

Planned Sampling of Spatially Varying BRDFs

Hendrik P.A. Lensch¹, Jochen Lang¹, Asla M. Sá² †, and Hans-Peter Seidel¹

¹ MPI Informatik, Saarbrücken, Germany

² Instituto Nacional de Matemática Pura e Aplicada - IMPA, Brazil

Abstract

Measuring reflection properties of a 3D object involves capturing images for numerous viewing and lighting directions. We present a method to select advantageous measurement directions based on analyzing the estimation of the bi-directional reflectance distribution function (BRDF). The selected directions minimize the uncertainty in the estimated parameters of the BRDF. As a result, few measurements suffice to produce models that describe the reflectance behavior well. Moreover, the uncertainty measure can be computed fast on modern graphics cards by exploiting their capability to render into a floating-point frame buffer. This forms the basis of an acquisition planner capable of guiding experts and non-experts alike through the BRDF acquisition process. We demonstrate that spatially varying reflection properties can be captured more efficiently for real-world applications using our acquisition planner.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism Virtual Reality I.4.1 [Computer Vision]: Digitization and Image Capture, Reflectance

1. Introduction

In the field of 3D object acquisition progress has been made both in the area of geometry and appearance acquisition. Appearance or reflection properties are in most approaches measured by capturing a number of samples of the BRDF of the object. The samples are commonly acquired by a sensor (a digital camera in our set-up) and a point-light source. One pair of light source and camera position (called a view collectively in the remainder of this paper) captures a single reflectance sample for each point that is visible and lit.

A number of researchers have built special gantries to perform a robot controlled dense sampling of the reflection properties^{8, 28, 10, 30}. Others position the camera and the light source manually^{23, 25, 19, 20}.

The basic question for both, the automatic and the manual approach is: How to sample the reflection properties in an efficient way? The acquisition of reflection properties needs to be planned in order to measure efficiently, failing to plan may result in insufficient data for the modeling task or lead to highly redundant over-sampling. Measurement is




Figure 1: Comparison of Spatially Varying BRDF Models. The model¹⁹ on the left contains holes in the BRDF due to undersampling. The model on the right obtained from the same number of views suggested by our planner samples the surface evenly.

typically an involving task and efficiency in the process is of paramount importance.

In this paper we present a method that assesses the uncertainty in the parameters of a Lafortune model¹⁶ (The underlying method is however also applicable to other parametric models and may be adapted to non-parametric models as well.) Based on this uncertainty measure we develop an acquisition planning algorithm that computes from where to

† This work was conducted at MPI Informatik.

sample next in order to minimize the uncertainty in the parameters, i.e., where to place the camera *and* the light source with respect to a set of previously acquired views. For 3D objects we have to evaluate and combine the predicted uncertainty of each single surface point. A good set of views will measure each point on the surface several times with varied viewing and lighting directions sampling a highlight at each point. The view planning is influenced by a number of further constraints including the 3D shape. The shape limits the number of visible and lit surface points in a view.

One of the goals of acquisition planning is to perform measurements efficiently. Time spent on the planning itself therefore has to be reasonable. We compute the uncertainty measure in modern graphics hardware with floating-point frame buffers. The evaluation is performed directly on the texture atlas of the object.

The measurement theory behind our approach is well established in other fields; physicists and other natural scientists apply it quite routinely to their measurement tasks. Our contribution in this area is to adapt some of the natural sciences' measurement theory to the task of measuring the BRDF for computer graphics. Our paper makes three main contributions:

- the definition of a function to measure the reduction in uncertainty added by one view (camera and light source position),
- a view planning algorithm that combines this function with geometric constraints imposed by a 3D object to predict the next best view for efficient measurement, and
- a hardware-accelerated implementation for evaluation of the objective function directly on the texture atlas.

In the next section we discuss related work before we present an overview over the acquisition planning and the measurement process in Section 3.

2. Related Work

Work related to the automated acquisition of reflection properties of complete objects can be found in different fields including computer graphics, computer vision, robotics and visual metrology. We start our review with a brief summary of work in computer vision, followed by a discussion of automatic scanning of 3D models including some theoretical issues, and we conclude our review with work on BRDF acquisition and representation in computer graphics.

The task of exploring unknown spatially-varying BRDF of an object relates directly to viewpoint control in computer vision, e.g., in minimizing uncertainty of 3D object representations⁴⁵ and in scene exploration¹⁵. Related to BRDF acquisition are also visual metrology tasks which are reviewed by Tarabanis and Tsai⁴². In metrology the planning task is to position a sensor to satisfy some sensing quality criterion, e.g., work by Cowan and Kovcsi⁶ and Mason and

Grün²⁶. The quality criterion of interest in our work is the certainty in the acquired BRDF. We aim at achieving high quality by choosing advantageous viewpoints of camera and light source positions.

The placement of a camera relative to an object has been planned for geometric model acquisition with a robotic facility^{36,17} and with an automated commercial scanner³³. In practice, view planning is often started after an initial rough acquisition of the object's geometry from pre-set viewpoints, as in the geometry model acquisition system by Reed and Allen^{36,37}. The task in this situation is to fill holes in the existing model stemming from tight visibility restrictions. Filling holes is an example where an automated planner can be very beneficial.

The incremental *Next Best View* planning strategy of Whaitte and Ferrie⁴⁵ is closely related to ours. Their strategy explores the *geometry* of a 3-D model with *a priori* unknown shape. They apply a synthesis approach which is based on a probabilistic model of an object's geometry to be explored. The approach minimizes uncertainty of a parametric object model. Objects are represented by sets of superquadrics. Using an active sensing strategy a next view is selected. The selected view minimizes the current uncertainty of the object model. The algorithm is applied to explore the environment of a robotic agent enabling object recognition and manipulation.

An alternative approach to view planning is to delegate the (computationally complex) task to a human but to provide real-time feedback. This has been demonstrated successfully by Rusinkiewicz et al.³⁸ for geometry scanning of objects. Planning the acquisition of reflectance properties in this fashion is however impossible since humans can generally not reason about the four-dimensional BRDF on the surface of complex objects in real-time.

In computational geometry, visibility in polygonal environment is considered⁵ with applications in geographical data processing, security and military²². Acquiring realistic reflectance properties of a 3D object requires imaging its complete surface. This completeness constraint restricts the set of possible solutions in acquisition planning. Finding the minimal set of views which cover the complete surface of an object is a NP-hard problem⁴³. A planned set of views may still fail in practice due to positioning or modeling errors and off-line plans need to consider uncertainty to ensure coverage of the complete surface when actually executed^{43,40}.

In computer graphics, the imaging of visible surfaces is relevant in radiosity, ray-tracing, scene walk-throughs and texturing of surfaces⁴¹ as well as for image-based rendering^{13,4,44}. In the BRDF measurement and representation field, quite a number of articles have been published but, to the best of our knowledge, so far only Kay and Caelli¹⁴ have investigated the dependency between measurement and the quality of the estimated BRDF. McAllister³⁰ computes




Figure 2: View planning interacts with the entire pipeline of appearance measurements: An optimal view is proposed and captured. Manual placement of camera and light source requires registration with the 3D object. Visibility and shadows are computed and the data is resampled. From the resampled data of all views the BRDF parameters per cluster are updated which again influence the planning of the next view.

the number of samples required to densely cover the hemisphere above a point with a light source. Ramamoorthi and Hanrahan³⁵ examine the problem of global inverse illumination within a signal processing framework. Existing BRDF measurement approaches can be classified by the number of images necessary and the generality of the estimated BRDF model.

A family of methods measure reflectance properties by performing a very dense sampling. Some of them are designed for flat surface samples only^{7,30} while others can deal with 3D objects^{8,27,28,29,10}. Since a dense sampling allows to render directly from the measured data, the reflection properties are unrestricted in those approaches. The quality of the outcome when no BRDF model is fit mainly depends on the sampling density.

A variety of techniques estimate the appearance of an object from a sparse set of images. In many of these approaches the specular part of the BRDF is restricted to be constant over a small patch or even over the entire object. Measurements in these approaches are typically performed either by using a point light source^{39,25,24,31,19,20,21} or by performing inverse global illumination^{46,11,1,35,32}. Lensch et al.^{19,20} and Li et al.²¹ demonstrated how to estimate a varying specular part from a sparse set of views. In both approaches clustering techniques are applied to the BRDF parameters over the object's surface. In this paper, we also follow a sparse sampling approach but at the same time assuring the quality of the model. The fully spatially varying model in our work is identical to the one employed by Lensch et al.²⁰.

Next, we introduce our method which is able to select a sparse set of measurements and ensures the quality of the BRDF model at the same time.

3. Acquisition Loop

The measurement of reflection properties is executed as a number of successive steps which are shown in Figure 2. Prior to the acquisition of reflectance properties the 3D geometry of the object has to be acquired. The acquisition starts with the planning of the first view. Each new view is planned based on the current estimate of the BRDF parameters and the visibility and shadowing constraints imposed by the 3D geometry of the object.

In the second step the planned view is acquired. Next, the

view is registered with the 3D object¹⁸, since the real camera and light source position may deviate from the proposed view. The recovered positions are used to determine the regions of the object's surface which are visible and lit. The valid pixels are resampled into a texture atlas.

For the planning a coarse texture atlas is used to speed up the process. Local viewing and lighting direction are computed per pixel. From the resampled data a new set of BRDF parameters is estimated. As the number of measurements for a single point are too few to obtain credible BRDF parameters, clusters of points are used to fit the parameters. Points can be clustered either based on their diffuse color or using the method by Lensch et al.²⁰. The estimated spatially varying BRDF parameters are then used in the next execution of the planning step.

After the capturing process is complete, all measurement data is resampled again using a high resolution texture atlas and a final spatially varying reflection model is estimated.

Planning is performed by minimization of an objective function that takes into account the previously acquired data, the geometry of the object and the currently estimated BRDF parameters. The objective function is based on co-variance matrices as an uncertainty measure. The co-variance matrices are compact and summarize all necessary information about the views acquired so far. Their storage cost and the computation time of planning algorithm is constant and independent of the number of acquired views. In particular, we do not have to store each local viewing and lighting direction per pixel for each measurement.

In the next section we detail the relationship between the co-variance matrix and measurement uncertainty with an object consisting of a single pixel as a tutorial example.

4. One-Pixel Objects

Measuring the BRDF of a single point on an object is already a task which involves some effort. It is necessary to understand how measurements of the reflectance of a single pixel influence the reliability of parameters of the BRDF model we are going to fit. We briefly summarize background material on parameter estimation and measurement theory related to our fitting approach. We conclude that not all possible measurements contribute in the same way to the quality of the fitted model parameters.

For the purpose of the discussion in this and the next section, the prior information I for our measurement task is that we would like to estimate a Lafortune model f_r with one isotropic specular lobe (see Equation 1) for an object consisting of one-pixel.

$$f_r(\beta; \omega) = \rho_d + (c_x \omega_i \omega_{o_x} + c_y \omega_i \omega_{o_y} + c_z \omega_i \omega_{o_z})^N \quad (1)$$

where ρ_d denotes the diffuse reflectance, N the exponent of the specular lobe, c_x , c_y and c_z the weighting coefficients of the dot product between ω_i and ω_o ; ω_i is the incident light direction and ω_o the exitant lighting direction. In the following, we denote these parameters collectively as β . The model $M(\beta, \omega)$ calculates the reflectance for a given ω_i and ω_o (collectively ω). Equation 2 shows the standard regression problem for m measurements with an additive error term ϵ resulting in m reflectance samples R_m . For the purpose of this discussion, we assume that the error terms are independent and identically distributed (iid) samples from a distribution centered at 0.

$$R_m = M(\beta, \omega_m) + \epsilon_m \quad (2)$$

A principled approach to solve this model fitting task is by employing Bayes theorem

$$P(M|DI) = \frac{P(M|I)P(D|M)}{P(D|I)}.$$

Bayes theorem describes how to obtain the posterior probability distribution of the reflectance model $P(M|DI)$. It depends on our prior belief of possible model parameters $P(M|I)$, the measurement or predictive probability of a measurement $P(D|I)$ with acquired data D and our model of the measurement process $P(D|M)$ which here is Equation 2. Measurements are a principled way to change one's prior beliefs. If the observations provide strong evidence, the data term dominates while with a lack of evidence the prior remains unchanged. The certainty in the model is described by the full distribution $P(M|I)$. The distribution is in the dimensions of the model and requires a summary for interpretation. The most probable parameter value and confidence intervals are common summaries. (See, e.g., Bretthorst² or Hastie et al.¹² for a more complete introduction to Bayesian model estimation).

An approach to find the most probable model is minimization. The sum of squared errors

$$Q = \sum_m (R_m - M(\beta, \omega_m))^2 \quad (3)$$

is the most common error measure to minimize. This error measure coincides with the most probable model in the Bayesian approach given the model is linear in the parameters β , the noise ϵ in the measurements is Gaussian and our prior beliefs are uninformative (flat priors)¹². Under these circumstances, the uncertainty in the most probable model parameters depend linearly on the co-variance matrix \mathbf{CoV} . The co-variance matrix \mathbf{CoV} is the inverse of the Hessian matrix \mathbf{H} with entries $\mathbf{H}_{i,j} = \frac{\partial^2 Q}{\partial \beta_i \partial \beta_j}$ (see Appendix A for the

derivatives). The singular values σ of the co-variance matrix define the length of the major axes of the hyperellipsoid of a given Q . In linear models this hypervolume bounded by the hyperellipsoid for a given quadratic error measure Q is directly related to the posterior probability.

In the non-linear Lafortune reflectance model the simple relationship between the least-squares residual Q and the uncertainty in the parameters does not hold. The most popular way to proceed is to employ a non-linear least-square solver to minimize Q but analyze the fit with the co-variance matrix \mathbf{CoV} . The co-variance matrix is only strictly valid for linear models, however, employing \mathbf{CoV} is justifiable if Q is well approximated by a quadratic near the minimum³⁴. We performed some Monte-Carlo bootstrap analysis^{9,3} in order to confirm the validity of the linear approximation when fitting the non-linear Lafortune model. We do not report the details here but in summary, our conclusion is that the non-linearity of the Lafortune reflectance model prevents us from stating confidence intervals based on the linear approximation. However, the co-variance matrix is a good indicator of parameter uncertainty and if we would like to obtain measurements in a way such that we are most confident in the estimated parameters, minimizing the co-variance matrix is a sensible strategy. This conclusion is also consistent with the reasoning of Whaite and Ferrie⁴⁵. The co-variance matrix of the Lafortune model depends only on the chosen incident and exitant light direction given a fixed estimate $\hat{\beta}$. A recipe of how to choose the light and viewing directions for the one-pixel object is described in the next section.

5. Uncertainty Minimization

The uncertainty in the estimated parameters for a set of views is minimum if the co-variance matrix is minimal. If we are exploring unknown reflectance properties, we only learn our model parameters as we acquire new views, gaining information incrementally. The certainty gain of a view is therefore the reduction in the co-variance matrix from the previous view to the current one and hence, the objective function

$$F = \|\mathbf{CoV}(\hat{\beta}_v, \omega_{1,\dots,v})\| - \|\mathbf{CoV}(\hat{\beta}_v, \omega_{1,\dots,v+1})\|. \quad (4)$$

We would like to maximize the certainty gain, i.e., we have to maximize F . A greedy strategy selects after each view v the one which maximizes the expected gain of the next view $v+1$. This greedy strategy is optimal in the linear case with fixed estimates $\hat{\beta}$ because then the co-variance is independent of the order of views. In our scenario this approximation becomes more appropriate as the certainty in the estimates increases.

The selection strategy must also deal with a rank-degenerate Hessian matrix \mathbf{H} during the views $v = 1 \dots L$, where L is the number of parameters in the Lafortune model f_r , in our case $L = 4$. These initial views have to be chosen in order to arrive at a small $\|\mathbf{CoV}\|$ after view $v = L$. While

$\|\mathbf{CoV}\| = \infty$ for $v < L$, we can calculate the Pseudo-Inverse of the Hessian matrix instead. Figure 3(a) shows the singular values σ_k of the pseudo-inverse which approximately increase exponentially with k . We maximize the information gain based on this pseudo co-variance until the Hessian matrix reaches full rank. The gain can be considered infinite with each rank gained, i.e., a choice of $F(k \leq L) = \text{Inf} - \sigma_k$ suggests itself. We illustrate this strategy in Figure 3 for the one-pixel object. (We pick $\text{Inf} = 10^9$ and the parameters $\beta = (\rho_d = 0.3, N_s = 10, c_x = c_y = -0.8, c_z = 0.8)$ for the Lafortune model (Equation 1)). A better continuation of the objective function towards the first view can be obtained with $F(k \leq L) = 10^{L-k} * (\text{Inf} - \sigma_k)$ which is shown in Figure 3(b). The difference in the two considered choices is only of importance in multi-pixel objects when the information gain at one pixel has to be compared to the information gain at another pixel. Under these circumstances, the exponentially increasing function will strongly favour low rank updates over high rank updates.




Figure 3: Objective Function F for Infinite Co-Variance. The singular values σ_k of the \mathbf{CoV} show an exponential increase with k . Figure 3(b) shows two choices for F when $k \leq L$: $F(k \leq L) = 10^{L-k} * (\text{Inf} - \sigma_k)$ in red and $F(k \leq L) = \text{Inf} - \sigma_k$ ($\text{Inf} = 10^9$) in blue.

5.1. Maximization

The task of view planning is to maximize the objective function F in Equation 4. Our observation is that the objective function is partially smooth depending on the visibility of a given view. Globally, it can have many discontinuities due to shadows and visibility. In order to maximize the objective function we apply a two-step procedure with randomization:

- Search for the best pair of initial positions of camera and light source on a discretized sphere. Randomize the orientation of the discretization before each new view.
- Use a local continuous optimizer to improve the discrete solution found.

The random rotation of the sphere discretization improves the discrete search, eventually evaluating all views. This procedure allows us to use a coarse discretization in the discrete

search achieving acceptable coverage at least over multiple views.

We have tested two downhill non-gradient-based optimizers: the Simplex method and Powell's method with non-gradient based linear search³⁴. The Simplex method is initialized with the starting value of the discrete search plus random points within its neighborhood. Although both methods may converge to local maxima, we still achieve good results due to the discrete initialization. In Figure 4, we show the norm of the co-variance matrix achieved by maximizing F in Equation 4 with both, the Simplex and Powell's method. We have observed that Powell's method requires typically less function evaluations but achieves slightly worse results than the Simplex method. Figure 4 also compare the behavior of the optimization with the 2-norm $\|\mathbf{CoV}\|_2$ versus the Frobenius norm $\|\mathbf{CoV}\|_{Frob}$. The difference is negligible due to the exponential rate of decay in the size of the singular values of the co-variance matrix. Besides planning




Figure 4: Norm of the Co-Variance Matrix during Optimization. Green $\|\mathbf{CoV}\|_2$ and red $\|\mathbf{CoV}\|_{Frob}$ are obtained with Powell's method, while blue $\|\mathbf{CoV}\|_2$ and cyan $\|\mathbf{CoV}\|_{Frob}$ are obtained by the simplex method. The optimization is applied to the one-pixel object.

the next best view, the objective function can be employed to determine when to stop acquiring more views. Observe that the norm decays roughly exponentially as the number of samples of the one-pixel object increases. It indicates that one can stop the optimization after the gain in confidence is below a pre-defined threshold.

6. Multi-Pixel Objects

We are now generalizing our insights from the one-pixel object to 3D objects. The reflectance of real-world 3D objects cannot be modeled with sufficient accuracy by a single point. The complete surface area of a 3D object is not visible from any single given camera and light source position. Thus, the geometric shape influences the objective function and plays an important role in the uncertainty minimization.

6.1. Homogeneous vs. Spatially Varying BRDFs

Calculation of model uncertainty for a real 3D object requires the computation of the the Hessian matrices \mathbf{H}_i considering all previous measurements at each point i on the




Figure 5: *Texture Atlas.* Lighting and the derivatives of an isotropic Lafortune BRDF (see Appendix A) are computed in a texture atlas.

surface. Assuming all points on the surface have identical reflectance then the overall co-variance matrix of all surface points is $\|\mathbf{CoV}\|_{Homogenous} = \|(\sum_i \mathbf{H}_i)^{-1}\|$. In this norm there is no relationship between samples and surface location. It does not ensure an even sampling of the surface and thus is only applicable if the object's surface consists of a single homogeneous material without the slightest spatial variation.

$$\|\mathbf{CoV}\|_{obj} = \sum_i \|\mathbf{H}_i^{-1}\|. \quad (5)$$

An approach more applicable to real-world 3D objects is to treat each surface point individually, assuming a different BRDF for each pixel. The uncertainty measure in this case is Equation 5. This uncertainty measure can only be decreased by imaging each surface point repeatedly. This is required if the goal of the measurement process is to model each surface point with a different BRDF but as well if surface points are eventually to be clustered and represented by a smaller set of BRDFs.

6.2. Real-World Constraints

While performing a real measurement one has to take into account a number of additional constraints. The planning algorithm must prevent placement of light source and camera at the same location. In order to achieve this, we simply set the objective function to zero in these cases. Furthermore, the reliability of collected samples is not independent of the viewing angles for real measurements. Samples at grazing angles are typically hard to measure since the registration of the 3D object with the 2D image is not perfect. We follow the approach by Lafortune et al.¹⁶ and weight the influence of each sample by the cosine of the angle between the surface normal and the viewing direction and between the normal and the lighting direction. The weight is set to zero for all points where one of the angles is larger than 80 degrees.

7. Implementation

We have defined the objective function F in Section 5 and constraints in the previous section. Here, we detail the implementation of an efficient optimization which results in

a usable view planning algorithm. The view planning has to simulate the next view in order to evaluate the objective function. We calculate the norm of the co-variance matrices for each point on the surface. This computation has to be very efficient since it is executed several hundred times during one step of the optimization. We achieve this efficiency by exploiting newly available graphics hardware with floating-point precision frame buffers.

7.1. Texture Atlas

Since all quantities need to be computed for all surface points we represent the object's surface by a texture atlas (see Figure 5(b)). We construct the texture atlas from the 3D mesh. All subsequent calculations are performed directly on the texture atlas. Using a vertex program it is very easy to perform calculations on the texture atlas: the final vertex position is set to the texture coordinates of the vertex while the original vertex position is used for other computations, e.g. lighting and shadowing.

7.2. Visibility and Shadows

The texture atlas shows all points on the object's surface at the same time, but the derivatives may only be computed for those points which are visible and lit for a given view. We accomplish this by first computing two depth maps of the 3D object, one depth map for the camera view and one from the view of the light source. As in traditional shadow mapping the point is visible and lit if the transformed pixel has the same or smaller depth than the corresponding points in the depth maps. Figure 5(c) shows the resulting texture atlas of all visible and lit points of the corresponding view Figure 5(a).

7.3. Derivatives and Matrix Norms

Given the 3D mesh and some BRDF Parameters we setup a fragment program which computes the derivatives of the BRDF model with respect to its parameters for the valid texels on the graphics board (see Figure 5(d)–5(g)). We take the spatially varying parameters obtained by the clustering into account. The calculated derivatives are then downloaded from a floating-point frame buffer to main memory. Hessian




Figure 6: Comparison of Planner and Human Expert. First and third rows show results with the planner, while results obtained by the human expert are shown in the second and fourth row. (Images shown are after every 3rd view.) Notice how the planner selects samples to cover the object's surface evenly. The human expert acquires redundant samples of some surface areas (front) while other areas remain undersampled (back and bottom). The images show $\log(\|\mathbf{CoV}\|)$ color-coded in matlab jet style. Blue means high confidence in the estimated BRDF parameters while the uncertainty increases towards red. Black regions have not been updated to full rank so far.

matrices for the current view are calculated and added to the accumulated Hessian matrices of previous views. The result is inverted using singular value decomposition to obtain the norms of the co-variance matrices. The final value of the ob-

Task	$\partial f/\partial\beta$	Download	\mathbf{H}	SVD	Total
Time [s]	0.119	0.477	0.021	0.523	1.138

Table 1: Time consumed for computing the derivatives, to download the results to main memory, to add to the Hessian matrices and to perform the remaining calculations (SVD) in software in order to evaluate the objective function of one view on a 512x512 texture atlas. Note that on average only one fourth of the 147444 valid pixels were visible and lit.

jective function is then computed as the sum of the objective functions of each pixel. All software computations are done only at those pixels in the texture atlas which are visible and lit. A total of approximately 440 evaluations of the objective function are performed in the optimization of one view.

Table 7.3 lists how much time is spent for each of the computation steps. A considerable amount of time is unfortunately consumed by downloading the frame buffer. To save bandwidth we currently use a monochromatic isotropic Lafortune BRDF model with one lobe (4 parameters) for the

view planning. The method can however be easily extended to work with more complex or multi-lobe models.

8. Measurement Results

We compare the views selected by our method to the views selected by a human expert prior to our method. The 3D object for which the spatially varying BRDF is acquired, are the angels shown in Figure 1. The reflectance model of the

Rank k	# Pixels Planner	# Pixels Expert	$\bar{\sigma}_{L-k+1}$ Planner	$\bar{\sigma}_{L-k+1}$ Expert
0	10	857	-	-
1	300	2013	0.06587	0.04482
2	1755	973	2.3804	0.9823
3	1412	583	39.170	75.075
4	5767	4818	745.16	2308.94

Table 2: Comparison of $\|\mathbf{CoV}\|$ Obtained by Planner and Human Expert after 27 views. Shown are the singular values averaged over all pixels. The planner acquired more pixel with higher rank and higher confidence. (Note the considerably smaller singular value in the last row).

angels has previously been captured with a set of 27 views (not shown). Figure 9 shows the images acquired with our planner. The planner selects camera and light source position in order to collect samples of each surface point under

various angles. A comparison of the resulting uncertainty $\log(\|\mathbf{CoV}\|)$ shows that the planner covers the surface of the object quite evenly (see Figure 6). The human expert samples the front of the model frequently but misses areas on the rear, bottom and sides of the model (see Figure 1). The even sampling with our planner is even more apparent by studying the smallest singular value σ_1 of the \mathbf{CoV} in Figure 8. The average singular values are compared in Table 2. The planner increases the rank of the \mathbf{CoV} per pixel evenly. Table 2 shows that after 27 views the \mathbf{CoV} has reached full rank at 62.4% of pixels. The lower rank (pseudo) \mathbf{CoV} have already mostly gained rank 2 or 3. The \mathbf{CoV} obtained by the human expert are mainly either, rank 4 (52.1%) or still rank 1 and rank 0 (unobserved). Additionally, the norms $\|\mathbf{CoV}\|$ with high rank are considerably larger than the number obtained with the planner, i.e., the measurements result in higher uncertainty.

Highlights have to be observed in order to update the rank of one pixel beyond rank 1 because three of the derivatives evaluate to nearly zero for non-highlight directions (Appendix A and Figure 5). For a human it is hard to keep track of the observed highlight areas and to reason on how to place the camera and the light source in order to observe a highlight at a specific surface region. The planner automatically considers this by maximizing the proposed objective function (see Figure 9).

The objective function evaluated after each view is shown in Figure 7. The series with the planner has been continued in simulation to 50 views. Note, that while local minima are encountered by the planner, it successfully reduces the error measure in each step and the steps decrease approximately exponentially as expected. (Note, that $\log(F_{obj})$ in Equation 5 is a difference and the numbers have to be accumulated for absolute values.) The optimization per view




Figure 7: Objective Function for Angel Acquisition. The red curve is obtained with the uncertainty minimization planner while the blue dashed curve is the objective function evaluated for the series taken by a human expert.

took less than 2 minutes, computed on a 2.4GHz Pentium PC with an ATI Radeon 9700. The computations were carried out on a texture atlas with resolution 128x128. This time is constant and does not depend on the number of previously




Figure 8: Smallest Singular Value: Comparison of Planner and Human Expert. Top row show results with the planner, while the bottom row shows results obtained by the human expert. Notice how the planner selects samples to cover the object's surface evenly. The human expert acquires redundant samples of some surface areas (front) while other areas remain unobserved (back and bottom). The images show $\log(\|\sigma_1\|)$ after 27 views color-coded in matlab jet style.

seen views, since the information is accumulated in the per-pixel Hessian matrix \mathbf{H} .

In summary, the planner helps to sample the angels' surface more evenly resulting in higher certainty in the BRDF parameters. The number of views (27) is generally too low for being able to judge the fit of BRDF parameters at each pixel. The results however are visually already quite pleasing over the entire surface (Figure 1). The simulation suggests that after 48 views on more than 90% of the surface the \mathbf{CoV} would reach full rank. Over 90% of the surface have already been imaged once after 5 views (> 99% after 10 views) in the actual measurements. The actual capturing process has been slightly slowed down by the planning but mainly because view registration has to be performed during acquisition when using the planner. The registration also revealed that some additional cues would be very helpful in setting up camera and light source. A low quality real-time registration algorithm would integrate nicely into the planner.

9. Conclusion

We have presented a novel method to analyze and plan the acquisition of realistic reflection models of 3D objects. Central part of the method is a measure of uncertainty which allows one to assess the quality of the sampling so far and to select from where to view the object. This uncertainty measure can be evaluated efficiently in graphics hardware with floating point precision and has been integrated into a view planner for BRDF acquisition. The performance of the planner compares favorably to the view selection by a human expert. It is extremely hard for humans to reason about the 4D BRDF on the surface of 3D objects. Consequently, it is very difficult for them to select good camera and light source positions in order to obtain a high-quality reflectance model. Our view planning algorithm can assist experts and




Figure 9: Views Planned for the Angels. Our algorithm positions both the camera and the light source to minimize BRDF uncertainty. As a result a highlight is observed at each surface point.

enables novices to measure the BRDF of 3D objects. Automatic measurements acquiring densely sampled BRDFs may also benefit from this method since the same quality can be achieved with less planned views. Since the automatic setup does not require a registration step acquisition of one view and planning of the next may be run in parallel.

Appendix A: Hessian Matrix

The Hessian matrix for Lafortune model (Equation 1) is given by the following equation:

$$\frac{\partial^2 Q}{\partial \beta^2} = -2 \left(\frac{\partial L}{\partial \beta} \right)^T \frac{\partial L}{\partial \beta}. \quad (6)$$

The derivatives $\frac{\partial L}{\partial \beta}$ of the single specular lobe Lafortune model are $\frac{\partial L}{\partial \rho_d} = 1$,

$$\frac{\partial L}{\partial c_{xy}} = (c_{xy} \omega_i \omega_{o_x} + c_{xy} \omega_y \omega_{o_y} + c_z \omega_z \omega_{o_z})^{N-1} N (\omega_i \omega_{o_x} + \omega_y \omega_{o_y}),$$

$$\frac{\partial L}{\partial c_z} = (c_{xy} \omega_i \omega_{o_x} + c_{xy} \omega_y \omega_{o_y} + c_z \omega_z \omega_{o_z})^{N-1} N (\omega_z \omega_{o_z}), \text{ and}$$

$$\frac{\partial L}{\partial N} = (c_{xy} \omega_i \omega_{o_x} + c_{xy} \omega_y \omega_{o_y} + c_z \omega_z \omega_{o_z})^N N \cdot \log(c_{xy} \omega_i \omega_{o_x} + c_{xy} \omega_y \omega_{o_y} + c_z \omega_z \omega_{o_z}).$$

References

1. S. Boivin and A. Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. *Computer Graphics Proceedings, Annual Conference Series*, pages 107–116. ACM SIGGRAPH, August 2001.
2. G.L. Bretthorst. *Bayesian spectrum analysis and parameter estimation*. Springer-Verlag, New York, 1988.
3. J. Carpenter and J. Bithell. Bootstrap confidence intervals: when, which, what? a practical guide for medical statisticians. *Statistics in Medicine*, 19:1141–1164, 2000.
4. J.X. Chai, S.C. Chan, X. Tong, and H.Y. Shum. Plenoptic sampling. In *Computer Graphics, Annual Conference Series*, New Orleans, USA, Jul 2000. ACM SIGGRAPH.
5. V. Chvátal. A combinatorial theorem in plane geometry. *J. Combin. Theory (B)*, 18:39–41, 1975.
6. C. K. Cowan and P. D. Kovesi. Automatic sensor placement from vision task requirements. *IEEE Trans. on Pattern Recognition and Machine Intelligence*, 10(3):407–416, 1988.
7. K. Dana, B. van Ginneken, S. Nayar, and J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. on Graphics*, 18(1):1–34, January 1999.
8. P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the Reflectance Field of a Human Face. *Computer Graphics Proceedings, Annual Conference Series*, pages 145–156. ACM SIGGRAPH, July 2000.
9. B. Efron and R. Tibshirani. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, 1(1):54–75, 1986.
10. R. Furukawa, H. Kawasaki, K. Ikeuchi, and M. Sakauchi. Appearance based object modeling using texture database: Acquisition compression and rendering. In *Thirteenth Eurographics Workshop on Rendering*, pages 267–276, June 2002.
11. Simon Gibson, Toby Howard, and Roger Hubbard. Flexible image-based photometric reconstruction using virtual light sources. *Computer Graphics Forum*, 20(3), 2001.
12. T. Hastie, R. Tibshirani, and J.H. Friedman. *The elements of statistical learning: data mining, inference and prediction*. Springer-Verlag, 5 edition, 2002.
13. V. Hlavac, A. Leonardis, and T. Werner. Automatic selection of reference views for image-based scene representations. In *4th European Conference on Computer Vision, LCNS 1064/1065*, pages 525–526, Cambridge, Uk, 1996. IEEE, Springer-Verlag, Berlin.
14. G. Kay and T. Caelli. Inverting an illumination model from range and intensity maps. *CVGIP: Image Understanding*, 59(2):183–201, March 1994.
15. K.N. Kutulakos and C.R. Dyer. Recovering shape by purposive viewpoint adjustment. *International Journal of Computer Vision*, 12:113–136, 1994.
16. E. Lafortune, S. Foo, K. Torrance, and D. Greenberg. Non-Linear Approximation of Reflectance Functions. *Computer*

- Graphics Proceedings, Annual Conference Series, pages 117–126. ACM SIGGRAPH, August 1997.
17. J. Lang and M.R.M. Jenkin. Active object modeling with VIRTUE. *Autonomous Robots*, 8(2):141–159, 2000.
 18. H.P.A. Lensch, W. Heidrich, and H.-P. Seidel. Silhouette-based algorithm for texture registration and stitching. *Graphical Models*, 63(4):245–262, July 2001.
 19. H.P.A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially varying materials. In *12th Eurographics Workshop on Rendering*, pages 103–114, June 2001.
 20. 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. on Graphics*, 22(2), April 2003. in print.
 21. Y. Li, S. Lin, S.B. Kang, H. Lu, and H.-Y. Shum. Single-Image Reflectance Estimation for Relighting by Iterative Soft Grouping. In *Pacific Graphics '02*, pages 483–485, October 2002.
 22. M. Marengoni, B. Draper, A. Hanson, and R. Sitaraman. Placing observers to cover a polyhedral terrain in polynomial time. *Vision and Image Computing*, 18(10):773–780, 2000.
 23. S. Marschner. *Inverse rendering for computer graphics*. PhD thesis, Cornell University, 1998.
 24. S. Marschner, B. Guenter, and S. Raghupathy. Modeling and Rendering for Realistic Facial Animation. *11th Eurographics Workshop on Rendering*, pages 231–242, June 2000.
 25. S. Marschner, S. Westin, E. Laforge, K. Torrance, and D. Greenberg. Image-based BRDF Measurement Including Human Skin. In *10th Eurographics Workshop on Rendering*, pages 131–144, June 1999.
 26. S.O. Mason and A. Grün. Automatic sensor placement for accurate dimensional inspection. *Computer Vision and Image Understanding*, 61(3):454–467, 1995.
 27. V. Masselus, P. Dutré, and F. Anrys. The Free-form Light Stage. In *Thirteenth Eurographics Workshop on Rendering*, pages 257–266, June 2002.
 28. W. Matusik, H. Pfister, A. Ngan, P. Beardsley, and L. McMillan. Image-Based 3D Photography Using Opacity Hulls. *Computer Graphics Proceedings, Annual Conference Series*, pages 427–437. ACM SIGGRAPH, July 2002.
 29. W. Matusik, H. Pfister, R. Ziegler, A. Ngan, and L. McMillan. Acquisition and Rendering of Transparent and Refractive Objects. In *Thirteenth Eurographics Workshop on Rendering*, pages 277–288, June 2002.
 30. D. McAllister. *A Generalized Representation of Surface Appearance*. PhD thesis, University of North Carolina, 2002.
 31. K. Nishino, Y. Sato, and K. Ikeuchi. "eigen-texture method: Appearance compression and synthesis based on a 3d model". *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(11):1257–1265, nov 2001.
 32. K. Nishino, Z. Zhang, and K. Ikeuchi. "determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image synthesis". In *in Proc. of Eighth IEEE International Conference on Computer Vision ICCV '01*, pages 599–606, july 2001.
 33. R. Pito. A solution to the next best view problem for automated surface acquisition. *IEEE Trans. on Pattern Recognition and Machine Intelligence*, 21(10):1016–1030, 1999.
 34. W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical recipes in C: the art of scientific computing*. Cambridge University Press, 2 edition, 1992.
 35. R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. *Computer Graphics Proceedings, Annual Conference Series*, pages 117–128. ACM SIGGRAPH, August 2001.
 36. M.K. Reed and P.K. Allen. 3-d modeling from range imagery: An incremental method with a planning component. *Image & Vision Computing*, 17:99–111, 1999.
 37. M.K. Reed and P.K. Allen. Constraint-based sensor planning for scene modeling. *IEEE Trans. on Pattern Recognition and Machine Intelligence*, 22(12):1460–1467, 2000.
 38. S. Rusinkiewicz, O. Hall-Holt, and M. Levoy. Real-time 3d model acquisition. In *ACM Trans. on Graphics*, volume 21, pages 438–446, San Antonio, USA, Jul 2002. ACM SIGGRAPH.
 39. Y. Sato, M. Wheeler, and K. Ikeuchi. Object Shape and Reflectance Modeling from Observation. *Computer Graphics Proceedings, Annual Conference Series*, pages 379–388. ACM SIGGRAPH, August 1997.
 40. W. Scott, G. Roth, and J.-F. Rivest. View planning with a registration constraint. In *3rd International Conference on 3-D Digital Imaging and Modeling*, pages 127–134, Québec City, Canada, 2001. IEEE.
 41. W. Stürzlinger. Imaging all visible surfaces. In *Graphics Interface*, pages 115–122, June 1999.
 42. K. Tarabanis, R. Tsai, and P.K. Allen. The mvp sensor planning system for robotic tasks. *IEEE Trans. on Robotics and Automation*, 11(1):72–85, 1995.
 43. G. H. Tarbox and S. N. Gottschlich. Planning for complete sensor coverage. *Computer Vision and Image Understanding*, 61(1):84–111, 1995.
 44. P.P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich. Image-based modeling using viewpoint entropy. In *Computer Graphics International*, 2002.
 45. P. Whaite and F.P. Ferrie. Autonomous exploration: Driven by uncertainty. *IEEE Trans. on Pattern Recognition and Machine Intelligence*, pages 193–205, 1997.
 46. Y. Yu and J. Malik. Recovering Photometric Properties of Architectural Scenes from Photographs. *Computer Graphics Proceedings, Annual Conference Series*, pages 207–218. ACM SIGGRAPH, July 1998.