The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage
VAST (2005)
M. Mudge, N. Ryan, R. Scopigno (Editors)

Rapid Synchronous Acquisition of Geometry and Appearance of Cultural Heritage Artefacts

Gero Müller, Gerhard H. Bendels, Reinhard Klein

Institut für Computergraphik, Universität Bonn, Germany
{gero,bendels,rk}@cs.uni-bonn.de

Abstract

In order to produce visually appealing digital models of cultural heritage artefacts, a meticulous reconstruction of the 3D geometry alone is often not sufficient, as colour and reflectance information give essential clues of the object's material. Standard texturing methods are often only able to overcome this fact under strict material and lighting condition limitations. The realistic reconstruction of complex yet frequently encountered materials such as fabric, leather, wood or metal is still a challenge. In this paper, we describe a novel system to acquire the 3D-geometry of an object using its visual hull, recorded in multiple 2D images with a multi-camera array. At the same time, the material properties of the object are measured into Bidirectional Texture Functions (BTF), that faithfully capture the mesostructure of the surface and reconstruct the look-and-feel of its material. The high rendering fidelity of the acquired BTF texture data with respect to reflectance and self-shadowing also alleviates the limited precision of the visual hull approach for 3D geometry acquisition.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Picture/Image Generation]: Digitizing and scanning I.3.7 [Three-Dimensional Graphics and Realism]: Color, shading, shadowing, and texture I.4.1 [Digitization and Image Capture]: Reflectance

1. Introduction

Three-dimensional digitisation using laser range scanners has a long tradition in industrial applications, such as reverse engineering, quality control and part inspection, where the geometric error of the acquired data compared to the original model is the ultimate criterion. More recently, with hardware becoming available at more reasonable costs, the use of 3D data acquisition has expanded to other application fields like cultural heritage and entertainment industries. In these application fields, however, the faithful reconstruction of the visual appearance is often much more important than the geometric error alone.

In archeological and museum applications, where digital archives of cultural heritage artefacts become more and more commonplace, the demand for realistic reconstruction of an artwork's material is particularly high, as this often gives critical hints about important characteristics of the object, such as for instance the location, epoch, or circumstances under which the artwork was created.

A standard approach to digitise a 3D cultural heritage artefact would be to acquire the geometry first (up to a certain degree of detail) using a Laser Range Scanner, structured light, tactile sensors or even a volumetric approach on the basis of MRT- or CT-data. Given an appropriate parametrisation, additionally recorded 2D-images can then be projected onto the geometry as diffuse textures to represent colour and fine detail information — at the disadvantage that the illumination is fixed to the conditions under which the photographs were taken. Fixing the lighting conditions, unfortunately, severely limits the use of 3D digital representations in cultural heritage applications, because moving, rotating and seeing an artefact in different environments is an important operation that often gives additional insight. Traditional texturing methods are not capable of transporting the characteristic appearance of complex, inhomogeneous, yet often encountered materials like fabric, leather, or wood. Another problem of this standard approach, common to all of the above geometry acquisition methods is their inability to handle transparent, reflective or fuzzy materials. In ad-

Figure 1: *Leather Material, rendered using standard texturing (left), bump mapping (middle), and BTF-rendering (right).*

dition, geometry acquisition usually involves sophisticated procedures during, and after, recording (such as path planning in order to ensure detection of the complete surface, registration, reconstruction and parametrisation).

As a consequence, 3D photography (as it is also called) is often perceived by end-users in the museums as a complicated and time-consuming procedure with possibly non-satisfying results – in one word *impractical* for large classes of objects and applications.

With this situation in mind, we present in this paper a novel high fidelity acquisition system to synchronously capture an object's 3D geometry *and* material properties in a very time-efficient and user-friendly way. Our system exploits images from the multi-camera array depicted in Figure 3 first to reconstruct an artefact's coarse to medium scale geometry using a GPU-based visual hull technique, resulting in a closed triangle mesh. In parallel, the images are also used to capture the object's appearance into so-called bidirectional texture functions (BTF) – a 6-dimensional texture representation introduced by Dana et al. [DNvGK97] which extends the common textures by dependence on light- and view-direction, and thereby allows for photo-realistic rendering of an objects micro- and mesostructure (cf. Figure 1). The key contribution of this paper is a system that

- fully automatically acquires 3D-data, capturing an object's geometry *and* its visual appearance in form of bidirectional textures
- faithfully reconstructs the object's mesostructure using BTF-techniques and therefore effectively overcomes the limited accuracy of the visual hull technique
- is time efficient and very easy to use.

The reasoning behind our decision to also reconstruct the explicit geometry (compared to a purely image based and geometry-free approach) is that in addition to efficiency reasons, having direct access to the geometrical properties appears to be a virtue in itself for various purposes, such as

- Statistical Shape Analysis,
- Interaction with other objects (collisions, shadows, ...)
- Multimodal interaction (sound, haptics)
- Modelling applications (surface editing)

The rest of this paper is organised as follows: In the following section we will shortly review the relevant literature related to our approach presented in this paper, an overview over which is presented in section 3. The setup of our multi-camera acquisition device used to capture the required images will be described in section 4. Our modified Visual Hull algorithm is discussed in section 5.1, whereas section 6 deals with the techniques used to record and render the BTF equipped model.

2. Related Work

The desire to create synthetic photo-realistic images of objects under new lighting conditions and for new view points has a long tradition in computer graphics.

Traditionally the geometry of a surface is modelled explicitly (e.g. with triangles) only up to a certain scale, while the remaining surface properties responsible for the reflectance behaviour of a material are simulated using relatively simple analytical bidirectional reflectance distribution function (BRDF) models (e.g. [LFTG97]). Fine textural surface detail is typically captured by photographic images or procedurally generated textures which are projected onto the geometry.

Although this traditional approach has produced remarkable results for numerous examples, it becomes unfeasible for large classes of materials where the so-called mesostructure is essential for faithfully reproducing important characteristic material properties.

Numerous researchers have proposed image-based methods to overcome these hindrances. Independently, Gortler et al. [GGSC96] and Levoy and Hanrahan [LH96] introduced light-field rendering, an efficiently renderable 4D representation of the plenoptic function. In their approaches, an object's complete appearance is stored in a fourdimensional function (called *Light Field* in [LH96] and *Lumigraph* in [GGSC96]), that describes the flow of light at the sampled positions in the sampled directions. While these approaches allow photographic images – interpreted as a sampling of the complete light field – to be interpolated to generate images from new viewpoints, the illumination has to remain constant.

To overcome this drawback, several researchers proposed representing an object in terms of its reflectance field [DHT*00], [HCD01]. Recently, Hawkins et al. also proposed a novel, dual way of measuring an object's reflectance properties [HED05]. In particular [HCD01] demonstrated remarkably realistic synthesised images of cultural heritage artefacts under new lighting conditions, although the generation of novel viewpoints remained basic.

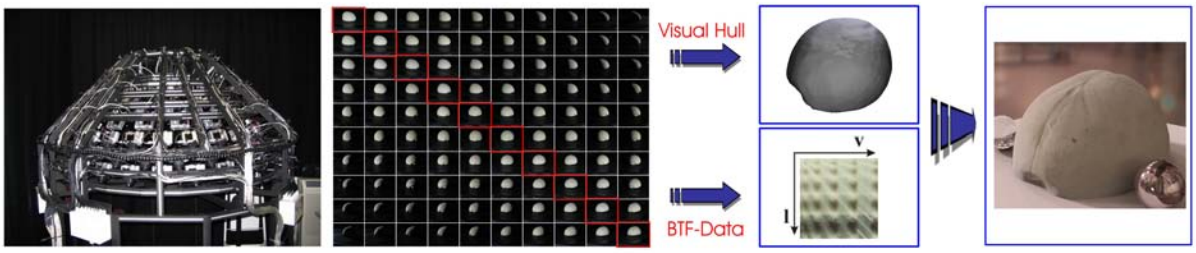

Figure 2: *The Acquisition Pipeline: The multi-camera array records 151 images per camera per light direction, resulting in 22801 images. Here, only the first ten cameras (from left to right) and the first ten light directions (from top to bottom) are shown. From the full set of these images, the BTF is constructed, while only a subset (the diagonal in this matrix notation) is used for the geometry reconstruction.*

In the above approaches, no explicit geometry is present, only renderings from new view points and/or under synthesised lighting conditions are feasible. In recent years, though, researchers have found the classical two-fold representation of an object as *geometry* + *material* to be superior in terms of time and memory efficiency [WAA*00][CHLG05][Pul97]. We also rely on such a two-fold representation but our material representation captures both light *and* view dependent appearance variation. Therefore, in concept our approach is most similar to the methods of Furukawa et al. [FKIS02] and Lensch et al. [LKG*03], who also construct a representation that allows rendering objects from novel viewpoints under arbitrary lighting. However, as they employ laser range scanners to record the 3D geometry of the object, still large classes of objects, in particular those with complex reflectance behaviour, can not be handled.

In our approach we use an array of fixed cameras with fixed light sources, all mounted on a hemispherical gantry. Although similar acquisition setups have been used by numerous research groups, see e.g. [MPZ*02], [FKIS02], [HCD01], [MGW01], our setup benefits from massive parallelisation and the fact that no moving parts are required – a fact that renders tedious and time consuming tasks like recalibration, registration, etc. unnecessary.

As in our work, the simultaneous reconstruction of geometry and appearance is also an inherent part of the numerous methods aiming at 3D-reconstruction from uncalibrated image-sequences (e.g. [PGV*04]). But these techniques are neither designed nor capable of performing the highly accurate reflectance measurements possible with a fully computer controlled and calibrated measurement device.

3. Overview

The basic concept of our algorithm is as follows: In the acquisition phase, we record a dense set of photographs from pre-defined viewpoints, regularly sampled over the hemisphere. For each viewpoint $v_1,\ldots,v_m$, we illuminate the object from light sources likewise regularly positioned at $l_1,\ldots,l_n$. This recording is massively parallelised (as will be described in the following section), resulting in a set $\{I_{ij}\}_{i=1\ldots m, j=1\ldots n}$ of $m\times n$ Images, where I_{ij} denotes the image taken from viewpoint v_i of the object to be digitised under illumination from l_j.

With these images at hand, the first step is to reconstruct a triangulated surface representation from the object. To this end, we apply a GPU-based modification of the classical volume carving method to transform the visual hull representation derived from the dense set of images to a volumetric representation, from which the final triangle mesh can then be reconstructed. For this step, we only exploit the subset $\{I_{ii}\}_{i=1\ldots m}$ of the total set of images, i.e. those images where light and viewing direction is coincident.

After parametrisation of the resulting triangle mesh, the next step is to exploit the full set of images to define the BTF. Efficient extraction, compression and storage of the BTF data is described in section 6. Figure 2 illustrates the whole process.

4. Multi-Camera Grid

For fast acquisition of the images required to measure the BTF and to reconstruct the geometry of the object to be digitised, we use an array of 151 commodity digital still cameras mounted on a hemispherical gantry (see Figures 3, 4). A similar gantry with mounted light sources was used by Malzbender et al. to capture *Polynomial Texture Maps* [MGW01]. By arranging the cameras into this array, the acquisition of the images required to construct the BTF textures is parallelised and no moving parts (e.g. a rotating stage) are needed. Therefore, the positions of the image sensors and the light sources can be calibrated in a preprocessing step which only has to be carried out if a camera has been replaced or after

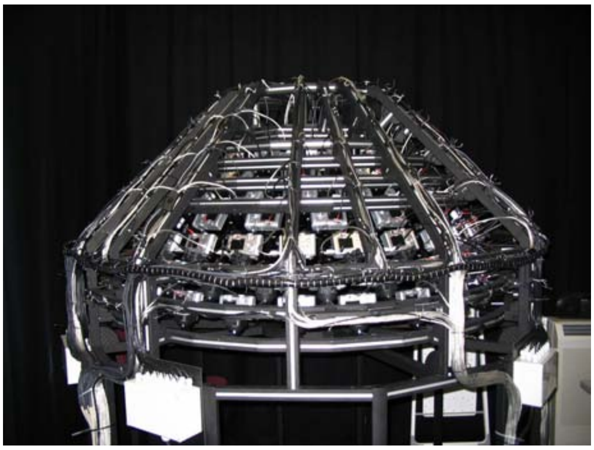

Figure 3: The Dome - *151 cameras mounted on a hemispherical gantry.*

Figure 4: *Detail view of the dome. Only commodity class, off-the-shelf cameras have been used.*

the whole setup has been transported. The low-level post-processing (geometric correction, colour correction) is fast enough to be done in parallel to the measurement.

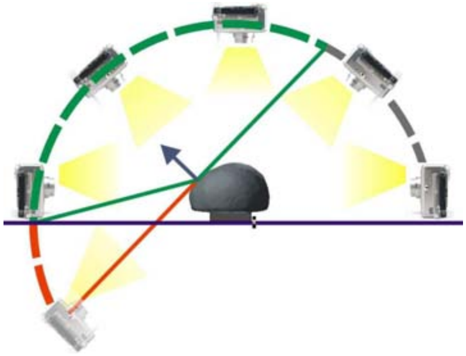

Figure 5: *Hemispherical setups – in contrast to full spheres of cameras and light sources – produce a subsampling of the full set of possible light and view directions.*

For non-planar objects, our hemispherical setup can only deliver a subsampling of the full set of possible light and view directions. Figure 5 illustrates this effect. For geometry and appearance acquisition, however, this subsampling can easily be completed to cover the full sphere by repositioning the object. To this end, a registration of the recorded images after repositioning with the original images is required. To solve this task, we apply the automatic registration technique given in [BDW*04]. It is worth noting, that with a full automatic technique at hand, this repositioning approach can also be used to deliberately increase the sampling rate, that would otherwise be limited by the finite number of cameras mounted on the gantry.

For our setup, we used consumer-class, off-the-shelf cameras with an image resolution of 2048 × 1536 pixels, coated with a diffuse black finish to reduce scattered light effects in the images and controlled by a network of eight commodity personal computers. As light sources, the built-in flash lights of the cameras are used, no further light sources are required. Formally, we therefore have $n = m$ and roughly $v_i = l_i$ for all $i = 1, \ldots, n$ in the above formulation. It is worth noting, however, that as a consequence of the massively parallelised setup, we are able to capture a full dataset of $151 \times 151 = 22801$ images in about 40 minutes.

5. Geometry Acquisition

One of the major advantages of our acquisition device without any moving parts is that camera and light source positions are known a-priori, hence so are the transformations required to transfer the different views into a common, global coordinate system. Thus, the registration of the multiple views is straight forward.

5.1. Visual Hulls

The idea to exploit a set of registered photographs of an object to derive its three-dimensional layout is far from new [Bau74]. Due to its conceptual simplicity and availability – only classical 2D photographic images are required – numerous authors have developed algorithms to reconstruct the 3D geometry of an object from images, long before 3D acquisition devices as laser range scanners have become commonly available. In our approach, we apply a volumetric approach in the concept of *shape from silhouette* [MA83].

In our current implementation, we first extract the objects silhouettes from the acquired 2D photographs by simple thresholding. As we use no backdrop in our setup, we set every pixel with a brightness of less then a certain threshold to be *outside*, the remaining pixel are *inside* (see Figure

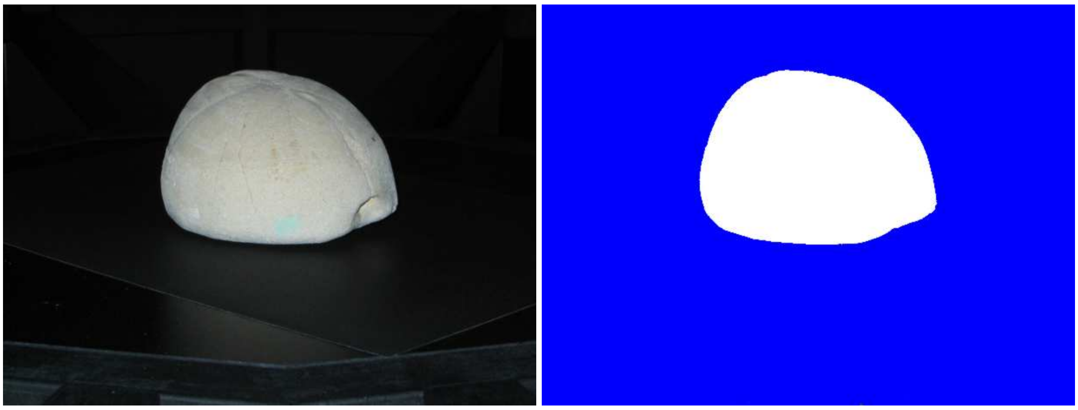

Figure 6: *One example image* $\in \{I_{ii}\}_{i=1,\ldots,n}$ *(for which light and view direction are coincident) and the corresponding binary* inside/outside-*image after thresholding*

6). In this step, we only exploit a subset of the set of all acquired images, namely those images where light and view directions are identical, i.e. $\{I_{ii}\}_{i=1,\ldots,n}$, resulting in a set of binary images $\{J_i\}_{i=1,\ldots,n}$.

Together with the viewpoint information known from our setup, every outside-pixel in each image now defines a ray in scene space that is known *not* to intersect the object, whereas the inside-pixel define rays that intersect the surface at some unknown distance to the viewpoint. In the continuous case (pixel width $\rightarrow 0$) the union of all these intersecting rays would define a generalised cone that is guaranteed to contain the object.
As this fact holds for all acquired images, the intersection of all these generalised cones (the *Visual Hull*, [Lau94][MBR*00]) describes a tight volume in space in which the complete object must lie.

We make use of this guarantee by applying volume carving to a regular grid, i.e. we (conceptually) traverse the grid following the *outside*-rays and mark any cell that is encountered during traversal as *empty*. The triangle mesh is then reconstructed using the well-known marching cubes approach [LC87].

5.2. Efficient Evaluation on the GPU

The large number of acquired images and the (potential) need for finer grids make it impractical to actually traverse the grid following the *outside*-rays. Instead, we use a hardware-supported approach based on projective texture mapping, that we will shortly describe here for completeness.

Suppose, we have an axis-parallel grid of dimension $X \times Y \times Z$, corresponding to the axis directions $\mathbf{e}_x$, $\mathbf{e}_y$, and $\mathbf{e}_z$, respectively. For reconstruction, we have to assign to each grid point 0 (for outside) or 1 (for inside). Let $\mathbf{v}_i$ be the viewing direction from viewpoint v_i onto the object, and let (without loss of generality) $\mathbf{e}_x$ be the direction such that the scalar product $|\langle \mathbf{e}, \mathbf{v}_i \rangle|$ is maximal.

We interpret the grid to be a stack of binary 2D textures $\{T_k\}_{k=1,\ldots,X}$, where each T_k has a resolution of $Y \times Z$ pixels, see Figure 7. The *inside/outside*-information is then efficiently collected by projecting every source image J_i to each texture T_k. We perform bitwise AND-operations during this projective texture mapping, to set a pixel in F_k to 0 if at least one J_i indicates so.

6. Bidirectional Texture Functions

The image-based approach to capturing the appearance of an object for later rendering is to take dense sets of images under controlled viewing and lighting conditions in order to sample its reflectance field appropriately. As mentioned, Malzbender et al. [MGW01] captured the lighting variability of a texture by taking images lit from different directions and compressed the data to a compact representation called Polynomial Texture Maps. As we also want to be able to vary the viewing direction, we rely on BTFs. Mathematically the BTF can be expressed as a measured 6D-slice of the general 8D-reflectance field

$$\mathbf{RF}_{rgb}(\mathbf{x}_i \rightarrow \mathbf{x}_0, \omega_i \rightarrow \omega_0)$$

parameterised over a base surface S:

$$\mathbf{BTF}_{rgb}(\mathbf{x}, \omega_i \rightarrow \omega_0) = \int_S \mathbf{RF}_{rgb}(\mathbf{x}_i \rightarrow \mathbf{x}, \omega_i \rightarrow \omega_0) d\mathbf{x}_i$$

Please note that fixing the position $\mathbf{x}$ restricts the BTF at this position to a 4-dimensional BRDF, which is often called *apparent BRDF*, because this function still contains local self-occlusions, scattering, etc. [MMS*05].

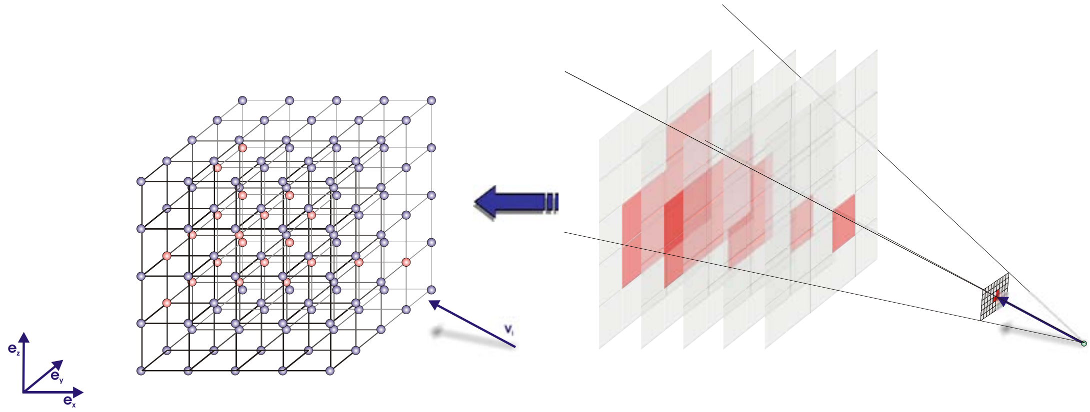

Figure 7: *The 3D grid is interpreted as set of binary images.*

6.1. Compression

One major issue of high fidelity BTF measurements are the huge memory requirements. A raw BTF data set easily requires multiple Gigabytes of memory, requirements that cannot be fulfilled even by the most state-of-the-art graphics hardware. Therefore, BTF data is usually compressed either by fitting standard BRDF models to the data or by using statistical analysis methods, such as principal component analysis (PCA) or clustered PCA. The latter is used in our current implementation, and the resulting data can easily be decompressed and evaluated using graphics hardware [MMK04].

Statistical analysis, however, requires that data entries in the BTF are semantically correspondent – an assumption that holds for the raw data only under the assumptions of planarity, orthographic projection and directional light sources. This is not the case for our system since the dimensions of our acquisition setup cannot be considered "large" compared to the dimensions of the objects to be digitised. Therefore, we perform resampling of the raw BTF data based on a planar parameterisation of the reconstructed triangle mesh computed with the method of Degener et al. [DMK03] *before compression*.

6.2. Incomplete BTF data for non-planar objects

To capture the reflectance of a material independently of the specific geometric layout, most common approaches for BTF acquisition record images of a small planar material sample. Then a projection of the images on a common plane typically suffices. For digitising 3D cultural heritage artefacts, though, we cannot rely on planarity. Hence, we have to deal with effects like self-shadowing and occlusion.

Measuring BTFs generally consists of recording for every point on the surface its reflectance from each view direction under each light direction. For non-flat surfaces, however, the reflectance for some light and viewing directions will be zero (or close to) simply because of occlusion and/or self-shadowing. Using standard approaches, this *missing* information would be misinterpreted as a property of the material. Instead, we apply a different technique to interpolate the occluded data.

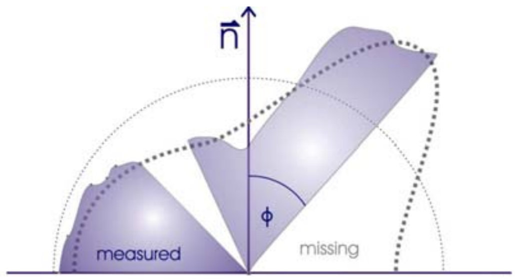

Figure 8: *Polar plot illustration of the measured BTF, that might be incomplete for some directions Φ due to occlusion and has to be completed, e.g. using statistical analysis methods.*

We first identify from all the points on the object surface those points that have an incomplete BTF measurement, i.e. points which are occluded for at least one light source or camera position. Figure 8 illustrates this situation.

For these points, the BTF has to be completed. In our current implementation, we first perform statistical analysis, here LPCA, of the measured BTF of all object surface points for which this is complete. The eigenvectors to the largest eigenvalues of the corresponding covariance matrices span a lower dimensional subspace approximating the original measured data. This way each individual bidirectional reflection function $\mathbf{BTF}_{rgb}(\mathbf{x}, \omega_i \to \omega_0)$ in a surface point $\mathbf{x}$ can then be approximated by

$$\mathbf{BTF}_{rgb}(\mathbf{x}, \omega_i \to \omega_0) = \sum c_k u_k,$$

where the u_k are the basis (apparent) BRDFs and the c_k the corresponding weights.

The incomplete BRDFs are completed individually, depending on the specific values that are missing. Let π be the projection of a full BRDF-vector $\mathbf{BTF}(\mathbf{x})$ to a lower dimensional subspace that only contains the *present* measurements and neglects the missing values. Then we find a set of coefficients $\{c_k\}_{k=1,\ldots,K}$ such that

$$||\pi(\mathbf{BTF}(\mathbf{x})) - \sum c_k \pi(u_k)||$$

is minimal. The reconstructed vector $\sum c_k u_k$ is a reasonable completion of $\mathbf{BTF}(\mathbf{x})$ and hence used in our approach. For complex materials this process can be iterated, while taking into account the already completed vectors.

7. Results & Conclusions

Our approach to reconstruct the geometry from the acquired images using visual hulls computed on the GPU is reliable and fast. Of course, identifying a nonconvex object using a silhouette-based approach inherently and inevitably implies neglecting some features of its surface geometry. Despite this general seemingly inaptness of the visual hull reconstruction, we were able to produce realistic images of our captured objects because the neglected surface features are well-captured in their appearance using the BTF texturing techniques. Figure 9 demonstrates this effect. The defect on the front size of this echinite from the Goldberg collection is not detected by the visual hull algorithm and therefore not represented in the geometric reconstruction. Its appearance is nevertheless preserved.

In order to use this appearance preservation such that users can explore an object and detect little concavities as depicted in Figure 9, it is important that users can arbitrarily choose viewing *and* lighting directions, which is possible using our system.

To further improve the geometric reconstruction, any of the numerous extensions of the visual hull reconstruction algorithm as e.g. presented in [IS03], [SCM*04], [LMS03a], [LMS03b], or [GSD03] can naturally be incorporated into our system. A promising direction for future work is also the incorporation of reconstruction techniques based on photometric correspondences as proposed in [FLP05].

8. Acknowledgements

This work was partially funded by the European Union under the project RealReflect (IST-2001-34744). Some of the used high-dynamic range lighting environments were taken from Paul Debevecs webpage. Special thanks belong to Ralf Sarlette for allowing us to use the acquisition dome and Dirk Koch for invaluable help with the acquisition.

References

[Bau74] BAUMGART B. G.: Geometric modeling for computer vision. Technical Report AIM-249, Artificial Intelligence Lab, Stanford University, October 1974. 4

[BDW*04] BENDELS G. H., DEGENER P., WAHL R., KÖRTGEN M., KLEIN R.: Image-based registration of 3d-range data using feature surface elements. In *5th Int. Symp. on Virtual Reality, Archaeology and Cultural Heritage (VAST)* (December 2004). 4

[CHLG05] COOMBE G., HANTAK C., LASTRA A., GRZESZCZUK R.: Online reconstruction of surface light fields. In *Eurographics Symposium on Rendering* (June 2005). 3

[DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2000), ACM Press/Addison-Wesley Publishing Co., pp. 145–156. 2

[DMK03] DEGENER P., MESETH J., KLEIN R.: An adaptable surface parametrization method. In *The 12th International Meshing Roundtable 2003* (September 2003). 6

[DNvGK97] DANA K., NAYAR S., VAN GINNEKEN B., KOENDERINK J.: Reflectance and texture of real-world surfaces. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Proceedings* (1997), 151–157. 2

[FKIS02] FURUKAWA R., KAWASAKI H., IKEUCHI K., SAKAUCHI M.: Appearance based object modeling using texture database: acquisition, compression and rendering. In *EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2002), Eurographics Association, pp. 257–266. 3

[FLP05] FURUKAWA Y., LAZEBNIK S., PONCE J.: Carved visual hulls for high-accuracy image-based modeling. In *Siggraph Sketches* (2005). 7

[GGSC96] GORTLER S. J., GRZESZCZUK R., SZELISKI R., COHEN M. F.: The lumigraph. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1996), ACM Press, pp. 43–54. 2

[GSD03] GRAUMAN K., SHAKHNAROVICH G., DARRELL T.: A bayesian approach to image-based visual hull reconstruction. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (June 2003), pp. 187–194. 7

[HCD01] HAWKINS T., COHEN J., DEBEVEC P.: A photometric approach to digitizing cultural artifacts. In *VAST '01: Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage* (New York, NY, USA, 2001), ACM Press, pp. 333–342. 2, 3

[HED05] HAWKINS T., EINARSSON P., DEBEVEC P.: A dual light stage. In *Eurographics Symposium on Rendering* (June 2005). 2

[IS03] ISIDORO J., SCLAROFF S.: Stochastic refinement of the visual hull to satisfy photometric and silhouette consistency constraints. In *Proc. 9th International Conference of Computer Vision* (Nice, France, October 2003), pp. 1335–1342. 7

[Lau94] LAURENTINI A.: The visual hull concept for silhouette-based image understanding. *IEEE Trans. Pattern Anal. Mach. Intell. 16*, 2 (1994), 150–162. 5

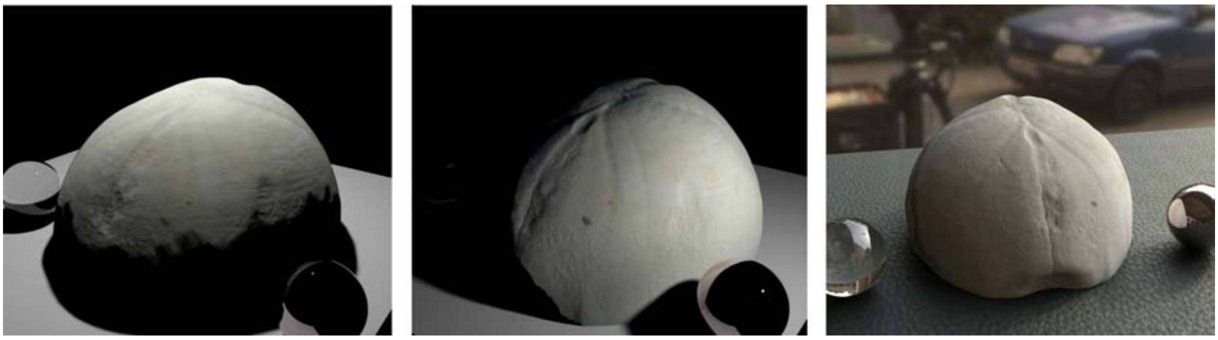

Figure 9: *Raytraced renderings of a captured and reconstructed echinite under novel lighting and viewing conditions. The left and middle image are rendered with a small area light source and demonstrate the fine geometric details captured in the BTF. The right image shows a relighting of the echinite with a complex image-based lighting environment captured in front of the Pädagogische Fakultät in Bonn. The ground floor in the right image is covered with a synthetic leather BTF courtesy of DaimlerChrysler AG.*

[LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1987), ACM Press, pp. 163–169. 5

[LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1997), ACM Press/Addison-Wesley Publishing Co., pp. 117–126. 2

[LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1996), ACM Press, pp. 31–42. 2

[LKG*03] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph. 22*, 2 (2003), 234–257. 3

[LMS03a] LI M., MAGNOR M., SEIDEL H.-P.: Hardware-accelerated visual hull reconstruction and rendering. In *In Proceedings of GIŠ03* (Halifax, Canada, 2003). 7

[LMS03b] LI M., MAGNOR M., SEIDEL H.-P.: Improved hardware-accelerated visual hull rendering. In *Vision, Modeling, and Visualization 2003* (Germany, 2003). 7

[MA83] MARTIN W. N., AGGARWAL J. K.: Volumetric description of objects from multiple views. *IEEE Trans. Pattern Analysis and Machine Intelligence 5*, 2 (1983), 150–158. 4

[MBR*00] MATUSIK W., BUEHLER C., RASKAR R., GORTLER S. J., MCMILLAN L.: Image-based visual hulls. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2000), ACM Press/Addison-Wesley Publishing Co., pp. 369–374. 5

[MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), ACM Press, pp. 519–528. 3, 5

[MMK04] MÜLLER G., MESETH J., KLEIN R.: Fast environmental lighting for local-pca encoded btfs. In *Computer Graphics International 2004 (CGI 2004)* (June 2004), IEEE Computer Society, pp. 198–205. 6

[MMS*05] MÜLLER G., MESETH J., SATTLER M., SARLETTE R., KLEIN R.: Acquisition, synthesis and rendering of bidirectional texture functions. *Computer Graphics forum 24*, 1 (March 2005), 83–109. 5

[MPZ*02] MATUSIK W., PFISTER H., ZIEGLER R., NGAN A., MCMILLAN L.: Acquisition and rendering of transparent and refractive objects. In *EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2002), Eurographics Association, pp. 267–278. 3

[PGV*04] POLLEFEYS M., GOOL L. V., VERGAUWEN M., VERBIEST F., CORNELIS K., TOPS J., KOCH R.: Visual modeling with a hand-held camera. *International Journal of Computer Vision 59*, 3 (2004), 207–232. 3

[Pul97] PULLI K. A.: *Surface reconstruction and display from range and color data*. PhD thesis, 1997. Chairperson-Linda G. Shapiro. 3

[SCM*04] SLABAUGH G. G., CULBERTSON W. B., MALZBENDER T., STEVENS M. R., SCHAFER R. W.: Methods for volumetric reconstruction of visual scenes. *Int. J. Comput. Vision 57*, 3 (2004), 179–199. 7

[WAA*00] WOOD D. N., AZUMA D. I., ALDINGER K., CURLESS B., DUCHAMP T., SALESIN D. H., STUETZLE W.: Surface light fields for 3d photography. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2000), ACM Press/Addison-Wesley Publishing Co., pp. 287–296. 3

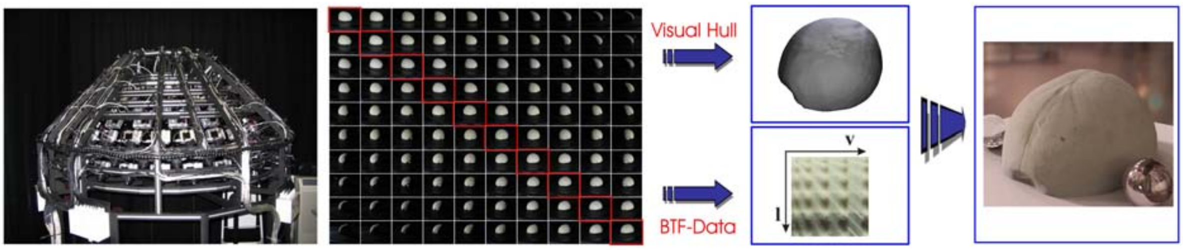

Figure 2: *The Acquisition Pipeline: The multi-camera array records 151 images per camera per light direction, resulting in 22801 images. Here, only the first ten cameras (from left to right) and the first ten light directions (from top to bottom) are shown. From the full set of these images, the BTF is constructed, while only a subset (the diagonal in this matrix notation) is used for the geometry reconstruction.*

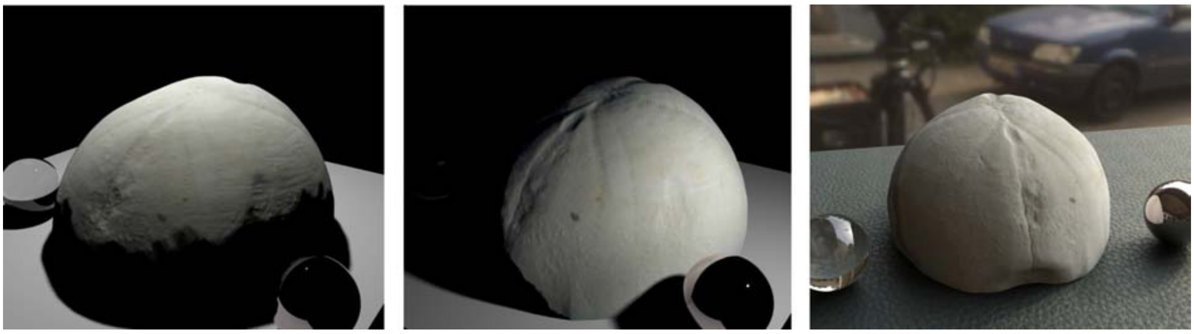

Figure 9: *Raytraced renderings of a captured and reconstructed echinite under novel lighting and viewing conditions. The left and middle image are rendered with a small area light source and demonstrate the fine geometric details captured in the BTF. The right image shows a relighting of the echinite with a complex image-based lighting environment captured in front of the Pädagogische Fakultät in Bonn. The ground floor in the right image is covered with a synthetic leather BTF courtesy of DaimlerChrysler AG.*