

Surface Reconstruction by Integrating 3D and 2D Data of Multiple Views

Maxime LHUILLIER,
LASMEA UMR 6602
Université Blaise Pascal/CNRS, France

Long QUAN
Department of Computer Science
Hong Kong University of Science and Technology

Abstract

Surface representation is needed for almost all modeling and visualization applications, but unfortunately, 3D data from a passive vision system are often insufficient for a traditional surface reconstruction technique that is designed for densely scanned 3D point data. In this paper, we develop a new method for surface reconstruction by combining both 3D data and 2D image information. The silhouette information extracted from 2D images can also be integrated as an option if it is available. The new method is a variational approach with a new functional integrating 3D stereo data with 2D image information. This gives a more robust approach than existing methods using only pure 2D information or 3D stereo data. We also propose a bounded regularization method to implement efficiently the surface evolution by level-set methods. The properties of the algorithms are discussed, proved for some cases, and empirically demonstrated through intensive experiments on real sequences.

1. Introduction

Given a set of uncalibrated 2D images of a 3D scene acquired by a hand-held camera, obtaining a surface representation of the objects in the scene has always been one of the most challenging and fundamental problems of 3D computer vision. Although effective for computing camera geometry, most of recent approaches [10, 6] reconstruct only 3D 'point cloud' of the scene, whereas surface-based representations are indispensable for most current modeling and visualization applications. Surface reconstruction is a natural extension of the point-based geometric methods. But unfortunately 3D data from such a passive system are often insufficient for a direct surface reconstruction method that is designed for scanned 3D data. The main difficulties are that the 'passive' 3D points are noisy, sparse, irregularly distributed, and missing in many big parts. These difficulties have motivated us to develop a new approach to constructing surface representations from 3D stereo data [15], but using extra 2D image information that is still available from a passive system.

Surface reconstruction from 3D data Surface reconstruction from scanned 3D point data has been a very traditional research topic in both computer vision and graphics. Szeliski et al. [22] used a particle-based model of deformable surfaces; Hoppe et al. [11] presented a signed distance for implicit surfaces; Curless and Levoy [3] described a volumetric method; and Tang and Medioni [24] introduced a tensor voting method. Most recently, Zhao et al. [26] developed a level-set method based on a variational method of minimizing a weighted minimal surface. Similar work to [26] has also been developed by Whitaker [25] using a MAP framework. Surface reconstruction from depth data obtained from stereo systems is more challenging than that from scanned 3D data as the stereo data are usually much sparser and less regular. Fua [8] used a system of particles to fit the stereo data. Kanade et al. [17] and Fua and Leclerc [7] proposed a deformable mesh representation to match the multiple dense stereo data. These methods that perform reconstruction by deforming an initial model or tracking the discretized particles to fit the data points are both topologically and numerically limited compared to modern dynamic implicit surface approaches.

Surface reconstruction from 2D images Recently, several new volumetric algorithms have been proposed [20, 14, 5, 13] that simultaneously reconstruct the surface and obtain dense correspondence. The method of space carving or voxel coloring [20, 14] directly works on discretized 3D space, voxels, based on their image consistency and visibility. Kolmogorov and Zabih [13] proposed a direct discrete minimization formulation that is solved by graph cuts. A more general theory has been laid down by Faugeras and Keriven [5]. It is a variational method implemented by level sets, which is intrinsically a multiple view and handles naturally the topology changes and occlusion problems based on their earlier work [4, 19]. The approach has great potential, and some results have been presented. However, it is not clear under what conditions their methods converge as the actual proposed functional seems highly non-convex.

Our contributions Some ideas in our approach are inspired by these methods, but fundamentally these methods either solely operate on 3D data or on 2D data. Our ap-

proach tries to bridge these approaches by combining both 3D and 2D data. This is possible because of a unified functional based on a minimal surface formulation. We believe that the combined functional will have far less local minima than the one derived from 2D data alone, and that this will result in more stable and more efficient algorithms. For the efficient evolution of surfaces, we also propose a bounded regularization method based on level-set methods. Its stability is also proved. Intensive experiments are also carried out to demonstrate the validity of our method.

2. Problem statement

Given a set of calibrated 2D images, a set of 3D points derived from the given images, and optionally a set of silhouettes extracted from the given images, the goal is to reconstruct a surface representation of the objects in the scene.

The problem is different from surface reconstruction from a set of calibrated images as addressed in [5, 20, 14] in which only 2D images are used without any 3D information. It is also different from surface reconstruction from scanned 3D data without 2D image information [11, 22, 3, 8, 17, 26, 24].

It is very important to emphasize that the set of 3D points considered here is not from scanned data, but derived from the given set of images by stereo and bundle-adjustment methods [15]. The 'passive' 3D points are difficult to reconstruct as they are sparser and irregularly distributed and have more missing parts than the 'active' data. The major motivation of this study is to improve the insufficiency of 3D stereo data by using original 2D image information. The obtainment of 3D stereo data will be discussed in Section 7.

3. General Approach

The general methodology that we will follow is a variational approach inspired by the work of Faugeras and Keriven [5], Caselles et al. [1, 2], Zhao et al. [26], and many others. An intrinsic functional as a kind of weighted minimal surface is defined to integrate both 3D point data and 2D image data. The object surfaces are represented as a dynamic implicit surface $u(\mathbf{x}) = 0$ in \mathbf{R}^3 which evolves in the direction of the steepest descent provided by the variation calculation of the functional we define to minimize. The intrinsic nature of the functional (i.e., independent of any surface parametrization) makes the implementation of surface evolution by the level-set method possible, which in turn handles the surface topology changes.

Our contribution is twofold. We first introduce a new intrinsic functional which takes into account both 3D data points and 2D original image information, unlike previous works that consider either only 2D image information [5] or only scanned 3D data [26] were considered. The new

functional is expected to have much smaller number of local minima and better convergence. Secondly we propose a bounded regularization method that is more efficient than the usual full regularization methods.

4. Defining the functional

By analogy to 2D geodesic active contours [1] whose nice mathematical properties have been established, the weighted minimal surface formulation was introduced by Caselles et al. [2], and Kichenassamy et al. [12] for 3D segmentation from 3D images, i.e., the 3D surfaces they seek are those minimizing the functional $\int \int w ds$ using the weight $w = g(\nabla I)$ where ds is the infinitesimal surface element and g is a positive and decreasing function of the 3D image gradient ∇I .

Faugeras and Keriven [5] developed a surface reconstruction from multiple images by minimizing the functional $\int \int w ds$ using a weighting function w that measures the consistency of the reconstructed objects reprojected onto 2D images. This measure is usually taken to be a function of the correlation functions $\rho(\mathbf{x}, \mathbf{n})$ between pairs of 2D images, i.e., $w(\mathbf{x}, \mathbf{n}) = g(\rho(\mathbf{x}, \mathbf{n}))$. The correlation function is dependent not only on the position \mathbf{x} of the object surface, but also its orientation \mathbf{n} . Reference [5] and the extended technical report also derived all the fundamental evolution equations provided by the Euler-Lagrange equations of the functional for the weighting function involving the surface normal. A potentially general and powerful reconstruction approach was therefore established. But the existence and uniqueness of a solution for the proposed functional have not yet been elucidated.

In a different context for surface reconstruction from sufficiently dense and regular sets of scanned 3D point data, Zhao et al. [26] proposed to minimize the functional $\int \int w ds$ using a new weighting function w to be the distance function of any surface point \mathbf{x} to the set of 3D data points. Given a set of data points \mathcal{P} and $d(\mathbf{x}, \mathcal{P})$ the Euclidean distance of the point \mathbf{x} to \mathcal{P} , the weighting function is simply $w(\mathbf{x}) = d^p(\mathbf{x}, \mathcal{P})$. It gives interesting results with good 3D data points.

For our purpose of surface reconstruction, we have both 3D data points and 2D image data. But it is interesting to observe that the variational formulation mentioned above in different contexts is all based on the minimal surface. This makes it possible to define a unifying functional taking into account data of a different nature. Thus, we propose to minimize the functional $\int \int w ds$ using a new weighting function for the minimal surface formulation consisting of two terms $w(\mathbf{x}, \mathbf{n}) = d^p(\mathbf{x}, \mathcal{P}) + \lambda e(\mathbf{x}, \mathbf{n}, I)$ where the first $d(\mathbf{x}, \mathcal{P})$ is the 3D data attachment term that allows the surface to be attracted directly onto the 3D points; and the second $e(\mathbf{x}, \mathbf{n}, I)$ is a consistency measure of the reconstructed object in the original 2D image space. The consistency mea-

sure might be taken to be any photo-consistency or correlation function. The functional to minimize is given by

$$p(\mathbf{x}) = \int \int (d^p(\mathbf{x}, \mathcal{P}) + \lambda \epsilon(\mathbf{x}, \mathbf{n}, I)) ds.$$

The silhouette information is also a useful source of information for surface construction [23]. It is not sufficient on its own as it gives only an approximate visual hull, but it is complementary to other sources of information. It amends the distance function of the weighting function as

$$d(\mathbf{x}, \mathcal{P} \cup \mathcal{S}) = \min(d(\mathbf{x}, \mathcal{P}), \epsilon + d(\mathbf{x}, \mathcal{S})),$$

where d is the 3D Euclidean distance function; \mathcal{P} is the set of 3D points; \mathcal{S} is the surface of the intersections of the cones defined by the silhouettes, i.e., the visual hull; and ϵ is a small constant favoring 3D points over the visual hull in the neighborhood of 3D points.

5. Solving the variational problem

The solutions of the minimizing functional are given by a set of PDEs: the Euler-Lagrange equation designated $\nabla p = 0$, and obtained from the functional $p = \int \int w ds$ to be minimized. The Euler-Lagrange equation is often impossible to solve directly. One common way is to use an iterative and steepest-descent method by considering a one-parameter family of smooth surfaces $\mathbf{x}(t) : (u, v, t) \mapsto (x(u, v, t), y(u, v, t), z(u, v, t))$ as a time-evolving surface \mathbf{x} parametrized by time t . The surface moves in the direction of the gradient of the functional p with the velocity $-\nabla p$, according to the flow $\mathbf{x}_t = \frac{\partial \mathbf{x}(u, v, t)}{\partial t} = -\nabla p$. This is the Lagrangian formulation of the problem that describes how each point on the dynamic surface moves in order to decrease the weighted surface. The final surface is then given by the steady state solution $\mathbf{x}_t = 0$. The problem with this approach is well known [21] because it does not handle the topology change. However it is important to notice that though the derivation has been based on a parametrization, the various quantities including the velocity for the steepest descent flow are intrinsic, i.e., independent of any chosen parametrization that makes the computation possible. This paves the way for the well-known and powerful level-set formulation [18, 21] that regards the surface as the zero level-set of a higher dimensional function. As the flow velocity $-\nabla p$ is intrinsic (it has been demonstrated for a general w depending also on surface normal in [5]), we may easily embed it into a higher dimensional smooth hyper-surface $u(t, \mathbf{x}) = 0$ which evolves according to $u_t = -(\nabla p \cdot \mathbf{n}) \|\nabla u\|_2$ and the normal $\mathbf{n} = -\frac{\nabla u}{\|\nabla u\|_2}$. Topological changes, accuracy, and stability of the evolution are handled by using the proper numerical schemes developed by Osher and Sethian [18].

6. A bounded regularization method

The Bounded Regularization Method The Euler-Lagrange expression ∇p might be complicated if the

weighting function $w(\mathbf{x}, \mathbf{n})$ depends also on the normal of the surface [5]. Unfortunately or fortunately it seems that the complication by this dependency on the surface normal is rather unnecessary in practice [9]. We therefore assume a weighting function independent of the surface normal. Thus, the expression $\nabla p \cdot \mathbf{n}$ consists simply of two terms like the geodesic active contour case $\nabla w \cdot \mathbf{n} + w \nabla \cdot \mathbf{n}$, in which the first is the data attachment term and the second the regularization term. By using $\mathbf{n} = -\frac{\nabla u}{\|\nabla u\|_2}$ on the level-set function, the surface evolves according to

$$\frac{\partial u}{\partial t} = \nabla w \nabla u + w \|\nabla u\|_2 H$$

where $H = \nabla \cdot \frac{\nabla u}{\|\nabla u\|_2}$ is the sum of the two principal curvatures (twice its mean curvature). When w is taken to be the correlation functions, it is the simplified version of [5] presented in [9]. And when w is taken to be the 3D distance function, it is the first method proposed in [26]. However the curvature-based regularization $w \|\nabla u\|_2 H$ over-smooths, resulting in a loss of geometric details and in slow convergence as the time step has to be $\Delta t = O(\Delta x^2)$ for a stable solution.

In [27], a convection model is also proposed to simply ignore the regularization term $w \|\nabla u\|_2 H$ and speed up the procedure, but this is only envisageable for applications where data quality is sufficient, for instance, for synthetic and high quality scanned data [27].

Motivated by the need for regularization of noisy data and the inefficiency of the curvature-based regularization, we propose an intermediate bounded regularization method. It has a ‘‘bounded’’ regularization term $\min(w, w_{max}) \|\nabla u\|_2 H$ instead of the ‘‘full’’ regularization term $w \|\nabla u\|_2 H$. The corresponding evolution equation is given as:

$$\frac{\partial u}{\partial t} = \nabla w \nabla u + \min(w, w_{max}) \|\nabla u\|_2 H.$$

The following remarks can be made:

- The fully regularized surface evolution is obtained when $w_{max} \geq \|w\|_\infty$.
- The unregularized one is obtained when $w_{max} = 0$.
- As $0 \approx w \leq w_{max}$ in the vicinity of the steady surface for any w , it is expected that the fully regularized and the bounded regularized evolutions behave the same manner in this region.

Efficiency of the bounded regularization method The efficiency of our proposed bounded regularization method is evaluated by estimating the maximum time step Δt_{max} for stability computation. We are currently yet unable to quantify Δt_{max} of the bounded regularization method for the general curvature-based regularization, but we are able to prove it for a simplified isotropic regularization using a Laplacian operator. This is motivated by the fact that

the curvature/anisotropic regularization term $\|\nabla u\|_2 H = \|\nabla u\|_2 \text{div} \frac{\nabla u}{\|\nabla u\|_2}$ and the Laplacian/isotropic one Δu are equal when $\|\nabla u\|_2 = 1$ is enforced, which is the case in practice to avoid too flat and too steep variations of u . It is therefore tempting to simplify the evolution equation as

$$\frac{\partial u}{\partial t} = \nabla w \nabla u + \min(w, w_{max}) \Delta u.$$

The following result can be established:

Assuming that the stability condition is the same for curvature-based and Laplacian-based regularizations, the stability $\|u^{n+1}\|_\infty \leq \|u^n\|_\infty$ is achieved if $\Delta t \leq \Delta t_{max}$ with

$$\Delta t_{max} = \frac{\Delta x^2}{6w_{max} + \|\Delta x(|d^{0x}w| + |d^{0y}w| + |d^{0z}w|)\|_\infty}.$$

This can be proved as follows. Let assume $w \geq 0$, $\|w\|_\infty \geq w_{max} \geq 0$. Take $\Delta x = \Delta y = \Delta z$ as the space step, and Δt the time step. Also denote the centered (resp. forward and backward) differences for the x -axis at the 3D grid point ijk by d_{ijk}^{0x} (resp. d_{ijk}^{+x} , d_{ijk}^{-x}). Further, let $a_x = \frac{\Delta t}{\Delta x} d_{ijk}^{0x} w$ and $b = \min(w_{ijk}, w_{max}) \frac{\Delta t}{(\Delta x)^2}$ with similar notations a_y, a_z .

We may choose the simplest up-wind discretization scheme for $\nabla w \nabla u$ and centered discretization scheme for Δu . Then the equation $\frac{\partial u}{\partial t} = \min(w, w_{max}) \Delta u + \nabla w \nabla u$ is discretized by $u_{ijk}^{n+1} = u_{ijk}^n + \Delta t(l_{ijk}^n + c_{ijk}^{xn} + c_{ijk}^{yn} + c_{ijk}^{zn})$ such that $l_{ijk}^n \Delta t$ is equal to

$$b(u_{i+1jk}^n + u_{i-1jk}^n + u_{ij+1k}^n + u_{ij-1k}^n + u_{ijk+1}^n + u_{ijk-1}^n - 6u_{ijk}^n)$$

and $c_{ijk}^{xn} = d_{ijk}^{0x} w ((d_{ijk}^{0x} w < 0) ? d_{ijk}^{-x} u : d_{ijk}^{+x} u)$ with similar notations for c_{ijk}^{yn} and c_{ijk}^{zn} .

Now we only need to prove for the case $0 \leq a_x, a_y \leq 0$, $0 \leq a_z$, as other cases can be performed in a similar manner. In this case, we have $u_{ijk}^{n+1} = u_{ijk}^n + a_x(u_{i+1jk}^n - u_{ijk}^n) + a_y(u_{ij+1k}^n - u_{ijk}^n) + a_z(u_{ijk+1}^n - u_{ijk}^n) + b(u_{i+1jk}^n + u_{i-1jk}^n + u_{ij+1k}^n + u_{ij-1k}^n + u_{ijk+1}^n + u_{ijk-1}^n - 6u_{ijk}^n)$. We notice that u_{ijk}^{n+1} is a weighted sum of $u_{i'j'k'}^n$, such that the sum of weights is 1. It is also easy to check that all these weights are in $[0, 1]$ iff $|a_x| + |a_y| + |a_z| + 6b \leq 1$. This condition is satisfied if $\Delta t \leq \Delta t_{max}$. In this case, u_{ijk}^{n+1} is in the convex hull of $u_{i'j'k'}^n$, and thus we have proved that $\forall ijk, |u_{ijk}^{n+1}| \leq \|u^n\|_\infty$. \diamond

We choose w_{max} to be proportional to Δx for our bounded regularization method, i.e., fixing $w_0 = \frac{w_{max}}{\Delta x}$, and obtain

$$\Delta t_{max} = \frac{\Delta x}{6w_0 + \|\|\nabla w\|_1\|_\infty}.$$

Under this condition, the complexity of Δt_{max} is given by $\Delta t_{max} = \Theta(\Delta x)$, the same for the bounded regularized and unregularized evolutions, much better than $\Delta t_{max} = \Theta(\Delta x^2)$ for the fully regularized evolution.

In practice, the time step $\Delta t = \Delta t_{max}$ is always used for surface evolution of all our examples with the bounded and curvature-based regularization.

7. Implementation and experiments

7.1. Data Acquisition

Acquisition of 3D stereo data We choose to implement a general reconstruction method from an uncalibrated sequence [15]. First points of interest are detected and an initial sparse correspondence is established for each pair of images. Then the disparity map by best first propagation is computed for each pair. Next, we take each regular grid point in one image plane and its transformation by the grid-induced homography in the other images as a new corresponding pair of points. All these corresponding grid points are used for further N-view geometry computation. Finally, all the camera geometry and point positions are optimized over the whole sequence using a bundle adjustment method. The difference with the standard uncalibrated approach [10, 6] is that we use a much denser set of points re-sampled from the disparity map instead of points of interest, increasing the robustness, reducing the uncertainties, and obviously increasing the number of 3D points for surface reconstruction.

To model a complete object, we usually make a full turn around the object by capturing about 30-35 images to compute the geometry of the sequence.

Surface initialization from 3D stereo data The stereo data points are segmented into the foreground object and the background. The background includes obviously outliers like the isolated and distant points from the majority. The points of the foreground object are obtained as the largest connected component of the graph neighborhood of all points such that the distance between any two ‘‘edge’’-points of this graph should be smaller than a multiple of the uncertainty median of the points. The surface initialization is then obtained as follows. The object points are regularly sliced into sections along the major direction of the point cloud. A 2D-convex hull is computed for each section and these convex hulls are used to define the successive sections of a truncated cone as the bounding volume of the object. The initialization of all examples shown in this paper is automatically obtained using this method. One example of the initialization for the Bust sequence is shown on the top-left of Figure 7. We note that the initialization procedures proposed in [26, 27] can not be applied here because of the too big holes without 3D points, especially at the object bottom.

Also all 3D points are rescaled into a $150 \times 150 \times 150$ voxel space for all examples by applying a similarity transformation. The resulting voxel size is of the same order of magnitude as the uncertainty median of the 3D points.

7.2. Description of different methods

BR3D is the Bounded Regularization method by taking the weighting function w to be only the 3D distance from the set of the reconstructed 3D points: $w(\mathbf{x}) = d(\mathbf{x}, \mathcal{P})$. The number of iterations is always 100 with $w_0 = 0.1$.

BR2D is the Bounded Regularization method by taking the weighting function w to be the image correlation function ρ . More details are given in Section 7.3.

BR3D2D is the method using the weighting function w as a combination of a 3D distance function and a 2D image consistency measure using a bounded regularization method: $w(\mathbf{x}) = d(\mathbf{x}, \mathcal{P}) + \epsilon e(\mathbf{x}, I)$, where $e = 0.2\sqrt{\sigma_r^2(\mathbf{x}) + \sigma_g^2(\mathbf{x}) + \sigma_b^2(\mathbf{x})}$ and $\sigma(\mathbf{x})$ is the standard deviation of the reprojected voxel in each of three color channels in $[0, 1]$. The consistency measure e is similar to the photo-consistency of the space-carving method. The basic idea is to avoid surface evolution in the immediate neighborhood of the reconstructed points where the surface previously obtained by BR3D is assumed to be correct. It also inflates the surface elsewhere and stops in surface portions having inconsistent reprojections, mainly due to the difference between the object and the background colors. Thus, we use the following evolution equation

$$\frac{\partial u}{\partial t} = \nabla w \nabla u + \min(w, w_{max}) \|\nabla u\|_2 (c + H),$$

where $\frac{\partial}{\partial t} c$ is an inflating constant introduced and used in segmentation works [16, 2]. Note that the term $\min(w, w_{max}) \|\nabla u\|_2 c$ is negligible in areas where $w \approx 0$, i.e. in the close neighborhood of reconstructed points. This is a much desired outcome. We choose

- $\epsilon = 0$ in the immediate neighborhood of reconstructed points $d(\mathbf{x}, \mathcal{P}) < 2\Delta x$ in the unit cube $[0, 1]^3$ and $\epsilon = 1$ elsewhere;
- $c = -5$ and $w_0 = 0.1$ with $u < 0$ inside the current surface $u = 0$.

BR3DS The mixed method **BR3DS** combines both 3D points and the silhouette information by using a weighting function, $w(\mathbf{x}) = \min(d(\mathbf{x}, \mathcal{P}), \epsilon + d(\mathbf{x}, \mathcal{S}))$. We choose $\epsilon = 2\Delta x$ to favor the 3D points \mathcal{P} over the visual hull \mathcal{S} in the immediate neighborhood of reconstructed points.

BR3D+2D The BR3D and BR3D2D methods are sequentially applied. Fifty iterations are used with BR3D2D.

BR3D+S The BR3D and BR3DS methods are sequentially applied. Fifty iterations are used with BR3DS.

BR3D+BR2D The BR3D and BR2D methods are sequentially applied. Fifty iterations are used with BR2D.

Freeze plane To avoid the convergence of the dynamic surface to the empty surface, a freeze plane is often introduced to stop/freeze the surface evolution in one of the two delimited half spaces. The freeze plane is manually put to fill in the biggest gap, often on the bottom or on the back of the object if the sequence is not complete.

7.3. Results, comparisons, and discussions

The reconstructed surfaces and experiments summary are shown in Figure 1 on many image sequences taken by a hand-held still digital camera, except the Lady sequence taken by a special device. Each row first shows 3 images of the given sequence, which are followed by the reconstructed stereo points, a Gouraud-shaded and a textured-mapped view of the surface, both from the same viewpoint.

BR3D vs. BR3D+2D Combining 2D image information using BR3D2D can significantly improve the final reconstruction results as using only a 3D distance function may fail when there are no sufficient reconstructed points on some parts of the surface. This is illustrated in Fig. 2.

3D distance vs. image correlation Using only image correlation as suggested in [5, 9] makes converge very difficult for low-textured objects. Here we take a reasonably textured object, the bust, to test BR2D method and compare it with the others.

The surface initialization is shown on the top-left of Figure 7 and is obtained with the method described in Section 7.1. Figure 3 shows the results by BR2D method with $w = 0.1(1 - \rho)$ and $\rho \in [-1, 1]$ for 400 (top) and 1000 (bottom) iterations with $w_0 = 0.1$ (left), $w_0 = 0.5$ (middle), and $w_0 = 1$ (right), using a 9×9 ZNCC-window. The low bound $w_0 = 0.1$ gives a too noisy surface, see the pyramid part. The big bound $w_0 = 1$ gives a too smooth surface, see a too flat nose. The intermediate bound gives a compromise between the two. The original correlation [5, 9] with full regularization is even smoother than the big bound $w_0 = 1$ case. Also the convergence is extremely slow, it is still not done around the intersection of the concave part between the cube and the pyramid after 800 iterations. We have also found that the original correlation method is really slower than BR2D, since its time step Δt_{max} is 380 smaller.

We experimented with the two-step method BR3D+BR2D. The results are shown in Figure 7. The results are similar to the previous case, and not very satisfactory. However this method is more efficient: the 100 steps of BR3D-iterations takes only about 5 min. on a P4 2.4 GHz (including initialization), compared with the 20 (resp. 50) min. for 400 (resp. 1000) BR2D steps.


Figure 4 shows the difference between BR3D (middle) and BR3D+BR2D methods (right), with the best previous bound $w_0 = 0.5$ and only 50 iterations for BR2D. Still the nose is smoothed too much, and the chin is also degraded.

Isotropic vs. anisotropic smooth Using Laplacian/isotropic Δu instead of the curvature/anisotropic smooth $\|\nabla u\|_2 H = \|\nabla u\|_2 \text{div} \frac{\nabla u}{\|\nabla u\|_2}$ leads to faster evolution, as the level-set function update $\|\nabla u\|_2 = 1$ is done twice as frequently for the anisotropic smoothing than for the isotropic smoothing, which has a smaller discretization neighborhood. It is also important to observe that no apparent difference occurs between these two different smoothings in the final surface geometry, as shown in Figure 5 (idem with $w_0 = 0.5$ or 1). This suggests that the benefit of using curvature-based smooth is negligible in our context.


With vs. without silhouette Figure 6 shows results obtained by BR3D, BR3D+2D, BR3D+S and the pure silhouette method S for the Man 3 sequence. Using only 3D points by BR3D will miss the low-textured cheeks, and using only

Bust

#C=26
 #P=28065
 R=480x640
 M=BR3D
 F=bottom

**Man 1**


#C=38
 #P=53406
 R=480x640
 M=BR3D
 F=bottom

**Apple**


#C=38
 #P=43414
 R=450x450
 M=BR3D+S
 F=none

**Man 2**

#C=28
 #P=15862
 R=640x512
 M=BR3D
 F=bottom

**Lady**

#C=20
 #P=26823
 R=768x512
 M=BR3D
 F=back

**Bottle**

#C=26
 #P=33156
 R=504x752
 M=S
 F=none

**Man 3**

#C=34
 #P=28538
 R=480x640
 M=BR3D+2D
 F=bottom




Figure 1. Each row shows the results for one example. First, experimental details. Second, 3 frames of the sequence. Third, reconstructed 3D points. Fourth, Gouraud-shaded surface geometry. Fifth, textured-mapped view of the geometry. In the details, #C the number of cameras, #P the number of points; R image resolution, M the method used for the example, and F location of the freeze plane. Times are about 5 & 3 min. for BR3D(S) & BR3D2D with a P4 2.4GHz.

the visual hull by S will miss many important concavities on the surface like in the areas of the ears and nose. Combining the two gives excellent final results.

Adding silhouette information improves the pure 3D results, but automatic extraction of silhouette from unknown backgrounds is difficult and is done manually for this case.

8. Conclusion

We have presented a surface reconstruction method from 3D stereo data acquired by a hand-held camera. The approach has been based on the variational calculus. We have introduced a new functional integrating both 3D stereo data points and 2D image information. It builds a bridge between the pure 2D image-based approach (having too many local minima to be efficient in practice) and the pure 3D point driven approach (often fails due to missing data). We have also presented an efficient bounded regularization method for implementing the level-set evolution. The new bounded regularization method is significantly faster than the previous and full regularization. The stability has been proved for the simplified isotropic smoothing case. The methods have been intensively tested on many real sequences and very convincing results have been shown.

Acknowledgements

This project is partially supported by the Hong Kong RGC grant HKUST 6188/02E. We thank Dayton Taylor for providing the Lady sequence.




Figure 2. Surface geometry obtained by BR3D in (a) and BR3D+2D in (b). There are many missing 3D points in the low-textured cheeks (cf. Fig. 1), so BR3D using only 3D information gives poor results while BR3D+2D gives better results by adding 2D information.

References

- [1] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *IJCV*, 22(1):61–79, 1997.
- [2] V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert. Minimal surfaces based object segmentation. *TPAMI*, 19(4):394–398, 1997.
- [3] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *SIGGRAPH'96*.
- [4] R. Deriche, S. Bouvin, and O. Faugeras. A level-set approach for stereo. *SPIE Conference 2942: Investigative Image Processing*, 1996.
- [5] O. Faugeras and R. Keriven. Complete dense stereovision using level set methods. *ECCV'98* (and INRIA report 3021).
- [6] O. Faugeras, Q. Luong, and T. Papadopoulos. *The Geometry of Multiple Images*. The MIT Press, 2001.
- [7] P. Fua. Parametric models are versatile: The case of model based optimization. *ISPRS WG III/2 Joint Workshop*, 1995.
- [8] P. Fua. From multiple stereo views to multiple 3d surfaces. *IJCV*, 24(1):19–35, 1997.
- [9] J. Gomes and O. Faugeras. Pde-based stereo applied to human faces capture. *IMPROOFS 23-515 Project*, 1998.
- [10] R.I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. The CU Press, 2000.
- [11] H. Hoppe, T. Derose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Computer Graphics*, 26:71–77, 1992.
- [12] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. *ICCV'95*.
- [13] V. Kolmogorov and R. Zabih. Multi-camera scene reconstruction via graph cuts. *ECCV'02*.
- [14] K.N. Kutulakos and S.M. Seitz. A theory of shape by space carving. *ICCV'99*.
- [15] M. Lhuillier and L. Quan. Quasi-dense reconstruction from image sequence. *ECCV'02* (see papers and demos at www.lasmea.univ-bpclermont.fr/Personnel/Maxime.Lhuillier).
- [16] R. Malladi, J.A. Sethian, and B.C. Vemuri. Shape modeling with front propagation: A level set approach. *TPAMI*, 17(2):158–175, 1995.
- [17] P.J. Narayanan, P.W. Rander and T. Kanade. Constructing Virtual Worlds using Dense Stereo. *ECCV'98*.
- [18] S. Osher and J.A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
- [19] L. Robert and R. Deriche. Dense depth map reconstruction: a minimization and regularization approach which preserves discontinuities. *ECCV'96*.
- [20] S.M. Seitz and C.R. Dyer. Photorealistic scene reconstruction by voxel coloring. *CVPR'97*.
- [21] J.A. Sethian. *Level-set methods and fast marching methods*. Cambridge University Press, Cambridge, UK, 1999.
- [22] R. Szeliski, D. Tonnesen, and D. Terzopoulos. Modelling surfaces of arbitrary topology with dynamic particles. *CVPR'93*.
- [23] R. Szeliski. Rapid octree construction from image sequences. *Computer Vision, Graphics and Image Processing*, 58(1):23–32, July 1993.
- [24] C.K. Tang and G. Medioni. Curvature-augmented tensor voting for shape inference from noisy 3d data. *TPAMI*, 24(6):858–864, 2002.
- [25] R. Whitaker. A level-set approach to 3d reconstruction from range data, *IJCV*, 29(3):203–231, 1998.
- [26] H.K. Zhao, S. Osher, B. Merriman, and M. Kang. Implicit and non-parametric shape reconstruction from unorganized data using a variational level set method. *CVIU*, 2000.
- [27] H.K. Zhao, S. Osher, and R. Fedkiw. Fast surface reconstruction using the level set method. *Proc. of IEEE Workshop on Variational and Level Set Methods in Computer Vision*, 2001.




Figure 3. Surfaces obtained with the BR2D method after 400 (top) and 1000 (bottom) iterations with $w_0 = 0.1$ (left), $w_0 = 0.5$ (middle), and $w_0 = 1$ (right).




Figure 6. Surfaces obtained with BR3D in (a), with BR3D+S in (b), with BR3D+2D in (c), and with S in (d).




Figure 4. One original image (left), surfaces obtained with BR3D (middle), and BR3D+BR2D (right) with $w_0 = 0.5$.




Figure 5. Surfaces obtained with curvature-based (left) and Laplacian-based (right) smoothing for BR3D method.




Figure 7. Surface initialization (a), computed by BR3D method (b), BR3D+BR2D with $w_0 = 0.1$ (c), BR3D+BR2D with $w_0 = 0.5$ (d), and BR3D+BR2D with $w_0 = 1$ (e).