Computers & Graphics 73 (2018) 26–36

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Technical Section

Rapid material capture through sparse and multiplexed measurements☆

Dennis den Brok*, Michael Weinmann, Reinhard Klein

Institute of Computer Science II, University of Bonn, Endenicher Allee 19a, Bonn 53115, Germany

ARTICLE INFO

Article history:
Received 11 October 2017
Revised 9 February 2018
Accepted 6 March 2018
Available online 20 March 2018

Keywords:
Appearance acquisition
Sparse acquisition
Illumination multiplexing
Appearance modeling

ABSTRACT

Among the many models for material appearance, data-driven representations like bidirectional texture functions (BTFs) play an important role as they provide accurate real-time reproduction of complex light transport effects such as interreflections. However, their acquisition involves time-consuming capturing of many thousands of bidirectional samples in order to avoid interpolation artifacts. Furthermore, high dynamic range imaging including many and long exposure steps is necessary in the presence of low albedo or self-shadowing. So far, these problems have been dealt with separately by means of sparse reconstruction and multiplexed illumination techniques, respectively. Existing methods rely on data-driven models learned on data that has been range-reduced in a way that made their simultaneous application impossible. In this paper, we address both problems at once through a novel method for learning data-driven appearance models, based on moving the dynamic range reduction from the data to the metric. Specifically, we learn models by minimizing the relative L_2 error on the original data instead of the absolute L_2 error on range-reduced data. We demonstrate that the models thus obtained allow for faithful reconstruction of material appearance from sparse and illumination-multiplexed measurements, greatly reducing both the number of images and the shutter times required. As a result, we are able to reduce acquisition times down to the order of minutes from what used to be the order of hours.

1. Introduction

Analytical material reflectance models such as (spatially varying) bidirectional reflectance distribution functions ((SV)BRDFs) [1], which model reflectance per surface point depending on the incident and outgoing light directions, can nowadays be obtained efficiently, as simply as by taking two photographs with a cellphone camera [2]. Many applications require a higher degree of accuracy than what these models are able to deliver, or real-time rendering including meso-scale light transport effects that require solving global illumination, e.g. interreflections and self-shadowing. Image-based representations such as data-driven SVBRDFs or bidirectional texture functions (BTFs) [3] provide these advantages. Capturing them, however, demands much more effort, up to days for a single material [4]. This can largely be attributed to two factors: high-frequency features and dynamic range. The former can be caused e.g. by specularity, parallax and shadows. In order to avoid visible sub-sampling artifacts in rendering, often tens of thousands of images need to be obtained. The latter is a consequence particularly of specularity, but also of shadows or low albedo. The more prominent these effects, the greater the number of exposure steps and the maximum exposure time necessary to capture a material's full dynamic range.

By now there are a number of approaches to solving these problems separately. Sparse acquisition techniques are applied when only a small subset of the desired dense sampling is actually measured; the remaining data is then interpolated by means of linear models learned from an existing material database (e.g. den Brok et al. [5], Nielsen et al. [6]). Conversely, illumination multiplexing exploits the linearity of the superposition of light by illuminating the material sample not with a single light source but with patterns of light sources, which increases the amount of light on the sample and eliminates shadows, thereby reducing dynamic range (see Fig. 3). The desired representation with one active light source per image can then be obtained by solving an appropriate linear system, a process that is, however, known to be detrimental to the signal-to-noise ratio (SNR). The models used in sparse acquisition have been shown to also help mitigate the noise problems [7].

Either way, acquisition times can be reduced significantly, down to the range of at most a few hours [4], but still far from what acquisition devices for analytical SVBRDFs are capable of. As the approaches are completely orthogonal, the question arises whether

☆ This article was recommended for publication by H. Rushmeier.

* Corresponding author.
E-mail address: denbrok@cs.uni-bonn.de (D. den Brok).

https://doi.org/10.1016/j.cag.2018.03.003
0097-8493/

the two paradigms can be combined. So far, this has been impossible: the linear bases used as models in the above approaches rely heavily on range-reduction techniques such as logarithmic transformations applied to the training data. These transformations, however, do not commute with multiplexing; i.e., we cannot infer the transformed data from a multiplexed measurement without prior de-multiplexing. But de-multiplexing requires images for *all* multiplexing patterns, which we wish to avoid in sparse acquisition.

In this paper, we present, to the best of our knowledge, the first approach to accurate reflectance acquisition which simultaneously exploits sparse acquisition and multiplexed illumination, enabling faithful BTF acquisition in several minutes. Specifically, we propose a different approach to dynamic range reduction in model learning: rather than the *absolute* L_2 error on non-linearly transformed data as a metric, we minimize the *relative* L_2 error on untransformed data, which ultimately allows for sparse multiplexed acquisition of BTFs. As obtaining a basis this way is not as straightforward as simply computing a truncated singular value decomposition (SVD), we provide an efficient alternating least-squares approach to compute a suitable basis. As demonstrated by our results, combined sparse and multiplexed acquisition allows for a reduction of acquisition time from the order of hours/days required for brute-force measurements down to only several minutes, significantly outperforming both sparse acquisition and multiplexed acquisition.

We evaluate the performance in the sparse and multiplexed case, both separately and combined, and compare against the state-of-the-art. In our evaluation, we find that a method recently presented by Nielsen et al. [6], which had only been tested on BRDFs and flat SVBRDFs so far, also works on material BTFs and slightly outperforms the state-of-the-art in this field.

In summary, our paper presents the following key contributions:

- a novel basis for measured material appearance based on minimizing the *relative* L_2 error on the *untransformed* data instead of the *absolute* L_2 error on non-linearly transformed data,
- an evaluation of our basis' performance as a model for the appearance of typical real-world materials in the context of sparse or multiplexed acquisition.
- *En passant*, we find that a recently presented sparse acquisition method only known so far to work for BRDFs and flat SVBRDFs also lends itself to arbitrary material BTFs and slightly outperforms the state-of-the-art in this field.
- We demonstrate that our basis is designed to take advantage of both sparse acquisition and multiplexed illumination at once, resulting in an overall acquisition speed-up of up several orders of magnitude in comparison to a full measurement, and a still significant speed-up of the acquisition process in comparison to sparse or multiplexed acquisition, while maintaining perceptually accurate results.

2. Related work

In this section, we briefly review related work on modeling surface appearance including fine surface details. Furthermore, we discuss previous work on fast appearance acquisition based on the aforementioned concepts of sparse acquisition and illumination multiplexing.

2.1. Acquisition and modeling of material appearance

Detailed surveys on appearance acquisition and modeling can be found in the literature [8–11]. Widely used reflectance models that capture spatially varying material characteristics under varying viewing and illumination conditions include spatially varying bidirectional reflectance distribution functions (SVBRDFs) [1] and bidirectional texture functions (BTFs) [3]. In contrast to SVBRDFs, BTFs are not necessarily defined with respect to the material's actual surface. Indeed, often a planar reference geometry is assumed, as for many materials like irregular fabrics it is difficult or practically impossible to determine the exact surface geometry with conventional acquisition setups. As a result, SVBRDFs do not accurately capture the light exchange for such materials. Moreover, BTFs do not impose restrictions regarding energy conservation on the per-texel BRDFs and simply encode the *appearance* of the material at one particular coordinate on the reference geometry, which is why they are known as *apparent* BRDFs (ABRDFs) [12]. Together with the parametrization over a flat geometry this allows capturing non-local effects such as interreflections, self-occlusions and self-shadowing. As measured SVBDRFs can be considered a subclass of BTFs, we shall focus on BTFs in this work. Due to their generality, BTFs are impossible to model, which is why one typically retreats to image-based representations that can be evaluated through a (possibly interpolated) table look-up. Measured BTFs have natural representations as matrices $\mathbf{B} \in \mathbf{R}^{n_{lv} \times n_{tx}}$, where the rows correspond to linearized light- and view-dependent 2D textures, the columns to linearized per-texel apparent BRDFs (cf. Fig. 2), n_{tx} denotes the number of texels (incorporating color channels for brevity) and n_{lv} the number of pairs of incoming and outgoing light directions under consideration. Note that in order to avoid interpolation artifacts, it is desirable that n_{lv} be large, in the order of thousands or tens of thousands, which in practice translates to the expensive process of acquiring tens of thousands of images of a given material. Given the matrix representation, both existing methods to mitigate this problem and the proposed method can be written concisely in terms of matrix operations, as we shall detail on in the following.

2.2. Sparse reflectance acquisition

Seminal work on sparse reflectance acquisition has been introduced by Matusik et al. [13] with the introduction of a new reflectance model that represents materials in terms of linear combinations from a set of densely sampled BRDF measurements. In subsequent work, Matusik et al. [14] approach the sparse reconstruction of isotropic BRDFs based on both a wavelet basis and a linear model obtained from the MERL database of isotropic BRDFs. They, however, did not investigate generalizations of these results to more complex reflectance models. In a closely related work, BTF compression was approached by Koudelka et al. [15], where single linear models for apparent BRDFs have been computed per-material. So far, the only technique that focuses on sparse reconstruction of BTFs has been presented by den Brok et al. [5]. Similar to the technique presented by Matusik et al. [13], a linear model is derived by applying singular value decomposition on a (logarithmically transformed) database of ABRDFs. By fitting these models to small BTF patches, non-local effects of light exchange are taken into account and BTFs have been reconstructed from only 6% of the typically used view-light conditions without a reduction of the resolution determined by the acquisition setup. Furthermore, manifold bootstrapping has been introduced by Dong et al. [16], where a manifold is constructed from analytical BRDFs fitted to BRDF measurements of certain manually selected surface positions on the material sample and used for the reconstruction of anisotropic SVBRDFs from sparse measurements. While a generalization to previously unseen materials might be achieved, the significant increase of the dimensionality of the manifold of per-texel reflectance distribution functions makes this technique impractical for BTFs. Nielsen et al. [6] present an approach for BRDF reconstruction from an optimized sparse sampling, where an improved logarithmic mapping of the MERL database is employed to obtain