [10]. The scalar quantities κ_{M_P} and $\kappa_{\mathcal{M}_P}$ are similarly referred to as the principal curvatures at P . Following the convention of [55, 56], we refer to $\mathcal{D}(P) = (P, M_P, \mathcal{M}_P, N_P, \kappa_{M_P}, \kappa_{\mathcal{M}_P})$ collectively as the augmented Darboux frame at P .

Assuming that the system providing measurements of the surface is influenced by random or systematic noise, initial estimates of the augmented Darboux frames at each surface sample point will be noisy. The curvature consistency algorithm is then

used to iteratively refine the components of the extended Darboux frame in such a way that two neighbouring frames finally represent the same continuous surface patch at different positions. This will occur when the position, orientation and curvature vary smoothly across the fitted surface.

3.2 Initial Estimates

The components of $\mathcal{D}(P)$ can be estimated either directly from measurements of the surface gradient, or indirectly from the parameters of some appropriate local approximation of S (e.g. a quadric patch). Several methods exist to provide these estimates [1, 10, 11, 18, 25] which use operator-based methods, feature analysis or functional analysis. In our reconstruction method we use a least-squares fit of a parabolic quadric to each point P and its neighbours, to estimate each component of $\mathcal{D}(P)$.

The surface measurements are acquired as a grid of discrete points $z(i, j)$. To fit an osculating parabolic quadric of the form of the form

$$h(u, v) = au^2 + buv + cv^2, \quad (3.2)$$

to a point P_{ij} , first the $I \times J$ neighbourhood around P_{ij} needs to be transformed into a local coordinate system centered at this point. The coordinate system is aligned with the normal $N_{P_{ij}}$ to the surface at point P_{ij} and the tangent plane $T_{P_{ij}}$, shown in Figure 3.1. Notice that (3.2) has an additional parameter b as compared to (3.1). This is to account for the unknown rotation of the local frame with respect to $N_{P_{ij}}$.

The tangent plane $T_{P_{ij}}$ is obtained by a planar fit to the $I \times J$ window centered at P_{ij} . A linear transformation is then used to map each point in the $I \times J$ neighbourhood to the coordinate frame of P_{ij} .

Using least-squares methods, a parabolic quadric is fit to the points projected into P_{ij} 's coordinate system. Once the a , b , and c parameters are determined then $M_P, \mathcal{M}_P, \kappa_{M_P}$ and $\kappa_{\mathcal{M}_P}$ are determined by

$$\begin{aligned}
M_P &= (a - c + \sqrt{(a - c)^2 + b^2}, b) \quad a \geq c \\
&\quad (b, - (a - c + \sqrt{(a - c)^2 + b^2})) \quad a < c \\
\mathcal{M}_P &= (-b, a - c + \sqrt{(a - c)^2 + b^2}) \quad a \geq c \\
&\quad (a - c - \sqrt{(a - c)^2 + b^2}, b) \quad a < c \\
N_P &= (0, 0, 1) \\
\kappa_{M_P} &= a + c + \sqrt{(a - c)^2 + b^2} \\
\kappa_{\mathcal{M}_P} &= a + c - \sqrt{(a - c)^2 + b^2}.
\end{aligned} \tag{3.3}$$

However as can be expected, the frames $\mathcal{D}(P)$ at each point will be subject to the usual effects of noise and quantization error. Various techniques can be used to optimize estimation in the presence of noise, but they largely fall short when it comes to estimating directions [56]. We now consider how to improve on these errors through an iterative refinement of $\mathcal{D}(P)$ using the curvature consistency algorithm.

3.3 Iterative Refinement Using Curvature Consistency

The algorithm operates by iteratively minimizing the difference between the description of S at P given by $\mathcal{D}(P)$, and that predicted by its local neighbours Q_α according to a local model that describes how the frames at Q_α should appear at P when transported along the surface of that model (Figure 3.2). The particular considerations in formulating the algorithm are:

- The form of the transport model. This explicitly embeds the expected structure of the local surface by constraining how $\mathcal{D}(P)$ changes as it is moved across it. Ideally, the transport model is chosen to enforce the locally constant curvature assumption [46], e.g. by using a torroidal transport model [37]. However, for computational simplicity, a quadric is often used in approximation of this constraint.




Figure 3.2: The local transport model determines how a frame at a neighbouring point Q appears when it is extrapolated over to P .

- The form of the minimization functional. Let ξ_{Q_α} denote the set of frames corresponding to each neighbour of P and ξ_{P_α} the set of frames that results when each element of ξ_{Q_α} is transported across the surface to P as shown in Figure 3.2. This functional describes the least-squares difference between $\mathcal{D}(P)$ and ξ_{P_α} . A set of updating functions for each component of $\mathcal{D}(P)$ is determined by minimizing this functional.
- Determination of convergence. At each iteration, $\mathcal{D}(P)$ is updated and a residual error is computed by taking $\sum_\alpha \|\mathcal{D}(P) - \xi_{P_\alpha}\|^2$. The sum of these residuals over S gives an energy measure which stabilizes over time, thus providing a means of determining when to halt the updating process [55, 56, 15].

For further details on the above issues, the reader is referred to [14, 15, 56]. What is of importance are the functions updating $\mathcal{D}(P)$ from the neighbourhood predictions as they have a direct bearing on the behavior of the algorithm in the vicinity of surface discontinuities. They are summarized as follows for the P , N_P , κ_{M_P} and κ_{M_P} components of $\mathcal{D}(P)$:

$$P^{(k+1)} = \sum_{\alpha=1}^n \frac{P_\alpha^{(k)}}{n} \quad (3.4)$$

$$N_P^{(k+1)} = \frac{(\sum_{\alpha=1}^n N_{x_{P_\alpha}}^{(k)}, \sum_{\alpha=1}^n N_{y_{P_\alpha}}^{(k)}, \sum_{\alpha=1}^n N_{z_{P_\alpha}}^{(k)})}{\sqrt{(\sum_{\alpha} N_{x_{P_\alpha}}^{(k)})^2 + (\sum_{\alpha} N_{y_{P_\alpha}}^{(k)})^2 + (\sum_{\alpha} N_{z_{P_\alpha}}^{(k)})^2}}, \quad (3.5)$$

$$\kappa_{M_P}^{(k+1)} = \sum_{\alpha=1}^n \frac{\kappa_{M_P\alpha}^{(k)}}{n}, \quad (3.6)$$

$$\kappa_{\mathcal{M}_P}^{(k+1)} = \sum_{\alpha=1}^n \frac{\kappa_{\mathcal{M}_P\alpha}^{(k)}}{n}, \quad (3.7)$$

where the superscript k refers to the current iteration step.

The updating function for M_P is expressed in terms of an alternate parameterization which takes into account the fact that M_P is a *direction*, i.e., that there is a 180° ambiguity in orientation. If we express M_P in tangent plane coordinates as

$$M_P^{(k+1)} = \bar{b}_1 \cos \theta^{(k+1)} + \bar{b}_2 \sin \theta^{(k+1)}, \quad (0, 2\pi) \quad \text{such that} \quad \begin{aligned} 1) \quad & \bar{b}_1, \bar{b}_2 \in T_P \\ 2) \quad & \|\bar{b}_1\| = \|\bar{b}_2\| = 1 \\ 3) \quad & (\bar{b}_1 \cdot \bar{b}_2) = 0, \end{aligned} \quad (3.8)$$

then

$$\theta^{(k+1)} = \tan^{-1} \left[\frac{(A_{22} - A_{11}) + \sqrt{(A_{11} - A_{22})^2 + 4A_{12}^2}}{2A_{12}} \right], \quad (3.9)$$

where

$$A_{ij} = \sum_{\alpha=1}^n (M_{P\alpha}^{(k)} \cdot \bar{b}_i)(M_{P\alpha}^{(k)} \cdot \bar{b}_j). \quad (3.10)$$

Note that this also determines the solution for $\mathcal{M}_P^{(k+1)}$ since $\mathcal{M}_P = M_P \times N_P$.

3.4 Applying Curvature Consistency to Range Images

Range images are the easiest for surface reconstruction algorithms to be applied to, as opposed to intensity images. This is because the pixel position information is provided directly, and does not have to be derived from other image properties, e.g., by shading analysis. The errors introduced in estimation of surface patches are then mainly due to the systematic errors inherent to the range camera. No Shape-from-X [37] methods need to be applied to infer the coordinates of pixelized surface samples.

Since the local curvature variation is accounted for in the minimization process, the curvature consistency algorithm lends itself very well to reconstruct locally continuous

C^2 surfaces. Figure 3.3 demonstrates its application to samples from such a surface. Figure 3.3(a) shows the samples generated from an ideal underlying surface following a sinusoidal trace, with amplitude 5 (pixel height units). In Figure 3.3(b) Gaussian noise of amplitude 3 (pixel height units) has been added. The initial estimate of the Extended Darboux frames at each sample point is obtained in Figure 3.3(c). The initial estimate of normals at each image sample is obtained using a local fit of planes. The curvature is estimated by fitting local parabolic quadrics. A 5×5 neighbourhood size was used to obtain the initial estimate. Figures 3.3(d) and (e) show the recovered surface after 3 and 10 iterations respectively of the curvature consistency algorithm using a 5×5 neighbourhood. As can be seen an almost perfect fit is obtained. The additive Gaussian noise has been smoothed out, recovering the underlying surface. Further iterations will not affect the surface fit much. The error of fit tends closer to zero with each iteration. The iterations can be stopped after the energy residual passes a certain threshold. The plot of the energy residual is shown in figure 3.3(f) against the number of iterations.

Figure 3.4 shows the application of the algorithm on natural range images from a laser range camera. The imaged object is a toy car. As can be noticed, the noise level in the data is not as severe as the synthetic image in Figure 3.3. The most significant noise-effected regions are on either side of the windshield area, where the range camera has obtained bad estimates of depth. Figure 3.4(a) shows the range image samples in a grid form and Figure 3.4(b) shows the initial estimate. The reconstruction algorithm provides a progressively smoother surface fit at iteration 3 in Figure 3.4(c), iteration 5 in Figure 3.4(d), and iteration 10 in Figure 3.4(e). Notice that unlike Gaussian smoothing or moving average filters, the curvature information is preserved while smoothing out the noise. The plot of the energy residual is shown in Figure 3.4(f).

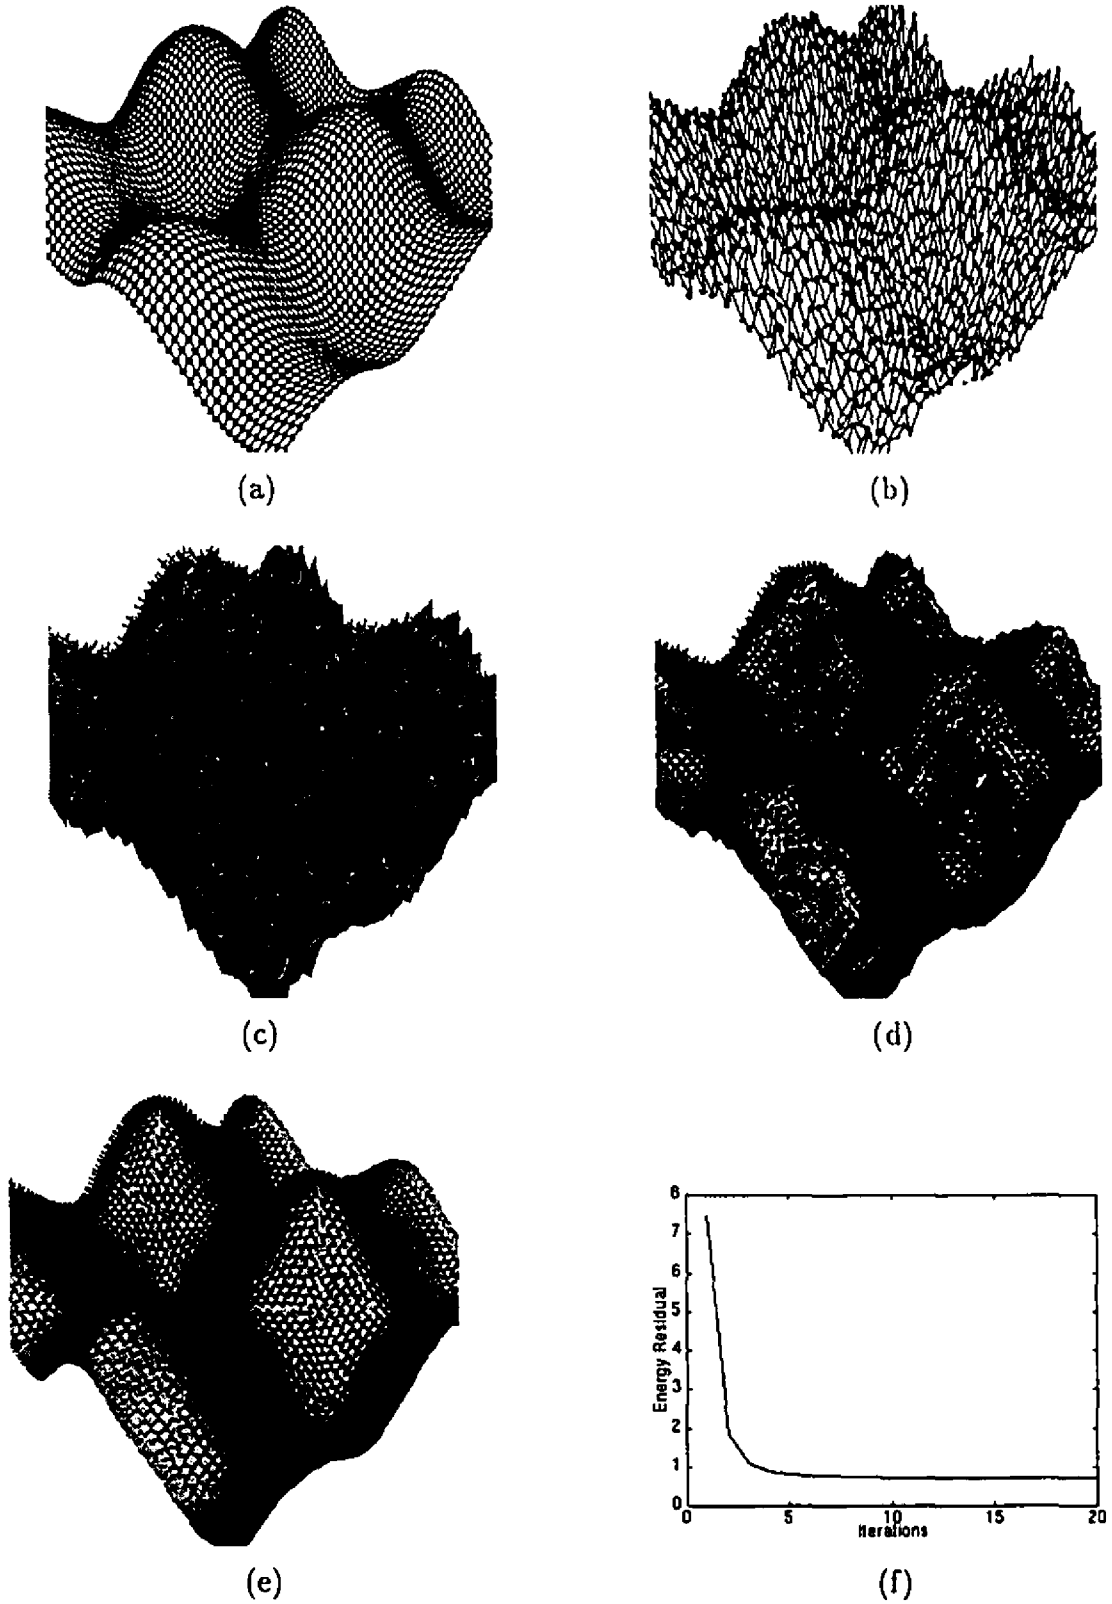


Figure 3.3: (a) A synthetic range image. (b) Noise added to the range points (SNR 5/3). (c) The initial estimate of the darbbox frames and surface patches. Darbbox frames and recovered surface after (d) 3 iterations, (e) 10 iterations of the curvature consistency algorithm. (f) Energy residual plot.

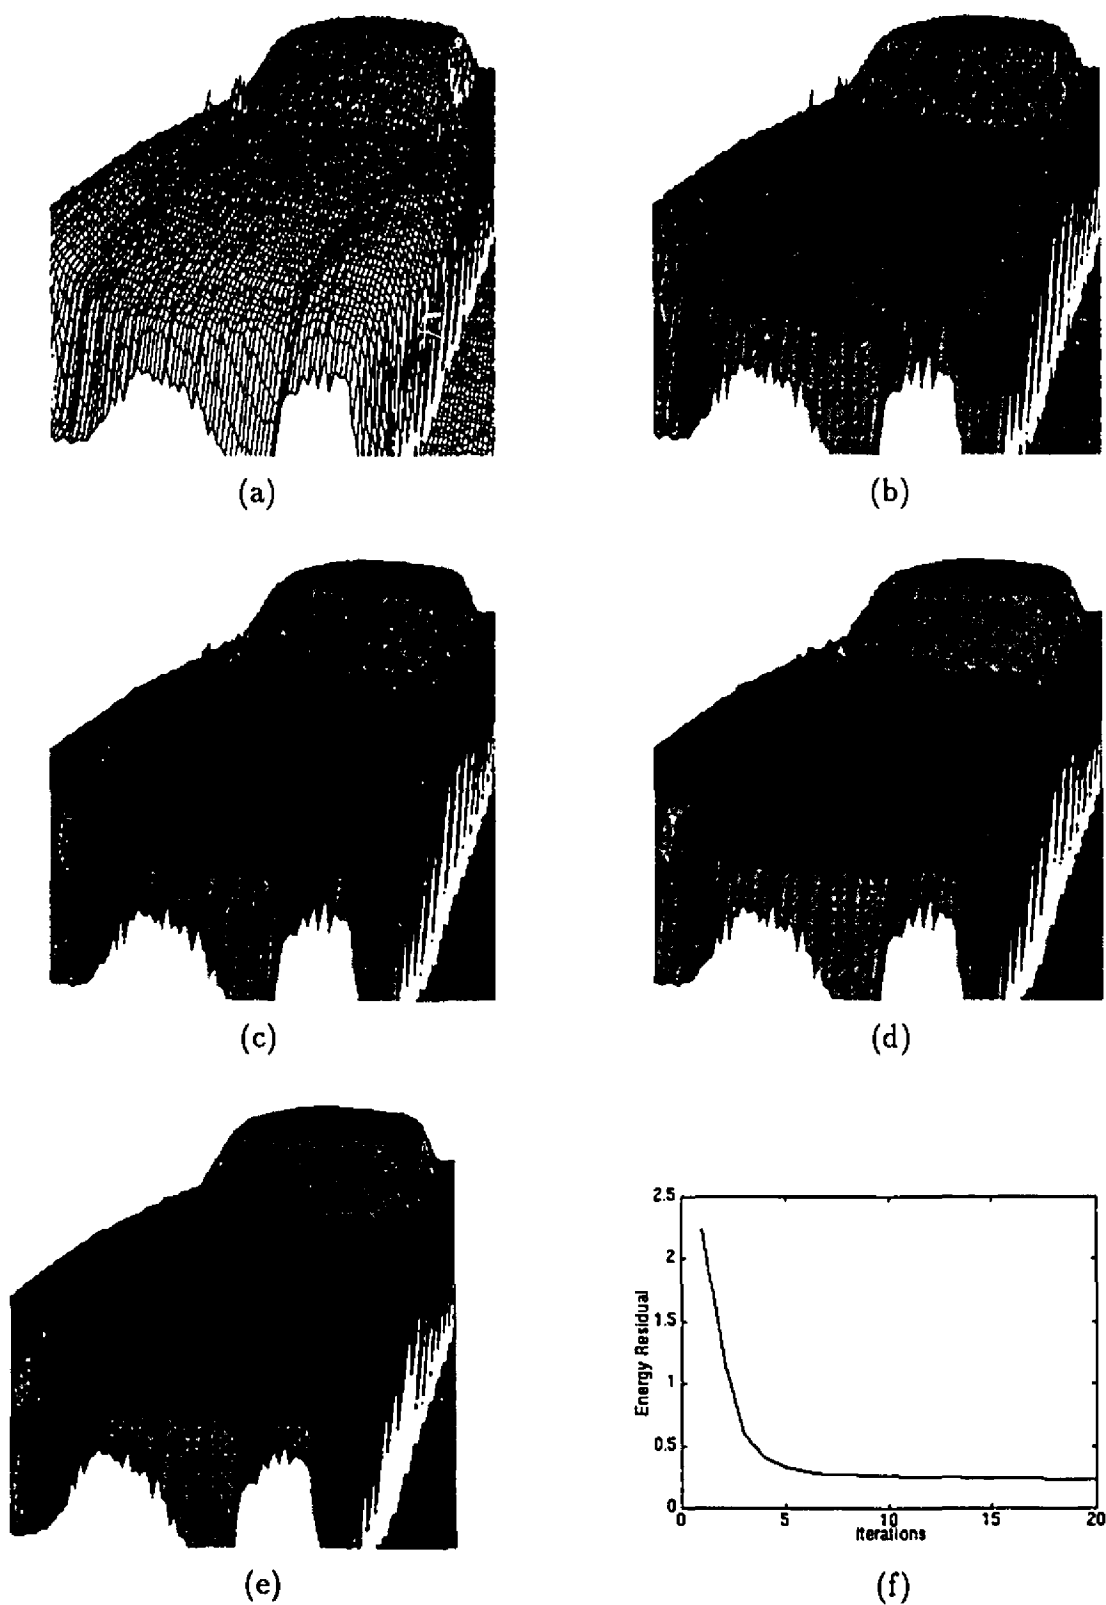


Figure 3.4: (a) A real range image of a toy car. (b) Initial estimate of the surface patches. Recovered surface after (c) 3 iterations, (d) 5 iterations, (e) 10 iterations of the curvature consistency algorithm. (f) Plot of energy residual.

3.5 Chapter Summary

In this chapter the theory behind the curvature consistency surface reconstruction algorithm was presented along with some examples of its application to range image reconstruction. From the theory it can be seen that the curvature consistency algorithm is basically a relaxation process, where the continuity information is propagated through the image with each iteration. A goal of this thesis is to show that in a local neighbourhood, this relaxation process is equivalent to a multiple measurement fusion problem.

The effort in the rest of the thesis is now to show how to obtain an *optimal* estimate of these parameters using the inputs from the various neighbours. In the next chapter a similar multiple-measurement fusion problem is set up and solved purely in a control theory, optimal estimation context. Later in this thesis it will be shown how the curvature consistency algorithm falls under the same optimal estimation framework. Using this framework, and by redefining the concept of “noise” in the image, a new set of update equations will be derived which will result in discontinuity localization while providing a smooth continuous fit to noisy range image data.

The last few chapters established the basis for the surface reconstruction algorithm used in this thesis. In this chapter, we now move on to establish the optimal estimation theory on which the discontinuity localization solution is based.

Estimation theory offers a technique to integrate information from different sources. A more specific and popular example of an estimator is the Kalman filter. With close similarities between the fields of digital signal processing and estimation theory, the use of the Kalman filter has seen tremendous use in signal estimation, image restoration, de-blurring, multi-sensor fusion, and other fields of image processing. The approach in this thesis is to show that the same optimal estimation process can be used to model interactions between pixels in a local neighbourhood in the curvature consistency surface reconstruction algorithm. To understand the optimal estimation theory taking place in this case, it is instructive to consider it separately, out of the context of the reconstruction algorithm.

In Section 4.1, an introduction to optimal estimation theory is presented. A simple example of dual-channel, noisy information fusion is set up and an optimal estimate derived in Section 4.2. In the same section, the example is extended to include several sources providing measurements of the quantity being estimated. This example provides an optimal estimate only from single measurements from each information source. In Section 4.3, a state-space formulation of the same example is derived. In this case, multiple noisy measurements from each source are now combined to provide the optimal estimate. A recursive method of estimating the noise variance of each channel is also developed. Finally a special case is considered, where the noise properties of the sources could change over time. In Section 4.4, the lessons learnt from the optimal estimation example are summarized. These lessons will be useful when applying the same estimation theory to the curvature consistency algorithm in the next chapter.

4.1 Optimal Estimation Theory

Given sets of noisy measurements from multiple channels, over multiple instances of time, if the problem is to estimate an underlying quantity from them, the solution is in the form of an optimal estimator:

“An optimal estimator is a computational algorithm that processes measurements to deduce a minimum error estimate of the state of a system by utilizing: knowledge of system and measurement dynamics, assumed statistics of system noises and measurement errors, and initial condition information.” [23]

The Kalman filter is one such optimal estimator. Given the following information:

1. a linear system model
2. measurement of the behaviour of the system
3. statistical models of system and measurement errors
4. initial condition information

the Kalman filter enables the processing of the measurement data to estimate the system state.

An advantage of the Kalman filter is that it generates its own error analysis. It tracks the accuracy of the estimates that it generates. It is classified as a *recursive filter* since there is no need to store past measurements for the purpose of computing present estimates. All previous information is embodied in a *prior estimate*.

The Kalman filter equations will not be derived here, since it is beyond the scope of this thesis. For more details on the development and use of the filter see [23] [71]. More appropriate would be an illustration of some of the concepts behind the filter and its use which are pertinent to the solution of the discontinuity preserving surface reconstruction problem. An optimal estimation problem is used to instantiate these concepts.

4.2 An Optimal Estimation Problem

Consider the following problem: Two sources of measurement are available, each making a single measurement $q_i (i = 1, 2)$ of a constant but unknown quantity, p , in the presence of random, independent, unbiased measurement errors¹, $v_i (i = 1, 2)$, with known variances. An estimator needs to be designed to combine the two measurements to produce an optimal estimate of p .

The measurements are described by

$$q_1 = p + v_1 \text{ and } q_2 = p + v_2. \quad (4.1)$$

In the absence of any other information, a linear function of the measurements may be obtained which gives an estimate of p ,

$$\hat{p} = \lambda_1 q_1 + \lambda_2 q_2, \quad (4.2)$$

where \hat{p} is an optimal estimate of the constant quantity p and λ_1, λ_2 need to be developed to provide this optimal estimate. Defining the estimation error, ϵ , as

$$\epsilon = \hat{p} - p. \quad (4.3)$$

The optimality criterion is one which would minimize the mean square value of ϵ . In addition, the choice of λ_1 and λ_2 should be independent of the value of p . This condition will hold if the estimate is unbiased, i.e., the expected value of the error

$$E[\epsilon] = E[\hat{p} - p] = 0. \quad (4.4)$$

Expanding (4.4) using (4.2)

$$E[\epsilon] = E[\lambda_1(p + v_1) + \lambda_2(p + v_2) - p] = 0 \quad (4.5)$$

¹In applying optimal estimation theory to surface reconstruction, the sources of error are not technically noise processes, but are treated as such.

4. An Optimal Estimation Problem

$$= \lambda_1 E(p) + \lambda_1 E(v_1) + \lambda_2 E(p) + \lambda_2 E(v_2) - E(p). \quad (4.6)$$

With $E[v_1] = E[v_2] = 0$ and $E[p] = p$, we obtain

$$\lambda_2 = 1 - \lambda_1. \quad (4.7)$$

Combining (4.1) to (4.7), the mean squared error is computed to be

$$E[\epsilon^2] = \lambda_1^2 \sigma_1^2 + (1 - \lambda_1)^2 \sigma_2^2, \quad (4.8)$$

where σ_1^2 and σ_2^2 are the variances of v_1 and v_2 respectively. To obtain optimal values of λ_1 and λ_2 , we differentiate the mean square error with respect to λ_1 , and set the value to zero,

$$2\lambda_1 \sigma_1^2 - 2(1 - \lambda_1) \sigma_2^2 = 0, \quad (4.9)$$

yielding

$$\lambda_1 = \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \quad \text{and} \quad \lambda_2 = \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}. \quad (4.10)$$

Using these values of λ_1 and λ_2 , we obtain the estimator

$$\hat{p} = \left(\frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \right) q_1 + \left(\frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} \right) q_2. \quad (4.11)$$

Put in another form, the estimator is

$$\hat{p} = \left(\frac{1/\sigma_1^2}{1/\sigma_1^2 + 1/\sigma_2^2} \right) q_1 + \left(\frac{1/\sigma_2^2}{1/\sigma_1^2 + 1/\sigma_2^2} \right) q_2. \quad (4.12)$$

The corresponding minimum mean square error of estimation is

$$E[\epsilon^2] = \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2} \right)^{-1}. \quad (4.13)$$

The optimality of the estimator is exhibited by the minimum mean square error (4.13), which is less than either of the mean square errors due to each measurement source.

These results can be easily generalized to the case where there are N sources of measurements each providing measurements q_i ($i = 1 \dots N$) of the static quantity p , each with known error variances σ_i^2 ($i = 1 \dots N$). The optimal estimator can be shown to be

$$\hat{p} = \left(\frac{1}{1/\sigma_1^2 + 1/\sigma_2^2 + \dots + 1/\sigma_N^2} \right) \sum_{i=1}^N (1/\sigma_i^2) q_i. \quad (4.14)$$

The corresponding minimum mean square error of estimation is

$$E[c^2] = \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2} + \dots + \frac{1}{\sigma_N^2} \right). \quad (4.15)$$

Using (4.12) and (4.14) we draw attention to an important characteristic of this optimal estimator: *Each measurement is weighted by a function of its error variance.* This variance-weighting prevents noisy measurements from entering into the estimate \hat{p} . The optimal estimator (4.14) makes sense in the various limits of interest:

- If $\sigma_1^2 = \sigma_2^2 = \dots = \sigma_N^2$, then the measurements are averaged.
- If one measurement q_k is perfect ($\sigma_k^2 = 0$), then the others are rejected.
- Measurements which have low error variances are given more weight over the ones which have worse noise properties.

4.3 State Space Formulation

In the example presented in Section 4.2, *single* measurements from different sources were combined in an optimal sense. The theory can now be extended to combine multiple error-prone measurements from each source. It was also assumed that the noise or error variances σ_i^2 are known a priori for each measurement source i . If we no longer assume this, we need a method of learning the error properties of each source using the multiple measurements provided through that source.

In this section, we will assume initially that the measurement error variance is an inherent property of the input source itself and its value does not change over time. Due to this, the measurement error variance for each source is a constant quantity

which needs to be estimated. At the end of the section we will consider the case where the noise properties of the measurement source can change over time. As an example of how this may occur in the physical world, measurement sensors are affected by external parameters like temperature, pressure changes etc. Assuming that the changes in noise properties are gradual, the error variance estimation method can be made adaptive to accommodate these changes

Note that now, two different processes will estimate two separate but related quantities at each time interval. One process combines measurements from different sources to produce an estimate of the constant quantity \hat{p} at time k . Variance weighting is used in this process to attenuate the contributions from the sources which are noisier than the others. The other process recursively estimates the measurement error variance associated with each measurement source, based on previous measurements through that source.

The noise property embodied in the variance estimate for each source is learnt using the prediction error of the constant quantity p from subsequent measurements. This results in an iterative process, where a better estimate of the constant quantity p , and the error variance σ_i^2 for each source, is obtained with each iteration.

4.3.1 Accommodating Sequences of Measurements

If we are given N sources, each provides a measurement of the constant quantity p of the form

$$q_i(k) = p + v_i(k) \quad i = 1 \dots N, \quad (4.16)$$

with error of measurement $v_i(k)$, at time (state) k . A linear combination of the measurements provides the estimate of p at time k

$$\hat{p}(k) = \sum_{i=0}^N \lambda_i(k) q_i(k), \quad (4.17)$$

where $\lambda_i(k)$ are the coefficients which are derived as in the previous sections, using the same optimality criterion. Again, each λ is just a function of the error variances of the measurement sources. However, since the error variances are not given quantities, but have to be estimated, the variances σ_i^2 have to be replaced by the revised estimate $\hat{\sigma}_i^2(k)$ at time k . $\lambda_i(k)$ is then given by

$$\lambda_i(k) = \frac{1/\hat{\sigma}_i^2(k)}{1/\hat{\sigma}_1^2(k) + 1/\hat{\sigma}_2^2(k) + \dots + 1/\hat{\sigma}_N^2(k)}. \quad (4.18)$$

4.3.2 Recursively Estimating the Noise Variance

The error of estimation for each measurement source i at time k is calculated by

$$\epsilon_i(k) = \hat{p}(k) - q_i(k). \quad (4.19)$$

The measurement error variance at time k for source i is then

$$\hat{\sigma}_i^2(k) = \frac{1}{k} \sum_{l=1}^k \epsilon_i^2(l). \quad (4.20)$$

When an additional measurement becomes available through that source, the new estimate of the variance is

$$\hat{\sigma}_i^2(k+1) = \frac{1}{k+1} \sum_{l=1}^{k+1} \epsilon_i^2(l). \quad (4.21)$$

This expression can be written in another form which makes explicit the prior estimate

$$\hat{\sigma}_i^2(k+1) = \frac{k}{k+1} \left(\frac{1}{k} \sum_{l=1}^k \epsilon_i^2(l) \right) + \frac{1}{k+1} \epsilon_i^2(k+1) \quad (4.22)$$

$$= \frac{k}{k+1} \hat{\sigma}_i^2(k) + \frac{1}{k+1} \epsilon_i^2(k+1). \quad (4.23)$$

Using Eq. (4.23) offers an advantage over Eq. (4.20), in that, past values of $\hat{\sigma}_i^2$ need not be stored. At time $(k+1)$ all previous values are embodied in the prior estimate of the variance $\hat{\sigma}_i^2(k)$. Eq. (4.23) represents a *recursive* estimator for the error variance.

Eq. (4.23) can be written in another form

$$\hat{\sigma}_i^2(k+1) = \hat{\sigma}_i^2(k) + \frac{1}{k+1}(c_i^2(k+1) - \hat{\sigma}_i^2(k)), \quad (4.24)$$

in which the new estimate is obtained by summing the prior estimate with an appropriately weighted quantity. The quantity $(c_i^2(k+1) - \hat{\sigma}_i^2(k))$ is the difference between the squared error of estimation obtained at the present time and the square of its expected mean. This quantity is the *error variance residual*. The error variance residual should typically decrease with time, as more measurements arrive through the source and a better estimate of the measurement error variance is obtained.

To start the estimate for each source i , an initial value is needed for $\hat{\sigma}_i^2$. At the start of the estimation process, there is no prior reason to believe that one source would give more accurate measurements of the constant quantity p than the other sources. Each source should be assigned equal weighting in the estimation of the constant quantity p . Due to this reason, the initial measurement error variance for each source is set to 1. For N measurement sources this gives

$$\lambda_i(0) = \frac{1}{N} \quad \forall i, i = 1 \dots N. \quad (4.25)$$

4.3.3 Sources With Changing Noise Properties

In accommodating sequences of measurements through different sources, we have so far assumed that the measurement error variance is constant over time and is an inherent property of the source itself. To add another twist to the problem, now we assume that the measurement error variance of the measurement sources changes over time. The change in noise properties could be any physical phenomenon beyond the control of the optimal estimation process. However, the estimation process has to adapt to the changes in these noise properties.

This can be done easily by making the estimation of the measurement error variance adaptive. In Eq. (4.20), rather than include all previous estimates of $\hat{\sigma}_i^2$, only the τ previous estimates are used. The expression to calculate the measurement error

variance is then

$$\hat{\sigma}_i^2(k) = \sum_{l=k'}^k \frac{1}{\tau} c_i^2(l) \quad \text{where,} \quad k' = \begin{cases} 1 & \text{if } k \leq \tau \\ k - \tau + 1 & \text{otherwise.} \end{cases} \quad (4.26)$$

Equations (4.23) and (4.24) then become

$$\hat{\sigma}_i^2(k+1) = \frac{\tau}{\tau+1} \hat{\sigma}_i^2(k) + \frac{1}{\tau+1} c_i^2(k+1) \quad (4.27)$$

$$\hat{\sigma}_i^2(k+1) = \hat{\sigma}_i^2(k) + \frac{1}{\tau+1} (c_i^2(k+1) - \hat{\sigma}_i^2(k)). \quad (4.28)$$

In some cases, we would want to give more recent measurements more emphasis in the measurement fusion process. In this case, the recursive estimation of the error variance can be adjusted so that the error variance from more recent measurements are given a higher weight than older ones. This can be done by using an exponential (Gaussian) falloff or a triangular falloff weighting function.

4.4 Chapter Summary

An optimal estimation problem was formulated to illustrate some of the concepts and use of recursive (or Kalman) filtering. The problem presented in this chapter was carefully chosen to be directly relevant to the problem of discontinuity preserving surface reconstruction, discussed in the next chapter.

The lessons learnt so far in this chapter which will directly contribute towards the design of the discontinuity preserving surface reconstruction method are summarized:

Given a set of measurement sources, each providing a sequence of measurements at discrete intervals of a constant quantity, and each with uncorrelated noise properties, optimal estimation theory offers the following attributes:

- It provides a framework to give an estimate of the constant quantity, which is optimal in the sense that it has a minimum mean squared estimation error.
- It tracks the total error of estimation over time, which indicates the confidence

in the estimate.

- Given an initial value, it is possible to learn the error variance attributed to each measurement source.
- An update stage uses the measurements from each source weighted by a function of the corresponding error variances to obtain a new estimate of the constant quantity.

In Chapter 1, the need for proper application of the continuity (or smoothness) constraint to the surface reconstruction problem was discussed. It was stated that the approach taken in this thesis is to apply the smoothness constraint selectively. In the last chapter, we developed a multi-sensor fusion problem and arrived at a method to control the measurement fusion process in such a way that noisy measurement sources would play less of a role. By construing the curvature consistency algorithm as an analogous multi-channel measurement fusion problem, we arrive at an automatic and adaptive mechanism to control the application of the smoothness constraint. This mechanism will prevent a blind application of the constant curvature constraint disregarding any discontinuities in the data.

In this chapter we show how the interaction between pixels in a local neighbourhood during the curvature consistency variational relaxation, or any relaxation process, can also be modeled as a multiple measurement fusion problem. Each neighbour in a local neighbourhood acts as a measurement source. We call these “pixel-sources”. Each pixel-source provides a “measurement” of a constant quantity, which here is the set of Extended Darboux Frame parameters which describe the quadric patch surface interpolating the image samples in that neighbourhood. It was assumed initially that the noise in the image is Gaussian. Due to this noise, measurements provided by each pixel-source are also noisy and follow a Gaussian distribution. The optimal estimation solution developed in the last chapter then applies directly to the curvature consistency variational relaxation smoothing problem. The variational relaxation stage of the curvature consistency can now be modified to offer better estimates of the quadric patch interpolants through the image data.

In addition to demonstrating the link between optimal estimation theory and inter-pixel interaction during relaxation processes, another key concept is introduced which helps in accommodating discontinuities in the image sample data: The concept

of noise in a pixel-source is extended to include another noise component due to structural discontinuities in the local neighbourhood. Using this component, optimal estimation theory offers a framework to localize discontinuities while performing noise smoothing in piece-wise continuous surface regions. The extension of the noise model to include a discontinuity noise model is an important contribution of this thesis.

In developing the solution to the measurement fusion problem in the last chapter, it was assumed that the noise between the different sources was identically and independently distributed. No grouping of sources according to noise variance was considered. However, in the presence of a discontinuity in a neighbourhood, pixel sources are correlated in terms of their noise variances. The discontinuity divides this neighbourhood in two subregions, each subregion defined by a similarity in the noise variance of its constituent pixels. The correlation results from the fact that pixels belonging to the same subregion in the discontinuous neighbourhood demonstrate similar estimation error properties while estimating the Darboux Frame parameters of the interpolating surface patch. The estimation error properties are however different between subregions. To accommodate this pixel grouping, the optimal estimation solution needs to be modified. This is done by introducing a new weighting function which still falls under the optimal estimation framework, but also takes the pixel grouping in discontinuous regions into account. With this form of weighting function true localization of discontinuities can be demonstrated while recovering a smooth continuous surface covering elsewhere.

5.1 Control Processes at the Pixel Level in Curvature Consistency

In the field of image processing, it is common to see estimation theory used to combine multiple images arriving from different sources, at different times, or from different positions. It is interesting to note that the same information fusion theory can also be used to describe information flow processes between neighbouring pixels on a single image in some of the well known image processing algorithms which use relaxation or




Figure 5.1: Direction of information flow into a 3×3 neighbourhood during a relaxation process

diffusion. Rather than multiple images acting as sources of new information, multiple pixels in a region act as information channels.

5.1.1 Information Flow in Relaxation Processes

In relaxation processes, due to the interaction of neighbouring pixels, information is transmitted across the data set. The state of a local neighbourhood is influenced by its constituent pixels. Each pixel then acts as a channel of information for that neighbourhood. Figure 5.1 shows the direction of information flow into a 3×3 neighbourhood from the rest of the image. Hence, given a $I \times J$ discrete neighbourhood of image samples centered at pixel P , there are $I \times J$ separate sources of information Q_i , ($i = 1 \dots I \times J$), each contributing a “measure” of the state of that neighbourhood.

In the curvature consistency algorithm, neighbouring pixels interact to determine each other's surface patch parameters. It also operates at a $I \times J$ pixel neighbourhood level. Each neighbouring pixel Q_i in the $I \times J$ neighbourhood of P provides a measure $D_i(P)$ of a constant quantity, which is the Extended Darboux Frame $D(P)$ of the surface patch centered at P . $D_i(P)$ is obtained using a local transport model that describes how the Extended Darboux Frame at Q_i should appear at P when transported along the surface of the model. $D(P)$ then represents the optimum quadric function which interpolates the data in that neighbourhood. Hence, in information

fusion terms, a measurement model similar to (4.1) can be set up as

$$D_i(P) = D(P) + v_i, \quad i = 1, 2, \dots, I \times J \quad (5.1)$$

where v_i is an error vector assigned to each neighbourhood pixel-source. The components of this error vector are the individual error scalars related to each parameter of the Extended Darboux Frame. The error is due to noise in the measurements. A thorough description of this noise follows next, but before that, equation (5.1) will be extended to include the fact that relaxation processes are iterative. At each iteration, a new measure of $D(P)$ will be provided by the neighbours. Hence similar to (4.16), a state space equation can be set up. At time (state) k , each neighbour of P provides an estimate of the Extended Darboux Frame $D(P)$ with an error of measurement $v_i(k)$.

$$D_i(P)(k) = D(P) + v_i(k), \quad i = 1, 2, \dots, I \times J \quad (5.2)$$

As will be shown in the next section, the error of measurement v_i is not necessarily constant, but may vary with each iteration, assuming there is either Gaussian noise in the image or some discontinuities in the surface structure, or both.

5.1.2 Defining the Noise Property

Given the variation of the data in the neighbourhood set, each pixel will give a different estimate of what the interpolating surface patch parameters should be. Hence, the estimated quantity from a particular neighbour will differ from the optimal estimate obtained by combining the different estimates (or measurements) from all the neighbours. Relaxation processes are iterative, in which the overall state of the image converges with each iteration to a stable underlying state. Since at each iteration, a pixel has new information flowing in and affecting its own state, its contribution to the neighbouring pixels changes too. Hence, the measurement error will vary as the state of the neighbour pixel changes with time. A “noise” property can then be

attached to each pixel which causes this error variance. In the curvature consistency algorithm, this noise property determines the accuracy of measurement of the surface patch parameters from a neighbour.

Here we come to a crucial area of discussion. How exactly do the errors in the measurements arise?

We propose that the variation in the information flow during relaxation is due to two factors:

1. The additive Gaussian noise assumed to be in the image.
2. Discontinuities on the surface which was sampled.

Each reason contributes independently to the error value v in (5.1). The error v can be broken down into two separate components

$$v = v_g + v_d, \tag{5.3}$$

where v_g represents the error component due to the Gaussian noise in the image, and v_d represents the component due to the presence of discontinuities. Each component can be best analyzed by setting the other's contribution to zero.

Setting v_d to zero implies that there are no discontinuities on the surface being sampled. The only reason for error in the measurement of $D(P)$ by the neighbours is due to the additive Gaussian noise in the image. The additive Gaussian noise provides variation in the pixel values over the whole image, which results in the pixel sources in a neighbourhood giving measurements which are different from each other. In a local neighbourhood then, each pixel can be assumed to be sampled from a continuous underlying patch. In providing measurements of the Extended Darboux Frame $D(P)$ over several iterations, every pixel in that neighbourhood is expected to have similar error variances. A straightforward relaxation process is enough to smooth out the Gaussian noise and to give a maximum likelihood estimate of the underlying continuous surface. If σ_i^2 ($i = 1 \dots I \times J$) is the noise variance associated with each pixel Q_i , then in the case where the samples at these pixel positions come

from continuous surface regions,

$$\sigma_1^2 = \sigma_2^2 = \dots = \sigma_i^2 = \dots = \sigma_{I \times J}^2. \quad (5.4)$$

On the other hand setting v_g to zero in (5.3), means that the pixels are sampled from surfaces containing discontinuities but no random Gaussian noise. In that case, the surface can be modeled by piecewise continuous surface patches with combinations of discontinuities in depth, orientation and curvature components. The variation in the pixel-source measurements during relaxation is now due to the presence of these discontinuities. If there is a discontinuity within a neighbourhood, image samples would lie on one or the other continuous patch across the discontinuity. In this discontinuous region, pixels on opposite sides of the discontinuities provide different measurements of $D(P)$ from each other, while within a continuous patch there is no source of error in measurement. If a neighbour Q_i provides a measurement of $D(P)$, when Q_i and P lie on the same continuous patch, the measurement is going to be exact, i.e., $v = 0$. However, if Q_i lies on the other side of the discontinuity as P , the measurement is not going to be exact, and can be represented by:

$$D_i(P) = D(P) + v_{g_i}. \quad (5.5)$$

For the pixels on the same side of the discontinuity in a local neighbourhood, the error variances are going to be equal. They will however, have different variance values from the pixels on the opposite side.

A typical set of image samples will probably contain both additive Gaussian noise and surface discontinuities. If in the reconstruction process, only the Gaussian noise is accounted for, then discontinuities will be considered as being due to the Gaussian noise process too. If a simple relaxation process is used to eliminate this noise, then the discontinuities will be distorted as a consequence. On the other hand, if only the discontinuities are accounted for, then variation in the surface due to the Gaussian noise will be inferred as discontinuities as well, and in a discontinuity localization process, there will be edges found everywhere.

The shortcoming of the original curvature consistency algorithm is that it only accounts for the v_g error component in (5.3) and simple relaxation is used to smooth out this noise. Therefore, it will perform well in recovering surfaces from data containing only additive Gaussian noise but will fail to preserve any discontinuities.

This situation can be remedied by using a cooperative method, where the process of evaluating neighbouring depth samples on the same continuous patch complements the process of evaluating a discontinuity between them. By this method, the measurement operation (5.2) can be restricted to occur between image samples which are perceived to be in the same subregion in the discontinuous neighbourhood. Within each continuous subregion, the measurements from each neighbour will presumably contain only the v_g error component, which can be easily smoothed out by the relaxation process, without the influence of neighbours from the other side of the discontinuity.

To restrict the measurement operation in this manner requires knowledge of the presence of discontinuities in the local neighbourhood. Since discontinuities are not given a priori, this cooperative process cannot be applied straight off. However, starting from an initial estimate, with each iteration of curvature consistency the (dis)continuous nature of the local surface can be learnt. At each iteration, the measurements from each neighbour can be weighted according to the current belief that it is on the same continuous patch as the pixel at which the interpolating surface is centered. In the next section, it is shown how it is possible to use the optimal estimation techniques discussed in Chapter 4 to learn the discontinuous nature of the surface between two neighbouring pixels and to apply curvature consistency in a piecewise continuous manner.

5.1.3 Optimally Combining Measurements From Neighbours

As has been already shown, the pixel level interactions taking place during the curvature consistency variational relaxation stage is very similar to the information fusion

example shown in the last chapter. Hence, given the measurement model of (5.2) at iteration k of the curvature consistency algorithm, using estimation theory and (4.17), the optimal estimate of $D(P)$ is of the form

$$\hat{D}(P)(k) = \sum_{i=0}^{I \times J} \lambda_i D_i(P)(k), \quad (5.6)$$

where each λ_i needs to be determined to provide this optimal estimate.

The value of λ_i can be seen to be the same as was derived in Section 4.3.1.

$$\lambda_i(k+1) = \frac{1/\hat{\sigma}_i^2(k)}{1/\hat{\sigma}_1^2(k) + 1/\hat{\sigma}_2^2(k) + \dots + 1/\hat{\sigma}_{NM}^2(k)} \quad i = 1 \dots I \times J, \quad (5.7)$$

where $\hat{\sigma}_i^2(k)$ is the current estimate of the measurement error variance for neighbour Q_i in estimating $D(P)$ ¹. At each iteration, the error of measurement is calculated as

$$\epsilon_i(k) = \hat{D}(P)(k) - D_i(P)(k), \quad (5.8)$$

which is then used to calculate the current estimate of the error variance

$$\hat{\sigma}_i^2(k) = \frac{1}{k} \sum_{l=1}^k \epsilon_i^2(l). \quad (5.9)$$

As shown in Section 4.3.2, a recursive estimation process can be used to update $\hat{\sigma}_i^2(k)$ instead of (5.9)

$$\hat{\sigma}_i^2(k) = \hat{\sigma}_i^2(k-1) + \frac{1}{k}(\epsilon_i^2(k) - \hat{\sigma}_i^2(k-1)). \quad (5.10)$$

This eliminates the need to store all the values for ϵ_i over all the iterations. Only the error variance $\hat{\sigma}_i^2(k-1)$ obtained in the last iteration is needed.

Many researchers [49] [54] [70] have used the first derivative (difference between neighbouring pixel quantities) as an estimate of the continuity of the surface. Equation 5.8 is similar to taking a spatial first derivative in the discrete form (difference

¹There is a slight change in notation between (5.7) and (4.18) regarding the index of λ_i . Since the estimate of the measurement error variance used in calculating λ_i is actually obtained in the previous iteration, the index of λ_i is changed to $k+1$, rather than k .

equation), except that rather than taking the difference between the neighbouring quantities, we are taking the difference between the current estimate of the quantity at pixel P , and the estimate by neighbour Q_i of the quantity at P .

One other property of relaxation processes needs to be modeled in control theory terms. This property is associated with the change in the noise properties of pixel sources over time according to new information content flowing into the neighbourhood. During a relaxation process, information travels across the image. There is travel time associated with the information which depends on the neighbourhood size being used. If $L_{nbd} = I \times J$ is a dimension of the local neighbourhood size being used in the relaxation process then, for information to travel a distance of L pixels it would take L/L_{nbd} number of iterations. If drastically new information reaches a neighbourhood from a far region through one of the neighbourhood pixel sources, the measurement error variance could change. Hence, the neighbourhood pixels should be treated as sources with changing noise properties. Fortunately, this situation has been already considered in Section 4.3.3. The same solution is applied to the case of curvature consistency relaxation process. The variance estimation process should be made adaptive by considering only the τ previous estimates in (5.9), i.e.,

$$\hat{\sigma}_i^2(k) = \sum_{l=k'}^k \frac{1}{l} \epsilon_i^2(l) \quad \text{where,} \quad k' = \begin{cases} 1 & \text{if } k \leq \tau \\ k - \tau + 1 & \text{otherwise.} \end{cases} \quad (5.11)$$

This results in an equation similar to (4.28) which gives a current estimate of the noise variance of the pixel source whose noise properties changes over time

$$\hat{\sigma}_i^2(k) = \hat{\sigma}_i^2(k-1) + \frac{1}{\tau} (\epsilon_i^2(k) - \hat{\sigma}_i^2(k-1)). \quad (5.12)$$

As mentioned in Section 4.3.3, if more recent measurements need to be favored over older measurements, then a Gaussian or triangular falloff weighting function can be applied to (5.11). The more recent the measurement, the higher the weight given in its contribution to the overall estimate of the noise variance. In this thesis, rather

than weighting the variance contribution over time, we set $\tau = 1$, which implies

$$\hat{\sigma}_i^2(k) = \epsilon_i^2(k) = [\hat{D}(P)(k) - D_i(P)(k)]^2. \quad (5.13)$$

This effectively disregards the history of the error variance, but provides computational simplicity.

5.1.4 Optimal Estimation in the Curvature Consistency Formulation

It is now worthwhile to analyze how the original curvature consistency formulation fits into the optimal estimation framework and vice-versa. In Section 3.3 it was shown that through a surface transport model (Fig. 3.2) each neighbour in a local neighbourhood centered at pixel P provides measures $D_i(P)$ of the Extended Darboux Frame at P . Equations (3.4), (3.5), (3.6), (3.7), and (3.9) show how the individual components of $D_i(P)$ can be combined to give an estimate of $D(P)$. Analyzing these equations, it can be seen that the Extended Darboux Frame measurements $D_i(P)$ from each neighbour are being combined in the form:

$$\hat{D}(P) = \frac{1}{I \times J} \sum_{i=1}^{I \times J} D_i(P). \quad (5.14)$$

This is similar to the form in equation (5.6), which was derived using an optimal estimation and measurement fusion framework. Comparing equations (5.6) and (5.14), the value of λ_i used in (5.14) can be seen to be a constant:

$$\lambda_i = \frac{1}{I \times J} = \frac{1}{n}. \quad (5.15)$$

What this implies is that each neighbour q_i is being given the same weight in the measurement fusion equation (5.14) in each iteration. This is a natural consequence of the assumption that each neighbour lies on a continuous underlying surface. Comparing equations (5.15) and (5.7), it can be seen that the former equation can be

derived from the latter by setting

$$\sigma_1^2 = \sigma_2^2 = \dots = \sigma_i^2 = \dots = \sigma_{I \times J}^2. \quad (5.16)$$

From the discussion in Section 5.1.2, it is known that this situation arises when Gaussian noise is assumed to be present in the image, but no discontinuities. If discontinuities were present in the image, then they would be distorted by using the update equation in (5.14).

5.1.5 Modification to the Update Equations

In a previous section it was shown that in regions of discontinuities, not every neighbour can be expected to give measurements of $D(P)$ with the same accuracy due to the presence of the v_d error component in (5.3). A neighbour Q_i which is on the same side of the discontinuity² as P will possibly give consistently better estimates of $D(P)$ than a neighbour Q_j which is on the opposite side. Due to the more accurate measurements, Q_i would have a lower error variance than Q_j .

If the value of λ_i is computed as (5.7) rather than (5.15), an optimal estimate of $D(P)$ would be obtained from all the measurements. It is then possible to weight the measurements of each neighbour according to its error variance. This kind of variance weighting encourages measurements from neighbours which are on the same side of the discontinuity as P and inhibits the contributions from neighbours which are across the discontinuity. This achieves the goal discussed in Section 5.1.2 of performing the curvature consistency relaxation process only between neighbours which are on the same continuous subregion in the discontinuous neighbourhood, hence eliminating the v_d error component in Eq. (5.3).

At the beginning of the curvature consistency iterations, the error variance values $\hat{\sigma}_i^2(k)$ are not available. Hence, initial values are needed for them. There is no prior reason to believe that one neighbour would give more accurate measurements of $D(P)$

²This reasoning applies best in an $I \times J$ neighbourhood to step and roof edges when I and J are odd numbers. This way the central point is always one side or the other of the discontinuity

than the others. In this case each neighbour should be assigned equal weighting in the measurement update equations. The initial error variance for each source is set to 1, i.e.,

$$\lambda_i(0) = \frac{1}{I \times J} \quad \forall i, i = 1 \dots I \times J. \quad (5.17)$$

What this achieves in the modified curvature consistency algorithm is that at first, full variational relaxation is allowed to take place over the whole image, smoothing out the Gaussian noise, and admittedly, starting to distort the discontinuities. However, as the iterations progress, the error variance of each neighbour in a local neighbourhood is learnt, and just after a few iterations the discontinuity localization process starts playing a major role. The consequent iterations then follow a stable pattern where the curvature consistency relaxation process takes place only within piecewise continuous regions and not across them.

At this point, it should be remembered that $D(P)$ is actually a vector composed of the Extended Darboux Frame components, i.e., $\mathcal{D}(P) = (P, M_P, \mathcal{M}_P, N_P, \kappa_{M_P}, \kappa_{\mathcal{M}_P})$. Hence in calculating $D(P)$ from the neighbourhood measurements Q_i , each component of $D(P)$ needs to be determined through this optimal estimation process. Each component of $D(P)$ would have an independent λ_i component associated with it, i.e., $\lambda_i = (\lambda_{iP}, \lambda_{iM_P}, \lambda_{i\mathcal{M}_P}, \lambda_{iN_P}, \lambda_{i\kappa_{M_P}}, \lambda_{i\kappa_{\mathcal{M}_P}})$. In a similar vein, $\sigma_i^2 = (\sigma_{iP}^2, \sigma_{iM_P}^2, \sigma_{i\mathcal{M}_P}^2, \sigma_{iN_P}^2, \sigma_{i\kappa_{M_P}}^2, \sigma_{i\kappa_{\mathcal{M}_P}}^2)$. In other words, each component of $D(P)$ is updated independently using the corresponding component of λ_i , which in turn is determined by the corresponding component of σ_i^2 .

Using the above discussion, we arrive at the modified curvature consistency update equations which use optimal estimation theory as a basis. Equations (3.4), (3.5), (3.6), (3.7), and (3.9), which are the update equations for each of the components of the Extended Darboux Frame $D(P)$, are now written as follows:

$$P(k+1) = \sum_{i=1}^n \lambda_{iP}(k) P_i(k) \quad (5.18)$$

$$N_P(k+1) = \frac{(\sum_{i=1}^n \lambda_{iN_P} N_{x_{P_i}}(k), \sum_{i=1}^n \lambda_{iN_P} N_{y_{P_i}}(k), \sum_{i=1}^n \lambda_{iN_P} N_{z_{P_i}}(k))}{\sqrt{(\sum_i \lambda_{iN_P} N_{x_{P_i}}(k))^2 + (\sum_i \lambda_{iN_P} N_{y_{P_i}}(k))^2 + (\sum_i \lambda_{iN_P} N_{z_{P_i}}(k))^2}} \quad (5.19)$$

$$\kappa_{M_P}(k+1) = \sum_{i=1}^n \lambda_i \kappa_{M_P} \kappa_{M_P n}(k) \quad (5.20)$$

$$\kappa_{\mathcal{M}_P}(k+1) = \sum_{i=1}^n \lambda_i \kappa_{\mathcal{M}_P} \kappa_{\mathcal{M}_P n}(k) \quad (5.21)$$

$$M_P(k+1) = \bar{b}_1 \cos \theta(k+1) + \bar{b}_2 \sin \theta(k+1), \quad (0, 2\pi) \quad \text{such that} \quad \begin{aligned} 1) \quad & \bar{b}_1, \bar{b}_2 \in T_P \\ 2) \quad & \|\bar{b}_1\| = \|\bar{b}_2\| = 1 \\ 3) \quad & (\bar{b}_1 \cdot \bar{b}_2) = 0, \end{aligned} \quad (5.22)$$

$$\theta(k+1) = \tan^{-1} \left[\frac{(A_{22} - A_{11}) + \sqrt{(A_{11} - A_{22})^2 + 4A_{12}^2}}{2A_{12}} \right], \quad (5.23)$$

where

$$A_{ij} = \sum_{i=1}^n \lambda_i M_P(k) \cdot \bar{b}_i (M_P(k) \cdot \bar{b}_j). \quad (5.24)$$

Each component of λ_i is calculated as

$$\lambda_{ij}(k+1) = \frac{1/\hat{\sigma}_{ij}^2(k)}{1/\hat{\sigma}_{1j}^2(k) + 1/\hat{\sigma}_{2j}^2(k) + \dots + 1/\hat{\sigma}_{nj}^2(k)}, \quad \forall j, j \in (P, M_P, \mathcal{M}_P, N_P, \kappa_{M_P}, \kappa_{\mathcal{M}_P}), \quad (5.25)$$

where

$$\hat{\sigma}_{ij}^2(k) = \hat{\sigma}_{ij}^2(k-1) + \frac{1}{\tau} (\epsilon_{ij}^2(k) - \hat{\sigma}_{ij}^2(k-1)), \quad \forall j, j \in (P, M_P, \mathcal{M}_P, N_P, \kappa_{M_P}, \kappa_{\mathcal{M}_P}), \quad (5.26)$$

and

$$\epsilon_{ij}(k) = \hat{D}_j(P)(k) - D_{ij}(P)(k), \quad \forall j, j \in (P, M_P, \mathcal{M}_P, N_P, \kappa_{M_P}, \kappa_{\mathcal{M}_P}). \quad (5.27)$$

The coefficients λ developed in (5.7) and more specifically for each component of $D(P)$ in (5.25), are the weighting functions assigned to each pixel in a neighbourhood. The properties of the optimal estimator using this kind of variance weighting were presented in the last chapter in Section 4.2. The variance weighted curvature

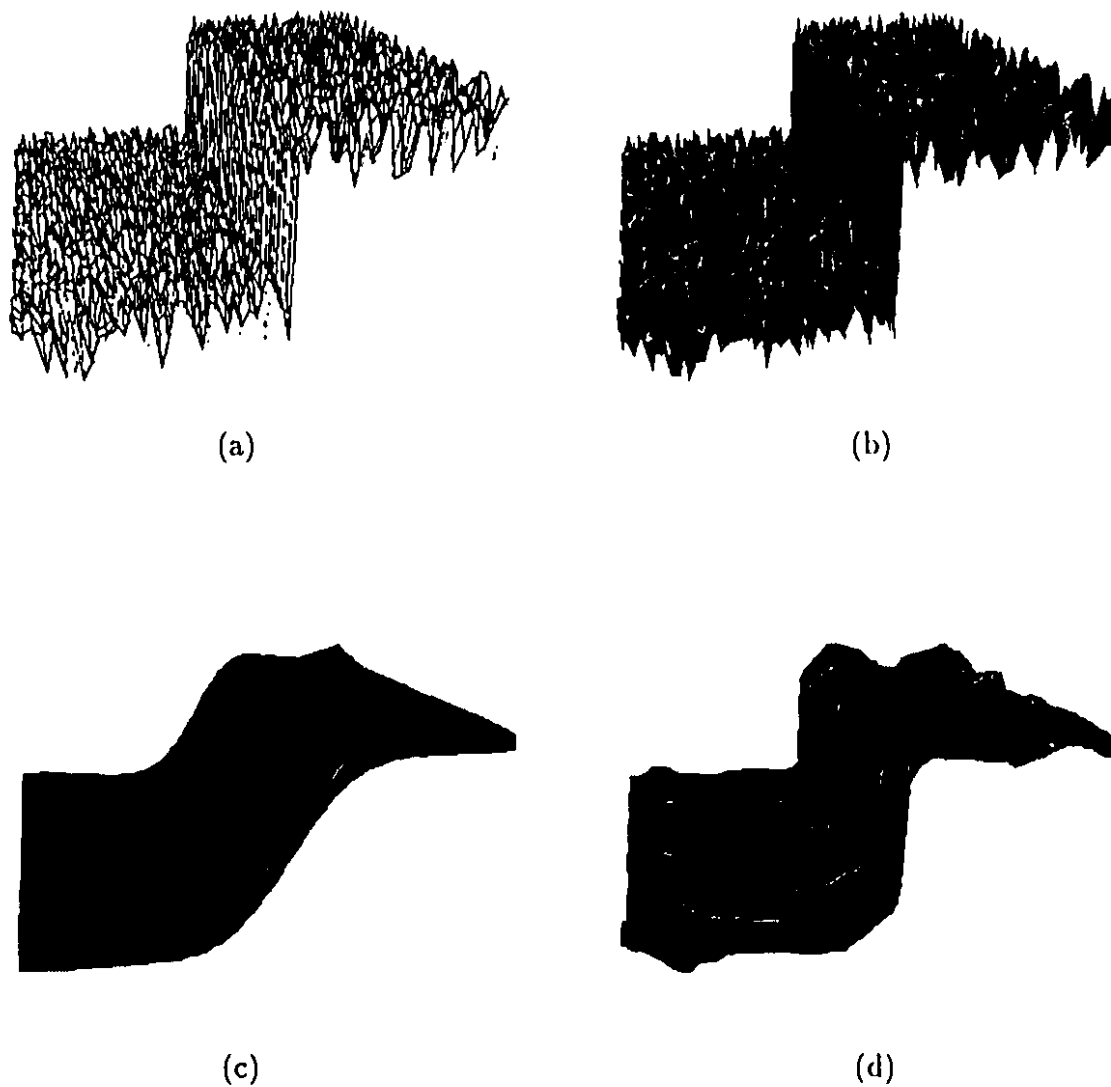


Figure 5.2: (a) A simulated noisy range image of a step junction. (b) Initial estimate of the surface patches. (c) Underlying surface recovered using the original curvature consistency algorithm after 40 iterations. (d) Surface reconstruction using the curvature consistency modified to incorporate variance weighting.

consistency update equations demonstrate similar properties:

1. Neighbours which demonstrate low error variances in estimating the interpolant patch parameters are given more weight over the ones which have high error variances.
2. If $\sigma_1^2 = \sigma_2^2 = \dots = \sigma_N^2$, for all neighbours Q_i , then the measurements are averaged. This is the case in regions of ideal continuity.
3. If one measurement $D_i(P)$ is perfect ($\sigma_i^2 = 0$) from a neighbour Q_i , then the contributions from the others are rejected.

Figure 5.2 shows the results of applying the modified curvature consistency update equations to a synthetic range image of a noisy step junction. Figure 5.2(c) shows the results of applying the original curvature consistency algorithm using a 5×5 neighbourhood area after 40 iterations. As expected, the step edge is distorted completely while obtaining a smooth surface fit. Figure 5.2(d) shows the results of applying the variance weighted update equations after 40 iterations. From this figure it can be seen that although discontinuity localization has been achieved, *a smooth continuous surface is not obtained*. After just a few iterations the surface patches congeal to form many piecewise continuous local neighbourhoods. This result shows that there is something else we need to take into consideration while modifying the curvature consistency algorithm to fit the optimal estimation framework. The concept missing here is of pixel groupings in discontinuous regions. In other words, in the optimal estimation framework, we need to consider non-IID measurement sources.

5.2 Accommodating Pixel Groupings in Discontinuous Regions

In the optimal estimation framework developed in Chapter 4, and applied to curvature consistency in the last section, an assumption was made that the noise in the different sources was IID (Identically and Independently Distributed). No correlation between the different sources was assumed with respect to their noise properties. However, in

the case of surface samples on a discontinuous region of a surface there is a grouping according to similarity in estimation error variances. In this case, pixels on the same side of the discontinuity are correlated in their error variances. This correlation comes about due to the inclusion of the noise component v_d due to discontinuities in the noise model defined in (5.3). The v_g noise component can be safely assumed to be uncorrelated between neighbours over the whole image, however in the presence of an edge, pixels on the same side of the edge are correlated in the v_d noise component. This correlation necessitates the inclusion of the concept of pixel groupings in the neighbour weighting process.

The variance weighting method developed in the last section gave us only partial results. It achieved the goal of discontinuity localization, but did not give a smooth surface description. To propagate a continuity constraint, relaxation needs to occur within pixel neighbourhoods. However, in using the variance weighting method from the last section, this relaxation process was being curtailed in continuous regions. In the simple variance weighting process it is possible for a single pixel source to take full control of the relaxation process in a local neighbourhood. This could happen if initially this pixel source provided more accurate predictions than the rest of the pixels in the neighbourhood. This pixel source would then get maximum weight in the next iteration and possibly in the consecutive iterations. Soon, there would be no inflow of information into that neighbourhood through the relaxation process. The surface patch covering that neighbourhood would “freeze” and take on the Extended Darboux Frame parameters that this influential pixel source estimates.

In the presence of a discontinuity, the continuity constraint needs to be propagated through relaxation only within the subregions partitioned by the discontinuity, but not across the discontinuity. If pixels in one subregion give better estimates of the surface patch parameters than the other subregions, then the *whole* set of pixels should be given higher weights than the rest. Within this highly weighted continuous subregion however, the distribution of weights amongst the pixels should be relatively uniform. The inter-region difference in weights will prevent the continuity constraint from being applied across the discontinuity. The intra-subregion similarities in weights




Figure 5.3: A 5×5 neighbourhood mask straddling a (a) step discontinuity, (b) a roof discontinuity. The mask is centered at P . The discontinuity divides the neighbourhood into two subregions **A** and **B**.

would encourage relaxation to take place within the subregions, hence providing a continuous surface fit within each partitioned subregion.

The above operation can be illustrated with Fig. 5.3. Figure 5.3(a) shows a 5×5 neighbourhood centered at P straddling a step discontinuity, and Figure 5.3(b) shows 5×5 neighbourhood straddling a roof discontinuity. In both cases, pixels in subregion **A** consistently provide better estimates of the surface patch covering, centered at point P than the pixels in subregion **B**. All the pixels in this subregion should have higher weights than the pixels in subregion **B** during the weighted curvature consistency process. However, amongst themselves, pixels in subregion **A** should have similar weights, and similarly pixels within subregion **B** should have similar weights. This would stop the flow of the continuity constraint information across the discontinuity, but relaxation would take its normal course within each continuous subregion.

5.2.1 Mapping Error Variances to a New Function

A new weighting functions needs to be used to accommodate grouping of pixels. This weighting function should still be under the optimal estimation framework, in the




Figure 5.4: Distribution of weights according to a Gaussian error function.

sense that the mean square error of estimation after combining information from all the pixel sources should be less than the mean square error due to each pixel source separately. In addition, to accomplish pixel groupings in discontinuous regions, weights should be assigned in such a manner that groups of pixels rather than individual pixels should be favored in these regions.

The desired weighting function should now have the following properties:

1. The lower the error variance, more the weight given to the source.
2. The function should be bounded as the error variance approaches zero.
3. Groups of pixels providing more accurate predictions should be given higher weights.
4. Within continuous regions or subregions (across discontinuities) if the prediction errors are similar between the pixel sources, the weighting should be similar too.

A mapping function which encapsulates these desired properties is one which follows a Gaussian trace (Fig. 5.4)

$$y = e^{(-\frac{x^2}{\gamma})}, \quad (5.28)$$

where γ is a smoothing control parameters which will be discussed later in this section.

Instead of using simple variance weighting, each source's error variances should be mapped onto this modified Gaussian function to generate a value, which is normalized

across the sources to obtain the final weight.

The weights for the elements of $\mathcal{D}(P)$ to be used at each iteration step are then calculated as

$$W_i(k) = e^{(-\frac{\delta^2(k)}{\gamma})} \quad (5.29)$$

$$\lambda_i(k) = W_i(k) / \sum_{i=1}^n W_i(k) \quad (5.30)$$

This function is monotonically decreasing, as shown in Fig. 5.4. Property 1 would then be satisfied, i.e., the lower the error variance, the lesser the weight given to the pixel. Property 2 is satisfied since the function is bounded at the value 1 as the error variance approaches zero. The γ parameter controls the spread of the weight distribution. This parameter is adaptive over the local neighbourhoods, and is calculated for each iteration, and for each local neighbourhood. The adaptive method to derive values for γ is discussed next, where it will also be shown how this new function satisfies properties 3 and 4.

5.2.2 Making γ Adaptive

The γ parameter controls the amount of smoothing. As used in (5.30) it determines the spread of the weighting function. The standard deviation of the Gaussian function is determined by $\delta = \gamma/2$. The spread of the Gaussian weighting is determined by the spread in the estimation error variance values of all the pixel sources in the neighbourhood.

Some other researchers ([54] [49] [41]) have used functions similar to (5.30) in their adaptive smoothers, but have left the parameter which determines the spread of the Gaussian as a user-set parameter. According to their reasoning, by setting high values for this parameter, more smoothing rather than discontinuity localization would occur, and by setting low values, discontinuity preservation would be favored over smoothing. This parameter being user-set in some ways defeats the whole purpose of having an “adaptive” data smoother. We intend to make this an adaptive parameter in our implementation. The value of this parameter is determined by taking the mean

of all the error variance values,

$$\gamma = \frac{2}{I \times J} \sum_{i=1}^{I \times J} \hat{\sigma}_i^2(k). \quad (5.31)$$

In fact, therein lies the strength of our adaptive smoother. With the spread of the function directly related to the mean of the error variances, there would always be pixels sources on both sides of the standard deviation. This means that unless each pixel source has exactly the same error variance as the others, there will always be some which will be given more weight than the others. If however, the γ parameter is kept constant as cited above, there is a danger of setting it too low, in which case noise may be interpreted as discontinuities, or too high, in which case legitimate discontinuities may be smoothed over.

The effect of using a Gaussian mapping function is illustrated in Fig. 5.5. The figure shows the weight distribution over two 3×3 neighbourhoods taken from the range image shown in Fig. 5.2(a). The samples in the first neighbourhood have been obtained from a continuous surface region. The second set of image samples are obtained from a neighbourhood which contains a step edge discontinuity.

In the neighbourhood containing samples from the continuous region, each pixel source shows similar estimation error properties. These error variances represented by the 'x' symbols on the plots, are mapped onto the X-axis in Fig. 5.5(a). The error variances can be seen to be clustered close together over a small range. After obtaining the value of γ using (5.31), the weights for each pixel source are obtained using the Gaussian weighting function in (5.30). Since the estimation error properties between pixel sources in this neighbourhood are similar, the weights assigned to them are relatively similar too.

Pixel sources in one subregion of the discontinuous neighbourhood show better estimation error properties than the other. As can be seen from Fig. 5.5(b), the error variance is distributed over a larger range on the X-axis. All the samples from the subregion which give worse estimates are mapped on the higher end of the X-axis. After mapping them on the Gaussian trace it can be seen that between the




Figure 5.5: Distribution of weights for the depth component according to a Gaussian function amongst pixel sources in a 3×3 (a) continuous neighbourhood, and (b) discontinuous neighbourhood. The 'x' symbols on the X-axis represent the error variance estimate of the different pixel sources; the '*' symbols represent the corresponding weights assigned according to the Gaussian function; the 'o' symbol on the X-axis is the mean of error variances of all the pixel sources; the '+' symbol shows the corresponding mapping of this mean onto the Gaussian trace.




Figure 5.6: (a) A simulated noisy range image of a step junction. Surface reconstruction using a 5×5 neighbourhood size after 40 iterations (b) using the original curvature consistency algorithm, (c) using variance weighting in the algorithm, (d) using Gaussian weighting.

two subregions there is a large difference in the assigned weights. However within each subregion (mapped on either sides of the '+' symbol), the weights are relatively similar.

Using Fig. 5.5 as an illustration, it can be seen that the Gaussian weighting function proposed in (5.30) does satisfy properties 3 and 4 presented in the last section. In continuous regions or subregions divided by discontinuity, pixel sources are assigned similar weights and they are grouped together in their weighting.

Having established the legitimacy of the new weighting function, we modify the curvature consistency algorithm and apply it to the same image of a noisy step junction as in Fig. 5.2. The result is presented in Fig. 5.6. For reference purposes, the surface recovered using the original curvature consistency algorithm is shown in Fig. 5.2(b), the surface recovered using variance weighting is shown in Fig. 5.2(c) and finally the surface recovered using the Gaussian weighting is shown in Fig. 5.2(d). As can be seen from Fig. 5.2(d), the Gaussian weighted curvature consistency algorithm provides us with a surface fit with the smoothness of the original curvature consistency algorithm, and the discontinuity preservation features of the variance weighting method.

5.3 Chapter Summary

In this chapter, two diverse areas of research: surface reconstruction theory and optimal estimation theory, were linked and the curvature consistency algorithm modified to perform discontinuity localization while smoothing in continuous regions of a noisy range image.

It was shown that during the curvature consistency iterations, each pixel in a local neighbourhood acts as an information channel during variational relaxation. At each iteration, the surface patch parameters are updated using contributions from each pixel in a neighbourhood. Similarities were shown between the curvature consistency update equations and the multiple measurement fusion equations.

The concept of noise in the pixel sources was redefined. Two separate components

of the noise property were identified. One of these components was associated with discontinuities present in the neighbourhood. The optimal estimation framework provides a method to weight down contributions from measurement sources with high noise levels. Using this framework, a pixel source weighting method was incorporated into the curvature consistency update equations. With the new inverse variance weighting method, the continuity constraint was prevented from being applied across discontinuities. Since, pixel sources in a neighbourhood containing a discontinuity are correlated in their noise properties, another weighting function was developed. This function accomplished pixel grouping by mapping the error variances onto a Gaussian trace.

The discontinuity localization method which has resulted from this chapter has been made completely adaptive. In fact the only parameter to set is the neighbourhood mask size. Even this parameter can be made constant by using the same neighbourhood mask for most reconstruction operations. In the next chapter it will be demonstrated that the discontinuity localization property of the modified surface reconstructor remains invariant with different mask sizes. It will also be shown that the algorithm is robust over different scales of discontinuities and adaptive over regions with different noise properties.

In this chapter some of the practical issues related to the modified curvature consistency algorithm are discussed. A detailed analysis is performed on some of the parameters related to the surface fitting and discontinuity localization process. Several results of applying this algorithm to synthetic and real range images are presented to validate the robustness of the algorithm.

It is assumed that the surface sampled in the range image contains continuous regions separated by discontinuities, which are mainly in depth and orientation. The performance of the modified curvature consistency algorithm on all these regions will be analyzed. The changes in the surface patch estimates over the iterations of the modified curvature consistency algorithm are characterized in a general manner. This is followed by a more detailed analysis of the changes in surface patch parameters, and weighting values in selected neighbourhoods over different images. Pixel neighbourhoods straddling step edges and roof edges are analyzed separately due to the difference in the behaviour of the algorithm in their vicinity. This is followed by results of applying the algorithm to images containing more complex artificially generated surfaces, and to real range data.

One practical issue dealt with in this chapter is iteration control and convergence. To get a good surface fit, each image requires a different number of iterations of the modified curvature consistency algorithm. Each iteration produces a new global error of fit over the images. The global error of fit at each iteration provides a measure of the amount of change the surface patches have gone through from the previous iteration. The method of computing the global error of fit in the original curvature consistency algorithm is modified to take into account the fact that all pixels may not contribute equally to the update of the interpolant surface patch parameters. A new method of calculating the global error of fit is then developed for the modified curvature consistency algorithm.

Using the global error of fit, convergence properties of the modified algorithm are compared to the convergence properties of the original algorithm for different kind of surfaces. A discussion is also presented on the state of the surface patches at ideal convergence, i.e., when the global error of fit is zero. The conditions necessary to achieve this state are also presented.

A claim which is made in this thesis about the modified curvature consistency algorithm is that it is adaptive over different image properties. The discontinuity localization feature is shown to be robust and adaptive over changes in scale and noise properties in an image. To validate the claims that this algorithm is robust over scale-space, experiments are conducted using different mask sizes and by using discontinuities of different scales. To demonstrate adaptability over noise, the algorithm is applied to images containing regions with varying noise properties and using masks of different sizes.

6.1 Qualitative Analysis

In this section details and experiments are presented to demonstrate the operation of the modified curvature consistency algorithm. It is shown how the Gaussian weighting function encourages relaxation to take place only within continuous regions of the range image samples, or within continuous subregions partitioned by discontinuities. Over several iterations of the modified curvature consistency algorithm, the changes in the surface estimates are tracked. In addition, certain relevant neighbourhoods are chosen, and the prediction error variance of each neighbour and the corresponding weights are tracked over these iterations for certain interpolant patch parameters. Using this method, the modified curvature consistency algorithm is analyzed qualitatively.

Step edges and roof edges are treated separately in this analysis. In a curvature consistency framework, with the presence of a step edge in a neighbourhood, the surface patch parameter which is effected most by the discontinuity is the position (P) parameter (depth in a range image context). For roof edges, the parameters

affected most are the normals (N) or one of the principal curvatures (κ_1 or κ_2). Due to the difference in the effected parameters, the algorithm operates differently in the vicinity of the two kinds of discontinuities. It is worth analyzing the changes in the surface patch parameters separately for the two. First, a general behavioral pattern of the algorithm is established which can be observed on all kinds of surface data.

6.1.1 General Behaviour of the Algorithm

In the first iteration of the algorithm, all pixels in a local neighbourhood play an equal role in the update of the interpolant patch parameters. Due to the equal confidence given to all neighbours, at first the emphasis is on applying the constant curvature constraint. During these initial iterations, some of the additive Gaussian noise is smoothed out, and the algorithm performance is very similar to the original curvature consistency algorithm. Since the estimation error properties of the constituent pixels have not been learnt yet, surface patches at pixels straddling discontinuities undergo some deformation. However, after just a few iterations the estimation error variance of each neighbour is learnt and the parameters of patches at discontinuity points start being corrected. The surface patch estimates at the discontinuity points start being influenced by neighbours in one subregion more than others. With further iterations, the variational relaxation process continues only within continuous neighbourhoods, and within continuous subregions partitioned by the discontinuities. The propagation of the constant continuity constraint is not allowed across discontinuities due to the weighting between neighbouring subregions.

6.1.2 Localization of a Step Edge

The behaviour of the algorithm as described above can be seen in action in the surface plots shown in Figure 6.1 for an image containing a step discontinuity. The two key features of the algorithm are seen to be satisfied: In continuous regions of the surface data, the surface patches are smoothed by applying the continuity constraint; In the discontinuous region of the image samples, the continuity constraint is inhibited to

give true discontinuity localization.

The most interesting region to watch is the one straddling the step edge. The surface patches in the vicinity of the edge initially get distorted when the continuity constraint is applied with almost equal weights given to each neighbour. As can be seen in Figures 6.1(b) and (c), at iteration 2 and 4 respectively, the positions of the set of pixels straddling the top of the edge are effected by neighbours on both sides of the discontinuity almost equally. It can be seen from the figures that the positions of pixels on top of the step edge are pulled down and the ones at the bottom of the edge are pulled up. This state of affairs lasts only for those few iterations. Once the estimation error properties of each neighbour of those pixels are learnt and the weighting in the update equations starts taking effect, they are now influenced by pixels in one subregion more than others. The pixels near the top are influenced more by their neighbours on top of the step edge, and the pixels near the bottom start being influenced more by the neighbours at the bottom of the edge. With iterations 10, 20 and 50 in Figures 6.1(d), (e) and (f) respectively, these pixels move back to the top or bottom of the edge, depending on the subregion which supports them the most.

For pixels close to the step edge, the Extended Darboux Frame parameter which undergoes the maximum change during the iterations is the position parameter P . Figures 6.2 and 6.3 show the estimation error variance and weight distribution amongst pixels in two different neighbourhoods during the update of the position parameter. The first neighbourhood is a 5×5 neighbourhood similar to the one shown in Figure 5.3(a) straddling the step edge discontinuity. The other neighbourhood is a 5×5 neighbourhood in a continuous region of the range samples, away from the discontinuity.

Figure 6.2(a) shows the estimation error variance of each pixel in the neighbourhood containing the discontinuity after the first iteration. Recall that for the first iteration, the weights assigned to all the pixels were equal. The error variance values are mapped on the X-axis, and the corresponding weights mapped using the Gaussian weighting function presented in Section 5.2 of the last chapter. At this point,




Figure 6.1: Changes in the surface structure around discontinuities with progressive iterations of the modified curvature consistency algorithm, using a 5×5 mask, after (a) initial patch estimate, (b) 2 iterations, (c) 4 iterations, (d) 10 iterations, (e) 20 iterations, (f) 50 iterations

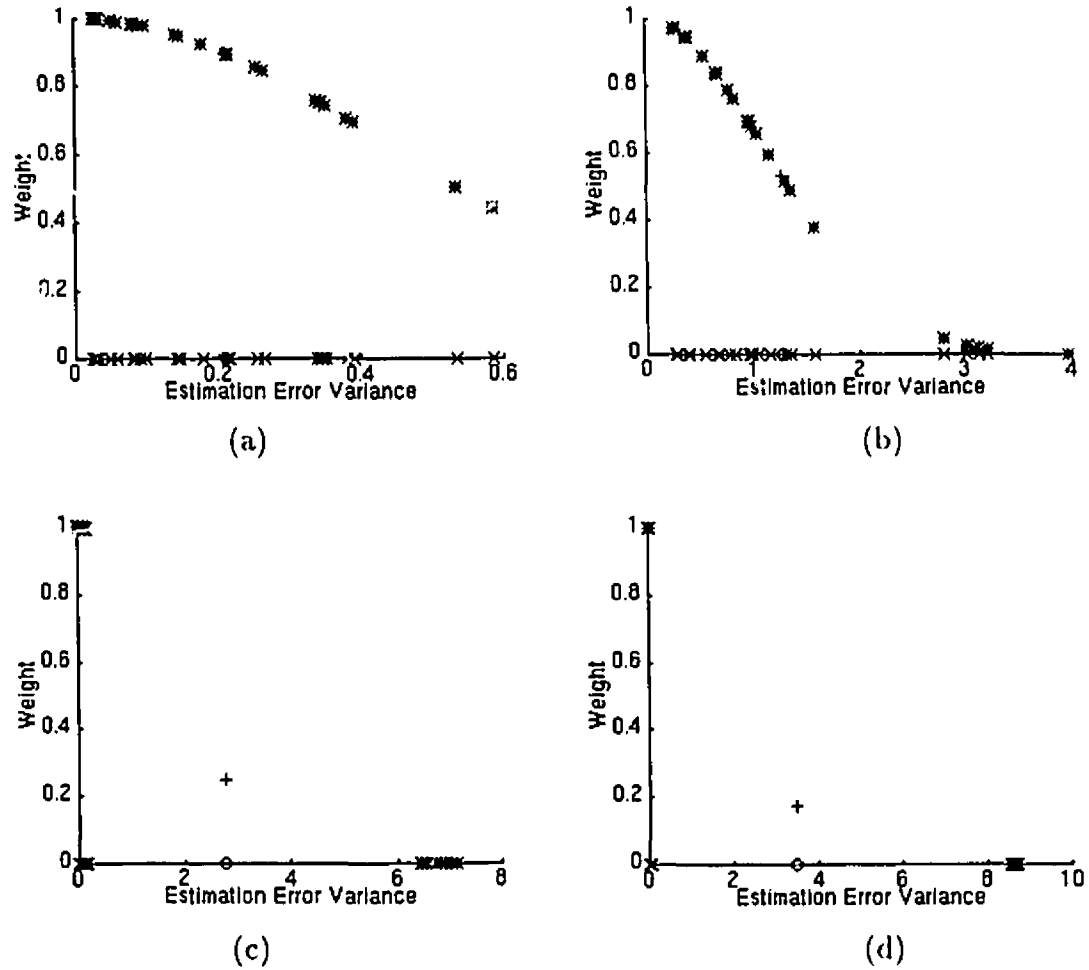


Figure 6.2: The estimation errors and the corresponding weights for the position (depth in range image) component of the Extended Darboux Frame at a pixel straddling the discontinuity using the modified curvature consistency algorithm. Using a 5×5 mask after (a) 1 iteration, (b) 4 iterations, (c) 10 iterations, (d) 20 iterations.

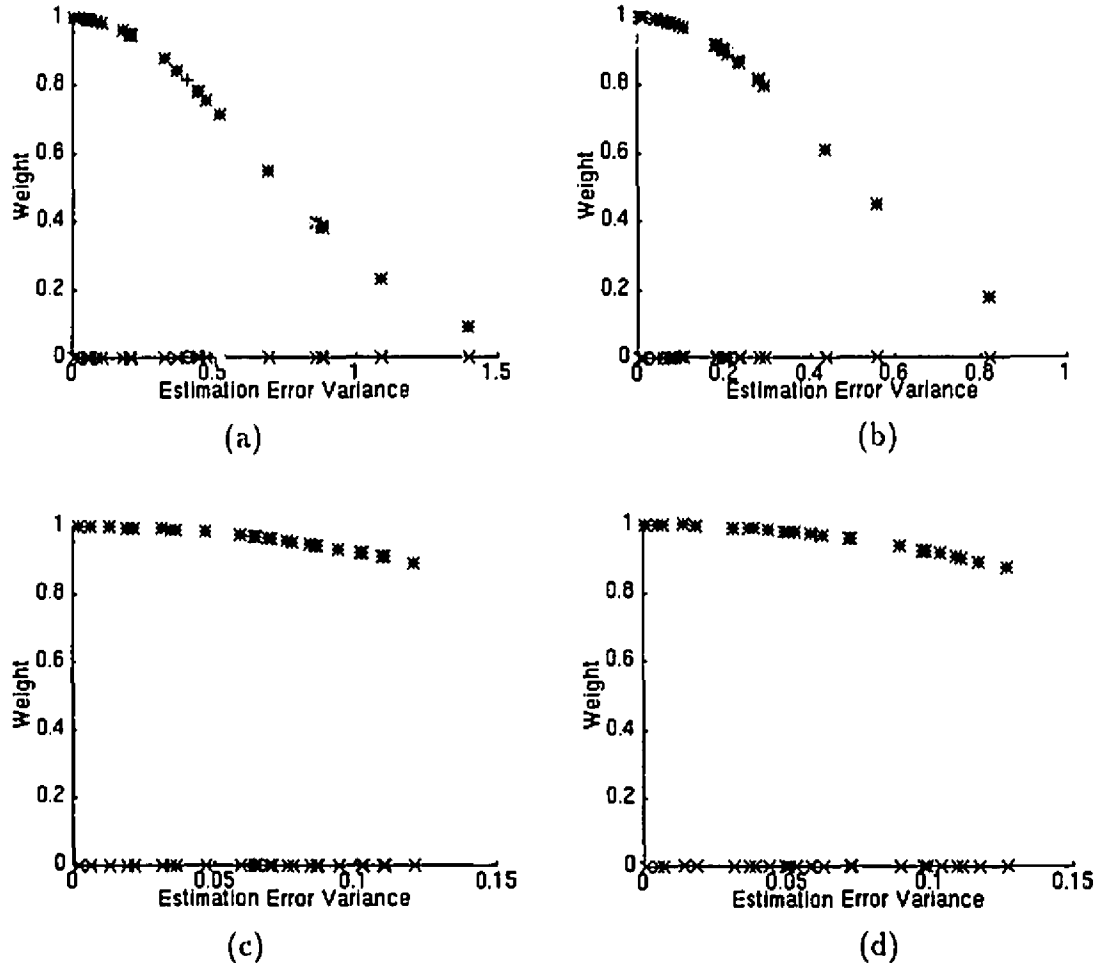


Figure 6.3: The estimation errors and the corresponding weights for the position (depth in range image) component of the Extended Darboux Frame at a pixel in a continuous region of the surface data using the modified curvature consistency algorithm. Using a 5×5 mask, after (a) 1 iteration, (b) 4 iterations, (c) 10 iterations, (d) 20 iterations.

the segregation of pixels according to their estimation error properties starts taking place. The neighbours which belong to the subregion at the bottom of the step edge provide the worst estimate of the position parameter of the surface patch. Using the Gaussian mapping function they are given low weighting for the update stage in the next iteration. After the 4th iteration (Figure 6.2(b)), the segregation of subregions according to their estimation properties is well underway. The constant curvature constraint starts being applied only between the continuous subregions. By iteration 10 (Figure 6.2(c)), discontinuity localization has been achieved. At this stage and after (Figure 6.2(d)), the curvature consistency algorithm continues to perform smoothing, but only within continuous subregions of the neighbourhood.

Figure 6.3 demonstrates the effect of the same algorithm on the estimation error variance and weighting of pixels in a continuous neighbourhood of the sample data. Due to the presence of high amount of additive Gaussian noise, initially the pixels in the neighbourhood provide estimates of the position parameter with varying degrees of accuracy (Fig. 6.3(a)). However, the estimation error properties of all the pixels start to show some similarities as the surface patch parameters are refined even more (Figure 6.3(b)). Soon they all start giving similar estimates of the surface patch position parameter, and hence are assigned similar weights (Figure 6.3(c) and (d)). At this stage the continuity constraint is being propagated freely over the whole neighbourhood, with each neighbour playing an almost equally important role in refining the interpolant surface patch parameters.

6.1.3 Localization of a Roof Edge

The behaviour of the modified curvature consistency for roof edges is similar to the general behaviour described before. Whereas the parameter of the surface patch effected most in the presence of a step edge is the position P , in the presence of a roof discontinuity, the surface normal N or one of the principal curvatures (κ_1 or κ_2) go through the maximum change.

The effect of the original curvature consistency algorithm and of the modified one can be seen in Figure 6.4. Figures 6.4(c) and (d) show the surface structure after




Figure 6.4: Surface reconstruction of a simulated noisy range image containing a roof discontinuity. (a) The original image, (b) initial estimate of surface patches. Using a 5×5 mask, distortion of the edge after (c) 10 iterations and (d) 20 iterations of the original curvature consistency. Preservation of the edge after (e) 10 iterations and (f) 20 iterations of the modified algorithm.

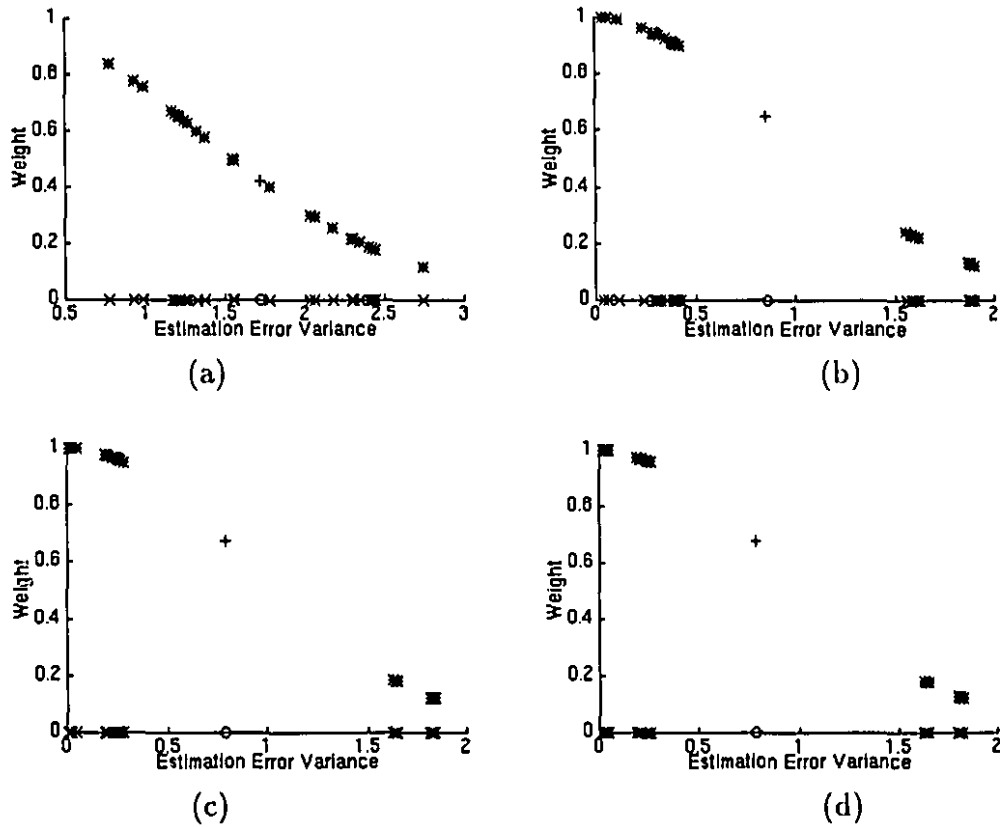


Figure 6.5: The estimation errors and the corresponding weights for the Normal component of the Extended Darboux Frames at a pixel straddling the roof discontinuity while applying the modified curvature consistency algorithm. Using a 5×5 mask, the estimation error variance and corresponding weights after (a) 1 iteration, (b) 5 iterations, (c) 10 iterations, (d) 20 iterations.

10 and 20 iterations respectively, of the original algorithm. It can be seen that the surface patches close to the roof edge are distorted in their normals and curvature components.

Initially, the surface patches at pixels straddling the discontinuity have a high curvature component, while the ones away from the discontinuity have a lower curvature component. The directions of the normals are different between pixels across the discontinuity. Due to the continuity constraint, in the relaxation process of the original algorithm, the high curvature values of pixels straddling the discontinuity are forced lower and the low curvature values of their neighbours which are further from the discontinuity, are forced higher. This process coupled with the averaging of the normal component results in surface patch estimates which are more continuous in their normals and curvatures. However, this continuity constraint also destroys the roof edge by making it bulbous.

From Figures 6.4(e) and (f), it can be seen that the modified curvature consistency algorithm applies the curvature and normal continuity constraint only between pixels which are on the same side of the roof discontinuity. Figure 6.5 tracks the weighting of pixels in a neighbourhood straddling the roof discontinuity during the refinement of the normal (N) component of the interpolant surface patch. The neighbourhood analyzed is similar to the one shown in Figure 5.3(b). In accordance with the general behaviour of the modified curvature consistency algorithm, each neighbour plays an equal role initially in the refinement of the normal component. This results in a slight distortion of the surface patch estimate in the discontinuous neighbourhood. However, after the first iteration, the estimation error properties of the neighbours, with respect to the normal component, are quickly learnt (Figure 6.5(a)). In Figure 6.5(b), it can be seen that after the 5th iteration the grouping of weights according to subregions is well established. At iterations 10 and 20 (Figures 6.5(c) and (d)) and subsequent iterations, the discontinuity localization is fully achieved. During the update of the normal component, pixels in only one subregion play a major role. The influence of the pixels in the subregion across the discontinuity is curtailed due to the low weights assigned to them.




Figure 6.6: Surface reconstruction of an image containing discontinuity in position, as well as curvatures and normals. (a) Original data of a discontinuous sinusoidal trace in grid form, (b) initial fit of surface patches to the image. Using a 3×3 mask, (c) surface reconstruction after 20 iterations of the original curvature consistency algorithm, (d) surface reconstruction after 20 iterations of the modified algorithm.

6.1.4 Discontinuity Localization in Complex Images

Surfaces in real images are assumed to be composed of continuous areas and a combination of step and roof discontinuities. The neighbouring surface patches may be discontinuous not just in their positions, but also in their normals, curvatures and principal directions. So far, the images analyzed have contained either step discontinuities or roof discontinuities. Using images containing combinations of discontinuities of different types and range images obtained from a laser range finder, it will be demonstrated that the modified algorithm performs well on more complex images.



(a)



(b)



(c)

Figure 6.7: (a) Range image of an owl statue acquired with the NRCC/McGill laser range-finder. (b) Reconstructed surface after 8 iterations of the original algorithm. (c) Reconstructed surface after 8 iterations of the modified algorithm.




Figure 6.8: Surface reconstruction of an image of two overlapping rocks, acquired using the NRCC/McGill University laser range-finder; (a) Original range data in grid form; (b) Initial fit of surface patches to the image; (c) Surface reconstruction after 50 iterations of the original curvature consistency algorithm; (d) Surface reconstruction after 50 iterations of the modified algorithm.

Figure 6.6(a) shows an artificially generated image of a surface following a sinusoidal trace. A step discontinuity is introduced in the image. The surface is discontinuous not just in position, but also in the normals and curvatures across the discontinuity. After noise is added to the ideal image data, the first estimate of the surface patches is shown in Figure 6.6(b). Using a 3×3 mask the result of applying the original curvature consistency algorithm (Figure 6.6(c)) is contrasted with the result of applying the modified algorithm (Figure 6.6(d)). It can be visually verified that the modified algorithm is better at localizing the discontinuities in depth, normals and curvatures.

To demonstrate the application of the original and the modified algorithms on range data obtained from real objects, experiments were conducted on images gathered from the NRCC/McGill laser range finder camera [52]. Figure 6.7(a) shows a range image of an owl statue. The surface in this image consists of several natural discontinuities like the step edge discontinuity around the rim of the eyes, and a roof discontinuity at the nose position. In addition there are concave discontinuities between the body of the owl and the wings. If segmentation of the owl surface is required, then it is imperative to preserve discontinuities between different surface regions. After 8 iterations, the original algorithm blurs away most of the edges and completely loses some of the finer curvature discontinuities (Figure 6.7b). The modified algorithm, on the other hand, preserves most of the edges and structure (Figure 6.7c).

A similar result is demonstrated for a range image of two overlapping jagged rocks in Figure 6.8(a). The jaggedness and the boundary between the two rocks are seen to be better preserved by using the modified algorithm (Figure 6.8(c)) rather than the original one (Figure 6.8(b)).

6.2 Iteration Control

One problem in using the original curvature consistency algorithm is to establish when to stop iterating. As described earlier in Section 3.3, one approach is to track

the global error of fit between iterations [15, 37, 56]; when it falls below a particular threshold, the process is stopped. The same approach can be taken for the modified algorithm as well. However the convergence behaviour is slightly different between the original and the modified curvature consistency algorithms.

In the original curvature consistency, the error of fit of a surface patch in a single $I \times J$ neighbourhood centered at P at iteration number k is:

$$e_p(k) = \sum_i^{I \times J} ((\hat{P}(k) - P_i(k))^2 + (\hat{M}_P(k) - M_{P_i}(k))^2 + \quad (6.1)$$

$$(\hat{\mathcal{M}}_P(k) - \mathcal{M}_{P_i}(k))^2 + (\hat{N}_P(k) - N_{P_i}(k))^2 + \quad (6.2)$$

$$(\kappa_{\hat{M}_P}(k) - \kappa_{M_{P_i}}(k))^2 + (\kappa_{\hat{\mathcal{M}}_P}(k) - \kappa_{\mathcal{M}_{P_i}}(k))^2) \quad (6.3)$$

where $\hat{D}(P)(k) = (\hat{P}(k), \hat{M}_P(k), \hat{\mathcal{M}}_P(k), \hat{N}_P(k), \kappa_{\hat{M}_P}(k), \kappa_{\hat{\mathcal{M}}_P}(k))$ are the updated estimates of the Extended Darboux Frame parameters of the interpolant patch after iteration k and $D_i(k) = (P_i(k), M_{P_i}(k), \mathcal{M}_{P_i}(k), N_{P_i}(k), \kappa_{M_{P_i}}(k), \kappa_{\mathcal{M}_{P_i}}(k))$ are the estimates of these parameters provided by neighbour Q_i . To calculate the global error of fit, the e_p for all neighbourhoods in the image are summed.

$$e(k) = \sum_j^{width \times height} e_{p_j} \quad (6.4)$$

In the modified curvature consistency algorithm, the neighbours are weighted during the update stage for each parameter. In the presence of a discontinuity, this weighting will have an impact on the error of fit calculation. In providing the updated patch parameters $\hat{D}(P)(k)$, each neighbour Q_i is assigned a certain weight $\lambda(P)(k) = (\lambda_P(k), \lambda_{M_P}(k), \lambda_{\mathcal{M}_P}(k), \lambda_{N_P}(k), \lambda_{\kappa_{M_P}}(k), \lambda_{\kappa_{\mathcal{M}_P}}(k))$. Similarly each neighbour should be assigned the same set of weights while calculating the error of fit in the neighbourhood. Equation (6.1) then becomes:

$$e_p(k) = \sum_i^{I \times J} ((\hat{P}(k) - P_i(k))^2 * \lambda_{P_i}(k) + \quad (6.5)$$

$$(\hat{M}_P(k) - M_{P_i}(k))^2 * \lambda_{M_{P_i}}(k) + \quad (6.6)$$

$$(\hat{\mathcal{M}}_P(k) - \mathcal{M}_{P_i}(k))^2 * \lambda_{\mathcal{M}_{P_i}}(k) + \quad (6.7)$$

$$(\hat{N}_P(k) - N_{Pi}(k))^2 * \lambda_{N_{Pi}}(k) + \quad (6.8)$$


$$(\kappa_{\hat{M}_P}(k) - \kappa_{M_{Pi}}(k))^2 * \lambda_{\kappa_{M_{Pi}}}(k) + \quad (6.9)$$

$$(\kappa_{\hat{\mathcal{M}}_P}(k) - \kappa_{\mathcal{M}_{Pi}}(k))^2 * \lambda_{\kappa_{\mathcal{M}_{Pi}}}(k) \quad (6.10)$$


The global error of fit in the modified curvature consistency algorithm can be expected to converge more rapidly to zero than in the original algorithm. In the modified algorithm, due to the weighting process, the surface patches may become more dependent on one subregion of the neighbourhood more than others. In that case the surface patch parameters are influenced by less numbers of neighbours, and hence go through less of a change from one iteration to another. Since the global error of fit is just a measure of the change taking place in the image from one iteration to the next one, it converges to a zero value much faster. The original algorithm suffers from a slower convergence purely due to the fact that there are more neighbours involved in influencing the surface patch estimates, hence there is more potential for change in the surface patch parameters from one iteration to another.

The global error of fit for the step image and the roof image, shown in Figure 6.1 and Figure 6.4 respectively, are tracked for both the original and the modified curvature consistency algorithm. The results are shown in Figure 6.9. The modified curvature consistency algorithm out-performs the original one with respect to rate of convergence. For the step image (Fig. 6.9(a)), in the early iterations both algorithms are matched in the rate of convergence, but in the later iterations the modified algorithm provides a lower global error of fit. For the roof image (Fig. 6.9(b)), the global error of fit in the original algorithm does not even converge. In this case however, the modified algorithm actually forces the global error of fit to convergence.

With both versions of the algorithm the iterations can proceed infinitely. The convergence only tends to zero (real valued) but never reaches it. Practically, it is therefore prudent to stop the iteration when the global error of fit reaches a certain tolerable value. Iterative algorithms for surface reconstruction converge very slowly in general [63]. A local Fourier analysis of the error function from one iteration to another is done in [6]. To paraphrase Terzopoulos in [63]:



(a)



(b)

Figure 6.9: Comparison of global error of fit when the original (solid line) and modified (dashed line) curvature consistency algorithm are applied to the image containing (a) the step discontinuity, (b) the roof discontinuity.

Such an analysis shows that the high frequency components of the error—those components with wavelengths on the order of the grid spacing—are short lived. On the other hand, low frequency components persist through many iterations. Thus a typical behaviour of an iterative algorithm is that the error decreases quickly during the first few iterations while the high-frequency components are being liquidated, but then settles down to a slow, asymptotic behaviour when only low-frequency components remain. The point is that although relaxation is inefficient at completely liquidating the error, it is very efficient at smoothing it out.

In the ideal case of convergence, all neighbours should lie on the same piecewise continuous surface. Since we are dealing with a patchwork of local surface fits, this implies that all the local surface patches are C^2 continuous with each other at the extents of the neighbourhood support. Depending on the surface transport model we choose in the curvature consistency algorithm, it is possible that at ideal convergence each neighbour lies on the same continuous patch.

An iso-surface check can easily determine whether two points are on the same surface or not. The idea behind using the iso-surface check is as follows. If a neighbour Q_i has prediction error zero for the surface patch parameters centered at P , then after applying the transport model, P should fall on the trace of the surface patch centered at Q_i . If ideal convergence has been reached, then the reverse should also be true: When the transport model is applied to the patch centered at P , then Q_i should also fall on the trace of this surface patch.

6.3 Adaptiveness of the Modified Algorithm

In certain images it is possible for image properties to vary across the image. The image may contain discontinuities of different scales and of different types (step, roof or a combination of both). The noise over the image may vary across the image. This may happen for example in a range image taken with the camera at a small angle to the surface being imaged. In this case, the laser beam from the laser range-finder

camera would not be perpendicular to the surface, but would subtend an angle with the surface. Due to this slight foreshortening of the surface, samples from the surface further away from the laser receiver optics may be more noisy than the ones closer to it. In these cases, it is imperative that the surface reconstruction algorithm adapt to the changing properties of the surface and give a consistent surface fit. In this section, the discontinuity localization property of the modified curvature consistency is shown to be robust over changes in these image properties.

In the curvature consistency algorithm, all the computations are done locally, using data in a neighbourhood defined by a mask size. These local computations provide the ability to use, compute and store information related to each neighbourhood separately from the other surrounding neighbourhoods. If image properties are different between two neighbouring regions, the algorithm can adapt to this difference. Due to this local processing the algorithm has the capability to adapt to changes in surface structure, noise level and scales over the same image.

6.3.1 Scale Space Robustness

The idea of discontinuities in data is intuitively connected to the concept of scale. At one scale some surface structure may be considered important while at a larger scale, to be undesirable noise and smoothed over. In several applications such as volumetric fitting to range data, segmentation (parts decomposition) is performed on the reconstructed surface. One approach to segmentation is to look for surface features marked by extremal values of curvature (e.g. negative local minima, concave discontinuities, etc.) [15, 17, 31]. Parts are then segmented along boundaries comprised of these features. It is thus important that the reconstruction algorithm be largely scale invariant over the same image, i.e., surface features of all scales be preserved during the reconstruction process. The robustness of the segmentation algorithm depends on these features.

The concept of scale is also related to the neighbourhood mask size used in a convolution process. The mask size determines the granularity of the features being measured. In the convolutions used for the curvature consistency variational relax-

ation process, the size of the mask determines the number of neighbours providing estimates of the interpolant surface patch parameters. In the original curvature consistency algorithm, increasing mask sizes distort discontinuities at higher rates. Larger neighbourhood sizes distort discontinuities even more rapidly than the smaller ones. However, in the modified algorithm this is not true. The discontinuity preservation property of the modified algorithm remains intact between differently sized masks.

The idea of scale space filtering in 2D images has been developed by Witkin in [72]. Koendrik [36] showed the equivalence between scale space filtering with heat conduction or diffusion equations. He stated two criteria for the diffusion equation formulation. These criteria were extended by Perona and Malik in [49]. Any paradigm for generating multi-scale “semantically meaningful” description of images must satisfy:

1. Causality: No spurious details should be generated passing from finer to coarser scale.
2. Immediate Localization: At each resolution, the region boundaries should be sharp and coincide with the semantically meaningful boundaries at that resolution.
3. Piecewise Smoothing: At all scales, intra-region smoothing should occur preferentially over inter-region smoothing.

From the results in this section it will be seen that all the above criteria are satisfied by the modified curvature consistency algorithm when applied to surface range data. No spurious details are introduced when the mask size is increased. The region boundaries are sharp and coincide with the boundaries in the original uncorrupted image. At all scales, the smoothing takes place within subregions separated by discontinuities but not across them.

The localization properties of the original and modified algorithms are demonstrated in Figures 6.10 and 6.11 respectively. The image displayed in these figures contains discontinuities of varying scales, with the scale becoming smaller with height. Four different neighbourhood sizes have been used in the reconstruction process. After

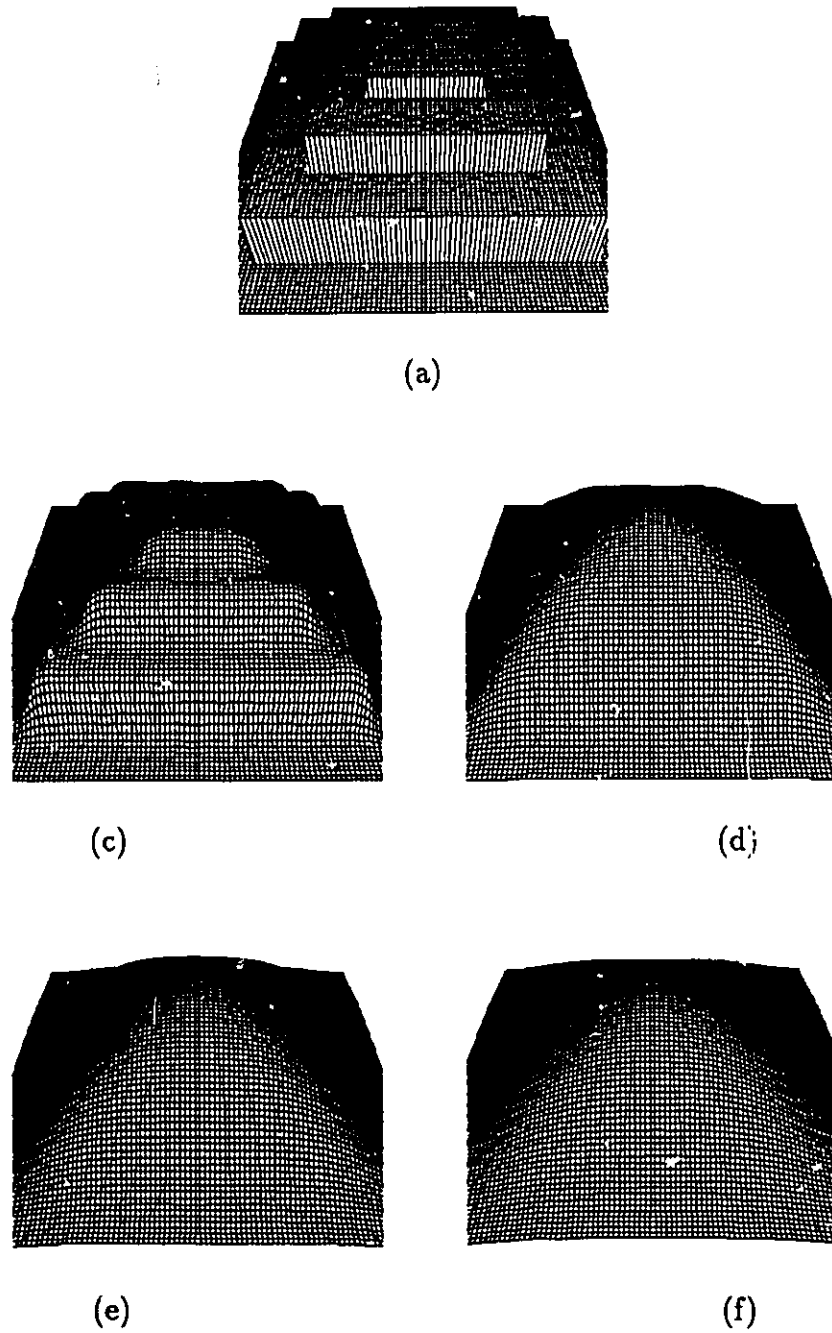


Figure 6.10: Scale space performance of the original curvature consistency algorithm on an image containing discontinuities of varying scales. (a) Original image, surface reconstruction using a (b) 3×3 mask, (c) 5×5 mask, (d) 7×7 mask, (e) 9×9 mask.




Figure 6.11: Scale space performance of the modified curvature consistency algorithm on an image containing discontinuities of varying scales. (a) Original image, surface reconstruction using using a (b) 3×3 mask, (c) 5×5 mask, (d) 7×7 mask, (e) 9×9 mask.

the same number of iterations with increasing mask size the distortion of discontinuities is progressively worse (Figure 6.10) when the original algorithm is applied. In Figure 6.11, the discontinuity localization property of the modified algorithm can be seen to remain stable over the different scales of discontinuities and using different mask sizes. For the same number of iterations, the step edge is localized with each mask size. The major difference in the use of the different mask size is the smoothness of the continuous data regions.

The mask size basically determines how fast information travels across the image during the relaxation process. With smaller masks, the rate of propagation is lower than with larger masks. Therefore, to achieve the same smoothness of the surface, a reconstruction process using a smaller mask would need more iterations than one using a larger size mask.

6.3.2 Noise Level Robustness

Due to the fact that the computations taking place in the original curvature consistency algorithm are local to a neighbourhood, it exhibits an adaptive behaviour over noise, i.e, if the noise properties differ across the image, the original curvature consistency algorithm adapts to the noise variance. In each neighbourhood, the curvature consistency algorithm provides a maximum-likelihood estimate of the interpolant surface patch, using information supplied by pixels only in that neighbourhood. The same adaptive behaviour over noise can be expected from the modified curvature consistency algorithm. However, it remains to be shown that the discontinuity localization property of the modified algorithm remains invariant over the different noise levels within the same image.

In Figure 6.12, the noise level robustness of the modified curvature consistency algorithm is demonstrated. The image contains 4 separate regions, each with different levels of Gaussian noise added to the range positions. Each step discontinuity separating the four regions are of the same height. The region in the middle has the highest amount of noise with standard deviation of 5. Figure 6.12(b) shows the reconstructed surface after several iterations. It can be clearly seen that the algorithm




Figure 6.12: Noise adaptation of the modified curvature consistency algorithm. (a) Original image with noise of standard deviation 1 (outermost), 2, 3 and 4 (innermost) added to the surface regions. (b) Surface reconstruction after 25 iterations using a 5×5 mask.

adapts to the changing noise levels in the image, and performs the same discontinuity localization regardless of the magnitude of noise.

It is also interesting to evaluate the noise-adaptive properties of the algorithm over scale space. Figures 6.13(a) shows an image with a much higher variation in the standard deviation of the noise. The middle region of the image has a SNR of 10. Figures 6.13(b) to (e) show the reconstructed surface after applying 50 iterations of differently sized masks. It can be seen that with each change in scale, the algorithm manages to localize the discontinuities to a different degrees. The 3×3 mask in Figure 6.13(b) has the most trouble in providing a smooth estimate of the surface in the noisiest region. The discontinuity localization works in most neighbourhoods, but in certain regions artifactual discontinuities are localized. In some other regions, the discontinuities are destroyed when they shouldn't be. Due to the high level of noise and a small neighbourhood support, by the time the weighting process starts having an effect, smoothing of the surface samples according to the curvature consistency algorithm has not taken place fully. Noisy samples start being localized as surface discontinuities after a few initial iterations. This problem is reduced somewhat in regions with lower noise levels, or with increasing neighbourhood mask sizes. With decreasing noise levels, the first few iterations achieve most of the smoothing, before

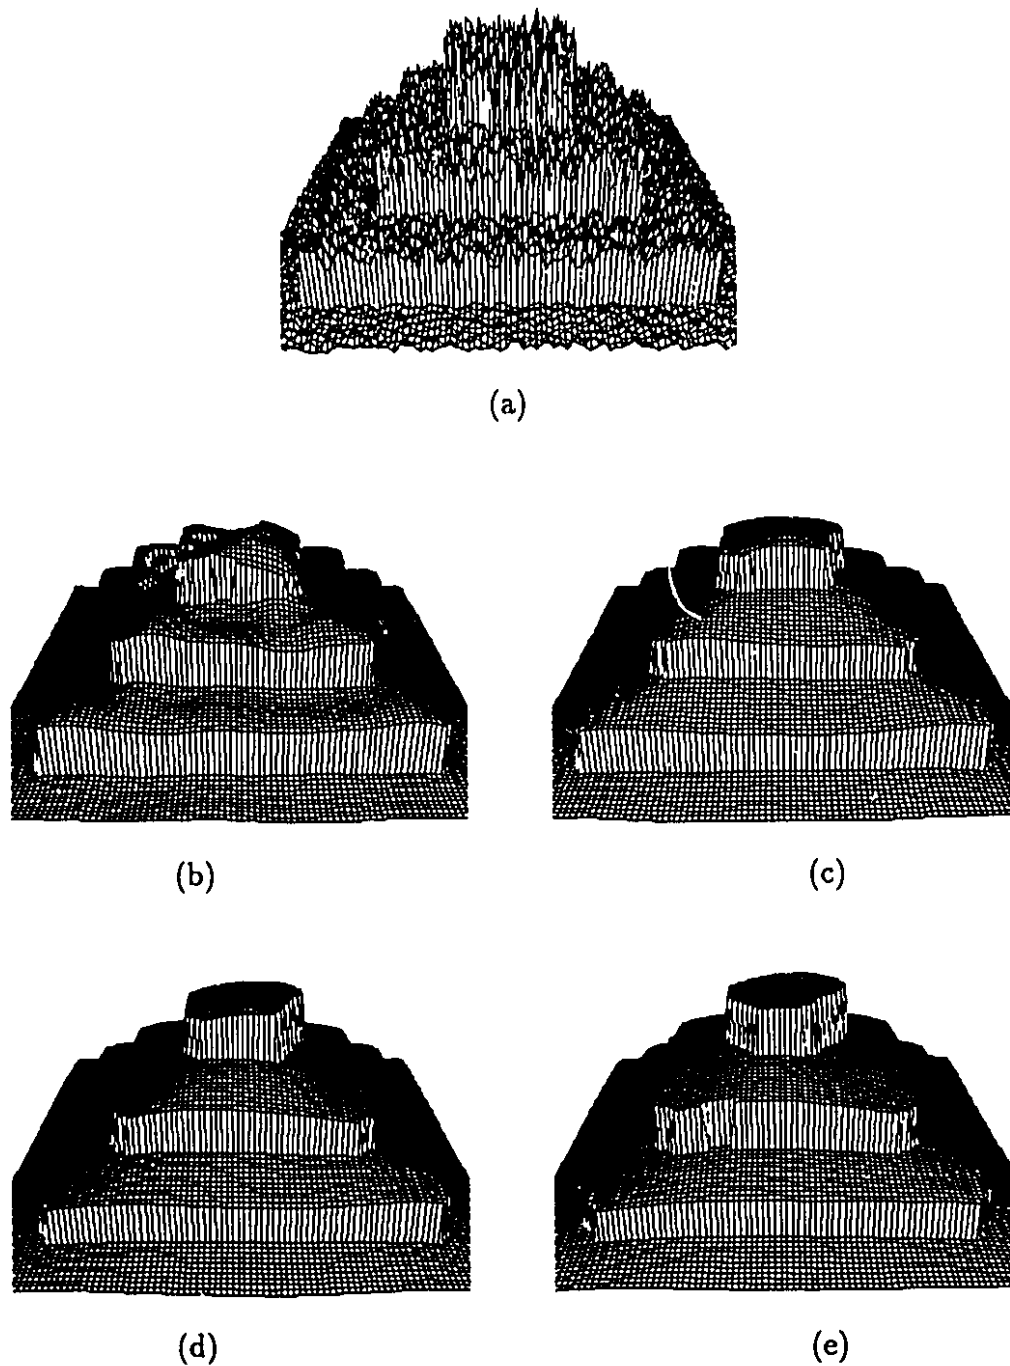


Figure 6.13: Noise space adaptation of the modified curvature consistency algorithm at multiple scales. Image contains regions with signal to noise ratios of 10 (outermost region), 3.33, 2, and 1 (middle region). (a) Original image, surface reconstruction using a (b) 3×3 mask, (c) 5×5 mask, (d) 7×7 mask, (e) 9×9 mask.

the weighting process takes over the discontinuity localization. With a larger mask size more samples play a role in the surface patch estimation hence, most of the smoothing has already taken place before the weighting process kicks in.

When the noise levels in the image are low, a mask of any size would perform equally well. However, when the SNR starts approaching 1, a larger mask size will obviously perform better than a smaller one. Without an estimate of the maximum noise standard deviation in the image, it is difficult to establish the mask size to be used. However, by trying to estimate the noise variance a priori defeats the purpose of making the algorithm completely automatic and adaptive. An alternative which could be the subject of further research is to take a multi-grid approach similar to the one proposed by Terzopoulos in [65]. He recognized that in relaxation methods constraint propagation is very slow between widely separated processing elements. By processing the surface at multiple scales, coarser representation of the surface can be used to constrain the finer ones and allow finer representations to constrain and improve the accuracy of coarser ones.

In this thesis a link is demonstrated between relaxation processes and multiple-measurement fusion problems. Optimal estimation theory is shown to apply to both areas. Using optimal estimation theory as a basis, a new method is proposed in the curvature consistency framework to estimate interpolant surface patch parameters from neighbouring pixels. The main difference between the estimation theoretic method of integrating neighbourhood information and the original algorithm's update equations is that in the former, an inverse-variance weighting is performed to give more emphasis to neighbours which provide better estimates of the interpolant surface patch parameters.

By including discontinuities as a cause of noise in the pixel-source, the discontinuity localization feature of the optimal estimator is justified. Since some correlation exists between pixel-sources in subregions partitioned by a discontinuity, the optimal estimator is adjusted to take this pixel grouping into account. The variance weighting is changed to a Gaussian weighting, with an automatically set γ parameter adaptively controlling the spread of the Gaussian. The smoothing of surface patch parameters is now performed in such a fashion that intra-region smoothing is encouraged, and smoothing across discontinuities (inter-region smoothing) is discouraged.

The scale-space robustness of the discontinuity localization method was demonstrated. It was shown to meet all the requirements enumerated by Perona and Malik in [49] to determine the robustness of filter behaviour over scale space. It was demonstrated that the algorithm adapts to changing noise properties within the *same* image. This was made possible mainly due to the automatic setting of the γ parameter.

As claimed at the beginning of the thesis, this algorithm is free of user set parameters. In fact the only parameter a user may need to set is the mask size. However it has been shown that if the noise levels are not perceptually very high (no detailed quantitative analysis of the effect of noise levels has been done) different mask

sizes give the same discontinuity localization and smoothing performance, the only difference being the number of iterations required to achieve the same result. To accommodate high levels of noise in the image, a multi-grid approach similar to the one proposed in [65] may be investigated as an extension to this research.

An attractive feature of this algorithm is that the (dis)continuity of the local surface region is learnt over the iterations. This information is stored in the error variance associated with each pixel-source in a given neighbourhood. To avoid computational complexity, the parameter τ was set to 1 in equation (5.11). This effectively disregards the history of the error variance. As an extension to this research, the effect of setting τ to a value higher than 1 may be investigated. Another extension would be the study of various weighting functions and their effects in weighting down older contributions to $\hat{\sigma}_i^2(k)$ in (5.11).

The mandate of the algorithm presented in this thesis has been *adaptation* to discontinuities, not their explicit detection. Since the weighting function acts on every parameter of the extended darboux frame $D(P)$, discontinuities in depth, orientation curvatures, and the principal directions should be localized. As an interesting extension of this thesis, it would be worthwhile to study the use of the error variances associated with each parameter in explicit discontinuity detection, and in region partitioning.

It is my hope that this research and specially the new estimation theoretic approach to looking at relaxation processes, contributes a new direction in the study of adaptive noise elimination techniques. It is my hope to derive a more general framework for the use of optimal estimation theory in relaxation processes and to provide a basis for a more general adaptive smoother which can be applied to any sort of data.

Bibliography

- [1] P. Besl and R. Jain. Segmentation through symbolic surface description. In *Proceedings IEEE Conf. Computer Vision and Pattern Recognition*, pages 77–85, Miami Beach, Florida, June 1986.
- [2] P. J. Besl and R. C. Jain. Segmentation through variable-order surface fitting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-10(2):167–192, Mar. 1988.
- [3] A. Blake. Comparison of the efficiency of deterministic and stochastic algorithms for visual reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-11(1):2–12, Jan. 1989.
- [4] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, Cambridge, Massachusetts, 1987.
- [5] R. Bolle and B. Vemuri. On three-dimensional surface reconstruction methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(1):1–13, Jan. 1991.
- [6] A. Brandt. Multi-level adaptive solutions to boundary-value problems. *Math. Comp.*, 31:333–390, 1977.
- [7] T. Broida and R. Chelappa. Kinematics and structure of a rigid object from a sequence of noisy images. In *Proc. IEEE Workshop on Motion: Representation and Analysis*, pages 95–100, Charelstone, S.C., May 1986. Computer Society of the IEEE, IEEE Computer Society Press.
- [8] R. T. Chin and C. Yeh. Quantitative evaluation of some edge-preserving noise smoothing techniques. *Computer Vision, Graphics, and Image Processing*, 23:67–91, 1983.

- [9] L. S. Davis and A. Rosenfeld. Noise cleaning by iterated local averaging. *IEEE Transactions on Systems, Man, and Cybernetics*, 8(9):705-710, Sept. 1978.
- [10] M. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1976.
- [11] T. Fan, G. Medioni, and R. Nevatia. Description of surfaces from range data using curvature properties. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 86-91, Miami Beach, Florida, June 1986. Computer Society of the IEEE, IEEE Computer Society Press.
- [12] F. P. Ferrie and J. Lagarde. On computing stable surface descriptions from range images. In *Proceedings 5th International Conference on Image Analysis*, Positano, Italy, September 20-22 1989.
- [13] F. P. Ferrie and J. Lagarde. Robust estimation of shape from shading. In *Proceedings 1989 Topical Meeting on Image Und. and Machine Vision*, pages 24-27, Cape Cod, Massachusetts, June 1989.
- [14] F. P. Ferrie and J. Lagarde. Curvature consistency improves local shading analysis. In *Proceedings 10th International Conference on Pattern Recognition*, pages 70-76, Atlantic City, New Jersey, June 1990.
- [15] F. P. Ferrie, J. Lagarde, and P. Whaite. Darboux frames, snakes, and superquadrics: Geometry from the bottom-up. In *Proceedings IEEE Workshop on Interpretation of 3D Scenes*, pages 170-176, Austin, Texas, Nov. 27-29 1989. Computer Society of the IEEE, IEEE Computer Society Press. IEEE Trans. PAMI - to appear.
- [16] F. P. Ferrie, J. Lagarde, and P. Whaite. Towards sensor-derived models of objects. In *Proceedings. Vision Interface '89*, London, Ontario, June 19-23 1989.

- [17] F. P. Ferrie, J. L. Lagarde, and P. Whaite. Recovery of volumetric descriptions from laser rangefinder images. In *Computer Vision – ECCV 90*, pages 387–396, Antibes, France, 23–27 Apr. 1990. INRIA, France, Springer-Verlag.
- [18] F. P. Ferrie and M. D. Levine. Deriving Coarse 3D Models of Objects. In *IEEE Comp. Soc. Conf. on Computer Vision and Pattern Recognition*, pages 345–353, University of Michigan, Ann Arbor, Michigan, June 1988.
- [19] F. P. Ferrie and M. D. Levine. Where and why local shading analysis works. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(2):198–205, Feb. 1989.
- [20] F. P. Ferrie, S. Mathur, and G. Soucy. Feature extraction for 3-d model building and object recognition. In A. Jain and P. Flynn, editors, *3D Object Recognition Systems*. Elsevier, Amsterdam, 1993.
- [21] B. R. Frieden. A new restoration algorithm for the preferential enhancement of edge gradients. *Journal of the Optical Society of America*, 66:280–283, 1976.
- [22] D. Geiger and F. Girosi. Parallel and deterministic algorithms for mrfs: surface reconstruction and integration. AI Memo 1114, MIT Artificial Intelligence Lab, Cambridge, MA, June 1989.
- [23] A. Gelb. *Applied Optimal Estimation*. MIT Press, Cambridge, Mass., 1974.
- [24] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741, 1984.
- [25] G. Godin and M. Levine. Structured edge map of curved objects in a range image. In *Proceedings IEEE Comp. Soc. Conf. on Computer Vision and Pattern Recognition*, San Diego, California, June 4-8 1989.
- [26] R. C. Gonzalez and P. Wintz. *Digital Image Processing*. Addison-Wesley, Reading, MA, 1977.

- [27] W. Grimson. *From Images to Surfaces*. MIT Press, Cambridge, MA., 1981.
- [28] W. E. L. Grimson. An implementation of a computational theory of visual surface interpolation. *Computer Vision, Graphics, and Image Processing*, 22:39-69, 1983.
- [29] J. Heel. Temporally integrated surface reconstruction. In *Proceedings, 3RD International Conference on Computer Vision*, pages 292-295, Osaka, Japan, Dec. 1990. Computer Society of the IEEE, IEEE Computer Society Press.
- [30] E. Hildreth. Edge detection. In Shapiro, S.C., editor, *Encyclopedia of A.I., Vol 1*. John Wiley & Sons Inc., 1987.
- [31] D. Hoffman and W. Richards. Parts of recognition. *Cognition*, 18:65-96, 1984.
- [32] B. Horn. *Robot Vision*. MIT Press, Cambridge, Massachusetts, 1986.
- [33] T. Huang, G. Y. Yang, and G. Y. Tang. A fast two-dimensional median filtering algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-27(1):13-18, 1979.
- [34] R. Hummel and S. Zucker. On the foundation of relaxation labeling processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-5(3):267-287, 1983.
- [35] V. K. Ingle and J. W. Woods. Multiple model recursive estimation of images. In *Proc. ICASSP '79*, pages 642-645, Washington D.C., Apr. 1979.
- [36] J. Koenderink. The structure of images. *Biological Cybernetics*, 50:363-370, 1984.
- [37] J. Lagarde. Constraints and their satisfaction in the recovery of local surface structure. Master's thesis, Dept. of E.E., McGill Univ., 1989.
- [38] J. S. Lee. Digital image enhancement and noise filtering by use of local statistics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-2(2):165-168, 1980.

- [39] A. Lev, S. W. Zucker, and A. Rosenfeld. Iterative enhancement of noisy images. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-7(6):435–442, June 1977.
- [40] M. D. Levine. *Vision in Man and Machine*. McGraw-Hill Book Co., New York, 1985.
- [41] S. Li. Reconstruction without discontinuities. In *Proceedings, 3RD International Conference on Computer Vision*, pages 709–712, Osaka, Japan, Dec. 1990. Computer Society of the IEEE, IEEE Computer Society Press.
- [42] G. A. Mastin. Adaptive filters for digital noise smoothing: An evaluation. *Computer Vision, Graphics, and Image Processing*, 31:103–121, 1985.
- [43] L. Mattheis, T. Kanade, and R. Szeliski. Kalman Filter-based Algorithms for Estimating Depth from Image Sequences. *International Journal of Computer Vision*, 3:209–236, 1989.
- [44] L. Matthies, R. Szeliski, and T. Kanade. Incremental estimation of dense depth maps from image sequences. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 366–374–000, Ann Arbor, Michigan, June 1988. Computer Society of the IEEE, IEEE Computer Society Press.
- [45] M. Nagao and T. Matsuyama. Edge preserving smoothing. *Computer Graphics and Image Processing*, 9:394–407, 1979.
- [46] P. Parent and S. Zucker. Curvature consistency and curve detection. *J. Opt. Soc. Amer., Ser. A*, 2(13), 1985.
- [47] A. Pentland. Fractal-based description of natural scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):661–674, 1984.
- [48] A. Pentland. Shading into texture. *Artificial Intelligence Journal*, 29:147–170, June 1986.

- [49] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-12(7):629–639, 1990.
- [50] T. Poggio, V. Torre, and C. Koch. Computational vision and regularization theory. *Nature*, 317(26):314–319, Sept. 1985.
- [51] W. K. Pratt. Median filtering. *Semiannual Report, Image Processing Institute, Univ. Southern California*, pages 116–123, Sept. 1975.
- [52] M. Rioux and L. Cournoyer. The NRCC three-dimensional image data files. National Research Council of Canada, CNRC No. 29077, June 1988.
- [53] A. Rosenfeld and A. C. Kak. *Digital Picture Processing*. Academic Press, New York, 1976.
- [54] P. Saint-Marc, J. S. Chen, and G. Medioni. Adaptive smoothing: A general tool for early vision. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 618–624, Ann Arbor, Michigan, June 1989. Computer Society of the IEEE, IEEE Computer Society Press.
- [55] P. Sander. *Inferring Surface Trace and Differential Structure from 3-D Images*. PhD thesis, Dept. Elect. Eng., McGill University, Montréal, Québec, Canada, 1988.
- [56] P. Sander and S. Zucker. Inferring differential structure from 3-d images: Smooth cross sections of fiber bundles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-12(9):833–854, 1990.
- [57] L. Schumaker. Fitting surfaces to scattered data. In Lorentz, G.G., Chui, C.K., Schumaker, L.L., editor, *Approximation Theory II*, pages 203–267. Academic Press, New York, 1976.
- [58] A. Singh. *Image-flow computation: An estimation theoretic framework, unification and integration*. PhD thesis, Dept. Computer Science, Columbia University, New York, 1990.

- [59] S. S. Sinha and B. G. Schunck. Discontinuity preserving surface reconstruction. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 229–234, Ann Arbor, Michigan, June 1989. Computer Society of the IEEE, IEEE Computer Society Press.
- [60] R. L. Stevenson and E. J. Delp. Viewpoint invariant recovery of visual surfaces from sparse data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(9):897–909, Sept. 1992.
- [61] R. L. Stevenson and E. J. Delp. Three-dimensional surface reconstruction: Theory and implementation. In A. Jain and P. Flynn, editors, *3D Object Recognition Systems*. Elsevier, Amsterdam, 1993.
- [62] R. Szeliski. Estimating motion from sparse range data without correspondence. In *Proceedings, 2ND International Conference on Computer Vision*, pages 207–215, Tampa, Florida, Dec. 1988. Computer Society of the IEEE, IEEE Computer Society Press.
- [63] D. Terzopoulos. Multilevel computational processes for visual surface reconstruction. *Computer Vision, Graphics, and Image Processing*, 24:52–96, 1983.
- [64] D. Terzopoulos. *Multiresolution Image Processing and Analysis*, chapter Multi-Level Representation of Visual Surfaces: Variational Principles and Finite Element Representations. Springer-Verlag, New York, 1983.
- [65] D. Terzopoulos. Image analysis using multigrid relaxation methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8:129–139, 1986.
- [66] D. Terzopoulos. Regularization of inverse visual problems involving discontinuities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(4):413–424, July 1986.
- [67] A. N. Tikhonov and V. A. Arsenin. *Solution of Ill Posed Problems*. Winston & Sons, Washington D.C., 1977.

- [68] S. Tomita, F. Tsuji. Extraction of multiple regions by smoothing in selected neighbourhoods. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-7:107-109, 1977.
- [69] J. W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, Reading, Mass., 1971.
- [70] D. C. C. Wang, A. H. Vagnucci, and C. C. Li. Gradient inverse weighted smoothing scheme and the evaluation of its performance. *Computer Graphics and Image Processing*, 15:167-181, 1981.
- [71] A. S. Willsky. *Digital Signal Processing and Estimation Theory*. MIT Press, Cambridge, Mass., 1979.
- [72] A. P. Witkin. Scale-space filtering. In *Int. Joint Conf. Artificial Intelligence*, pages 1019-1021, 1983.
- [73] J. W. Woods and C. Radewan. Kalman filtering in two dimensions. In *IEEE Trans. Inform. Theory.*, volume IT-23, pages 473-482, July 1977.