A Variable-Resolution Probabilistic Three-Dimensional Model for Change Detection

Daniel Crispell, *Member, IEEE*, Joseph Mundy, and Gabriel Taubin, *Fellow, IEEE*

***Abstract*—Given a set of high-resolution images of a scene, it is often desirable to predict the scene's appearance from viewpoints not present in the original data for purposes of change detection. When significant 3-D relief is present, a model of the scene geometry is necessary for accurate prediction to determine surface visibility relationships. In the absence of an *a priori* high-resolution model (such as those provided by LIDAR), scene geometry can be estimated from the imagery itself. These estimates, however, cannot, in general, be exact due to uncertainties and ambiguities present in image data. For this reason, probabilistic scene models and reconstruction algorithms are ideal due to their inherent ability to predict scene appearance while taking into account such uncertainties and ambiguities. Unfortunately, existing data structures used for probabilistic reconstruction do not scale well to large and complex scenes, primarily due to their dependence on large 3-D voxel arrays. The work presented in this paper generalizes previous probabilistic 3-D models in such a way that multiple orders of magnitude savings in storage are possible, making high-resolution change detection of large-scale scenes from high-resolution aerial and satellite imagery possible. Specifically, the inherent dependence on a discrete array of uniformly sized voxels is removed through the derivation of a probabilistic model which represents uncertain geometry as a density field, allowing implementations to efficiently sample the volume in a nonuniform fashion.**

***Index Terms*—Computer vision, data structures, remote sensing.**

I. Introduction

HIGH-RESOLUTION aerial imagery is quickly becoming a ubiquitous source of information in both defense and civil application domains due to the advancing technology of remote sensing and increasing availability of aerial platforms (including unmanned aerial vehicles). As the immense volume of produced imagery data grows over time, automated processing algorithms become ever more important.

A. Change Detection

One application wherein the collected imagery is frequently used as input for is change detection, where a new image is collected and must be compared with the "expected" appearance of the scene given the previously observed images. Some change detection algorithms operate at large scales, typically indicating changes in land-cover type (e.g., forests, urban, and farmland). Due to the increasing availability of high-resolution imagery, however, interest in higher resolution and intraclass change detection is growing. The precise definition of "change" is application dependent in general, and in many cases, it is easier to define what does *not* constitute valid change [1], [2]. Typically, changes in appearance due to illumination conditions, atmospheric effects, viewpoint, and sensor noise are not desired to be reported. Various classes of methods for accomplishing this have been attempted, a survey of which was given by Radke *et al.* [1] in 2005 (which includes the joint histogram-based method [3] used for comparison by Pollard *et al.* [2]). One common assumption relied on by most of these methods, as well as more recent approaches [4], [5], is an accurate registration of pixel locations in the collected image to corresponding pixel locations in the base imagery. When the scene being imaged is relatively flat or the 3-D structure is known *a priori*, "rubber sheeting" techniques can be used to accomplish this. When the scene contains significant 3-D relief viewed from disparate viewpoints, however, techniques based on this assumption fail due to their inability to predict occlusions and other viewpoint-dependent effects [2]. High-resolution imagery exacerbates this problem due to the increased visibility of small-scale 3-D structure (trees, small buildings, etc.) which is ubiquitous over much of the planet. In this case, a 3-D model of the scene is necessary for accurate change detection from arbitrary viewpoints. There have been some previous works using 3-D models for change detection. Huertas and Nevatia [6] matched image edges with projections of 3-D building models with some promise but relied on the manual step of model creation. Eden and Cooper's method [7] based on the automatic reconstruction of 3-D line segments avoids the manual model creation step, although it is unable to detect change due to occlusion since only a "wireframe" (as opposed to surface) model is constructed. Pollard *et al.* [2] proposed a probabilistic reconstruction method for change detection and is the basis for the work presented here.

B. Probabilistic Models

Computing exact 3-D structure based on 2-D images is, in general, an ill-posed problem. Bhotika *et al.* [8] characterized the sources of difficulty as belonging to one of two classes: scene *ambiguity* and scene *uncertainty*. Scene ambiguity exists due to the existence of multiple possible photo-consistent

Manuscript received December 1, 2010; revised March 18, 2011; accepted April 22, 2011.

D. Crispell is with the National Geospatial-Intelligence Agency, Springfield, VA 22150 USA (e-mail: dancrispell@gmail.com).

J. Mundy and G. Taubin are with Brown University, Providence, RI 02912, USA (e-mail: mundy@lems.brown.edu; taubin@brown.edu).

Digital Object Identifier 10.1109/TGRS.2011.2158439

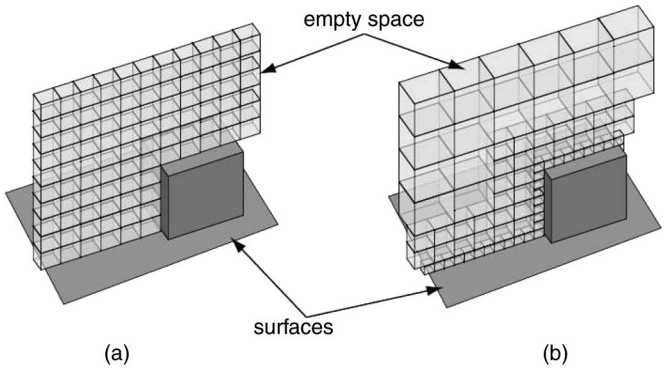

Fig. 1. Variable-resolution models (b) such as octrees allow for high-resolution representation where needed (i.e., near surfaces) with far less data than required in (a) fixed grid models. Representing surface probabilities using the proposed method allows for variable-resolution models to be used for probabilistic 3-D modeling approaches to change detection.

scenes and is a problem even in the absence of any sensor noise or violations of assumptions built into the imaging and sensor model. In the absence of prior information regarding the scene structure, there is no reason to prefer one possible reconstruction over another. The term "scene uncertainty," on the other hand, is used to describe all other potential sources of error including sensor noise, violations of certain simplifying assumptions (e.g., Lambertian appearance), and calibration errors. The presence of scene uncertainty typically makes reconstruction of a perfectly photo-consistent scene impossible. Probabilistic models allow the scene ambiguity and uncertainty to be explicitly represented, which, in turn, allows the assignment of confidence values to visibility calculations, expected images, and other data extracted from the model. A probabilistic model can also be used to determine which areas of the scene require further data collection due to low confidence.

C. Contributions

The work presented in this paper generalizes previous probabilistic reconstruction models in such a way that multiple orders of magnitude savings in storage are possible, making precise representation and change detection of large-scale outdoor scenes possible. Specifically, the inherent dependence on a discrete array of uniformly sized voxels is removed through the derivation of a new probabilistic representation based on a density field model. The representation allows for implementations which nonuniformly sample the volume, providing high-resolution detail where needed (e.g., near surfaces) and coarser resolutions in areas containing little information (e.g., in empty space) (Fig. 1). Additionally, it represents the first probabilistic volumetric model to provide a principled way to take viewing ray/voxel intersection lengths into account, enabling higher accuracy modeling and rendering. The proposed model combined with the reconstruction and associated algorithms comprise a practical system capable of automatically detecting change and generating photo-realistic renderings of large and complex scenes from arbitrary viewpoints based on image data alone.

D. Outline

The remainder of this paper is laid out as follows. In Section II, a brief survey of related work in the fields of 3-D reconstruction is given. Section III describes the theoretical foundations of the proposed model. Sections IV and V describe the implementation using an octree data structure and the reconstruction algorithms, respectively. The application of the reconstructed models for change detection are discussed in Section VI, followed by the paper's conclusion in Section VII.

II. Related Work

There is a large body of previous work in the computer vision community involving the automatic reconstruction of 3-D models from imagery, a brief overview of which is given here. The bulk of the representations used are not probabilistic in nature and are discussed in Section II-A. Existing probabilistic methods are discussed in Section II-B.

A. Deterministic Methods

Three-dimensional reconstruction from images is one of the fundamental problems in the fields of computer vision and photogrammetry, the basic principles of which are discussed in many texts including [9]. Reconstruction methods vary both in the algorithms used and the type of output produced.

Traditional stereo reconstruction methods take as input two (calibrated) images and produce a depth (or height) map as output. A comprehensive review of the stereo reconstruction literature as of 2002 is given by Scharstein and Szeliski [10]. While highly accurate results are possible with recent methods [11], [12], the reconstruction results are limited to functions of the form $f(x, y)$ and cannot completely represent general 3-D scenes on their own.

Many multiview methods are capable of computing 3-D point locations as well as camera calibration information simultaneously using the constraints imposed by feature matches across multiple images (so called "structure from motion"). One example of a such a method is presented by Pollefeyes *et al.* [13], who use tracked Harris corner [14] features to establish correspondences across frames of a video sequence. Brown and Lowe [15] and Snavely *et al.* [16] use scale invariant feature transform features [17] for the same purpose with considerable success. Snavley *et al.* have shown their system capable of successfully calibrating data sets consisting of hundreds of images taken from the Internet. The output of feature-based matching methods (at least in an initial phase) is a discrete and sparse set of 3-D elements which are not directly useful for the purpose of appearance prediction since some regions (e.g., those with homogeneous appearance) will be void of features and, thus, also void of reconstructed points. It is possible to estimate a full surface mesh based on the reconstructed features [18], [19], but doing so requires imposing regularizing constraints to fill in "holes" corresponding to featureless regions. Methods based on dense matching techniques avoid the hole-filling problem but are dependent on

smoothness and ordering assumptions to perform the matching. The methods presented in this paper are not dependent on any assumptions regarding the 3-D scene geometry yet produce a dense model suitable for high-resolution change detection.

B. Probabilistic Methods

As discussed in Section I-B, probabilistic 3-D models have the desirable quality of allowing a measure of uncertainty and ambiguity in the reconstructed model to be explicitly represented. Probabilistic methods are also capable of producing a complete representation of the modeled surfaces while making no assumptions about scene topology or regularizing constraints.

There exists in the literature several distinct methods for reconstructing a probabilistic volumetric scene model based on image data, all based on discrete voxel grid models. Although the methods vary in their approach, the goal is the same: to produce a volumetric representation of the 3-D scene, where each voxel is assigned a probability based on the likelihood of it being contained in the scene. The algorithms grew out of earlier "voxel coloring" algorithms [20]–[23] in which voxels are removed from the scene based on photometric consistency and visual hull constraints. Voxel coloring methods are prone to errors due to scene uncertainty; specifically, violations of the color consistency constraint often manifest themselves as incorrectly carved "holes" in the model [24]. To combat these errors, probabilistic methods do not "carve" voxels but rather assign each a probability of existing as part of the model. Broadhurst *et al.* [25] assign a probability to each voxel based on the likelihood that the image samples originated from a distribution with small variance rather than make a binary decision. Similarly, Bhotika *et al.* [8] carve each voxel with a probability based on the variance of the samples in each of a large number of runs. The final voxel probability is computed as the probability that the voxel exists in a given run.

In addition to uncertainty due to noise and other unmodeled phenomenon, any reconstruction algorithm must also deal with scene ambiguity, the condition which exists when multiple photo-consistent reconstructions are possible given a set of collected images. If certain *a priori* information about the scene is available, the information may be used to choose the photo-consistent reconstruction which best agrees with the prior. The reconstruction algorithm presented in this paper is as general as possible and, thus, does not utilize any such prior. Another approach is to define a particular member of the set of photo-consistent reconstructions as "special" and aim to recover that member. This is the approach taken by Kutulakos and Seitz [21] and Bhotika *et al.* [8]. Kutulakos and Seitz define the *photo hull* as the tightest possible bound on the true scene geometry, i.e., the maximal photo-consistent reconstruction. They show that, under ideal conditions, the photo hull can be recovered exactly, while Bhotika *et al.* present a stochastic algorithm for probabilistic recovery of the photo hull in the presence of noise. The photo hull provides a maximal bound on the true scene geometry but does not contain any information about the distribution of possible scene surfaces within the hull. A third approach is to explicitly and probabilistically represent the full range of possible scene reconstructions. Broadhurst *et al.* [25] aim to reconstruct such a representation, as well as Pollard *et al.* [2] for the purpose of change detection. Pollard *et al.* use an online Bayesian method to update voxel probabilities with each observation. Because a model which fully represents both scene ambiguities and uncertainties and is capable of change detection is desired, the model and algorithms presented in this paper are based on this approach.

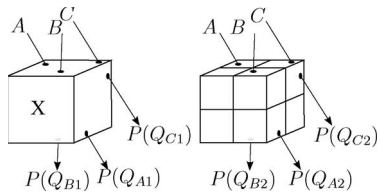

Fig. 2. (Left) Three viewing rays pass through a voxel, which (right) is then subdivided to achieve higher resolution. Assuming that no further information about the voxel is obtained, the occlusion probabilities $P(\mathcal{Q}_{A1})$, $P(\mathcal{Q}_{B1})$, and $P(\mathcal{Q}_{C1})$ should not depend on the level of subdivision and should be equal to $P(\mathcal{Q}_{A2})$, $P(\mathcal{Q}_{B2})$, and $P(\mathcal{Q}_{C2})$, respectively.

One quality that current volumetric probabilistic reconstruction methods all share is that the voxel representation is inherently incapable of representing the true continuous nature of surface location uncertainty. Using standard models, occlusion can only occur at voxel boundaries, since each voxel is modeled as being either occupied or empty. A side effect of this fact is that there is no principled way to take into account the length of viewing ray/voxel intersections when computing occlusion probabilities, which limits the accuracy of the computations. These limitations are typically handled in practice by the use of high-resolution voxel grids, which minimize the discretization effects of the model. Unfortunately, high-resolution voxel grids are prohibitively expensive, requiring $O(n^3)$ storage to represent scenes with linear resolution n.

1) Variable-Resolution Probabilistic Methods: The benefits of an adaptive variable-resolution representation are clear: In theory, a very highly effective resolution can be achieved without the $O(n^3)$ storage requirements imposed by a regular voxel grid. One hurdle to overcome is the development of a probabilistic representation which is invariant to the local level of discretization. A simple example is informative.

Consider a single voxel pierced by three distinct viewing rays, as shown in Fig. 2 (left). After passing through the single voxel, the viewing rays have been occluded with probabilities $P(\mathcal{Q}_{A1})$, $P(\mathcal{Q}_{B1})$, and $P(\mathcal{Q}_{C1})$, respectively. Given that no further information about the volume is obtained, the occlusion probabilities should not change if the voxel is subdivided to provide finer resolution, as shown in Fig. 2 (right). In other words, the occlusion probabilities should not be inherently tied to the level of discretization.

Using traditional probabilistic methods, $P(\mathcal{Q}_{A1}) = P(\mathcal{Q}_{B1}) = P(\mathcal{Q}_{C1}) = P(X \in S)$, where $P(X \in S)$ is the probability that the voxel belongs to the set S of occupied voxels. Upon subdivision of the voxel, eight new voxels are created, each of which must be assigned a surface probability $P(X_{child} \in S)$. Whatever the probability chosen, it is assumed to be constant among the eight "child" voxels since there is no reason for favoring one over any of the others. Given that

rays A, B, and C pass through four child voxels, two child voxels, and one child voxel, respectively, the new occlusion probabilities are computed as the probability that any of the voxels passed through belong to the set S of surface voxels. This is easily solved using De Morgan's Laws by instead computing the complement of the probability that all voxels passed through are empty

$$P(\mathcal{Q}_{A2}) = 1 - (1 - P(X_{\text{child}} \in S))^4 \tag{1}$$

$$P(\mathcal{Q}_{B2}) = 1 - (1 - P(X_{\text{child}} \in S))^2 \tag{2}$$

$$P(\mathcal{Q}_{C2}) = P(X_{\text{child}} \in S). \tag{3}$$

Obviously, the three occlusion probabilities cannot be equal to the original values, i.e., $P(X \in S)$, except in the trivial cases $P(X \in S) = P(X_{\text{child}} \in S) = 0$ or $P(X \in S) = P(X_{\text{child}} \in S) = 1$. This simple example demonstrates the general impossibility of resampling a traditional voxel-based probabilistic model while maintaining the semantic meaning of the original. This presents a major hurdle to generalizing standard probabilistic 3-D models to variable-resolution representations.

The methods proposed in this paper solve the problems associated with resolution dependence by modeling surface probability as a density field rather than a set of discrete voxel probabilities. The density field is still represented discretely in practice, but the individual voxels can be arbitrarily subdivided without affecting occlusion probabilities since the density is a property of the points within the voxel and not the voxel itself. Occlusion probabilities are computed by integrating the density field along viewing rays, providing a principled way to take voxel/viewing ray intersection lengths into account. The derivation of this density field is presented in Section III.

III. Occlusion Density

In order to offset the prohibitively large storage costs and discretization problems of the regular voxel grid on which traditional probabilistic methods are based, a novel representation of surface probability is proposed in the form of a scalar function termed the *occlusion density*. The occlusion density at a point in space can be thought of as a measure of the likelihood that the point occludes points behind it along the line of sight of a viewer, given that the point itself is unoccluded. More precisely, the occlusion density value at the point is a measure of occlusion probability per unit length of a viewing ray which is passing through the point.

If the occlusion density is defined over a volume, probabilistic visibility reasoning can be performed for any pair of points within the volume. In the case where surface geometry exists and is known completely [e.g., scenes defined by a surface mesh or digital elevation model (DEM)], the occlusion density is defined as infinite at the surface locations and zero elsewhere.

Given a ray in space defined by its origin point $\mathbf{q}$ and a unit direction vector $\mathbf{r}$, the probability of each point $\mathbf{x}$ along the ray being visible from $\mathbf{q}$ may be computed. It is assumed here that $\mathbf{q}$ is the position of a viewing camera and $\mathbf{r}$ represents a viewing ray of the camera, but the assumption is not necessary in general. Points along the line of sight may be parameterized by s, the distance from $\mathbf{q}$

$$\mathbf{x}(s) = \mathbf{q} + s\mathbf{r}, \qquad s \geq 0. \tag{4}$$

Given the point $\mathbf{q}$ and viewing ray $\mathbf{r}$, a proposition $\mathcal{V}_s$ may be defined as follows:

$$\mathcal{V}_s \equiv \text{“The point along } \mathbf{r} \text{ at distance } s \text{ is visible from } \mathbf{q}\text{.”} \tag{5}$$

The probability $P(\mathcal{V}_s)$ is a (monotonically decreasing) function of s and can be written as such using the notation $\text{vis}(s)$

$$\text{vis}(s) \equiv P(\mathcal{V}_s). \tag{6}$$

Given a segment of r with arbitrary length ℓ beginning at the point with distance s from $\mathbf{q}$, the *segment occlusion probability* $P(\mathcal{Q}_s^\ell)$ is defined as the probability that the point at distance $s + \ell$ is not visible, given that the point at distance s is visible

$$\begin{aligned} P\left(\mathcal{Q}_s^\ell\right) &= P(\bar{\mathcal{V}}_{s+\ell} | \mathcal{V}_s) \\ &= 1 - P(\mathcal{V}_{s+\ell} | \mathcal{V}_s). \end{aligned} \tag{7}$$

Using Bayes' theorem

$$P\left(\mathcal{Q}_s^\ell\right) = 1 - \frac{P(\mathcal{V}_s | \mathcal{V}_{s+\ell}) P(\mathcal{V}_{s+\ell})}{P(\mathcal{V}_s)}. \tag{8}$$

Substituting $\text{vis}(s)$ for the visibility probability at distance s and recognizing that $P(\mathcal{V}_s | \mathcal{V}_{s+ds}) = 1$

$$P\left(\mathcal{Q}_s^\ell\right) = 1 - \frac{\text{vis}(s+\ell)}{\text{vis}(s)}$$

$$P\left(\mathcal{Q}_s^\ell\right) = \frac{\text{vis}(s) - \text{vis}(s+\ell)}{\text{vis}(s)}. \tag{9}$$

If an infinitesimal segment length $\ell = ds$ is used, (9) can be written as

$$P\left(\mathcal{Q}_s^{ds}\right) = \frac{-\partial \text{vis}(s)}{\text{vis}(s)} \tag{10}$$

$$\frac{P\left(\mathcal{Q}_s^{ds}\right)}{ds} = -\frac{\text{vis}'(s)}{\text{vis}(s)}. \tag{11}$$

The left-hand side of (11) is a measure of occlusion probability per unit length and defines the occlusion density at point $\mathbf{x}(s)$. The estimation of the occlusion density value is discussed in Section V

$$\alpha\left(\mathbf{x}(s)\right) \equiv -\frac{\text{vis}'(s)}{\text{vis}(s)}. \tag{12}$$

2) Visibility Probability Calculation: The visibility probability of the point at distance s along a viewing ray can be derived in terms of the occlusion density function along the

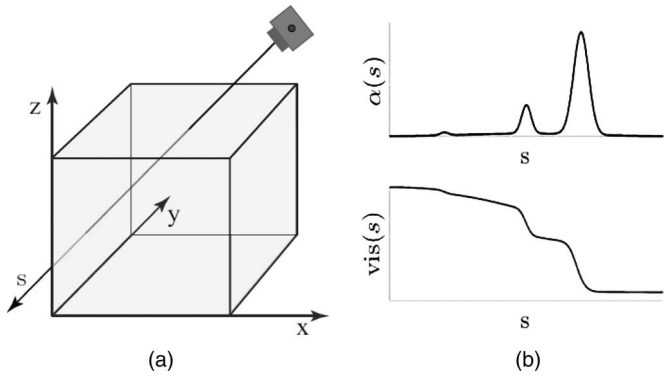

Fig. 3. (a) Viewing ray originating at the camera pierces the volume of interest. (b) (Top) (Arbitrary) Example plot of occlusion density along the viewing ray. (Bottom) Resulting visibility probability along the viewing ray.

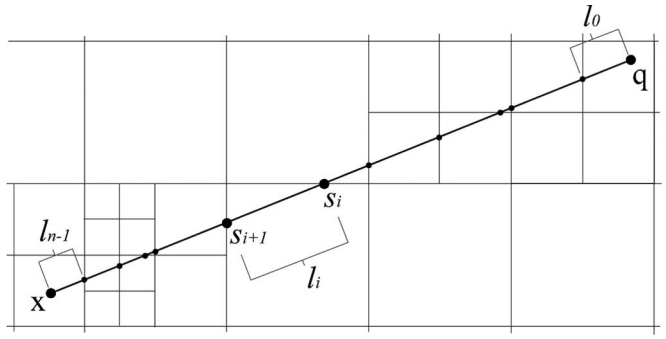

Fig. 4. Camera ray parameterized by s cuts through cells in an octree. Both the occlusion density and appearance model are approximated as being constant within each cell.

ray by integrating both sides of (12) with respect to a dummy variable of integration t

$$\int_0^s \alpha(t)dt = \int_0^s \frac{-\partial \mathrm{vis}(t)}{\mathrm{vis}(t)}$$

$$-\int_0^s \alpha(t)dt = \left[\ln\left(\mathrm{vis}(t)\right) + c\right]_0^s$$

$$-\int_0^s \alpha(t)dt = \ln\left(\mathrm{vis}(s)\right) - \ln\left(\mathrm{vis}(0)\right). \tag{13}$$

Finally, by recognizing that $\mathrm{vis}(0) = 1$ and placing both sides of (13) in an exponential

$$\mathrm{vis}(s) = e^{-\int_0^s \alpha(t)dt}. \tag{14}$$

Equation (14) gives a simple expression for the visibility probability in terms of the occlusion density values along the viewing ray.

An example of corresponding occlusion density and visibility functions is shown in Fig. 3, which depicts a camera ray piercing a theoretical volume for which the occlusion density is defined at each point within the volume. The value of the occlusion density $\alpha(s)$ as a function of distance along the camera ray is plotted, indicating two significant peaks in surface probability. The resulting visibility probability function is plotted directly below it.

3) Occlusion Probability: Substituting (14) back into (9), a simplified expression of a segment's occlusion probability is obtained

$$P\left(\mathcal{Q}_s^\ell\right) = 1 - e^{-\int_s^{s+\ell} \alpha(t)dt}. \tag{15}$$

4) Relationship With Discrete Voxel Probabilities: The key theoretical difference between the discrete voxel probabilities $P(X)$ of existing methods and the preceding formulation of occlusion density is the interpretation of the probability values. Because existing methods effectively model the probability that a voxel boundary is occluding (whether they are defined as such or not), the path length of the viewing ray through the voxel is irrelevant. By contrast, the occlusion probabilities $P(\mathcal{Q}_s^\ell)$ represent the probability that the viewing ray is occluded at any point on the interval $[s, s+\ell]$. The path length ℓ becomes important when one moves away from high-resolution regular voxel grids to variable-resolution models because its value may vary greatly depending on the size of the voxel and the geometry of the ray-voxel intersection.

IV. Implementation: Octree Representation

In order to make practical use of the probabilistic model described in Section III, a finite-sized representation which is able to associate both an occlusion density and appearance model with each point in the working volume is needed. Details are presented in this section of an octree-based implementation which approximates the underlying occlusion density and appearance functions as piecewise constant.

Most real-world scenes contain large slowly varying regions of low occlusion probability in areas of "open space" and high quickly varying occlusion probability near "surfacelike" objects. It therefore makes sense to sample $\alpha(\mathbf{x})$ in a nonuniform fashion. The proposed implementation approximates both $\alpha(\mathbf{x})$ and the appearance model as being piecewise constant, with each region of constant value represented by a cell of an adaptively refined octree. The implementation details of the underlying octree data structure itself are beyond the scope of this paper; the reader is referred to Samet's comprehensive treatment [26]. Fig. 4 shows a viewing ray passing through a volume which has been adaptively subdivided using an octree data structure and the finite-length ray segments that result from the intersections with the individual octree cells.

Each cell in the octree stores a single occlusion density value α and appearance distribution $p_A(\mathbf{i})$. The appearance distribution represents the probability density function of the pixel intensity value resulting from the imaging of the cell. The occlusion density value and appearance distribution are assumed to be constant within the cell. Note that this piecewise-constant assumption can be made arbitrarily accurate since, in theory, any cell in which the approximation is not acceptable can always be subdivided into smaller cells. In practice, however, the amount of useful resolution in the model is limited by the resolution of the input data used to construct it.

A. Appearance

In addition to the occlusion density, it is necessary to define appearance information for points in space in order to perform modeling and rendering operations. Here, "appearance" describes the predicted pixel value of an imaging sensor, given that the point is visible from the sensor. In general, many factors contribute to this value: lighting conditions, viewing angle, particularities of the imaging sensor, and others. There is a large body of work in the computer graphics literature that is focused on modeling these factors' effect on an object's appearance, a comprehensive survey of which is given by Dorsey *et al.* [27]. For practical reasons, the appearance model used in this work is relatively simple; it is independent of viewing direction and is modeled using a mixture of Gaussian distributions. It is assumed that for a given point $\mathbf{x}$, a single distribution $p_A(c, \mathbf{x})$ describes the probability density of the imaged value c of $\mathbf{x}$. In the case of grayscale imagery, c is a scalar value and the distribution is 1-D. In the case of multi-band imagery (e.g., RGB color), c is an n-dimensional vector and the distribution is defined over the n appearance dimensions. When dealing with data such as satellite imagery in which vastly different lighting conditions may occur, a separate appearance model is used for each predefined range of illumination conditions (i.e., sun positions). When performing change detection, only the relevant appearance model for the given sun position is considered. This is the approach taken by Pollard *et al.* [2].

V. Reconstruction From Imagery

Pollard *et al.* [2] estimate the probability that each voxel X of the model belongs to the set S of "surface voxels." This "surface probability" is denoted as $P(X \in S)$ or simply $P(X)$ for convenience. The voxel surface probabilities are initialized with a predetermined prior and updated with each new observation using an online Bayesian update equation (16). The update equation determines the posterior surface probabilities of each of a series of voxels along a camera ray, given their prior probabilities $P(X)$ and an observed image $\mathcal{D}$

$$P(X|\mathcal{D}) = P(X)\frac{P(\mathcal{D}|X)}{P(\mathcal{D}_t)}. \tag{16}$$

The marginal $(P(\mathcal{D}_t))$ and conditional $(P(\mathcal{D})|X))$ probabilities of observing $\mathcal{D}$ can be expanded as a sum of probabilities along the viewing ray. In practice, a single camera ray is traversed for each pixel in image t and all pixel values are assumed to be independent.

Rather than a camera ray intersecting a series of voxels, the equation can be generalized to a series of N intervals along a ray parameterized by s, the distance from the camera center. The ith interval is the result of the intersection of the viewing ray with the ith octree cell along the ray, and it has length ℓ_i. (See Fig. 4.) The interval lengths resulting from the voxel–ray intersections are irrelevant in the formulation of Pollard *et al.* because the occlusion probabilities are fixed and equal to the discrete voxel surface probabilities $P(X_i)$. The surface probability of the ith voxel is replaced by $P(\mathcal{Q}_{s_i}^{\ell_i})$, the occlusion probability (15) of the ith interval. The integral is replaced by a multiplication due to the assumption that occlusion density has constant value α_i within the cell

$$P\left(\mathcal{Q}_{s_i}^{\ell_i}\right) = 1 - e^{-\alpha_i \ell_i}. \tag{17}$$

The probability that the start point of the ith interval is visible from the camera is denoted by the term $\text{vis}(s_i)$ and can be simplified using the piecewise-constant model as follows:

$$\text{vis}(s_i) = \exp\left(-\sum_{j=0}^{i-1} \alpha_j \ell_j\right). \tag{18}$$

The ith posterior occlusion probability $P(\mathcal{Q}_{s_i}^{\ell_i}|\mathcal{D})$ is computed by following closely the formulation of Pollard *et al.* [2] and generalizing to a series of intervals with varying lengths

$$P\left(\mathcal{Q}_{s_i}^{\ell}|\mathcal{D}\right) = P\left(\mathcal{Q}_{s_i}^{\ell}\right)\frac{P\left(\mathcal{D}|\mathcal{Q}_{s_i}^{\ell}\right)}{P(\mathcal{D})}. \tag{19}$$

In order to simplify computations, the term pre_i, which represents the probability of observing $\mathcal{D}$ taking into account segments 0 to $i-1$ only is defined as follows:

$$\text{pre}_i \equiv \sum_{j=0}^{i-1} P\left(\mathcal{Q}_{s_j}^{\ell}\right) \text{vis}(s_j) p_{A_j}(c_{\mathcal{D}}). \tag{20}$$

The term $p_{A_j}(c_{\mathcal{D}})$ is the probability density of the viewing ray's corresponding pixel intensity value $c_{\mathcal{D}}$ given by the appearance model of the ith octree cell along the ray. Equation (19) can then be generalized as

$$\text{vis}_\infty \equiv \prod_{i=0}^{N-1}\left[1 - P\left(\mathcal{Q}_{s_i}^{\ell}\right)\right] \tag{21}$$

$$P\left(\mathcal{Q}_{s_i}^{\ell}|\mathcal{D}\right) = P\left(\mathcal{Q}_{s_i}^{\ell}\right)\frac{\text{pre}_i + \text{vis}(s_i)p_{A_i}(c_{\mathcal{D}})}{\text{pre}_\infty + \text{vis}_\infty p_{A_\infty}(c_{\mathcal{D}})} \tag{22}$$

where pre_∞ represents the total probability of the observation based on all voxels along the ray. The probability of the ray passing unoccluded through the model is represented by vis_∞ and is computed based on (21). The term $p_{A_\infty}(c_{\mathcal{D}})$ represents the probability density of the observed intensity given that the ray passes unoccluded through the volume and can be thought of as a "background" appearance model. In practice, portions of the scene not contained in the volume of interest may be visible in the image. In this case, the background appearance model represents the predicted intensity distribution of these points and is nominally set to a uniform distribution. Note that the denominator of (22) differs from the update equation of Pollard *et al.* [2] due to consolidation of the "pre" and "post" terms into pre_∞ and the addition of the "background" term $\text{vis}_\infty p_{A_\infty}(c_{\mathcal{D}})$.

Equation (22) provides a method for computing the posterior occlusion density of a viewing ray segment but must be related back to the cell's occlusion density value to be of use. This is

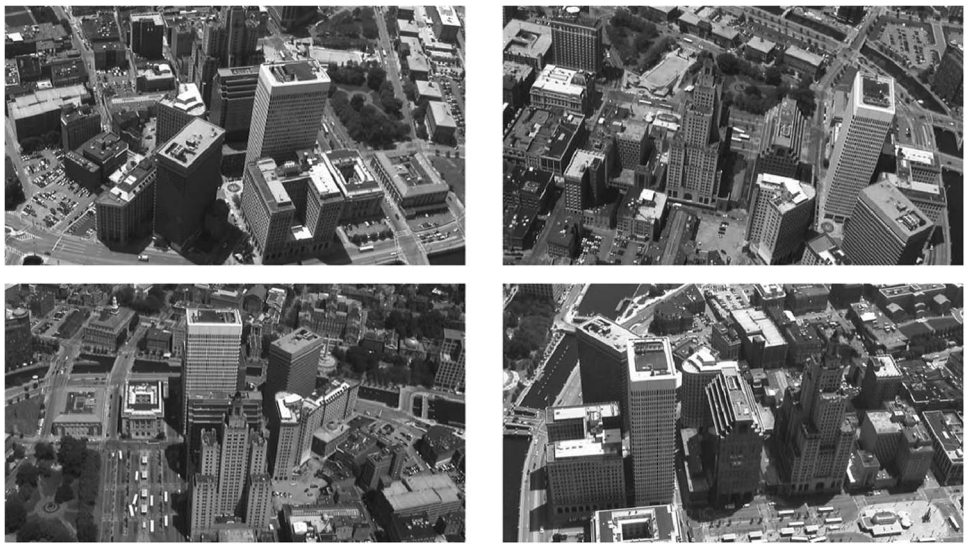

Fig. 5. Four representative frames from the "downtown" video sequence filmed over Providence, RI. Images courtesy of Brown University. Because of the short duration of the sequence, moving objects such as vehicles dominate the changes from image to image.

accomplished by assuming that $\alpha_i = \alpha(t)$ is constant within the segment $(s_i <= t < s_i + \ell_i)$ and by solving (17) for α

$$\alpha_i = \frac{-\log\left(1 - P\left(\mathcal{Q}^{\ell}_{s_i}\right)\right)}{\ell_i}. \tag{23}$$

In general, however, multiple viewing rays from an image will pass through a single octree cell. (This happens any time a cell projects to more than one pixel in the image.) In order to reconcile the multiple posterior probabilities, an occlusion density value is calculated by computing the probability $P(\bar{\mathcal{Q}})$ of occlusion by *any* of the K viewing ray segments within a cell and relating it to the total combined length $\bar{\ell} = \sum_k \ell_k$ of all of the segments. This has the effect of giving more weight to observations with longer corresponding segment lengths

$$P(\bar{\mathcal{Q}}) = 1 - \prod_{k=0}^{K-1} \left(1 - P(\mathcal{Q}_k)\right). \tag{24}$$

The posterior occlusion probability corresponding to the segment of the kth camera ray passing through the cell is denoted $P(\mathcal{Q}_k)$ for convenience

$$\bar{\alpha} = \frac{-\log\left(\prod_{k=0}^{K-1}\left(1 - P(\mathcal{Q}_k)\right)\right)}{\sum_{k=0}^{K-1} \ell_k}$$

$$\bar{\alpha} = \frac{-\sum_{k=0}^{K-1} \log\left(1 - P(\mathcal{Q}_k)\right)}{\sum_{k=0}^{K-1} \ell_k}. \tag{25}$$

The previous occlusion density value of each cell is replaced with the $\bar{\alpha}$ value computed using all K viewing rays of the given image which pass through it.

A. Appearance Calculation

In addition to computing the occlusion density of a given octree cell, a reconstruction algorithm must also estimate the cell's appearance model. Pollard *et al.* update the Gaussian mixture model of each voxel with each image using an online approach based on Stauffer and Grimson's adaptive background modeling algorithm [28], and a similar approach is used here. Because the proposed multiresolution model allows for the possibility of large cells which may project to many pixels of the image, a weighted average of pixel values is computed for each cell prior to the update, with the weights proportional to the segment lengths of the corresponding viewing ray/octree cell intersections. The cell's Gaussian mixture model is then updated using the single weighted average pixel value.

B. Adaptive Refinement

Upon initialization, the model consists of a regular 3-D array of octrees, each containing a single (i.e., root) cell. At this stage, the model roughly corresponds to a regular voxel grid at a coarse resolution. As more information is incorporated into the model, however, the sampling of regions with high occlusion density may be refined. The proposed implementation subdivides a leaf cell into its eight children when its maximum occlusion probability $P(\mathcal{Q}_{\max_i})$ reaches a global threshold. The maximum occlusion probability of cell i is a function of the longest possible path through the cell (i.e., between opposite corners) and the cell's occlusion density

$$P\left(\mathcal{Q}_{max_i}\right) = 1 - \exp\left(-\ell_{max_i}\alpha_i\right). \tag{26}$$

The occlusion densities and appearance models of the refined cells are initialized to the value of their parent cell. This process is executed after each update to the model.

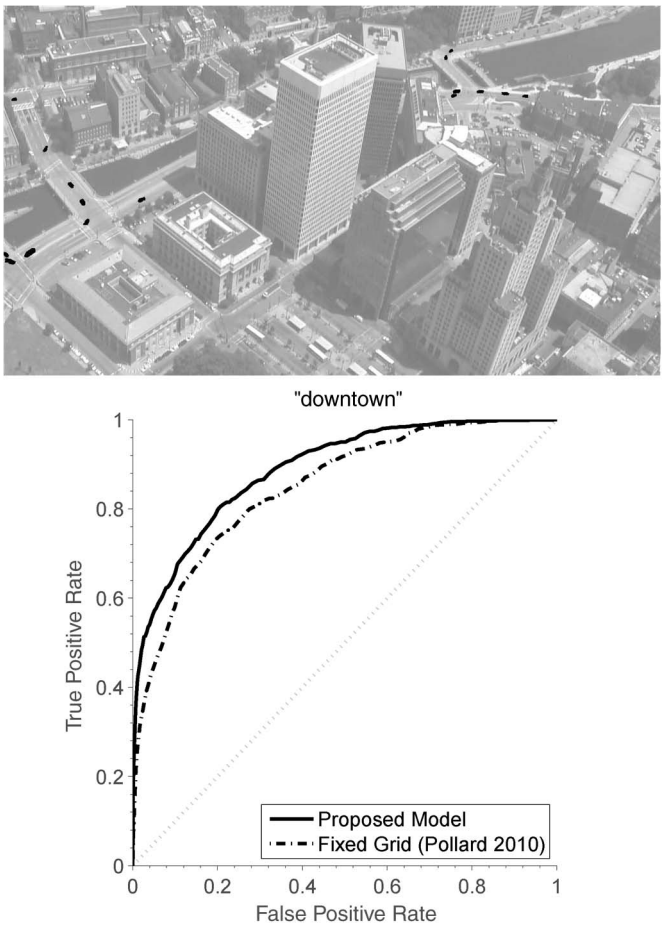

Fig. 6. (Top) Ground truth segmentation used to evaluate the change detection algorithms. Original image courtesy of Brown University. (Bottom) ROC curves for change detection using the proposed 3-D model and a fixed grid voxel model.

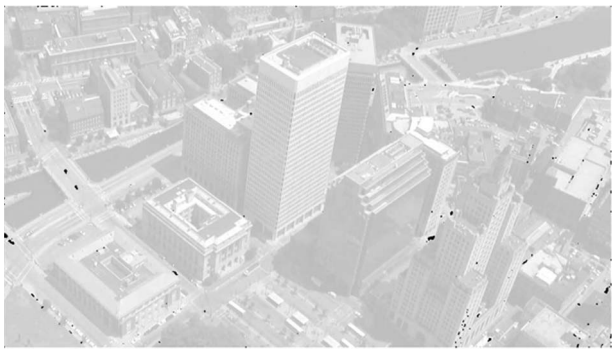

Fig. 7. Changes detected by the proposed algorithm are marked in black. Note that the moving vehicles have low probability, as well as some false detects around building edges, presumably due to small errors in the camera model. Original image courtesy of Brown University.

VI. Change Detection

Given a new image of a previously modeled scene, it is often useful to detect regions in the image which represent deviations from the expected intensity values. Pollard *et al.* [2] demonstrated their system, which was the first to use probabilistic 3-D information for the purposes of change detection,

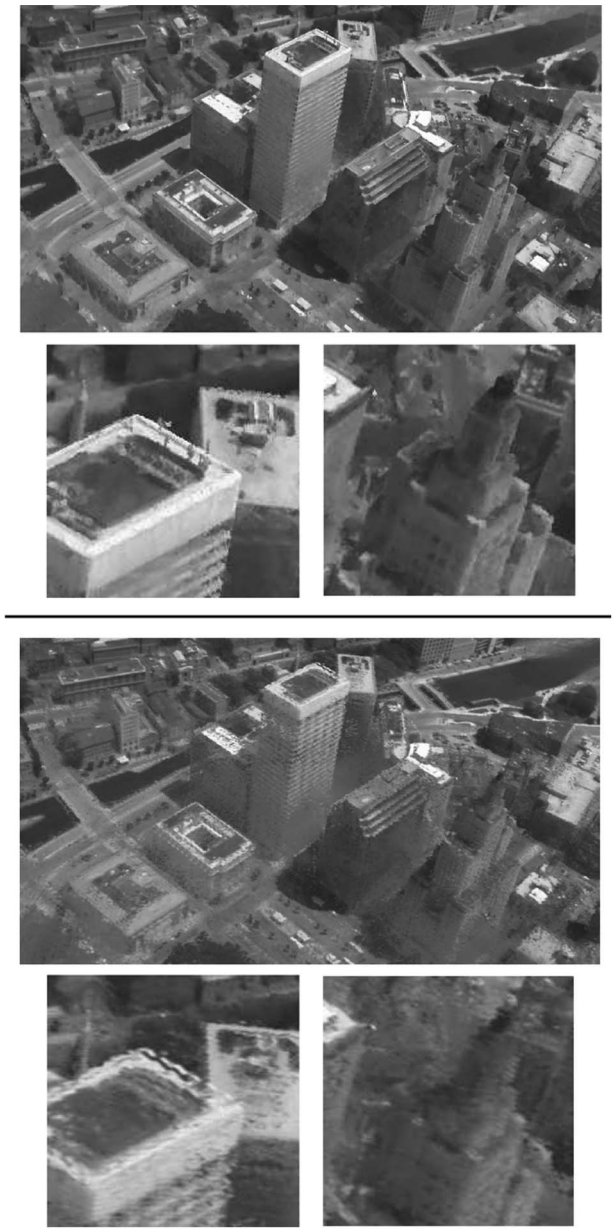

Fig. 8. Expected images generated from the viewpoint of the image used to evaluate change detection. Because there is a low probability of a moving vehicle being at a particular position on the roads, they appear empty. (Top) Image generated using the proposed 3-D model. (Bottom) Image generated using a fixed grid voxel model. The proposed variable-resolution model allows for finer details to be represented, as is visible in the expanded crops.

to be superior to previous 2-D approaches. In particular, the capability to model a 3-D structure allowed apparent changes due to viewpoint-based occlusions to be correctly ignored. A popular 2-D approach [3] consistently marked these regions as change, leading to high false positive rates. The proposed system utilizes the estimated probabilistic 3-D model in a similar fashion, but with the added advantage of higher resolution capability over broader areas that the multiresolution representation provides.

Given a model constructed using the methods proposed in Section V and a viewpoint defined by a camera model, a probability distribution $p_A(c)$ for each pixel in the image can

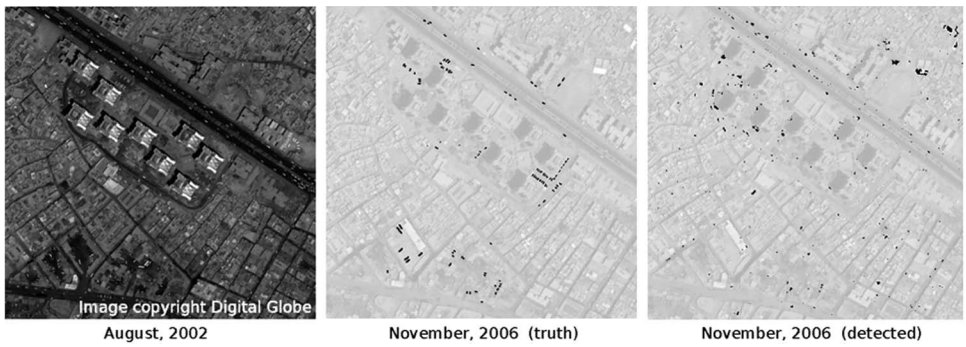

Fig. 9. (Left) Training image of the Haifa Street region from 2002. (Middle) Ground truth changes marked in one of the test images from 2007. Most of the changes are a result of moving vehicles on the roads. (Right) Changes detected using the proposed algorithm. Original image copyright Digital Globe.

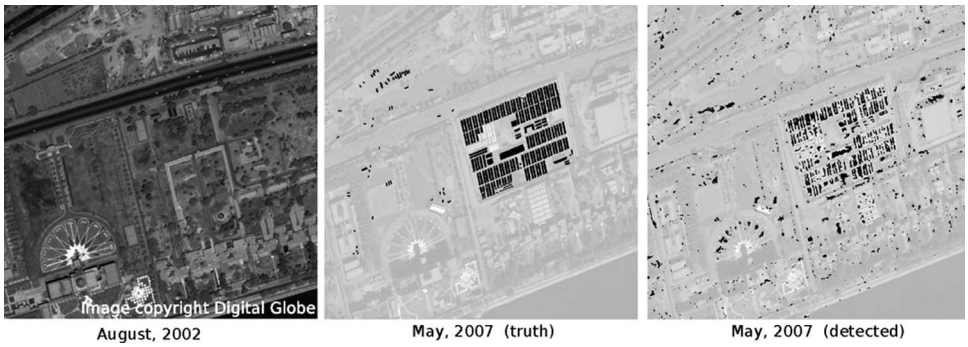

Fig. 10. (Left) Training image of the region surrounding the construction of the U.S. embassy from 2002. (Middle) Ground truth changes marked in one of the test images from 2007. There is a variety of changes resulting in the placement of large storage containers and other construction apparatus. (Right) Changes detected using the proposed algorithm. Original image copyright Digital Globe.

be generated by integrating the appearance information of each voxel along the pixel's corresponding viewing ray R

$$p_A(c) = \sum_{i \in R} \text{vis}(s_i)(1 - e^{-\alpha_i \ell_i}) p_{A_i}(c). \tag{27}$$

Pixel values with a low probability density are "unexpected." A binary "change mask" can therefore be created by thresholding the image of pixel probability densities. The receiver operating characteristic (ROC) curves in Figs. 6, 11, and 12 show the rate of ground-truth change pixels correctly marked as change versus the rate of pixels incorrectly identified as change. The plotted curves indicate how these values vary as the threshold τ is varied.

Two distinct collection types are investigated: full motion video collected from an aerial platform and satellite imagery. The full motion video was collected using a high-definition (1280×720 pixels) video camera from a helicopter flying over Providence, RI, U.S. A few representative frames are shown in Fig. 5. The model was updated using 175 frames of a sequence (in random order) in which the helicopter made a full 360° loop around a few blocks of the downtown area, and the change detection algorithm was then run on an image (not used in the training set) that contains some moving vehicles (Fig. 7).

The ROC curve in Fig. 6 demonstrates the advantage that higher resolution 3-D modeling provides to the change detection algorithm. In order to better visualize the higher resolution capability of the proposed model, a synthetic image can be generated by computing the expected intensity value $E[p_A(c)]$ of each pixel based on the computed probability density distributions of each pixel. Fig. 8 shows expected images generated using the proposed model and the fixed grid model. Small features such as individual building windows and rooftop air-conditioning units are visible using the proposed model but blend into the background using the fixed grid model. The resolution of the fixed grid model can, in theory, always be increased to match the capabilities of the variable-resolution model, but doing so quickly becomes impractical due to the $O(n^3)$ storage requirements. The fixed grid model of the "downtown" sequence is nearly 50% larger than the variable-resolution model while providing half the effective resolution. A fixed grid model equaling the effective resolution of variable-resolution model would require approximately 18 GB (1000% larger than the variable-resolution model).

The satellite imagery used for experimentation was collected by Digital Globe's Quickbird satellite over Baghdad, Iraq, from 2002 to 2007 and are the same data sets used by Pollard *et al.* [2] in their change detection experiments. Two areas of interest are focused on the following: a region with some high-rise buildings along Haifa Street and the region surrounding the construction site of the new U.S. embassy building. A manually determined translation is applied to bring the supplied camera model to within approximately one pixel of reprojection error. The images have a nominal resolution of approximately 0.7-m GSD. Although the 3-D structure is less pronounced than in the video sequence, it is still sufficient to pose a challenge to 2-D change detection algorithms, as shown by Pollard *et al.* The "haifa" and "embassy" models were updated with each of

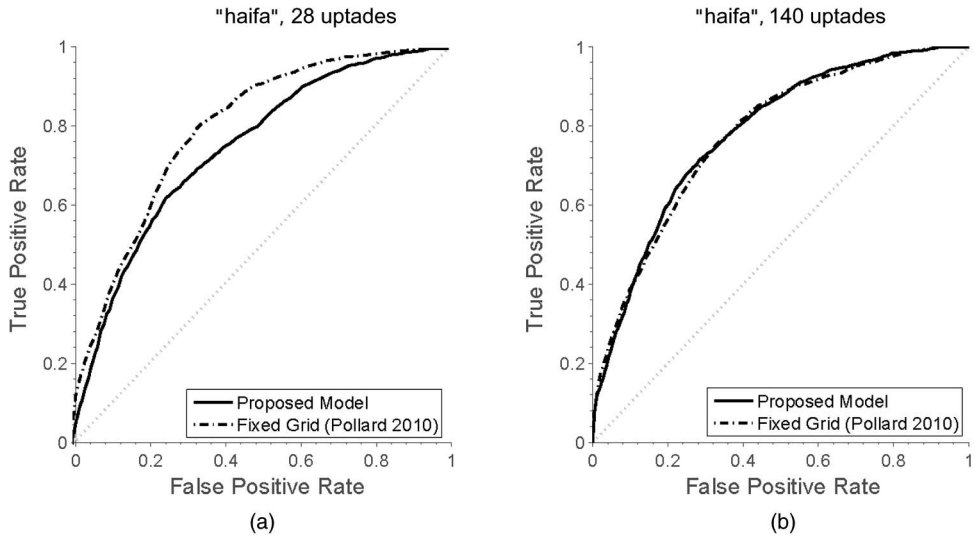

Fig. 11. (a) Change detection ROC curve for the "haifa" data set after (a) one pass and (b) five passes of the 28 training images. The additional passes allow the octrees in the variable-resolution model time to reach their optimal subdivision levels.

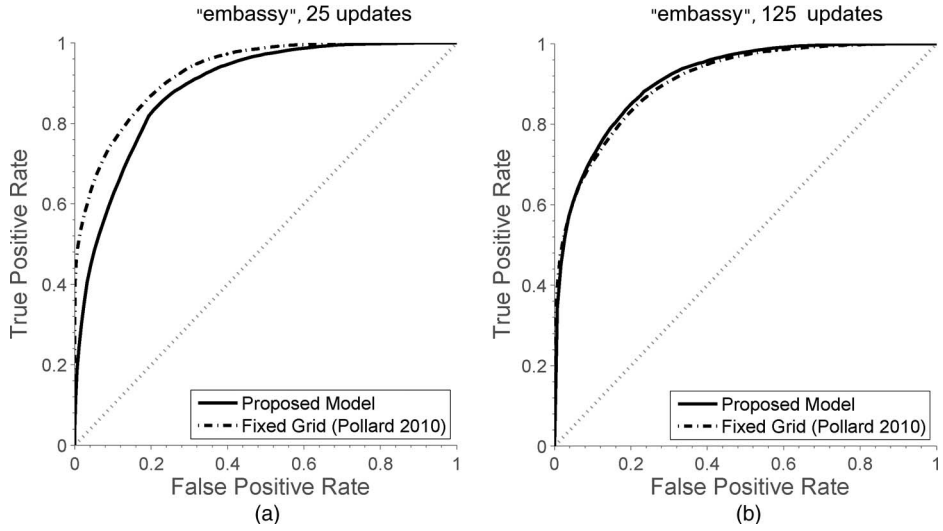

Fig. 12. (a) Change detection ROC curve for the "embassy" data set after (a) one pass and (b) five passes of the 25 training images. After the additional passes, the variable-resolution model has again reached the effective resolution of the fixed grid model.

28 and 25 images, respectively, taken between 2002 and 2006. Although this number of images appears to be sufficient to train the fixed-grid model, a larger number of images allows the cells of the variable-resolution model time to reach optimal subdivision levels. In order to accommodate this requirement, five passes of the images were used to update the models to simulate a larger number of collected images. The change detection algorithm was then run on previously unobserved images taken during 2007. Figs. 9 and 10 show a representative training image and the ground truth changes for one of the 2007 images for the two experiments. Figs. 11 and 12 show the resulting change detection ROC curves resulting after training the model using one pass of the images and after five passes. Because the variable-resolution models are initialized at a much coarser level than the fixed grid model, their performance suffers in the period before the proper subdivision at surface voxels occurs. After five passes, however, the effective resolution has reached that of the fixed grid models while needing a fraction of the memory and disk space. Table I lists the area modeled, effective resolution, and storage requirements of each of the models discussed. Fig. 13 provides a visualization of the storage

TABLE I
COVERAGE AND MODEL SIZES FOR THE AERIAL VIDEO AND SATELLITE IMAGE TEST SETS. THE RESOLUTION LISTED FOR THE VARIABLE-RESOLUTION MODELS INDICATES THAT OF THE FINEST OCTREE SUBDIVISION LEVEL

dataset	coverage	resolution	size
"downtown" (fixed)	0.3 km^2	512 × 512 × 128	2.25 GB
"downtown" (var.)	0.3 km^2	1024 × 1024 × 256	1.53 GB
"haifa" (fixed)	0.51 km^2	512 × 512 × 110	5.58 GB
"haifa" (var.)	0.51 km^2	512 × 512 × 108	1.04 GB
"embassy" (fixed)	0.51 km^2	512 × 512 × 80	4.06 GB
"embassy" (var.)	0.51 km^2	512 × 512 × 80	0.86 GB

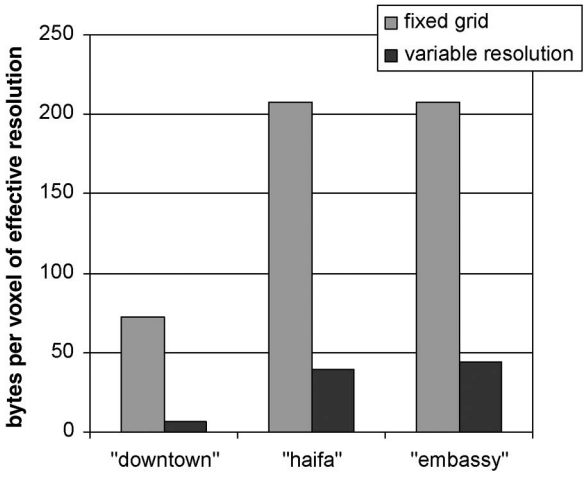

Fig. 13. Storage costs of the models normalized by the number of voxels of effective resolution. The "downtown" model requires less memory per voxel in both cases because only a single appearance distribution is needed to model the constant illumination condition across video frames.

costs of the models, normalized by their respective effective resolutions.

VII. Conclusion

The proposed novel probabilistic 3-D model allows for representations which sample the 3-D volume in a nonuniform fashion by providing a principled method for taking into account voxel/viewing ray intersection lengths. The model is a generalization of previous probabilistic approaches which makes feasible the exploitation of currently available high-resolution remote sensing imagery. Experiments have shown that this capability provides a distinct advantage for the application of change detection over previously existing models. Future work will focus on appearance models that can more efficiently model large changes in illumination conditions, as well as methods to extract surface information (e.g., DEMs) from the probabilistic models.

References

[1] R. J. Radke, S. Andra, O. Al-Kofahi, and B. Roysam, "Image change detection algorithms: A systematic survey," *IEEE Trans. Image Process.*, vol. 14, no. 3, pp. 294–307, Mar. 2005.

[2] T. Pollard, I. Eden, J. Mundy, and D. Cooper, "A volumetric approach to change detection in satellite images," *Photogramm. Eng. Remote Sens.*, vol. 76, no. 7, pp. 817–831, Jul. 2010.

[3] M. Carlotto, "Detection and analysis of change in remotely sensed imagery with application to wide area surveillance," *IEEE Trans. Image Process.*, vol. 6, no. 1, pp. 189–202, Jan. 1997.

[4] C. Benedek and T. Szirányi, "Change detection in optical aerial images by a multilayer conditional mixed Markov model," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 10, pp. 3416–3430, Oct. 2009.

[5] T. Celik and K.-K. Ma, "Unsupervised change detection for satellite images using dual-tree complex wavelet transform," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 3, pp. 1199–1210, Mar. 2010.

[6] A. Huertas and R. Nevatia, "Detecting changes in aerial views of man-made structures," in *Proc. 6th ICCV*, 1998, pp. 73–80.

[7] I. Eden and D. B. Cooper, "Using 3D line segments for robust and efficient change detection from multiple noisy images," in *Proc. Comput. Vis.—ECCV*, vol. 5305, *Lecture Notes in Computer Science*, D. Forsyth, P. Torr, and A. Zisserman, Eds., 2008, Springer: Berlin, Germany.

[8] R. Bhotika, D. J. Fleet, and K. N. Kutulakos, "A probabilistic theory of occupancy and emptiness," in *Proc. Eur. Conf. Comput. Vis.*, 2002, pp. 112–132.

[9] R. Hartley and A. Zisserman, *Multiple View Geometery in Computer Vision*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2003.

[10] D. Scharstein and R. Szeliski, "A taxonomy and evaluation of dense two-frame stereo correspondence algorithms," *Int. J. Comput. Vis.*, vol. 47, no. 1–3, pp. 7–42, Apr.–Jun. 2002.

[11] Z.-F. Wang and Z.-G. Zheng, "A region based stereo matching algorithm using cooperative optimization," in *Proc. Comput. Vis. Pattern Recognit.*, 2008, pp. 1–8.

[12] Q. Yang, L. Wang, R. Yang, H. Stewénius, and D. Nistér, "Stereo matching with color-weighted correlation, hierarchical belief propagation, and occlusion handling," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 3, pp. 492–504, Mar. 2009.

[13] M. Pollefeys, L. V. Gool, M. Vergauwen, F. Verbiest, K. Cornelis, J. Tops, and R. Koch, "Visual modeling with a hand-held camera," *Int. J. Comput. Vis.*, vol. 59, no. 3, pp. 207–232, Sep./Oct. 2004.

[14] C. Harris and M. Stephens, "A combined corner and edge detector," in *Proc. 4th Alvey Vis. Conf.*, 1988, pp. 147–151.

[15] M. Brown and D. Lowe, "Unsupervised 3D object recognition and reconstruction in unordered datasets," in *Proc. 5th Int. Conf. 3DIM*, 2005, pp. 56–63.

[16] N. Snavely, S. M. Seitz, and R. Szeliski, "Modeling the world from internet photo collections," *Int. J. Comput. Vis.*, vol. 80, no. 2, pp. 189–210, Nov. 2008.

[17] D. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, Nov. 2004.

[18] M. Goesele, N. Snavely, B. Curless, H. Hoppe, and S. M. Seitz, "Multi-view stereo for community photo collections," in *Proc. ICCV*, 2007, pp. 1–8.

[19] Y. Furukawa and J. Ponce, "Accurate, dense, and robust multi-view stereopsis," in *Proc. Comput. Vis. Pattern Recognit.*, 2007, pp. 1–8.

[20] S. M. Seitz and C. R. Dyer, "Photorealistic scene reconstruction by voxel coloring," in *Proc. CVPR*, 1997, pp. 1067–1073.

[21] K. N. Kutulakos and S. M. Seitz, "A theory of shape by space carving," *Int. J. Comput. Vis.*, vol. 38, no. 3, pp. 199–218, 2000.

[22] G. Slabaugh, B. Culbertson, T. Malzbender, and R. Schafer, "A survey of methods for volumetric scene reconstruction from photographs," in *Proc. Int. Workshop Volume Graph.*, 2001, pp. 81–100.

[23] G. G. Slabaugh, W. B. Culbertson, T. Malzbender, M. R. Stevens, and R. W. Schafer, "Methods for volumetric reconstruction of visual scenes," *Int. J. Comput. Vis.*, vol. 57, no. 3, pp. 179–199, May/Jun. 2004.

[24] J. S. De Bonet and P. Viola, "Roxels: Responsibility weighted 3d volume reconstruction," in *Proc. Int. Conf. Comput. Vis.*, 1999, pp. 418–425.

[25] A. Broadhurst, T. Drummond, and R. Cipolla, "A probabilistic framework for space carving," in *Proc. Int. Conf. Comput. Vis.*, 2001, pp. 388–393.

[26] H. Samet, *Applications of Spatial Data Structures: Computer Graphics, Image Processing, and GIS*, M. A. Harrison, Ed. Reading, MA: Addison-Wesley, 1990.

[27] J. Dorsey, H. Rushmeier, and F. Sillion, *Digital Modeling of Material Appearance*. San Mateo, CA: Morgan Kaufmann, 2007.

[28] C. Stauffer and W. Grimson, "Adaptive background mixture models for real-time tracking," in *Proc. Comput. Vis. Pattern Recognit.*, 1999, pp. 246–252.

Daniel Crispell (M'11) received the B.S. degree in computer engineering from Northeastern University, Boston, MA, in 2003 and the M.S. and Ph.D. degrees in engineering from Brown University, Providence, RI, in 2005 and 2010, respectively.

While at Brown University, his research focused on camera-based devices for 3-D geometry capture, aerial video registration, and 3-D modeling and rendering from aerial and satellite imagery. He is currently a Visiting Scientist at the National Geospatial-Intelligence Agency, Springfield, VA.

Joseph Mundy received the B.S. and Ph.D. degrees in electrical engineering from Rensselaer Polytechnic Institute, Troy, NY, in 1963 and 1969, respectively.

He joined General Electric Global Research in 1963. In his early career at GE, he carried out research in solid state physics and integrated circuit devices. In the early 1970s, he formed a research group on computer vision with emphasis on industrial inspection. His group developed a number of inspection systems for GE's manufacturing divisions, including a system for the inspection of lamp filaments that exploited syntactic methods in pattern recognition. During the 1980s, his group moved toward more basic research in object recognition and geometric reasoning. In 1988, he was named a Coolidge Fellow, which awarded him a sabbatical at Oxford University, Oxford, U.K. At Oxford, he collaborated on the development of theory and application of geometric invariants. In 2002, he retired from GE Global Research and joined the School of Engineering, Brown University, Providence, RI. At Brown University, his research is in the area of video analysis and probabilistic computing.

Gabriel Taubin (M'86–F'01) received the Licenciado en Ciencias Matemáticas degree from the University of Buenos Aires, Buenos Aires, Argentina, and the Ph.D. degree in electrical engineering from Brown University, Providence, RI.

In 1990, he joined IBM, where during a 13-year career in the Research Division, he held various positions, including Research Staff Member and Research Manager. In 2003, he joined the School of Engineering, Brown University, as an Associate Professor of Engineering and Computer Science. While on sabbatical from IBM during the 2000–2001 academic year, he was appointed Visiting Professor of Electrical Engineering at the California Institute of Technology, Pasadena. While on sabbatical from Brown during the spring semester of 2010, he was appointed Visiting Associate Professor of Media Arts and Sciences at Massachusetts Institute of Technology, Cambridge. He serves as a member of the editorial board of the *Geometric Models* journal. He has made significant theoretical and practical contributions to the field now called Digital Geometry Processing: to 3-D shape capturing and surface reconstruction and to geometric modeling, geometry compression, progressive transmission, signal processing, and display of discrete surfaces. The 3-D geometry compression technology that he developed with his group was incorporated into the MPEG-4 standard and became an integral part of IBM products.

Prof. Taubin is the current Editor-in-Chief of the IEEE COMPUTER GRAPHICS AND APPLICATIONS MAGAZINE and has served as Associate Editor of the IEEE TRANSACTIONS OF VISUALIZATION AND COMPUTER GRAPHICS. He was named IEEE Fellow for his contributions to the development of 3-D geometry compression technology and multimedia standards, won the Eurographics 2002 Günter Enderle Best Paper Award, and was named IBM Master Inventor.