# Binary-Space-Partitioned Images for Resolving Image-Based Visibility

Chi-Wing Fu, Student Member, IEEE Computer Society, Tien-Tsin Wong, Member, IEEE Computer Society,
Wai-Shun Tong, Student Member, IEEE Computer Society, Chi-Keung Tang, Member, IEEE Computer Society, and Andrew J. Hanson, Member, IEEE Computer Society

Abstract—We propose a novel 2D representation for 3D visibility sorting, the *Binary-Space-Partitioned Image (BSPI)*, to accelerate real-time image-based rendering. BSPI is an efficient 2D realization of a 3D BSP tree, which is commonly used in computer graphics for time-critical visibility sorting. Since the overall structure of a BSP tree is encoded in a BSPI, traversing a BSPI is comparable to traversing the corresponding BSP tree. BSPI performs visibility sorting efficiently and accurately in the 2D image space by warping the reference image triangle-by-triangle instead of pixel-by-pixel. Multiple BSPIs can be combined to solve "disocclusion," when an occluded portion of the scene becomes visible at a novel viewpoint. Our method is highly automatic, including a tensor voting preprocessing step that generates candidate image partition lines for BSPIs, filters the noisy input data by rejecting outliers, and interpolates missing information. Our system has been applied to a variety of real data, including stereo, motion, and range images.

Index Terms—Visibility sorting, binary space partitioning, image-based rendering, segmentation.

#### **1** INTRODUCTION

Real-TIME rendering of complex scenes has long been a challenge for computer graphics. Image-based modeling and rendering (IBMR) techniques approach this problem by using image data directly to decouple the rendering time from the scene complexity. Input images acquired from fixed viewpoints are then warped and composited to generate novel views.

The rendering problem then reduces to that of synthesizing acceptable novel views. A typical solution is to warp the image in a pixel-by-pixel manner [13], [31], [51]. A potentially faster way is to exploit image space coherence by triangulating the image and warping the unstructured triangles instead of pixels. We show in this paper that, by subdividing the image using 2D binary space partitioning (BSP), we obtain a hierarchical image representation, namely, the *binary-space-partitioned image* (*BSPI*), that allows us to perform real-time visibility sorting entirely in the 2D image space.

The BSP representation has long been exploited in graphics systems, especially for game applications [1], due to its efficiency in solving the visibility [22], [21], [42], collision detection, and shadow generation [14], [15] problems.

- T.-T. Wong is with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Shatin, Hong Kong. E-mail: ttwong@acm.org.
- W.-S. Tong and C.-K. Tang are with the Department of Computer Science, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong. E-mail: {cstws, cktang}@cs.ust.hk.

Manuscript received 13 June 2002; revised 2 Dec. 2002; accepted 5 Dec. 2002. For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number 116751.

Methods for generating good BSP trees are typically studied as a part of computational geometry (see, e.g., [6], [44], [2], [16], [46], [61], [9]). On the other hand, the BSP representation has also been used to encode two-dimensional images [50], [48], [49] or even volume data [55]. In this paper, we show that the 2D BSPI is an efficient approach to visibility computation for image-based rendering.

The primary contribution of this paper is twofold:

- A hierarchical image-based representation, the *Binary-Space-Partitioned Image (BSPI)*, is proposed to efficiently encode data coherence in image space and to resolve visibility in linear time. The 2D BSPI is actually an embedding of a 3D BSP tree. Hence, visibility sorting can be performed to correctly and efficiently render novel images by traversing the embedded BSP tree. To handle scenes with occlusions, multiple BSPIs can be used.
- By using *epipolar geometry*, we show that visibility sorting can be performed completely in 2D (rather than 3D) without using any 3D information such as depth or disparity. No depth buffering [10] is required.

The secondary contribution consists of automating the whole process, improving the rendering quality, and also providing the capability to handle real and noisy data obtained from stereo, motion, and sensing devices:

• We make use of *tensor voting* [40] to automatically refine noisy reference images and segment them into a tree structure based on region partition boundaries. Tensor voting determines the edge locations by separating coherent pixel groups, thus subdividing the image into triangles rather than pixels. This is

<sup>•</sup> C.-W. Fu and A.J. Hanson are with the Computer Science Department, Indiana University, Bloomington, IN 47405. E-mail: {cwfu, hanson}@cs.indiana.edu.



Fig. 1. An example of a binary space partitioning. (a) The scene (top view). (b) The BSP tree yielding the visibility order corresponding to the eye in (a).

achieved by the robust detection of discontinuity curves [60].

 The BSPI is a unified approach for a variety of image data, including range images, stereo image pairs, and image sequences. To demonstrate its wide applicability and practicability, all examples used in this paper are real rather than synthetic data.

Our work is closely related to the work by Subramanian and Naylor [55] on representing images using a BSP tree. The focus of our work is on resolving visibility in image space.

#### 2 RELATED WORK

#### 2.1 Visibility in Geometry-Based Computer Graphics

Hidden-surface removal is a classical problem in computer graphics. Analytic studies of visibility algorithms have been carried out in [37], [41]. Related methods can be roughly subdivided into three categories: image-space, object-space, and hybrid methods.

*Image-space* methods resolve visibility after projecting objects onto the screen. Depth buffering [10], [29] keeps track of a per-pixel depth value to determine the closest object fragment. It is the most popular approach used in current graphics hardware design. On the other hand, while scanline algorithms [65], [7], [8], [64], [52], [17] exploit scanline coherence, the required data structure is relatively complex.

*Object-space* algorithms resolve visibility in the 3D object space. The 3D scene is first partitioned and represented by various hierarchical structures such as cells [3], [59], [58], octrees [26], [56], kd-trees [33], or BSP trees [22], [21], [12] in a preprocessing stage. During rendering, the visibility is resolved by traversing the hierarchical structures. A back-to-front ordering can also be generated in some cases [45], [54], [22], [21] and fed to a Painter's algorithm for rendering. Moreover, the hierarchical nature of these representations also facilitates fast culling of invisible objects in a complex scene in order to achieve interactive frame rates [3], [23], [24], [59], [58].

Among these structures, the BSP tree is widely used due to its robustness and efficiency. Our proposed BSPI method can be regarded as an image-based version of the classical BSP-tree-based visibility algorithm. For a brief description of BSP trees, consider the example in Fig. 1. In each iteration, a polygon is chosen as the root node of the current binary subtree. It is then extended to form a partition plane (hyperplane) that divides the scene into two half-spaces. The resultant BSP tree is shown in the upper part of Fig. 1b. In general, there is no restriction on the selection of partition planes. The plane need not contain a polygon. To render the view at the eye position in Fig. 1a, a variant of the inorder traversal (O(n) running time where n is number of nodes in the tree) of the BSP tree can be used to generate the drawing order from back to front (Fig. 1b, bottom). Note that this back-to-front drawing order is essential for rendering transparent objects correctly, which depth buffering cannot accommodate. In addition, the algorithm is robust and view independent. Its efficiency explains why it is favored in the design of real-time computer games [1]. Variants can also be applied to shadow generation [14], [15], surface modeling [42], and visibility sorting in dynamic scenes [61].

*Hybrid methods* utilize information in both image space and object space. Zhang et al. [66] used both the bounding volume in object space and the occlusion map in image space to compute visibility. By combining hierarchical polygon tiling and a hierarchical visibility algorithm, Greene [28] developed a fast algorithm, based on Warnock's work [63], to render highly complex scenes. Naylor [43] proposed a 2D BSP image tree to represent the 2D projection of 3D polyhedral models in a continuous form. The representation proposed in his paper resembles ours since we also construct a 2D BSP tree. While visibility in Naylor's approach is computed in 3D object space, we perform visibility sorting entirely in 2D image space using 2D input images rather than 3D geometric models.

#### 2.2 Visibility in Image-Based Computer Graphics

IBMR systems acquire images as inputs. In light field rendering [34] and the lumigraph method [27], the visibility problem is treated as an interpolation problem. When the image is equipped with depth, visibility issues can be resolved with depth-buffering [13]. If the images (sprites) are also associated with underlying 3D geometric models, visibility sorting [33], [54] can be performed in object space. Sometimes images are associated with coarse geometric primitives [18]. In this case, visibility can be computed with a standard depth-buffering algorithm.

However, there is no guarantee that the depth information or geometric models will always be available. Alternatively, disparity or depth can be recovered [31], [19] from stereo pairs or motion data. This information cannot be



Fig. 2. The drawing order between two pixels that lie on the same epipolar line.

directly used in the geometry-based visibility algorithms above to resolve visibility. In some cases, an intermediate surface model (*pseudo-object*), has to be reconstructed, introducing extra overhead. An alternate approach is to directly use the disparity to approximate the true visibility solution using a tailor-made rendering engine [35].

Starting from basic epipolar geometry, McMillan and Bishop [38] found that, even when the true depth is not available, visibility-correct novel views can still be generated if the epipole is known. Fig. 2 illustrates the basic idea. Positions c and e are the reference and the desired viewpoints, respectively. Reference image  $I_c$  contains two pixels,  $i_1$  and  $i_2$ , and their corresponding points in 3D space are  $p_1$  and  $p_2$ , respectively. The projection of e onto the reference image is the epipole (positive in this example, denoted by a plus sign). When we generate the novel view at e, even though exact positions of  $p_1$  and  $p_2$  may not be known, the correct visibility can still be ensured if we always draw  $i_1$  before  $i_2$ . Notice that  $\mathbf{p}_1$  cannot occlude  $\mathbf{p}_2$ since  $i_1$  is farther away from the positive epipole than  $i_2$ ; thus, we can determine an occlusion-compatible drawing order for pixels from the position of epipole even though no absolute depth information is available.

The same visibility algorithm is also used in later imagebased rendering systems [53], [11], [36]. However, since this original occlusion-compatible drawing order is only applicable to pixel-sized image entities, the rendering is less efficient than it could be. Fu et al. [20] extended this drawing-order strategy to arbitrarily sized triangles and made use of texture-mapped triangles to speed up rendering. However, cycles may exist in the triangle drawing order in some rare cases, leading to incorrect visibility determination. The underlying reason is that the trianglebased ordering rule cannot guarantee a partial order.

#### **3** OVERVIEW

Our system overview is shown in Fig. 3. The input consists of either a range image, a stereo pair, or a video sequence. A BSPI is then produced and rendered at novel views. The major steps of the system are as follows:

- 1. Automatic preprocessing estimates and cleans up the given depth or disparity maps (Section 5.1).
- 2. Feature extraction by tensor voting (Section 5.2).
- 3. BSPI generation (Section 5.3).
- 4. BSPI rendering (Section 5.4).

The rest of this paper is organized as follows: In Section 4, we first explain how the BSPI ensures correct hidden surface removal, the main theoretical contribution of this paper. We detail the four steps in Section 5. Finally, we present results on a variety of input images in Section 6. A review of the tensor voting method is given in the Appendix.

# 4 THE BINARY-SPACE-PARTITIONED IMAGE REPRESENTATION

IBMR research motivates us to study the efficient rendering of image-based data that may not contain geometry information. To synthesize the desired images, we need to



Fig. 3. Overall approach: (a) the flowchart and (b) a running example ANGEL. (Input image and initial depth map courtesy of the Signal Analysis and Machine Perception Laboratory, The Ohio State University).



Fig. 4. Forming the partition plane. (a) The reference image and (b) the partition plane formed by connecting the reference viewpoint and an edge in the image.

solve two subproblems: 1) where image pixels move to and 2) which pixel is in the front if more than one pixel moves to the same place.

Subproblem 1) can be solved if depth information is available. The new pixel positions can be calculated, for example, by reprojection [13]. This problem can also be solved if the disparity of each pixel is known. One can warp the reference image pixel-by-pixel [38], [53], [11], [36] or triangle-by-triangle [20] while exploiting pixel coherence.

Subproblem 2) also can be solved if accurate depth information is known. Depth-buffering is the simplest method. However, efficient rendering of high-resolution images depends on the performance of graphics hardware as a per-pixel depth comparison is needed. If only disparity information can be recovered, the usual way is to fit a 3D surface (pseudo-object) and use this geometric surface for rendering. Recovering and rendering this intermediate pseudo-object introduces additional overhead.

In this paper, we focus on subproblem 2), the visibility problem, with the assumption that subproblem 1) can be handled by pixel reprojection or disparity interpolation. An efficient algorithm based on the BSP visibility algorithm is proposed. No intermediate pseudo-object needs to be constructed and the algorithm is also independent of the resolution of the final desired image. It operates on convex polygonal primitives (not necessarily only triangles), hence taking advantage of the coherence between adjacent pixels. The efficiency of the BSP algorithm is evidenced by its common applications in computer game kernels [1], despite the wide availability of depth-buffer-based graphics hardware. We will show in the following sections that the proposed algorithm preserves the efficiency of the original geometry-based BSP algorithm and achieves real-time visibility sorting.

#### 4.1 BSP in 2D

Every image reveals the visible surfaces seen from its corresponding viewpoint. Fig. 4a shows the visible surfaces of a cube. Taking a 2D edge in the image plane, we can create an extended partition plane containing the reference viewpoint (the eye) and the edge (the bold line) on the image plane (Fig. 4b). There is a one-to-one correspondence between the partition plane and the 2D edge on the reference image. They are projection-equivalent with respect to the reference viewpoint. By recursively partitioning the 2D reference image into 2D convex polygons (Fig. 5),



Fig. 5. Recursive partition of the 2D reference image.

we can create a corresponding BSP tree in 3D. We call the 2D binary-space-partitioned image the *BSPI*.

Note that the original BSP visibility algorithm does not restrict the selection of partition planes. What we have done in the BSPI is to constrain all partition planes to intersect at the reference viewpoint, as illustrated in Fig. 6, which depicts the top view of a simplified scene. The gray region indicates the space inside the field of view. The thick lines radiating from the reference viewpoint are the partition planes. The BSPI representation contains an implicitly embedded 3D BSP tree and, hence, it inherits the corresponding visibility sorting capability.

#### 4.2 Visibility Sorting in 2D

Given the desired viewpoint as shown in Fig. 6, we need to warp the reference image to generate the related novel image. Since we have subdivided the reference image into a set of 2D convex polygons, warping is a process of moving and distorting these convex polygons. To ensure correct visibility, we can derive a drawing order from the BSPI. Just like traditional 3D BSP visibility, the visibility drawing order can be derived by traversing the BSP tree embedded in the BSPI. Recall that, during the tree traversal, we need to check to see on which side of the partition plane the desired viewpoint is found. Interestingly, we can carry out all such checks in 2D by projecting the desired viewpoint onto the reference image plane. The projection is the *epipole*.

To illustrate the idea, we show the 2D analogy (Fig. 7). In 2D, the 3D partition plane degenerates to a partition line (the thick blue line in Fig. 7). Similarly, the image plane degenerates to a line (symbolized by a horizontal line in Fig. 7). The desired viewpoint e is projected onto the reference image plane. In Fig. 7a, the projection is a positive epipole, since the desired viewpoint is in front of the



Fig. 6. Partition planes (thick dark (blue) lines) are all constrained to pass through the reference viewpoint (top view). The gray region indicates the space inside the field of view.



Fig. 7. Positive and negative epipoles defined relative to the reference viewpoint (center of projection) and the desired viewpoint.

reference viewpoint c.<sup>1</sup> Since all partition planes (lines in 2D) are constrained to intersect at the reference viewpoint, we can always determine the drawing order by checking on which side of the 2D partition lines the positive epipole resides. Those convex polygons (projections of objects on the image plane) on the opposite side of the positive epipole should be drawn first as these objects can never occlude objects on the same side as the positive epipole.

On the other hand, as shown in Fig. 7b, when the desired viewpoint e is "behind" the reference viewpoint c (in the case of the negative epipole), the desired viewpoint e is actually on the other side of the partition plane (the thick blue line) because all partition planes must pass through the reference viewpoint c. Hence, the drawing order is reversed. The rationale here is that, by projecting everything (including the viewpoints and the partitioning lines/planes) onto the image plane, we can effectively compute the visibility order of image space triangles in the 2D image space.

#### 4.3 Comparison with Related BSP Algorithms

Our BSPI method is inspired by McMillan and Bishop's occlusion-compatible drawing order [39], which also makes use of epipolar geometry. In another classic work [4], drawing orders and grid splitting are used to ensure visibility correctness. Table 1 presents a comparison between the properties of the occlusion-compatible drawing algorithm and the BSPI.

#### 5 THE BSPI RENDERING SYSTEM

Now that we have outlined the theory behind the BSPI method, we next describe the four stages of the BSPI system implementation. The system flowchart in Fig. 3 provides an overview, and we use the ANGEL image in that figure as the running example to facilitate discussion.

# 5.1 Preprocessing (Depth Estimation and Noise Filtering)

If depth information is unavailable, we perform a preprocessing step to estimate the scene depth. Note that depth is used in preprocessing only and no depth sorting is needed in subsequent novel-view rendering. We consider two common cases for which we have images but no initial depth information: 1) a stereo pair and 2) an image sequence. Computer vision algorithms for stereo and motion-sequence data can be applied to estimate the corresponding depth maps. Alternatively, if a noisy initial depth map is available (see, e.g., ANGEL's noisy depth data, Fig. 3b), it can be cleaned up by rejecting outliers and inferring missing details (see, e.g., ANGEL's refined depth map, Fig. 3b).

In this work, tensor voting [40] is used to estimate the depth or disparity maps from stereo or motion data, and to clean up noisy depth maps. We will not repeat here the details of customizing general voting frameworks to infer depth maps from stereo and motion data, but refer interested readers to [32] for depth from stereo and [25] for depth from motion. We will, however, supply an intuitive illustration of the generic methodology that is used to infer missing details and reject outliers. Additional information on tensor voting can be found in [40]. Detailed pseudocode for the methods we employ is provided in the Appendix; the core C++ implementations of the pseudocode segments (about 200 lines of code) are available in the library TVlib (see Section 7). The tensor voting algorithm can be implemented efficiently and it runs in O(nk) time, where n is the number of points (pixels) and k is the size of neighborhood or the scale of the analysis.

Tensor voting is a computational framework for feature extraction and segmentation. Tensors are used for data representation and they "communicate" with each other via a voting algorithm. In essence, we want to answer the following question: Given a point *P* with a finite neighborhood, what is the most likely normal direction at *P*, if there is any? Consider an analogy in particle physics, two particles vibrate and each particle emits a waveform to communicate with the other particle. Their vibrations reinforce each other if the particles are vibrating at *compatible* frequencies, resulting in a maximum (or resonance). The *ball voting field*, described in the Appendix, is used to mimic this communication process.

Fig. 8a shows a scenario with an outlier E and a missing data point. Suppose points A, B, C, D, F, G, H, and I all lie on an underlying smooth curve. Inspired by the particle physics analogy, we want the tangent directions at each point to be *compatible* to produce a smooth curve; for example, the tangent directions at all points on a straight line should be equal.

How do we find the compatible tangent direction, say  $\theta$ , at a particular point *C*? In tensor voting, all points communicate with each other by the ball voting field, which suggests desirable tangent directions. After *C* has collected a set of suggested tangent directions, principal component analysis is performed to determine an optimal  $\theta$ . The collected votes can be understood as a histogram: In Fig. 8b, a maximum occurs at  $\theta$  after vote collection for the point *C* (the situation is, of course, oversimplified in this illustration; in practice, we see a distribution of angles such as that suggested by the dotted curve).

To detect outliers, we look for atypical tensor voting histograms like that shown in Fig. 8c. Since E receives a set of incompatible directions from the other eight points, no salient maximum can be found in the histogram.

At x = 6, there is a missing data point in Fig. 8a. All the eight points cast directional votes to the set of locations

<sup>1.</sup> Whether the epipole is positive or negative depends on whether the projection is formed by intersection of positive epipolar ray (the vector from the reference viewpoint toward the desired viewpoint) or negative epipolar ray (the vector from the desired viewpoint toward the reference viewpoint) with the reference image plane.

-

	McMillan's method [39]	BSPI	
Drawing order	Based on epipolar geometry	Based on BSP in image space	
Drawing primitive	Pixel	Polygon	
Number of primitives	Fixed (image size)	O(image size), depending on how	
		much image space coherence is ex-	
		ploited by tensor voting. The num-	
		ber of polygons is usually much	
		smaller than the number of pixels in	
		the image.	
Depth buffering	N	0	
Hardware acceleration	Hardware-accelerated rendering of	Hardware-accelerated rendering of	
	points	textured polygons	
Gaps in rendered views	Yes (due to forward mapping)	No	
Multiple views	Possible with LDI or LDI trees	Layered BSPI	

 TABLE 1

 Comparison between the Occlusion-Compatible Drawing Algorithm and the BSPI

 $\{(x,d)|d \in [0,D] \text{ is the set of quantized depths}\}$ . By the continuity constraint imposed by the ball voting field, (6,17) is the solution. A salient maximum should be found in its corresponding histogram after tensor voting. Consequently, we not only reject outliers, but can also fill in missing details and infer probable tangent (or normal) directions.

#### 5.2 Feature Extraction

Once a (cleaned) depth map is available, we perform feature extraction on the depth map so that illumination, colors, and textures are no longer problematic. The output feature curves indicate the loci of the depth discontinuities (see ANGEL's depth discontinuities, Fig. 3), along which we may choose to place the partition lines of the BSPI.

Next, we illustrate the use of tensor voting to localize feature curves (for details see [60]). We use the RENAULT example in Fig. 9 to illustrate an underlying depth surface that is piecewise smooth but not closed. Our goal is to label the *depth discontinuity curves* automatically (drawn in bright red in Fig. 9).

Consider a finite neighborhood around a point P = (x, y, d), where (x, y) are the image coordinates and d is the depth. Note that the normal direction at P, and thus its tangent plane, is available after preprocessing. We collect the points within P's neighborhood such that they lie on a smooth surface. This collection is projected onto P's tangent plane. To check if a neighboring point is smoothly connected to P, we arrange for the points to communicate with one another using the *stick voting field* (defined in the Appendix).

Finally, we consider the vote distribution of the projected points. If the distribution is uneven, we label the point as a boundary point (Fig. 10).

# 5.3 BSPI Generation

Since we apply tensor voting in the 3D *x-y-d* volume in order to enforce surface smoothness, we perform the following steps before building a BSPI in the 2D image domain (see, e.g., ANGEL's BSPI, Fig. 3):

- 1, Project and quantize the discontinuity curves obtained in Section 5.2 onto the 2D image domain. Let the projected set be  $X = \{(x, y)\}$  as illustrated in Fig. 11a.
- 2, Apply tensor voting to *X* and collect directions at *every* pixel in the image domain. This process generates our dense 2D *CMap* map. Each pixel in a *CMap* has an associated tuple  $(s_p, \hat{t}_p)$ , where  $s_p$  is a scalar indicating the saliency or confidence level for the presence of a salient curve and  $\hat{t}_p$  is a unit vector indicating the 2D tangent direction if a curve exists (see Fig. 11b).

Before calling the recursive procedure **BSPI** whose pseudocode is given below, we initialize the structure *bsptree* to be empty, and set the structure *polygon* to be the rectangular 2D image domain.

**BSPI** (polygon, bsptree, CMap) {

Pick a candidate set *Y*:

 $Y = \{p = (s_p, \hat{t}_p) | s_p \ge \text{ some saliency value and} \\ p \in polygon\}$ 

if  $(Y = \emptyset)$  return

Compute  $score_p$  for all  $p \in Y$  (score is defined below) Partition polygon into  $P_1 \cup P_2 \cup \{\ell\}$  at pixel  $q \in Y$ ,

such that  $score_q$  is maximum.

 $\ell$  is the partition line, whose gradient is given by  $\hat{t}_q$ .  $bsptree \leftarrow bsptree \cup \{\ell\}$ 



Fig. 8. (a) Rejecting outliers and inferring missing details by tensor voting. (b) Histogram for point C after tensor voting:  $\theta$  is the *compatible* tangent direction at point C producing a smooth curve. (c) Histogram for point E after tensor voting: No maximum is produced since E is an outlier.



Fig. 9. Two views of the depth surface of the RENAULT model. The bright red curves are depth discontinuity and boundary curves.

Inhibit pixels in CMap close to  $\ell$ BSPI ( $P_1$ , bsptree, CMap) BSPI ( $P_2$ , bsptree, CMap)

}





Fig. 10. Automatic depth discontinuity detection by tensor voting: uneven distribution of surfel votes.

candidate partition line  $\ell_p$ . We quantize  $\ell_p$  in the image domain and set R = I - Y, where I is the set of all image pixels. R is therefore the set of region pixels and corresponds to smooth depth surfaces. We penalize  $\ell_p$  if it trespasses on the interior of an object, i.e., if  $\ell_p \cap R \neq \emptyset$ . Therefore,  $score_p$  is defined as

$$score_p = \alpha \sum_{e \in E} d_{p,e}e - \beta \sum_{r \in R} r,$$

where  $\alpha, \beta > 0$ .  $d_{p,e}$  is the inner product of  $\ell_p$  and  $\hat{t}_e$ , indicating their orientation consistency.

When we have computed  $\ell$  that gives the maximum score, an envelop of pixels surrounding the quantized  $\ell$  is inhibited in order to avoid producing elongated partitions in later recursions. Convex polygons at leaf nodes are then triangulated to support the rendering operation.



Fig. 11. (a) Projecting the discontinuity curves. (b) CMap is produced after tensor voting. Here,  $score_p > score_q$  since  $\ell_q \cap R \neq \emptyset$ . Here,  $\ell_p$  is chosen as the partition line.

#### 5.4 BSPI Rendering

Rendering a BSPI consists of two steps: visibility sorting and warping. Visibility sorting begins at the root node of the BSPI and the order of subtree traversal depends on the location of the novel view, as discussed in Section 4. Note that, since the BSPI is a static data structure, no write-back to memory is needed when resolving visibility as is the case for depth buffering. In contrast, straightforward comparison at internal nodes is sufficient to produce the correct warping order. Two novel views of the rendered result and the underlying drawing orders of triangles of ANGEL are shown in Fig. 3, using a color code to be explained below.

During the BSP tree traversal, whenever we visit a leaf node, we warp and draw the polygon that corresponds to the current leaf node. Warping can be done either using the 2D warping equation [36] or by reprojecting vertices onto the image planes of the novel view (Fig. 12). Consequently, we can draw the warped polygons with the reference images texture-mapped onto them. Finally, we apply a standard feathering technique that adjusts the alpha value at pixels along edge discontinuities to avoid undesirable artifacts and enhance the image quality.

#### 6 RESULTS

To demonstrate the effectiveness and efficiency of the proposed method, we create BSPIs from three different kinds of real images: single-view range images with corresponding color images (Fig. 16), stereo image pairs without depth (Fig. 17), and motion sequences without



Fig. 12. Reprojecting vertices of polygons to synthesize novel views.

depth (Fig. 18 and Fig. 19). Readers are encouraged to visit the Web page listed in Section 7 for the video that explains our algorithm with running examples and animations.

Fig. 16a shows the single-view image of a human FACE with its given depth in Fig. 16b (captured by a laser range finder). Although the captured depth data is quite accurate, holes exist due to the scanner's inability to capture surfaces with particular reflectivities. Tensor voting is applied to fill these gaps and remove noise. The resulting BSPI is shown in Fig. 16c. To illustrate the order of partition steps, we color-code the partition lines. Partition lines generated in earlier iterations are coded in red (brighter), while those in later iterations are in blue (darker). As shown in Fig. 16c, the first few partitions are done at the silhouette of the face, where higher depth discontinuities are located. Fig. 16d and Fig. 16e show the rendered novel views. Their corresponding drawing orders are illustrated in Fig. 16f and Fig. 16g by color-coded triangles. Triangles that are to be drawn earlier are darker, while those drawn at a later stage are brighter. Hence, foreground triangles are brighter than those in the background. Note how the nose occludes the other parts of the face in the rendered images. The average frame rate for the rendering shown under the figures is measured on a PC with a Pentium III 750MHz processor and 128MB RAM without any graphics accelerator installed. Thus, our method offers real-time rendering on ordinary PCs. It is worth noting that, due to the hierarchical nature of the BSPI, we can naturally introduce level-of-detail control to further speed up the rendering. Another example of a range image is the angel sculpture ANGEL shown in Fig. 3b.

Fig. 17 shows an input stereo image pair for the RENAULT piece, a manufacturing part. No depth is given. Disparity is recovered from the stereo pair. The recovered depth may not be Euclidean, but can be used for generating novel views from different viewpoints. As before, we color-code the partition lines (Fig. 17c) and the triangles (Fig. 17f and Fig. 17g) to show the partition process and drawing orders, respectively.

Fig. 18 demonstrates the use of multiple BSPIs to handle occlusion inherent in motion data, the FLOWER GARDEN sequence in our case. A portion of the flower garden scene is occluded by a tree trunk in the reference frames. The initial depth map is obtained from three frames [25]. First, a layered and textured description involving the foreground

	ANGEL	Face	FLOWER	Flower
			GARDEN	GARDEN
			(background)	(foreground)
No. of pts in feature extraction (Section 5.2)	14,670	20,281	76,161	10,922
Time for feature extraction	72 sec	95 sec	251 sec	39 sec
Time for BSPI generation (Section 5.3)	12 sec	26 sec	103 sec	5 sec
Total time for feature extraction				
and BSPI generation	84 sec	121 sec	354 sec	44 sec
Height of the BSP tree generated	33	44	38	24

TABLE 2 Sample Preprocessing Times for BSPI Generation

and the background is obtained by tensor voting. Texture synthesis in the presence of occlusion is detailed in [30]. Note that this texture synthesis preprocessing step is also automatic. Two BSPIs, which we call layered BSPIs, are obtained. During rendering, the background BSPI is traversed first, followed by the foreground BSPI. Our system can be readily scaled up to handle additional layered BSPIs. Since the drawing order has been resolved by the layered descriptions, rendering multiple BSPIs does not need any depth sorting. Note that our method alleviates the problem of "disocclusion," when a previously occluded part of the scene becomes visible at a novel viewpoint. No gaping hole is observed in our result. The same example has been used by many researchers (layered representation in [62] and sprites in [5], among others). To synthesize a novel view, the depth layers or sprites are rendered at a given viewpoint. No depth sorting is required in the BSPI rendering approach.

Note that, using tensor voting, constructing a BSPI from the raw depth or disparity data takes only a few minutes. This varies according to the noisiness and complexity of the data set. Like the BSPI warping, all the processing is done in 2D, with the exception of the feature curve inference. The latter is performed in disparity space in order to enforce the local smoothness constraint. Still, this can be done very efficiently since we do not fit any pseudo-object or disparity surface for curve inference. The preprocessing step is performed only once, including the BSPI generation. Once the BSPI is generated, it can be used by the BSPI warper to yield novel views efficiently.

Table 2 tabulates some sample preprocessing times of the results we presented. All preprocessing times were measured on an IBM ThinkPad PIII 850 MHz with 512M RAM. The mean preprocessing rate for depth estimation and noise filtering (Section 5.1) is 10,000 points per minute. Processing the statement of the section of th

sing times for feature extraction (Section 5.2) and BSPI generation (Section 5.3) are shown. These typical times show that our automatic preprocessing, which is run once, takes only a few minutes. The height of the BSP tree generated is also reported in the table.

#### 7 WEB AVAILABILITY

Interested readers are encouraged to visit the following Web site for an accompanying video explaining the BSPI algorithm with running examples presented in this paper: http://www.cse.cuhk.edu.hk/~ttwong/papers/bspi/ bspi.html. For the detailed implementation of tensor voting, readers can obtain the core library of C++ code from http:// www.cs.ust.hk/~cktang/TVlib.zip.

#### 8 CONCLUSION

In this paper, we present the BSPI—a novel IBMR approach to represent and render range images, stereo image pairs, and motion data. No intermediate pseudo-objects need to be constructed for the rendering. The BSPI representation implicitly embeds a projection-equivalent 3D BSP tree into the 2D image plane and thus actually supports a version of the general 3D BSP visibility algorithm. We utilize epipolar geometry to perform visibility sorting completely in 2D image space rather than 3D object space. Our results demonstrate the efficiency and robustness of the method. The BSP representation of images also facilitates efficient image-based data compression [49] and image retrieval [47]. To handle noise in the input data, we extend the tensor voting algorithm for the data cleanup process and incorporate the results into the recursive BSP partitioning process. This technique enhances the noise tolerance in the BSPI partitioning stage. To handle disocclusion, we combine





Fig. 13. Ellipse: a second order symmetric tensor in 2D.

Fig. 14. The design rationale of the stick voting field.



Fig. 15. The directions and strengths of the stick and ball voting fields.



(b)





(b)







(f)

(g)

Fig. 16. Results on FACE. Average frame rate: 20.844 fps and triangle count: 2,703. Input image and initial depth map courtesy of the Signal Analysis and Machine Perception Laboratory, The Ohio State University. (a) Reference image, (b) recovered depth, (c) BSPI, (d) novel view, (e) novel view, (f) drawing order, (g) drawing order.

multiple BSPIs from the same reference viewpoint, giving a layered representation of the scene. We also obtain this representation using tensor voting in the automatic preprocessing step.



Fig. 17. Results on RENAULT. Average frame rate: 52.098 fps and triangle count: 1,450. Input images and initial depth map courtesy of Lee and Medioni. (a) Stereo pair, (b) recovered depth, (c) BSPI, (d) novel view, (e) novel view, (f) drawing order, (g) drawing order.

In the future, we plan to extend the layering method to combine multiple BSPI's from multiple viewpoints. Layering techniques similar to the Layered Depth Image (LDI) method [53] can be exploited to combine multiple BSPIs corresponding to different viewpoints. This is useful for fast object surface reconstruction using data from single-view laser range scanners. Currently, we assume surfaces are Lambertian; view-dependent texture mapping could be used to remove this limitation. In addition, it would be beneficial to apply the BSPI partitioning method directly to the optic-flow field, thus avoiding the loss of information when converting to depth.



Fig. 18. Results on FLOWER GARDEN. Layered BSPIs are used to address "disocclusion." Average frame rate: 33.459 fps and triangle count: 744. Input images and initial velocity map courtesy of Gaucher and Medioni. (a) Image sequence, (b) recovered background depth, (c) recovered foreground depth, (e) background BSPI, (f) foreground BSPI.

# **APPENDIX** A

#### **TENSOR REPRESENTATION**

A tensor encapsulates both direction certainty *and* uncertainty. Consider a point in 2D space. It belongs to a curve (absolute certainty in a single direction), is a point junction (absolute *un*certainty in all directions), or is an outlier. The 2D stick and ball tensors are used, respectively, to encode the two extreme cases. The whole continuum can therefore be described as an ellipse or, equivalently, as a second order symmetric tensor in 2D: Mathematically, it is a  $2 \times 2$  eigensystem, with two unit eigenvectors  $\hat{e}_1$  and  $\hat{e}_2$  and two eigenvalues  $\lambda_1 \geq \lambda_2$  (see Fig. 13).

The eigenvectors represent the tangent direction and the eigenvalues indicate the magnitudes of orientation (un)certainties since the latter encode the size of the ellipse. The equation of a second order symmetric tensor is  $(\lambda_1 - \lambda_2)\mathbf{S} + \lambda_2\mathbf{B}$ , where  $\mathbf{S} = \hat{e}_1\hat{e}_1^T$  defines a stick tensor and  $\mathbf{B} = \hat{e}_1\hat{e}_1^T + \hat{e}_2\hat{e}_2^T$  defines a ball tensor. If normal or tangent information is unavailable initially, we encode the input points as ball tensors. Take the 2D case as an example. Then, each point  $(x, y)^T$  is associated with an eigensystem with  $\lambda_1 = \lambda_2 = 1$ ,  $\hat{e}_1 = \begin{bmatrix} 1 & 0 \end{bmatrix}^T$ ,  $\hat{e}_2 = \begin{bmatrix} 0 & 1 \end{bmatrix}^T$ . After encoding, the input points communicate with one another via the ball voting field.

#### **APPENDIX B**

#### THE STICK AND BALL VOTING FIELDS

We want to answer the following question: For a given point P in space and normal  $\vec{N}$  at the origin, what is the most likely normal at P if there is a smooth surface interpolating O (the origin) to P, and perpendicular to  $\vec{N}$ ? Fig. 14 illustrates the situation in 3D. We claim that the osculating circle connecting O and P is the most likely curve since it keeps the curvature constant along the hypothesized circular arc. See [40] for detailed mathematical derivation. The most likely normal direction at P is thus given by the normal to the osculating circle at P. The length of the normal vector at P, which represents the surface saliency of the vote, is inversely proportional to the arc length OP and also to the curvature of the underlying circular arc. In spherical coordinates, the decay of the field takes the following form:

$$\overline{DF}(r,\varphi,\sigma) = e^{-\left(\frac{r^2 + c\varphi^2}{\sigma^2}\right)},\tag{1}$$

where *r* is the arc length *OP*,  $\varphi$  is the curvature, and  $\sigma$  is the size of neighborhood or scale of analysis. The set of normals at all *P*s in the 2D space constitutes the stick voting field *V*<sub>S</sub>.

The ball voting field  $V_B$  is obtained by rotating and integrating vote contributions of the 2D *stick voting fields* since all directions are equally likely:  $V_B = \int_0^{\pi} V_S d\theta$ . Fig. 15 shows the shapes, directions, and strengths of the stick and the ball voting fields in 2D.



Fig. 19. Results on FLOWER GARDEN (continued). Average frame rate: 33.459 fps and triangle count: 744. Input images and initial velocity map courtesy of Gaucher and Medioni. (g) Novel view, (h) novel view, (i) background drawing order, (j) background drawing order, (k) foreground drawing order, (I) foreground drawing oder.

# **APPENDIX C**

### ALGORITHMS OF 2D TENSOR VOTING

This section describes the general tensor voting algorithm [57]. The voter makes use of GENTENSORVOTE (shown in Fig. 20) to cast a tensor vote to vote receiver (votee). Normal direction votes generated by GENNORMALVOTE (shown in

lgorithm 1 GENTENSORVOTE (voter, votee)
/* It uses GENNORMALVOTE to compute the most likely normal
direction vote at the votee. Then, plate and ball tensors are
computed, by integrating the resulting normal votes cast by voter. */
for all $0 \le i, j < 2$ , outTensor[i][j] $\leftarrow 0$
voterSaliency[0] $\leftarrow$ voter[ $\lambda_1$ ] - voter[ $\lambda_2$ ]
voterSaliency[1] $\leftarrow$ voter[ $\lambda_2$ ]
/* Generate the stick component */
if (voterSaliency $[0] > 0$ ) then
$vecVote \leftarrow GENNORMALVOTE$ (voter, votee)
/* Compute stick component */
COMBINE (outTensor, vecVote)
end if
/* Compute ball component */
if (voterSaliency[1] $> 0$ ) then
/* count is a sufficient number of samples uniformly distributed on a unit circle */
while $(\operatorname{count} \neq 0)$ do
/* Generate a random point on a unit circle. */
voter[direction] $\leftarrow$ random[direction] $\leftarrow$ GENRANDOMUNIFORMPT()
$vecVote \leftarrow GENNORMALVOTE$ (voter, votee)
COMBINE (outTensor, vecVote, voterSaliency[1])
$\operatorname{count} \leftarrow \operatorname{count} - 1$
end while
end if
return outTensor

Fig. 20. Algorithm 1: GENTENSORVOTE.

Algorithm 2 GENNORMALVOTE (voter, votee)
A vote (vector) on the most likely normal direction is returned.
$v \leftarrow votee[position] - voter[position]$
/* voter and votee are connected by high curvature? */
if (angle(voter[direction], v) $< \pi/4$ ) then
return NULL {smoothness constraint violated}
end if
stickvote[position]   votee[position]
/* voter and votee on a straight line, or voter and votee are the same point */
if (angle(voter[direction], $v$ ) = $\pi/2$ ) or (voter = votee) then
stickvote[direction] $\leftarrow$ voter[direction]
stickyote[length] $\leftarrow e^{-\frac{r^2}{\sigma^2}}$ {equation (1)}
return stickvote
end if
Compute center and radius of the osculating hemisphere
/* assign stick vote */
stickvote[direction] ← center - votee[position]
stickvote[length] $\leftarrow e^{-(\frac{p^2+c\varphi^2}{2})}$ {equation (1)}
return stickvote
Algorithm 3 COMBINE (tensorvote, stickvote, weight)
/* It performs tensor addition, given a stick vote. */
for all $i, j$ such that $0 \le i, j < 2$ do
$tensorvote[i][j] \leftarrow tensorvote[i][j] + weight \times stickvote[i] \times stickvote[j]$
end for



Fig. 21) are accumulated using COMBINE (also shown in Fig. 21). A  $2 \times 2$  outTensor is the output. The votee thus receives a set of *outTensor* from voters within its neighborhood. All *outTensors* are summed up by tensor addition, which is simply ordinary  $2 \times 2$  matrix addition. This resulting matrix is equivalent to an ellipse.

#### **ACKNOWLEDGMENTS**

The authors would like to thank Dr. Michael S. Brown for his audio help in making the video, which is available at our Web site. Also, they would like to thank Gaucher, Lee, and Medioni as well as the Signal Analysis and Machine Perception Laboratory, The Ohio State University for various data sets used in this research work. This research was supported in part by the Hong Kong Research Grant Council Cooperative Research Centres (CRC) Scheme (Project no. CRC 4/98) and RGC Earmarked Grants (Project No. CUHK 4186/00E, HKUST6089/99E, HKUST6193/02E, HIA98/ 99.EG01, and AOE/E-01/99).

#### REFERENCES

- M. Abrash, "Inside Quake: Visible-Surface Determination," [1] Dr. Dobb's Sourcebook, pp. 41-45, Jan./Feb. 1996.
- P.K. Agarwal, J. Erickson, and L.J. Guibas, "Kinetic Binary Space Partitions for Intersecting Segments and Disjoint Triangles," Proc. Ninth Ann. ACM-SIAM Symp. Discrete Algorithms, pp. 107-116, Jan. 1998.
- [3] J.M. Airey, J.H. Rohlf, and F.P. Brooks, "Towards Image Realism with Interactive Update Rates in Complex Virtual Building Environments," ACM SIGGRAPH Special Issue 1990 Symp. Inter-
- active 3D Graphics, pp. 41-50, 1990.
  [4] D.P. Anderson, "Hidden Line Elimination in Projected Grid Surfaces," ACM Trans. Graphics, vol. 1, no. 4, 274-288, Oct. 1982.
- [5] S. Baker, R. Szeliski, and P. Anadan, "A Layered Approach to Stereo Reconstruction," Proc. IEEE Computer Vision and Pattern Recognition 1998 (CVPR '98) Conf., 1998.
- [6] M. De Berg, M.M. De Groot, and M.H. Overmars, "Perfect Binary Space Partitions," CGTA: Computational Geometry: Theory and Applications, vol. 7, pp. 81-91, Jan. 1997.
  [7] W. Bouknight, A Procedure for Generation of Three-Dimensional
- Halftoned Computer Graphics Presentations, 1970.
- [8] W. Bouknight and K. Kelly, An Algorithm for Producing Halftone Computer Graphics Presentations with Shadows and Movable Light Sources, 1970.

- [9] T. Cassen, K.R. Subramanian, and Z. Michalewicz, "Near-Optimal Construction of Partitioning Trees by Evolutionary Techniques," *Proc. Graphics Interface* '95, W.A. Davis and P. Prusinkiewicz, eds., pp. 263-271, 1995.
- [10] E.E. Catmull, "A Subdivision Algorithm for Computer Display of Curved Surfaces," PhD thesis, Dept. of Computer Science, Univ. of Utah, 1974.
- [11] C.-F. Chang, G. Bishop, and A. Lastra, "LDI Tree: A Hierarchical Representation for Image-Based Rendering," SIGGRAPH '99 Conf. Proc., pp. 291-298, Aug. 1999.
- [12] H.-M. Chen and W.-T. Wang, "The Feudal Priority Algorithm on Hidden-Surface Removal," Proc. SIGGRAPH '96, pp. 55-64, Aug. 1996.
- [13] S.E. Chen and L. Williams, "View Interpolation for Image Synthesis," *Computer Graphics (SIGGRAPH '93 Proc.)*, vol. 27, pp. 279-288, Aug. 1993.
- [14] N. Chin and S. Feiner, "Near Real-Time Shadow Generation Using BSP Trees," SIGGRAPH '89 Conf. Proc., vol. 23, pp. 99-106, July 1989.
- [15] N. Chin and S. Feiner, "Fast Object-Precision Shadow Generation for Areal Light Sources Using BSP Trees," *Computer Graphics* (1992 *Symp. Interactive 3D Graphics)*, vol. 25, pp. 21-30, Mar. 1992.
- [16] J.L.D. Comba, "Kinetic Vertical Decomposition Trees," PhD thesis, Dept. of Computer Science, Stanford Univ., 1999.
- [17] G. Crocker, "Invisibility Coherence for Faster Scan-Line Hidden Surface Algorithms," SIGGRAPH '84 Conf. Proc., vol. 18, pp. 95-102, July 1984.
- [18] P.E. Debevec, C.J. Taylor, and J. Malik, "Modeling and Rendering Architecture from Photographs: A Hybrid Geometry- and Image-Based Approach," *SIGGRAPH '96 Conf. Proc.*, H. Rushmeier, ed., pp. 11-20, Aug. 1996.
- [19] O. Faugeras and L. Robert, "What Can Two Images Tell Us about a Third One?" technical report, INRIA, July 1993.
- [20] C.-W. Fu, T.-T. Wong, and P.-A. Heng, "Computing Visibility for Triangulated Panorama," Proc. 10th Eurographics Rendering Workshop, June 1999.
- [21] H. Fuchs, G. Abram, and E. Grant, "Near Real-Time Shade Display of Rigid Objects," SIGGRAPH '83 Conf. Proc., vol. 17, pp. 65-72, July 1983.
- [22] H. Fuchs, Z.M. Kedem, and B.F. Naylor, "On Visible Surface Generation by A Priori Tree Structures," SIGGRAPH '80 Conf. Proc., vol. 14, pp. 124-133, July 1980.
- [23] T.A. Funkhouser, C.H. Séquin, and S.J. Teller, "Management of Large Amounts of Data in Interactive Building Walkthroughs," ACM SIGGRAPH Special Issue 1992 Symp. Interactive 3D Graphics, pp. 11-20, 1992.
- [24] T.A. Funkhouser and C.H. Séquin, "Adaptive Display Algorithm for Interactive Frame Rates during Visualization of Complex Virtual Environments," *SIGGRAPH '93 Conf. Proc.*, J.T. Kajiya, ed., vol. 27, pp. 247-254, Aug. 1993.
- [25] L. Gaucher and G. Medioni, "Accurate Motion Flow Estimation with Discontinuities," Proc. Seventh Int'l Conf. Computer Vision (ICCV '99), pp. 695-702, Sept. 1999.
- [26] A.S. Glassner, "Space Subdivision for Fast Ray Tracing," IEEE Computer Graphics and Applications, vol. 4, no. 10, pp. 15-22, Oct. 1984.
- [27] S.J. Gortler, R. Grzeszczuk, R. Szeliski, and M.F. Cohen, "The Lumigraph," Computer Graphics Proc., Ann. Conf. Series, SIG-GRAPH '96, pp. 43-54, Aug. 1996.
- [28] N. Greene, "Hierarchical Polygon Tiling with Coverage Masks," SIGGRAPH '96 Conf. Proc., pp. 65-74, Aug. 1996.
- [29] N. Greene and M. Kass, "Hierarchical Z-Buffer Visibility," Computer Graphics Proc., Ann. Conf. Series, pp. 231-240, 1993.
- [30] J. Jia and C.-K. Tang, "Texture Synthesis in the Presence of Discontinuities from Noisy Data," Proc. Eurographics Rendering Workshop 2002, submitted.
- [31] S. Laveau and O. Faugeras, "3-D Scene Representation as a Collection of Images," Proc. 12th Int'l Conf. Pattern Recognition (ICPR '94), pp. 689-691, Oct. 1994.
- [32] M.-S. Lee and G. Medioni, "Inferring Segmented Surface Description from Stereo Data," Proc. IEEE CS Conf. Computer Vision and Pattern Recognition (CVPR '98), pp. 346-352, June 1998.
- [33] J. Lengyel and J. Snyder, "Rendering with Coherent Layers," Computer Graphics Proc., Ann. Conf. Series, SIGGRAPH '97, pp. 233-242, Aug. 1997.

- [34] M. Levoy and P. Hanrahan, "Light Field Rendering," Computer Graphics Proc., Ann. Conf. Series, SIGGRAPH '96, pp. 31-42, Aug. 1996.
- [35] M. Lhuillier and L. Quan, "Image Interpolation by Joint View Triangulation," Proc. Computer Vision and Pattern Recognition (CVPR), 1999.
- [36] W.R. Mark and G. Bishop, "Memory Access Patterns of Occlusion-Compatible 3D Image Warping," *Proc.* 1997 Siggraph/Eurographics Workshop Graphics Hardware, pp. 35-44, Aug. 1997.
- [37] M. McKenna, "Worst-Case Optimal Hidden-Surface Removal," ACM Trans. Graphics, vol. 6, no. 1, pp. 19-28, Jan. 1987.
- [38] L. McMillan and G. Bishop, "Plenoptic Modeling: An Image-Based Rendering System," Computer Graphics Proc., Ann. Conf. Series, SIGGRAPH '95, pp. 39-46, Aug. 1995.
- [39] L. McMillan, "An Image-Based Approach to Three-Dimensional Computer Graphics," PhD thesis, Dept. of Computer Science, Univ. of North Carolina at Chapel Hill, Chapel Hill, 1997.
- [40] G. Medioni, M.-S. Lee, and C.-K. Tang, A Computational Framework for Feature Extraction and Segmentation. Amsterdam: Elseviers Science, 2000.
- [41] K. Mulmuley, "An Efficient Algorithm for Hidden Surface Removal," Computer Graphics (Proc. SIGGRAPH '89), vol. 23, no. 3, pp. 379-388, July 1989.
- [42] B. Naylor, J. Amanatides, and W. Thibault, "Merging BSP Trees Yields Polyhedral Set Operations," *SIGGRAPH '90 Conf. Proc.*, vol. 24, pp. 115-124, Aug. 1990.
- [43] B.F. Naylor, "Partitioning Tree Image Representation and Generation from 3D Geometric Models," *Proc. Graphics Interface* '92, pp. 201-212, May 1992.
- [44] B.F. Naylor, "Constructing Good Partitioning Trees," Proc. Graphics Interface '93, pp. 181-191, May 1993.
- [45] M.E. Newell, R.G. Newell, and T.L. Sancha, "A Solution to the Hidden Surface Problem," Proc. ACM Ann. Conf., vol. I, pp. 443-450, 1972.
- [46] M.S. Paterson and F.F. Yao, "Efficient Binary Space Partitions for Hidden-Surface Removal and Solid Modeling," Discrete & Computational Geometry, vol. 5, pp. 485-503, 1990.
- [47] G. Qiu and S. Sudirman, "Color Image Coding, Indexing and Retrieval Using Binary Space Partitioning Tree," Proc. IEEE Int'l Conf. Image Processing, Oct. 2001.
- [48] H.R. Rabiee, R.L. Kashyap, and H. Radha, "Multiresolution Image Compression with BSP Trees and Multi-Level Block Truncation Coding," *Proc. IEEE Second Int'l Conf. Image Processing*, pp. 600-603, 1995.
- [49] H. Radha, M. Vetterli, and R. Leonardi, "Image Compression Using Binary Space Partitioning Trees," *IEEE Trans. Image Processing*, vol. 5, no. 12, pp. 1610-1624, Dec. 1996.
- [50] H. Rahda, R. Leonardi, M. Vetterli, and B.F. Naylor, "Binary Space Partitioning Tree Representation of Images," *Visual Comm. and Image Representation*, vol. 2, no. 3, pp. 201-221, Sept. 1991.
- [51] G. Schaufler and M. Priglinger, "Efficient Displacement Mapping by Image Warping," Proc. 10th Eurographics Workshop Rendering, pp. 175-186, 1999.
- [52] S. Sechrest and D.P. Greenberg, "A Visible Polygon Reconstruction Algorithm," ACM Trans. Graphics, vol. 1, no. 1, pp. 25-42, Jan. 1982.
- [53] J. Shade, S. Gortler, L.-W. He, and R. Szeliski, "Layered Depth Images," SIGGRAPH '98 Conf. Proc., Ann. Conf. Series, July 1998.
- [54] J. Snyder and J. Lengyel, "Visibility Sorting and Compositing without Splitting for Image Layer Decomposition," Proc. SIGGRAPH '98, pp. 219-230, July 1998.
- [55] K.R. Subramanian and B.F. Naylor, "Converting Discrete Images to Partitioning Trees," *IEEE Trans. Visualization and Computer Graphics*, vol. 3, no. 3, pp. 273-288, July-Sept. 1997.
- [56] O. Sudarsky and C. Gotsman, "Output-Sensitive Visibility Algorithms for Dynamic Scenes with Applications to Virtual Reality," *Computer Graphics Forum*, vol. 15, no. 3, pp. 249-258, Aug. 1996.
- [57] C.-K. Tang, G. Medioni, and M.-S. Lee, "N-Dimensional Tensor Voting, and Application to Epipolar Geometry Estimation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 8, pp. 829-844, Aug. 2001.
- [58] S. Teller and P. Hanrahan, "Global Visibility Algorithms for Illumination Computations," Proc. SIGGRAPH '93, pp. 239-246, Aug. 1993.

- [59] S.J. Teller and C.H. Séquin, "Visibility Preprocessing for Interactive Walkthroughs," *Computer Graphics (Proc. SIGGRAPH '91)*, T.W. Sederberg, ed., vol. 25, pp. 61-69, July 1991.
- [60] W.-S. Tong and C.-K. Tang, "First Order Tensor Voting, and Application to 3D Scale Analysis," Proc. IEEE CS Conf. Computer Vision and Pattern Recognition 2001 (CVPR '01), Dec. 2001.
- [61] E. Torres, "Optimization of the Binary Space Partition Algorithm (BSP) for the Visualization of Dynamic Scenes," *Proc. Eurographics* '90, pp. 507-518, Sept. 1990.
- [62] J.W.A. Wang and E.H. Adelson, "Layered Representation for Motion Analysis," Proc. IEEE Computer Vision and Pattern Recognition (CVPR '93) Conf., pp. 361-366, 1993.
- [63] J. Warnock, "A Hidden-Surface Algorithm for Computer Generated Half-Tone Pictures," Technical Report TR 4-15, NTIS AD-753 671, Computer Science Dept., Univ. of Utah, 1969.
- [64] G.S. Watkins, "A Real-Time Visible Surface Algorithm," Technical Report UTECH-CSc-70-101, Salt Lake City, Utah, 1970.
- [65] C. Wylie, G. Romney, D. Evans, and A. Erdahl, "Half-Tone Perspective Drawings by Computer," *Proc. AFIPS* 1967 FJCC, vol. 31, pp. 49-58, 1967.
- [66] H. Zhang, D. Manocha, T. Hudson, and K.E. Hoff III, "Visibility Culling Using Hierarchical Occlusion Maps," *Proc. SIGGRAPH* '97, pp. 77-88, Aug. 1997.



**Chi-Wing Fu** received the BSc and MPhil degrees in computer science and engineering from the Chinese University of Hong Kong in 1997 and 1999, respectively; he is now working toward the doctorate degree in the Computer Science Department of Indiana University at Bloomington. His research interests include image-based modeling, rendering, and relighting, medical visualization, as well as visualization and navigation of large-scale astrophysical tudent member of the IEEE Computer Society.

environment. He is a student member of the IEEE Computer Society.



**Tien-Tsin Wong** graduated from the Chinese University of Hong Kong in 1992 with the BSc degree in computer science. He obtained the MPhil and PhD degrees in computer science from the same university in 1994 and 1998, respectively. He is an assistant professor in the Department of Computer Science and Engineering at the Chinese University of Hong Kong. His main research interest is computer graphics, including image-based modeling and rendering,

medical visualization, natural phenomena modeling, and photorealistic and nonphotorealistic rendering. He is a member of the IEEE Computer Society



Wai-Shun Tong received the BEng and MPhil degrees from The Hong Kong University of Science of Technology (HKUST) in computer science. He was a visiting student at the University of Southern California (USC), Los Angeles, in 2001. In 2002, he was awarded the Microsoft Fellowship (Asia). His research interests include low-to-mid-level vision such as segmentation, correspondence, shape analysis, and, in particular, Tensor Voting. He is a student computer Society.

member of the IEEE Computer Society.



**Chi-Keung Tang** received the MS and PhD degrees in computer science from the University of Southern California (USC), Los Angeles, in 1999 and 2000, respectively. He has been with the Computer Science Department at the Hong Kong University of Science and Technology since 2000, where he is currently an assistant professor. He is also an adjunct researcher at the Visual Computing Group of Microsoft Research, China, working on various exciting

research topics in computer vision and graphics. His research interests include low-to-mid-level vision such as segmentation, correspondence, shape analysis, and vision and graphics topics such as image-based rendering and medical image analysis. He is a member of the IEEE Computer Society.



Andrew J. Hanson received the BA degree in chemistry and physics from Harvard College in 1966 and the PhD degree in theoretical physics from the Massachusetts Institute of Technology in 1971. He is a professor in the Computer Science Department at Indiana University. He did research in theoretical physics from 1971 until 1980, when he began working in machine vision, graphics, and visualization, first with the perception research group at the SRI Artificial

Intelligence Center and then at Indiana University starting in 1989. His current research interests include computer graphics, perception, navigation and collaboration in virtual environments, the visualization of abstract concepts in mathematics and physics, and virtual astronomy. He is a member of the IEEE Computer Society, ACM Siggraph, the American Mathematical Society, the American Physical Society, and Sigma Xi.

▷ For more information on this or any computing topic, please visit our Digital Library at http://computer.org/publications/dlib.