

Statistical Acquisition of Texture Appearance

Addy Ngan

Frédo Durand[†]

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology

Abstract

We propose a simple method to acquire and reconstruct material appearance with sparsely sampled data. Our technique renders elaborate view- and light-dependent effects and faithfully reproduces materials such as fabrics and knitweaves. Our approach uses sparse measurements to reconstruct a full six-dimensional Bidirectional Texture Function (BTF). Our reconstruction only require input images from the top view to be registered, which is easy to achieve with a fixed camera setup. Bidirectional properties are acquired from a sparse set of viewing directions through image statistics and therefore precise registrations for these views are unnecessary. Our technique is based on multi-scale histograms of image pyramids. The full BTF is generated by matching the corresponding pyramid histograms to interpolated top-view images. We show that the use of multi-scale image statistics achieves a visually plausible appearance. However, our technique does not fully capture sharp specularities or the geometric aspects of parallax. Nonetheless, a large class of materials can be reproduced well with our technique, and our statistical characterization enables acquisition of such materials efficiently using a simple setup.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction

The creation of photorealistic images has been greatly facilitated by dramatic advances in 3D geometry scanning and rendering algorithms. In contrast, the acquisition and reproduction of real material appearance remains a critical challenge. The most common solution is the use of photographic textures, but plain texture mapping is only a crude approximation for most real world materials as it is only suitable for representing perfectly smooth surfaces with albedo variation and is unable to simulate the effect of any underlying mesostructure of the material. The Bidirectional Texture Function (BTF) [DvGNK99] is a 6D function that describes the full light/view dependence of an image patch. Unfortunately, the measurement of BTFs is a formidable task that requires the imaging of a single sample for all pairs of view and light directions. A robotic setup is typically used, and precise calibration is necessary to align all these images, a task made harder by the various moving parts.

In this work, we propose a method to capture complex

materials such as wool knitwear in minutes, with a focus on visual faithfulness rather than geometric and photometric accuracy. Simplicity of implementation is our main goal, and we therefore seek to minimize the reliance on precise registration and alleviate the need for robotic parts. While our approach cannot handle highly specular materials such as metals and does not fully measure parallax, for a large class of materials it produces photorealistic BTFs using a setup that is much simpler than comprehensive approaches. In addition, while disparity aspects of parallax are not captured, its statistical effects such as the view-dependent changes of color distribution and sharpness are reproduced, which results in compelling materials.

Key to our approach is a novel BTF reconstruction technique that takes sparse measurements and reconstructs the full BTF. Our reconstruction is based on the interpolation of statistical properties, thereby alleviating the need for precise registration and avoiding cross-fading artifacts that could be caused by simple data interpolation.

[†] email: addy, fredo@csail.mit.edu




Figure 1: Indoor scene rendered using 4 textures acquired and reconstructed with our technique.

Contributions This paper introduces the following contributions:

- We propose a new algorithm to reconstruct BTFs from *sparsely sampled* and *unaligned* measurements.
- This enables a simple and low-cost acquisition setup that allows for the simultaneous capture of multiple views and does not require robotically-controlled moving parts.

1.1. Related Work

We give a brief overview of work on BTFs and material measurement and refer the reader to the survey by Müller et al. [MMS*05] for a comprehensive study of recent works on BTFs. We focus on BTF acquisition and reconstruction and our work is orthogonal to issues such as BTF rendering [SSK03, SBLD03, MMK04].

Since Dana et al. [DNvGK97] introduced the notion of the Bidirectional Texture Function (BTF) and published the first BTF database [Cur], a small number of teams have performed measurements using robotically-controlled setups [SSK03, KMBK03]. These robotic setups involve a moving material sample, as a result, measurements at different view/light directions are not always perfectly aligned, as also observed by Filip and Haindl [FH05]. Despite these problems, these dense measurements are invaluable and provide us with references to evaluate our reconstruction technique.

Recently, a number of setups have been described to capture material appearance more efficiently by simultaneously capturing multiple views using multiple cameras [MBK05] or mirrors [HP03, DW04]. Our goal is to enable simpler setups and minimize the time and effort of measurements by introducing a statistical reconstruction of BTFs.

In particular, we want to leverage the increased resolution of digital cameras and perform multiple measurements

within a single picture using spatial multiplexing as done for BRDF by Marschner et al. [MWL*99] and Ngan et al. [NDM05]. However, this is more challenging for BTFs because they are by definition not spatially uniform, and it might not be possible to put different samples in perfect correspondence. This motivates our use of statistics that are robust to registration.

A number of acquisition techniques have been developed to capture simplified versions of the BTF. Kautz et al. [KSS*04] only capture variation due to light elevation and obtain convincing results when light azimuth and view direction are not important. Malzbender et al. [MGW01] capture the effect of light dependence and approximate them with smooth polynomials. Our approach provides increased accuracy for the light sampling and adds view-dependent effects such as low-frequency BRDFs and the intricate appearance exhibited by fuzzy materials such as fabrics.

The BTF synthesis work by Liu et al. [LYS01] is closely related to our work. They use a sparse set of images to estimate an approximate height field by shape-from-shading, and synthesize new geometry that is statistically similar to the acquired sample. Pixel samples from the input are then copied to the synthesized image based on feature matching to reconstruct the bidirectional appearance. Their technique is limited to stochastic textures that can be described as height fields. Our work is different as we do not make the height field assumption, and we do not rely on geometric information for reconstruction.

Our work is related to efforts on BTF compression [SSK03, KMBK03, VT04] since it reconstructs a BTF from a sparse subset. However, most compression approaches seek to optimize decompression at rendering time, a feature that is not directly possible with our method. Nevertheless, the compression technique by Filip and Haindl [FH05] is closely related to our work. They fit per-*texel* Laforune lobes together with per-view/light histogram remapping functions to achieve compression and fast rendering. The context of BTF compression is different because they are given the full BTF to perform parameter estimation, while we need to fill in missing data. In addition, our statistical characterization is more elaborate and includes frequency content.

Our work is inspired by studies in vision that analyze bidirectional texture from a statistical point of view. Leung and Malik [LM01] have shown that view-dependent masking effects can significantly change the color distribution, which explains why sunflower fields look more yellow at grazing angles. Cula and Dana [CD01] use multi-scale features from BTF for material recognition. Pont and Koenderink [PK05] use a simple micro-facet Lambertian model to predict texture contrast at different lighting and viewing configurations. We build on experiments by Ginneken et al. [vGKD99] who investigate the pixel histograms of a wider class of materials in the CURET data set as a function of view and light direction. In particular, in the case of surfaces with uniform albedo,




Figure 2: Acquisition setup - both the camera and the measured target are fixed, and a handheld wireless flash is used as the light source. During measurement, the user moves the flash source around to roughly cover all possible directions, and remotely triggers the camera shutter to take pictures.

they show that a single texture modulated with histogram matching offers improved material appearance.

Dana and Nayar’s work [DN98, DN99b] is most closely related to our reconstruction technique. They study BTF under varying light/view configurations and propose analytical histogram and correlation models in the special case of random isotropic surfaces with Lambertian reflectance. They also propose BTF synthesis based on histogram transfers [DN99a]. The top-view image with correct lighting is first synthesized and then the image is transformed to an arbitrary view through pixel histograms transfer. They showed that the technique works well for a sample Lambertian material with gaussian height distributions. Our reconstruction technique extends this idea to a wider class of materials by using multi-scale statistics. In particular, we show that the idea of pyramid matching traditionally used to synthesize textures from white noise [HB95] can be adapted to modify a base texture and enforce the most salient visual variation due to a material’s appearance. In addition, our reconstruction scheme is specifically designed for sparsely sampled data and allows for meaningful histogram interpolations and extrapolations.

1.2. Overview

Our method seeks to reconstruct a full BTF based on measurements that are easy to obtain. We only require the images of the frontal view of the material sample to be aligned, which is easy to achieve with a fixed camera and sample, and a moving light source. The images with different view directions do not need to be registered with the frontal view. In particular, we can combine multiple views per images by using multiple samples of the material placed at different angles (see Fig. 2), which greatly accelerates acquisition.

Our reconstruction algorithm is the central contribution that makes it possible to obtain BTFs from unregistered views. It is based on the idea of characterizing textures by

their multi-scale statistics [HB95]. Our reconstruction uses the naturally-aligned images from the top view as base textures. View-dependent effects are then transferred to these base textures using the histograms of the pyramid subbands. For a large class of materials, our reconstruction provides a visually plausible approximation to the true BTF.

2. Acquisition

We first present our acquisition setup to make the input data of our reconstruction technique concrete. However, different acquisition setups could be used with our reconstruction as long as they provide aligned images for at least one view direction.

Our acquisition setup (Fig. 2) exploits the ability of our reconstruction technique to work with view images that are not aligned. It typically captures 13 views of a material at a time using different samples at various orientations. The camera is fixed and the light direction is sampled by moving a light source.

In practice, we paste a number of planar patches of the material onto square backing boards with known dimensions, which are then positioned to form a pyramid-like target. The arrangement provides 4 views at about 30° incidence angle, 8 views at about 60° and the top view. We use an 8 megapixel digital SLR camera with a hand-held electronic flash. The camera is set up on a tripod at a fixed position about 1 meter above the measurement target, and the size of the top-view patch in the image is roughly 500×500 pixels. We put specular spheres (billiard balls) around the target for light position estimation. The user holds the flash directed at the target from various directions, and a remote control is used to trigger the camera. We take about 100 pictures for each material in about 10 – 15 minutes.

The camera is calibrated automatically using several images of a checkerboard [Zha00]. For each sequence, the user manually marks the position of the corners of each square backing boards. With the known dimensions of the boards, we compute the position and orientation of each patch using a least-square optimization. The flash source is small, and we approximate it as a point source. For each image, the mirror peak on each specular sphere is located automatically, if available, and the light source position is estimated with a least-square fit.

For each measured image, we resample each lit patch into a texture image at a resolution of 512×512 . As the extent of a single patch is relatively small, we assume the light and view directions (ω_i, ω_o) are constant across each patch. The resampled texture is then normalized by the estimated irradiance to form a 2D spatial slice of the BTF. Even though the flash power is manually set to constant power, in practice the flash output can deviate from the specified power significantly. We put diffuse gray cards next to the pyramid to get a reliable estimate of the flash power. However, for some

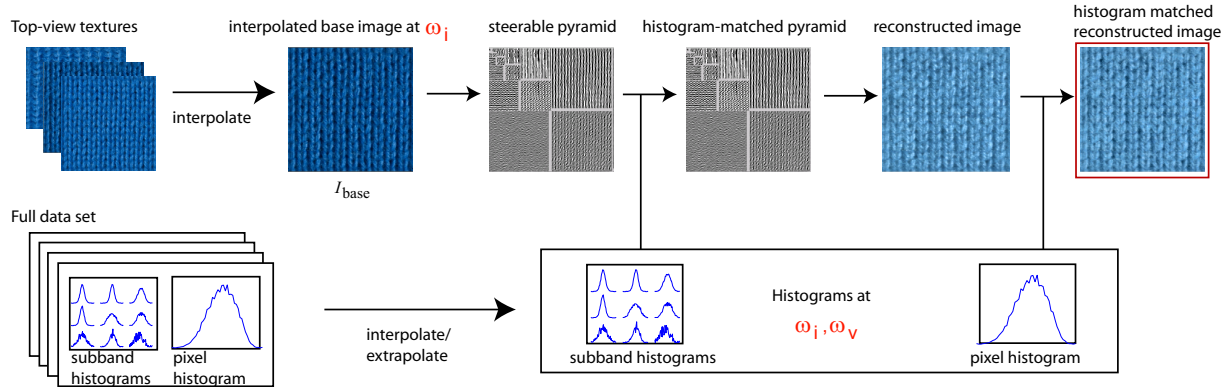


Figure 3: Reconstruction pipeline for producing a texture at light/view direction ω_i, ω_v . From the set of top-view textures, we interpolate the neighbors of ω_i to produce a base texture I_{base} . From the full data set comprising all light/view combinations, we locate the closest neighbors of (ω_i, ω_v) . We interpolate (and potentially extrapolate) the subband histograms and pixel histogram to form the desired histograms of our target texture. The base texture I_{base} is decomposed into a multi-scale oriented pyramid, and each subband is matched to the desired histogram. We then collapse the pyramid to reconstruct the image, which is then matched to the desired pixel histogram. This produces our final texture.

measured images the gray cards are shadowed and in those cases the mean flash power is assumed. To further compensate for the flash power estimation error, we compute the average BRDF of the measured data by averaging over each normalized texture. We then smooth this BRDF and use it to re-scale our measured data so that it is consistent with the smoothed BRDF. In practice the BRDF blurring does not degrade the quality of measurement as our ability to measure specular materials is limited by the sparse sampling of the view.

3. Reconstruction

The central idea of our reconstruction is to use the actual images only from the top view and to characterize the visual properties of other views using alignment-insensitive statistics. We reconstruct a BTF one texture at a time, for each pair of view-light directions. A texture is reconstructed in two steps (Fig. 3). We first use our set of top-view images, which are assumed to be aligned, and interpolate them linearly based on the light direction to obtain a *base texture*. This texture contains appropriate shadowing effects but might exhibit cross-fading artifacts and does not include view-dependent effects such as masking, BRDF, and asperity scattering. Our second step improves interpolation quality and reproduces these effects by enforcing statistics corresponding to the appropriate view-light directions.

3.1. Histogram Statistics

Before presenting our full reconstruction pipeline, we discuss our choice of statistics to characterize material appearance. We observe that the pixel histogram encodes variation

in the color distribution due to effects such as shadowing and masking. Fig. 4(a) shows a BTF slice of the measured knitwear at $[\omega_v = (60, 0), \omega_i = (60, 180)]$.[†] Fig. 4(b) shows the *base texture* from the top-view with approximately the same light direction ω_i . We observe prominent shadowing due to the low elevation of the light in the base texture, but the effect is significantly reduced when viewed from the side in (a), both due to masking and the increased scattering path length. By matching the pixel histogram of the base image to the target image, we can recover the overall distribution of intensity, and as a result shadows are mostly eliminated (Fig. 4(c)). This pixel histogram matching technique has been proposed by Dana and Nayar [DN99a]. However, pixel histogram matching does little to the structure of the image as it is insensitive to multi-scale effects such as blurring.

To further improve the appearance transfer, we employ the steerable pyramid in an approach similar to Heeger and Bergen [HB95]. The base and target image are both decomposed into an image pyramid with multiple scaled and oriented subbands, and the coefficient histogram of each subband is matched independently. Both pixel histograms and pyramid coefficient histograms can be computed without image registration and are well suited for our goal. The subband histogram matching captures effects such as the fuzziness that some materials exhibit at grazing angle. In Fig. 4(d), we first transfer the subband histograms from the target image to the base image, followed by the pixel histogram matching.

Heeger and Bergen [HB95] generate new textures starting

[†] We represent direction vectors ω in spherical coordinates (θ, ϕ) , where θ and ϕ are the incidence and azimuth angles respectively.




Figure 4: Multi-scale statistics transfer. We seek to reproduce the statistics of a target texture (a) starting from a top-view base texture (b). Note how our full technique (d) improves the reconstruction quality beyond pixel histogram matching (c). Without the spatial structure provided by the base texture, texture synthesis technique by Heeger and Bergen [HB95] (e) is unable to reproduce the target texture.

from noise and iteratively enforcing the histogram of pixel color and pyramid coefficients. As a result, their method is generally unable to reproduce extended spatial structure. Fig. 4(e) shows the result when we directly apply their texture synthesis technique with five iterations using statistics from (a). In contrast, we use the multi-scale statistics to *modify* the base texture obtained for a given light direction, which already contains a very similar spatial structure to the target.

Our reconstruction makes use of both pixel and subband histograms to statistically characterize the BTF. For each BTF image slice in the measured data set, we decompose the image into oriented subbands with 4 orientations and all levels of scale. We compute and store the histogram for each subband ready for the texture generation step.

3.2. Texture Generation

To reconstruct the full BTF we reconstruct 2D image slices at discretely sampled view/light directions. The pipeline for reconstruction is shown in Fig. 3. To generate the texture at a particular light/view configuration (ω_i, ω_v) , we first find neighbors of ω_i in the set of top-view textures. We first gather all the textures from the top view, and we project the light direction of each texture in the unit hemisphere onto the unit disk. We perform a Delaunay triangulation of the set, and find the three neighboring textures of ω_i by searching the containing triangle. If ω_i is outside the convex hull, we project ω_i radially onto the convex hull. We blend the corresponding textures using the barycentric weights to form I_{base} , which roughly corresponds to the texture lit from the desired direction ω_i but viewed from the top.

To obtain the desired histogram statistics, we need to find the close neighbors and determine the appropriate interpolation weights. As the set of view and light directions are separable ($2D \times 2D$), we can perform a two-level interpolation: 1) first find the weights for view interpolation, 2) for each neighboring view we find the interpolation weights for the neighboring light directions. As the view direction is assumed to be sparsely sampled and roughly structured,




Figure 5: View interpolation/extrapolation in spherical coordinates. The view directions are grouped into classes with similar incidence angle θ and form rings on the hemisphere. With this semi-uniform structure, the view direction can be interpolated or extrapolated in a bilinear fashion.

we use a different interpolation scheme to avoid skewed triangles from triangulations, and to allow for meaningful extrapolation. We represent the view directions in spherical coordinates (θ, ϕ) and group them into classes based on the incidence angle θ (Fig. 5). For our setup we have three classes $\theta \approx 0, \theta \approx 30$ and $\theta \approx 60$. We join the view directions within each class with a polyline, and we assume the polylines do not cross each other. These polylines divide the spherical domain into rings, with which a bilinear interpolation/extrapolation can be well-defined. The second-level interpolation of the light direction for each view is performed using the same strategy as the base texture.

We interpolate the pixel and subband histograms according to the computed blending weights. Histograms are interpolated by linear blending of the inverse CDFs (cumulative density function), as proposed by Matusik et al. [MZD05]. In practice such interpolation gives much more natural transitions. Also, inverse CDFs can be readily extrapolated with negative weights that sum to one, while directly extrapolating histograms can lead to invalid negative density.

We decompose I_{base} with the same steerable pyramid, and we match each subband to the corresponding blended his-

togram. We then collapse the pyramid to recover an image I'_{base} . Finally, we match the pixel histogram of I'_{base} to the interpolated pixel histogram to produce our final reconstructed texture.

4. Results

We evaluate our approach in two different ways. First, we perform new measurements using the simple setup described in Section 2, and show that we are able to capture important visual characteristics of realistic textures. Second, we use publicly-available BTF databases measured using robotic setups as ground truth and perform our reconstruction using a subset of the data.

Acquisition and reconstruction We have measured 16 materials including different kinds of carpets and fabrics and reconstructed the full BTF from the measurements; rendered images are shown in Fig. 13. The cloth geometry used in the rendered images is generated through standard cloth simulation without considering physical properties of the rendered materials. As a result, the amount of folding and stretching may not be consistent with the real materials. In addition, our texture representation inherits the limitation of BTFs: silhouettes are not captured, and effects due to surface curvature are ignored. Nonetheless, in most cases our reconstruction is able to reproduce the visual quality of real materials faithfully. **Please refer to our supplemental video for animated sequences of the acquired BTFs.**

Acquisition time is dependent on the sampling density of the light directions; in practice, we capture about 100 images for a material in about 10 – 15 minutes. Next the images are processed and individual patches (512×512 pixels) are resampled. This processing takes 20 – 30 minutes on a P4 3.0GHz PC and about 800 slices of the BTF is typically captured.

We decompose each texture slice of the captured BTF into a steerable pyramid as described in Section 3, and we compute and store a histogram for each subband, in addition to the pixel histogram of the image. This step takes about 2 – 3 hours on our PC. Texture reconstruction takes about 20 seconds per image, and it takes 36 hours to reconstruct a full BTF with 81×81 views and lights. Our reconstruction is implemented using Matlab and could be greatly optimized.

Validation To further validate our reconstruction method, we compare our results against a number of materials in the Bonn BTF database [SSK03]. Each material in the Bonn database is measured at 81 light directions \times 81 view directions, for a total of 6561 texture images. To test the capability of our reconstruction, we pick 13 views roughly corresponding to our acquisition setup, including 8 views at 60° incidence, 4 at 30° and the top view. We use all 81 light directions for each of the 13 views. The input to our reconstruction thus includes the 81 light-dependent textures from the

top view, and 13×81 sets of pixel and subband histograms. This is roughly equal to the number of samples we collect in our acquisition, despite the fact that our sample directions are not as uniform.

We reconstruct the BTF at each of the original 6561 direction pairs with our technique. The first two columns of Fig. 12 shows the comparison between the reconstruction and original BTF for the materials corduroy, wool and proposte.

We also compare our reconstruction with other approximation techniques. The third column of Fig. 12 shows renderings when only the top-view textures are used, without any statistics transfer. Notice that this can be seen as an upper bound for the approximation quality of view-independent methods, e.g. [MGW01]. The fourth column uses a single texture (light and view both from the top), modulated by the average BRDF computed from the original Bonn data (i.e. 6561 samples). We do the same comparison for our measured materials *knitwear-1* and *green-knitwear* in Fig. 11. Note the complex visual features not reproduced in the texture-mapped and top-view versions, notably the reduction of the shadowed texels due to masking, and how the texture becomes more blurry at grazing angles, revealing the fluffiness of the materials. Interestingly, we have observed that, in the absence of context, casual observers often prefer the top-view version because it exhibits more contrast and sharpness. This is a well-known bias in image quality evaluation, e.g. [JF00]. We however emphasize that this excessive sharpness results in hyper-realistic images that look artificial in complex scenes. Contrast/sharpness reduction is an important effect that many real materials exhibit and that it is particularly critical to reproduce fuzzy materials like fabric. In addition reduced contrast due to masking is also important in more solid materials such as tree bark and plaster [PK05].

Statistical characterization In Figs. 4 and 6, we show results of our reconstruction for a few example slices viewed from 60° incidence. Our reconstruction provides good visual matches to the target images in terms of brightness, color, contrast, blurriness and other subtle multi-scale effects. Notice that these side-view images are rectified to the normal view for display: for $\omega_v = (60, 0)$, the width of the texture would be halved when used for actual rendering. As a result, any discrepancy along the horizontal dimension would be diminished.

Our method captures effects such as the Fresnel term to a first-order approximation. As described in Section 3.2, we extrapolate the histograms for slices beyond 60 degrees, and in practice, it is able to provide plausible increase of the brightness at grazing angle (Figure 7). It is in particular important that we extrapolate the *inverse cumulative* histograms, which faithfully increases brightness when appropriate and avoids numerical problems such as invalid histograms. However, because we perform linear extrapolation, we can underestimate the increase of brightness when the




Figure 6: Texture reconstruction at $\omega_v = (60, 0)$. Row 1: Carpet from Koudelka et al. [KMBK03], Row 2: Measured material carpet-1.




Figure 7: Comparing the interpolated pixel histograms to the Bonn measurement of the material propose. Fixing the light direction at $(75, 0)$, we compare histograms for six viewing directions. All views except $(0, 0)$ are not present in the reconstruction data set. The $(75, 150)$ view is extrapolated as our most inclined views are at 60 degrees.

function is concave (Figure 7). Unfortunately, we have found it difficult to reliably quantify the extrapolation errors, since available measurements are typically unreliable near grazing angle (e.g. the 75 degree measurements of the Bonn dataset).

Sampling Density Experimenting with reducing the sampling density in the light direction, we found that for some materials, high fidelity can be achieved with a surprisingly low number of lights. In Fig. 8(b) we show that our reconstruction yields very good results even when only 13 lights per view are used. In comparison, when the same number of texture images are directly used for rendering via interpolation, the result exhibits strong cross-blending artifacts (Fig. 8(a)).




Figure 8: Comparing different sampling density of light directions: (a) Direct interpolation of 13×13 (view \times light) textures from Bonn data, (b) Our reconstruction with 13×13 textures, (c) Our reconstruction with 13×81 textures and (d) Original BTF with $81 \times 81 = 6561$ textures.

Limitations Our acquisition is limited to materials without sharp specularities due to the sparse sampling of the view directions. Strong parallax effect is also difficult to reproduce with histogram statistics as spatial structure is not directly encoded. For example, the Lego material measured by Koudelka et al. [KMBK03] exhibits strong parallax due to its relatively steep depth, while having an otherwise simple geometry that lacks significant scattering: it is the worst-case scenario for our approach. In Fig. 9(a) we show a slice of the BTF viewed at 60° incidence. Compared to the base image from the top view, the bumps are offset due to disparity, while shadows are stretched due to perspective. Our reconstruction is unable to capture the disparity, but its multi-scale nature allows it to partially reproduce the elongation of shadows. In Fig. 10 we compare the original and the reconstructed Lego BTF rendered on a sphere. Note that the light fall-off near the texture boundary in both the Lego BTF and the sponge BTF shown in the video are artifacts from the original measurements. In summary, while we consider the Lego material as a failure case for our approach, the reconstruction still looks surprisingly good and exhibits more 3D effects than a flat texture, in particular in the video.

5. Discussion

We have presented a simple method to acquire the appearance of materials such as fabric and knitwear that present rich spatial and angular variation. Our main contribution




Figure 9: Failure case for the Lego BTF [$\omega_v = (60, 0)$]. Our reconstruction is unable to reproduce the disparity, but the shadow elongation is partially captured.

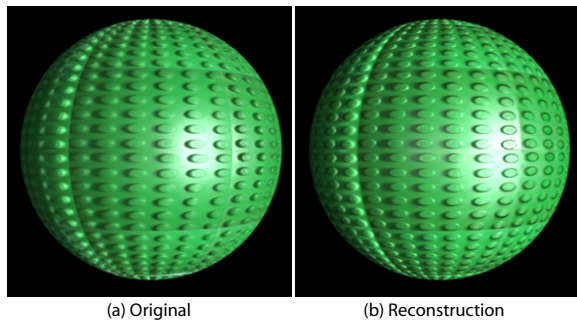


Figure 10: Failure case for the Lego BTF: original vs our reconstruction.

is a reconstruction algorithm that generates a full bidirectional texture function (BTF) from a sparse set of measurements. Different views of the material sample do not need to be aligned because we characterize view-dependent effects using alignment-insensitive statistics, namely marginal and multi-scale histograms.

As we have limited resolution in the view dimension, it is clear that our technique cannot capture high-frequency effects such as highly specular materials. Our statistical characterization does not handle the geometric effect of parallax but it reproduces some of its effects such as masking. Our statistical reconstruction tends to work best on materials with complex spatial structure (e.g. wool, propose), as the high-frequency content and the statistical variation dominate the visual appearance. For such materials, it fits an important gap since realistic fuzzy fabric and knitwear appearance has been challenging to measure with simple means.

Our technique inherits from the limitations of the BTF concept. While BTFs capture statistical effects of multiple scattering, they do not model the spatial component that causes blurred shadows and light bleeding. Moreover, BTFs do not capture the appearance of material silhouettes, a critical visual factor for which appropriate acquisition techniques are needed.

We believe that our statistical reconstruction has potential beyond our simple setup, notably to exploit partial views

such as the one obtained with the kaleidoscope acquisition setup. We also believe that statistical material modeling can be applied to the capture of appearance from a single photograph. Domain-specific knowledge or priors on BTFs might also enable a better use of the higher light sampling rate to improve reconstruction along the view direction and to better handle parallax. Finally, our statistical reconstruction from sparse samples suggests that the 6D BTF can be avoided altogether, and we want to develop a compact representation based on statistics that will enable real-time rendering with a small memory footprint.

Acknowledgement We would like to thank Jan Kautz, Jiawen Chen, Erik Hubo, Neel Joshi, Bill Freeman, Sara Su, Paul Green and the anonymous reviewers for helpful comments. Tom Buehler produced the accompanying video, and Robert Wang created the cloth mesh. We acknowledge BTF Database Bonn and Koudelka et al. for sharing their BTF measurements. This work was supported by a National Science Foundation CAREER award 0447561 “Transient Signal Processing for Realistic Imagery.” Frédo Durand acknowledges a Microsoft Research New Faculty Fellowship.

References

- [CD01] CULA O., DANA K.: Compact Representation of Bidirectional Texture Functions. In *Conference on Computer Vision and Pattern Recognition (CVPR 2001)* (Dec. 2001), vol. 1, pp. 1041–1047.
- [Cur] Curet: Columbia-Utrecht reflectance and texture database. <http://www.cs.columbia.edu/CAVE/curet/>.
- [DN98] DANA K. J., NAYAR S. K.: Histogram model for 3D textures. In *CVPR '98: Proc. of the 1998 Conference on Computer Vision and Pattern Recognition* (1998), IEEE Computer Society, p. 618.
- [DN99a] DANA K. J., NAYAR S. K.: 3D textured surface modeling. In *CVPR workshop on the Integration of Appearance and Geometric Methods in Object Recognition* (1999).
- [DN99b] DANA K. J., NAYAR S. K.: Correlation model for 3d texture. In *ICCV '99: Proc. of the International Conference on Computer Vision-Volume 2* (1999), IEEE Computer Society, p. 1061.
- [DNvGK97] DANA K. J., NAYAR S. K., VAN GINNEKEN B., KOENDERINK J. J.: Reflectance and texture of real-world surfaces. In *CVPR '97: Proc. of the 1997 Conference on Computer Vision and Pattern Recognition* (1997), IEEE Computer Society, p. 151.
- [DvGNK99] DANA K. J., VAN GINNEKEN B., NAYAR S. K., KOENDERINK J. J.: Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* 18, 1 (1999), 1–34.
- [DW04] DANA K., WANG J.: Device for convenient mea-




Figure 11: Comparing approximations to the measured materials knitwear-1 and green-knitwear. First column: single texture modulated by acquired BRDF, second column: light-varying textures from top view, and third column: our reconstruction.

- surement of spatially varying bidirectional reflectance. *J. Opt. Soc. Am. A* 21, 1 (2004), 1–12.
- [FH05] FILIP J., HAINDL M.: Efficient image-based bidirectional texture function model. In *Texture 2005: 4th International Workshop on Texture Analysis and Synthesis* (Beijing, China, October 2005), pp. 7–12.
- [HB95] HEEGER D. J., BERGEN J. R.: Pyramid-based texture analysis/synthesis. In *Proceedings of SIGGRAPH 1995* (1995), ACM Press, pp. 229–238.
- [HP03] HAN J. Y., PERLIN K.: Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Trans. Graph.* 22, 3 (2003), 741–748.
- [JF00] JOHNSON G., FAIRCHILD M.: Sharpness rules. In *IST/SID 8th Color Imaging Conference* (2000), pp. 24–30.
- [KMBK03] KOUDELKA M. L., MAGDA S., BELHUMEUR P. N., KRIEGMAN D. J.: Acquisition, compression, and synthesis of bidirectional texture functions. In *Texture 2003: Third Int. Workshop on Texture Analysis and Synthesis* (Nice, France, October 2003), pp. 59–64.
- [KSS*04] KAUTZ J., SATTLER M., SARLETTE R., KLEIN R., SEIDEL H.-P.: Decoupling BRDFs from surface mesostructures. In *GI '04: Proceedings of the 2004 conference on Graphics interface* (2004), pp. 177–182.
- [LM01] LEUNG T., MALIK J.: Representing and Recognizing the Visual Appearance of Materials using Three-dimensional Textons. *International Journal of Computer Vision* 43, 1 (2001), 29–44.
- [LYS01] LIU X., YU Y., SHUM H.-Y.: Synthesizing bidirectional texture functions for real-world surfaces. In *Proceedings of SIGGRAPH 2001* (2001), ACM Press, pp. 97–106.
- [MBK05] MÜLLER G., BENDELS G., KLEIN R.: Rapid Synchronous Acquisition of Geometry and BTF for Cultural Heritage Artefacts. In *The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST)* (November 2005), pp. 13–20.
- [MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *Proceedings of SIGGRAPH 2001* (2001), ACM Press, pp. 519–528.
- [MMK04] MESETH J., MÜLLER G., KLEIN R.: Reflectance field based real-time, high-quality rendering of bidirectional texture functions. *Computers and Graphics* 28, 1 (February 2004), 103–112.
- [MMS*05] MÜLLER G., MESETH J., SATTLER M., SARLETTE R., KLEIN R.: Acquisition, synthesis, and rendering of bidirectional texture functions. *Computer Graphics Forum* 24, 1 (March 2005), 83–109.
- [MWL*99] MARSCHNER S., WESTIN S., LAFORTUNE E., TORRANCE K., GREENBERG D.: Image-based




Figure 12: With 3 materials from the Bonn BTF database (Wool, Corduroy and Proposte), we compare our reconstruction against the original data. The scene is lit with a main directional light from the right, and a second directional light from the front. The third column shows the results when only the top-view textures are used, while the fourth column shows a single texture (top view, top light), modulated by the per-slice BRDF averaged over the texture.

- BRDF measurement including human skin. In *Rendering Techniques '99* (1999), pp. 139–152.
- [MZD05] MATUSIK W., ZWICKER M., DURAND F.: Texture design using a simplicial complex of morphable textures. *ACM Trans. Graph.* 24, 3 (2005), 787–794.
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering* (2005), Eurographics Association, pp. 117–126.
- [PK05] PONT S. C., KOENDERINK J. J.: Bidirectional texture contrast function. *International Journal of Computer Vision* 62, 1-2 (2005), 17–34.
- [SBLD03] SUYKENS F., BERGE K., LAGAE A., DUTRE P.: Interactive rendering with bidirectional texture functions. *Comput. Graph. Forum* 22, 3 (2003), 463–472.
- [SSK03] SATTLE M., SARLETTE R., KLEIN R.: Efficient and realistic visualization of cloth. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering* (2003), Eurographics Association, pp. 167–177.
- [vGKD99] VAN GINNEKEN B., KOENDERINK J. J., DANA K. J.: Texture histograms as a function of irradiation and viewing direction. *Int. J. Comput. Vision* 31, 2-3 (1999), 169–184.
- [VT04] VASILESCU M. A. O., TERZOPOULOS D.: Tensor textures: multilinear image-based rendering. *ACM Trans. Graph.* 23, 3 (2004), 336–342.
- [Zha00] ZHANG Z.: A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.* 22, 11 (2000), 1330–1334.




Figure 13: Rendered images of the measured materials. First row: pink-knitwear, fleece, denim, carpet-2, second row: carpet-3, velcro, knitwear-1, green-knitwear and third row: carpet-1, pattern-3, pattern-2 and pattern-1.