

The Impact of Dense Range Data on Computer Graphics

Lars Nyland, David McAllister, Voicu Popescu, Chris McCue, Anselmo Lastra,
Paul Rademacher, Manuel Oliveira, Gary Bishop, Gopi Meenakshisundaram, Matt Cutts, and Henry Fuchs
University of North Carolina, Chapel Hill, NC 27599-3175

Abstract

The ability to quickly acquire dense range data of familiar environments has been met with enthusiastic response in our graphics laboratory. In this paper, we describe our prototype range collection system, based on a scanning laser rangefinder and a high-resolution digital color camera, that allows us to create panoramic color photographs where every pixel has an accurate range value. To accommodate occlusions, the data acquisition process is repeated from multiple locations and the results are registered with software. We discuss the acquisition system and review how the data from this prototype system have been used in existing graphics projects. We speculate about future improvements to range acquisition hardware and how those improvements impact graphics applications. The data acquired from our system is just a hint of what will be available in the future and we conclude with speculation about the impact of such data on graphics hardware and algorithms, since prevailing graphics hardware and software do not support this data well.

1. Introduction

Dense range data, where dense means at least one range sample per milliradian, has the potential to change our view of 3D computer graphics. Our particular area of interest has been in enhancing image-based rendering by providing color images of real scenes where each pixel has accurate range information. Thus, in addition to range image acquisition, we have also collected high-resolution color images from the same point-of-view, and built a prototype software system that builds high-resolution color photographs with depth suitable for image warping. An example of the range and color data is shown in figure 1.

The availability of range data collection has also been useful in advancing a variety of other graphics and geometry projects in our lab. For example, the collection of range

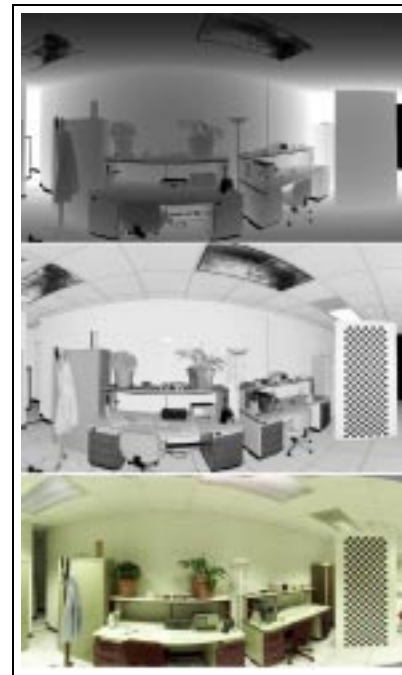


Figure 1. Samples of the range and color data collected from our hardware teaching lab. At the top is a gray-scale range image, below that is the reflected laser intensity image, and at the bottom is a color panorama of the scene. Images such as these are used to create color images with depth.

data in the “Office of the Future” gives static information about the geometry of projection surfaces for immersive telecollaboration. As range image acquisition systems improve, dynamic range data will allow the collaborators to walk about, viewing the collaboration scene in true 3D.

A well-established research area is building geometric models from range data (point clouds and range images), and we have some new algorithms that quickly consume the range data, simplifying it to only a few percent of its original complexity.

Two additional image-based rendering applications use the dense range data. The first is our *multiple-center-of-*

Please visit <http://www.cs.unc.edu/~ibr/pubs/nyland-mvview99/> for a full-color version of this paper.

projection images where, potentially, each range sample is taken from a different, but known, location. This allows us to acquire data throughout an environment, substantially reducing occluded areas in a single data set, producing distorted images that are correctly reprojected.

The second is our work on *image-based objects*. Range data of an object is acquired from several locations, the images are then registered and warped to a single point-of-view. This representation has the advantage that occluded areas are drastically reduced while preserving properties of single-image warping, such as occlusion compatible ordering of rendering to preserve proper occlusion.

The final area discussed is the registration (both user-assisted and automatic) of range images taken from different locations in the same environment. One automatic technique improves upon the *iterated-closest-point method* [3] by taking into account the presence of shadows. A separate method for exploring automatic registration looks at 3D Hough transformations of separate range images.

The hardware system we have assembled is a proof-of-concept, and as such, it allows us to speculate about the future impact of high-density range images. Current hardware and algorithms in computer graphics have not considered the existence of high-density range images, and with the growing availability of similar commercial devices, new methods of handling such data need to be considered. For instance, what can be done with 100 million color samples that have positions in three dimensions? Currently, not much more than simplification, but perhaps it is time to think about new trends in rendering hardware, rendering algorithms, model-building and illumination based on high-density color and range data.

This paper describes our prototype range acquisition system, the calibration procedures, the specifics of the data collected, registration of multiple range images, the process of matching range data with color images, and the impact that it has had on our research projects. We conclude with goals for future acquisition systems, citing potential impact on graphics research.

2. Hardware System

Our prototype data collection system consists of a commercial laser rangefinder, a high-resolution digital camera, a panning unit, with a PC to control the devices and collect the color and range data. The components are all mounted on a cart to provide mobility.

Scanning, Laser Rangefinder. Our rangefinder is a commercially available scanning laser rangefinder from Acuity Research [21]. It is a time-of-flight rangefinder that can acquire range data at up to 50K samples/sec. It can accurately

determine range from 0 to 15 meters, with better accuracy at closer ranges.

The scanning mirror sweeps the laser beam in a vertical plane with a 60° occlusion (aimed downward). The scanning mirror has a 4096-position shaft encoder to aid in determining where the laser is aimed.

Pan-Tilt Unit. A pan-tilt unit from Directed Perception [17] controls the orientation of the scanning plane. The rangefinder is placed on the pan-tilt unit such that the position where the laser reflects off the scanning mirror is on the panning axis. The panning device has 14,000 steps in 180°; thus, to acquire samples every milliradian, we typically move 3 or 4 steps between each successive scan.

The Color Camera. To collect color data, we use a Canon/Kodak EOS D2000 digital camera. It has high resolution (1728x1152 pixels), accessibility to the raw CCD data, and FireWire communication. The lens in use is the Canon 14mm flat-field lens, chosen to acquire a wide field-of-view, considering that the camera's digital CCD is smaller than 35mm film.

During acquisition, we disable automatic exposure and focus, set the aperture at $f/11$ to get a depth-of-field from 0.5m to ∞ , and set the exposure time as necessary to accommodate the aperture for the entire scene. The wide field of view of any panorama typically leads to very large differences in illumination, making a single exposure setting problematic. Currently, we use additional lighting to more evenly illuminate the panorama, but are working on acquiring high dynamic range images, as in [4].

3. Calibration

Since the system is a prototype cobbled together from many parts, calibration procedures were necessary to determine where the laser beam is pointing. These include calibration of the mirror's angle, the scanning motor's position, and the panning device's positioning. The camera also required calibration to determine lens distortion.

Calibrating the Range. The rangefinder is well calibrated by the manufacturer to return accurate measurements over a wide range of values (0 to 15m). Still, there are a few steps a user can take to improve the range values read. First, the rangefinder can return a more confident value if it has a longer period to make a measurement. Knowing this, we typically set the collection rate to 1/2 to 1/3 of the peak rate, and have seen dramatic improvements in the data.

The second is limiting the maximum range value that the rangefinder will return. The rangefinder modulates the laser

light to create an interference pattern, so to avoid harmonics, it must look at the longest possible distances first. If this distance is set to be shorter than the maximum, the device can settle on a modulation frequency more quickly, returning a better indication of the range.

Calibrating the Latitudinal Angle. The scanning mirror controls the latitudinal angle (ϕ), and we read its position from the attached 4096-position shaft encoder. Since we are typically taking 5,000–10,000 samples per revolution, it is clear that we must interpolate the shaft position, since subsequent readings may report identical shaft positions.

To do this, we assume constant motor speed over the time required to collect a 1K buffer of samples. It is not possible to know exactly when the encoder moves from one position to the next, but since we have a large number of samples, we can estimate where the transitions occur by performing a least-squares line fit using the points around the transitions.

Calibrating the Longitudinal Angle. Determining the actual angle around the polar axis relies not only upon the panning motor position, but on the angle of the 45° mirror as well. We devised a simple experiment to determine both at the same time.

In a room, we aim the laser horizontally at a wall so it is roughly perpendicular, marking its position. We then move the scanning mirror over the pole 180° , aiming at the opposite wall, and mark that position. We then pan the device 180° , where it should coincide with the first mark. We mark that position, and again move the scanning mirror over the pole to point at the second mark. We mark its position. If all the hardware were perfect, the two marks on each wall would be coincident, but due to errors, they are not. From the separation of the points and knowing the distance from the rangefinder to the points, the panning error and mirror error are both determined. The values we found are 14,039 steps in 180° , and 44.89° (the mirror error also affects the latitude). At a distance of 4m, this moves the data sample more than 3cm.

4. Data Collected

The range data collected consist of quad-tuples of range, longitude and latitude angles (θ and ϕ), and the intensity of the reflected laser light. Two visualizations of the data are shown in figure 1, which are spherical panoramic images (not fisheye) showing the range and reflected laser intensity values at regular latitude-longitude positions. After the data is collected, it can be processed to correct for the calibration values found.

The color data simply consists of a panoramic set of images taken with the camera's nodal point coincident with

that of the rangefinder's (a custom mounting bracket ensures this). The camera is rotated (using the pan-tilt unit) stopping every 24° to acquire a $55^\circ \times 77^\circ$ image.

5. Combining Range and Color Data

After the data is collected, substantial processing is required to combine range data with color information. The result is a color image with accurate range information for every pixel in an image.

Undistorting Color Images. All camera lenses have distortions from the pinhole model they are designed to emulate. Fortunately, lens distortion is well studied, and free software exists to determine parameters that aid in undistorting an image [23]. We acquired several dozen photographs and performed the analysis, using the determined parameters to resample the images into an undistorted form. This is necessary, as the distortion placed some pixels 30 pixels away from their ideal pinhole position.

Range Data Resampling and Cleanup. The data from the rangefinder is not on a regular grid, as there is no control between the scanning motor and the sampling hardware. We project all of the range samples onto a spherical grid, apply some error removal and hole filling heuristics, and then produce a spherical image of the range and intensity values.

If the laser beam spans two disparate surfaces during a sampling period, the resulting range is usually between the two surfaces (though not always). We use a voting scheme on our projection grid that looks at the range of the 8 nearest neighbors. If at least 4 are within some tolerance, the value is deemed to be valid, otherwise it is removed. This has the effect of removing all floating samples.

Since the rangefinder's ability to determine distance depends on the amount of light reflected, we cannot acquire range information for very dark or specular objects. Objects such as glossy (or even semi-glossy) furniture, dark metal, rubber or plastic objects (wall trim, electronic equipment, plastic trim on furniture), or metallic frames and light fixtures all cause problems.

We use a variation of the Splat-Pull-Push algorithm [6] to place the range data on a regular grid and fill in the holes. The algorithm was designed to perform well on sparse data, but also works very well on dense data like that from the laser rangefinder. The splat portion of the algorithm performs most of the work since the samples are about as dense as the image pixels. The pull and push phases interpolate the samples to fill in places that were not scanned well by the laser. We output two images from this process—a range image and an infrared laser intensity image. These images are used to align the color camera images with the range data.

Alignment. The goal of the alignment procedure is to find the orientation of the camera image that best correlates with the rangefinder image. The laser intensity image cannot be directly correlated to the color image (or even the red channel of the color image) because the illumination of the two images is so different that straightforward image correlation gives poor results. The laser image is illuminated directly by the laser and nothing else (ambient light is removed). The laser image has no shadows, the entire scene is equally illuminated, and specularities occur in different places than in the normally lit scene.

Instead, we perform the alignment on the edges in the images. Edge-detection algorithms work well on the data from the laser rangefinder, but tend to show the high frequency noise in the color images. To solve this problem, we apply a variable conductance diffusion operation (VCD) [24] to the color images. Edge detection on the blurred image then finds only the salient edges. The edge pixels are undistorted according to the distortion parameters found in the camera calibration.

Edge detection is performed on both the range and intensity images from the rangefinder. The edges in these images are then blurred by convolving them with a kernel that has wide support, but whose gradient increases near the center of the kernel. This enhances the search by giving nearby solutions a small error value, but not nearly as small as an exact solution.

We only search over the three angles of registration between the spherical range image and planar color image since we know that the two images share the same center-of-projection. The error value is computed as the degree of edgeness from the rangefinder image that corresponds with the edges in the color image. We use a simulated annealing search strategy, which works well when presented with a reasonable starting point.

Having found values for the 3 angles, we return to the original color images, correct for distortion, and determine the proper distance information. To do so, we project the range information onto the planar grid, making a list of range values for each pixel. Resampling range data is problematic, so we perform a clustering algorithm on the range values for each pixel, setting the range as the average of the largest cluster. This method avoids the error of blending samples across multiple surfaces.

Result. The output of this process is a variation of TIFF with one extra layer (disparity, related to inverse depth) and the camera parameters. This is our standard file format for image-based warping reference data [15].

6. Use in Graphics Projects

This section describes how the color and range data is used in current graphics projects in our lab.

Image-Based Rendering by 3D Warping. The primary motivation in producing the data described here is to support our image-based rendering project [1]. Many different aspects of IBR are being studied, such as representation, visibility, reconstruction, multiple views, hardware acceleration, and hybrid CG systems, and all require source images to render.

The registered color and range images lead naturally to an image-warping walk-through application that renders them with as few artifacts as possible. If the images were rendered as triangle meshes, errors would occur at silhouettes such as table edges, doorways, and other spatial discontinuities, where the mesh would be stretched across the spatial discontinuities.

One of the first steps to perform is silhouette edge detection. While many sophisticated methods for performing this exist, it turns out that simple heuristics perform nearly as well and are extremely easy to compute. One method computes the dot product of the viewing ray with the normal vector of the triangles in the mesh. Silhouettes (and badly sampled surfaces) will have values close to 0, and thus the mesh can be broken at these locations.

We have developed a simple application [9] that uses OpenGL and runs on our Onyx2 hardware as well as the custom PixelFlow hardware [5]. The user interface allows the user to move around the environment arbitrarily, using multiple panoramic source inputs. The effect is very real—during demonstrations, many people believe that we have either taken photographs from a significant number of positions or are somehow showing a video. Images from a walk-through sequence with two panoramas are shown in figure 2. The performance is near real-time.

Optimizations have been made to improve rendering performance by using the PixelFlow graphics hardware. For instance, it is possible to perform incremental warping calculations where the warping arithmetic that applies to groups of pixels is performed once and only the pixel-specific arithmetic is performed for each pixel. We have also developed a custom rendering primitive called an image tile. We can cull at the image tile level, providing a dramatic improvement in rendering. The rendering of an image tile is also where the incremental arithmetic is performed.

As a further extension, we have also developed a new point primitive for rendering, which we call the Voronoi-region primitive. It is basically a cone-shaped point rather than a flat disc, aimed at the viewer and falling off in z as its radius increases. When several of these primitives are



Figure 2. Some sample images of a walk-through of our reading room. The input data is composed of panoramas taken from 2 locations and consists of 10 million samples.

displayed from a planar surface, they implicitly compute the Voronoi regions of the samples. See [9] for full details.

Immersive, 3D Telecollaboration: the Office of the Future. An immersive, telecollaboration project [20] is underway that seeks to share workspaces by projecting images of remote locations onto nearly every surface possible, including irregular and dynamic surfaces. The ceiling of the office is populated with projectors, and cameras are strategically placed in unobtrusive locations for two-way interaction.

Currently, range information provided by the system described here could be used to provide the range of the static structures in the shared environment. This information is used to not only locate all the projection surfaces but to locate the projectors as well, so that the position of each pixel from each projector in the room can be computed.

Multiple-Center-of-Projection Images. MCOP images correctly reconstruct 3D scenes where every pixel in the source image was potentially acquired from a different, but known, position [19]. MCOP images have the advantage that objects can essentially be scanned, but the data is still a single image. If a strip camera is the acquisition model, then pose information is only required for each column of data.




Figure 3. For multiple-center-of-projection images, we show the path taken to acquire the data at the top, the acquired data (distorted middle view), and the appropriate reconstruction of the data.

The MCOP project was the first client of the ranging system described here. We attached the UNC Hi-Ball tracking system to the rangefinder, made the acquisition software network-aware, and slowly rolled the range acquisition cart in the tracked environment.

In figure 3, we show a photo of our lab environment with the rangefinder's path superimposed. We then show the resulting image that was collected, colored by hand to distinguish the important parts of the scene. Finally, we show the reconstructed image.

Image-Based Objects. Image-based objects [16] are composed of six layered-depth images [22] of an object properly registered. An IBO can be displayed from an arbitrary viewpoint where it will be rendered correctly. Applications of image-based objects include virtual museums, web-based catalogs and computer games. Multiple IBO's can be composed together and still be properly rendered, since McMillan's results about occlusion compatible ordering [12, 11] for rendering apply to one or more IBO's.




Figure 4. On the left are the 4 input images of a coin-bank. On the right are views of the reconstructed IBO from a variety of viewpoints.

One benefit of image-based objects is that of bandwidth. Sending an IBO to a web-browser is not much more costly than sending an animated GIF, yet with an IBO, the user can move the viewpoint arbitrarily. Additionally, an IBO can be rendered interactively, giving the user instant feedback.

An example of 4 of the images used to build an IBO and their rendering are shown in figure 4. The images only use the laser light, but future objects will also have color.

Reconstruction of Geometry. A fast, memory efficient, linear time algorithm that generates a manifold triangular mesh has been developed [13]. It has been used on standard point cloud data sets as well as the data sets acquired by our prototype range system. The speed of the algorithm is derived from a projection-based approach used to determine the incident faces on a point.

An example of the output from this algorithm is shown in figure 5. The original input has 6.5 million samples, while the result has only $\sim 140k$ points. The processing time taken to build the mesh was 88 seconds on an Oynx2. The rendering of the reconstructed geometric model is made by applying the intensity image as a texture map onto the created triangle mesh. Despite the drastic reduction in complexity, the model retains the important features of the collected data set.

Registration of Multiple Range Images. We have three methods for registering range image data. One is user-assisted, and was used for all the warped data sets described in this paper. The other two are automatic.

In the user-assisted process, the user selects points from 3 corresponding planes in each data set [10]. The data is shown in a 3D reprojection that can be translated and rotated for easy selection of the points. Once the 3 planes are selected, the fundamental matrix can be found. Error metrics are given, and the process can be repeated as needed.



Figure 5. On the left is a texture-mapped rendition of simplified geometry from our reading room data. On the right, we have zoomed in and revealed the mesh structure of the simplified data.

Other selection techniques have been suggested for plane selection using the 2D display of the range images, including using ordinary box, lasso, and spray paint models. These may be quicker, as it is simple to interact with 2D data displayed on a monitor and manipulated with a mouse.

One automatic method is called the Empty Space Registration Method [18], which is a variant of the Iterated Closest Point Algorithm [3]. In ICP, the data being registered is assumed to be “full,” that is, there are no shadows or occlusions and the data is sampled similarly in each data set. In data sets of real environments, there are bound to be occlusions. The empty space registration method considers empty space to the first visible surface and shadow volumes explicitly, knowing that nothing can occupy the empty space and anything can be in the shadows. Results of the search are shown in figure 6, where 2 source views of a computer on a table are shown. The merging shows their almost-correct registration, despite the fact that very little of the scene is shared between the views.

Our second automatic registration method is under development, and is based on a 3D Hough transform of the range data. The edge-detection step is not necessary with the rangefinder data, as the data collected represents the first “edge” in 3D (surface). This method takes each sample and performs the standard Hough transform operation by incrementing all possible buckets of r , θ , and ϕ for each sample. A sample of plane detection is shown in figure 7, where the pixels showing the ceiling of the room have been recovered from the Hough transform of the range data.

7. Additional Applications of Range Data Systems

As-Built Models. Commercial scanners are used to create *as-built models* from range data. One example is an Atlas launch tower was scanned from several locations, all images were inter-registered, and a model was produced for




Figure 6. The top two images show range scans of a workstation. The bottom two images show the automatic alignment.

renovation [8]. The areas scanned were difficult to reach, apparently, as the scanning hardware was held hundreds of feet off the ground by a crane. Another example of acquiring as-built models is that of sets for Hollywood movies [7]. The Cyrax scanner was used to scan a cave set from the movie “Starship Troopers.” 3D CAD models were produced, which allowed computer graphics artists to superimpose animations in the proper locations.

Forensic Recording. During a crime scene investigation, many photos are taken, and in some cases, the scene is frozen or even disassembled and reassembled elsewhere for the duration of the legal process. All of this is done to allow the participants to re-examine the site to determine what happened.

This is a case where a static imagery of range and color is exactly what is needed. There is no need for dynamic data collection. It is also a case where the measurements need to be very accurate.

Remanufacturing. A range scanner from the Canadian National Research Council has a high accuracy and real-time acquisition rates device that captures range over a limited volume [2]. The error in range is 10–100 microns, the maximum scanning rate is 10 million samples per second, and the volume is a cube several centimeters on a side at a distance of some tens of centimeters. The high-accuracy of such a system makes it ideal for scanning an existing part




Figure 7. Automatic detection of planes is shown in this image of our hardware lab. The range data from the ceiling created the highest peak in the Hough transform, and by performing the inverse operation, we have highlighted the pixels that contributed to this plane with a cross-hatch pattern.

with the goal of reproducing that part, especially when used with a CCM milling machine to accurately place the device.

Remote Walk-Through. The K²T company has demonstrated, in conjunction with the Robotics Institute at CMU, the ability to build remote walk-throughs [8]. They scanned an abandoned research facility, collected color images, and also videotaped a helicopter fly-over of the site to reconstruct the surrounding topology. The result is a model that can be viewed from arbitrary locations.

3D Movies and Television. A obvious use of dense range data is the creation of true 3D movies where a viewer could move to any position desired to view the movie. This is in contrast to the so-called 3D movies that are stereo projections that give the user a 3D view from the camera’s location. This application requires real-time range acquisition as well as hardware acceleration to display the images in real-time. Research in this area is underway at CMU’s Robotics Institute [14].

In its fullest form, a viewer could walk anywhere and look in any direction, viewing a properly reconstructed scene. For sports broadcasts, the freedom of motion holds immense appeal, in that a viewer could be on the field or court with the players, moving with the action.

Of course, movie and television directors have made their careers choosing the best presentation for the viewers, and are bestowed with awards when they do this well. To preserve this, the viewer might be restrained to a small volume, but still be allowed to move, gaining a true sense of the 3D nature of the scene.

8. Conclusions

We are excited about building a successful range acquisition system, and the success of our system has been shown by the desire of others to use the data in their projects. As a custom system, we are able to quickly adapt our data collection methods to that required by different projects, an advantage not found in commercial systems. We feel in our own work and in the work of others that there is a large future for dense range data.

The ability to match color data with range data has also been successful, although fraught with significant processing requirements. The combined data enables realistic walk-throughs of real environments that would be far too complex (or impossible) to model.

9. Future Work

Our immediate goals for the future have to do with reliable automatic registration methods, and better calibration of the rangefinder system. We are also exploring the development of hardware devices that will allow the simultaneous collection of range and color data.

Our future work is narrow compared to all that could be done to provide better support for dense range images with color. Obviously, real-time acquisition hardware from multiple locations would increase the demand for real-time rendering, so hardware acceleration both in acquisition and rendering are important areas to explore, despite the difficulty involved. The large number of points is, in itself, a demand for faster rendering hardware. In addition, dealing with view-dependent artifacts such as specular highlights could be explored (and has been [25]), where recovery of the surface properties could be performed to more realistically yield the artifacts.

References

- [1] Image-based rendering at unc. <http://www.cs.unc.edu/~ibr/>, 1999.
- [2] J.-A. Beraldin, M. Rioux, F. Blais, J. Domey, and L. Cournoyer. Registered range and intensity imaging at 10-mega samples per second. *Opt. Eng.*, 31(1):88–94, 1992.
- [3] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE PAMI*, 14(2):239–256, 1992.
- [4] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH '97*, pages 369–178, Los Angeles, 1997.
- [5] J. Eyles, S. Molnar, J. Poulton, T. Greer, A. Lastra, N. England, and L. Westover. Pixelflow: The realization. In *SIGGRAPH/Eurographics Workshop on Graphics Hardware*, pages 57–68, Los Angeles, CA, 1997.
- [6] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The lumigraph. In *SIGGRAPH 96*, 1996.
- [7] C. Inc. The cyrax 3d laser imaging system. <http://www.cyra.com/>, 1999.
- [8] K^2T . K^2T home page. <http://www.k2t.com/>, 1999.
- [9] D. McAllister, L. Nyland, V. Popescu, A. Lastra, and C. McCue. Real-time rendering of real world environments. Technical TR99-019, UNC Computer Science, 1999.
- [10] C. McCue. Multi-function point viewer. <http://www.cs.unc.edu/~ibr/projects/multiple>, 1999.
- [11] L. McMillan. *An Image-based Approach to Three-Dimensional Computer Graphics*. Ph.d. dissertation, University of North Carolina at Chapel Hill, 1997. also available as UNC Technical Report TR97-013.
- [12] L. McMillan and G. Bishop. Plenoptic modeling: An image-based rendering system. In *Proceedings of SIGGRAPH 95*, pages 39–46, Los Angeles, CA, 1995.
- [13] G. Meenakshisundaram and S. Krishnan. A fast and efficient projection based approach for surface reconstruction. Technical Report TR99-023, UNC Computer Science, 1999.
- [14] P. J. Narayanan, P. W. Rander, and T. Kanade. Constructing virtual worlds using dense stereo. In *Proceedings of Sixth IEEE International Conference on Computer Vision (ICCV'98)*, pages 3–10, Bombay, India, 1998.
- [15] M. Oliveira. Ibr tiff file format. http://www.cs.unc.edu/~ibr/src/ibr_tiffib/, 1997.
- [16] M. M. Oliveira and G. Bishop. Image-based objects. In *Interactive 3D*, Atlanta, GA, 1999.
- [17] D. Perception. Manufacturers of innovative devices and software for the intelligent control of sensors and sensor processing. <http://www.dperception.com/>, 1998.
- [18] P. Rademacher. Range image registration via consistency of empty space. <http://www.cs.unc.edu/~ibr/projects/emptyspace/>, 1999.
- [19] P. Rademacher and G. Bishop. Multiple-center-of-projection images. In *SIGGRAPH98*, Orlando, FL, 1998.
- [20] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image-based modeling and spatially immersive displays. In *SIGGRAPH98*, pages 179–188, Orlando, FL, 1998.
- [21] A. Research. Acuity research scanners and sensors. <http://www.acuityresearch.com/>, 1998.
- [22] J. Shade, S. Gortler, L.-w. Hey, and R. Szeliski. Layered depth images. In *SIGGRAPH 98*, Orlando, FL, 1998.
- [23] R. Y. Tsai. An efficient and accurate camera calibration technique for 3d machine vision. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 364–374, Miami Beach, FL, 1986.
- [24] T. S. Yoo and J. M. Coggins. Using statistical pattern recognition techniques to control variable conductance diffusion. In H. H. Barrett and A. F. Gmitro, editors, *Information Processing in Medical Imaging (Lecture Notes in Computer Science 687)*, pages 495–471, Berlin, 1993. Springer-Verlag. Also available as UNC Technical Report TR94-047.
- [25] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *SIGGRAPH 99*, August 1999.