

Virtualizing Real-World Objects

Hendrik P.A. Lensch Jan Kautz Michael Goesele Jochen Lang Hans-Peter Seidel
Max-Planck-Institut für Informatik
Saarbrücken, Germany
{hpseidel,lensch,jnkautz,goesele,lang}@mpi-sb.mpg.de

Abstract

High quality, virtual 3D models are quickly emerging as a new multimedia data type with applications in such diverse areas as e-commerce, online encyclopaedias, or virtual museums, to name just a few.

This paper presents new algorithms and techniques for the acquisition and real-time interaction with complex textured 3D objects and shows how these results can be seamlessly integrated with previous work into a single framework for the acquisition, processing, and interactive display of high quality 3D models.

In addition to pure geometry, such algorithms also have to take into account the texture of an object (which is crucial for a realistic appearance) and its reflectance behavior. The measurement of accurate material properties is an important step towards photorealistic rendering, where both the general surface properties as well as the spatially varying effects of the object are needed. Recent work on the image-based reconstruction of spatially varying BRDFs enables the generation of high quality models of real objects from a sparse set of input data.

Efficient use of the capabilities of advanced PC graphics hardware allows for interactive rendering under arbitrary viewing and lighting conditions and realistically reproduces the appearance of the original object.

1 Introduction

The rapid advances of consumer level graphics hardware make it possible to render increasingly complex and accurate models in real time. Computer-generated movies are getting more and more realistic and users will soon demand a similar level of realism in a wide range of everyday applications such as computer games, digital libraries and museums, or encyclopedias. Being able to efficiently generate, process and display the necessary models will become a more and more important part of computer vision and computer graphics.

A high quality representation must capture two different things: the shape of the object represented as a geometric description of its surface and the appearance of the materials it is made of, e.g. the object's color, texture, or reflection properties. This data must be integrated into a single digital model for processing and display trying to meet conflicting requirements (such as realism versus interactive speed). As more and more visual complexity is demanded, it is often infeasible to generate these models manually. Automatic and semi-automatic acquisition methods are becoming increasingly important.

The virtualization pipeline relies on computer vision and computer graphics techniques: To obtain the geometry of an object, a 3D scanner and various mesh processing tools are used. The surface appearance properties are acquired by taking a number of images with constrained lighting. After registration to the 3D geometry the object's texture, spatially varying reflection properties, and microstructure (normal maps) can be extracted. A compact representation of the object can be obtained that allows for accurately shaded, photorealistic rendering from new viewpoints under arbitrary lighting conditions. In addition, this model may be used for object recognition and material investigation.

This article highlights some recent results on the acquisition and interactive display of high quality 3D models. It shows how computer vision and computer graphics techniques can be seamlessly integrated into a pipeline for capture, processing, and interactive display of real-world objects.

2 The Virtualization Pipeline

This paper focuses on the generation of high quality 3D models containing the object's geometry and surface appearance. This data is needed for many computer graphics or computer vision applications.

Fig. 1 shows an overview of the steps necessary to acquire such a model. First, the geometry and the texture of the object are acquired – typically using different techniques and acquisition devices which makes a separate reg-

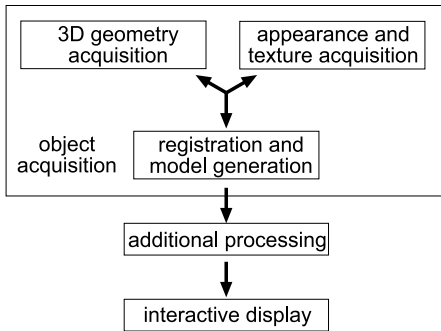


Figure 1. The virtualization pipeline. Depending on the applied techniques geometry acquisition, texture and appearance acquisition, and registration depend on each other in different configurations.

istration step necessary in order to align both data sets. In subsequent processing steps the required information such as reflection properties or normal maps is extracted from the input data. Once a complete model is created, it can be resampled, converted to a different data representation, or compressed to make it suitable for a particular application scenario. Finally, the target application should be able to display the model interactively without omitting important information.

In the following sections we give a detailed description of all the steps of the 3D object pipeline. We start with an overview over the acquisition of 3D geometry in Section 3. In Section 4 we present image-based acquisition techniques followed by acquisition techniques for appearance properties in Section 5. A technique to register texture and image data is described in Section 6. Section 7 introduces several methods to display the acquired models interactively. We present some examples of acquired models in Section 8 before we conclude with Section 9.

3 3D Geometry Acquisition

In most cases no high quality 3D geometry models of real world objects are available. It is therefore normally necessary to acquire the geometry of objects using a 3D scanner. Several research groups including [20, 1] have built their own 3D scanner – some of them tailored to specific requirements. Furthermore, there is a broad range of commercial products made by companies like Cyberware, Minolta, or Steinbichler.

There are several different approaches to acquire the 3D geometry of an object (for an overview see [5]) but most of the systems for small or medium sized objects are based

on an active stereo structured light approach. One or several patterns are projected onto the object with a computer controlled projection system (e.g., a video projector, a color coded flash stripe projector, or a laser beam). The projected light patterns on the object are observed by a digital camera which is rigidly connected to the projection system. The 3D location of a point on the surface of an object is then defined by the intersection of a ray from the projected pattern with the viewing ray that corresponds to the pixel in the digital image that observed this ray.

The position of these rays in space is determined in a separate calibration step: The patterns are projected onto a calibration target – typically a flat board or a three-dimensional structure with a regular pattern whose geometric properties are exactly known. The acquired images are analyzed to recover the parameters of the projection system and the camera using standard camera calibration techniques.

Using the active stereo approach most objects cannot be acquired with a single scan – either because front and back part of the object cannot be scanned with a single scan or because for a given configuration not all parts of the object are visible from both the position of the projection system and the digital camera. Therefore several scans have to be registered against each other in order to combine them into a single set of surface points. This is commonly done using a variant of the iterative closest point method (ICP) [2, 26]. The resulting point cloud is triangulated leading to a single triangular mesh using one of a large variety of methods (for an overview see [5]). Further processing steps include smoothing to reduce noise (e.g., using [32, 16]) and editing of the resulting mesh for which a huge selection of tools is available.

4 Image-Based Acquisition Techniques

Image-based techniques become increasingly popular to acquire models of complex objects. A digital camera captures images from which various properties of the object can be derived. The large number of measurements that can be made in parallel (i.e., every pixel is a measurement) lead to efficient methods to sample complex functions such as four-dimensional reflectance functions. These measurements are only meaningful if the appropriate equipment is used, the properties of the devices are known, and if the relevant parts are calibrated.

4.1 Acquisition and Lighting Equipment

Although both analog and digital cameras can be used for measurement purposes, digital cameras are used in basically all current systems. They combine high accuracy, speed, and flexibility with the possibility to control them re-



Figure 2. A photo studio covered with black, diffuse reflecting material.

motely by a computer. A wide range of techniques is available to measure and improve their performance.

For most algorithms that reconstruct the appearance properties of an object from images, it is important to control the lighting conditions exactly. A point light source, i.e., a light source where all light is emitted from a single point, is ideal for many of the techniques mentioned above. A perfectly constant and diffuse lighting is ideal to capture the color of an object. Both lighting conditions are rarely used in normal photography and special solutions have to be implemented.

Care should be taken that the surrounding of the object has little to no influence on the measurements. This can be achieved by covering the surrounding with black, diffusely reflecting material, as shown in Fig. 2.

A more technical and in-depth discussion of camera and lighting issues can be found in [7].

4.2 Camera Calibration

When using a camera as a measurement device, various aspects should be calibrated in order to achieve precise results and repeatable measurements.

The properties of the camera transformation which describes how an object is projected onto the camera's image plane should be recovered using geometric camera calibration tools such as [34, 36]. These methods determine camera parameters such as the focal length of the lens, the principal point, and various distortion coefficients. Once this information is known, a ray in space can be assigned to each pixel in an image.

The dynamic range of a camera is for most cameras much smaller than the dynamic range of a scene. Some techniques must therefore be used to capture the full dynamic range of such a scene.

Several authors proposed methods to extend the dynamic range of digital images by combining multiple images of the

same scene with different exposure times. Debevec and Malik [6] and Robertson et al. [27] recover the response curve of the imaging system and linearize the input data before combining them into a single high dynamic range image. Goesele et al. [8] have shown a way to combine color calibration with high dynamic range imaging in order to improve the color fidelity of appearance acquisition methods.

5 Appearance Acquisition

The appearance of an object consists of several surface properties including color, texture, reflection properties, and normal directions. Due to their large number and complexity they are difficult to acquire but nevertheless necessary to generate convincing looking images of objects. It is therefore justifiable to put a lot of effort into this step.

The appearance of an object is traditionally captured using a variety of special devices [11]. But many surface properties can also be acquired using a camera and controlled lighting. Captured images can for example be used to color the geometry model during rendering. The digital pictures are simply projected onto the model as image textures using texture mapping [9]. To obtain more precise surface properties than just a single color value, more elaborate techniques are needed.

5.1 Reflection Properties

Most surfaces have a very complex micro-structure which determines their appearance and reflection properties. The *bidirectional reflectance distribution function* (BRDF or reflectance model) is used to describe how a surface reflects light. The BRDF yields the fraction of light arriving at a point from one direction to the light that is reflected off the surface at the same point into an exitant direction. It is a four-dimensional function $f_r(\hat{\omega}_o, \hat{\omega}_i)$ that depends on the incident light direction $\hat{\omega}_i$ and the viewing direction $\hat{\omega}_o$. A number of analytical BRDF models such as [17, 33, 35] have been developed to approximate the reflection properties of real materials.

5.2 Measuring Reflection Properties

There are special devices available to to measure real-world BRDFs: The most general approach is to use a gonireflectometer which measures the light that is emitted in every direction when the object is illuminated from a given direction. This measurement procedure can however be time consuming and captures only the properties of a single point on the surface of an object.

Marschner [21] used a digital camera to determine a single BRDF for an object by combining all the pixel data.

Compared to a gonireflectometer this technique is considerably faster, but it still assumes that the entire object consists of a single material. To allow for variations of the reflectance properties over the object’s surface, Marschner et al. [22] extracted the purely diffuse part (albedo map) of the object’s texture for each visible point. The resulting texture includes only view-independent color information and no specular reflection.

5.3 Measuring Spatially Varying BRDFs

Based on Marschner’s approach, Lensch et al. [18] developed a technique that is able to reconstruct spatially varying reflection properties by just a few images (around 15–25). The key idea here is that most objects typically consist of a small number of materials, i.e., many points on the object’s surface have approximately the same reflection properties. By clustering points consisting of the same material, a large number of BRDF samples can be collected by just a few images. After generating the BRDFs for clusters of points, separate reflection properties for each single point are determined to account for subtle details and small changes. The BRDF for each point is determined as a weighted sum of the clusters’ BRDFs. Thus, a high quality and very compact representation of the original object can be obtained with moderate acquisition effort.

5.3.1 Data Acquisition

The entire procedure is as follows: The geometry of the object is obtained as triangle mesh. In order to capture the reflection properties a small number of high dynamic range images of the object are taken showing the object lit by a single point light source. In a next step the camera position (see Section 6) as well as the light source position relative to the geometric model are recovered for all images.

For every point on the object’s surface all available data (position, normal, and a list of reflectance samples together with the corresponding viewing and lighting directions) is collected from the different views in a data structure called *lumitexel*. The reflectance samples are obtained by resampling the color value at the position of the surface point projected into the image.

5.3.2 Clustering of Materials

Because only a limited number of different views and lighting directions is acquired, a single lumitexel does not carry enough information to reliably fit a BRDF model to the reflectance samples. To provide more data from which the parameters can be derived, the lumitexels are grouped into clusters of similar materials. Starting with a single cluster containing all lumitexels, the parameters of an average BRDF are fitted using the Levenberg-Marquardt algorithm




Figure 3. The clustering process at work. In every image a new cluster is created. The object is reshaded using only the single BRDFs fitted to each cluster before the projection into a basis of multiple BRDFs.

to perform a non-linear least square optimization. In order to separate the distinct materials the initial cluster has to be split. Given the average BRDF, two new sets of parameters are generated by varying the fitted parameters along the direction of maximum variance, yielding two slightly different BRDFs. The lumitexels of the original cluster are then assigned to the nearest of these BRDFs, forming two new clusters. A stable separation of the materials in the clusters is obtained by repeatedly fitting BRDFs to the two clusters and redistributing the original lumitexels. Further splitting isolates the different materials until the number of clusters matches the number of materials of the object as illustrated in Fig. 3.

5.3.3 Spatially Varying Behavior

After the clustering the same reflection behavior is assigned to all lumitexels in one cluster. However, small features on the surface and smooth transitions between adjacent materials can only be represented if every lumitexel is assigned its own BRDF. In this algorithm, this BRDF is a weighted sum of the BRDFs recovered by the clustering procedure. The spatially varying reflection properties can then be represented by a set of basis BRDFs for the entire model plus a set of weighting coefficients for each lumitexel.

The weighting coefficients are found by projecting the lumitexel’s data into a basis of per-cluster BRDFs. An optimal set of weighting coefficients minimizes the error between the measured reflectance and the weighted sum of reflectance values obtained by evaluating the basis BRDFs for the viewing and lighting direction of the measured sample. To recover the coefficients, the least square solution of the corresponding system of equations is computed (see [18] for more details).

5.4 Normal Maps

The resolution of the acquired geometry of an object is typically limited by the 3D scanning device. Additional processing and filtering of the 3D data further erase fine




Figure 4. Left: Normals of the original mesh. Right: Normals optimized using spatially varying BRDFs

scale geometric detail. When reconstructing the object using a coarse geometric model, small features in the surface structure such as bumps, cracks or wrinkles can be simulated with normal or bump maps [3]. These textures store a perturbation of the surface normal for each surface point.

Normal maps can be acquired for real world objects: Rushmeier et al. [30] calculated normal directions from a set of images showing the same view of the object illuminated by a point light source placed at different, known positions. The surface is assumed to be perfectly diffuse and thus its color can be represented by an albedo map [29]. The restriction of a purely diffuse surfaces can be removed if appearance acquisition techniques are used to first measure the approximate reflection properties at each surface point.

To measure the exact normal at a point, an initial normal is obtained from the triangular mesh. Given the viewing and lighting directions for the reflectance samples in world coordinates, the current estimate of the normal is used to transform them into the local coordinate frame. Then, the error between the measured reflectance values and the reconstructed reflectance values is computed. The reconstructed reflectance values are obtained by evaluating the measured BRDF using the transformed directions. If enough reflectance samples are provided for each point the actual normal direction at the point can be found by minimizing this error using a non-linear least square optimization technique. Fig. 4 compares the quality of the reconstructed normals to the normals of the original mesh.

6 Registration of Geometry and Texture Data

Since texture and geometry are typically acquired by two different processes, the collected data has to be merged afterwards. This requires the alignment of the geometry data and the captured images.




Figure 5. Measuring the difference between photo (bottom) and one view of the model (top) by the area occupied by the XOR-ed foreground pixels.

In order to align or register the 3D model to the texture data one has to recover the parameters of the camera transformation that maps points in 3-space (the 3D geometry) onto the 2D image, i.e., the camera position, orientation and focal length. Further intrinsic parameters are assumed to be already known.

A simple approach to recover the camera position and orientation is to manually select corresponding points on the geometric model and in the picture [28]. If enough correspondences are established, the transformation can be determined directly.

In order to simplify the registration process some semi-automatic approaches have been published [23, 25]. The user is asked to roughly align the 3D model to the image. The algorithm then tries to optimize for the camera parameters by minimizing the distance between the outline of the 3D model rendered with the current set of camera parameters and the outline of the object found in the image. For each tested set of camera parameters the distance between the outlines has to be computed which is computationally expensive.

In [19], Lensch et al. proposed a method to compute the distance between a view of the 3D model and the 2D image by comparing silhouettes instead of outlines. The geometry is first rendered in front of a black background in a monochrome color and afterwards combined with the segmented image using the XOR-operation as shown in Fig. 5. The resulting image is black except for those pixels covered by only one of the silhouettes. The number of remaining pixels is a measure for the distance between the silhouettes. It can be counted by evaluating the histogram. Minimizing the number of remaining pixels leads to an optimal set of camera parameters. Note that all of these steps can be performed using graphics hardware to speed up the optimization. Additionally, it is also possible to automatically find a

rough initial guess for the camera parameters.

6.1 Texture Preparation

Given the camera transformation for an image, it can be stitched onto the surface of the 3D model. For a triangular mesh this is done by computing texture coordinates for every vertex of the model visible in the image. Texture coordinates are calculated by projecting the 3D coordinates of the vertices into the image plane using the recovered camera transformation. The exact transformation for projecting surface points into the images is now known. This information is required when collecting all reflectance samples for one point on the objects surface into a lumitexel (compare Section 5.3.1). In order to obtain the highest quality with respect to the input images, lumitexels are generated at surface points that correspond to the pixels in the original image with the highest sampling rate for a particular triangle.

Since every lumitexel is assigned to a triangular region within one of the HDR images, it is possible to construct a 2D texture of lumitexels. This texture will unfortunately consist of a large number of separate triangles. Instead of treating these regions as independent textures, it is more convenient to pack the regions into a single image, e.g., using the technique proposed by Rocchini et al. [28]. This not only reduces the size of the data set but makes it also more suitable for compression and progressive transmission.

7 Interactive Display

After measuring the reflection properties of an object, we explain now how it can be displayed interactively.

We first investigate the case of one homogeneous material, i.e., one BRDF per object. Standard OpenGL only supports the empirical and physically implausible Phong model, which makes surfaces always look “plastic”-like.

In order to render surfaces with other BRDFs, two similar approaches [10, 13] can be used. Both approaches decompose the four-dimensional BRDF $f_r(\hat{\omega}_o, \hat{\omega}_i)$ into a product of two two-dimensional functions $g(\hat{\omega}_o)$ and $h(\hat{\omega}_i)$. These functions are stored in two texture maps. Rendering with these decomposed BRDFs is done by simply re-multiplying the texture maps using blending. Unfortunately, these techniques cannot simulate small surface bumps or spatial variation of reflection properties.

7.1 Rendering with Normal Maps and Spatially Varying BRDFs

Blinn [3] has shown how wrinkled surfaces can be simulated without changing the underlying surface by only perturbing the normal vector. The perturbed normal is then



Figure 6. Left: A spatially varying BRDF applied to a sphere. Right: A vinyl bust illuminated by a sunset environment.

used for the lighting calculations instead of the original surface normal. This technique is generally called normal or bump mapping.

Early algorithms [15] for rendering bump maps at interactive rates use texture maps containing per-pixel normals, which are used to perform the lighting calculations instead of per-vertex normals. These early methods only support the simple Blinn-Phong model [4], whose material properties are assumed to be constant across the surface, i.e., these methods do not allow for spatial variation.

If the reflection properties vary across the surface, spatially varying BRDFs must be considered. For rendering the captured models we use an algorithm based on Kautz et al. [12] (see Fig. 6), which we extended so that the acquired data can be rendered interactively (similar to [24]). For each basis BRDF the object is rendered once, and all the contributions are summed up in the frame buffer. Using this method, objects can be rendered interactively and allow for changing illumination. More recent work [14, 31] can even include area lighting instead of point lights as well as self-shadowing; an example is shown in Fig. 6.

8 Examples

In this section we describe some examples for high quality 3D object acquisition. Geometry and reflection data have been acquired for a bronze bust of Max Planck, a clay bird, and a painted models of two angles. Some statistics are listed in Table 1.

The triangle mesh of the angels was generated by extracting an isosurface of a CT scan, the bust and the bird were acquired using a structured light 3D scanner.

The images for the textures and reflection properties were taken with a professional digital camera. Several views with varying light source positions were captured for each model (see Table 1). For each view around 15

model	Δ	views	lumi- texels	rad. samples	clus- ters	basis BRDFs
angels	47k	27	1.6M	7.6	9	6
bird	14k	25	1.9M	6.3	5	4
bust	50k	16	3.6M	4.2	3	4

Table 1. This table lists the number of triangles of each model (Δ), the number of views used to reconstruct the spatially varying BRDFs, the number of acquired lumitexels and the average number of radiance samples per lumitexel, the number of partitioned material clusters, and the number of basis BRDFs per cluster.

photographs were necessary: two for recovering the light source position, one to extract the silhouette of the object for the 2D–3D registration, and the rest to provide the necessary high dynamic range.

Fig. 7 shows on the left only the per-cluster BRDFs. Details are missing and the object looks rather flat. After performing the projection every lumitexel is represented as a linear combination in a basis of four BRDFs, now resulting in a much more detailed and realistic appearance. The bust in Fig. 8 shows another reconstructed object with very different reflection properties. The bronze look is very well captured.

A comparison between an object rendered with an acquired BRDF (using the presented method) and a photograph of the object is shown in Fig. 9. They are very similar, but differences can be seen in highlights and in places where not enough reflectance samples were captured. Capturing



Figure 7. Left: Last result of the clustering step. Right: Bird with the spatially varying BRDF determined by projecting each lumitexel into a basis of BRDFs. Note the subtle changes of the materials making the object look realistic.



Figure 8. A bronze bust rendered with a spatially varying BRDF, which was acquired with the presented reconstruction method.

more samples will increase the quality. The difference in the hair region is due to missing detail in the triangle mesh. Those would be resolved by recovering the normal map for the object.

For all the models, generally only a few clusters were needed to accurately represent all the materials since the projection takes care of material changes. In our experiments even Lafortune BRDFs [17] consisting of a single lobe were sufficient to form good bases for the clustering and projection.

9 Conclusion

We presented some recent advances in capturing, processing, and interactive display of real-world objects. The acquired models are fairly complete representations of the geometry and surface properties of a large class of real world objects. In order to achieve the highest possible qual-




Figure 9. Left: Photograph of model. Right: Model with acquired BRDF rendered from the same view with similar lighting direction.

ity, state-of-the-art computer vision and computer graphics techniques need to be used throughout the 3D object pipeline.

Given such a detailed model, many computer vision algorithms such as reconstructing surface normals [30] or material detection can be improved or extended to new classes of objects. Common assumptions about the characteristics of the object (e.g., purely diffuse reflection) are no longer necessary.

References

- [1] F. Bernardini, J. Mittleman, and H. Rushmeier. Case study: Scanning michelangelo's florentine pieta. In *Course Notes for SIGGRAPH 1999*, August 1999.
- [2] P. Besl and N. McKay. A method for the registration of 3-D shapes. *IEEE Trans. PAMI*, 14(2):239–258, 1992.
- [3] J. Blinn. Simulation of Wrinkled Surfaces. In *Proc. SIGGRAPH*, pages 286–292, Aug. 1978.
- [4] J. Blinn. Models of Light Reflection For Computer Synthesized Pictures. In *Proc. SIGGRAPH*, pages 192–198, July 1977.
- [5] B. Curless and S. Seitz. 3D Photography. In *Course Notes for SIGGRAPH*, July 2000.
- [6] P. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proc. SIGGRAPH*, pages 369–378, August 1997.
- [7] M. Goesele, W. Heidrich, H. P. Lensch, and H.-P. Seidel. Building a Photo Studio for Measurement Purposes. In *Proc. VMV 2000*, November 2000.
- [8] M. Goesele, W. Heidrich, and H.-P. Seidel. Color calibrated high dynamic range imaging with ICC profiles. In *Proc. CIC-9*, Scottsdale, USA, 2001.
- [9] P. Haeberli and M. Segal. Texture Mapping As A Fundamental Drawing Primitive. In *EG Rendering Workshop*, pages 259–266, June 1993.
- [10] W. Heidrich and H.-P. Seidel. Realistic, hardware-accelerated shading and lighting. In *Proc. SIGGRAPH*, pages 171–178, August 1999.
- [11] R. S. Hunter and R. W. Harold. *The measurement of appearance*. Wiley, 2. ed. edition, 1987.
- [12] J. Kautz and H.-P. Seidel. Towards interactive bump mapping with anisotropic shift-variant brdfs. In *Proc. EG/SIGGRAPH Workshop on Graphics Hardware*, pages 51–58, Aug. 2000.
- [13] J. Kautz and M. McCool. Interactive Rendering with Arbitrary BRDFs using Separable Approximations. In *EG Rendering Workshop*, pages 281–292, June 1999.
- [14] J. Kautz, P.-P. Sloan, and J. Snyder. Fast, Arbitrary BRDF Shading for Low-Frequency Lighting Using Spherical Harmonics. In *13th Eurographics Workshop on Rendering*, pages 301–308, June 2002.
- [15] M. Kilgard. *A Practical and Robust Bump-mapping Technique for Today's GPUs*. NVIDIA Corporation, April 2000.
- [16] L. Kobbelt. Discrete fairing. In *Proc. Seventh IMA Conf. on the Mathematics of Surfaces*, pages 101–131, 1996.
- [17] E. Lafortune, S.-C. Foo, K. Torrance, and D. Greenberg. Non-Linear Approximation of Reflectance Functions. In *Proc. SIGGRAPH*, pages 117–126, August 1997.
- [18] H. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially varying materials. In *EG Rendering Workshop*, pages 104–115, London, Great Britain, 2001. Springer.
- [19] H. P. A. Lensch, W. Heidrich, and H.-P. Seidel. A silhouette-based algorithm for texture registration and stitching. *Graphical Models*, 63(4):245–262, July 2001.
- [20] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital michelangelo project: 3D scanning of large statues. In *Proc. SIGGRAPH*, pages 131–144, July 2000.
- [21] S. R. Marschner. *Inverse rendering for computer graphics*. PhD thesis, Cornell University, 1998.
- [22] S. R. Marschner, B. Guenter, and S. Raghupathy. Modeling and rendering for realistic facial animation. In *EG Rendering Workshop*, pages 231–242, June 2000.
- [23] K. Matsushita and T. Kaneko. Efficient and handy texture mapping on 3d surfaces. *Computer Graphics Forum*, 18(3):349–358, September 1999.
- [24] D. McAllister, A. Lastra, and W. Heidrich. Efficient Rendering of Spatial Bi-directional Reflectance Distribution Functions. In *Proceedings Graphics Hardware*, pages 79–88, September 2002.
- [25] P. J. Neugebauer and K. Klein. Texturing 3d models of real world objects from multiple unregistered photographic views. *Computer Graphics Forum*, 18(3):245–256, September 1999.
- [26] K. Pulli, T. Duchamp, H. Hoppe, J. McDonald, L. Shapiro, and W. Stuetzle. Robust meshes from multiple range maps. In *Proc. IEEE 3DIM97*, 1997.
- [27] M. A. Robertson, S. Borman, and R. L. Stevenson. Dynamic Range Improvement Through Multiple Exposures. In *Proc. ICIP*, pages 159–163. IEEE, Oct. 1999.
- [28] C. Rocchini, P. Cignoni, and C. Montani. Multiple textures stitching and blending on 3D objects. In *EG Rendering Workshop*. Eurographics, June 1999.
- [29] H. Rushmeier, F. Bernardini, J. Mittleman, and G. Taubin. Acquiring input for rendering at appropriate levels of detail: Digitizing a pieta. *EG Rendering Workshop*, pages 81–92, June 1998.
- [30] H. Rushmeier, G. Taubin, and A. Guézic. Applying shape from lighting variation to bump map capture. *EG Rendering Workshop*, pages 35–44, June 1997.
- [31] P.-P. Sloan, J. Kautz, and J. Snyder. Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments. In *Proceedings SIGGRAPH*, pages 527–536, July 2002.
- [32] G. Taubin. A signal processing approach to fair surface design. *Proc. SIGGRAPH*, pages 351–358, 1995.
- [33] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the OSA*, 57(9):1105–1114, Sept. 1967.
- [34] R. Tsai. A versatile camera calibration technique for high accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE Journal of Robotics and Automation*, 3(4), Aug. 1987.
- [35] G. Ward. Measuring and modeling anisotropic reflection. In *Proc. SIGGRAPH*, pages 265–272, July 1992.
- [36] Z. Zhang. A Flexible New Technique for Camera Calibration. Technical Report MSR-TR-98-71, Microsoft Research, 1999. Updated version of March 25, 1999.