

Textured mesh surface reconstruction of large buildings with multi-view stereo

Chen Zhu · Wee Kheng Leow

© Springer-Verlag Berlin Heidelberg 2013

Abstract There are three main approaches for reconstructing 3D models of buildings. Laser scanning is accurate but expensive and limited by the laser's range. Structure-from-motion (SfM) and multi-view stereo (MVS) recover 3D point clouds from multiple views of a building. MVS methods, especially patch-based MVS, can achieve higher density than do SfM methods. Sophisticated algorithms need to be applied to the point clouds to construct mesh surfaces. The recovered point clouds can be sparse in areas that lack features for accurate reconstruction, making recovery of complete surfaces difficult. Moreover, segmentation of the building's surfaces from surrounding surfaces almost always requires some form of manual inputs, diminishing the ease of practical application of automatic 3D reconstruction algorithms. This paper presents an alternative approach for reconstructing textured mesh surfaces from point cloud recovered by patch-based MVS method. To a good first approximation, a building's surfaces can be modeled by planes or curve surfaces which are fitted to the point cloud. 3D points are resampled on the fitted surfaces in an orderly pattern, whose colors are obtained from the input images. This approach is simple, inexpensive, and effective for reconstructing textured mesh surfaces of large buildings. Test results show that the reconstructed 3D models are sufficiently accurate and realistic for 3D visualization in various applications.

Keywords Building reconstruction · Multi-view stereo · Mesh surface reconstruction

1 Introduction

Reconstruction of 3D models of buildings has many applications. It can be used to create 3D models for virtual reality, games, and digital special effects. It can also be used for preserving heritage architectures and city planning. There are three main approaches for reconstructing 3D models of buildings. Laser scanning [1, 3, 7, 11] is accurate but expensive and limited by the laser's range. Structure-from-motion (SfM) [4–6, 12, 13, 21, 22] and multi-view stereo (MVS) [8, 10, 15, 18] recover 3D point clouds from multiple views. In particular, patch-based MVS (PMVS) [8, 10, 15, 16] can achieve higher density than do SfM methods. Since the point clouds recovered by SfM and MVS are unordered, sophisticated algorithms [8, 10, 15] need to be applied to the point clouds to construct mesh surfaces. Moreover, the recovered point clouds can be sparse in areas that lack features for accurate reconstruction, making recovery of complete surfaces difficult.

Regardless of the technology used, reconstruction of an entire building cannot be achieved in practice by a single laser scan or single point cloud of a building. It is always necessary to repeat the process to reconstruct and merge various parts of the building, making laser scanning more cumbersome to apply. Moreover, segmentation of the building's surfaces from surrounding surfaces almost always requires some form of manual inputs, diminishing the ease of practical application of automatic 3D reconstruction algorithms.

This paper presents an alternative approach for reconstructing textured mesh surfaces from point cloud recovered

C. Zhu · W.K. Leow (✉)
Department of Computer Science, National University
of Singapore, Computing 1, 13 Computing Drive,
Singapore 117417, Singapore
e-mail: leowwk@comp.nus.edu.sg

C. Zhu
e-mail: zhuchen@comp.nus.edu.sg

by patch-based MVS method. It regards a building's surfaces as effectively flat surfaces which can be modeled by planes or curve surfaces. This is a good first approximation when the scene depth is much smaller than the distance of the building to the camera, which is true of many modern architectural design. In this case, it is not necessary to apply sophisticated algorithms to reconstruct detailed mesh surfaces from point cloud. Instead, simple surfaces are fitted to the point cloud, and 3D points are resampled on the fitted surfaces in an orderly pattern, whose colors are obtained from the input images. In this way, complete textured mesh of building surfaces can be reconstructed. This approach is thus simple, inexpensive, and effective for reconstructing textured mesh surfaces of large buildings.

2 Related work

Many methods have been proposed over the past decades for the acquisition of 3D models of objects [19]. Among these methods, laser scanning, structure-from-motion (SfM), and multi-view stereo (MVS) have been successfully demonstrated for 3D reconstruction of buildings and street scenes. Laser scanning is very accurate and efficient, but requires an expensive laser scanner, which is limited by the laser's range. In reconstructing large historic sites, laser scanning is often used together with other data such as engineering drawing, aerial images, ground-plane images, etc. [1, 3, 7, 11].

SfM methods recover 3D data and camera positions and orientations from a sequence of images [4–6, 12, 13, 21, 22]. The camera parameters can be calibrated explicitly or computed by self-calibration methods.

MVS methods apply various geometric constraints to determine point correspondence for triangulation of 3D points given known camera parameters. In particular, patch-based MVS (PMVS) can recover denser 3D point cloud than do SfM methods [8, 10, 15, 16]. SfM and MVS may be used together in a pipeline with MVS serving as the refinement stage [5, 17]. Sophisticated algorithms are required to reconstruct mesh surfaces from point cloud. For example, [10] uses Poisson surface reconstruction method and [15] employs graph cut algorithm with geometric constraints.

3 Textured surface reconstruction algorithm

Our proposed algorithm continues from where PMVS ends. It consists of the following main stages:

1. Recover a 3D point cloud of a building using PMVS.
2. Reconstruct main surfaces from 3D point cloud by robust surface fitting (Sect. 3.1) and splitting (Sect. 3.2).
3. Resample 3D points to refine surfaces (Sect. 3.3).

First, the Bundler algorithm [20] is applied to the input images to extract matching feature points and camera parameters. Next, the PMVS algorithm [9, 10] is applied to recover a 3D point cloud with color information (Fig. 5, 6). Then, robust surface splitting is performed on the 3D point cloud to split it into multiple parts each corresponding to a single surface, and the surfaces are reconstructed by robust surface fitting. Next, 3D points are resampled over the entire fitted surfaces and their colors are obtained from the input images. Finally, a 3D mesh model that is complete with color texture is constructed from the resampled 3D points to represent the building surfaces.

For a large building, a single point cloud cannot cover the full extent of the building. In this case, multiple point clouds are recovered for different parts of the building, and our algorithm is applied to reconstruct different parts of the building's surfaces, which are aligned and merged together.

3.1 Robust surface fitting

First, we consider the case of fitting a single surface to 3D points. Without loss of generality, we assume that the surface's normal is not perpendicular to the Z -axis, a degenerate case of our algorithm. For a degenerate case, we can rotate the surface to align its normal to the Z -axis, perform the computation, and then rotate the solution back to the surface's original orientation.

A plane in 3D space can be defined by the equation

$$Z = S(X, Y) = a_1X + a_2Y + a_3. \tag{1}$$

Given a set of n 3D points $\mathbf{X}_i = (X_i, Y_i, Z_i)$, the parameters (a_1, a_2, a_3) of the plane can be recovered as the linear least square solution of the equation:

$$\begin{bmatrix} X_1 & Y_1 & 1 \\ \vdots & \vdots & \vdots \\ X_n & Y_n & 1 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix} = \begin{bmatrix} Z_1 \\ \vdots \\ Z_n \end{bmatrix}. \tag{2}$$

A curved surface can be recovered in a similar way. For example, a quadratic surface of the form

$$Z = a_1X^2 + a_2Y^2 + a_3XY + a_4X + a_5Y + a_6 \tag{3}$$

can be recovered as the linear least square solution of the equation:

$$\begin{bmatrix} X_1^2 & Y_1^2 & X_1Y_1 & X_1 & Y_1 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ X_n^2 & Y_n^2 & X_nY_n & X_n & Y_n & 1 \end{bmatrix} \begin{bmatrix} a_1 \\ \vdots \\ a_6 \end{bmatrix} = \begin{bmatrix} Z_1 \\ \vdots \\ Z_n \end{bmatrix}. \tag{4}$$

Other forms of curved surfaces may be recovered in the same manner. Our current implementation restricts to planar and quadratic surfaces which are sufficient for the test cases. The type of surface is specified by the user.

The surface recovered in this way is severely affected by the outliers in the point cloud. It is too tedious to manually identify the inliers and exclude the outliers. So, a robust




Fig. 1 Robust surface fitting. (a) Convergence curve. (b) 3D points around a curved surface. (c) Inliers identified by the robust fitting algorithm lie very close to the fitted surface

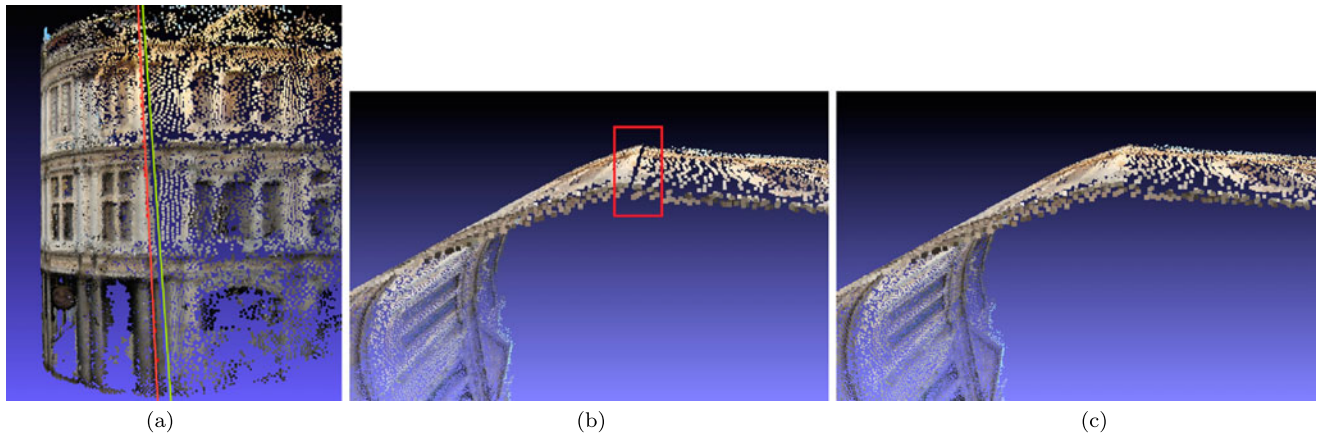


Fig. 2 Splitting and alignment of surfaces. (a) Splitting of point cloud. (Green) Initial splitting line. (Red) Refined splitting line. (b) Before alignment, there is a gap (in the red box) between the two surfaces. (c) After alignment, the gap is removed

surface fitting algorithm is adopted to automatically identify the inliers. The algorithm adopts a robust error measure $E = \text{median}_i r_i$, where r_i is the residual error of a point \mathbf{X}_i with respect to the fitted surface:

$$r_i = (Z_i - S(X_i, Y_i))^2. \quad (5)$$

It iteratively identifies and excludes outliers from the point cloud. The robust surface fitting algorithm can be summarized as follows:

Robust surface fitting

1. Initialize $P \leftarrow$ input 3D point cloud, $E \leftarrow \infty$.
2. Repeat while $E > \tau_s$:
 - (a) Fit a surface S to the points in P .
 - (b) Compute robust error E of S on P .
 - (c) Remove from P the points with residual $r_i \geq E$.

Empirical tests show that a threshold of $\tau_s = 5 \times 10^{-7}$ yields good results. Figure 1 shows that the algorithm can converge within a small number of iterations and the solution has very small error. Although a standard robust algorithm such as RANSAC may be used, the above algorithm is much more efficient than RANSAC. It is stable and accurate because of

the large number of inliers available in the point cloud for accurate fitting.

3.2 Robust surface splitting

The above algorithm fits a single surface to a 3D point cloud. Now, we illustrate how to split the point cloud into multiple parts for fitting different surfaces.

Two non-parallel surfaces of the form $Z = S_1(X, Y)$ and $Z = S_2(X, Y)$ intersect along a line or, more generally, a curve. Define $D(X, Y)$ as the absolute difference between the two surfaces:

$$D(X, Y) = |S_1(X, Y) - S_2(X, Y)|. \quad (6)$$

Then, the equation $D(X, Y) = 0$ gives the intersection of the surfaces, and $D(X, Y)$ measures the distance of a surface point \mathbf{X} to the intersection.

Let us project the 3D data points along the Z -axis onto the X - Y plane, and let $l(\mathbf{X})$ denote the equation of a line manually drawn on the X - Y plane to divide the data points into two parts (Fig. 2(a)):

$$l(\mathbf{X}) = b_1X + b_2Y + b_3 = 0. \quad (7)$$

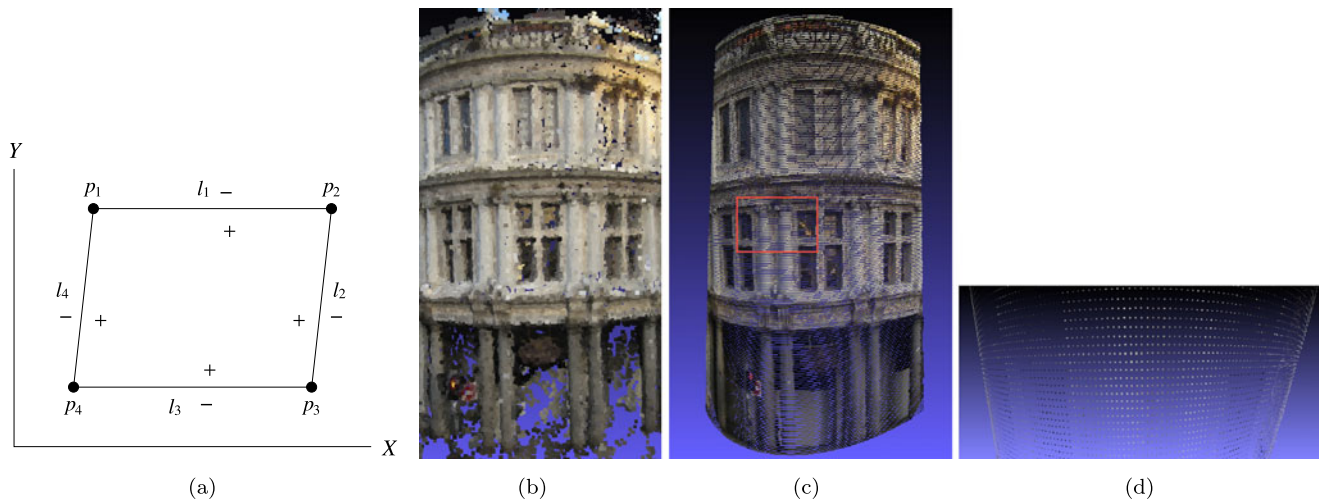


Fig. 3 Resampling of surface points. (a) Resampling region. The lines of the resampling region are computed such that points in the region have positive signs. (b) PMVS point cloud is sparse in some regions.

(c) Resampled 3D points cover the surface completely. (d) Zoomed-in view of the region in the red box

Then, points \mathbf{X}_i not lying on the line have non-zero values $l(\mathbf{X}_i)$. So, a set of 3D points \mathbf{X}_i can be split into two subsets according to the sign of $l(\mathbf{X}_i)$. After splitting, a surface can be fitted to each of the two subsets using the robust surface fitting algorithm. The intersection of these surfaces induces a change of $l(\mathbf{X})$, and the whole process can be iterated to obtain optimal splitting. So, the robust surface splitting algorithm can be summarized as follows:

Robust surface splitting

1. Let P be a 3D point cloud to be split and l be the initial splitting line.
2. Repeat:
 - (a) Split P into P_1 and P_2 according to the sign of $l(\mathbf{X}_i)$, for all points $\mathbf{X}_i \in P$.
 - (b) Perform robust fitting of surface S_1 to P_1 and S_2 to P_2 .
 - (c) Compute $D(\mathbf{X}_i)$ of each point \mathbf{X}_i and select a subset Q of points with the smallest $D(\mathbf{X}_i)$, which are near the intersection.
 - (d) Fit l to the points in Q using linear least square method.

This algorithm is iterated for a fixed number of iterations to split the input point cloud and reconstruct surfaces. In the current implementation, the size of Q is 50, and empirical tests show that two iterations are sufficient. The same algorithm is repeatedly applied to different parts of the point cloud to robustly split and reconstruct multiple surfaces of a building. To define the boundaries of a quadrilateral surface, four splitting lines need to be manually drawn. For a surface with more complex boundaries, more splitting lines are needed.

As the surfaces are fitted to their 3D points independently, they may not join perfectly at the intersection, resulting in the appearance of a gap or seam (Fig. 2(b)). This problem is resolved by aligning the surfaces as follows. Recall that Q is the set of 3D points near the intersection. Then, for each point \mathbf{X}_i in Q , we can obtain the corresponding points \mathbf{X}_{1i} and \mathbf{X}_{2i} on the surfaces S_1 and S_2 as follows:

$$\begin{aligned} X_{1i} = X_{2i} = X_i, & \quad Y_{1i} = Y_{2i} = Y_i, \\ Z_{1i} = S_1(X_i, Y_i), & \quad Z_{2i} = S_2(X_i, Y_i). \end{aligned} \tag{8}$$

Given the corresponding points \mathbf{X}_{1i} and \mathbf{X}_{2i} , we apply the algorithm of [14] to compute the best similarity transformation to align the surfaces (Fig. 2(b)).

3.3 Resampling of surface points

The 3D point cloud recovered by PMVS typically does not cover a surface completely. To obtain complete color texture for a surface, it is necessary to resample the color information from the input images.

First, the set P of 3D points of a surface are projected onto the X - Y plane. Next, corner points of the desired resampling region are manually marked on the X - Y plane. A line segment $l(\mathbf{X})$ is computed between two connected corners such that $l(\mathbf{X})$ is positive for points inside the sampling region and negative outside (Fig. 3(a)). This property holds for all convex regions and allows for easy decision of whether a point is in the resampling region. Finally, 3D points are sampled at regular spacing within the resampling region using the equation of the surface (Eqs. (1), (3); Fig. 3). The resampling rate is determined by the user according to the resolution required and it is independent of the sampling density of the point cloud.




Fig. 4 Removal of distorted color texture. (a) Color texture with gross distortion. (b) Retrieving colors from a frontal-view image removes distortion but induces a visible seam. (c) Color blending removes the seam

Table 1 Results of mesh surface reconstruction. Execution times are measured in minutes, excluding manual inputs

Building	No. of images	No. of point clouds	No. of surfaces	Bundler run time	PMVS run time	Our algo run time	Total run time
1	27	1	3	20	12	6	38
2	105	2	3	79	47	14	140

The colors of the resampled surface points are obtained from the images. Among the multiple views of a building, the image that corresponds to the frontal view of a building surface is used. Let $\tilde{\mathbf{X}} = (X, Y, Z, 1)^\top$ and $\tilde{\mathbf{x}} = (x, y, 1)^\top$ denote the homogeneous coordinates of 3D point \mathbf{X} and image point \mathbf{x} , respectively. Then, the 3D point projects to the image point according to the perspective projection equation:

$$\rho \tilde{\mathbf{x}} = \mathbf{P} \tilde{\mathbf{X}} \tag{9}$$

where ρ is a scaling parameter and \mathbf{P} is the projection matrix computed by PMVS algorithm. Denoting each row of \mathbf{P} as \mathbf{P}_k^\top , and rearranging Eq. (9) yields

$$x = \frac{\mathbf{P}_1^\top \mathbf{X}}{\mathbf{P}_3^\top \mathbf{X}}, \quad y = \frac{\mathbf{P}_2^\top \mathbf{X}}{\mathbf{P}_3^\top \mathbf{X}}. \tag{10}$$

So, the color of \mathbf{X} can be obtained from the image at pixel location $\mathbf{x} = (x, y)^\top$. Typically, \mathbf{x} has real-valued coordinates. Its color should be derived from those of the four nearest pixels around it using bilinear interpolation.

In some cases, an image may contain building surfaces at oblique view. Retrieval of pixel colors from the oblique view may result in gross distortion (Fig. 4a). To correct for the distortion, the image that presents the surface at frontal view should be used. When multiple images are used for different parts of a surface, seams can appear at the intersections of the images. In this case, color blending should be applied across the seams to remove them (Fig. 4c). Finally, 3D mesh model of the recovered surfaces is constructed by applying

ball-pivoting algorithm [2]. Ball-pivoting algorithm is adequate in producing a good mesh because the 3D points are resampled densely at regular spacing.

4 Test results and discussions

Two large buildings with curve surfaces were used as the test cases: one with a convex curved surface, the other with a concave curved surface. Multiple images of the buildings were taken from various viewing angles. For each test case, PMVS algorithm was executed to recover 3D point clouds, and our surface reconstruction algorithm was executed to reconstruct the surfaces.

Table 1 tabulates the results of running the algorithms. Building 2 had an extended wall and required two separate point clouds to capture the building's surfaces. Bundler and PMVS took significant amounts of time to compute the matching feature points, camera parameters, and point clouds.

Figure 5 and 6 show samples input images, point clouds recovered by PMVS, and reconstructed textured mesh models of the buildings. Notice that the point clouds are sparse in some regions of the surfaces. Nevertheless, our algorithm can resample the color textures in those regions from the input images and reconstruct complete textured mesh of the surfaces.

In the current implementation, our algorithm does not differentiate between the buildings and the occluders, such as the trees, in front of the buildings. So, the occluders are



Fig. 5 Reconstructed mesh models of building 1. (Row 1) Sample images from various views. (Row 2) Point cloud recovered by PMVS and (Row 3) reconstructed textured mesh model at various viewing angles



Fig. 6 Reconstructed mesh models of building 2. (Row 1) Sample images from various views. (Row 2) Point clouds recovered by PMVS and (Row 3) reconstructed textured mesh model at various viewing angles

regarded as part of the buildings' textures. To remove the occluders, it is necessary to capture the images with the camera located in between the buildings and the occluders. With the camera located at a close distance from the buildings, it is necessary to construct the building's surfaces in multiple parts and then merge them together. This is technically possible but practically tedious to perform.

5 Conclusions

This paper presented a simple, inexpensive, and effective method for reconstructing textured mesh surfaces of large buildings with curved surfaces. The method applies PMVS to recover point clouds from multiple images of a building. Then, robust surface splitting and fitting algorithms are applied to fit multiple surfaces to different parts of the point clouds. These surfaces are then aligned, merged and color blended to produce a single textured mesh model of the building. The mesh model is already segmented from the surrounding and can be used directly in various applications. Test results show that the building models reconstructed by our algorithm are sufficiently accurate and realistic for 3D visualization in various applications.

References

- Allen, P.K., Troccoli, A., Smith, B., Murray, S., Stamos, I., Leordeanu, M.: New methods for digital modeling of historic sites. *IEEE Comput. Graph. Appl.* **23**(6), 32–41 (2003)
- Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C., Taubin, G.: The ball-pivoting algorithm for surface reconstruction. *IEEE Trans. Vis. Comput. Graph.* **5**(4), 349–359 (1999)
- Blais, F.: Review of 20 years of range sensor development. *J. Electron. Imaging* **13**(1), 232–240 (2004)
- Cornelis, N., Cornelis, K., Van Gool, L.: Fast compact city modeling for navigation pre-visualization. In: Proc. CVPR (2006)
- Cornelis, N., Leibe, B., Cornelis, K., van Gool, L.J.: 3D urban scene modeling integrating recognition and reconstruction. *Int. J. Comput. Vis.* **78**(2–3), 121–141 (2008)
- Dellaert, F., Seitz, S.M., Thorpe, C.E., Thrun, S.: Structure from motion without correspondence. In: Proc. IEEE CVPR (2000)
- Frueh, C., Zakhor, A.: Constructing 3D city models by merging aerial and ground views. *IEEE Comput. Graph. Appl.* **23**(6), 52–61 (2003)
- Furukawa, Y., Curless, B., Seitz, S., Szeliski, R.: Towards Internet-scale multi-view stereo. In: Proc. CVPR (2010)
- Furukawa, Y., Ponce, J.: Patch-based multi-view stereo software. <http://www.di.ens.fr/pmvs/>
- Furukawa, Y., Ponce, J.: Accurate, dense, and robust multi-view stereopsis. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(8), 1362–1376 (2010)
- Guidi, G., Micoli, L., Russo, M., Frischer, B., De Simone, M., Spinetti, A., Carosso, L.: 3D digitization of a large model of imperial Rome. In: Proc. Int. Conf. 3D Imaging and Modeling, pp. 565–572 (2005)
- Havlena, M., Torii, A., Knopp, J., Pajdla, T.: Randomized structure from motion based on atomic 3D models from camera triplets. In: Proc. CVPR (2009)
- Havlena, M., Torii, A., Pajdla, T.: Efficient structure from motion by graph optimization. In: Proc. ECCV (2010)
- Horn, B.K.P., Hilden, H.M., Negahdaripour, S.: Closed-form solution of absolute orientation using orthonormal matrices. *J. Opt. Soc. Am. A* **5**(7), 1127–1135 (1988)
- Jancosek, M., Pajdla, T.: Multi-view reconstruction preserving weakly-supported surfaces. In: Proc. CVPR (2011)
- Labatut, P., Pons, J.P., Keriven, R.: Efficient multi-view reconstruction of large-scale scenes using interest points, Delaunay triangulation and graph cuts. In: Proc. CVPR (2007)
- Liang, Y., Sheng, Y., Song, H., Wang, Y.: 3D Reconstruction of Architecture from Image Sequences. World, Cleveland (2013)
- Mičušík, B., Košecák, J.: Piecewise planar city 3D modeling from street view panoramic sequences. In: Proc. CVPR (2009)
- Moons, T., Van Gool, L., Vergauwen, M.: 3D reconstruction from multiple images. Part 1. Principles. *Found. Trends Comput. Graph. Vis.* **4**(4), 287–398 (2008)
- Snavely, N., Bundler: Structure from motion (SfM) for unordered image collections. <http://phototour.cs.washington.edu/bundler/>
- Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring image collections in 3D. In: Proc. SIGGRAPH (2006)
- Xiao, J., Fang, T., Zhao, P., Lhuillier, M., Quan, L.: Image-based street-side city modeling. In: Proc. SIGGRAPH Asia (2009)



Chen Zhu obtained his B.Comp. in Computer Science from National University of Singapore in 2012. He received the title of Ingénieur diplômé de l'École Polytechnique in France in 2011. His current research interest is 3D Building Reconstruction.



Wee Kheng Leow obtained his B.Sc. and M.Sc. in Computer Science from National University of Singapore in 1985 and 1989, respectively. He pursued Ph.D. study at The University of Texas at Austin and obtained his Ph.D. in Computer Science in 1994. He is currently an Associate Professor at the Department of Computer Science, National University of Singapore. His research interests include computer vision, medical image analysis and surgery planning and simulation.