Reconstruction of high resolution 3D visual information

Marc Berthod, Hassan Shekarforoush, Michael Werman, Josiane Zerubia

INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Reconstruction of High Resolution 3D Visual Information

Marc BERTHOD, Hassan SHEKARFOROUSH,
Michael WERMAN, Josiane ZERUBIA

N° 2142
Novembre 1993

PROGRAMME 4

Robotique, image
et
vision

Reconstruction Of High Resolution 3D Visual Information

Marc Berthod[1], Hassan Shekarforoush[2], Michael Werman[3], Josiane Zerubia[4]

Programme 4 - Robotique, image et vision

Projet PASTIS

Research Report - Novembre 1993

Abstract: Given a set of low resolution camera images, it is possible to reconstruct high resolution luminance and depth information, specially if the relative displacements of the image frames are known. We have proposed iterative algorithms for recovering high resolution albedo and depth maps that require no a priori knowledge of the scene, and therefore do not depend on other methods, as regards boundary and initial conditions. The problem of surface reconstruction has been formulated as that of Expectation Maximization (EM) and has been tackled in a probabilistic framework using Markov Random Fields (MRF). As for the depth map, our method is directly recovering surface heights without refering to surface orientations, while increasing the resolution by camera jittering. Conventional statistical models have been coupled with geometrical techniques to construct general models of the world and the imaging process.

Key-words: super-resolution, image enhancement, shape from shading, 3D surface reconstruction, Markov Random Fields, regularization.

Acknowledgements : We would like to thank our collegue F. Ployette for preliminary implementions of the algorithm and also J. L. Lotti for providing the real aerial images for experimentation. Finally, we would like to thank l'Association Franco-Israëlienne pour la Recherche en Science et en Technologie (AFIRST) for providing a partial financial support to the project.

1 berthod@sophia.inria.fr*

2 hshekar@sophia.inria.fr*

3 werman@cs.huji.ac.il**

4 Zerubia@sophia.inria.fr*

* Unité de recherche INRIA Sophia Antipolis
2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex (France)
Téléphone : (33) 93 65 77 77- Télécopie : (33) 93 65 77 65

** Institute of Computer Science
The Hebrew University of Jerusalem
91904 Jerusalem, Israel

Reconstruction d'Information Visuelle 3D à Haute Résolution

Marc Berthod[1], Hassan Shekarforoush[2],
Michael Werman[3], Josiane Zerubia[4]

Programme 4 - Robotique, image et Vision

Projet PASTIS

Rapport de Recherche - Novembre 1993

Résumé : Etant donné un ensemble d'images à basse résolution, il est possible de reconstruire l'information lumineuse et la carte de profondeur, si le recalage relatif entre les images est connu. Nous avons proposé des algorithms itératifs pour récupérer la carte d'albedo aussi bien que la carte de profondeur à une haute résolution. Ces algorithmes n'exigent pas la connaissance a priori de la scène, et par conséquent sont indépendants des résultats qui seraient obtenus par d'autres méthodes telles que stéréovision et analyse de mouvement pour les conditions au bord et les conditions initiales. Le problème de reconstruction de surface est formalisé comme un problème de Maximisation de probabilité A Posteriori (MAP) et il est traité dans un cadre probabiliste en utilisant les Champs de Markov. Nous avons résolu le problème de profondeur directement sans faire référence à l'orientation de surface. En même temp, nous avons augmenté la résolution grâce à la multiplicité des vues. L'utilisation conjointe des modèles statistiques et des techniques géométriques nous a permis de construire un modèle général du monde extérieur.

Mots-clé : super-résolution, amélioration d'image, shape from shading
reconstruction de surface 3D, Champs de Markov, régularision.

Remerciements : Nous tenons à remercier F. Ployette pour son aide dans la mise en oeuvre des premiers algorithmes, ainsi que J. L. Lotti pour la fourniture des images aériennes. Enfin, nous remercions l'Association Franco-Israëlienne pour la Recherche en Science et en Technologie (AFIRST) pour le soutien financier partiel apporté à ce projet de recherche.

1 berthod@sophia.inria.fr*

2 hshekar@sophia.inria.fr*

3 werman@cs.huji.ac.il**

4 Zerubia@sophia.inria.fr*

* Unité de recherche INRIA Sophia Antipolis
2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex (France)
Téléphone : (33) 93 65 77 77 - Télécopie : (33) 93 65 77 65

** Institute of Computer Science
The Hebrew University of Jerusalem
91904 Jerusalem, Israel

Contents

1 Introduction

Image resolution depends on the physical characteristics of the sensor: the optics, the density and the spatial response of the sensing elements. Increasing the resolution by sensor modification may not always be an available option. An increase in the sampling rate can, however, be achieved by obtaining more samples of the same scene from a sequence of displaced images. Therefore, the most important application of computing super-resolution images is when the available sensor resolution is inadequate and when it is more economical to capture multiple images rather than purchase a sensor of higher resolution, if such sensor exists at all. We, therefore, intend to develop theoretical and practical tools that would make it possible to reconstruct images at a resolution higher than that of any existing sensor.

In this report we propose an algorithm for reconstructing high resolution images given a set of low resolution observations. Markov Random Field's (MRF)[1][2][3] have been used for modeling the properties of the processed images. Use of MRF which became popular with the ingenious paper of Geman & Geman[10] is indeed equivalent to regularization techniques used by Marroquin, Grimson, Terzopoulos and many others [21][25][33]. Therefore, in what follows MRF's will be directly interpreted in terms of regularization and the corresponding algorithms will be standard optimization ones such as the gradient descent and the conjugate gradient. In the case of discrete data, the use of MRF's gives rise to combinatorial optimization algorithms. We propose here to work on some fundamental applications of optimization in the field of image reconstruction. There is, for example, an extensive literature on extracting 3D information, which is based mainly on stereopsis and motion analysis [30]. Here we intend to investigate the possibility of direct retrieval of 3D high resolution visual information from a sequence of displaced low resolution intensity images.

Early research on super-resolution was carried out by Tsai and Huang [15], who used frequency domain methods, disregarding the blurring effect of the imaging process. Later Gross[11] assumed an exact knowledge of the imaging process and the relative shifts between the input pictures. He then constructed a single blurred picture of higher spatial sampling rate by merging the low resolution images over a finer grid using interpolation. The merged picture was then deblurred by means of a restoration filter using standard pseudo-inverse techniques applied to the blurring operator. As in the work of Tsai and Huang, only interframe translations were considered.

Peleg and Keren[20][28] estimated an initial guess of the higher resolution image, and simulated the imaging process to obtain a set of estimated low resolution camera images. Using an error criterion to minimize the difference between the estimated and actual sensor images they devised an iterative algorithm to reconstruct the higher resolution image. The method however suffered from lack of speed and sensitivity to noise. Irani and Peleg [18][19] describe a method based on the

resemblance of the presented problem to the reconstruction of a 2D object from its 1D projections in Computer Aided Tomography (CAT). The high resolution image is constructed using an approach similar to the back-projection method used in CAT.

A better model introduced recently by Zevin and Werman [36] is taking into account the 3D features of the world with cameras in general positions and the projection considered as being perspective. The model has quite good geometric aspects although different reflectance properties have been ignored.

Here we will assume a quite general reflectance model and hence the formulation shares common points with that of height from shading [14] [23]. We specifically intend to use Expectation Maximization (EM) [31]. Initially we will concentrate on designing robust registration algorithms. Future extensions will include other aspects such as robustness to noise, discontinuities, occlusions and other reflectance laws. We will also look at different relaxation algorithms. Tests have been carried out on both synthetic and real data (aerial images). The results of the project are expected to have an important impact in the area of satellite and aerial image processing and provide an intermediate processing level for applications such as feature extraction.

2 Image Formation

We will now consider a model for computing the irradiance value at any image point. Rigorous analysis of this problem was first carried out by Horn [12][13] in early 70's. Horn introduced optical models of image acquisition in the area of computational vision which were based on projecting surface points onto the image plane.

Our primary assumption here is that the focal point is sufficiently far from the scene, and hence all rays from a reasonably extended surface to the image plane are parallel in which case the projection onto the image plane can be regarded as being orthographic. In standard orthographic projection this will also lead to interchangeable coordinate frames, and hence the use of separate reference frames will be somewhat redundant.

Two decisions have to be made: one regards the reflection model and the other the sensor model. As for the sensors, we will assume that they are pin-hole cameras located at infinity and hence the appropriate geometric projection to be used will be an orthographic one. Here our aim is to take into account radiometric and structural characteristics of the scene that are explicitly related to features captured in an image. We will also integrate this method with conventional statistical image modeling techniques. The basic premise underlying this approach is the assumption that brightness variations depend, in a deterministic way, on physical properties of the scene such as surface albedo, orientation and distance. However, one can also assume that invariant pictorial features (eg. regions of homogeneous brightness) correspond to semantically meaningful objects or surfaces of objects.

When an image is formed, the light intensity recorded at an image point depends on three main factors: the incident illumination, the surface reflectance and the surface orientation. We will assume that the surface is an ideally diffusing one and hence the registered luminance energy can be determined by Lambert's law [35]:

$$G = Ig\cos i \tag{2.1}$$

where I is the incident illumination, g is the surface albedo and i is the incident angle. It can be seen from this simple formula that photometrically an infinite number of illumination, reflectance and orientation combinations can lead to the same image light intensity. Therefore, recovering scene characteristics and structures is an underconstrained problem and requires additional constraints for solution. The most usual constraint imposed here is the continuity. This means that brightness variations are usually smooth and depend on physical phenomena of the outside world. The exceptions to this are the boundary discontinuities, which can cause brightness edges. If we lay down a continuity constraint then the illumination and the reflectance will be both uniform so that the image brightness will be directly proportional to the change in orientation [35]:

$$\frac{dG}{G} \propto \frac{d(\cos i)}{\cos i} \tag{2.2}$$

The problem is now sufficiently constrained and could lead to a solution.

2.1 Computing the Intensity

As mentioned earlier three major factors are contributing to the amount of light intensity received by the camera at any image point. These factors are determined by the intrinsic properties of the surface and its geometric structure. The reflected light intensity, however, can be degraded prior to formation of the final image. This can be due to different reasons such as spatial blurring, random noise, chromatic or temporal effects, aberrations of the optical system and relative motion of the camera. At this stage we will assume that our intensity images are only degraded by a Gaussian blurring operator [31]:

$$H(u,v) = \frac{1}{2\pi\sigma^2} exp\left(-\frac{u^2+v^2}{2\sigma_\varepsilon^2}\right) \tag{2.3}$$

where:

$$u = a_{11}X + a_{12}Y + a_{13}Z + a_{14} \tag{2.4}$$

$$v = a_{21}X + a_{22}Y + a_{23}Z + a_{24} \tag{2.5}$$

(X,Y,Z) are the world coordinates, a_{ij} 's the camera parameters and σ_ε is the standard deviation of the kernel.

Pixel intensities are, therefore, given by:

$$I_{(k,l)} = \sum_{(i_1,i_2) \in w} H((u-k),(v-l))\, G_{(i_1,i_2)} \tag{2.6}$$

where k and l are the shifts along the image coordinate axes, $G_{(i_1,i_2)}$ is the energy received by the sensor at the corresponding pixel and w is the contributing square surface patch.

The advantages of using a Gaussian operator are its symmetry and its causality. We can, therefore, expect isotropic shift invariant degradations, when the noise is negligible. Hence the final intensity image is obtained by convolving the reflected flux with the so-called Point Spread Function (PSF) over a surface patch whose contribution to the pixel intensity is non-trivial. The usual approach is to perform the convolution process in the continuous domain. Direct discretization of the surface proposed here implies that the energy reflected in the direction of the sensor is constant in the elementary patch corresponding to each surface pixel.

During the process of image formation the centre of the contributing patch or the so called receptive field, moves along the axes of the camera with an amount proportional to the sampling rate along these axes (ie. S_u and S_v) and to the rotation angle of the image plane. Assuming that the camera motion is composed of translations and rotations only, we will obtain the following camera parameters:

$$a_{11} = a_{22} = S_u \cos\theta \tag{2.7}$$

$$a_{12} = a_{21} = S_v \sin\theta \tag{2.8}$$

where θ is the rotation angle. When dealing with square images with equal sampling along the two image axis:

$$S_u = S_v = \frac{h}{H} \tag{2.9}$$

where h and H are the low resolution and the super-resolution image widths, respectively. As for a_{13} and a_{23}, we have set these to zero while a_{14} and a_{24} are given by the translation of the camera along the world axes.

An important concern here is the granularity of the discretization. In order to be able to recover a sufficiently fine image resolution, the area of the receptive field has to be kept small enough so that the discretized kernel would be a smooth approximation of the ideal Gaussian operator. The area of the receptive field depends on the orientation of its normal and its maximum vertical extension is inversely proportional to the cosine of the angle of this normal with the vertical, ie. on the local gradient of the surface. Applying the continuity constraint would always guaranty a

projected patch smaller than the receptive field and hence the projected patch can be made as small as possible by refining the resolution of the surface.

2.2 Surface Reflection

Various reflectance models have been used in the area of artificial vision systems. There are two main approaches in modeling the reflection of incident light from a surface [35]: wave optics (based on electromagnetic wave theory) using Maxwell's equations to study the behaviour of light; geometrical optics which is used to explain the macro-behaviour of light assuming that the wavelength is much larger than the micro-structures of the surfaces involved.

Here we will only consider the latter approach and for the time being we will only take into account the case where we have only diffuse reflection (Lambertian). Other cases such as specular reflection and interreflections will not be considered at this stage.

We will assume that, after discretization, our surface is made up of elementary units of constant direction. The smoothness of the surface is kept as high as possible by refining the surface resolution to the required level. The percentage flux reflected from a natural surface depends on the micro-structures of the surface and the distribution of the illuminating light. Here after discretization, only the latter plays an important role. Now let us construct the tangent planes to the surface elements. We can see that the distribution of the incident light is taking place over a hemisphere. Therefore, in the most general case one can compute the reflected flux from each direction and then superimpose the results. Here we will assume a single light source sufficiently far from the surface and hence a fixed direction of illumination. In this simple case the geometry of the system comprises of three angles: the incident, the emittance and the phase angles (see Fig 2-1).

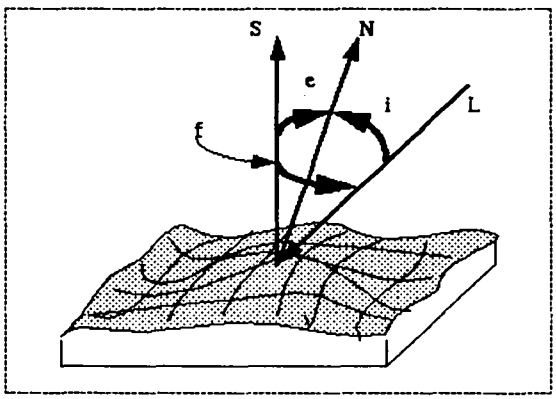

Fig 2-1 A geometric view of the imaging process. N: local normal, S: viewing direction, L: illumination vector, i:incident angle, e: emittance angle, f : phase angle.

These, as shown in Fig 2-1, are the angles between the illumination vector and the local normal, the emittance direction and the local normal and the illumination vector and the emittance vector, respecively. The fraction of illuminating light reflected per unit surface area and per unit solid angle in the direction of the viewer is refered to as the reflectivity function. and is denoted by $R(i,e,f)$. For an illumination of I (flux/area) and a surface luminance of B (flux/steradian/projected area) in the direction of the viewer the reflectivity is given by B/I. Here the projected area is inversely proportional to the cosine of the emittance angle and refers to the foreshortened area as seen from the viewing position.

The basic assumption underlying our approach henceforth is the image irradiance equation:

$$I_{\hat{i}}(u,v) = gR(\hat{n}) \tag{2.10}$$

where $\hat{n}$ is the unit normal vector to the surface element $\hat{i}$ that projects from the receptive field to the image point (u,v) . We have used hereafter the vector notation $\hat{i}$ or $\hat{j}$ for surface elements with coordinates (i_1,i_2) and (j_1,j_2) respectively and also $\hat{k}$ for the image element (pixel) (k_1,k_2) . In our model the actual image intensity is the superposition of the reflected flux from each surface element of the receptive field weighted by the Gaussian kernel, ie. a discrete convolution by the Gaussian kernel (see equation (2.6) above). Note that R in practice contains many other variables which may or may not be known. An important assumption here is that the image irradiance is only a function of the surface orientation and not the surface position [13]. Suppose the surface is specified in the world coordinate frame by $z = f(x,y)$. One can easily see that the vectors $\left(-1,0,-\frac{\partial z}{\partial x}\right)$ and $\left(0,-1,-\frac{\partial z}{\partial y}\right)$ are tangent to the surface at a given point. The vector product yields:

$$\left(-1,0,-\frac{\partial z}{\partial x}\right)\Lambda\left(0,-1,-\frac{\partial z}{\partial y}\right) = \left(\frac{\partial z}{\partial x},\frac{\partial z}{\partial y},-1\right) \tag{2.11}$$

which is the the normal to the tangent plane and hence to the surface. Let $p = \frac{\partial z}{\partial x}$ and $q = \frac{\partial z}{\partial y}$.

The ordered pair (p,q) are in fact the corresponding gradient space coordinates. Hence, the gradient space can be found by directly refering to the surface orientation. We have, therefore, a mapping from the surface orientation to the

gradient space by means of the normal $(p, q, -1)$, where (p, q) is a point in the gradient space.

In general we can analyse our reconstruction problem more easily in the gradient space. It can be easily seen, for example, that parallel planes would map into a single point in the gradient space. We have, therefore, a translational ambiguity in our recovery of the absolute visual depth. All planes perpendicular to the viewing vector are mapped at the origin of the gradient space. Hence, the distance from the origin equals to the tangent of the emittance angle :

$$\sqrt{p^2 + q^2} = \tan(e) \tag{2.12}$$

In the case of a single light source near the viewer we have $e = i$ and therefore:

$$\cos(i) = \frac{1}{\sqrt{1 + tan^2(e)}} = \frac{1}{\sqrt{1 + p^2 + q^2}} \tag{2.13}$$

Alternatively, using the scalar product of the unit normal and the unit viewing vector:

$$\cos(i) = \frac{(p, q, -1)}{|(p, q, -1)|} \cdot \frac{(0, 0, -1)}{|(0, 0, -1)|} \tag{2.14}$$

Now let us consider the general case where the light source is not near the viewer. Here the reflectance is a function of the angles of incident, emittance and phase. It is common practice to work with cosines of these angles, since these can be easily calculated using the scalar products of the unit vectors. For a single distant light source given by the unit vector $\vec{L}(\alpha_L, \beta_L, \gamma_L)$ we will have:

$$\cos(i) = \frac{\vec{N} \cdot \vec{L}}{\sqrt{1 + p^2 + q^2}} \tag{2.15}$$

where $\vec{N}$ is the surface normal. Assuming a uniform illumination:

$$G \propto \frac{\vec{N} \cdot \vec{L}}{\sqrt{1 + p^2 + q^2}} \tag{2.16}$$

Since in the present case, digitization occurs on the surface rather than in the image plane, the number of elementary patches in the receptive field will be inversely proportional to the cosine of the reflection angle and hence:

$$G \propto \vec{N} \cdot \vec{S} \tag{2.17}$$

$\hat{S} = (\alpha_s, \beta_s, \gamma_s)$ is the unit vector along the viewing direction.

Now Lambert's law will directly lead to the following equation for the energy received from an elementary patch on the receptive field:

$$G = g\frac{\vec{N}\cdot\vec{L}\times\vec{N}\cdot\vec{S}}{\sqrt{1+p^2+q^2}} = gR \tag{2.18}$$

When the incident angle is constant, for a distant light source and an orthogonal projection, we obtain a second order polynomial in terms of p and q [35]. The loci of all directions that have the same incident angle are conic cross sections as shown in Fig 2-2.

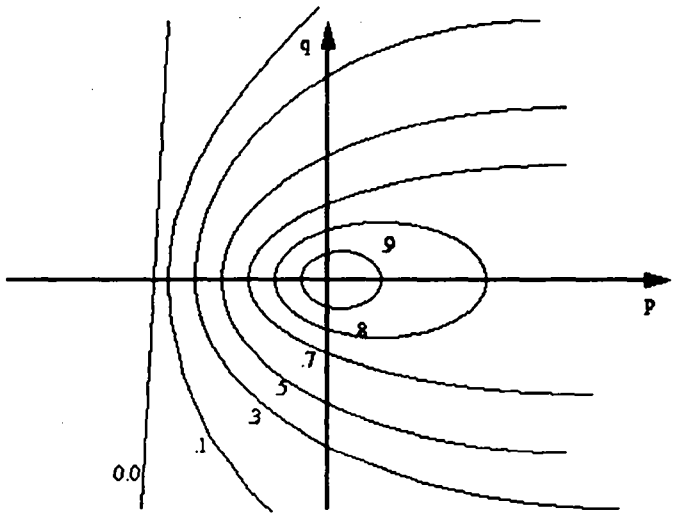

Fig 2-2 Loci of constant illuminating direction contours (isophotes)

Solving the polynomial equation for p and q will lead to two sets of solutions and hence gives rise to the gradient ambiguity for a given intensity value. Our aim here will be to avoid this underconstrained problem by solving directly for the height while increasing our resolution by camera jittering [17]. The common approach is to impose smoothness and integrability constraints [8][14] to overcome the multiplicity of solutions.

2.3 Uniqueness Of The Image Irradiance

The major difficulty in our problem of reconstructing 3D information from a sequence of intensity images is the multiplicity of the solution to the image irradiance

equation. We have already mentioned some ambiguities that could arise due to photometric or geometric features of the surface. Bruss [6] found some results on this issue by assuming that the light source and the boundary information (usually in terms of occluding contours) are known. She assumed that the image irradiance equation is eikonal:

$$I_i(u, v) = p^2 + q^2 \tag{2.19}$$

She then demonstrated that if boundary conditions are specified and (2.19) has at most one stationary point then the solution is unique up to inversion. The special case of a lambertian surface with the light source near the viewer leads to the above results. Blake et al. [5] present a proof for the uniqueness of the solution in the case when we have a general viewpoint, by assuming that the normal vectors on the boundaries are known. In a more recent attempt, Oliensis [27] showed that a single point on the surface whose normal is parallel to the light source would be sufficient to reconstruct the surface from the image. Leclerc et al. [23] argue that discretization of the problem adds more complication by making some continuous analyses inapplicable. For example the method of computing the discrete derivatives can affect the expected values of the image intensity.

Multiplicity of the solution mainly arises due to the fact the problem is underconstrained or ill-conditioned. In our solution we will impose a smoothness constraint which will be discussed further in the next section. Note that in our model we have used a smoothly varying albedo which brings in additional complexity as against the case where the albedo is assumed to be constant over the entire surface. The light source has, however, been kept constant and is assumed to be known throughout this paper. We have also reduced our solution space by directly working on the height. In this sense our work shares common points with the method proposed by Leclerc et al. [23] although they work on the basis of a constant surface albedo.

3 An Algorithm For Super-resolution

We will now consider the problem of refining photometric and structural information extracted from a sequence of intensity images. We will assume that the imaging process is well known and that a sequence of low resolution observations are available with pixels on each frame registered at subpixel accuracy. This requires a subpixel displacement between picture frames. Initializing the super-resolution image to an arbitrary estimate, we will then simulate the imaging process to obtain a set of estimated low resolution pictures. Comparing these with the observed sequence of low resolution frames would allow us to minimize a penalty function iteratively and hence update the initial guess until a stop criterion is met (see Fig 3-1).

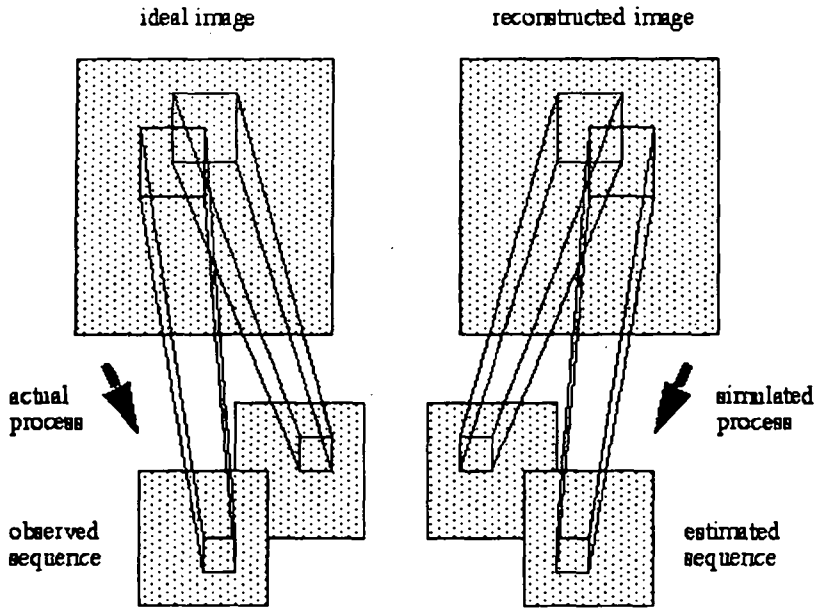

Fig 3-1 Schematic diagram of the iterative super-resolution algorithm.

This approach is sometimes refered to as camera jittering [17]. Another possibility would be using several sensors simultaneously. We will assume that the reflection model is known everywhere on the surface. We will also assume that both intrinsic and extrinsic parameters of the camera as well as the PSF are known. Now, we have to specify the nature of the initialization process, the penalty function and the updating scheme, although these will be discussed in more details in the following sections. In most reconstruction schemes the initialization depends directly on the optimization method used and on the order of the penalty function. Some optimization algorithms, for example, are known to be sensitive to the way the initial conditions have been chosen. The proposed method here, as we will see later, is based on a deterministic method which does not require a good initial estimate for converging to a solution, provided that the penalty functional is convex. Non-convexity could be overcome by several well known algorithms, such as the Simulated Annealing (SA) [10] , Mean Field Annealing (MFA) [9] and the Graduated Non-Convexity (GNC) [4], at the expense of high computational time. To update the super-resolution image on each iteration we will use a correction vector (see equation (4.1)) which will be computed using a gradient based optimization algorithm.

The criterion to be minimized is of standard form: it contains two terms, one for the estimation error and the other a constraint such as smoothness to ensure a stable convergence to a unique solution. The system of equations is usually providing a means of controlling the trade off between these two terms. If the brightness error term is set aside then the solution degenerates to pure interpolation. However in practice, in order to be able to regularize the solution, some terms reflecting the

constraints are coupled with the estimation error. The degree of regularization required here can depend on the actual image being processed. Some authors[4][23] advocate empirical tuning of the weighting parameters at each iteration. As for the updating scheme, we will use the gradient of the penalty functionals to specify the direction and the amount by which each pixel value has to be modified in order to minimize the penalty.

4 Regularization Of The Solution

We face the problem of solving a set of non-linear first order differential equations. The approach adopted here is to compute a correction vector C on each iteration and subtract it from the current estimate of the solution S to obtain a new set of values in S [24]:

$$S^{i+1} = S^i - C \tag{4.1}$$

On each iteration we can then compute a vector of error measurements between the reconstructed image and the observed data. The aim now is to solve for eliminating this error. Assuming local linearity we have:

$$JC = \varepsilon \tag{4.2}$$

where J is the Jacobian matrix.

Minimizing the quadratic norm of the residual is in fact equivalent to solving (4.2). It should be, however, noted that this scheme is based on the assumption that the original non-linear equation is locally linear given a range of typical errors. Therefore, the above method can be potentially unstable and lead to ill-conditioned solutions. This problem can be solved by introducing *a priori* knowledge in the form of constraints, which will provide a stable convergence towards the solution.

A problem that arises here is the lack of trade-off between the criterion of closeness to data and that of smoothness. The most common approach is to introduce weighting factors on each row of the above equation so that the components on the right hand side have the same standard deviation. This means that the contribution of each constraint will be proportional to the standard deviation from its expected value. To scale each row of the a priori model we will introduce a control parameter together with a diagonal matrix in which the scaling components are inversely proportional to the standard deviation [24] (ie.):

$$\begin{bmatrix} J \\ \lambda\omega \end{bmatrix} C = \begin{bmatrix} \varepsilon \\ \lambda\omega d \end{bmatrix} \tag{4.3}$$

Solving this system is equivalent to minimizing:

$$\|JC - \varepsilon\|^2 + \lambda^2 \|\omega (C - d)\|^2 \tag{4.4}$$

which is the standard regularization using Tikhonov-Arsenin [34] stabilizing functional. Note that, as the parameter λ increases, the method corresponds to the gadient descent with decreasing step sizes, whereas for decreasing λ it corresponds more to Newton's method with fast quadratic convergence near the solution but the possibility of divergence when initializing too far away.

In what follows we will propose a stochastic method for regularizing our problem. The advantage of this approach is that the *a priori* knowledge will not restrict the solution space and instead, it will be represented in terms of a probability distribution and hence the class of solutions can be much larger than in the standard regularization [33]. By its very nature it will also be easier in this approach to incorporate discontinuity information directly. The choice between this methodology and standard regularization is due to flexibility of the former in handling the varying conditions of the shape of the surface, the albedo and the degree of super-resolution. However, this flexibility is obtained at the expense of additional computational complexity.

4.1 A Probabilistic Framework

At this stage we will sample both g and z (ie. the albedo and the altitude respectively) on a super-resolution grid size and impose a probabilistic model on them. We will, therefore, try to solve a classical Expectation Maximization (EM)[31] problem, directly in a finite dimensional space, but possibly with real values. We will use the nonlinear technique of Maximum A Posteriori (MAP) [1], which will enable us to take into account the nonlinearities associated with the image acquisition process. This will also allow us to assume a nonstationary mean for the image random field.

The class of stochastic models that meet our criteria is the Markov Random Fields (MRF) on finite lattices [2]. We will assume that the random fields associated with g and z are independent. Let (g, z) depict the pair of random fields to which our observed images belong. Also let (g*, z*) represent the ideal superresolution images. We need to construct the a posteriori density function p(g*, z*) given (g, z) , that is to construct the conditional density function $p((g^*,z^*)/(g, z))$, where $I_1 \ldots I_n$ are the observations using n sensors. Let $\hat{I}_1 \ldots \hat{I}_n$ be the values of expected I_n at which the above conditional density takes its largest value. This means that $\hat{I}_1 \ldots \hat{I}_n$ are representing the most probable values of the observation fields. As we saw earlier the technique used here amounts to regularizing a cost functional. A further step now is to ensure that the MAP estimate will minimize the cost function.

We can now see that the determination of the MAP estimate requires the knowledge of constructing the *a posteriori* density function. Now, let us see how this *a posteriori* density function can be determined.

Using Bayes law gives:

$$p((g,z)/I_1...I_n) = kp(I_1...I_n/(g,z))p(g,z) \tag{4.5}$$

where $k = 1/p(I_1...I_n)$ which is assumed to be a constant independent of (g,z).

Assuming that I_n's are independent ie. the noise process is independent from one view to another:

$$p(I_1...I_n/(g,z)) = \prod_n p(I_n/(g,z)) \tag{4.6}$$

For the sake of simplicity we will consider the case for g only and extend the method to z. MRF's have the property that their probability distribution can be expressed in the form of a Gibbs distribution [26]. Therefore:

$$p(g) = \frac{1}{Z}exp\left(-\frac{1}{2}U_g^2 C_{gg}^{-1}T^{-1}\right) \tag{4.7}$$

where $C_{gg} = E\{(g-\mu_g)^2\}$ is the covariance of g, $\mu_g = E\{g\}$ is the mean of g and T is a parameter set to 1 here.

Z called the partition function is a normalizing constant [10] and U_g is a function of local potentials of the field given by:

$$U_g = \sqrt{\sum_{j \in \vartheta_i} (g_i - g_j)^2} \tag{4.8}$$

ϑ_i is a set of pixels defined by the neighbourhood structure of the random field. Note that if we allow μ_g to vary randomly (ie. associate different mean values to each pixel) then the random field will be automatically nonhomogeneous, which is usually the case with real pictures. At this stage we will however make an assumption to simplify the matter. We know that different regions in an image can have different average albedos. However, the variations from these averages and the correlation properties of these variations can be assumed to have the same order of magnitude. This is equivalent to saying that the covariance matrix of the registered image is of Toeplitz form where the elements on the leading diagonal are all equal and are given by the variance of the albedo.

Now, assuming that the noise process is a white gaussian then the density of the observation error is given by the following distribution:

$$p(\varepsilon) = \left(\frac{1}{(2\pi)C_{\varepsilon\varepsilon}}\right)^{1/2} exp\left(-\frac{1}{2}\varepsilon^2 C_{\varepsilon\varepsilon}^{-1}\right) \tag{4.9}$$

where $C_{\varepsilon\varepsilon}$ is the covariance of ε : $C_{\varepsilon\varepsilon} = E\{\varepsilon^2\}$.

Since the conditional density depends on the structure of the noise we can compute its distribution function as follows:

$$p(I_n/g) = \left(\frac{1}{(2\pi)C_{\varepsilon\varepsilon}}\right)^{1/2} exp\left(-\frac{1}{2}F^2 C_{\varepsilon\varepsilon}\right) \tag{4.10}$$

F is a function of brightness errors given by $F = \sqrt{\sum_{\hat{k}\in w}\left(I_{n,\hat{k}} - \hat{I}_{n,\hat{k}}\right)^2}$, where $\hat{I}_n$ is an estimate of the intensity given by (2.6). The size of the window w depends on the size of the PSF. Substituting (4.7) and (4.10) in (4.5) and taking logarithms of both sides, we get:

$$lnp(g/I_n) = lnk - \frac{1}{2}U_g^2 C_{gg}^{-1} - lnZ - \frac{1}{2}F^2 C_{\varepsilon\varepsilon}^{-1} - \frac{1}{2}ln(2\pi)C_{\varepsilon\varepsilon}^{-1} \tag{4.11}$$

Knowing that the logarithm is a monotonically increasing function, the same values of g would maximize both $p(g/I_n)$ and $lnp(g/I_n)$. Therefore, we can look for the MAP, simply by looking for the maximum of the right hand side of (4.11), which is equivalent to minimizing the following energy function:

$$\Psi(g) = \frac{1}{2}U_g^2 C_{gg}^{-1} + \frac{1}{2}F^2 C_{\varepsilon\varepsilon}^{-1} \tag{4.12}$$

For the general case where z is also taken into account:

$$\Psi(g,z) = U_g^2 C_{gg}^{-1} + U_z^2 C_{zz}^{-1} + F^2 C_{\varepsilon\varepsilon}^{-1} \tag{4.13}$$

The minimum of this function is obtained when its partial derivatives with respect to g and z vanish and hence: $\nabla\Psi_g = 0$ and $\nabla\Psi_z = 0$.

Assuming that the covariance matrices (ie. $N \times N$ matrices whose elements are equal to the covariance of corresponding pixel albedo or altitude values) are of Toeplitz form then the MRF's associated with g and z will be homogeneous. A similar assumption about the noise process will imply that the PSF is shift invariant (isotropic). This will lead us to minimizing the following energy function defined at pixel level:

$$E = \sum_{j\in\vartheta_i}\frac{(g_i - g_j)^2}{2\sigma_g^2} + \sum_{j\in\vartheta_i}\frac{(z_i - z_j)^2}{2\sigma_z^2} + \sum_{\hat{k}\in w}\frac{\left(I_{n,\hat{k}} - \hat{I}_{n,\hat{k}}\right)^2}{2\sigma_\varepsilon^2} \tag{4.14}$$

where the σ's are the standard deviations corresponding to each term in the above equation.

Let us now examine each term in the above equation more carefully. The first two terms are measuring how severely g and z are distorted. This is equivalent to the elastic energy of a membrane model when the neighbourhood structure of the field does not exceed that of the second order MRF. The last term, however is a measure of faithfulness to data.

The gradient based optimization was proven to be efficient for the albedo. However, we found it not quite so when working on altitude values. We, therefore, devised an algorithm which is based on minimizing the above energy function locally within some neighbourhood of each pixel. The advantage is then the possibility of regularizing the system of equations locally and hence penalizing errors is only taking place on a local level. As will be seen later the algorithm would then require no a priori knowledge for the depth map either.

4.2 Computational Considerations

Due to non-linearities involved in image formation one can easily see that the energy function given by (4.18) may not be easy to minimize in practice. Two important issues arised here are the scale of the dimension of the problem (even for small images), and the non-convexity. Therefore, for real practical applications of image processing, a direct solution is not computationally feasable. However, one can use methods such as the Successive Over Relaxation (SOR) [4] to find the minimum of $\Psi(g, z)$ or equivalently E.

Non-convexity can be tackled by means of certain stochastic or deterministic optimization techniques such as the SA, MFA or the GNC. SA allows for random increase of energy function so that one can escape out of local minima. The cost is, however, usually very high in terms of the computational time and the convergence tends to be very slow. GNC is relatively more attractive, but can be applied only when certain conditions are satisfied by the energy function. In the case when a good *a priori* knowledge is available, one can even obtain the solution by using simple gradient based algorithms, if the algorithm is initialized near the optimal solution. As we will see later, for z we have initially considered this option and then developped an algorithm that requires no *a priori* knowledge. As for g the gradient based algorithm was found to converge to the optimal solution with no initial good estimates, and indeed very good results have been obtained both with synthetic and real images. We have, therefore, been able to use a SOR algorithm based on (4.1):

$$(g, z)_n^{m+1} = (g, z)_n^m - \frac{\omega}{T}\left(\frac{\partial E(g, z)}{\partial (g, z)}\right) \tag{4.15}$$

where T is a successively increasing parameter and ω is an initializing constant. We have used the case where $\omega = 1$ which is refered to as the Gauss-Seidel algorithm [4]. SOR can easily lend itself to a massively parallel implemetation.

Another simplification made at this stage is that z has been assumed to be known when minimizing for g and vice versa.

The problem is now reduced to simple computation of the partial derivatives of the energy function. Here is how we have calculated the partial derivative of (4.14) with respect to g:

$$\frac{\partial E}{\partial g_i} = \frac{\partial E_g}{\partial g_i} + \sum_{k \in w} \frac{\partial E_k}{\partial g_i} \tag{4.16}$$

where:

$$\frac{\partial E_g}{\partial g_i} = \sum_{j \in v_i} \frac{g_i - g_j}{\sigma_g^2} \tag{4.17}$$

and

$$\frac{\partial E_k}{\partial g_i} = \frac{1}{\sigma_\varepsilon^2}\left(\hat{I}_k - I_k\right) H_k R_i \tag{4.18}$$

Here v_i is a set of pixels determined by the neighbourhood structure of the field. σ_g and σ_ε are the standard deviations of the albedo and the noise process respectively and w is a window whose size is determined by the PSF.

The partial derivative with respect to z requires first p and q to be calculated. Here we simply have applied Prewitt operators:

$$p_i = \sum_{j \in w_p} P_{ij} z_j \tag{4.19}$$

$$q_i = \sum_{j \in w_p} Q_{ij} z_j \,. \tag{4.20}$$

where P_{ij} and Q_{ij} are the coeficients of the Prewitt array and w_p is the support of the array. Now differentiating (4.14) with respect to z we get:

$$\frac{\partial E}{\partial z_i} = \left(\frac{\partial E_z}{\partial z_i}\right) - \frac{1}{\sigma_\varepsilon^2} \sum_{k \in w} \left(I_{n,k} - \hat{I}_{n,k}\right) \frac{\partial \hat{I}_{n,k}}{\partial z_i} \tag{4.21}$$

Simplifying we get:

$$\frac{\partial R_{ij}}{\partial z_i} = -(\alpha_L P_{ji} + \beta_L Q_{ji})\left(\frac{\vec{N}\cdot\vec{S}}{\sqrt{1+p_i^2+q_i^2}}\right)$$

$$-(\alpha_s P_{ji} + \beta_s Q_{ji})\left(\frac{\vec{N}\cdot\vec{L}}{\sqrt{1+p_i^2+q_i^2}}\right)$$

$$-(p_i P_{ji} + q_i Q_{ji})\left(\frac{\vec{N}\cdot\vec{L}\times\vec{N}\cdot\vec{S}}{\left(1+p_i^2+q_i^2\right)^{3/2}}\right) \qquad (4.22)$$

As for the first term in (4.21) we have:

$$\frac{\partial E_z}{\partial z_i} = \sum_{j\in v_i} \frac{z_i - z_j}{\sigma_z^2} \qquad (4.23)$$

4.3 A Method For Computing z Using Local Energy Functions

The gradient method explained above can pose problems due to non-convexity of the energy function. In this section we propose a method based on minimizing the energy function locally with an implicit local continuity constraint imposed over small surface patches. Assuming that the energy function is locally convex, we can partition the overall surface into small overlapped windows and minimize the energy function locally. Any changes taking place in the overlapped areas affect the initial conditions in the neighbouring windows. This, in fact, will impose a local continuity constraint between overlapping surface patches (see Fig 5.2)

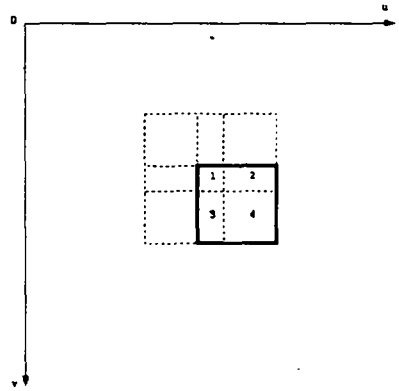

Fig 4.2. Minimizing energy functionals locally (see text for details).

On each iteration initial values of pixels in overlapping areas(regions 1, 2 and 3) of the square patch are affected by three adjacent windows: top, left and top left corner. Non-overlapping values (region 4) are given by those computed on the previous iteration. The value of the energy function in the square patch is then calculated prior to updating the pixel values. We call this E_i .We then update the values by small amounts (large steps can cause divergence). The local energy function is then calculated again to give the updated value E_u. Pixel changes are then accepted if $E_u \leq E_i$, otherwise they will be left untouched. The process is repeated sequentially until the whole image is covered. A mathematical formulation of the process is given below:

$$z_w^{n+1} = z_w^n + \delta\left(\frac{\tau}{T}\right) \tag{4.24}$$

where n is the iteration number, w is a sufficiently small window to maintain the spatial continuity, τ is a sufficiently small step size to maintain the altitude continuity($0 < \tau \leq 1$), T is a parameter for relaxation and δ is a function as below:

$$\delta(\Omega) = \begin{cases} \Omega \text{ if } E_u \leq E_i \\ \\ 0 \text{ otherwise} \end{cases} \tag{4.25}$$

To ease the task of updating we have assumed that the altitude is always positive (ie. the minimum altitude value is zero). This will amount to translating our observed surface along the z-axis until the minimum surface altitude coincides with the origin of the world coordinate frame.Updating will, then, take place only along the positive direction of the z-axis. This is equivalent to fixing one point in our surface and removing the translational ambiguity. Overlapped areas impose implicitly, horizontal, vertical and diagonal continuity constraints at the initialization level of each square patch. Values in these areas, however, are also subject to modifications on each iteration. When no patches can be updated over the entire image plane, the parameter T is increased and the process is resumed. We stop iterating when:

$$|E_i - E_u| \leq \eta \qquad \forall w_k \in \Im \tag{4.26}$$

where η is a small value, w_k is an elementary window and $\Im$ is the image plane. The smaller the value η the finer is the altitude map, but at the expense of a higher computational time

5 Experimental Results

In this section we present some experimental results on both synthetic and real images. Fig 6-1-(a) and 6-1-(b) show the ideal high resolution albedo and altitude. We have used our method to recover this high resolution albedo itratively, from a set of low resolution observations, with the assumption that the high resolution altitude is known. Below we can see the high resolution albedo recovered after 4 iterations.

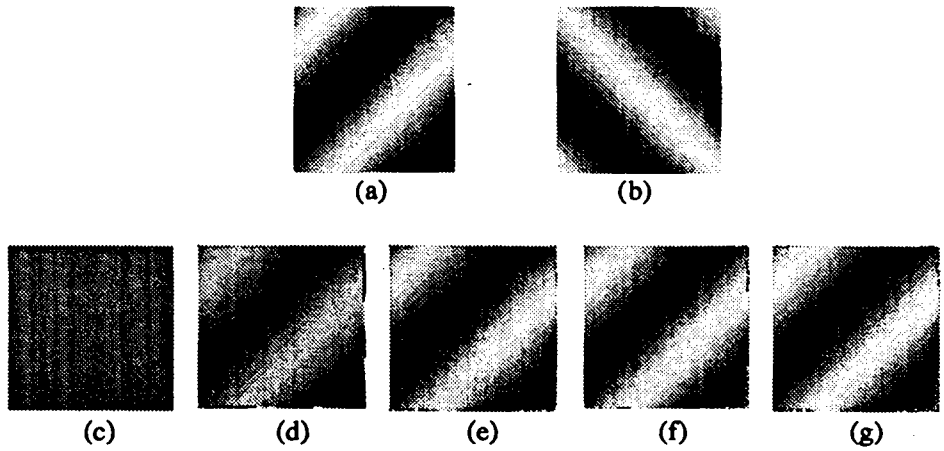

Fig 5-1 (a) and (b) high resolution albedo and altitude respectively; (c) initialization, (d), (e), (f) and (g) albedo recovered after the first four iterations.

The problem at the borders are due to the fact that the convolution process is performed after discretization by applying a Gaussian kernel. A better view of the efficiency of the algorithm can be obtained by looking at a cross section of the high resolution image after each iteration:

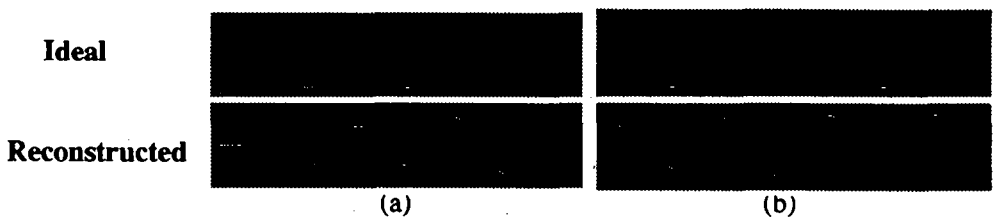

Fig 5-2 (a) Recovery after 2 iterations, (b) Recovery after 4 iterations

Here is an aerial image of Paris which we have used as our ideal albedo. Using this ideal albedo and the altitude map, we have simulated a set of nine low resolution observed intensity images, based on our image acquisition model. We then have run the algorithm using these observations and the altitude map to recover the ideal albedo. We have only shown one of the nine low resolution observations. The other eight differ only at subpixel level along the axis of the camera (see section 2 for details).

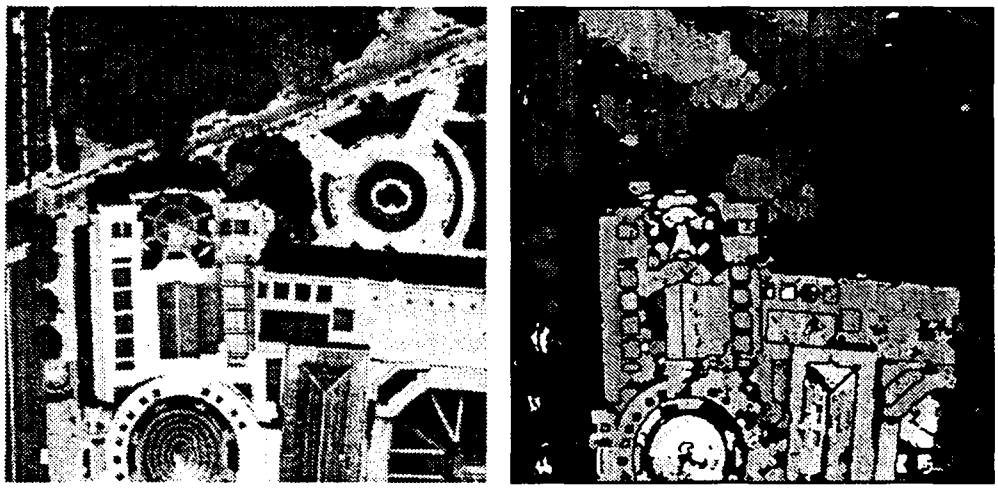

(a) (b)

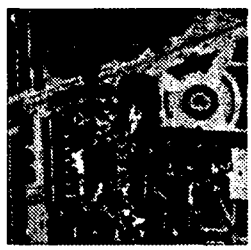

(c)

Fig 5-3 (a) ideal albedo (b) ideal altitude (c) one of the nine low resolution images

In Fig 5-4 we have initialized the algorithm with arbitrary albedo values. Results are shown after 6, 12 and 18 iterations:

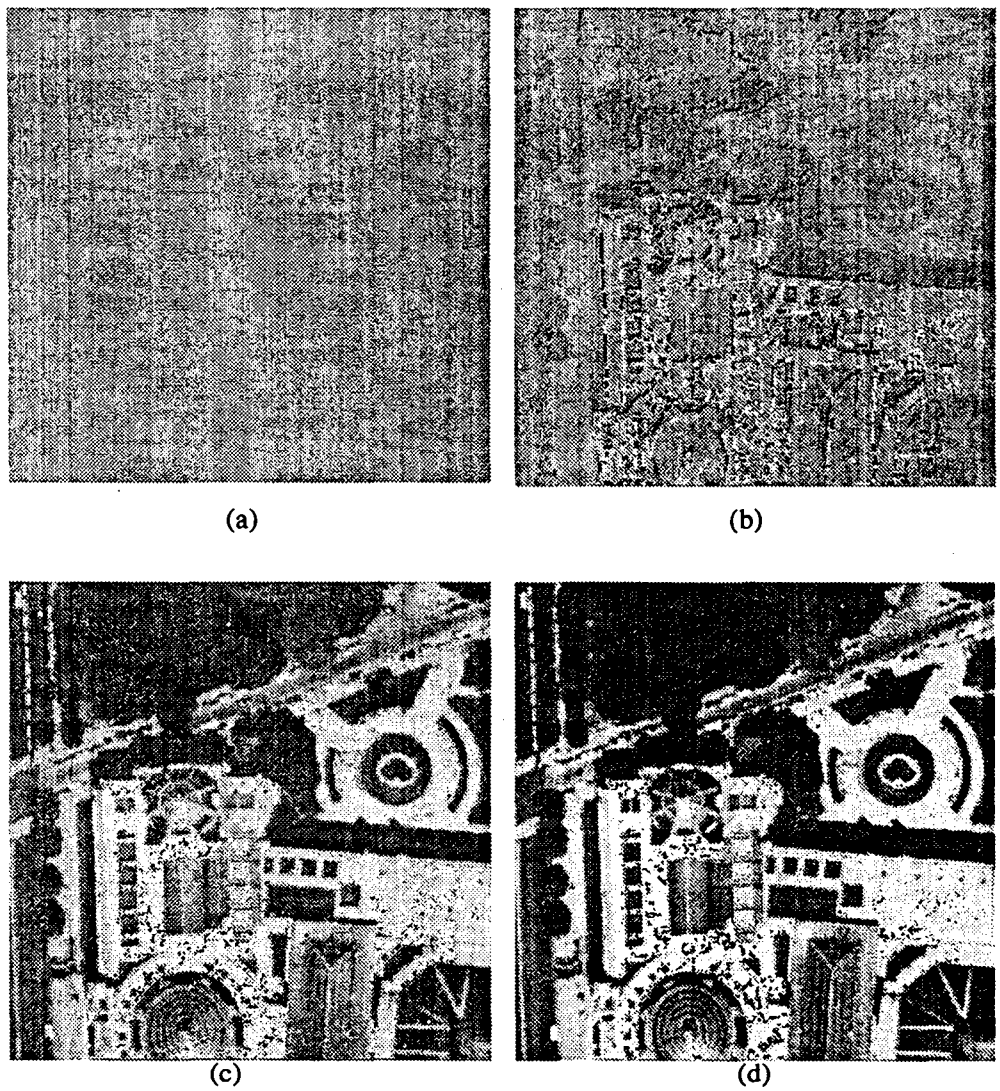

Fig 5-4 (a) initialization (b), (c) and (d) recovery of the albedo after 6, 12 and 18 iterations respectively (high resolution altitude assumed to be known)

Below is the altitude map recovered using MAP, assuming that the albedo is known. By initializing near the solution, one can attain the global minimum of the energy function. The algorithm is indeed very slow and would be more useful when coupled with stereo vision.

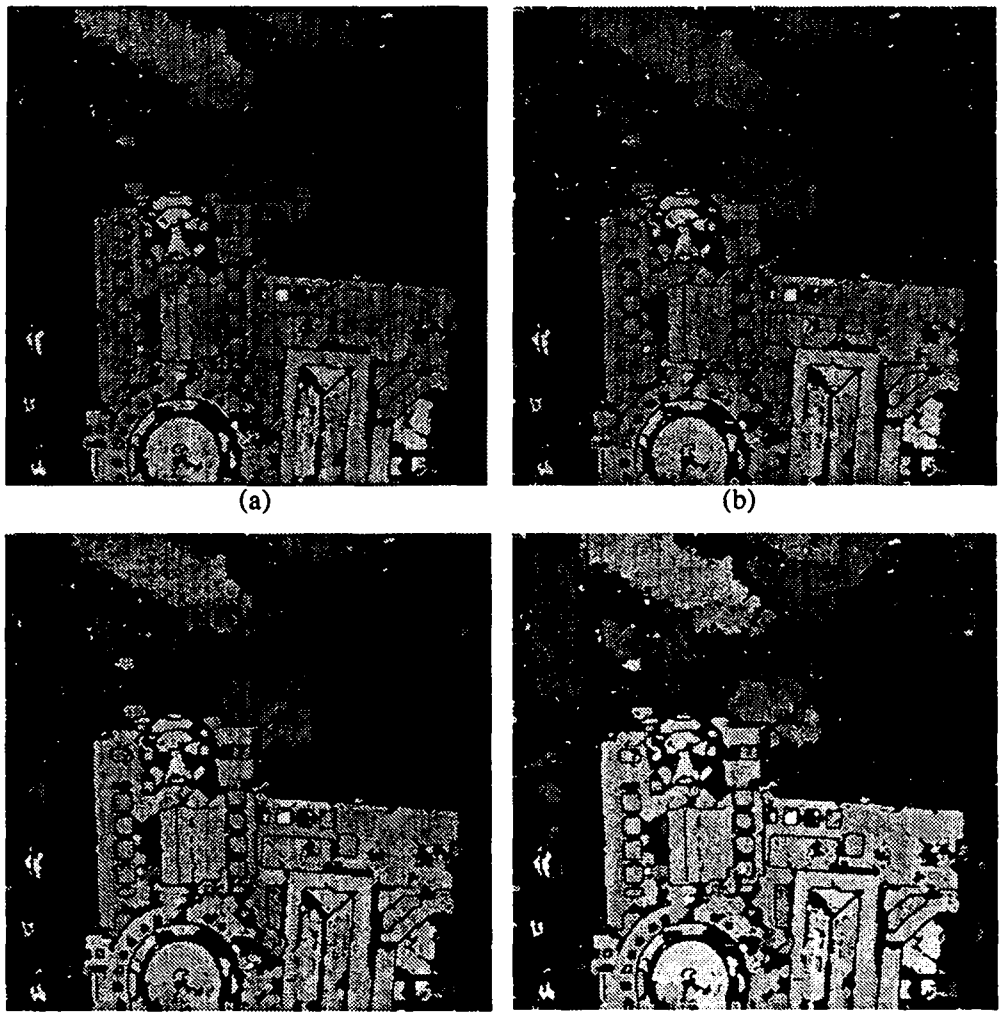

Fig 5-5 Recovery of depth map using MAP estimator: (a) initialization, (b), (c) and (d) are recoveries after 16, 28 and 40 iterations respectively.

We can see above that the gradient based algorithm shows poor convergence and initialization properties. We have overcome the problem by means of the algorithm described in section 4.2 which requires no *a priori* knowledge for initialization.

Fig 5-6 shows the recovery of the altitude map using the algorithm described in section 4.2. We have used only four low resolution observations while assuming that the ideal high resolution albedo is available.

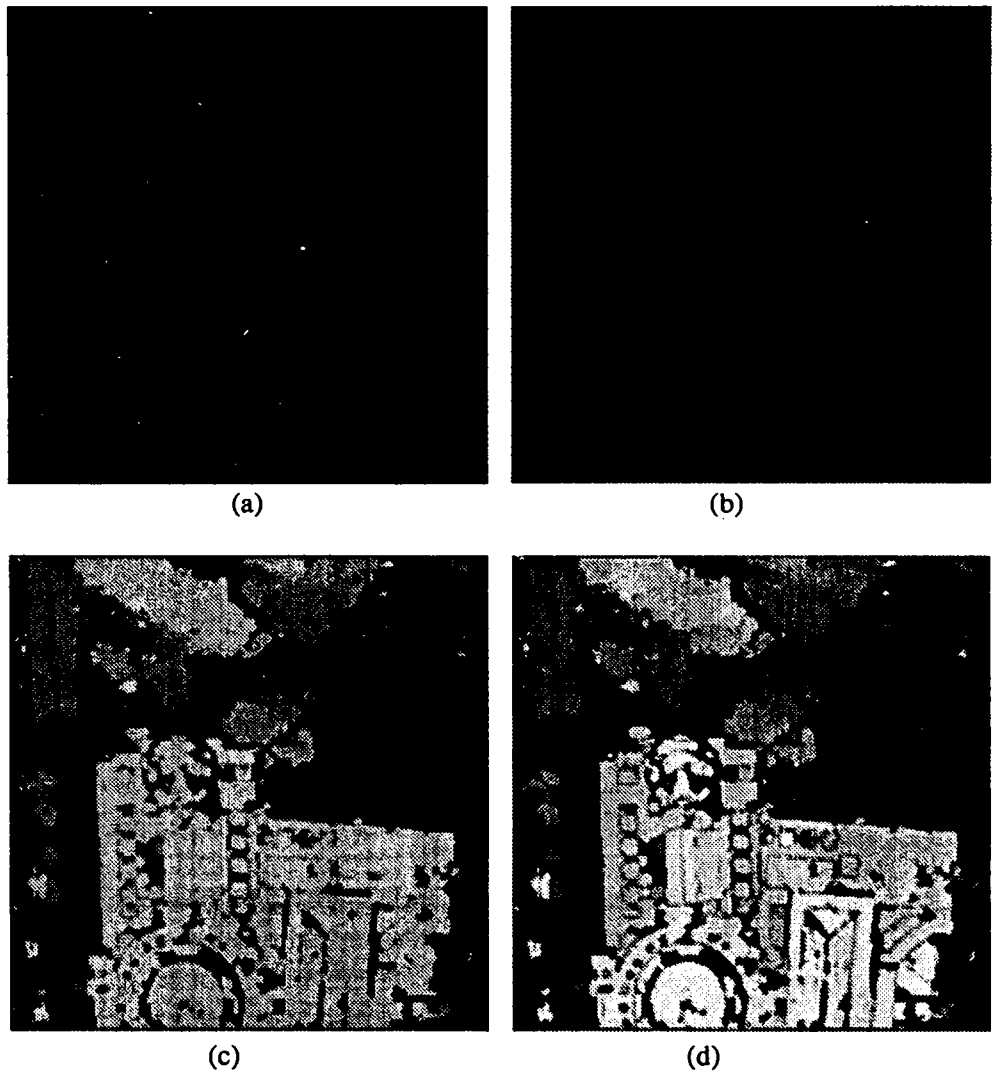

Fig 5-6 Recovery of altitude using 4 low resolution sensors: (a) initialization (b) recovery after one iteration (c) and (d) after 6 and 10 iterations respectively.

6 Concluding Remarks

We presented here an algorithm for recovering 3D information from a set of low resolution intensity images. It is important, however, to have interframe subpixel displacements [10][19] so that pixel information incorporated in our reconstructed high resolution images are determined by an oversampled surface. This, in essence, is equivalent to sampling the scene at a higher rate using a single high resolution camera. The advantage of the technique to pure interpolation is therefore the fact the the reconstructed image embodies interpixel values obtained directly from the actual scene using an error criterion, whereas, in interpolation interpixel values are computed using curve fitting techniques and carry no real information. Also, Change of spatial resolution using pure interpolation is usually obtained by trading the gray level quantization [19].

The problem for the albedo g requires no *a priori* knowledge for convergence, when using the gradient based optimization algorithm. By applying the method proposed in section 4.3. , it is also possible to recover the altitude map without any prior knowledge. Results are indeed very much promising.

In the case of g , the initialization of weighting parameters in SOR and the degree of relaxation can highly affect the speed of convergence, and the quality of final results, too. We have included two stopping conditions in our method. One is the usual MAP condition and the second one is when the algorithm does not minimize the penalty function after a fixed number of relaxations. Discretization prior to convolution can produce problems around the borders of the reconstructed images. These can be handled by means of extrapolation methods while still assuming continuity of the image within some neighbourhood width.

The method proposed here is equivalent to deblurring [8] for g and the computation of height from shading [23] for z if only one image frame is used in which case we can use it as an image restoration and height from shading algorithm. Note that it is, theoretically, possible to increase the resolution to any required level by increasing the number of frames. In practice, however, the number of interpixel samples can be limited by the low resolution camera and hence one can not achieve a resolution beyond those specified by the bandwidth of the channel, for example. We intend to extend our method to include other possible reflection models such as specular and inter-reflections. The robustness of the algorithm to noise will be scrutinized in future work and it will be implemented on a massively parallel SIMD architecture (Connection Machine CM200).

References

[1] R. Azencott, «Image Analysis and Markov Fields», Proc. of Int. Conf. Ind. & Appl. Math, SIAM, Paris, 1987.

[2] J. E. Besag, «Spatial Interaction and the Statistical Analysis of lattice Systems(with discussion)», Jl. Roy. Statis. Soc. B. 36, pp 192-236, 1974.

[3] J. E. Besag, «On the Statistical Analysis of Dirty Pictures», Jl. Roy. Statis. Soc. B.48,pp 259-302, 1986.

[4] A. Blake and A. Zisserman, «Visual Reconstruction, MIT Press, Cambridge - MA. 1987.

[5] A. Blake, A. Zisserman and G. Knowls, «Surface Description from Stereo and Shading»,Image & Vision Computing, vol. 3, no 4, pp183-191, 1985.

[6] A. R. Bruss, «The Eikonal Equation: Some Results Applicable to Computer Vision», Jl. of Math. Physics, vol. 23, no 5, pp 890-896, 1982.

[7] E. de Castro and C. Morandi, «Registration of Translated and Rotated Images Using Finite Fourier Transforms», IEEE Trans, PAMI-9, No 5, Sept 1987.

[8] R. T. Frankot and R. Chellappa, «A Method for Enforcing Integrability in Shape from Shading Algorithms», IEEE Trans. PAMI, vol. 10, no 4, pp 431-451, 1988.

[9] D. Geiger and F. Girosi, «Mean Field Theory for Surface Reconstruction and Visual Integration», A.I. memo no 1114, AI Laboratory, MIT, 1989.

[10] S. Geman and D. Geman, «Stochastic Relaxation, Gibbs Distributions and the Bayesian Restoration of Images», IEEE Trans. PAMI 6, pp721-741, 1984.

[11] D. Gross, «Super-resolution from Sub-pixel Shifted Pictures», Master's Thesis, Tel-Aviv University, Oct. 1986.

[12]. B. K. P. Horn, «Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View», Technical Report TR-79, Project MAC, MIT, 1970.

[13] B. K. P. Horn, «Understanding Image Intensities», Artificial Intelligence, vol. 8, no 2, pp 201-231, 1977.

[14] B. K. P. Horn, «Height and Gradient From Shading», Proc. of Image Understanding Workshop, DARPA, 1989.

[15] T. Huang and R. Tsai, «Multiframe Image Restoration and Registration», Advances in Computer Vision and Image Processing, volume 1, pp 317-339, JAI Press Inc., 1984.

[16] R. A. Hummel, B. Kimia and S. W. Zucker, «Debluring the Gaussian Blur», Computer Vision Graphics and Image Processing, 38, pp 66-80, 1987.

[17] D. Hutber, «Improvement of CCD Camera Resolution using a Jittering Technique», SPIE vol 849, Automated Inspection and High Speed Vision Architecture, 1987.

[18] M. Irani and S. Peleg. «Improving Resolution by Image Registration», Graphical Models and Image Processing, vol 53, No. 3, pp231-239, May 1991.

[19] M. Irani and S. Peleg. «Super-resolution from Image Sequences», Technical Report 89-7, Dept. of Computer Science, The Hebrew University of Jerusalem, June 1989.

[20] D. Keren, S. Peleg and R Brada, «Image Sequence Enhancement Using Sub-pixel Displacement», Proc. of CVPR, pp 742-746, Ann Arbor, Michigan, June 1988.

[21] D. Keren and M. Werman, «Variations on Regularization», Proc. of Int. Conf. on pattern Recognition, Atlantic City, 1990.

[22] Kim, Bose and Valenzuella, «Recursive Reconstruction of High Resolution Images from Noisy Undersampled Multiframes», IEEE Trans. on ASSP, 38(6), June 1990.

[23] Y. G. Leclerc and A. F. Bobick, «The direct Computation of Height from shading», Proc. of CVPR, 1991.

[24] D. G. Lowe, «Fitting Parametrized Three-Dimensional Models to Images», IEEE Trans, PAMI,vol 13, no 5, May 1991.

[25] J. Marroquin, «S. Mitter and T. Poggio, Probabilistic Solution of Ill-Posed Problems in Computational Vision», Jl. of American Statistical Association, 1987.

[26] J. Moussouris, «Gibbs and Markov Random Systems with Constraints», Jl of Stat. Physics, vol. 10, no 1, 1974.

[27] J. Oliensis, «Shape from Shading as a Partially Well-constrained Problem», CVGIP, Image Understanding, 1991.

[28] S. Peleg, D. Keren and L. Schweitzer, «Improving Image Resolution Using Subpixel Motion», Pattern Recognition Letters, pp 223-226, 1987.

[29] S. Peleg, M. Werman and H. Rom. «A unified Approach to the Change of Resolution: Spaceand Gray level», IEEE Trans PAMI, 11, pp 139-141,1989.

[30] J. Prager and M. A. Arbib, «Computing the Optic Flow: The Match Algorithm and Prediction», Computer Vision Graphics and Image Processing 24, pp 271-304, 1983.

[31] A. Rosenfeld and A. C. Kak, «Digital Picture Processing», vol. 1 and 2, 2nd Ed., Academic Press Inc., 1982.

[32] H. Shvayster and S. Peleg, «Inversion of Picture Operators», Pattern Recognition letters 5, pp 49-61, 1987.

[33] D. Terzopoulos, «Regularization of Inverse Visual Problems Involving Discontinuities», IEEE Trans. PAMI, vol 8, no 4, July 1986.

[34] A Tikhonov and A. Arsenin, «Solutions of Ill-posed Problems», 1977.

[35] L. B. Wolff, S. A. Shafer, G. E. Healey, «Radiometry, Physics Based Vision» (edited), Jones and Barlett Publishers, 1992.

[36] S. Zevin and M. Werman, private communications, Hebrew University of Jerusalem.

Unité de Recherche INRIA Sophia Antipolis
2004, route des Lucioles - B.P. 93 - 06902 SOPHIA ANTIPOLIS Cedex (France)

Unité de Recherche INRIA Lorraine Technopôle de Nancy-Brabois - Campus Scientifique
615, rue du Jardin Botanique - B.P. 101 - 54602 VILLERS LES NANCY Cedex (France)
Unité de Recherche INRIA Rennes IRISA, Campus Universitaire de Beaulieu 35042 RENNES Cedex (France)
Unité de Recherche INRIA Rhône-Alpes 46, avenue Félix Viallet - 38031 GRENOBLE Cedex (France)
Unité de Recherche INRIA Rocquencourt Domaine de Voluceau - Rocquencourt - B.P. 105 - 78153 LE CHESNAY Cedex (France)

EDITEUR
INRIA - Domaine de Voluceau - Rocquencourt - B.P. 105 - 78153 LE CHESNAY Cedex (France)

ISSN 0249 - 6399