ACCURATE RECOVERY OF A SPECULARITY FROM A FEW SAMPLES OF THE REFLECTANCE FUNCTION

Gilles Baechler, Ivan Dokmanić, Loïc Baboulaz, Martin Vetterli

School of Computer and Communication Sciences, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland e-mail: {gilles.baechler, ivan.dokmanic, loic.baboulaz, martin.vetterli}@epfl.ch

ABSTRACT

We present a new technique for estimating the specular peak of the bidirectional reflectance distribution function (BRDF) based on finite rate of innovation (FRI) sampling. The specular component of the BRDF varies rapidly, so it is challenging to acquire it by pointwise sampling. Yet, the knowledge of its precise location is key to render realistically complex materials. We show how to adapt the FRI framework to accurately determine the location of a single pulse when the sampling kernel is unknown. We use this result to determine the position of the specularity, and then estimate its shape by non-linear optimization. We demonstrate the feasibility of our approach in simulations and via a practical experiment using a custombuilt BRDF acquisition device.

Index Terms— Bidirectional Reflectance Distribution Function, Finite Rate of Innovation, Specularity, Sampling

1. INTRODUCTION

Traditional images, which are 2D projections of real scenes, do not capture all physical characteristics of the light. For example a 2D image does not allow relighting of a scene or a change of perspective. The bidirectional reflectance distribution function (BRDF) describes how materials interact with light and is one step towards richer representations of the real world. In its elementary form, a BRDF can be decomposed into two parts: the *diffuse* (or *Lambertian*) part — a low-frequency component — and the *specular* part — a high-frequency peak whose direction has the same angle as the incoming light with respect to the surface normal. The diffuse part is relatively easy to acquire and measure, but the acuity of the specular component requires a high sampling rate to ensure a faithful reconstruction.

We take inspiration from finite rate of innovation (FRI) sampling to develop a low sample density scheme to identify the specularity. FRI is a sampling framework allowing, from only a few samples, exact reconstruction of parametric functions that are not necessarily bandlimitted [1, 2]. In BRDF acquisition, we do not know the shape of the specularity beforehand, which would be required for traditional FRI techniques. We show in this paper how to generalize FRI to *unknown* sampling kernels, allowing us to estimate the specularity location. Even though we concentrate on the single pulse model, our goal is also to set up a sampling framework to enable acquisition of materials with *multiple* specularities, which typically appear on rough surfaces.

2. BACKGROUND ON LIGHT REFLECTION

The BRDF $\rho(\boldsymbol{w}_i, \boldsymbol{w}_o)$ is a four-dimensional function that measures the power of light that is reflected in the viewing direction $\boldsymbol{w}_o = [\theta_o, \varphi_o]^T$ given a light source with incoming direction $\boldsymbol{w}_i = [\theta_i, \varphi_i]^T$ (see Fig. 1). The BRDF is specific to the wavelength and in practice we often use a separate reflectance function for each spectral band acquired; for the sake of clarity however, we focus here on a single channel.



Fig. 1: Vectors involved in the BRDF function: ω_i is the incoming light direction, ω_o is the outgoing light direction, and n is the surface normal. The vectors ℓ_x and ℓ_y represent the projection of ω_i onto the xy plane.

A wide range of parametric models have been proposed to represent BRDFs [3]: in general, the diffuse part is described by a constant term and the specularity is approximated by a function with a few arguments. A relatively simple yet widely used BRDF model is due to Ward [4]. It exists in two flavors, one for isotropic and another for anisotropic specularities. The isotropic version is given by

$$\rho(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \frac{k_d}{\pi} + \frac{k_s}{4\pi\alpha^2 \sqrt{\langle \boldsymbol{n}, \boldsymbol{\omega}_i \rangle \langle \boldsymbol{n}, \boldsymbol{\omega}_o \rangle}} e^{-\frac{1}{\alpha^2} \left(\frac{1}{\langle \boldsymbol{n}, \boldsymbol{h} \rangle^2} - 1\right)},$$

where k_d is the diffuse reflection coefficient, k_s is the specular reflection coefficient, α is the specular roughness and $\mathbf{h} = \frac{\omega_i + \omega_o}{\|\omega_i + \omega_o\|}$ is the *halfway vector* between ω_i and ω_o .

To understand how the BRDF transforms the incoming light, we need one more ingredient: the *rendering equation*. It quantifies the amount of light L_o emitted from a point along the direction ω_o by integrating the product of the BRDF ρ and the incoming light L_i for

Authors are with LCAV-EPFL. This work was supported by the Swiss National Science Foundation grant : Inverse Problems Regularized by Sparsity, no. 20FP21151073.

all possible directions $\boldsymbol{\omega}$ on the hemisphere Ω :

$$L_o(\boldsymbol{\omega}_o) = \int_{\Omega} L_i(\boldsymbol{\omega}) \rho(\boldsymbol{\omega}, \boldsymbol{\omega}_o) \langle \boldsymbol{n}, \boldsymbol{\omega} \rangle d\boldsymbol{\omega}.$$
(1)

The surface normal of the material is denoted by n and the inner product $\langle n, \omega \rangle$, which is often absorbed in ρ , is called the *normal attenuation*. It accounts for the fact that objects appear darker when the angle between the light beam and the surface gets closer to $\frac{\pi}{2}$.

Despite its generality, the rendering equation does not take into consideration every aspect of the light's behavior. For instance, it does not capture inter-reflections, transmission, subsurface scattering or polarization. Nevertheless, (1) is a good approximation of the physics of light without being too complex.

2.1. BRDF acquisition

BRDF acquisition has been widely studied in the computer graphics community and is often performed using a *gonioreflectometer* a device that takes several photographs of the material or object of interest under different illuminations.

In most approaches [3, 5], the light source is pointwise, i.e., $L_i(\omega)$ is represented as a Dirac delta; this allows us to get rid of the integral in (1) and provides a direct reading of ρ for a specific incoming light direction ω_i . The main drawback of pointwise light sources is that the characterization of the narrow specular component requires a large number of samples. For example, in [3], more than a million values are used to accurately describe a single material.

One of the approaches to reduce the number of samples in BRDF acquisition and still faithfully capture the specularity is to extend the area illuminated by the light source. The earliest example of extended light sources shaped the radiance of the light as a Gaussian distribution and modeled the BRDF with a cosine and a Dirac pulse [6]. More recently, [7] suggested to acquire spatially-varying BRDFs of an almost flat surface using a *linear* light source and by only varying the inclination of the incoming light θ_i . Light sources can also take more elaborate patterns, and enable BRDF acquisition in a different domain, such as the spherical harmonic domain [8] or the Fourier domain [9].

3. FINITE RATE OF INNOVATION SAMPLING

The archetypal signal with a finite rate of innovation is a τ -periodic stream of several Diracs delta functions. Because of space limitations, we focus here on the *single* pulse case, i.e., signals of the form

$$x(t) = \sum_{k \in \mathbb{Z}} a_0 \delta(t - t_0 - k\tau),$$

where t_0 represents the location of the pulse and a_0 its amplitude. We now review how to recover t_0 and a_0 using at least 2 Fourier series coefficients:

$$X[m] = \frac{1}{\tau} a_0 e^{-i2\pi t_0 m/\tau}, \qquad m \in \mathbb{Z}.$$
 (2)

One of the key elements in the FRI theory is the search for an annihilating filter A such that $(A * X)[m] = 0, \forall m \in \mathbb{Z}$. The annihilating filter coefficients are found by solving the linear system of equations

$$A[0]X[m] + A[1]X[m-1] = 0,$$

for m = 1, ..., M - 1 and $M \ge 2$. It can be shown [1] that any signal X[m] of the form $a_0 u_0^m$, with $a_0 \in \mathbb{R}$ and $u_0 \in \mathbb{C}$ can be annihilated by the filter with z-transform

$$A(z) = A[0] + A[1]z^{-1} = 1 - u_0 z^{-1}.$$
 (3)

From (2), we see that the knowledge of the *root* $u_0 = e^{-i2\pi t_0/\tau}$ of the *z*-transform of the filter annihilating X[m] enables the recovery of the time location of the Dirac pulse:

$$t_0 = -\frac{\tau \angle u_0}{2\pi},\tag{4}$$

where $\angle u_0$ denotes the phase of u_0 . Once we know t_0 , we can retrieve a_0 using (2). When dealing with noisy data or model mismatch, two methods borrowed from the spectral estimation literature are generally used: Cadzow's method [10] and the matrix-pencil algorithm, which is based on *ESPRIT* [11].

FRI is not limited to Dirac pulses; an example is [12], where the Dirac pulse model is extended to include width and asymmetry. Furthermore, FRI signals can be sampled with various kernel shapes; as demonstrated in [2], signals sampled with polynomials or exponential reproducing kernels can be recovered with FRI. Note that the parameter estimation in these methods requires to know the signal model and the sampling kernel. In the next section, we show that t_0 can be recovered without necessarily knowing the kernel shape.

4. FRI WITH UNKNOWN SAMPLING KERNEL

Now let us assume that x(t) is convolved with an *unknown* τ -periodic filter $\phi(t)$; the only constraint we impose on $\phi(t)$ is that it must be *symmetric* (the reason of this is made clear later).

Let $y(t) = (x \otimes \phi)(t)$. By the convolution theorem, the Fourier series coefficients of y(t) are given by

$$Y[m] = \frac{1}{\tau} a_0 u_0^m \Phi[m], \qquad m \in \mathbb{Z},$$

where $\Phi[m]$ are the Fourier series coefficients of $\phi(t)$.

As in Equation (3), we can find an annihilating filter for Y[m] of length two: A[0] = 1 and $A[1] = -v_0$, where v_0 is the root of the filter. Using $M \ge 2$ Fourier series coefficients, we can estimate v_0 by solving

$$\boldsymbol{y_1} = \boldsymbol{v_0} \boldsymbol{y_0}, \tag{5}$$

with

$$\boldsymbol{y}_{0} = \begin{bmatrix} Y[0] \\ Y[1] \\ \vdots \\ Y[M-2] \end{bmatrix} \text{ and } \boldsymbol{y}_{1} = \begin{bmatrix} Y[1] \\ Y[2] \\ \vdots \\ Y[M-1] \end{bmatrix}.$$

The closest approximation of v_0 in (5) is simply given by the orthogonal projection of y_1 onto y_0 :

$$\begin{aligned} v_{0} &= \frac{\mathbf{y_{0}}^{*} \mathbf{y_{1}}}{\mathbf{y_{0}}^{*} \mathbf{y_{0}}} \\ &= \frac{\sum_{m=1}^{M-1} \left(\alpha_{0} \Phi[m-1] u_{0}^{m-1} \right)^{*} \left(\alpha_{0} \Phi[m] u_{0}^{m} \right)}{\sum_{m=0}^{M-2} \left(\alpha_{0} \Phi[m] u_{0}^{m} \right)^{*} \left(\alpha_{0} \Phi[m] u_{0}^{m} \right)} \\ &= \underbrace{\sum_{m=1}^{M-1} \left(\Phi[m-1]^{*} \Phi[m] |u_{0}^{m-1}|^{2} \right)}_{\sum_{m=0}^{M-2} \left(|\Phi[m]|^{2} |u_{0}^{m}|^{2} \right)} _{\text{purely real}} u_{0}. \end{aligned}$$

The key observation is that, since the sampling kernel is assumed to be symmetric, its Fourier series coefficients are real. Hence the expression multiplying u_0 is a real quantity, which implies that $\angle v_0 = \angle u_0$. In other words, the filter annihilating Y[m] allows us to retrieve the location t_0 from its root using Equation (4), as long as the unknown filter $\phi(t)$ is symmetric. As a final remark, note that even though the kernel shape is not necessary in the estimation of the location, it is needed for the recovery of the amplitude a_0 : it is simply impossible to differentiate the contribution of the kernel $\phi(t)$ from the contribution of the signal x(t) by observing y(t) only.

4.1. Extension to 2D signals

The generalization of the above result to 2-dimensional signals is straightforward if we assume that the sampling kernel is separable; i.e., $\phi = \phi_x * \phi_y$, where ϕ_x is a horizontal and ϕ_y a vertical filter (see [13] for more details). The separability of the sampling kernel combined with the convolution theorem yields the following expression for the 2D Fourier series coefficients:

$$Y[m,n] = \frac{1}{\tau^2} a_0 e^{-i2\pi x_0 m/\tau} e^{-i2\pi y_0 n/\tau} \Phi_x[m] \Phi_y[n], \quad m,n \in \mathbb{Z},$$

where (x_0, y_0) is the location of the Dirac pulse and a_0 its amplitude. We can reduce the problem to the 1-dimensional case by selecting the coefficients Y[m, 0] for $m \in \mathbb{Z}$:

$$Y[m,0] = \frac{1}{\tau^2} A_0 e^{-i2\pi x_0 m/\tau} \Phi_x[m],$$

with $A_0 = a_0 \Phi_y[0]$. In the same manner, we can estimate y_0 using the coefficients Y[0, n] for $n \in \mathbb{Z}$. We need $M \ge 2$ Fourier coefficients in each dimension to find the coordinates (x_0, y_0) .

4.2. Noisy case

We now assess the performance of the FRI estimation procedure in the presence of noise with the following simulation. We generate a single Dirac with a random location $t_0 \in [0, \tau)$, convolved with a Gaussian kernel of variance $\sigma^2 = 0.05$. We sample y(t) as

$$y_s[n] = \langle y(t), h(t/T - n) \rangle$$
 $n = 0, 1, ..., 2M - 1$

where $T = \frac{\tau}{2M}$ and h(t) is an ideal lowpass filter whose role is to avoid aliasing in the low-frequency Fourier coefficients. We then corrupt the samples $y_s[n]$ with additive white Gaussian noise (AWGN) of variance σ_{ϵ}^2 . We compare the efficiency of the FRI estimation procedure with the Cramér-Rao bound; given 2M discrete uniform samples, the variance var (\hat{t}_0) of any unbiased estimator \hat{t}_0 of t_0 is lower-bounded by CRB (\hat{t}_0) . In the AWGN case and with a single unknown parameter, the CRB is given by

$$\operatorname{CRB}(\hat{t}_0) = \frac{\sigma_{\epsilon}^2}{\sum_{n=0}^{2M-1} \left(\frac{\partial y_s[n]}{\partial t_0}\right)^2}$$

Results of the simulation can be seen in Fig. 2 for M = 2 and 8; we consider the FRI-based estimator to be fully efficient for an SNR above 10dB. Note that our FRI implementation is coupled with ESPRIT [11] to reduce the effect of the noise.

5. APPLICATION TO SPECULARITY ESTIMATION

We show here how to apply the FRI theory to reconstruct a BRDF from its filtered version and discuss the technicalities of the application of the algorithm.

First, we reduce the number of dimensions of the BRDF from 4 to 2 by fixing the viewing angle ω_o . Furthermore, we project the normalized incoming light vector ω_i onto the xy plane, so that we can make use of the Fourier theory — the foundation of FRI. We call the 2 components of the projection ℓ_x and ℓ_y (cf. Fig. 1).



Fig. 2: Sampling a single noisy τ -periodic Dirac convolved with a Gaussian kernel ($\tau = 1, \sigma = 0.05$). Using 2M = 2.2 and 2.8 samples, we measure the disparity $\Delta t_0/\tau$ between the FRI-estimated \hat{t}_0 and t_0 for a given input SNR (dotted lines) and compare it with the CRB (solid lines). Each data point is the average of 10,000 repetitions.

As stated in Section 2.1, the extended light source approach allows us to smoothen the specularity. The resulting acquired signal can be seen as the convolution of a Dirac pulse with two kernels: its (narrow) pulse shape and the extended light kernel. While we have some control on the shape of the extended light kernel (if it is not known, we can at least measure it), the specular pulse shape is *unknown*. Fortunately, both kernels are symmetric, which makes the specular peak a suitable candidate for our algorithm.

It should be noted that the BRDF of most materials is a mixture of a diffuse and a specular component. Since we are only interested in the specular part, we need to subtract the diffuse component from the acquired signal; luckily, the diffuse component is wide and can be easily estimated, for instance with a few pointwise light sources.

5.1. Simulations

We run numerical experiments in 1 and 2 dimensions to demonstrate the feasability of our method. First, we generate a BRDF function ρ that follows a Ward distribution with $k_d = 1$, $k_s = \frac{1}{1000}$, a random surface normal vector n, and $\alpha = 0.008$, which corresponds to a relatively narrow specular peak (see Fig. 3.A). Note that the specular peak of the Ward model is not a separable kernel, but our experiments have shown that it is close to it (it can be approximated with a Gaussian function, which is separable).

The signal ρ is then convolved with a Gaussian pulse to simulate the effect of a diffuse light and then sampled at a relatively low rate: 50 samples are taken in the 1D case and 40 samples in each dimension in the 2D case (Fig. 3.B). Recall that the theory states that the smoothing kernel should be periodic; this is not the case in our experiments. This difference can however be neglected as the kernels we use decay fast enough.

The diffuse component parameters (surface normal n and diffuse albedo k_d) are estimated with photometric stereo [14] from a lowpass version of ρ . After we subtract the (filtered) diffuse component, we are left with a low-frequency version of the specular peak (Fig. 3.C). We run the FRI algorithm on its DFT coefficients to obtain the location of the specularity.

We then estimate the width of the specularity; this is a nonlinear problem which we tackle with the Levenberg-Marquardt algorithm. Recall that this procedure requires to know the extended light shape. In practice, it can be measured by lighting a purely specular material with the desired light kernel from several locations. The rendering equation (1) shows that when ρ is a Dirac pulse, L_o directly provides pointwise samples of the shape of the light.

Finally, the BRDF is reconstructed using the recovered parameters (Fig. 3.D). Visually, the estimation matches closely the original



Fig. 3: Example of the estimation of the BRDF parameters with FRI in the 1D (top) and 2D (bottom) cases. (A) original BRDF ρ ; (B) signal lowpass-filtered with a Gaussian kernel ($\sigma = 0.5$ in 1D and $\Sigma = [0.3, 0; 0, 0.3]$ in 2D) and sampled; (C) lowpass specularity after subtraction of the diffuse part; (D) Reconstructed signal: the specularity is estimated with FRI and the diffuse component with photometric stereo. The parameter θ_i represents the inclination of the incoming light vector and (ℓ_x, ℓ_y) are the coordinates of its projection onto the *xy* plane.

shape. Quantitatively, we measure an SNR of 43.4 dB in the 1D case and 52.2 dB in the 2D case.

6. BRDF ACQUISITION WITH A CUSTOM-MADE DEVICE

To confirm the above results, we perform a practical acquisition with a custom-made light dome; it consists of an extensible motorized arm with a light source at its extremity. The arm can be rotated around the center, where a camera is installed (see Fig. 4). This enables us to position the incoming light at almost any location on the hemisphere.



Fig. 4: Light dome for BRDF acquisition: (left) draft and (right) photograph of it capturing the spatially varying BRDF of a painting.

We propose to localize the specularity of a flat mirror (its diffuse component is considered nil and it does not produce interreflections). We position it with an angle with respect to the camera plane such that t_0 does not correspond to an inclination of 0.

We focus the camera on the surface of the mirror and take 100 photographs of it; for each capture, we vary the inclination angle θ_i of the illumination. Moreover, we use both (almost) pointwise and extended light sources; the extended kernel is obtained by placing a diffusion filter of unknown shape in front of the light. For a single pixel selected in the middle of the frame, the two captured BRDFs are shown in Fig. 5. By taking only 20 samples from the filtered BRDF, we can effectively recover the location of the specularity with FRI; this is clearly impossible from 20 pointwise samples.



Fig. 5: 1D BRDFs of a mirror with pointwise (continuous line) and extended (dashed line) light source, where θ_i is the inclination of the light direction. The location (*) is estimated from 20 samples (•).

7. CONCLUSION

We have proposed a new FRI-based technique to estimate the location of a single Dirac pulse convolved with an *unknown* symmetric sampling kernel. We applied the method to the problem of identifying the specular component of the BRDF in numerical simulations and in a practical experiment. In both cases, we showed that the specularity is well approximated from a small number of samples. We also expect that the sampling framework we put in place will enable us to extend the method to materials with multiple specularities.

One of the limitations of the presented approach is that it essentially uses a projection of the light vectors onto the xy plane. Instead, it would be beneficial to directly perform the estimation in the spherical domain; recent work [15] has indeed shown that the FRI scheme can be applied on the sphere as well.

8. ACKNOWLEDGEMENTS

The authors would like to thank Gaël Soudan for his work on the light dome and Adam Scholefield for his comments and advices.

9. REFERENCES

- M. Vetterli, P. Marziliano, and T. Blu, "Sampling Signals with Finite Rate of Innovation," *IEEE Transactions on Signal Processing*, vol. 17, no. 6, pp. 1417–1428, 2002.
- [2] P. L. Dragotti, M. Vetterli, and T. Blu, "Sampling Moments and Reconstructing Signals of Finite Rate of Innovation: Shannon meets Strang-Fix," *IEEE Transactions on Signal Processing*, vol. 55, no. 5, pp. 1741–1757, 2007.
- [3] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A Datadriven Reflectance Model," ACM Transactions on Graphics, vol. 22, pp. 759, 2003.
- [4] G. J. Ward, "Measuring and Modeling Anisotropic Reflection," ACM SIGGRAPH Computer Graphics, vol. 26, no. July, pp. 265–272, 1992.
- [5] M. L. Koudelka, S. Magda, P. N. Belhumeur, and D. J. Kriegman, "Acquisition, Compression, and Synthesis of Bidirectional Texture Functions," *Texture 2003: Third International Workshop on Texture Analysis and Synthesis*, pp. 59–64, 2003.
- [6] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Determining Shape and Reflectance of Hybrid Surfaces by Photometric Sampling," *IEEE Transactions on Robotics and Automation*, vol. 6, no. 4, pp. 418–431, 1990.
- [7] A. Gardner, C. Tchou, T. Hawkins, and P. Debevec, "Linear Light Source Reflectometry," ACM Transactions on Graphics, vol. 22, pp. 749, 2003.
- [8] B. Tunwattanapong, G. Fyffe, P. Graham, J. Busch, X. Yu, A. Ghosh, and P. Debevec, "Acquiring Reflectance and Shape from Continuous Spherical Harmonic Illumination," ACM *Transactions on Graphics*, vol. 32, no. 4, pp. 109:1–109:12, 2013.
- [9] M. Aittala, T. Weyrich, and J. Lehtinen, "Practical SVBRDF Capture in the Frequency Domain," ACM Transactions on Graphics, vol. 32, no. 4, pp. 1, 2013.
- [10] J. A. Cadzow, "Signal Enhancement a Composite Property Mapping Algorithm," *IEEE Transactions on Acoustics, Speech,* and Signal Processing, vol. 36, no. I, 1988.
- [11] R. Roy and T. Kailath, "ESPRIT Estimation of Signal Parameters via Rotational Invariance Techniques," *IEEE Transactions* on Acoustics, Speech, and Signal Processing, vol. 37, no. 7, pp. 984–995, 1989.
- [12] G. Baechler, N. Freris, R. F. Quick, and R. E. Crochiere, "Finite Rate of Innovation Based Modeling and Compression of ECG Signals," *IEEE International Conference on Acoustics, Speech* and Signal Processing (ICASSP), pp. 1252–1256, 2013.
- [13] I. Maravić and M. Vetterli, "Exact Sampling Results for Some Classes of Parametric Nonbandlimited 2-D Signals," *IEEE Transactions on Signal Processing*, vol. 52, no. 1, pp. 175–189, 2004.
- [14] D. A. Forsyth and J. Ponce, Computer Vision: a Modern Approach, 2003.
- [15] I. Dokmanic and Y. Lu, "Sampling Sparse Signals on the Sphere: Algorithms and Applications," *IEEE Transactions on Signal Processing*, vol. 99, 2015.