

Rendering Deformable Surface Reflectance Fields

Tim Weyrich, Hanspeter Pfister, *Member, IEEE*, Markus Gross, *Member, IEEE*

Abstract—Animation of photorealistic computer graphics models is an important goal for many applications. Image-based modeling has emerged as a promising approach to capture and visualize real-world objects. Animating image-based models, however, is still a largely unsolved problem. In this paper, we extend a popular image-based representation called surface reflectance field to animate and render deformable real-world objects under arbitrary illumination. Deforming the surface reflectance field is achieved by modifying the underlying impostor geometry. We augment the impostor by a local parameterization that allows the correct evaluation of acquired reflectance images, preserving the original light model on the deformed surface. We present a deferred shading scheme to handle the increased amount of data involved in shading the deformable surface reflectance field. We show animations of various objects that were acquired with 3D photography.

Categories and Subject Descriptors (according to ACM CCS): I.3.3.e [Computer Graphics]: Image-based rendering

I. INTRODUCTION

Visualization and animation of realistic 3D computer graphics models are important for many applications, such as computer games, movies, advertisement, virtual environments, or e-Commerce. Broadly speaking, there are three approaches to reproduce the visual appearance of real objects: Explicit modeling with parametric representations, pure image-based approaches, and hybrid approaches that use a combination of both. In all cases, the reflectance properties of 3D objects are typically captured with methods of 3D photography.

For explicit appearance models, parametric BRDF models are fit to the acquired data. Parametric BRDFs can be efficiently rendered on modern graphics hardware, and the underlying geometry can be animated and deformed using well-developed techniques such as skinning and vertex blending. However, parametric BRDFs can not capture many effects of real-world materials, such as translucency, inter-reflections, self-shadowing, and subsurface scattering. Pure image-based techniques,

on the other hand, are well suited to acquire and represent complex object appearance. However, most of them impose restrictions on the viewpoints, and the lack of a 3D geometry model makes deformations very difficult or impossible.

Consequently, hybrid approaches have become very popular. They parameterize an image-based model on an impostor geometry that can be used for animation.

This representation is commonly known as a surface light field, introduced by Miller et al. [23]. Wood et al. [31] presented a simple technique for its animation. However, surface light fields can only capture an object under fixed illumination. This is a severe limitation if the object is rendered in a new environment or under dynamically changing lights. A more general hybrid representation is the *surface reflectance field*, which captures the object appearance for many possible light configurations. Objects with arbitrary reflectance properties can be rendered from any viewpoint under new illumination. To date, however, there has been no publication on the animation of surface reflectance fields.

In general, animating an image-based or hybrid appearance representation requires a scheme to evaluate the image-based data set to simulate varying object deformations. A key question is how to preserve the visual appearance of the deformed object surface, that is, how to preserve the perceptual impression of material properties under different lighting conditions.

In this paper we present a method for animating surface reflectance fields with arbitrary geometric deformations. We develop a shading scheme that aims to preserve the appearance of object materials during deformation. Our method uses a local parameterization of the impostor geometry that enables arbitrary warps. Our shading method uses this local parameterization to approximately preserve the spatially varying BRDFs of the undeformed object. We present a cache-optimized shading strategy to minimize computation time. Our technique is applicable to other hybrid representations that contain an impostor geometry.

II. PREVIOUS WORK

3D photography loosely describes methods that capture object shape and appearance from images. Some 3D photography approaches fit a parametric BRDF

T. Weyrich and M. Gross are with the Computer Graphics Laboratory, ETH Zentrum, 8092 Zurich, Switzerland. Email: {weyrich, gross}@inf.ethz.ch

H. Pfister is with MERL—Mitsubishi Electric Research Laboratories, Cambridge, MA 02139. Email: pfister@merl.com

model to the acquired appearance data [28], [32], [14], [13], [22], [16]. However, parametric BRDFs can not represent many of the reflectance properties of real-life objects [11]. In contrast, image-based appearance representations (with geometry) make no assumptions about the reflection property of materials.

The lumigraph [10] is the first method to combine an impostor geometry (a visual hull of the object) with an image-based representation. Images of the object are stored as a dense light field [17], and the impostor geometry is used to improve light field interpolation, i.e., to minimize blooming and ghosting. To reduce the amount of image data, view-dependent texture mapping [27], [8], [7], uses simple geometry and sparse texture data. This method is extremely effective despite the approximate 3D shape, but it has some limitations for highly specular surfaces due to the relatively small number of textures.

Surface light fields [23], [31], [3] are a more general and efficient representation because they parameterize the image data onto the object surface. Surface light fields can either be stored on accurate high-density geometry [31] or on coarse triangular meshes for objects with low geometric complexity [24]. Some techniques [25], [3] aggressively compress the data such that the models can be rendered in real-time on modern graphics hardware. To improve the appearance of complex object silhouettes, surface light fields can be combined with view-dependent opacity data into *opacity light fields* [29]. Unstructured lumigraph rendering [2] is a very effective method for rendering both surface and opacity light fields.

Although surface light fields are capable of reproducing important global effects such as inter-reflections and self-shadowing, they only show the object under fixed lighting. To overcome this limitation, recent approaches have used surface reflectance fields. The reflectance field of an object is the radiant light from a surface under every possible incident field of illumination. In practice, the reflectance field is sampled sparsely and interpolated during rendering. To further reduce the amount of data, most reflectance fields are acquired for a single view [5], [11], [15]. For approximately planar geometry and diffuse surfaces the data can be compressed further by fitting a parametric function to the reflectance field [18].

In our work we use the surface reflectance field data acquired by the 3D photography system of Matusik et al. [20]. The system acquires reflectance fields for over 400 views using cameras, turntables, and a rotating array of lights. The acquired data also includes view-dependent opacity information, called the *opacity reflectance field* (not considered in this paper). The impostor geometry is the visual hull of the objects.

Although it is an important aspect for many practical applications, the animation of image-based data has received very little attention in the literature. Wood et al. [31] describe arbitrary deformations on a surface light field and produce plausible renderings of the deformed model. However, their method does not deal properly with the diffuse component of the surface color, and it only works for purely reflective isotropic BRDFs. Feature-based light field morphing [33] morphs two light fields into each other, based on the concept of ray-correspondencies. The method requires substantial user input to specify corresponding feature polygons between the two objects, and it is not applicable to the general animation setting. Both methods work only for static illumination. Furukawa et al. [9] presented a scanning system to capture objects and spatially varying BRDFs, also called Bidirectional Texture Functions (BTFs) [4]. They use tensor product expansion to compress the BTF data and show results with surface deformations. Their system is the first to render deformations of a relightable, image-based object representation. However, they rely on a tight impostor geometry to support the BTF representation. Thus, the object geometry has to be acquired with a range scanning device. Moreover, their paper does not explicitly address appearance preservation under non-uniform, skewed deformations.

In this paper we describe an animation method for surface reflectance fields mapped onto *approximate geometry*. Our method allows to place the objects in new environments with arbitrary illumination, including dynamically changing lights. We carefully analyze the conditions that have to be met to preserve the appearance of the object (Section IV-A), and we present a novel method to approximately preserve spatially varying BRDFs during deformations (Section IV-B). We discuss the limitations of our approach (Section IV-D) and show results using objects that are difficult to handle for image-based approaches, including objects with specularities, transparency, and self-shadowing (Section VI).

III. OVERVIEW

A surface reflectance field consists of a large set of image-based reflectance data of an object, in conjunction with an impostor geometry used for rendering. The reflectance data is given by a collection of *reflectance images*. A reflectance image is a stack of camera images showing the object from the same viewpoint under varying directional illumination (Fig. 1). A surface reflectance field consists of reflectance images from many viewpoints around the object. For our objects we use approximately 400 reflectance images, each with 60




Fig. 1. Example reflectance image.

high dynamic range images of the object under different directional illumination.

Rendering a surface reflectance field can be understood as a two step procedure. First the reflectance images are used to compute an image of the object under the new illumination for each viewpoint (Fig. 2). These

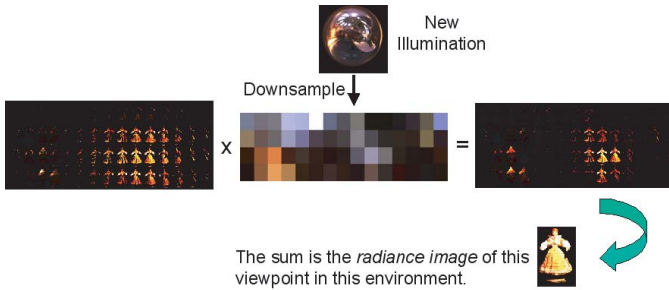


Fig. 2. Relighting using a reflectance image. The environment map of the new illumination is downsampled to the resolution of the reflectance image. Images are multiplied with the corresponding environment map color and added up to yield an image of the object under new illumination.

images are then rendered together using the impostor geometry with unstructured lumigraph interpolation [2]. Our impostor geometry is the visual hull of the object that can easily be determined from observed silhouette images. However, our method is independent of the choice of impostor geometry.

We animate and deform a surface reflectance field (SRF) by first applying a 3D warp to the impostor geometry. In order to preserve the warped object’s appearance, it becomes necessary to blend the reflectance images individually for each point of the impostor geometry. We developed a new look-up function to evaluate the reflectance images of the warped SRF. This function depends on the warp and is used to shade the warped impostor geometry during image generation (Fig. 3). Shading a point on the impostor surface requires to apply this look-up function to the surface point, the viewing ray, and the incident light direction. Once warped back into acquisition frame, the point can be shaded by blending the reflectance images according to the mapped light direction and viewing ray.




Fig. 3. Rendering a deformed surface reflectance field requires switching between reference frames. The acquired impostor geometry is warped into object space. During the shading process, queries to the reflectance images must be transformed back into the acquisition frame using a local parameterization of the impostor geometry.

In the remainder of this paper we develop this look-up scheme. We discuss some shading issues and present our implementation based on a point-based impostor geometry. Section VI shows some examples of deformed and animated surface reflectance fields.

IV. DEFORMABLE SURFACE REFLECTANCE FIELDS

Neglecting global light transport, a surface reflectance field can be understood as a discrete sampling of the object’s BRDFs, knowing neither the exact location of the surface, nor its normals. Ignoring wavelength and time, a BRDF is a scalar function $BRDF(\mathbf{l}, \mathbf{v})$, describing the fraction of light that is reflected in direction \mathbf{v} as the surface is lit from direction \mathbf{l} . Similar to the BRDF notation, we use $SRF(\mathbf{p}, \mathbf{l}, \mathbf{v})$ to denote the reflectance of the SRF for an impostor point \mathbf{p} , a light direction \mathbf{l} , and a viewing direction \mathbf{v} . This relationship is fundamental for our analysis of appearance preservation during surface reflectance field animation.

The initial step in animating an SRF is to deform the impostor geometry. In this section, we assume the impostor warp is defined by a differentiable warp function

$$\Psi: \mathbb{R}^3 \rightarrow \mathbb{R}^3. \tag{1}$$

Shading an object includes queries to the surface reflectance field to determine the object’s local reflectance. If the object is warped, the shading operation has to map a query $(\mathbf{p}^*, \mathbf{l}^*, \mathbf{v}^*)$ in object space to a query $(\mathbf{p}, \mathbf{l}, \mathbf{v})$ in the original acquisition space (Fig. 4). By intuition, the required mapping approximately follows the inverse warp Ψ^{-1} .




Fig. 4. During rendering, reflectance queries must be mapped into acquisition space. Given lighting and viewing direction \mathbf{l}^* and \mathbf{v}^* for an impostor point \mathbf{p}^* , an appropriate surface reflectance field query $SRF(\mathbf{p}, \mathbf{l}, \mathbf{v})$ with \mathbf{p} , \mathbf{l} , and \mathbf{v} in acquisition space has to be found. The look-up function $\mathcal{L}: (\mathbf{p}^*, \mathbf{l}^*, \mathbf{v}^*) \mapsto (\mathbf{p}, \mathbf{l}, \mathbf{v})$ should be appearance preserving.

However, applying Ψ^{-1} to the lighting and viewing directions is not necessarily appearance preserving, especially in the case of non-uniform deformations. Section IV-A develops a mapping \mathcal{L} to perform this operation, leading to a new impostor parameterization presented in IV-B. The final shading process is discussed in section IV-C.

A. Approximate BRDF Preservation

In general, a mapping from object space to SRF acquisition space that preserves all aspects of the object's appearance can not exist. This is due to the lack of exact object geometry and material properties in the SRF representation. Both would be needed to allow a prediction of complex non-local effects such as self-shadowing and inter-reflections.

Thus, we make the simplifying assumption, that the observed object can be described completely by its local BRDFs. We develop a look-up function that tries to preserve the main characteristics of the original object BRDFs. We do not aim at modeling changes in the BRDF due to the deformation of the material's microstructure. This would require precise knowledge of the structure, e.g. its microfacet distribution. Instead, we want to map the original BRDF to the deformed surface.

The desired look-up function is a mapping

$$\mathcal{L}: (\mathbf{p}^*, \mathbf{l}^*, \mathbf{v}^*) \mapsto (\mathbf{p}, \mathbf{l}, \mathbf{v}). \quad (2)$$

We enforce appearance preservation by imposing three conditions:

- 1) Suppose the viewing ray in object space $\mathbf{p}^* + s\mathbf{v}^*$, $s \in \mathbb{R}$ intersects the warped object geometry at a point \mathbf{q}^* (see Fig. 5). Then the viewing ray in acquisition space $\mathbf{p} + t\mathbf{v}$, $t \in \mathbb{R}$ should intersect the original object at the corresponding point $\mathbf{q} = \Psi^{-1}(\mathbf{q}^*)$ to ensure the reflected light




Fig. 5. A viewing ray $\mathbf{p}^* + s\mathbf{v}^*$, $s \in \mathbb{R}$, in object space intersects the warped object at \mathbf{q}^* . The corresponding viewing ray in the surface reflectance field's acquisition frame should intersect the object at the original object point \mathbf{q} . However, the impostor geometry of the SRF is typically not the same as the actual object geometry, so \mathbf{q}^* and \mathbf{q} remain unknown. Instead, we are using $\mathbf{p} = \Psi^{-1}(\mathbf{p}^*)$ as the origin of the viewing ray $\mathbf{p} + t\mathbf{v}$, $t \in \mathbb{R}$ in acquisition space.

to originate from the same surface point with the same BRDF. Unfortunately, this condition can not be guaranteed, since the true object geometry is typically not known, i.e., the exact location of the points \mathbf{q} and \mathbf{q}^* can not be determined. However, it is a reasonable approximation to force the two viewing rays to intersect corresponding points of the impostor geometry, since the impostor point \mathbf{p} is very likely close to \mathbf{q} . This can be achieved by choosing

$$\mathbf{p} := \Psi^{-1}(\mathbf{p}^*). \quad (3)$$

- 2) \mathbf{l} and the surface normal \mathbf{m} in \mathbf{q} have to enclose the same angle as \mathbf{l}^* and the normal \mathbf{m}^* in \mathbf{q}^* . This is a necessary condition to retain the reflectance characteristics of the object, as for example, the shape of the reflectance lobes of the corresponding BRDF. The same condition applies for \mathbf{v} and \mathbf{v}^* . Preserving the shape of the reflectance lobes requires to preserve the angle between \mathbf{l}^* and \mathbf{v}^* as well (see Fig. 6).
- 3) In order to preserve the effect of anisotropic object BRDFs, \mathbf{l} and \mathbf{v} should have the same azimuthal orientation relative to the object surface as \mathbf{l}^* and \mathbf{v}^* on the warped object. (See Fig. 6.)

As \mathbf{p} immediately follows from (3), the mapping $(\mathbf{p}^*, \mathbf{l}^*, \mathbf{v}^*) \mapsto (\mathbf{l}, \mathbf{v})$ remains to be found. This can be rewritten as a locally affine mapping

$$\mathcal{L}_{\mathbf{p}^*}: (\mathbf{l}^*, \mathbf{v}^*) \mapsto (\mathbf{l}, \mathbf{v}) \quad (4)$$

of lighting and viewing directions in the vicinity of \mathbf{p}^* . Note, that $\mathcal{L}_{\mathbf{p}^*}$ is a function of \mathbf{p}^* .

According to the second condition, $\mathcal{L}_{\mathbf{p}^*}$ has to be angle preserving. By convention, \mathbf{l} , \mathbf{v} , \mathbf{l}^* , and \mathbf{v}^* are unit vectors, so, $\mathcal{L}_{\mathbf{p}^*}$ needs to be length preserving as well. This implies that $\mathcal{L}_{\mathbf{p}^*}$ is an isometry, which means that the effect of \mathcal{L} on \mathbf{l}^* and \mathbf{v}^* can be described by a rotation

or a reflection, respectively, as a function of location \mathbf{p}^* . In the remainder of this paper, we assume $\mathcal{L}_{\mathbf{p}^*}$ to be a rotation. The special case of a reflection can be handled similarly and is left out for simplicity.

Using this observation and interpreting (3) as a translation of \mathbf{p}^* by $\Psi^{-1}(\mathbf{p}^*) - \mathbf{p}^*$, the total effect of \mathcal{L} on a local setting around \mathbf{p}^* can be expressed as a *rigid transformation*, translating the point \mathbf{p}^* onto \mathbf{p} while rotating lighting and viewing directions.

Condition 2 further restricts \mathcal{L} to rigid transformations that map the warped object normal at \mathbf{q}^* onto the normal at the original point \mathbf{q} . We do not know the exact object geometry. Instead, we refer to the impostor normal in \mathbf{p} as an approximation of the true surface normal. Consequently, the searched \mathcal{L} maps the tangential plane of the warped impostor point \mathbf{p}^* onto the tangential plane in \mathbf{p} .

According to condition 3, \mathcal{L} should be chosen to preserve \mathbf{l}^* 's and \mathbf{v}^* 's orientation inside the tangential plane. The following section presents a local impostor parameterization that enables one to track the transformation of the local tangential frame in order to find a mapping \mathcal{L} that fulfills the conditions.

B. Local Impostor Parameterization

We augment the impostor geometry with a local parameterization that allows us to determine the look-up function \mathcal{L} at each impostor point \mathbf{p}^* . The parameterization is independent from the geometric representation and can be applied to triangular meshes as well as to point sampled geometry. Depending on the representation, \mathbf{p}^* can be a mesh vertex or a non-connected surface point, respectively. This section develops the




Fig. 6. The bi-directional reflection distribution function (BRDF) describes the fraction of light from \mathbf{l} that is reflected towards \mathbf{v} . The spatial characteristics of a BRDF can be preserved by preserving the angles ϑ_l , ϑ_v , and γ when changing to acquisition space. Anisotropic materials require the preservation of the azimuthal orientations φ_l and φ_v relative to the surface frame. In an SRF, \mathbf{m} is not exactly known. Instead, the impostor normal \mathbf{n} is taken as an approximation.

Parameter	Description
\mathbf{p}_0	Original position of the impostor point \mathbf{p}^* in acquisition frame.
\mathbf{r}_0	Orientation R_0 of the tangential system of \mathbf{p}^* in the acquisition frame, stored as a Rodrigues vector.
(\mathbf{u}, \mathbf{v})	Local tangential frame of \mathbf{p}^* .

TABLE I

THE LOCAL IMPOSTOR PARAMETERIZATION ALLOWS FOR THE DETERMINATION OF THE LOOK-UP FUNCTION \mathcal{L} DURING RENDERING. THE PARAMETERIZATION CONSISTS OF A SET OF VECTORS $(\mathbf{p}_0, \mathbf{r}_0, \mathbf{u}, \mathbf{v})$ IN \mathbb{R}^3 , STORED AT EACH POINT OF THE DISCRETIZED IMPOSTOR GEOMETRY. \mathbf{p}_0 AND \mathbf{r}_0 ARE FIXED, WHILE \mathbf{u} AND \mathbf{v} HAVE TO BE ADAPTED WHENEVER THE GEOMETRY IS DEFORMED.

set of parameters that are stored at every point \mathbf{p}^* (see Table I).

Given an arbitrary warp $\Psi: \mathbb{R}^3 \rightarrow \mathbb{R}^3$, finding an explicit inverse function Ψ^{-1} is impossible in general. Constraining the SRF deformation to cases where an explicit inverse warp function exists would be too limiting. Consequently, another design goal for the parameterization was to determine \mathcal{L} without using a closed form of Ψ^{-1} .

As shown above, \mathcal{L} can be decomposed into a translational part and a rotation. The translation moves the warped impostor point back to its original position (see (3)). Storing the original position \mathbf{p}_0 in each impostor point \mathbf{p}^* allows the application of the translation without using Ψ^{-1} .

The rotational part $\mathcal{L}_{\mathbf{p}^*}$ (see (4)) aligns the warped tangential frame of the impostor point with the corresponding tangential system in the acquisition frame (see Fig. 7). To avoid the application of Ψ^{-1} , the tangential orientation R_0 in the acquisition frame must explicitly be stored in \mathbf{p}^* . R_0 is defined by the rotation matrix $(\mathbf{u}_0, \mathbf{v}_0, \mathbf{n}_0)$, built by the tangential system in acquisition space. \mathbf{u}_0 and \mathbf{v}_0 must be orthogonal and can be arbitrarily chosen. To minimize its memory footprint, R_0 can be stored as a Rodrigues vector [1]

$$\mathbf{r}_0 = \mathbf{k} \tan \frac{\vartheta}{2}, \quad (5)$$

given R_0 as a rotation of angle ϑ around an axis \mathbf{k} . This definition provides a minimal (i.e., three dimensional) parameterization of R_0 . It contains a singularity for 180° rotations that can be avoided in our context as R_0 can be arbitrarily chosen.

Additionally, a tangential coordinate system (\mathbf{u}, \mathbf{v}) is attached to \mathbf{p}^* . While \mathbf{p}_0 and R_0 remain unchanged during the deformation, (\mathbf{u}, \mathbf{v}) is subject to the same warp as the geometry: After a deformation by Ψ , (\mathbf{u}, \mathbf{v})




Fig. 7. For each point on the impostor geometry, the local impostor parameterization provides two coordinate systems: The orthogonal tangential system R_0 in \mathbf{p}_0 in the acquisition frame and its warped counter-part on the rendered impostor geometry in object space, defined by \mathbf{p}^* and (\mathbf{u}, \mathbf{v}) . During rendering, \mathbf{p}_0 , R_0 , and the orthogonalized tangential system $(\bar{\mathbf{u}}, \bar{\mathbf{v}})$ are used to determine the appearance preserving back-projection of the SRF query.

is set to $\Psi((\mathbf{u}_0, \mathbf{v}_0))$. Applying the warp function Ψ to a local tangential system is a standard problem in differential geometry. The tangential system has to be mapped by taking directional derivatives of Ψ . In our implementation, we use central differences to determine \mathbf{u} and \mathbf{v} :

$$\begin{aligned} \mathbf{u} &= \frac{1}{2\varepsilon}(\Psi(\mathbf{p}_0 + \varepsilon\mathbf{u}_0) - \Psi(\mathbf{p}_0 - \varepsilon\mathbf{u}_0)), \\ \mathbf{v} &= \frac{1}{2\varepsilon}(\Psi(\mathbf{p}_0 + \varepsilon\mathbf{v}_0) - \Psi(\mathbf{p}_0 - \varepsilon\mathbf{v}_0)) \end{aligned} \quad (6)$$

for a small $\varepsilon > 0$. In particular, this only restricts Ψ to be differentiable in a vicinity of the impostor geometry.

Knowing R_0 and (\mathbf{u}, \mathbf{v}) , the rotation $\mathcal{L}_{\mathbf{p}^*}$ can easily be reconstructed during rendering. Let

$$\bar{\mathbf{n}} = \frac{\mathbf{u} \times \mathbf{v}}{\|\mathbf{u} \times \mathbf{v}\|} \quad (7)$$

and

$$\bar{\mathbf{b}} = \frac{\mathbf{b}}{\|\mathbf{b}\|} \quad \text{with} \quad \mathbf{b} = \frac{\mathbf{u}}{\|\mathbf{u}\|} + \frac{\mathbf{v}}{\|\mathbf{v}\|} \quad (8)$$

be the normal and the normalized bisecting vector of the warped tangential system spanned by \mathbf{u} and \mathbf{v} , respectively (see Fig. 7). Then,

$$\begin{aligned} \bar{\mathbf{u}} &= \frac{1}{\sqrt{2}}(\bar{\mathbf{b}} + \bar{\mathbf{b}} \times \bar{\mathbf{n}}) \quad \text{and} \\ \bar{\mathbf{v}} &= \frac{1}{\sqrt{2}}(\bar{\mathbf{b}} + \bar{\mathbf{n}} \times \bar{\mathbf{b}}) \end{aligned} \quad (9)$$

build an orthogonalized tangential system that minimizes the squared angular differences between corresponding basis vectors of (\mathbf{u}, \mathbf{v}) and $(\bar{\mathbf{u}}, \bar{\mathbf{v}})$.

Using this orthogonalized system, the rotation $\mathcal{L}_{\mathbf{p}^*}$ is given by

$$\mathcal{L}_{\mathbf{p}^*}(\mathbf{x}) = R_0(\bar{\mathbf{u}}, \bar{\mathbf{v}}, \bar{\mathbf{n}})^\top \mathbf{x}. \quad (10)$$

Note that by choosing $(\bar{\mathbf{u}}, \bar{\mathbf{v}})$ as proposed, $\mathcal{L}_{\mathbf{p}^*}$ is an approximation of tangential orientation preservation in the sense of condition 3 of the previous section.

This approximation may introduce a slight rotation of the principal axis of an anisotropic BRDF when the

surface is sheared. The rotation may be counter-intuitive depending on the direction of the shear. In general, the effect of a shear on a real-world material effectively changes the original BRDF, depending on its microstructure. As the microstructure is not known, we decided to use the original BRDF. Thus, a tangential rotation of certain characteristics of the anisotropic BRDF can not be avoided.

It may not be surprising that \mathcal{L} turned out to be a rigid transformation. For rigid object transformations and cases where impostor geometry and object surface coincide, rotating the reflectance data according to the inverse object transformation is the appropriate choice. Consequently, this technique is used in many applications, e.g., for image-based BRDF measurements [30], [19], [16], for bump mapping, for BTFs [4], [9] and for various other texturing techniques. All these techniques assume relatively accurate impostor geometry.

We assume arbitrary, in particular non-uniform, skewed deformations of an approximate impostor geometry that is different from the real object surface. The central result of our analysis is that in these cases again a rigid transformation of the impostor surface frame meets the requirements of appearance preservation best. Our framework allows the derivation of \mathcal{L} for arbitrary deformations, using the local impostor parameterization. As the resulting transformation is rigid, material properties—including anisotropic BRDFs—stay the same, even for skewed object transformations.

C. Shading

Once \mathcal{L} was determined, an impostor point can be shaded as described in Section III. However, there are some issues that should be considered when lighting a deformed surface reflectance field.

When using an environment map to light an SRF, it needs to be filtered according to the spatial resolution of the reflectance images representing the SRF (see Fig. 1). Using an unfiltered environment map may lead to aliasing artifacts, if the map contains details that are finer than the spacing of the light sources used to acquire the SRF.

During rendering of a deformed SRF, $\mathcal{L}_{\mathbf{p}^*}$ is applied to all lighting directions incident to \mathbf{p}^* . This corresponds to a rotation of the environment map before evaluating the SRF with that environment. Thus, the rotated environment map needs to be re-filtered according to the reflectance field sampling. This operation would have to be performed for every surface point, which is impractical as filtering is an expensive operation for a non-uniform reflectance field sampling.

We present an alternative way to light a deformed surface reflectance field that comes without the need of re-filtering. Although an SRF is acquired using directional lighting, we can simulate lighting by *point light sources* with only little artifacts.

Lighting an impostor point \mathbf{p}^* by a point light source starts with the reflectance query $(\mathbf{p}^*, \mathbf{l}^*, \mathbf{v}^*)$ describing the viewing ray to \mathbf{p}^* , and the direction \mathbf{l}^* to the point light-source as seen from \mathbf{p}^* . Applying the look-up scheme $(\mathbf{p}, \mathbf{l}, \mathbf{v}) = \mathcal{L}(\mathbf{p}^*, \mathbf{l}^*, \mathbf{v}^*)$ transforms the query into acquisition space. $\text{SRF}(\mathbf{p}, \mathbf{l}, \mathbf{v})$ yields a reflectance coefficient that can be used to shade the impostor point \mathbf{p} with a color I_p :

$$I_p = \text{SRF}(\mathbf{p}, \mathbf{l}, \mathbf{v})A(\|\mathbf{p} - \mathbf{p}_l\|)I_l, \quad (11)$$

where I_l is the color of the point light source at \mathbf{p}_l , and $A(d)$ is the light attenuation factor depending on the distance to the light source. Using the impostor point \mathbf{p}^* instead of a point on the real object surface introduces an error in the incident lighting direction. But usually this error is small compared to the resolution of the SRF.

In contrast to environment mapping, this technique introduces SRF queries for lighting directions that are not present in the reflectance images. We rely on the SRF implementation to properly interpolate novel lighting directions, as discussed in Section V-B.

The proposed approximation of point light-sources allows for dynamic lighting effects when animating surface reflectance fields. However, as our look-up scheme is ignoring global illumination effects, these effects may appear wrong on deformed objects.

Point light sources can also be used to implement environment mapping without the need to re-filter the map for every impostor point. This can be done by properly sub-sampling the environment once for a dense set of directions. Then point light sources are defined at these directions, colored by the corresponding values of the environment. Provided the sub-sampling is dense enough, lighting the scene with these light sources leads to an appropriate reconstruction of the lighting environment. Note that light direction interpolation implicitly acts as a reconstruction filter. Thus, adaptively filtering the environment map is traded for interpolation. This alternative scheme is easier to implement and fits naturally into the framework of point light sources. For more flexible lighting effects, environment maps and point light sources can be combined.

D. Limitations

We presented a simple technique to preserve the appearance of a surface reflectance field under arbitrary

deformations. However, the look-up scheme contains some approximations affecting the proper reproduction of lighting effects. There are three classes of errors that may occur.

First, our BRDF-based derivation of the look-up function is ignoring non-local effects like inter-object reflection, self-shadowing, refraction, or sub-surface scattering. This may lead to false shadowing and erroneous refractions. However, modeling these effects requires incorporating the exact object geometry. As an SRF contains only approximate geometry information, these global lighting effects can not completely be preserved during the deformation.

The BRDF preserving approach itself contains some approximations. As \mathcal{L}_{p^*} is derived from the local warp around the observed impostor point instead of a point on the real object surface, the warped SRF may show some BRDFs that seem to be rotated relative to the surface frame. The effect grows as the local warp's rotational part around \mathbf{p}^* is different from the rotation of the observed, real surface point \mathbf{q}^* . Note that the error introduced by choosing the impostor normal to determine \mathcal{L}_{p^*} is comparatively low. In fact, the normal is only used to derive the rotational approximation of the local warp. Evaluating the SRF does not incorporate a normal anymore. In particular, this normal has no impact on the reflectance properties of the reproduced BRDFs.

A third error class affects object/impostor parallax in regions where the impostor geometry is distant to the object surface: For non-uniform stretches, the object texture may appear to be shifted. This happens if the corresponding viewing rays in object and acquisition space do not intersect the object at corresponding points \mathbf{q} and \mathbf{q}^* . Fig. 5 shows such an example.

Despite these limitations, the proposed scheme shows significant advantages. Its ability to preserve reflectance properties without explicit knowledge of the real surface normals is crucial for the deformation of surface reflectance fields. Moreover, it makes it suited for many other image-based applications where the exact object normals are not known.

An important property of this look-up scheme is that in the limit the BRDF preservation is exact, as the SRF geometry converges to the real object geometry. The proposed method directly benefits from improvements of the geometric representation, for example, for geometry acquired with laser range scanning.

V. IMPLEMENTATION

We implemented an animation system for surface reflectance fields similar to the system presented by Ma-




Fig. 8. Close-ups of a surfel geometry. The surfels are placed on a square grid that was anisotropically stretched. (a) and (b) depict surface splats with reduced support to show the surfel distribution. (a) Circular splats can not follow the deformation. Thus, holes in the surface may appear. (b) Adaptively deformed elliptical splats preserve a closed surface and provide anisotropic texture filtering. (c) Blending the elliptical splats at their original size produces a smooth texturing.

tusik et al. [20]. We improved the impostor representation, its rendering, and the reflectance field interpolation. Finally, we present a new shading scheme that speeds-up the rendering of warped SRFs.

A. Surfel Representation

Previous work on surface reflectance fields [20] used a point-based representation of the impostor geometry. In this approach the impostor was densely sampled by *surfels*. Surfels, as introduced by [26], are points in \mathbb{R}^3 , augmented by additional attributes, such as normal, radius, and some color properties. In their representation, the normal indicates the surface orientation, whereas the radius can be understood as an approximate, circular region of influence.

Subsequently, [35], [36] presented a framework to render surfels based on Heckbert’s elliptical weighted average (EWA) texture filtering, [12]. EWA surface splatting is a forward-mapping algorithm that allows for high quality, aliasing-free rendering of point sampled geometry.

Although featuring high image quality, the surfel renderer as used by [20] is not suited for our purposes. It organizes the surfels in a static layered depth cube (LDC) tree, a fixed, axis aligned spatial data structure, that can not be warped. Instead, we use a variant of the Pointshop3D renderer [34] which is an advanced version of EWA surface splatting. The Pointshop3D renderer is able to render unorganized point-sets and thus can cope with dynamically changing point sets.

Deformations of the impostor geometry may lead to visible holes in the point-based reconstruction. For reasonable deformations, this can be avoided by deforming the splat geometry appropriately. As proposed by [36], we are using elliptical surface splats. Fig. 8 shows the effect on a deformed patch of a surfel geometry. In contrast to circular splats that use a position/normal parametrization, our splats are defined by a position \mathbf{p}




Fig. 9. A surfel is defined by its position \mathbf{p} and its tangential system (\mathbf{u}, \mathbf{v}) . \mathbf{u} and \mathbf{v} span a skewed coordinate system that defines the surfel’s elliptical region of influence.

and two tangential vectors \mathbf{u} and \mathbf{v} , spanning a tangential coordinate system. \mathbf{u} and \mathbf{v} are explicitly allowed to be non-orthonormal. The surfel is associated with an elliptical region of influence defined by the tangential system (see Fig. 9).

When warping the impostor geometry, the warp has to be applied on each surfel’s tangential system by using an affine approximation of the local distortion. In our implementation, this approximation is determined similarly to (\mathbf{u}, \mathbf{v}) in section IV-B. This ensures surface coverage as long as the warp’s first derivative does not change excessively (see Fig. 8).

The surfel’s tangential vectors \mathbf{u} and \mathbf{v} provide us directly with the corresponding parameters of our local impostor parameterization. By adding the surfel’s original position \mathbf{p}_0 and orientation \mathbf{r}_0 to the set of surfel attributes, all impostor parameters required for the warped SRF rendering are encoded in the surfels.

B. Reflectance Image Interpolation

In general, each surface reflectance field look-up $\text{SRF}(\mathbf{p}, \mathbf{l}, \mathbf{v})$, must be interpolated from the fixed set of reflectance images. Depending on the viewing ray $\mathbf{p} + s\mathbf{v}$, $s \in \mathbb{R}$, the query must be handled by interpolating between the observing camera views of the SRF. View interpolation is a common problem in image-based rendering. Like Matusik et al. [20], we use unstructured lumigraph interpolation that was introduced with unstructured lumigraph rendering (ULR) [2].

ULR interpolation has proven to be a very flexible interpolation scheme that can be controlled by an arbitrary penalty function, rating each camera’s appropriateness. In our implementation, we chose the same penalty as in [20].

In previous work, the queried lighting direction \mathbf{l} was always chosen from the set of acquired illuminations, sampling a stationary environment map. However, since our point light-source implementation produces arbitrary lighting directions, we have to interpolate reflectance queries between the acquired lighting directions as well.

This problem has been addressed by Koudelka et al. [15] who intersected the queried light ray with a triangulated hull of the acquisition lightsource positions. The intersected triangle yielded the three closest light sources that were consequently used to interpolate between three different lighting conditions. However, this does not work for arbitrary positions of the acquisition light sources. We also found that for slim triangles in the hull light source interpolation was not smooth enough.

Instead, we are also using ULR to interpolate between light directions. Treating the light sources of the acquisition stage as cameras at infinite distance enables us to use ULR for light source interpolation. In our implementation, this is done by using the same penalty function as for camera interpolation.

Using ULR for camera selection affords consideration of occlusions in the scene. If the observed object point is not visible from a given camera, this camera should not be considered for interpolation. When interpolating between light sources, there are no visibility constraints, since shadowing effects are already captured in the acquired reflectance images. Moreover, taking occlusion into account would change the appearance of translucent or refractive materials.

C. Reordered Evaluation

Compared to the shading of static objects, substantially more reflectance images are involved during shading of a deformed model. This is because the rendered viewing rays tend to diverge in acquisition space when the impostor is warped (see Fig. 10).

The number of reflectance images involved in the shading process increases with deformations. Since SRF data usually exceeds the conventional main memory size, most SRF look-ups must be read from disk. The time needed for disk access exceeds all other computation times by orders of magnitude. This makes SRF look-up performance a critical issue, especially when producing animated SRFs.




Fig. 10. Viewing ray divergence due to back-projection.

Entry	Description
REF(S)	Reference to the target surfel.
(i_c, i_l)	Camera index and light source index for the reflectance image look-up.
(x_c, y_c)	Look-up position on the camera plane.
$(\omega_r, \omega_g, \omega_b)$	Color weights.

TABLE II

A shading operation REPRESENTS A BASIC REFLECTANCE FIELD LOOK-UP. IT CONTAINS THE DEPICTED ENTRIES THAT ARE USED TO ACCUMULATE REFLECTANCE VALUES IN THE SURFEL COLORS.

A SHADING OPERATION IS EXECUTED BY LOOKING-UP THE REFLECTANCE IMAGE (i_c, i_l) AT POSITION (x_c, y_c) . THE RESULTING RGB REFLECTANCE VALUE IS WEIGHTED BY $(\omega_r, \omega_g, \omega_b)$ AND ADDED TO THE COLOR OF SURFEL S .

We address this problem with two techniques: A smaller cache holds the most recently used blocks of data. Unfortunately, naively shading the geometry shows adverse cache-coherency. Thus, as a second technique, all shading operations are decomposed into minimal shading operations, each containing a single SRF look-up. Table II shows their detailed structure. The execution of the shading operation is deferred until a larger number of operations has been collected. Sorting them before execution in a cache-optimal way decreases shading times by an order of magnitude. Moreover, due to the the cache coherent shading execution, the cache can be kept very small. In our system, it was sufficient to dimension the cache large enough to hold three reflectance images at the time. To facilitate the deferred shading, we accumulate all surfel colors first before rendering them to frame buffer.

VI. RESULTS

For our experiments we use the two models shown in Fig. 11 that were scanned by Matusik et al. [20], [21]. The doll SRF is well suited for the analysis of appearance preservation, as it contains many different materials. The mug SRF was chosen to show the limitations of our approach, as it contains refractive effects that can not be handled correctly by our technique.

The reflectance images of the models were captured using the high dynamic range (HDR) technique of Debevec et al. [6]. For each viewpoint and lighting condition, four 1360×1032 pictures with different exposure time were taken. Using up to 432 different viewpoints and 60 light directions, this produced a total of $432 \times 60 \times 4 = 103680$ images. Uncompressed this would correspond to 407 GB of data. All reflectance images were compressed using a PCA compression that

achieved a reduction of at least a factor of 10. Fig. 11 lists the compressed size of both models.

With these optimizations, the vast part of the rendering time is spent for the ULR computation. Our current implementation assumes arbitrarily placed cameras and light sources during the acquisition. Using knowledge about the scanner geometry or a spatial data structure to locate the k -nearest light sources or cameras would speed up the rendering times significantly. However, this optimization has not been performed, resulting in rendering times of about 10 minutes per frame for the doll data set with 5 point light sources on a 1 GHz Pentium III.

View extrapolation is an important issue when deforming surface reflectance fields, as parts of the objects may become visible from directions where no data was acquired before. The extrapolation is implicitly covered by the ULR interpolation. Fig. 12 shows an example.

Fig. 13 shows the quality of appearance preservation at the example of the doll data set. Note the visual preservation of the different materials of the wooden base, the diffuse fabric, and the specular braid of the skirt. All scenes were lit by three colored point light sources (*red*, *green*, and *blue*) from static positions.

Fig. 14 shows various deformations of the beer mug model. Although our approach is not able to preserve refractive effects, the results appear realistic for reasonable deformations.

Our surfel renderer distorts surface splats in object space according to the local warp. This allows for rendering large deformations without visible holes (Fig. 15).

Section V-C discussed the viewing ray divergence coming with the deformation of a surface reflectance field. Fig. 16 shows an example of an SRF warp together with its corresponding camera and light source blending




Fig. 12. View extrapolation. (a) Original doll data set. (b) Occluded parts of the skirt become visible as one arm is lifted. The model is lit by two point light sources. (a) is chosen to be an original view to force the views used for view extrapolation to be at maximum distance from the current view.



Fig. 13. Preservation of material properties *Left*: The original doll surface reflectance field. *Center, Right*: A sinusoidal warp is applied to the data set.

fields. Both blending fields were visualized by assigning random colors to cameras and light sources, respectively. Subsequently, the colors were blended together according to the ULR interpolation coefficients. The blending fields clearly show the viewing ray divergence in distorted regions.

Fig. 17 shows a number of frames of a surface reflectance field animation including large deformations and varying lighting conditions.

VII. CONCLUSIONS

We developed a novel scheme to deform and relight surface reflectance fields. Our technique approximately preserves all material properties of the captured object during deformation. The method uses an enclosing impostor geometry of the rendered object. An appearance preserving data look-up scheme employs a local parameterization of the impostor geometry without relying on proper surface normals. Our method is especially suited for relighting of real world models with complex surface properties for which normals are hard to determine.



Model	Viewpoints	Lights	Surfels	PCA Data
Doll	432 (360°)	60	532 935	4.4 GB
Mug	216 (180°)	44	757 760	1.8 GB

Fig. 11. The presented models. *Left*: The surface reflectance field of a doll. *Right*: A half filled beer mug. The table shows the complexity of each data set.




Fig. 17. Selected frames of a surface reflectance field animation. The model is continuously deformed, while viewpoint and lighting conditions change. The model is lit by colored light sources, denoted by spheres of the light color.



Fig. 14. Deformation of refractive objects. Different warps were applied to the mug data set. All glasses are rendered from the same viewpoint and lit by five point light sources (red, green, blue, 2xwhite). The effect of the deformations on specular highlights can easily be observed.

We presented an extended surfel representation to render deformable surface reflectance fields. Using elliptical splats that follow the local deformation in object space allowed for seamless rendering of deformations while preserving optimal texture quality. At the same time, the surfel representation complements naturally the local impostor parameterization. In principle, the presented look-up scheme can be used for any other image-based




Fig. 15. Even under large deformations, the surfel renderer displays a closed surface. (a) The depicted warp anisotropically stretches parts of the doll data (b) The stretch anisotropy is color encoded showing the ratio of the principal axes of the elliptical surface splats. Splats in blue regions remain circular, green areas show a stretch ratio of 1:2, while red denotes a local stretch of more than 1:4.

approach that uses approximating impostor geometry.

VIII. FUTURE WORK

Simulating global lighting effects during the deformation requires knowledge of the object geometry. While surface reflectance fields only provide an approximate geometry in general, this geometry is usually close to the real object surface. Exploiting this fact may lead to approximate methods to analyse and reproduce global illumination on deformable surface reflectance fields.

Another way to preserve global effects would be to collect additional information during the acquisition process of the surface reflectance field. Matusik et al [20]



Fig. 16. Viewing ray divergence. *Top*: The doll's skirt was distorted. *Left*: The corresponding camera blending field displays the viewing ray divergence in distorted regions. *Right*: The light source blending field.

proposed the *opacity hull* to explicitly capture translucency and refraction through an SRF. This information could be used to reconstruct the internal light transport. In particular, opacity hulls may be used to detect and manipulate self-shadowing inside a surface reflectance field.

Shifted texture parallax can be alleviated by enhancing the impostor geometry. Stereo algorithms could be used on textured objects to improve the geometric approximation.

IX. ACKNOWLEDGEMENTS

We would like to thank Matthias Zwicker, Remo Ziegler, Wojciech Matusik, and Addy Ngan for many valuable discussions.

REFERENCES

- [1] Simon L. Altman. Hamilton, grassmann, rodrigues, and the quaternion scandal—what went wrong with one of the major mathematical discoveries of the nineteenth century? *A Mathematical Association of America journal*, December 1989.
- [2] C. Buehler, M. Bosse, L. McMillan, S. Gortler, and M. Cohen. Unstructured lumigraph rendering. In *Computer Graphics, SIGGRAPH 2001 Proceedings*, pages 425–432, Los Angeles, CA, 2001.

- [3] W-C. Chen, J-Y. Bouguet, M. H. Chu, and R. Grzeszczuk. Light Field Mapping: Efficient Representation and Hardware Rendering of Surface Light Fields. *ACM Transactions on Graphics*, 21(3):447–456, July 2002. ISSN 0730-0301 (Proceedings of ACM SIGGRAPH 2002).
- [4] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink. Reflectance and texture of real world surfaces. *ACM Transactions on Graphics*, 1(18):1–34, 1999.
- [5] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Computer Graphics, SIGGRAPH 2000 Proceedings*, pages 145–156, July 2000.
- [6] P. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Computer Graphics, SIGGRAPH 97 Proceedings*, pages 369–378, Los Angeles, CA, 1997.
- [7] P. Debevec, C. Taylor, and J. Malik. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *Computer Graphics, SIGGRAPH 96 Proceedings*, pages 11–20, August 1996.
- [8] P. Debevec, Y. Yu, and G. Borshukov. Efficient view-dependent image-based rendering with projective texture-mapping. In *Proceedings of the 9th Eurographics Workshop on Rendering*, pages 105–116, Vienna, Austria, June 1998.
- [9] R. Furukawa, H. Kawasaki, K. Ikeuchi, and M. Sakauchi. Appearance based object modeling using texture database: Acquisition, compression and rendering. In *Proceedings of the 13th Eurographics Workshop on Rendering*, Pisa, Italy, June 2002.
- [10] S. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen. The lumigraph. In *Computer Graphics, SIGGRAPH 96 Proceedings*, pages 43–54, New Orleans, LS, August 1996.
- [11] T. Hawkins, J. Cohen, and P. Debevec. A photometric approach to digitizing cultural artifacts. In *2nd International Symposium on Virtual Reality, Archaeology, and Cultural Heritage*, Glyfada, Greece, November 2001.
- [12] P. Heckbert. Fundamentals of texture mapping and image warping. Master's thesis, University of California at Berkeley, Department of Electrical Engineering and Computer Science, June 17 1989.
- [13] W. Heidrich and H.-P. Seidel. Realistic, hardware-accelerated shading and lighting. In *Computer Graphics, SIGGRAPH 99 Proceedings*, pages 171–178, Los Angeles, CA, August 1999.
- [14] J. Kautz and M. McCool. Interactive rendering with arbitrary brdfs using separable approximations. In *Rendering Techniques '99 (Proceedings of the Tenth Eurographics Workshop on Rendering)*, pages 281–292, New York, NY, June 1999. Springer Wien.
- [15] M. Koudelka, P. Belhumeur, S. Magda, and D. Kriegman. Image-based modeling and rendering of surfaces with arbitrary brdfs. In *Proc. of Computer Vision and Pattern Recognition*, page in press, Kauai, HI, December 2001.
- [16] H. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially varying materials. In *Proceedings of the 12th Eurographics Workshop on Rendering*, pages 104–115, June 2001.
- [17] M. Levoy and P. Hanrahan. Light field rendering. In *Computer Graphics, SIGGRAPH 96 Proceedings*, pages 31–42, New Orleans, LS, August 1996.
- [18] T. Malzbender, D. Gelb, and H. Wolters. Polynomial texture maps. In *Computer Graphics, SIGGRAPH 2001 Proceedings*, pages 519–528, Los Angeles, CA, 2001.
- [19] S. Marschner, S. Westin, E. Lafortune, and K. Torrance. Image-based measurement of the bidirectional reflectance distribution function. *Applied Optics*, 39(16), 2000.
- [20] W. Matusik, H. Pfister, A. Ngan, P. Beardsley, R. Ziegler, and L. McMillan. Image-based 3d photography using opacity hulls. *ACM Transaction on Graphics*, 21(3):427–437, July 2002. ISSN 0730-0301 (Proceedings of ACM SIGGRAPH 2002).
- [21] W. Matusik, H. Pfister, R. Ziegler, A. Ngan, and L. McMillan. Acquisition and rendering of transparent and refractive objects.

In *Proceedings of the 13th Eurographics Workshop on Rendering*, Pisa, Italy, June 2002.

- [22] M. McCool, J. Ang, and A. Ahmad. Homomorphic factorization of BRDFs for high-performance rendering. *Computer Graphics*, 35(Annual Conference Series):171–178, 2001.
- [23] G. Miller, S. Rubin, and D. Ponceleon. Lazy decompression of surface light fields for precomputed global illumination. In *Proceedings of the 9th Eurographics Workshop on Rendering*, pages 281–292, Vienna, Austria, June 1998.
- [24] K. Nishino, Y. Sato, and K. Ikeuchi. Appearance compression and synthesis based on 3d model for mixed reality. In *Proceedings of IEEE ICCV '99*, pages 38–45, September 1999.
- [25] K. Nishino, Y. Sato, and K. Ikeuchi. Eigen-texture method: Appearance compression based on 3d model. In *Proc. of Computer Vision and Pattern Recognition*, pages 618–624, June 1999.
- [26] H. Pfister, M. Zwicker, J. van Baar, and M. Gross. Surfels: Surface elements as rendering primitives. In *Computer Graphics*, SIGGRAPH 2000 Proceedings, pages 335–342, Los Angeles, CA, July 2000.
- [27] K. Pulli, M. Cohen, T. Duchamp, H. Hoppe, L. Shapiro, and W. Stuetzle. View-based rendering: Visualizing real objects from scanned range and color data. In *Eurographics Rendering Workshop 1997*, pages 23–34, June 1997.
- [28] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *Computer Graphics*, SIGGRAPH 97 Proceedings, pages 379–387, 1997.
- [29] D. Vlasic, H. Pfister, S. Molinov, R. Grzeszczuk, and W. Matusik. Opacity light fields: Interactive rendering of surface light fields with view-dependent opacity. In *To appear in Proceedings of the Interactive 3D Graphics Symposium 2003*, 2003.
- [30] G. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics*, 26(Annual Conference Series):265–273, 1992.
- [31] D. Wood, D. Azuma, K. Aldinger, B. Curless, T. Duchamp, D. Salesin, and W. Stuetzle. Surface light fields for 3d photography. In *Computer Graphics*, SIGGRAPH 2000 Proceedings, pages 287–296, Los Angeles, CA, July 2000.
- [32] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Computer Graphics*, SIGGRAPH 99 Proceedings, pages 215–224, Los Angeles, CA, August 1999.
- [33] Z. Zhang, L. Wang, B. Guo, and H.-Y. Shum. Feature-based light field morphing. In *Computer Graphics*, SIGGRAPH 2002 Proceedings, pages 457–464, July 2002.
- [34] M. Zwicker, M. Pauly, O. Knoll, and M. Gross. Pointshop 3d: An interactive system for point-based surface editing. In *Computer Graphics*, SIGGRAPH 2002 Proceedings, pages 322–329, San Antonio, TX, July 2002.
- [35] M. Zwicker, H. Pfister., J. Van Baar, and M. Gross. Surface splatting. In *Computer Graphics*, SIGGRAPH 2001 Proceedings, pages 371–378, Los Angeles, CA, July 2001.
- [36] M. Zwicker, H. Pfister, J. Van Baar, and M. Gross. Ewa splatting. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):223–238, 2002.



Tim Weyrich received his MS degree in computer science from the University of Karlsruhe (TU), Germany in 2001. He is currently PhD student in the Computer Graphics Laboratory of the Institute of Technology (ETH) in Zürich. His research interests are point-based graphics, 3D reconstruction, and appearance modeling. He is member of ACM SIGGRAPH and GI.



Hanspeter Pfister is Associate Director and Senior Research Scientist at MERL – Mitsubishi Electric Research Laboratories – in Cambridge, MA. He is the chief architect of VolumePro, Mitsubishi Electric’s real-time volume rendering hardware for PCs. His research interests include computer graphics, scientific visualization, and computer architecture. His work spans a range of topics, including point-based graphics, 3D photography, 3D

face recognition, volume graphics, and computer graphics hardware. Hanspeter Pfister received his Ph.D. in Computer Science in 1996 from the State University of New York at Stony Brook. He received his M.S. in Electrical Engineering from the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland, in 1991. Dr. Pfister has taught courses at major graphics conferences including SIGGRAPH, IEEE Visualization, and Eurographics. He has been teaching introductory and advanced graphics courses at the Harvard Extension School since 1999. He is Associate Editor of the IEEE Transactions on Visualization and Computer Graphics (TVCG), chair of the IEEE Technical Committee on Visualization and Computer Graphics (TCVG), and has served as a member of international program committees of major graphics conferences. Dr. Pfister was the general chair of the IEEE Visualization 2002/03 conference. He is senior member of the IEEE, and member of ACM, ACM SIGGRAPH, the IEEE Computer Society, and the Eurographics Association.



Markus Gross is a professor of computer science and director of the Computer Graphics Laboratory of the Swiss Federal Institute of Technology (ETH) in Zürich. He received a Master of Science in electrical and computer engineering and a PhD in computer graphics and image analysis, both from the University of Saarbrücken, Germany. From 1990 to 1994, Dr. Gross worked for the Computer Graphics Center in Darmstadt, where he established and

directed the Visual Computing Group. His research interests include point-based graphics, physics-based modelling, multiresolution analysis, and virtual reality. He has been widely publishing and lecturing on computer graphics and scientific visualization, and he authored the book “Visual Computing”, Springer, 1994. Dr. Gross has taught courses at major graphics conferences including ACM SIGGRAPH, IEEE Visualization, and Eurographics. He is the associate editor of the IEEE Computer Graphics and Applications and has served as a member of international program committees of many graphics conferences. Dr. Gross has been a papers co-chair of the IEEE Visualization ’99, the Eurographics 2000, and the IEEE Visualization 2002 conferences. He will be chair of the papers committee of ACM SIGGRAPH 2005.