Detailed 3D Reconstruction of Large-Scale Heritage Sites with Integrated Techniques

Sabry F. El-Hakim, J.-Angelo Beraldin, Michel Picard, and Guy Godin
National Research Council of Canada

Many cultural heritage applications require 3D reconstruction of real-world objects and scenes. Over the past few years, it has become increasingly common to use 3D digitization and modeling for this purpose. This is mainly due to advances in laser-scanning techniques, 3D modeling software, image-based modeling techniques, computer power, and virtual reality. Many approaches are currently available. The most common are based on surveys and CAD tools and/or traditional photogrammetry with control points and a human operator. However, this approach is time-consuming and can be costly and impractical for large-scale sites. Modeling methods based on laser-scanned data and more automated image-based techniques have recently become available.

Our approach integrates several technologies based on our experience over more than a decade of trying to accurately and completely model large-scale heritage monuments and sites. Using both interactive and automatic techniques, we can model a highly detailed structure or site at various levels of detail. We use image-based modeling for basic shape and structural elements, and laser scanning for fine details and sculpted surfaces. To present the site in its proper context, we use image-based rendering for landscapes and surroundings. To apply this approach, we created hundreds of models from sites all over the world for documentation, walk-through movies, and interactive visualization. The results were compelling and encouraging.

Motivation and requirements

There are many motives for 3D reconstruction of heritage sites:

- documenting historic buildings and objects for reconstruction or restoration in case of fire, earthquake, flood, war, erosion, and so on;
- creating educational resources for history and culture students and researchers;
- reconstructing historic monuments that no longer or only partially exist;
- visualizing scenes from viewpoints impossible in the real world due to size or accessibility issues;
- interacting with objects without risk of damage; and
- providing virtual tourism and virtual museum exhibits.

An integrated approach for selecting the most effective technique for modeling large-scale heritage sites combines several technologies to create accurate and complete models.

In general, most applications specify eight requirements: high geometric accuracy, capture of all details, photorealism, high automation level, low cost, portability, application flexibility, and model size efficiency. The order of importance of these requirements depends on the application's objective—for example, whether it's for documentation or virtual tourism. A single system that can satisfy all eight requirements is still in the future. In particular, accurately capturing all details with a fully automated system for a wide range of objects and scenes remains elusive.

For small and medium objects, up to the size of an average adult person, range-based techniques such as laser scanners can provide accurate and complete details with a high degree of automation. However, being relatively new systems that aren't produced in large quantities, these scanners remain costly. They are also not portable enough for a single person to carry around and use like a video camera. Moreover, the resulting model can be inefficient for interactive visualization of large-scale scenes.

Image-based approaches use widely available hardware. The same system can potentially capture a wide range of objects and scenes. Image-based approaches are also capable of producing realistic models, and those based on photogrammetry have high geometric accuracy. The issues that remain in image-based modeling are the capture of details on unmarked and sculpted surfaces and the full automatic creation of the 3D models. Approaches such as image-based rendering[1] that skip the geometric-modeling step might suffice for visualization and walk-through. However, the lack of geometric modeling makes them unsuitable for documentation and reconstruction applications.

Most documented projects on cultural heritage have used one method or another; only a few have used a combination of techniques. For example, a group from IBM[2] combined structured-light 3D sensing and photometric stereo to model Michelangelo's *Florentine Pietà*. Researchers have also combined laser scanning with image-based modeling and rendering,[3] image-based modeling with image-based rendering,[4] and image-based rendering with laser scanning.[5]

Our approach integrates techniques as follows:

- We construct the basic shape and large regularly shaped details, such as columns, blocks, windows, and archways, from high-resolution digital images. This technique is based on advanced photogrammetry with several automated features that take advantage of properties found in classical architectures.
- We use laser scans to obtain fine geometric details, such as sculpted and irregularly shaped surfaces. Then we integrate this technique with the basic model created in the previous step.
- We obtain visual details in the geometric model from image textures and reflectance models.
- We use panoramas from aerial images to complete the surroundings and distant landscapes. This helps present the monument in its natural setting.

This combination of techniques satisfies most requirements except that the cost of laser scanning is not as low as that of a fully image-based system—at least for now.

Overview of 3D construction techniques

A standard approach to creating a 3D model is to build it from scratch using tools such as CAD software, which offers building blocks in the form of primitive 3D shapes. Some survey data or measurements from drawings and maps are also necessary. However, this geometry-based modeling technique is time-consuming, impractical, and costly for large-scale projects. Although many applications apply this approach—even TV programs in Europe use it to render sites that no longer exist—the created models look computer generated rather than realistic. They also don't include fine details or irregular and sculpted surfaces. Several recent techniques aim to increase the level of automation and realism by starting with actual images of the object or directly digitizing it with a laser scanner.

Image-based modeling

This technique involves widely available hardware, so the same system can potentially handle a broad range of objects and scenes. Such systems can also produce realistic models, and those based on photogrammetry have high geometric accuracy. Deriving 3D measurements from images naturally requires that interest points be visible in the image. Often, this is not possible, either because a region is hidden or occluded behind an object or surface or because there is no mark, edge, or visual feature to extract. In objects such as monuments in their normal settings, restrictions also stem from there being limited locations from which to take the images and from the existence of other objects, shadows, and uncontrolled illumination.

The ultimate goal of all 3D reconstruction methods is to satisfy the eight requirements listed earlier. Because this is difficult, these methods typically focus on some of the tasks at the expense of others. Efforts to increase the level of automation have become essential to widen the use of the technology. However, approaches to completely automate the process from taking images to a 3D model's output, although promising, are thus far not always successful. Some of the steps, mainly the automation of camera-pose estimation and computation of 3D pixel coordinates, have worked well in many cases.

This automation, now prevalent in computer vision,[6] starts with a sequence of images taken by an uncalibrated camera. The vision system extracts interest points (such as corners), sequentially matches them across views, and computes camera parameters and 3D coordinates of the matched points using robust techniques. This approach typically uses the first two images to initialize the sequence. Tracking the points over a long sequence is important for reducing error propagation. This all occurs in a projective geometry, usually followed by a bundle adjustment,[7] also in the projective space. Self-calibration to compute the intrinsic camera parameters—usually only the focal length—follows to obtain metric reconstruction, up to scale, from the projective reconstruction.[6] Again, the system typically applies bundle adjustment to the metric reconstruction to optimize the solution.

The next step, the creation of the 3D model, is more difficult to automate. This step typically requires some human interaction to define the topology and edit or postprocess the output. For large structures and scenes, the technique could require many images. Therefore, creating the model requires significant human interaction even if there was automatic camera-pose estimation and if many 3D point coordinates were fully computed automatically.

The most impressive results remain those achieved with highly interactive approaches. Rather than use full automation, Debevec developed an easy-to-use hybrid system known as Façade.[4] The method's main goal was to realistically create 3D models of architectures from a few photographs. First, the user interactively recovers the structure's basic geometric shape using models of polyhedral elements. In this step, the Façade system captures the actual size of the elements and camera pose, assuming the camera's intrinsic parameters are known. The second step is an automated matching procedure, constrained by the now-known basic model, to add geometric details. This approach proved effective for creating geometrically accurate and realistic models of architectures. The drawbacks are that it requires a high level of interaction and it's restricted to certain shapes. Also, because assumed shapes determine all 3D points and camera poses, the results are as accurate as the underlying assumption that the structure elements match those shapes.

Our method, although conceptually similar, replaces basic shapes with a few seed points to achieve more flexibility and a higher level of detail. In addition, we determine the camera poses and 3D coordinates without any assumption of the shapes, but rather through a full photogrammetric bundle adjustment with or without self-calibration, depending on the given configuration. This achieves higher geometric accuracy independent of assumptions about the object's shape.

The Façade approach has inspired several research activities to automate it. For example, Werner and Zisserman[8] proposed a fully automated Façade-like approach. Rather than using the basic shapes, their system creates the scene's principal planes automatically to assemble a coarse model. As in Façade, the coarse model guides a more refined polyhedral model of details such as windows, doors, and wedge blocks. Because this approach is fully automated, it requires feature detection, closely spaced images to ensure correct matching, and camera-pose estimation using projective geometry.

Range-based modeling

Three-dimensional measurement from images requires that interest points or edges be visible in the image, which is not always possible. Illumination or ambient light problems can also affect the extraction of such points and edges. Active sensors such as laser scanners avoid these limitations by creating features on the surface through controlled light projection. These sensors have the advantage of acquiring dense 3D points automatically.

Recent advances in lasers, charge-coupled device technology, and electronics make detailed shape measurements possible with accuracies better than 1 part per 5,000 at rates exceeding 10,000 points per second. The scanning and imaging configuration determine the point density. Many range sensors also produce organized points, in the form of an array or range image, suitable for automatic modeling. A single range image is usually insufficient to cover an object or structure. The amount of necessary images depends on the object's shape, the amount of self-occlusion and obstacles, and the object's size compared to the sensor range. We must then register the 3D data in a single coordinate system. Several registration techniques are available; most are based on the iterative closest point (ICP) approach. For this approach to converge to the correct solution, it needs to start with the images approximately registered. This requires either knowledge of sensor positions or manual registration using features. Once we've registered the range images in a single coordinate system, they're ready for modeling. The modeling step reduces the large number of 3D points into a triangular mesh, which preserves the geometric details and is suitable for fast rendering.[8] This process must integrate the areas where the images overlap to create a nonredundant mesh. Other requirements include filling holes and removing any outliers.

There are two main types of range sensors: triangular based and those based on the time-of-flight principle. Triangulation-based sensors project light in a known direction from a known position, and measure the direction of the returning light through its detected position. Measurement accuracy depends on the triangle base relative to its height. Because the triangle base is rather short (for practical reasons), such systems have a limited range of less than 10 meters (in fact, most are less than 3 meters). Sensors based on the time-of-flight principle measure the delay between emission and detection of the light reflected by the surface, and thus the accuracy does not rapidly deteriorate as the range increases. Time-of-flight sensors can provide measurements in the kilometer range.

Notwithstanding the advantages of range sensors, they do have some drawbacks. At the moment, accurate systems are costly and bulky, and surface-reflective properties and ambient light affect those that don't use lasers. Range sensors can also be complex to operate and calibrate. In addition, they're intended for a specific range, so one designed for close range is not suitable for long range, and vice versa. For large-scale environments, using a range sensor to model the entire scene can generate a huge amount of data and require considerable effort to register the many scans.

Image-based rendering

IBR uses images directly to generate new views for rendering without explicit geometric representation.[1] This has the advantage of creating realistic virtual environments at speeds independent of scene complexity. IBR relies on either accurately knowing the camera positions or using automatic stereo matching. In the absence of geometric data, success of the latter requires many closely spaced images. Object occlusions and discontinuities, particularly in large-scale and geometrically complex environments, also affect the output. The ability to move freely into the scene and view objects from any position can be limited, depending on the method used. It's therefore unlikely that IBR will be the approach of choice for purposes other than limited visualization. For tourists satisfied with general visualization, this approach might be adequate, but for historians and researchers, and of course for documentation, correct geometric details are essential.

Combining multiple techniques

From the preceding summary of current techniques, it's obvious that none can satisfy all the requirements of large-scale environments for applications such as culture heritage. Although laser scanning provides all the details, it's usually not practical as the only technique for every object and structure. Large buildings, for example,

1 Scanned section (shown enlarged) on the facade of the Abbey of Pomposa.

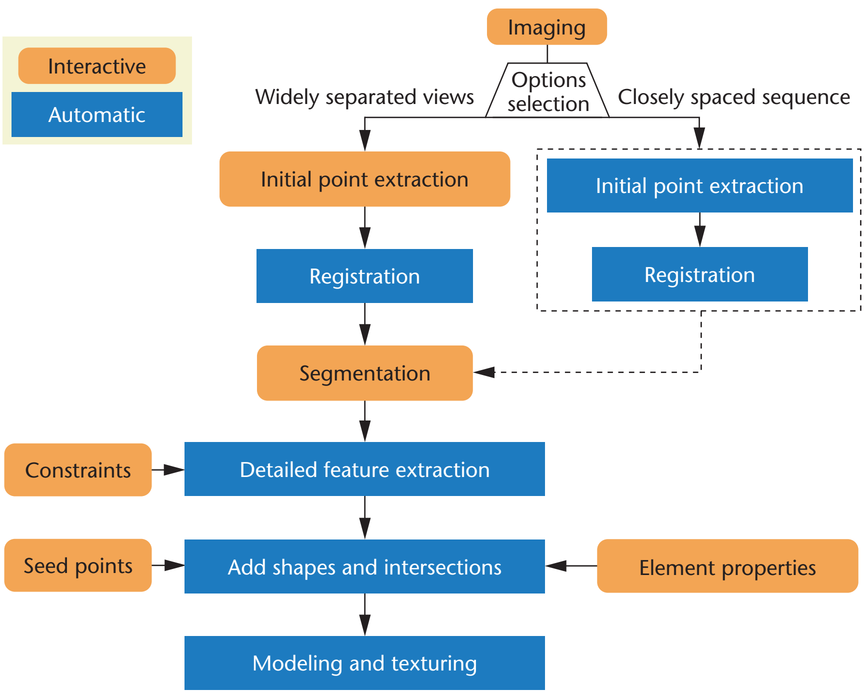

2 General procedure for image-based modeling.

require many scans and produce a huge number of points even for flat surfaces. Moreover, image-based modeling alone has difficulty with irregular and sculpted surfaces. It's also important to develop an approach that requires only a few widely separated views while simultaneously offering a high level of automation, and that can handle occluded and unmarked surfaces.

Therefore, the logical solution is to use image-based techniques to determine the basic shapes, and laser scanning to determine the fine details. In Figure 1, most of the structure is easy to model through images taken with a digital camera. However, parts of the surface (such as the enlarged section shown) contain fine geometric details that are difficult or impractical to model from images. It's best to acquire those parts using a laser scanner and then add them to the global model created from the images. This involves matching and integrating local 3D points obtained from the scanner with the global model. We measure several features, usually 8 to 10 points, using the images. Then we extract the 3D coordinates of the same features from the scanned data. We do this interactively using intensity images generated by the laser scanner. We then use the transformation parameters to register the two data sets in one coordinate system.

Details of our technique

Here we describe each of the approaches we developed to create models from digital images, range data, and the integration of the two.

Semiautomatic image-based modeling

We conceived this approach mainly for human-made objects such as classical architectures, whose designs are often constrained by proportion and configuration. Classical buildings are divided into architectural elements. These elements are logically organized hierarchically in space to produce the full structure. As long as we know some of a classical architecture's components, we can reconstruct it, even if images only partially reveal them. For example, a columnar element consists of the capital; a horizontal member on top; the column itself—a long, vertical tapered cylinder; and a pedestal or base on which the column rests. We can further divide each of these into smaller elements. In addition to columns, other elements include pillars, pilasters, banisters, windows, doors, arches, and niches. We can reconstruct each with a few seed points, from which we can build the rest of the element.

Our approach, which is photogrammetry based, aims neither to be fully automated nor to completely rely on a human operator. It provides enough automation to help the user without sacrificing accuracy or detail. Figure 2 summarizes the procedure and indicates which step is interactive and which is automatic (interactive operations are orange). The figure also shows an option for taking a closely spaced sequence of images, if conditions allow, and increasing the level of automation. Here, we discuss only the option of widely separated views, which is more practical for large-scale scenes. We focus on images with known camera setups. There should be a reasonable distance, or baseline, between the images to guarantee a strong geometric configuration. We interactively extract several features appearing in multiple images, usually 12 to 15 per image. We point to a corner and label it with a unique number, and the system can accurately extract the corner point. We use a Harris operator for its simplicity and efficiency. We base image registration and 3D coordinate computation on photogrammetric bundle adjustment because of its accuracy, flexibility, and effectiveness and because we wanted to provide the optimum solution compared to other structures from motion techniques.[7]

Advances in bundle adjustment eliminated the need for surveyed points or for manually entering initial approximate coordinates. (Photogrammetry has successfully tackled many other aspects required for high accuracy, such as camera calibration with full-distortion corrections, which we don't discuss here.) The bundle adjustment lets us register all the camera coordinates and orientations, as well as the 3D coordinates of the initial points, in one global coordinate system. The next interactive operation involves dividing the scene into connected segments to define the surface topology. This is followed by an automatic corner extractor and matching procedure across the images to add more points into each of the segmented regions. The epipolar condition and disparity range setup from the initial points' 3D coordinates constrains the matching within a segment. Bundle adjustment repeats with the newly added points to improve on previous results and recompute more accurate 3D coordinates of all points from all the images where they appear.

In addition to the multi-image approach, an approach to obtain 3D coordinates from a single image is essential because some parts of the scene appear only in one image. It's also necessary to cope with occlusions and a lack of features. Our approach uses several types of constraints for surface shapes such as planes and cylinders, and surface relations such as parallelism, perpendicularity, and symmetry. We determine the equations of some of the planes from seed points previously measured. We determine the equations of the remaining planes using the knowledge that they are either perpendicular or parallel to the planes already determined. With little effort, we can compute the equations of the structure's main planes and those to which other structural elements attach. From these equations and the known camera parameters for each image, we can determine 3D coordinates of any point or pixel from a single image, even without any marking on the surface. When some plane boundaries aren't visible, we can compute them using plane intersections. We can also apply this technique to surfaces such as quadrics or cylinders, whose equations we can compute from existing points. We can also use other constraints, such as symmetry and points with the same depth or same height. The general rule for adding points on structural elements and for generating points in occluded or symmetrical parts is to do the work in the 3D space to find new points to complete the shape, then project them on the images using the known camera parameters. Figure 3 shows the main steps.

As Figure 4a shows, the system constructs a cylinder after automatically determining its direction, radius, and position from four seed points. The user can set the ratio between the upper and the lower circle in advance. We set the default to about 0.85 to create a tapered column. From this information, the system can automatically generate in 3D the points on the column's top and bottom circles, resulting in a complete model, as Figure 4b shows.

For windows and doors, we need three (preferably four) corner points and one point on the main surface (see Figure 3). By fitting a plane to the corner points, and a plane parallel to it at the surface point, we can reconstruct the complete window or door. (An earlier article provides more details of this approach.[9]) Figure 5 (next

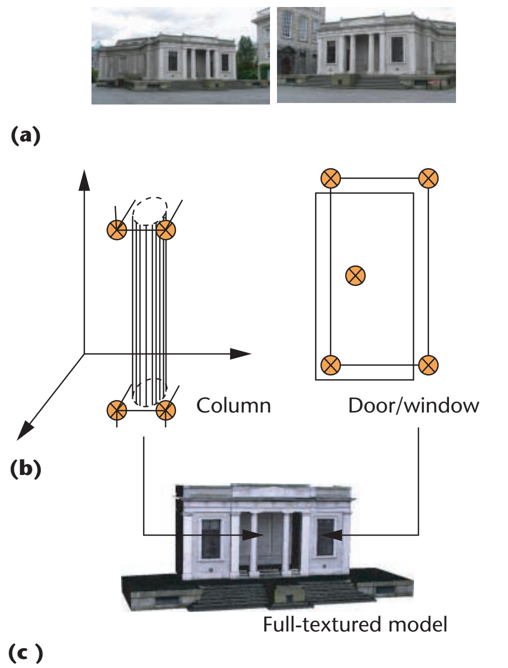

3 Main steps for constructing architectural elements semiautomatically (column and window examples): (a) extract in multiple image steps, match, and compute seed points' 3D coordinates; (b) in 3D space, reconstruct the object from the seed points; and (c) create a full 3D model and project the new 3D points onto the image for texture mapping.

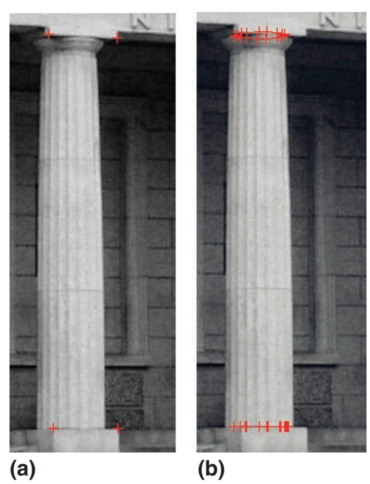

4 The system (a) extracts four seed points on the base and crown, then (b) automatically adds column points.

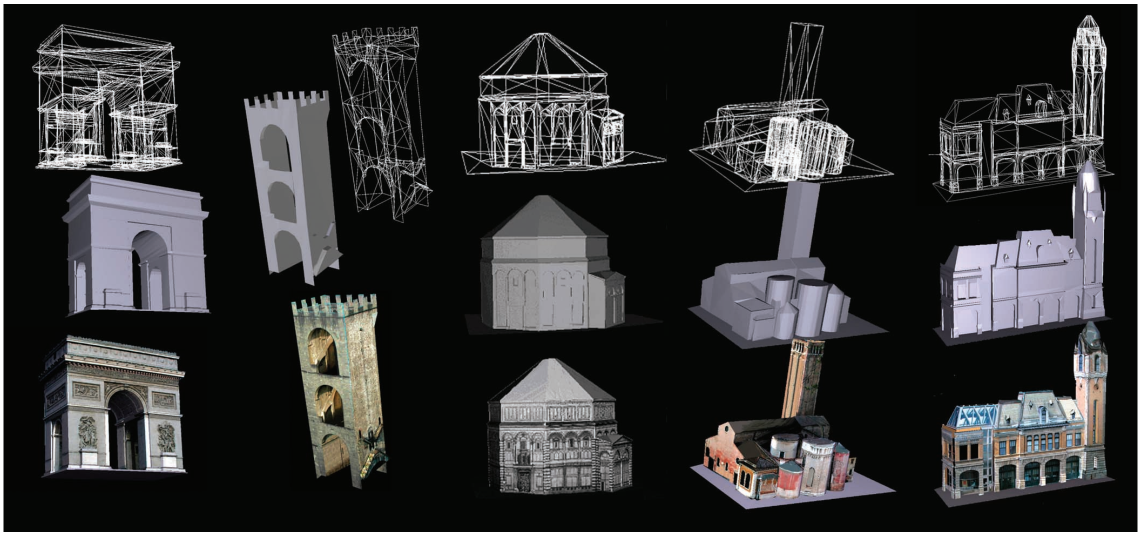

5 Models of structures from all over the world in wire frame, shaded solid, and textured solid.

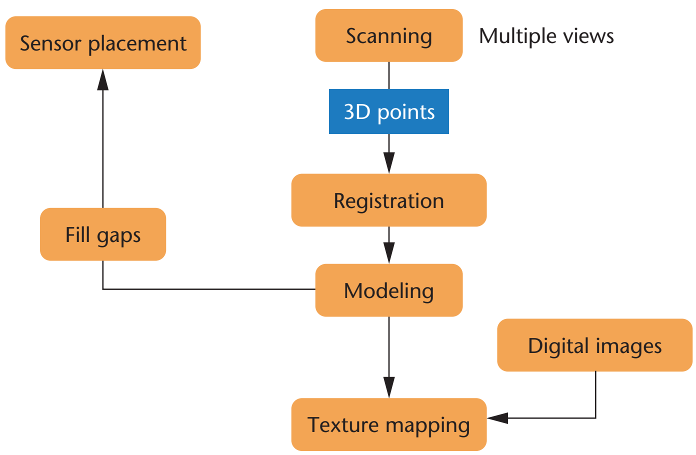

6 General procedure for range-based modeling.

page) shows examples of models created using this approach. We measured about 20 percent of these points interactively as seed points, and the system created the remaining 80 percent automatically.

Range-based modeling and texturing

Figure 6 outlines the steps for creating a triangular-mesh model from 3D images. If the 3D data is a set of registered images, it's easy to create a triangular mesh by simply triangulating each image. However, there is often sizeable overlap between the images from different views, so a mesh created this way has many redundant faces. Because we wanted to create a nonredundant mesh with no overlapping faces, we adopted a technique developed over the years at our laboratory and at InnovMetric Software.[10] InnovMetric Software implemented this technique in Polyworks commercial software.

Most laser scanners focus only on acquiring the geometry. They usually provide only a monochrome intensity value for each pixel as sensed by the laser. Acquiring a realistic look in the model requires texture maps obtained from a high-resolution, color digital camera. Some scanners have a color camera attached to them at a known configuration so that the acquired texture is always registered with the geometry. However, this approach might not provide the best results, because the ideal conditions for taking the images might not coincide with the best conditions for scanning. Therefore, our approach allows taking images at a different time than during scanning, and at whatever locations and lighting conditions (controlled illumination) are best for texture.[11]

Combining image- and range-based modeling

We use the semiautomatic image-based approach to model the entire structure without the fine details and sculpted surfaces. We can use the approach just described to separately model the sections requiring scanning. We use common points between the image- and range-based models to register them in the same coordinate system. We do this interactively using our own software, which can display and interact with images from various types of sensors and cameras. The next step is to automatically sample points from the range-based model along its perimeter and insert those into the image-based model. We adjust the image-based model's triangulated mesh on the basis of those new points to create a hole into which we add the range-based model so that there are no overlapping triangles.

Landscape visualization

When images of the entire scene taken at large distances, such as aerial images, are available, we can represent the landscape and integrate it with the model of the structures. This shows the structures in their natural setting and increases the level of realism. We determine the elevation of ground points between the main structures from aerial images. We then use cylindrical or spherical panoramas to represent the remainder of the landscapes and far objects, such as mountains. We use a few joint 3D points between the structures and the ground to register the ground elevation model and landscape panorama with the structures. The procedure is similar to the approach used by Sequeira et al.[3]

Modeling the Abbey of Pomposa

The Abbey of Pomposa, near Ferrara, Italy, is one of the most appealing Italian churches of the Romanesque period. Founded in the 7th century, it comprises several buildings that are part of one of the most important Benedictine monasteries. The abbey's core includes the refectory, the basilica, the capitulary hall, and the cloister. The bell tower was added in the 11th century. The abbey is architecturally simple with planar stone surfaces. Extensive artwork carved in marble decorates the facade. There are also three arches decorated with brick and stonework.

Data collection

Using the image-based technique, we completely modeled all of the abbey's structures, including the bell tower, except the carvings on the church's facade. We imaged the entire complex with an Olympus 4-megapixel digital camera. We acquired seven different sets of images, including one set from a low-flying airplane and one set from inside the church vestibule. Figure 7 shows the resulting seven models. We used the Biris 3D sensor (a short-range, submillimeter-accurate sensor developed in our laboratory) to scan details such as the left wheel *rosone* (rose window) and the peacock carvings on the left side, which are shown in Figure 8.

Results

We created seven individual models (see Figure 7) from the digital image, and several detailed ones from the scanned smaller regions. Figure 9 (next page) shows a close-up of the general model of the church building with the addition of eight new points from the wheel trim and the retriangulated mesh. The hole shown is where the model of the scanned wheel will fit. This is necessary because the wheel has many openings in its center section. The peacock model doesn't require cre-

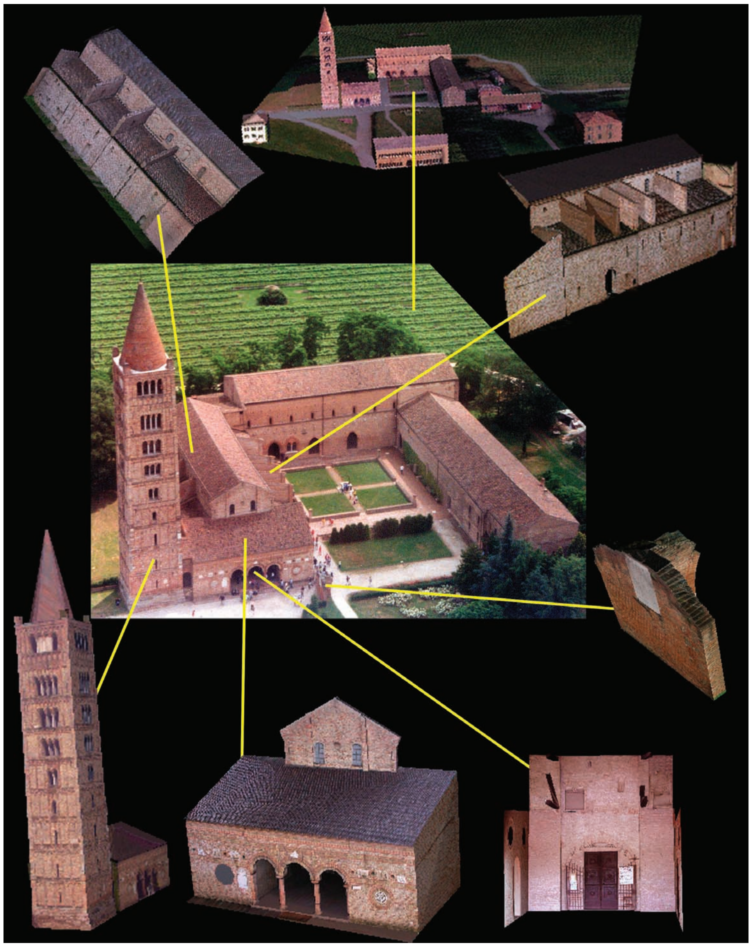

7 Image-based models that we created from the digital image.

8 Textured model of front building showing scanned regions.

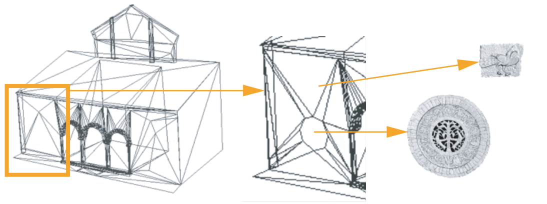

9 Integrating scanned and image-based models.

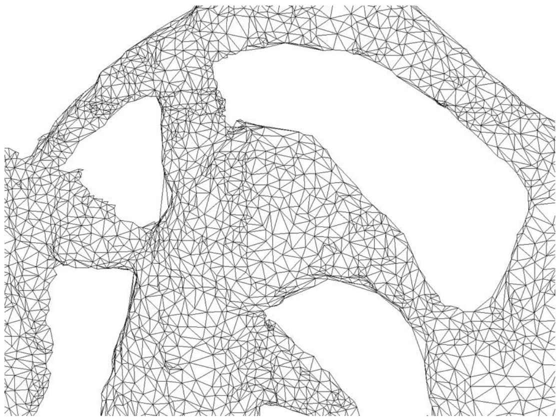

10 Close-up in the wire-frame model of part of the middle of a wheel (rose window).

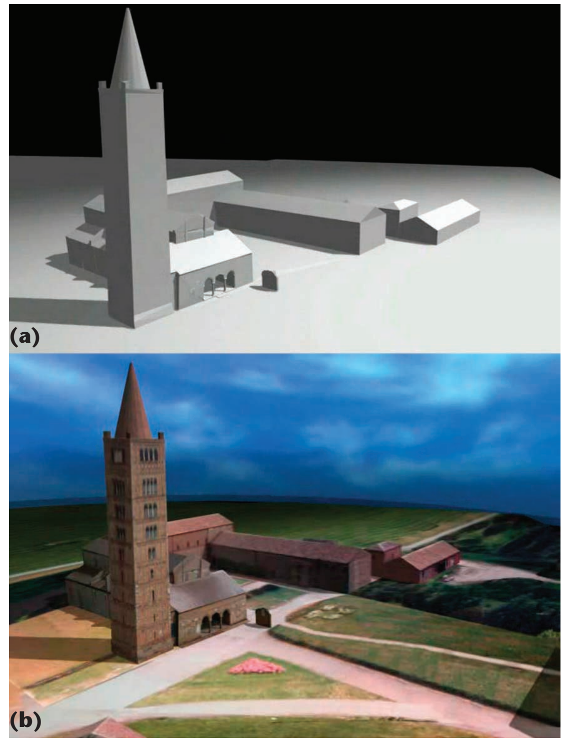

11 Final integrated model in (a) shaded and (b) textured rendering.

ating a hole in the main model, because it's completely solid and can simply attach to the plane of the wall model. Figure 10 shows a close-up of part of the middle of the wheel, displayed in a detailed wire frame. The scanned sections' level of detail, acquired at 0.5-mm resolution, is obviously far higher than that of the image-based regions, which lack the small geometric details. It's particularly more convincing when we view these sections up close while navigating through the model. Figure 11 shows snapshots from the walk-through movie: one without textures; and one fully textured, including all the landscapes.

Conclusion

We are now focusing on increasing the system's level of automation and ease of use. Another issue we are currently addressing is how to define the required accuracy from each component of the system to achieve acceptable geometric, visual fidelity and how to access the quality of the final model of large, complex sites. ■

Acknowledgments

We appreciate the help of Marco Gaiani (now at Politecnico di Milano), and architects Francesca Pozzi and the late Roberta Marchetti, who organized the survey campaign. We also thank architects Anna Maria Iannucci and Andrea Alberti for giving us the opportunity to demonstrate our technology in Pomposa.

References

1. H.-Y. Shum and S.B. Kang, "A Review of Image-Based Rendering Techniques," *Proc. Int'l Conf. Visual Comm. and Image Processing* (VCIP 00), SPIE, vol. 4067, 2000, pp. 2-13.
2. F. Bernardini et al., "Building a Digital Model of Michelangelo's *Florentine Pietà*," *IEEE Computer Graphics and Applications*, vol. 22, no. 1, Jan.-Feb. 2002, pp. 59-67.
3. V. Sequeira et al., "Hybrid 3D Reconstruction and Image-Based Rendering Techniques for Reality Modeling," *Proc. Int'l Conf. Videometrics and Optical Methods for 3D Shape Measurement*, SPIE, vol. 4309, 2001, pp. 126-136.
4. P. Debevec, "Image-Based Techniques for Digitizing Environments and Artifacts," *Proc. 4th Int'l Conf. 3-D Digital Imaging and Modeling* (3DIM 03), IEEE Press, 2003, pp. 234-242.
5. J. Wang and M.M. Oliveira, "Improved Scene Reconstruction from Range Images," *Proc. Ann. Conf. European Assoc. for Computer Graphics* (Eurographics 02), Blackwell Publishers, vol. 21, no. 3, 2002, pp. 521-530.
6. M. Pollefeys, R. Koch, and L. Van Gool, "Self-Calibration and Metric Reconstruction in Spite of Varying and Unknown Intrinsic Camera Parameters," *Int'l J. Computer Vision*, vol. 32, no. 1, Aug. 1999, pp. 7-25.
7. W. Triggs et al., "Bundle Adjustment for Structure from Motion," *Proc. Int'l Workshop Visual Algorithms*, *Vision Algorithms: Theory and Practice*, LNCS 1883, Springer-Verlag, 2000, pp. 298-372.
8. T. Werner and A. Zisserman, "New Technique for Automated Architectural Reconstruction from Photographs," *Proc. 7th European Conf. Computer Vision* (ECCV 02), LNCS 2351, Springer-Verlag, 2002, pp. 541-555.
9. S.F. El-Hakim, "Semi-automatic 3D Reconstruction of Occluded and Unmarked Surfaces from Widely Separated Views," *Proc. ISPRS Commission V Symp.*, Ziti Publishing, 2002, pp. 143-148.
10. M. Soucy et al., "Sensors and Algorithms for the Construction of Digital 3-D Colour Models of Real Objects," *Proc. Int'l*

Conf. Image Processing, IEEE Press, 1996, pp. 409-412.
11. J.-A. Beraldin et al., “Virtualizing a Byzantine Crypt by Combining High-Resolution Textures with Laser Scanner 3D Data,” *Proc. 8th Int’l Conf. Virtual Systems and Multimedia* (VSMM 02), Int’l Soc. Virtual Systems and Multimedia, 2002, pp. 3-14.

Sabry F. El-Hakim *is a senior research officer at the Visual Information Technology Group in the Institute of Information Technology at the National Research Council of Canada. His research interests include image-based modeling, multisensor data integration, camera calibration, and virtual reality. El-Hakim received an MSc and PhD in photogrammetry from the University of New Brunswick. He chairs the International Society of Photogrammetry and Remote Sensing (ISPRS) Working Group on Scene Modeling and Virtual Reality, and he chairs the SPIE series of conferences on videometrics.*

J.-Angelo Beraldin *is a senior research officer at the Visual Information Technology Group in the Institute of Information Technology at the National Research Council of Canada. His research interests include sensor systems engineering and signal processing (hardware and software) for 3D vision systems. Beraldin received a BEng from the University of Sherbrooke, Quebec, and an MS from the University of Ottawa, both in electrical engineering.*

Michel Picard *is a technical officer in the Institute for Information Technology at the National Research Council of Canada. His research interests include 3D imaging systems and 3D animations, especially virtualized environments. Picard received a diploma in architectural technology from La Cité Collégiale in Ottawa.*

Guy Godin *is a senior research officer in the Visual Information Technology Group in the Institute of Information Technology at the National Research Council of Canada. His research interests include 3D computer vision and image analysis, shape and appearance modeling, and interactive visualization. Godin received a BEng from Ecole Polytechnique de Montréal and an MEng from McGill University, both in electrical engineering.*

Contact Sabry F. El-Hakim at the Visual Technology Group, Inst. for Information Tech., Nat’l Research Council of Canada, 1200 Montreal Rd., M-50, Ottawa, Ontario K1A 0R6, Canada; sabry.el-hakim@nrc-cnrc.gc.ca.