

Seeing People in Different Light—Joint Shape, Motion, and Reflectance Capture

Christian Theobalt, *Member, IEEE*, Naveed Ahmed, Hendrik Lensch,
Marcus Magnor, *Member, IEEE*, and Hans-Peter Seidel

Abstract—By means of passive optical motion capture, real people can be authentically animated and photo-realistically textured. To import real-world characters into virtual environments, however, surface reflectance properties must also be known. We describe a video-based modeling approach that captures human shape and motion as well as reflectance characteristics from a handful of synchronized video recordings. The presented method is able to recover spatially varying surface reflectance properties of clothes from multiview video footage. The resulting model description enables us to realistically reproduce the appearance of animated virtual actors under different lighting conditions, as well as to interchange surface attributes among different people, e.g., for virtual dressing. Our contribution can be used to create 3D renditions of real-world people under arbitrary novel lighting conditions on standard graphics hardware.

Index Terms—3D video, dynamic reflectometry, real-time rendering, relighting.

1 INTRODUCTION

RECENT advances in graphics hardware and rendering algorithms enable the creation of images of unprecedented realism in real-time. In order to capitalize on these novel rendering possibilities, however, ever more detailed and accurate scene descriptions must be available. The price to be paid can be measured in working hours spent to create detailed geometry meshes, complex textures, convincing shaders, and authentic animations: Apparently, scene modeling is becoming a limiting factor in realistic rendering.

One alternative to avoid excessive modeling times consists of capturing suitable models directly from “the real thing.” Image and video-based rendering (IBR/VBR) approaches pursue this notion, aiming at automatically generating visually authentic computer models from real-world recorded objects and events [17]. Many of these techniques show how to interactively render photo-realistic views from real-world-captured, dynamic scenes (see also Section 2). While the ability to realistically display dynamic events from novel viewpoints by itself has a number of intriguing applications, the next step is to use objects that have been captured in the real world for augmenting virtual scenes. To import a real-world object into surroundings different from the recording environment, however, its appearance must be adapted to the new illumination situation. To do so, the bidirectional reflectance distribution function (BRDF) must be known for all object surface points. Data-driven [7], [26] as well as model-based [23],

[19] methods have been proposed to recover and represent the BRDF of real-world materials. Unfortunately, these methods cannot be directly applied to dynamic objects exhibiting time-varying surface geometry and constantly changing local illumination.

We present an approach that jointly captures the shape, motion, and time-varying surface reflectance of people (Fig. 1). An overview of our method is given in Section 3. As input to our algorithm, we only require a handful of calibrated and synchronized video recordings (Section 4). The algorithm returns subject-adapted 3D geometry and animation parameters by employing a markerless motion capture method. In a preprocessing step, multiview texture consistency is enhanced and a lateral shift of textiles is detected (Section 5). From the video input, we estimate a dynamic reflectance description comprising of a spatially-varying BRDF and a dynamic surface normal field (Section 6). PC graphics hardware-assisted rendering then allows us to photo-realistically visualize recorded people at interactive frame rates in changing lighting conditions and from arbitrary perspective (Section 7). We validate our algorithm (Section 8) and present results for several subjects wearing different clothes made of non-Lambertian textiles (Section 9). The paper concludes in Section 10.

The contributions of our paper are:

- an algorithm that warps input video images in order to guarantee multiview photo-consistency in conjunction with inexact object geometry,
- a method to detect and compensate for lateral shifting of textiles,
- dynamic reflectometry from multiview video, i.e.,
 - per-texel BRDF parameter estimation and
 - reconstruction of time-varying normal maps to capture the small, variable detail of surface geometry (e.g., wrinkles in clothing), and
- the integration of the acquisition, estimation, and rendering methods into a working system.

• C. Theobalt, N. Ahmed, H. Lensch, and H.-P. Seidel are with MPI Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany. E-mail: {theobalt, nahmed, lensch, hpseidel}@mpi-inf.mpg.de.

• M. Magnor is with TU Braunschweig, Computer Graphics Lab, Muehlenpfordtstr. 23, 38106 Braunschweig, Germany. E-mail: magnor@cg.tu-bs.de.

Manuscript received 6 May 2006; revised 20 Oct. 2006; accepted 24 Oct. 2006; published online 8 Jan. 2007.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org, and reference IEEECS Log Number TVCG-0064-0506. Digital Object Identifier no. 10.1109/TVCG.2007.1006.




Fig. 1. Our method jointly estimates the shape, motion, and dynamic surface reflectance of (a) a human actor from input video. A dynamic scene description based on (b) a template body model enables us to photo-realistically render moving actors from an arbitrary viewpoint and under both (c) captured real-world and (d) artificial illumination.

The following section highlights related work on free-viewpoint video and BRDF estimation.

2 RELATED WORK

We capitalize on previous research in many areas, but primarily pick up ideas from the fields of free-viewpoint video and image-based reflectance estimation.

Research in free-viewpoint video aims at developing methods for photo-realistic, real-time rendering of previously captured real-world scenes. The goal is to give the user the freedom to interactively navigate his or her viewpoint freely through the rendered scene. Early research that paved the way for free-viewpoint video was presented in the field of image-based rendering (IBR). Shape-from-silhouette methods reconstruct geometry models of a scene from multiview silhouette images or video streams. Examples are image-based [25], [36] or polyhedral visual hull methods [24], as well as approaches performing point-based reconstruction [14]. The combination of stereo reconstruction with visual hull rendering leads to a more faithful reconstruction of surface concavities [21]. Stereo methods have also been applied to reconstruct and render dynamic scenes [42], [17], some of them employing active illumination [34]. Alternatively, a complete parameterized geometry model can be used to pursue a model-based approach toward free-viewpoint video [6]. On the other hand, light field rendering [20] is employed in the 3D TV system [27] to enable simultaneous scene acquisition and rendering in real-time.

IBR methods can visualize a recorded scene only for the same illumination conditions that it was captured in. For correct relighting, it is inevitable to recover complete surface reflectance characteristics.

The estimation of reflection properties from still images has been addressed in many different ways. Typically, a single point light source is used to illuminate an object of known 3D geometry consisting of only one material. One common approach is to take HDR images of a curved object, yielding a different incident and outgoing directions per pixel and thus capturing a vast number of reflectance samples in parallel. Often, the parameters of an analytic BRDF model are fit to the measured data [33], [19] or a data-driven model is used [26]. Zickler et al. [41] proposed a scattered data interpolation method to reconstruct a reflectance model. Reflectance measurements of scenes

with more complex incident illumination can be derived by either a full-blown inverse global illumination approach [37], [12], [4] or by representing the incident light field as an environment map and solving for the direct illumination component only [38], [31], [28]. In our approach, we will approximate the incident illumination by multiple point light sources and estimate BRDF model parameters taking only direct illumination into account.

Reflection properties together with measured photometric data can also be used to derive geometric information of the original object [40]. Rushmeier et al. estimate diffuse albedo and normal map from photographs with varied incident light directions [32], [3]. A linear light source is employed by Gardner et al. [10] to estimate BRDF properties and surface normal. In [11], [13], reflectance and the shape of static scenes are simultaneously refined using a single light source in each photograph.

Instead of explicitly reconstructing a mathematical reflectance model, taking an image-based approach to relighting has also been tried. In [16], a method to generate animatable and relightable face models from images taken with a special light stage is described. Wenger et al. [35] extend the light stage device such that it enables the capturing of dynamic reflectance fields. Their results are impressive; however, it is not possible to change the viewpoint in the scene. Einarsson et al. [8] extend it further by using a large light stage, a treadmill where the person walks, and light field rendering for the display. Human performances can be rendered from novel perspectives and relit. Unfortunately, their method can only process single periodic motion, such as walking, and is only suitable for low frequency relighting. For our 3D video scenario, we prefer a more compact scene description based on parametric BRDFs that can be reconstructed in a fairly simple acquisition facility and that allows for arbitrary viewpoint changes.

Carceroni and Kutulakos present a volumetric method for simultaneous motion and reflectance capture for nonrigid objects [5]. They have shown nice results for spatially confined 3D scenes where they used a coarse set of surfels as shape primitives.

In this paper, we propose a model-based approach that captures the shape, motion parameters, and dynamic surface reflectance of the whole human body at high accuracy. As input, we only expect a handful of synchronized video streams showing arbitrary human motion. Our reconstructed



Fig. 2. The algorithmic workflow of our method.

dynamic scene description enables us to render virtual people in real time from arbitrary viewpoints and under arbitrary lighting conditions.

3 OVERVIEW

Fig. 2 illustrates the workflow between the components of our joint shape, motion, and reflectance capture approach. The input to our system consists of multiview video sequences that are recorded using eight synchronized color video cameras (Section 4). The reflectance estimation sequence (RES) is used to estimate surface reflectance properties. Arbitrary human motion is captured in the dynamic scene sequences (DSS) and it is these sequences that are later visualized and relit. In both types of sequences, the person wears identical clothes. The respective data paths for both input sequences are shown in Fig. 2. A generic body model is adapted to match the shape and proportions of the recorded person. Subsequently, human pose parameters are computed for all time frames by means of a silhouette-based marker-free motion capture approach (Section 5.1). To store all per-surface element data needed during reflectance estimation in texture space, we make use of surface parameterization methods. Multiview video (MVV) textures are generated by transforming each input video image into the texture domain (Section 5.2). To correct for photo-inconsistencies due to inexact body geometry, the input images can be warp-corrected prior to MVV texture generation (Section 5.3). Furthermore, texture registration problems due to shifting of the apparel over the body surface are prevented by means of an automatic cloth shift detection procedure (Section 5.4). From the RES video data, BRDF model parameter values are estimated for each surface texture element (texel) of the geometry model individually (Section 6.1). The recovered local reflectance properties then allow us to estimate the time-varying surface normal field in the DSS sequences (Section 6.2). The moving body model, its spatially-varying reflectance, and the time-varying normal field enable us to interactively render and instantaneously relight the DSS sequences from an arbitrary viewpoint and illumination direction (Section 7, Section 8, and Section 9).

4 ACQUISITION

As input to our system, we record multiview video (MVV) sequences. A multiview recording setup enables us to capture a volume of approximately $4 \times 4 \times 3$ m with eight externally synchronized video cameras. Seven of the

cameras are placed in an approximately circular arrangement around the center of the scene, the eighth camera records the set from an overhead position. Since we estimate both motion and reflectance properties, we have strict requirements concerning the spatial, temporal, and color resolution of our imaging devices. We employ Imperx MDC-1004 cameras that feature a $1,004 \times 1,004$ pixel CCD sensor with linear 12 bits-per-pixel dynamic range providing a sustained frame rate of 25 fps. The captured video data is streamed in real time to eight parallel RAID systems. The cameras are calibrated and intercamera color-consistency is ensured by applying a color-space transformation to each camera stream.

The lighting conditions in our studio are fully controllable. No exterior light can enter the recording area, and the influence of indirect illumination is minimized by covering up all of the walls and the floor with black cloth and carpet. We employ two K5600 Jokerbug 800 spot lights to illuminate our scenes. They are placed in opposite corners of our stage, and they are oriented toward the center of the recording area (Fig. 3). During estimation, we approximate the contribution of each spot light by a single point light source. Light source positions, intensities, and color response of the cameras are calibrated offline. Our flexible approach can also handle more general setups comprised of spot lights and diffuse illumination from the ceiling.

We successively record two MVV sequences for each person and each type of apparel. The first sequence, referred to as the reflectance estimation sequence (RES), serves as input to the BRDF estimation algorithm. The human subject performs, on the very spot, a simple 360 degree rotation around his vertical body axis which only takes a few seconds. The actor keeps the upper body in



Fig. 3. Illustration of camera and light source arrangement.



Fig. 4. Detection of relative cloth shift between two subsequent combined textures C_t and C_{t+1} (in parameterization II) by means of optical flow. Areas where shifting textiles were found are marked in red on the 3D model.

a static initialization posture. This way, we can safely assume that, for the duration of the RES, the apparel does not shift over the body. Prior to reflectance estimation, we fit our geometry model to each recorded body pose (Section 5.1). For each point on the model's surface, the RES contains (in the best case) as many different appearance samples as there are images depicting the respective point. Over time, the surface element's normal points in various directions with respect to light sources and imaging sensors. While surface normal orientation varies freely, our static camera and lighting setup allows for only a limited number of half vector directions $\hat{h}_j = \hat{l} + \hat{v}_j$, i.e., angular separations between spot light \hat{l} and camera directions \hat{v}_j . By placing the eight cameras nonsymmetrically with respect to the two spot lights, we gather samples for up to 16 different light-to-camera angles, which we found sufficient to robustly fit our isotropic BRDF models (Section 6).

The second type of input video footage, the dynamic scene sequences (DSS), capture the actor's performances from which the actual relightable 3D videos are to be generated. From the DSS, we also reconstruct a time-varying surface normal field (Section 6.2).

5 MOTION CAPTURE AND TEXTURE GENERATION

We use a model-based approach to represent the time-varying geometry of the actor. To estimate our model's shape and motion parameters, we perform markerless

motion capture. In order to merge the video data from multiple viewpoints, we compute static texture parameterizations, resample and align the input streams using a novel warp-correction technique, and temporally register the video streams by detecting shifting of the apparel.

5.1 Model-Based Motion Capture

To simultaneously recover dynamic 3D surface geometry as well as animation parameters, we use an extended version of the model-based, passive optical motion capture approach proposed in [6]. A kinematic body model consisting of 16 separate segments represented by individual triangle meshes is used for automatic silhouette-based motion estimation and shape customization. In a semi-automatic procedure, the segmented model can be transformed into a new model with closed single-skin surface geometry (Fig. 1b). To do so, the triangle meshes of adjacent segments are zippered and surface skinning weights are assigned to correctly model motion-dependent geometry deformation. While we use the segmented version of the model for motion estimation and shape individualization, we use the single-skin model for reflectance estimation and later rendering.

5.2 Texture Parameterization

We jointly apply two different static surface parameterizations of the single-skin mesh over the planar domain based on [39]. To this end, the mesh is cut along a predefined path. Parameterization I leaves the mesh boundary free and minimizes surface distortion (the left image in Fig. 5 shows a texture in parameterization I). We use it during BRDF and normal estimation since it leads to a fairly uniform distribution of samples on the model's surface. Parameterization II forces the boundary of the mesh to lie on the border of a square (Fig. 4 shows two textures in parameterization II). This causes larger distortion, but the textures with fixed boundary are ideal for cloth shift detection (Section 5.4).

Using our surface parameterizations, all data related to surface elements (e.g., image samples, normals, visibility information, light vectors) can be conveniently stored as textures. Throughout our experiments, we use $1,024 \times 1,024$ -texel texture maps. Before any further processing




Fig. 5. MVV texture generation (parameterization I) for camera 0: The color information for each surface point on the body model is not looked up in the original input video frame recorded from camera 0 (column with darker blue background). Instead, the texel color is taken from an image that has been obtained by reprojecting the model that has been textured with camera image 0 into the camera view that sees the surface point most head-on (second column of images). Given eight cameras, 56 warped reprojected images exist for one time step (image matrix in center).

commences, the graphics hardware is used to transform each video camera image into the texture domain. For each video time step, eight so-called multiview video textures (MVV textures) are created for both of the two described texture parameterizations.

5.3 Warp Correction

Although the body model initialization procedure yields a faithful representation of the person’s true geometry, small inaccuracies between the real human and its digital counterpart are inevitable. Due to these geometry inaccuracies, pixels from different input views may get mapped to the same texel position in different MVV textures, even though they do not correspond to the same surface element of the true body geometry.

One common strategy to enhance model-to-texture consistency is to deform the geometry until an overall photo-consistency measure is maximized. Geometry deformation-based optimization, however, tends to give unstable results, in particular due to nonlinear optimizations that are normally required. We take an alternative approach. Instead of moving surface elements to their correct locations in 3D, we move the image pixels within the 2D input image planes until they all become photo-consistent given the available geometry. The following example illustrates our modified MVV texture generation scheme (see Fig. 5 and video in [1]).

Let us assume we want to assemble an MVV texture from the video image $I_j(t)$ seen by camera j at time t . For texel k in the MVV texture, we find out which camera sees it best by searching for the minimal deviation between camera viewing vectors and the surface element normal. If the camera that sees the surface point best is j , the texel color is taken from $I_j(t)$. In case camera $i \neq j$ sees the point best and it is not occluded, we regard the video image $I_i(t)$ as the reference image. The model at time t is projectively textured with $I_j(t)$ and rendered into camera view i . The image of the reprojected textured model is warped such that it is optimally aligned with the reference image. The color of k is taken from the warped image. This way, all texel color values stem from the same physical camera image. The texel color, however, is always taken from a version of that camera image that has been optimally aligned with the camera view that sees the corresponding surface element most head-on.

Warped images are precomputed for all 56 possible combinations of i and j at every time step (Fig. 5). To establish per-pixel correspondences, the warping operation itself is based on the optical flow [22] between the reference image and the image of the reprojected textured model. A regular 2D triangle mesh is superimposed on the reprojected model image, per-vertex displacements are derived from the optical flow values, and the mesh is deformed accordingly via thin-plate spline interpolation [9]. Finally, the warped reprojected image is created on the GPU.

Optical flow is based on the assumption that all surfaces in the scene are diffuse. For reflectance estimation, though, we deliberately generate specular highlights on surfaces in the scene. Experiments show that our method nonetheless produces good results since, for the majority of garments, diffuse reflectance is the predominant contribution.

5.4 Cloth Shift Detection

Our BRDF estimation procedure (Section 6.1) yields a fixed assignment of BRDF parameter values to individual body surface locations which it stores as static textures. During dynamic normal estimation (Section 6.2) and rendering (Section 7), however, this rigid attribution of reflectance properties may not hold since, in a DSS, in contrast to the RES, it is likely that textiles undergo transversal shift over the body surface. It is thus necessary to detect cloth shifting and to take it into account for these two steps. We have developed the following procedure: For each time step of a DSS and the last time step of the corresponding RES, all MVV textures in parameterization II are weightedly blended into combined textures. Each combined texture describes the complete surface appearance of the person in a single square image (Fig. 4). The combined texture of the single RES time step is regarded as the reference texture since it captures the surface appearance from which BRDF parameters are estimated. Cloth shift in 3D corresponds to image motion in the texture domain, which we compute via optical flow (Section 5.3). For each time step of a DSS, we need to compute the motion of each texel with respect to its position in the BRDF parameter textures. Due to lighting changes, this absolute shift-motion cannot be computed by direct comparison to the reference texture. Instead, we first compute the relative shift-motion between all subsequent combined textures C_t and C_{t+1} (Fig. 4). The motion between the reference texture and C_0 is also computed. The absolute shift-motion of a texel in C_t is the sum of all its shift motions that have been detected at previous time steps. We use the GPU to render a texture-space representation of absolute shift-motion of each texel at each time step. To this end, we use the mesh-based warping technique described in Section 5.3 to warp the color-coded BRDF texture coordinates in parameterization II according to the absolute shift-motion.

6 DYNAMIC REFLECTOMETRY

Our reflectance estimation approach consists of two steps. In the first step, we determine BRDF parameter values per pixel from the reflectance estimation sequence. An iterative estimation process enables us to handle geometry inconsistencies between the real object and the much smoother human body model. In the second step, we compute even time-varying normal maps per frame to capture surface detail such as wrinkles in clothing whose shape and extent depend on the current pose of the person. The underlying technique is similar to [19], which we have extended in order to cope with multiple light sources, time-varying data, and interframe consistency.

6.1 BRDF Estimation

We estimate a set of spatially-varying BRDFs for each person and each outfit from the respective reflectance estimation sequence (RES) explained in Section 4. The pose parameters for the RES have been determined beforehand. For each surface element, the BRDF representation consists of an individual diffuse color component that is specific to the surface point and a set of specular parameters that are shared by all surface points belonging to the same material.




Fig. 6. Subsequent steps to estimate per-texel BRDFs.

Our framework is flexible enough to incorporate any parametric reflectance model. However, in the majority of our experiments, we employ the parametric BRDF model proposed by Phong [29]. We have also tested our method with the model proposed by Lafortune et al. [18] using two specular lobes.

In general, our estimation of BRDF parameters and, later, the estimation of the time-varying normals, is based on minimizing for each surface point \vec{x} the error $E(\vec{x}, \rho(\vec{x}))$ between the current model $\rho(\vec{x})$ and the measurements for this point from all cameras i at all time steps t :

$$E(\vec{x}, \rho(\vec{x})) = \sum_t^N \sum_i^8 \kappa_i(t) \left(S_i(t) - \left[\sum_j^J \lambda_j(t) (f_r(\hat{l}_j(t), \hat{v}_i(t), \rho(\vec{x})) \cdot I_j(\hat{n}(t) \cdot \hat{l}_j(t))) \right] \right)^2. \quad (1)$$

The term is evaluated separately in the red, green, and blue color channel. $S_i(t)$ denotes the measured color samples at \vec{x} from camera i and I_j denotes the intensity of light source j . The viewing directions $\hat{v}_i(t)$ and light source directions $\hat{l}_j(t)$ are expressed in the point's local coordinate frame based on the surface normal $\hat{n}(t)$. The visibility of the surface point with respect to each camera is given by $\kappa_i(t)$ and with respect to the light sources by $\lambda_j(t)$, both being either 0 or 1. f_r finally evaluates the BRDF. All information that is relevant for one texel can thus be grouped into an implicit data structure we call *dynamic texel* or *dysel*.

Using nonlinear optimization, this formula can in principle be used to determine a full BRDF and the surface normal at the same time. However, we apply an iterative approach. We carefully designed the reflectance estimation sequence to obtain a much more stable optimization. The subsequent steps of our iterative BRDF estimation scheme are *material clustering*, *first BRDF estimation*, *normal estimation*, and *refined BRDF estimation*, Fig. 6.

Instead of determining the specular part of the BRDF per pixel, we assume that there is only very little variation of the specular part within the same material, e.g., skin, hair, or the different fabrics. By combining the measurements of multiple surface points exhibiting the same material, we increase the number of samples and, more importantly, the variation in viewing and lighting directions in order to obtain a more faithful specular estimate. Please note that there are no specific requirements concerning the spatial layout and coherence of surface areas made of similar material. Although none of our test subjects wore strongly textured apparel, we believe that our approach can also

handle surfaces with a lot of fine detail as long as the camera resolution is sufficient to capture the spatial frequency. The *clustering* step determines to what material a surface element, i.e., each texel in the texture maps, belongs. The number of materials is preset a priori. We employ a straightforward color-based clustering approach that considers the raw texel color values.

During the *first BRDF estimation* pass, an optimal set of per-texel BRDF parameters is determined using the normals of the default geometry. The estimation itself consists of a nonlinear minimization of (1) in the BRDF parameters. For optimization, we make use of a Levenberg-Marquardt minimization scheme [30] in the same manner as [19]. First, we find an optimal set of parameter values for each material cluster of texels. To quantify the estimation error per material cluster, we sum the error term in (1) for all surface elements that belong to the cluster. Given the average BRDF for each material, we can render the model by applying only average specular reflectances. By subtracting this specular component from each sample, we generate new dynamic texels that contain purely diffuse reflectance samples. Using these purely diffuse samples, an individual diffuse component is estimated for each surface element (texel) by minimizing (1) over the diffuse color parameter. The output of the first BRDF estimation is then a set of spatially-varying BRDF parameters ρ_{first} .

The default normals of the human body model cannot represent subtle details in surface geometry, such as wrinkles in clothing. In a *normal estimation* step, we make use of the first set of estimated BRDF parameters ρ_{first} in order to reconstruct a refined normal field via photometric stereo. Please note that our algorithm is not photometric stereo in the classical sense since the viewing directions vary. However, we will continue to use this term for our algorithm and ask the reader to keep that difference in mind. In order to make this reconstruction tractable, we implicitly assume that the local normal directions do not change while the person is rotating in place. During normal estimation, we minimize the following extended version of the energy functional (1) in the local surface normal direction $\hat{n}(\vec{x})$ of each surface point \vec{x} :

$$E_{normal}(\vec{x}, \hat{n}(\vec{x})) = \alpha E(\vec{x}, \rho(\vec{x})) + \beta \Delta(\hat{n}(\vec{x}))^\gamma. \quad (2)$$

The additional regularization term $\Delta(\hat{n}(\vec{x}))$ penalizes angular deviation from the default normal of the body model. This way, we also enforce convergence to a plausible local minimum of the nonconvex energy functional. The coefficients α and β are weighting factors summing to one and the exponent γ controls the penalty impact. Appropriate values are found through experiments. Normal




Fig. 7. Two-dimensional illustration of robust time-varying normal map estimation. (a) The sequence subdivided into short chunks. For each chunk, one best-matching normal is derived per texel which is assigned to the chunk’s center time step (white). (b) Intermediate time steps are interpolated.

estimation robustness is further improved if only those color samples in a dyxel are used that come from the two best camera views. In the *second BRDF estimation*, we capitalize on the refined surface normals and compute a more accurate set of BRDF parameters, ρ_{final} . As output, we generate one floating-point texture for each BRDF model parameter. Through visual inspection, we found out that the BRDF estimation result that was obtained after one iteration does not change noticeably anymore in further iterations. Thus, it is safe to run just one iteration and, by this means, keep processing times as short as possible.

6.2 Estimation of Time-Varying Normals

The BRDF reconstructed in the previous step enables us to relight any dynamic scene in which the person wears the same apparel as in the respective RES. To generate a visually compelling rendition, however, we found that we need not only accurate reflectance, but also a representation of the small surface geometry details that appear and disappear while a person is moving. We are able to capture these geometry details by estimating a time-varying surface normal field for each DSS via photometric stereo.

Motion parameters for the DSS are found by means of our silhouette-based tracking approach (Section 5). The time-varying normal direction is estimated for each surface point individually. The estimation procedure is a nonlinear minimization of the regularized energy function (2) in the normal direction. Our cloth shift detection method (Section 5.4) has stored information about potential temporal displacement of textiles as a stream of warped BRDF texture coordinates. We account for this by performing BRDF parameter lookups using these temporally varying texture coordinates. In order to robustly perform photometric stereo and to minimize the influence of measurement noise, a sufficient number of samples has to be collected for each surface point. To achieve this purpose, we assume that changes in local normal direction within a short window in time can be neglected. The input sequence is split into subsequent chunks of frames (typically of length 5). For every point \vec{x} on the body surface, we fit an optimal normal \hat{n} to each chunk of video individually. After the time-varying normals have been estimated at this coarse scale, the normal directions between subsequent chunks are interpolated via spherical linear interpolation (Fig. 7). For some sequences, we are able to correct undersampling problems by adaptively enforcing consistency with the template model. To this end, in those surface areas

exhibiting a high fitting error, we rerun the normal estimation with a higher influence of the penalty term in (2). This way, we obtain a normal field that is smooth in the temporal domain and faithfully models subtle dynamic details in surface geometry.

7 RENDERING

Our approach outputs a relightable dynamic object description that consists of the animated geometry and the reconstructed material properties. The material properties consist of the time-independent BRDF textures, the dynamic normal maps, and a sequence of warped BRDF texture coordinates modeling cloth shifting.

During real-time playback, our renderer *Arena* displays the single-skin body model in the smoothly interpolated sequence of captured body poses. The final appearance of the body surface is determined on the GPU by BRDF shader programs implemented in Nvidia’s Cg language. We can currently demonstrate Phong and Lafortune shader implementations. Due to limited camera resolution, changing facial expressions cannot be faithfully captured by our dynamic reflectance representation. We therefore do not compute full BRDF lighting in the face, but only simulate diffuse surface appearance. To this end, we use the fragment colors from the weightedly blended input video streams as diffuse reflectance components. Our real-time renderer can illuminate the dynamic scene from multiple point or directional light sources with activated shadow mapping. Relighting from static environment maps that are approximated by discrete directional light sources via importance sampling is also feasible. On a 3.0 GHz Pentium4, we achieve 19 fps sustained rendering frame rate if a model with 90,000 triangles and four point lights with activated shadow mapping are applied. Using an environment map for relighting, we can still render the scene at 6 fps if 16 approximating lights are employed.

8 VALIDATION AND DISCUSSION

We have validated our approach by visual inspection and quantitative evaluation. In total, we have processed 15 different input sequences using Phong and Lafortune BRDFs. They cover three different human subjects, six different types of apparel, and comprise 150 to 350 frames each. For numerical verifications, we restrict ourselves to Phong sequences.

Warp correction (Section 5.3) leads to an average reduction in absolute image difference between reference and reprojected image in the range of 6 percent. These pairwise registration enhancements result in improvements in multiview texture-to-model consistency; see Fig. 8.

Cloth shift compensation also contributes to more realistic renditions, especially since the positions of seams of pieces of apparel are more reliably captured; see Fig. 9. Although both of these methods lead to visual improvements in locally confined surface areas, they have a very noticeable positive influence on the overall visual quality.

We have also measured the average peak-signal-to-noise-ratio with respect to one input video stream not used for reconstruction. By employing warp correction, we obtain a PSNR improvement of 0.2 dB. Cloth shift compensation typically accounts for an additional gain of




Fig. 8. (a) Result without warp-correction prior to reflectance estimation—ghosting due to misalignments along the stripes of the trousers are visible. (b) Result with warp-correction—ghosting artifacts have been significantly reduced due to better multiview consistency. Block artifacts are due to limited texture resolution.

0.1-0.2 dB. This numerically confirms the achieved reconstruction improvements. However, the global quality measure does not do full justice to the perceivable local visual improvements that we achieve.

For our captured test persons, we neither have ground truth BRDF parameters nor ground truth normal maps at our disposition. We therefore assess the quality of the reconstructed dynamic reflectance models by visual comparison to the input video footage. Our approach captures dynamic scene geometry reliably and jointly estimates dynamic surface reflectance of different materials (Fig. 12). Diffuse and specular reflectances are faithfully discriminated. The realistically reproduced specular reflection of the trousers of a male test subject is shown in the accompanying video that can be downloaded from [1]. Subtle details like time-varying wrinkles in clothing are faithfully captured in normal maps, Fig. 12a, and realistically relit under varying illumination conditions, Fig. 12b.

To get as-good-as-possible reflectance data for one RES, we generated a high-quality surface mesh using a laser range scanner and rotated the person on a turntable. We




Fig. 9. Zoom-in on renditions of the t-shirt's waist seam. Block artifacts are due to limited texture resolution. (a) shows the model textured with the blended input camera views as reference. (b) shows the appearance of the rendered seam if cloth shift detection has not been used during reconstruction. Parts of the trousers are erroneously displayed in orange color. In contrast, (c) shows the correctly rendered seam if textile motion has been considered during the estimation. Although the shift detection is not pixel-accurate, it still leads to a more realistic reproduction of the true appearance of the actor's apparel.

TABLE 1
Comparison of the Estimated Average Diffuse and Specular Phong BRDF Parameters that Have Been Estimated Using the Scanned and the Template Mesh (Values Are Given for Three of the Surface Materials)

Material	Scanned	Template
Shirt (Diffuse)	0.352	0.336
Shorts (Diffuse)	0.773	0.743
Skin (Diffuse)	0.448	0.445
Shirt (Specular)	0.012	0.013
Shorts (Specular)	0.014	0.010
Skin (Specular)	0.015	0.012

compared BRDF parameters and normal maps computed with our model to the ones obtained with the scanned model. The per-material BRDF parameters found with our approximate geometry only deviate minimally; see Table 1. Our surface normals nicely capture surface details, most importantly, wrinkles and seams of pieces of apparel. In some places, our normal field even exceeds the resolution of the scan (see Fig. 10).

To verify the rendering quality under novel lighting conditions, we have recorded one of our male test subjects under two different calibrated lighting setups, henceforth termed LC A and LC B. In each of the lighting setups, just one of the spot lights has been employed. The positions of the light sources in LC A and LC B are (angularly) approximately 45 degrees apart with respect to the center of the scene. We reconstructed the BRDF of the test subject using lighting setup LC A. Subsequently, we compared the ground truth images of the person illuminated by setup




Fig. 10. (a) Color-coded surface normals of a laser-scan of a person in shorts, (b) the normal field estimated with our method, and (c) the surface normals of our the template model. The three coordinates of the surface normal vectors in world space are encoded in the three RGB color channels. One has to note that our laser-scanner is not well-suited to full-body scanning and, thus, exact geometry of the head, the hands, and the feet could not be reconstructed. Furthermore, due to resolution limitations, geometry details were smoothed out and subtle wrinkles were not captured. Even though one can see that (b) our estimated normal field nicely reproduces the wrinkles that have been captured in the scan, these geometry features are not present in (c) the smooth template geometry. One can even see that our image-based approach does not smooth out details like the seam of the t-shirt and the seam of the shorts. Instead, it captures them accurately.




Fig. 11. Rendering under novel calibrated lighting conditions: (a) shows the ground-truth (recorded) appearance of the actor from one input view under calibrated lighting condition LC B. (b) shows the rendered model under LC B from the same viewpoint. For BRDF reconstruction, a different calibrated lighting condition LC A was used. The comparison shows that we are able to reproduce the true appearance of the actor even under novel lighting conditions.

LC B, Fig. 11a, to the renderings of the person under LC B, Fig. 11b. This comparison illustrates that our method realistically reproduces the true appearance of even subtle surface details under incident illumination that is different from the lighting conditions used during reconstruction.

Our approach is subject to a couple of limitations: Reflectance estimation is impossible if the shape estimation using the marker-less motion capture method already fails. There are certainly extreme postures, such as the fetal position, which could not be faithfully captured. However, we have already shown in [6] that our markerless motion estimation method can capture fairly complex motion, such as ballet dance, and can also handle poses where the arms come close to the body. In our test sequences, we show fairly general everyday motion, such as walking and a simple dance. Please note that, in these sequences, people are holding their arms away from the body in order to expose the otherwise occluded torso areas to the viewer. This way, we demonstrate that we can actually capture the appearance of areas possibly occluded to some cameras or light sources. It is thus not due to a limited operational range of our algorithm that these body poses were chosen. We believe that we can also handle faster and more complex motion in the same way.

Another limitation is that the single-skin surface model is generated in an interactive procedure. However, this is not a principal limitation of our method since model generation could be fully automated by a method similar to [2]. Furthermore, although we can handle normal everyday apparel, we cannot account for loose apparel, where surfaces can deviate almost arbitrarily from the body model.

Although cloth shift compensation and warp correction lead to visual improvements in the majority of cases, isolated local deteriorations are still possible. Cloth shift detection, for example, sometimes erroneously classifies evolving wrinkles as shifting of apparel. Also, in case of strongly incorrect body geometry, warp correction may induce noticeable discontinuities on the surface, e.g., due to changing reference cameras or visibility boundaries. Luckily, for the types of scene we intend to handle, body shape is already so close to the true geometry that these discontinuities play no significant role.

We nonetheless leave the decision if either of the two methods are used to the user.

Moreover, we currently neglect interreflections on the body. In the RES, they potentially play a role between the wrinkles in clothing. To prevent this effect from degrading the estimation accuracy, we have taken care to minimize the number of wrinkles in the RES.

Sometimes we observe small rendering artifacts due to undersampling (e.g., on the underneath of the arms). However, since we decouple the BRDF estimation from the dynamic normal map estimation, we can reduce undersampling problems significantly. First, undersampled areas during BRDF estimation can be minimized by recording an RES sequence showing several rotation motions with different body postures. If, even in this extended RES, a pixel is never seen by any of the cameras, we fill in reflectance properties from neighboring regions in texture space. This still leaves the possibility that, during dynamic normal estimation, a surface point is not exposed to any camera. In this case, we have to resort to the normal coming from the template model since no other evidence is available. We would like to point out that all of these undersampling issues are general problems of any passive vision system and not specific to our approach.

Finally, our current face rendering technique is a compromise that produces visually compelling but not physically correct results. Capturing the reflectance of the face separately with a more closely spaced camera setup would permit full BRDF relighting.

Despite these limitations, our method is an effective combination of algorithmic tools that allows for the creation of realistic relightable dynamic scene descriptions.

9 RESULTS

In the previous section, we have given evidence of the performance of individual algorithmic components of our joint motion and reflectance estimation scheme. In this section, we show that the interplay of these algorithmic ingredients enables the authentic rendition of moving real-world people in real-time. Our dynamic scene description allows us to photo-realistically render human actors under both the artificial, Fig. 1d, and real-world illumination that has been captured in high-dynamic range environment maps, Fig. 1c. We can also implant actors into virtual environments as they are commonly used in computer games, such as a little pavilion, Fig. 12c, or a pool house, Fig. 12d. Here, our dynamic surface reflectance description even enables the realistic display of the person's reflection in the mirroring floor and the water surface. Relightable 3D videos can be displayed from arbitrary viewpoints and even real-world lighting environments can be rotated interactively, see Figs. 12e, 12f, 12g, 12h, 12i, and the accompanying video (that can be downloaded from [1]).

Once we have estimated the BRDF for one type of clothing, we can also use the surface appearance description to change the apparel of a person even for motion sequences in which the person was originally dressed differently (see video).

The motion estimation process takes around 3 s per time step. Reflectance estimation takes at most 1 hour for the RES. Normal estimation takes approximately 50 s per time step, and it can be parallelized to bring the computation time down.




Fig. 12. (a) Input video frame (left) and color-coded rendition of estimated normal map (right). In the color-coded rendition, the three world-space coordinates of each surface normal have been used as RGB color components. Note that wrinkles are faithfully captured in the normal field. (b) Under varying incident illumination conditions, the wrinkles are authentically relit. (c) Actor dancing in a pavilion while being illuminated by white spot lights. Note the physically correct reflection in the mirroring floor. (d) In a scene that could be taken from a future computer game, the actor moves in a pool house lit by purple light sources. Humans also appear photo-realistically under captured real-world illumination, such as (e), (g), and (h), the St Peter's Basilica environment map, and (f) and (i), the Grace Cathedral environment map provided by Paul Debevec. Both viewpoint and environment lighting can be freely rotated in real time.

Optional input frame warping takes around 10 s for one pair of reference image and reprojected image. Cloth shift compensation accounts for an additional 35 s of computation time for one time step of video.

Even with the relighting option, our final data format is more compact than the traditional 3D video format [6], which is comprised of eight raw video streams for projective texturing and the moving model. As an example, for a 330 frame sequence, the data size of the complete relightable video (moving geometry + BRDF + normal maps) is 528 MB, whereas the total size of the raw input frames is 1,452 MB.

Our results demonstrate that we have developed an effective novel method for simultaneous capture of dynamic scene geometry, per-vertex BRDFs, and time-varying normal maps from multiview video. The acquired scene description enables realistic real-time rendition of relightable 3D videos.

10 CONCLUSION

Our video-based modeling approach jointly captures shape, motion, and surface reflectance of a person. From eight synchronized multivideo streams, we recover all information

necessary to photo-realistically render a recorded person from an arbitrary viewpoint and in arbitrary illumination. The ability to perform convincing relighting enables us to implant real-world, animated people into virtual surroundings. The abstract description of people's appearance in terms of geometry, animation, and surface reflectance further allows us to separate surface appearance from geometry. This way, we can interchange surface attributes among different people, e.g., for dressing one person with another person's clothes. Moreover, we employ a compact data format for our scene description that can be acquired with only a handful of imaging sensors.

Joint shape, motion, and reflectance capture can not only be applied to humans, but to any dynamic object whose motion can be described by a kinematic chain and for which a suitably parameterized geometry model is available. For BRDF parameter recovery, the proposed algorithm currently assumes that the subject is illuminated by a small number of spot lights. Our approach, however, can be extended to use more general lighting configurations, such as the light stage device [35].

ACKNOWLEDGMENTS

This project has been supported by EU 3DTV NoE project No. 511568. The authors would also like to thank Jan Petersen, Art Tevs, Natascha Sauber, Gernot Ziegler, and Edilson de Aguiar for their assistance. They are also grateful to Vlastimil Havran for providing his importance sampling code [15].

REFERENCES

- [1] <http://www.mpi-inf.mpg.de/~theobalt/RFVV/rfvv.avi>, 2007.
- [2] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rogers, and J. Davis, "SCAPE—Shape Completion and Animation of People," *ACM Trans. Graphics (Proc. ACM SIGGRAPH '05)*, vol. 24, no. 3, pp. 408-416, 2005.
- [3] F. Bernardini, I.M. Martin, and H. Rushmeier, "High-Quality Texture Reconstruction from Multiple Scans," *IEEE Trans. Visualization and Computer Graphics*, vol. 7, no. 4, pp. 318-332, Oct./Dec. 2001.
- [4] S. Boivin and A. Gagalowicz, "Image-Based Rendering of Diffuse, Specular and Glossy Surfaces from a Single Image," *Proc. ACM SIGGRAPH '01*, pp. 107-116, 2001.
- [5] R.L. Carceroni and K.N. Kutulakos, "Multi-View Scene Capture by Surfel Sampling: From Video Streams to Non-Rigid 3D Motion Shape & Reflectance," *Proc. IEEE Int'l Conf. Computer Vision (ICCV '01)*, pp. 60-67, 2001.
- [6] J. Carranza, C. Theobalt, M.A. Magnor, and H.-P. Seidel, "Free-Viewpoint Video of Human Actors," *Proc. ACM SIGGRAPH '03*, pp. 569-577, 2003.
- [7] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar, "Acquiring the Reflectance Field of a Human Face," *Proc. ACM SIGGRAPH '00*, pp. 145-156, 2000.
- [8] P. Einarsson, C.-F. Chabert, A. Jones, W.-C. Ma, B. Lamond, T. Hawkins, M. Bolas, S. Sylwan, and P. Debevec, "Relighting Human Locomotion with Flowed Reflectance Fields," *Rendering Techniques*, pp. 183-194, 2006.
- [9] G. Farin, *Curves and Surfaces for CAD: A Practical Guide*. Morgan Kaufmann, 1999.
- [10] A. Gardner, C. Tchou, T. Hawkins, and P. Debevec, "Linear Light Source Reflectometry," *ACM Trans. Graphics (Proc. ACM SIGGRAPH '03)*, vol. 22, no. 3, pp. 749-758, 2003.
- [11] A.S. Georgiades, "Recovering 3-D Shape and Reflectance from a Small Number of Photographs," *Proc. Eurographics Symp. Rendering*, pp. 230-240, 2003.
- [12] S. Gibson, T. Howard, and R. Hubbard, "Flexible Image-Based Photometric Reconstruction Using Virtual Light Sources," *Computer Graphics Forum*, vol. 20, no. 3, 2001.
- [13] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz, "Shape and Spatially-Varying BRDFs from Photometric Stereo," *Proc. IEEE Int'l Conf. Computer Vision (ICCV '04)*, pp. 341-448, 2004.
- [14] M.H. Gross, S. Würmlin, M. Näf, E. Lamboray, C.P. Spagno, A.M. Kunz, E. Koller-Meier, T. Svoboda, L.J. Van Gool, S. Lang, K. Strehlke, A.V. Moere, and O.G. Staadt, "Blue-c: A Spatially Immersive Display and 3D Video Portal for Telepresence," *ACM Trans. Graphics (Proc. ACM SIGGRAPH '03)*, vol. 22, no. 3, pp. 819-827, 2003.
- [15] V. Havran, M. Smyk, G. Krawczyk, K. Myszkowski, and H.-P. Seidel, "Importance Sampling for Video Environment Maps," *Proc. ACM SIGGRAPH '05*, Aug. 2005.
- [16] T. Hawkins, A. Wenger, C. Tchou, A. Gardner, F. Göransson, and P. Debevec, "Animatable Facial Reflectance Fields," *Proc. Eurographics Symp. Rendering*, pp. 309-319, 2004.
- [17] T. Kanade, P. Rander, and P.J. Narayanan, "Virtualized Reality: Constructing Virtual Worlds from Real Scenes," *IEEE Multimedia*, vol. 4, no. 1, pp. 34-47, Jan.-Mar. 1997.
- [18] E. LaFortune, S. Foo, K. Torrance, and D. Greenberg, "Non-Linear Approximation of Reflectance Functions," Aug. 1997.
- [19] H.P.A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel, "Image-Based Reconstruction of Spatial Appearance and Geometric Detail," *ACM Trans. Graphics*, vol. 22, no. 2, p. 27, 2003.
- [20] M. Levoy and P. Hanrahan, "Light Field Rendering," *Proc. ACM SIGGRAPH '96*, pp. 31-42, 1996.
- [21] M. Li, H. Schirmacher, M. Magnor, and H.-P. Seidel, "Combining Stereo and Visual Hull Information for On-Line Reconstruction and Rendering of Dynamic Scenes," *Proc. IEEE Multimedia and Signal Processing*, pp. 9-12, 2002.
- [22] B. Lucas and T. Kanade, "An Iterative Image Registration Technique with an Application to Stereo Vision," *Proc. Defense Advanced Research Projects Agency IU Workshop*, pp. 121-130, 1981.
- [23] S. Marschner, "Inverse Rendering for Computer Graphics," PhD thesis, Cornell Univ., 1998.
- [24] T. Matsuyama and T. Takai, "Generation, Visualization, and Editing of 3D Video," *Proc. First Int'l Symp. 3D Data Processing Visualization and Transmission (3DPVT '02)*, p. 234ff, 2002.
- [25] W. Matusik, C. Buehler, R. Raskar, S.J. Gortler, and L. McMillan, "Image-Based Visual Hulls," *Proc. ACM SIGGRAPH '00*, pp. 369-374, 2000.
- [26] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A Data-Driven Reflectance Model," *ACM Trans. Graphics (Proc. ACM SIGGRAPH '03)*, vol. 22, no. 3, pp. 759-769, 2003.
- [27] W. Matusik and H. Pfister, "3D TV: A Scalable System for Real-Time Acquisition, Transmission, and Autostereoscopic Display of Dynamic Scenes," *ACM Trans. Graphics (Proc. ACM SIGGRAPH '04)*, vol. 23, no. 3, pp. 814-824, 2004.
- [28] K. Nishino, Y. Sato, and K. Ikeuchi, "Eigen-Texture Method: Appearance Compression and Synthesis Based on a 3D Model," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1257-1265, Nov. 2001.
- [29] B.-T. Phong, "Illumination for Computer Generated Pictures," *Comm. ACM*, pp. 311-317, 1975.
- [30] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in C++*. Cambridge Univ. Press, 2002.
- [31] R. Ramamoorthi and P. Hanrahan, "A Signal-Processing Framework for Inverse Rendering," *Proc. ACM SIGGRAPH '01*, pp. 117-128, 2001.
- [32] H. Rushmeier, G. Taubin, and A. Guézic, "Applying Shape from Lighting Variation to Bump Map Capture," *Proc. Eurographics Workshop Rendering*, pp. 35-44, June 1997.
- [33] Y. Sato, M.D. Wheeler, and K. Ikeuchi, "Object Shape and Reflectance Modeling from Observation," *Proc. ACM SIGGRAPH '97*, pp. 379-388, 1997.
- [34] M. Waschbüsch, S. Würmlin, D. Cotting, F. Sadlo, and M. Gross, "Scalable 3D Video of Dynamic Scenes," *Proc. Pacific Graphics*, pp. 629-638, 2005.
- [35] A. Wenger, A. Gardner, C. Tchou, J. Unger, T. Hawkins, and P. Debevec, "Performance Relighting and Reflectance Transformation with Time-Multiplexed Illumination," *ACM Trans. Graphics (Proc. ACM SIGGRAPH '05)*, vol. 24, no. 3, pp. 756-764, 2005.
- [36] S. Würmlin, E. Lamboray, O.G. Staadt, and M.H. Gross, "3D Video Recorder," *Proc. IEEE Pacific Graphics*, pp. 325-334, 2002.

- [37] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs," *Proc. ACM SIGGRAPH '99*, pp. 215-224, Aug. 1999.
- [38] Y. Yu and J. Malik, "Recovering Photometric Properties of Architectural Scenes from Photographs," *Proc. ACM SIGGRAPH '98*, pp. 207-218, 1998.
- [39] R. Zayer, C. Rössl, and H.-P. Seidel, "Discrete Tensorial Quasi-Harmonic Maps," *Proc. Int'l Conf. Shape Modeling and Applications*, pp. 276-285, 2005.
- [40] R. Zhang, P.-S. Tsai, J. Cryer, and M. Shah, "Shape from Shading: A Survey," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 21, no. 8, pp. 690-706, Aug. 1999.
- [41] T. Zickler, S. Enrique, R. Ramamoorthi, and P.N. Belhumeur, "Reflectance Sharing: Image-Based Rendering from a Sparse Set of Images," *Proc. Eurographics Symp. Rendering*, pp. 253-264, 2005.
- [42] C.L. Zitnick, S.B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski, "High-Quality Video View Interpolation Using a Layered Representation," *ACM Trans. Graphics (Proc. ACM SIGGRAPH '04)*, vol. 23, no. 3, pp. 600-608, 2004.



Christian Theobalt received the MSc degree in artificial intelligence from the University of Edinburgh, Scotland, and the Diplom (MS) degree in computer science from Saarland University, Saarbrücken, Germany, in 2000 and 2001, respectively. From 2001 to 2005, he worked as a PhD student and research associate in Hans-Peter Seidel's Computer Graphics Group at MPI Informatik. In 2005, he received the PhD (Dr.-Ing.) degree from Saarland University. He is currently a postdoctoral researcher at the Max-Planck-Institut (MPI) Informatik, Germany. His research interests include free-viewpoint and 3D video, markerless optical motion capture, 3D computer vision, and image-based rendering. He is a member of the IEEE.



Naveed Ahmed received the BS and MSc degrees in computer science from the University of Karachi and Saarland University in 2001 and 2004, respectively. He is currently working as a PhD student in the Computer Graphics Group of the Max-Planck-Institut (MPI) Informatik, Saarbrücken, Germany. His research interests include video-based rendering and relighting.



Hendrik Lensch studied computer science at the University of Erlangen and the Royal Institute of Technology (KTH) in Stockholm, receiving the diploma degree in 1999. He then worked as a PhD student and research associate in Hans-Peter Seidel's Computer Graphics Group at the Max-Planck-Institut Informatik. In 2003, he received the doctorate degree from Saarland University. From 2004 to 2006, he was appointed to a visiting assistant professor position at Stanford University in the Computer Graphics Lab. He is currently the head of the research group on general appearance acquisition of the Max Planck Center for Visual Computing and Communication (Saarbrücken/Stanford) at the Max-Planck-Institut Informatik in Saarbrücken. His research interests include 3D appearance acquisition, BRDF reconstruction, image-based rendering, and computational imaging.



Marcus Magnor received the BA and MS degrees in physics from the University of Würzburg, Germany, and the University of New Mexico, respectively, and the PhD degree in electrical engineering from the University of Erlangen, Germany. For his postgraduate studies, he joined the Computer Graphics Lab at Stanford University before establishing his own independent research group at the Max-Planck-Institut (MPI) Informatik, Germany. In 2005, he received the *venia legendi* in computer science from Saarland University. He is currently a full professor of computer graphics at the Technical University of Braunschweig, Germany. His research interests encompass the entire visual information processing pipeline (visual computing). His recent research topics include video-based rendering, 3D-TV, and realistic, real-time rendering. He is a member of the IEEE.



Hans-Peter Seidel is the scientific director and chair of the computer graphics group at the Max-Planck-Institut (MPI) Informatik and a professor of computer science at Saarland University. He has published some 200 technical papers in the field and has lectured widely. He has received grants from a wide range of organizations, including the German National Science Foundation (DFG), the German Federal Government (BMBF), the European Community (EU), NATO, and the German-Israeli Foundation (GIF). In 2003, he was awarded the "Leibniz Preis," the most prestigious German research award, from the German Research Foundation (DFG). He is the first computer graphics researcher to receive this award.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.