



Physically Based Rendering of Simple Thin Volume Natural Nanostructures

Daljit Singh J. Dhillon^(✉)

School of Computing, Clemson University, Clemson, SC 29634, USA
djsingh@clemson.edu

Abstract. Thin volumes of semi-transparent nanostructures present on outer layers of organic embodiments often interact coherently with incident light waves to produce nuanced structural coloration. Such mechanisms are further complicated through incoherent scattering by accompanying micro-geometries. We present a simple physically based approach to directly use the sub-microscopic scans of quasi-periodic, one-dimensional modulations in such volumes to render them realistically. Our method relies on prior knowledge of quasi-periodicity to process the scan data in the Fourier space for recreating nuanced coloration effects. We demonstrate the working of our method with the actual scanned data of an egg-sac that shows coloration only when immersed in water. Proposed method can be used by bio-physicists for visual conformity of such mechanisms at a macro-scale as well as graphical rendering pipelines can employ it for scientific recreations or artistic renditions.

Keywords: Structural colors · Physically based rendering · Wave optics · Wave interference · Diffraction · Photonics · Natural phenomena · Sub surface scattering · Volumetric rendering

1 Introduction

Several natural surfaces are made up of organic matter consisting of thin, composite layers of translucent nanostructures (see Fig. 1). These thin-volumes exhibit a mix of coherent and incoherent interactions with the incident light waves. Many of these nanostructural compositions have geometric or statistical regularities to bring about stationary interference that result in magnificent, iridescent structural colors. Common examples include insect wings, bird feathers, animal or fruit skins and pearls. Devising photo-realistic models for reproducing their appearance characteristics is still largely an open research problem. Efficient computational methods find applications in domains ranging from computer graphics (VFX effects, animations, gaming) to scientific studies by bio-physicists and even in medicine.

In general, we need to solve Maxwell's equations to model the wave interactions under discussion for reproducing stationary interference patterns. More importantly, we need to represent and appropriately apply coherence spans (spatial as well as temporal) for the incident light over the nanostructural volumes




Fig. 1. Images showing structural coloration and microscopic scans for an egg-sac. Reproduced from [16] with permission from the Royal Society of Chemistry. (Color figure online)

to separate out the coherent and incoherent interactions for accurate modeling. Several methods such as those based on Kirchhoff integral theorem [4, 5, 7, 12] and Harvey-Shack model for diffraction [3, 6, 8, 15] have been proposed in existing literature to model simple diffraction mechanisms. Some even make good use of the spatial coherence properties of sunlight for photo-realistic reconstruction [4, 15]. However, for complex volumes of nanostructures with varying optical properties, existing methods mainly rely on computationally expensive *finite difference time domain* (FDTD) based formulations [10] or full-blown light transport frameworks with coherence/phase propagation mechanism [11, 13] for accurate reconstructions. Furthermore, evaluating such volumetric formulations involve an added challenge in generating or acquiring the volumetric data for natural compositions in order to model their nuanced light interactions accurately.

In this paper, we propose a physically based method to deal with a class of thin volume nanostructures that exhibit quasi-periodic modulations along a single dimension. Furthermore, these structures have a layered composition of similar modulations in each layer. The layering process may lend micro-geometric variations that cause incoherent light-scattering within the thin volume. This results in glossy specular highlights and blurring of the structural coloration patterns that emerge from coherent interference due to the nanostructural modulations. Our method uses prior information of the mean periodicity for the modulations to perform Fourier domain filtering of the structural information as explained later, in detail. It further accommodates multi-layer effects using simple image processing techniques. In summary, our main contribution in this paper is a novel physically based rendering method that: (a) adapts actual sub-microscopic scan data of a single representative layer, (b) for thin volumes with nanostructural layered modulations, (c) and simplifies application of Fourier optics to accurately reproduce their nuanced structural coloration, (d) at interactive speeds.

2 Thin Volume Nanostructures and Layer Surface Modulations

Thin volumetric nanostructural compositions may occur in several configurations. At a broader level, those are classified as 1D-, 2D- or 3D- photonic structures depending on the number of dimensions over which they exhibit observable modulations. In lack of any semblance of periodic modulations, they are generally referred to being stochastic. The most basic nanostructural compositions

are made up of a single layer that causes off-surface wave diffraction to produce structural colors. Such surfaces are easy to model. Unlike diffractive surfaces, thin volumes may be organized into multilayers of same or near similar structures. Furthermore, they may exhibit shape modulations with quasi periodicity along multiple dimensions to cause volumetric wave interactions and complex interference patterns. Apart from structural complexities, the material (optical) properties of the basic components in the composition may as well vary spatially. While one needs a generic implementation like an FDTD solver for an arbitrary composition, we can always exploit some prior knowledge about the overall structural layout to develop an efficient rendering model. With this generic strategy we devise a physically based model that can be used to render structural colors at interactive rates for thin volume nanostructures that satisfy the following criteria for its layout: (a) the whole thin volume is composed of basically the same material, (b) the volume is essentially laid out in multiple layers that are structurally similar to each other, (c) each layer exhibits quasi-periodic modulations at nanoscale, (d) the structure of an individual layer is available through scanning or construction, and (e) the range for quasi periodicities is known.

As a concrete example, we work with the data for an egg-sac (Hynobius Kimurae) provided by Zabuga et al. [16]. This egg-sac does not exhibit any structural coloration when placed in open air (see Fig. 1b). However, it shows nuanced blueish-green and greenish-yellow colors in peripheral regions when placed in water as shown in Fig. 1a. The egg-sac has a tubular structure and it is made up of translucent fibrous elements that are individually $2\text{--}5\mu\text{m}$ thick and tightly packed into layers ranging in their depth from $350\text{--}1000\text{nm}$. The total thickness of the egg-sac is about $100\mu\text{m}$. We assign a local frame of reference to each surface point on the sac where the X -axis is tangential to the circular cross-section, the Y -axis runs along the sac length (top to bottom) and the Z -axis traverses the sac depth. A closer, microscopic examination performed by Zabuga et al. [16] using a *focused ion beam-scanning electron microscope* (FIB-SEM) scanner reveals the volumetric structure as shown in Fig. 2. They discuss the nano-structure with painstaking details. Importantly, each layer of fibrous composition exhibits fine, quasi-periodic folds. These folds vary spatially in their




Fig. 2. Volumetric FIB-SEM scan slices for the egg-sac. White bar $\rightarrow 1\mu\text{m}$.




Fig. 3. An AFM scan of the inner surface of the egg-sac shown as a height-field. The grayscale for the inset images spans the height-range from 0–502 nm. (a) is the raw scan image. (b) Spatial low-pass filtering removes frequencies causing diffraction in the visible spectrum. (c) Spatially filtered nano-gratings.

mean periodicity from 170–230 nm. In this local frame of reference, the quasi-periodic 1D modulations are primarily oriented along the X -axis. An *atomic force microscope* (AFM) scan for the inner surface reveals these modulations as shown in Fig. 3.

3 Fourier Optics and Physically Based Rendering

For the type of quasi-periodic, multi-layer structures under discussion, we begin with the clarity that all layers are similar if not the same. We may thus think in terms of wave interference caused by these layers individually and the combined effect across them all. For natural compositions, the layers that lie close enough to cause coherent interactions, we found that such layers don't vary much structurally. We thus take an individual layer and consider it as the representative of local coherent wave interactions. This simplifies wave interference effects to be represented through a simple diffraction mechanism. This indeed is one of the main conclusions made by Zabuga et al. [16] in their detailed physical studies of such coloration effects in the case of egg-sacs. For diffraction caused by the roughness around a surface point, we can first measure and express surface bumps as a discretized height-field $h(x, y)$ in the local tangent plane at that point. We can then model the spectral *bidirectional reflectance distribution function* (BRDF) due to $h(x, y)$ with an application of Kirchhoff integral theorem [4, 12].

For layers that are mutually incoherent, their local coherent interference patterns add up in the Fourier power space [16]. We study the net impact of this Fourier power accumulation on the distribution profile for the Fourier power of the phasor function for Kirchhoff integral. We do these studies using the actual FIB-SEM data for the egg-sac case. Next, we define an image processing technique to emulate those effects over the Fourier power distribution computed from the height-field data for an individual layer. This physically based approach for incoherent mixing of coherent interference effects across crudely separated layer-sets that is both a computational simplification as well a data-reduction step. We thus can use a single 2D AFM scan for acquiring nanostructural details instead of requiring a full-blown 3D volumetric scan.




Fig. 4. Extracting nanostructural gratings through frequency-domain filtering.

Next, we consider the incoherent scattering due to micro-geometric variations in the layer’s height-field. They cause glossy specular highlights instead of sharp mirror like reflections. Figure 1(a) shows few such glossy highlights, for example. To this end, we propose to separate nano-structural variations from the micro-structural variations in a layer and use an approach similar to [6]. Next, we empirically approximate an isotropic Beckmann distribution through Blinn-Phong modeling of glossy reflectances due to micro-facets.

Finally, we model the incoherent scattering of the coherent diffractive collocation effects from lower layers as those reflections permeate through the micro-faceted structuring of higher-up layers. Thinking in terms of layered interactions, one may be motivated to accurately model these scatterings into a *bidirectional scattering distribution function* (BSDF). This would involve a convolution between the specular *bidirectional transmittance distribution function* (BTDF) and the diffractive BRDF forms [6]. In our method we propose using a simple light cone to emulate this convolution for mixing the back-scattered diffraction colors.

3.1 Proposed Method

We now present our proposed method in detail. Pseudo-code [1] gives the computational steps that a pixel shader needs to execute to render natural photonic structures under discussion. While this code-block is sequential, initial three steps and parameter estimation is performed in an offline manner to produce a lookup texture that can be used as at runtime. This allows us to achieve interactive performance for our method. Next, we explain the key steps in our proposed algorithm along with their application to the egg-sac case, in detail.

[STEP 1] Extracting Grating Structure. For the nanostructures under discussion, the key feature that result in nuanced appearances include natural meanderings of biologically laid out gratings that are quasi-periodic at large. Thus the variations in the periodicity and local orientations lend a richer, natural appearance to such compositions. To extract these features, we rely on prior knowledge of the mean periodicity of the key grating structures and variance in this periodicity measure. Using these two statistics, we devise a bandpass filter in the frequency space. While spatial domain filtering may produce a visually convincing representation of the key grating structures, operating in the

Pseudo-Code 1. Pixel Shader for Natural Layered Photonics

Input:

(a) Measured discrete height-field $h(x, y)$, (b) Light and view directions ω_i and ω_o , respectively, (c) Mean and variance in 1D/2D quasi-periodicities T_{mean} and T_{var} , (d) Surface normal ω_N , (e) Half-vector ω_H , (f) Medium and material refractive indices n_1, n_2 , respectively, and (g) Incident radiance $\mathbf{L}_R$ in RU units.

Initialize: $(u, v, w) = -(\omega_i + \omega_o)$.

Begin{Main}

— {*Offline Processing*}

Begin{Extract grating structure}

▷ See Section 3.1 Step(1) for details

1. Compute bandpass filter $\mathcal{B}(\omega_u, \omega_v)$ in *Frequency Space* using T_{mean} and T_{var} .

2. Filter height-field data in *Frequency Space* as: $\mathcal{H}_f(\omega_u, \omega_v) = \mathcal{B}(\omega_u, \omega_v) \cdot \mathcal{F}\{h(x, y)\}(\omega_u, \omega_v)$

End

Begin{Model surface diffractions with first-order approximations and a coherence window $\mathcal{W}$, See [4]}

$\mathbf{D}^\lambda(\omega_i, \omega_o) = |\mathcal{W} * \mathcal{H}_f(u/\lambda, v/\lambda)|^2$ ▷ ω_i :light direction, ω_o :view direction, λ : light wavelength

End

Begin{Consolidate multi-layer interference effects}

▷ See Section 3.1 Step(2) for details

1. Set blurr filter $\mathcal{M}(\omega_u, \omega_v)$

2. Compute $\mathbf{D}_m^\lambda(\omega_i, \omega_o) = (\mathcal{M} * \mathbf{D}^\lambda)(u/\lambda, v/\lambda)$

End— {*Of Offline Processing*}

Begin{Model specular highlight}

▷ See Section 3.1 Step(3) for details

1. Use (a) Blinn-Phong model or (b) Fit an isotropic Beckmann distribution for a micro-faceted specular BRDF: $\mathbf{S}(\omega_i, \omega_o) = (\omega_N \cdot \omega_H)^n$ or $\mathbf{S}(\omega_i, \omega_o) = \frac{\exp(-\tan^2 \alpha / \tan^2 \theta)}{\pi \tan^2 \theta \cos^4 \alpha}$, where $\alpha = \arccos(\omega_N \cdot \omega_H)$

End

Begin{Incorporate incoherent scattering of diffraction colors due to microfacet geometry} ▷ Sec.3.1 Step(4)

1. Estimate a Box blurr filter for light directions $\mathcal{G}(\omega_u, \omega_v)$ corresponding to $\mathbf{S}(\omega_i, \omega_o)$

2. Compute $\mathbf{D}_g^\lambda(\omega_i, \omega_o) = (w/\lambda)^2 (\mathcal{G} * \mathbf{D}_m^\lambda)(u/\lambda, v/\lambda)$ by sampling the light cone

End

Begin{Perform spectral rendering to compute pixel intensities}

1. Compute Fresnel factor using Schlick's approximation $\mathbf{F}_{raw} = R_0 + (1 - R_0)(1 - \cos \theta_i)^5$, where $R_0 = (n_2 - n_1)^2 / (n_2 + n_1)^2$, θ_i is the incident angle.

2. Normalize $F = F_{raw} / (R_0 \cdot w \cdot w)$ to work with relative units (RU) for radiance. ▷ See [4] for details.

3. Estimate geometric attenuation $\mathbf{G}(\omega_i, \omega_o)$ including a Cook-Torrance shadowing function

4. Set CIE-X,-Y,-Z color matching functions $C_i(\lambda)$, where $i \in \{X, Y, Z\}$

5. Set D65 spectral power distribution as $\mathbf{I}(\lambda)$

6. Compute pixel color in CIE-X,Y,Z color space as: $\mathbf{P}_i(\omega_i, \omega_o) =$

$$\int_{\Omega} \mathbf{L}_R(\omega_i) \cos \theta_i d\omega_i \int_{\lambda} \mathbf{I}(\lambda) \mathbf{C}_i(\lambda) (\mathbf{F} \cdot \mathbf{G} \cdot (\mathbf{S} + \mathbf{D}_g^\lambda))(\omega_i, \omega_o) d\lambda$$

7. Tone-map and gamma-correct pixel colors

End

End

frequency space is critical. For example, the FFT for spatially filtered grating structures shown in Fig. 3c still contains low frequency components as shown in Fig. 4c. Clearly, applying first-order approximations as used by [14] will result in systematic errors. We hypothesize that for the photonic thin-volumes of natural origin such lower frequency components that are far removed from the quasi-periodic frequency bands could in general be ineffective in producing stationary interference and need discounting. This hypothesis holds true for cases like the egg-sac. It is confirmed visually with the lack of any coloration while the sac is in air and through detailed physical experiments on it by Zabuga et al. [16]. For the representative case-study of the egg-sac, using mean periodicity $\mu = 230$ nm and standard deviation $\sigma = +/ - 72$ nm, we devise a frequency-domain, band-pass filter $\mathcal{B}$ with the pass-band $[\mu - \sigma, \mu + \sigma]$ radially. Figure 4d shows the filtered frequency response $\mathcal{H}_f$ for the egg-sac case. Using $\mathcal{H}_f$ as the first and only Fourier term in Equations (6) and (8) from [4] results in first-order approximation for diffraction effects from an individual layer as shown in Fig. 5b for the egg-sac data. Next, we look at consolidating these coloration effects across multiple layers.




Fig. 5. Rendering structural colors for the egg-sac with differently filtered height-fields. Filtered FFT is used in a spectral BRDF formulation [4]. Top row shows example BRDF slices for the light source at a grazing angle along X -axis. Refer to Sect. 4 (*Setup*) for more information on BRDF slices.

[STEP 2] Consolidate Coherent Interference from Multiple Layers.

For accumulating interference effects across multiple layers, we make note of two practical considerations: (a) for two layers close enough to cause coherent light wave interactions, their height-field profiles are very similar, and (b) for two layers that are far removed to cause coherent inter-layer light interactions, the spectral power distributions from their individual interference patterns add up to produce net effects like any two co-located, non-coherent light interaction events. For the class of photonic structures that have slow varying, 1D/2D quasi periodic layer surface modulations, both of the above assumptions hold true. These two assumptions also form the basis for the statistical analysis put to use in the actual physical experiments by [16]. With these two assumptions we are able to formulate a strategy for modeling multi-layer coherent interactions effectively. We assume that the whole volumetric mass can be approximated by non-overlapping slabs where the adjacent slabs only cause incoherent wave-interactions. Secondly, for each slab we assume that a single representative layer can be used to estimate the diffraction effects caused by the entire slab. These assumptions do not hold true in general. However, for the modulations that vary very slowly across the depth dimension, [16] empirically found these assumptions to be consistent with their physical experiments. In our method, we make use of the fact that the adjacent slabs are progressively but only slightly differing in their structural composition while remaining within the same range of quasi-periodic bandwidth. It implies that the Fourier transforms of the representative layers for individual slabs do not change much in the critical frequency band and adding up the resulting power spectrums is akin to low-pass filtering of the Fourier transform of an individual representative layer. We thus consolidate the

net effect of multiple layers by simply filtering the power distribution in the frequency space. See Pseudo-Code [1] and Sect. 4 for further details.

[STEP 3] Model Microfacet-Based Glossy Specular Highlight: $\mathbf{S}$. With the filtering process used in STEP_1 all micro-scale variations are removed from the height-field. [6] suggest that the micro-geometric and nano-geometric effects can be modeled separately and then superimposed in a BRDF. We note that for the achromatic, incoherent scattering behavior of our class of photonic structures, each individual layer is statistically similar and thus either of them is a good enough representative for the entire volume. We empirically fit a Blinn-Phong model to model the specular highlights.

[STEP 4] Accommodate Incoherent Scattering of Structural Colors. In one last final step, we take care of the nanostructural coloration effects getting modulated by the micro-structures. Since the extracted nano-structures are maneuvered by the underlying micro-structures, the diffraction colors arising from coherent scattering in the underlying layers are expected to undergo incoherent scattering as the diffracted light waves pass through the upper layers. As a simple method to accommodate these cascaded, mixed scattering events, we can expect the convolution of the achromatic, specular BRDF model $\mathbf{S}$ with the spectral diffractive coloration BRDF $\mathbf{D}_m^\lambda$ to suffice. To keep things even simpler, we instead emulate a low-pass filter kernel $\mathcal{G}$ that approximates $\mathbf{S}$ and convolve $\mathcal{G}$ with $\mathbf{D}_m^\lambda$.

4 Experiments and Results

In this section we present our various experiments for the filtering steps discussed in the previous section. We share the final implementation details that work the best and demonstrate the working of our method with the rendering results.

Implementation Details. Most of the implementation details are presented in the Pseudo-Code 1 for the pixel shader. Here we note that the diffraction BRDF is normalized to just saturate the reflectance from the perfectly planar surface of the same base material when viewed opposite to the surface normal with the light placed at the camera origin. This requires the incident light radiance to be expressed on a scale relative to this normalization step. We adopt the same nomenclature of expressing light intensity in *RU* or relative units as done by Dhillon et al. [4]. Likewise, for subjective comparability, we scale all final filtered FFT forms to have equal average power in the frequency band $\mu \pm \sigma$ for the mean periodicity μ and periodicity standard deviation σ . We use the total Fourier spectral power in the same band for a perfectly circular grating with periodicity same as given μ as the reference for scaling. We pre-filter all height-fields with a Blackman-Harris window to avoid windowing artifacts in the renderings. The height-field patch for the egg-sac is about 72×72 microns at a resolution of 3000×3000 pixels. We use the spatial coherence length with its *full width at half maximum* (FWHM) set at $65 \mu\text{m}$ for the sunlight [9].

We perform spectral integration at runtime with a sampling of 3–10 nm for most results.

Setup. Unless specified otherwise, we place a point light at the camera center itself for rendering surfaces. The tubular egg-sac is about 2 cm thick and 12.5 cm long. Its medial axis was traced in the reference image to reconstruct the surface mesh by hand. It is viewed from a close shot with the field of view of about 2° . The actual egg-sac is translucent. However, since our main focus is on iridescent back-scattering, we simplify the setup by assuming that it has an opaque base. Unless specified otherwise, we assume that the surface is submerged in water and use its refractive index $n = 1.33$ to shift the visible spectrum for computing correct diffraction effects. We show BRDF slices with fixed incident light directions ω_i as circular maps. The disk center corresponds to the view direction being fronto-parallel to the surface $\theta_o = 0$ for ω_0 . Each point on the map represents a view direction. The radial distance from this disk center to any point the map equals the sine of the angle for the view direction ($\sin \theta_0$) for that point. Also, the angle between the radial vector for that point and the X-axis equals the azimuth angle ϕ_o for the view direction represented at that point.

Experiments with Grating Approximations. We tried few simple experiments to study if the natural gratings can be reduced to some standard form. Using the known mean periodicity of $\mu = 190$ nm, we devised two regular gratings: (a) a 1D horizontal grating, and (b) a radial grating. We constructed their corresponding diffraction BRDFs and rendered the egg-sac geometry with it. Resulting BRDF maps are empty except for the specular highlight when the surrounding medium is air. Changing the medium to water only shows coloration in backscattering when the incident light is closer to the grazing angles. The light source for the BRDF slices shown in Fig. 6 is set horizontally at a grazing angle along the X-axis. For 1D grating, the scattered coloration is limited to only few view directions. As seen in Fig. 6a, it fails to produce any remarkable pattern over the egg-sac surface. A circular 2D grating produces back-scattered colors in a much larger range of view directions. It produces coloration in the peripheral regions for the egg-sac that are consistent with the photograph in Fig. 1a. However, the colors in Fig. 6b are too well-resolved and lack any nuances. Actual image shows meandered mixing of multiple wavelengths. We also changed the grating periodicity to 225 nm to match the variations in the inner curvature for the egg-sac. Figure 6c shows that doing so produces a greater range of visible colors when the specimen is in water. However, it lacks in wavelength mixing as well. We thus conclude that such simple structures cannot produce visually convincing renderings. The model needs to represent and use the large variations in natural gratings to produce realistic appearances. In our method, we are able to directly process the AFM height-fields for computing the diffraction part of the BRDF response as shown in Fig. 5c.

Experiments with the Bandpass Filtering Kernel $\mathcal{B}$. As discussed in Sect. 3.1 we need to filter the height-field $h(x, y)$ to extract the gratings. The importance of this step is discussed here along with several different kernels for




Fig. 6. Experiments with regular gratings as height-fields (Insets: BRDF slices).




Fig. 7. Experiments with the frequency domain filtering $\mathcal{B}$. Radial forms of ideal, Gaussian and Butterworth bandpass filters were devised with lower and upper cut-off frequencies derived from the mean periodicity μ and its variance σ^2 . (Color figure online)

the filter $\mathcal{B}$. Figure 5a shows that the raw data does not produce any similarity to the observed coloration effects when used directly in our method. Removing lower frequencies is impervious to coloration from the gratings. Also, not removing frequencies higher than those in the bandwidth defined by the mean periodicity μ and its standard deviation σ can produce unwanted visual artifacts. We thus device bandpass filters using known μ and σ statistics for the given height-field and apply them in Fourier space. We experimented with three different types of bandpass filters, namely: Ideal, Gaussian and Butterworth filters and our results are illustrated in Figs. 5 and 7. Figure 5b shows that an ideal filter with the bandwidth $\mu \pm \sigma$ provides good separation and coloration starts to resemble with observations. Although, the renderings are not soft enough and significant amount of streaking is produced. We increased the bandwidth to $\mu \pm 2\sigma$ but this causes the streaks to spread out unwantedly to adjacent regions (see Fig. 7a). Using a Gaussian filter with the lower cut-off corresponding to time-period $T = \mu + \sigma$ and the higher cut-off set at $T = \mu - \sigma$ fails completely. A Butterworth filter with the same stop-band works far better (Fig. 7c). However, its results are visually similar to those for the ideal filter as it does not extract any more useful information from the natural gratings that we work with.




Fig. 8. Experiments with the Gaussian kernel ($\mathcal{M}$) used in blurring of the frequency spectrum for consolidating multi-layer coherent scattering.

We thus resolve to use the simpler ideal filter in our pipeline. Figure 7d shows that for gratings with $\mu = 225$ nm we notice stronger yellowish colors at the egg-sac peripheries. This is the desired coloration for the inner curvature of the actual egg-sac which had a larger μ [16]. While extracting the gratings from a single layer works in principle, it shows a lot of color streaking. This is understandable as we know that the combined effect of multiple layers produces softer color mixes over the egg-sac surface.

Experiments with the Frequency-Space Blurring Kernel $\mathcal{M}$. As discussed in Sect. 3.1, we convolve the frequency spectrum magnitudes with a blurring kernel to consolidate for the variations and mixing in coherent interference coloration effects arising from different layers. To study these effects, we first used the volumetric data for the egg-sac. Slices of this data set are depicted in Fig. 2. We first compute the root-mean-squared (RMS) Fourier term magnitudes across all the $X - Y$ slices. The black solid curve in Fig. 8a shows this RMS Fourier spectral magnitude for the frequencies along the reference X-axis. The smoothness of this curve confirms structural similarity across the layers. Also, the peak seen at $f_s = \frac{1}{172} \text{ nm}^{-1}$ for the FIB-SEM sample, which shrank due to scanning preparation processes, is equivalent to $\frac{1}{210} \text{ nm}^{-1}$ for the living specimen. This peak is consistent and close enough to the mean grating period $T = 230$ nm for the outer curvature of the egg-sac (see [16] for details). Also, we use the volumetric data only to guide our multilayer related filtering process and not in the actual renderings. Next, we compute and show a similar curve for an individual layer in red color, in the same plot. Both the curves have similar profiles overall. The red curve has many exaggerated fluctuations that result in color streaking as seen in Fig. 7. We perform low-pass Gaussian filtering of the red curve with different σ_m values and smoother curves are shown in green, blue, magenta and orange colors. For $\sigma_m = 0.91 \text{ cpum}$, the (magenta) curve profile better matches with the reference (black). However, it is significantly lower in its peak and lack consolidated local fluctuations seen in the reference (black) curve. For, $\sigma_m = 0.41 \text{ cpum}$, the filtered (blue) profile tames and retains these local fluctuations while maintaining the peak at a reasonable

high. Deriving similar Gaussian filtering parameters from this study, we filtered the Fourier spectral magnitudes for the AFM height-field data and used resulting filtered FFTs $\mathcal{M} * \mathbf{D}^\lambda$ in rendering experiments. Figures 8b,c show softer color mixing due to such filtering. These renderings are photo-realistically much closer to observations. We found no significant advantage in using a Gaussian kernel $\mathcal{M}$ with $\sigma_m > 0.41\text{cp}\mu\text{m}$ and thus resolved to use this empirical value in our method. Finally, we note that filtering with $\mathcal{M}$ leaks some energy outside the stop-band for the filter $\mathcal{B}$. So we repeat bandpass filtering with $\mathcal{B}$ after this step.

Experiments with the Incoherent Scattering Model. For the specular highlight due to incoherent scattering, we fall-back to the empirical equivalence established between the Blinn-Phong NDFs and the Beckmann distribution to be used in a standard Cook-Torrance formulation for $\mathbf{S}$ (refer to Fig. 9.36 in the book *Real-Time Rendering* by Akenine-Möller et al. [2]). For the Blinn-Phong kernel, an exponent n between 250–2000 works well for the structures under consideration. We empirically set it to 500 for the egg-sac.

Also, to model the convolution between $\mathbf{S}$ and $\mathbf{D}_g^\lambda$ we use a screen-space approach by modeling the incident light with an elliptic cone. The span for the light cone along the incident plane is scaled to vary the projection of the light direction along the surface normal by a fixed proportion $p\%$. Similarly, the cone spans the off-plane perpendicular to the incident plane (while containing the light direction vector) to vary the projection of the light direction along the surface tangent-plane by a fixed proportion $p\%$. Apart from the original light direction, we draw 6 samples along the incident plane and 6 samples along the off-plane within the light cone. Setting p at 30% gave us the best smoothing of the remnant color streaks from the previous steps. Figure 9 shows the output of our incoherent scattering model for the egg-sac along side the actual photo. We positioned the specimen and the point light in the renderings as best to match the impressions from the photo. The photo has multiple light source while we use a single light source. Yet, the level of softness and color gamut for the back-scattered diffractions are subjectively similar between the rendering and the image. The specular highlights and the peripheral color bands occur in similar regions over the surface. We rendered the same specimen with several other positions, orientations, with alpha blending for a diffuse component and also with a different mean periodicity $T = 230\text{ nm}$. The renderings are shown in Fig. 9c.

Performance. We performed several runtime performance tests for our method on a Windows Surface Note Book 2 with an NVIDIA GPU GeForce GTX1050. It has a display memory (VRAM) of 1979MB and shared memory of 8150MB. We animated images similar to those in Fig. 9(c) including diffuse color blending and all the filtering steps discussed above and bench-marked performances using FRAPS [1]. Except for: (a) the spectral computations at a resolution of 10–50 nm, and (b) the screen-space light cone sampling with 13 samples, rest of the filtering operations are done offline for this exercise. Our display screen spanned a resolution of 2000×3000 pixels. With a simple rectangular animated object




Fig. 9. Final rendering results with all filtering steps incorporated in the method.

spanning between 800×600 sub-window, at spectral sampling of 50 nm, we achieved an average performance of 10.12 fps. With a resolution of 1300×900 pixels it dropped to 5.44 fps. With the spectral sampling rate of 10 nm it further dropped to 3.88 fps. For the egg-sac surface at the resolution of 1800×1600 it performed at about 2 fps. Here we note Dhillon et al. [4] use pre-computed lookup tables to collapse the spectral rendering integral. We claim that with their proposed pre-computations for our data, our method can perform in realtime.

5 Conclusion

In this paper, we have presented a physically based method for rendering simple thin-volume photonic structures on natural surfaces. Our method, employs separation of nano-gratings from the micro-geometries through frequency domain filtering of the nanostructures represented as height-fields. Our method consolidates the incoherent mixing of coherently scattered coloration effects from multiple layers with a simple low-pass filtering of Fourier domain magnitudes. We further model incoherent scattering of diffracted colors by the micro-geometries using simple light cones. We demonstrated the working of our method with an actual AFM scan of an example data set. Our method produces photo-realistic softness and gamut for the reference nanostructures. We also show that our method works interactively and can be further optimized for performance. Also, adapting our method for importance sampling will be an interesting challenge for

future. More importantly, we would like to adapt our proposed method to work with far more complex thin-volume surfaces or tubular compositions through efficient application of spatio-temporal light coherence properties and volumetric processing of microscopically acquired data for them.

Acknowledgments. I thank Prof. Milinkovitch, Ms. Arrigo and Dr. Zabuga from Univ. of Geneva (<https://www.lanevol.org>) for providing microscopic scans. I thank them along with Prof. M. Zwicker (UMD) and Dr. A. Ghosh (ICL) for having valuable discussions relating to this work. This work was partly supported by SNSF Early Postdoc. Mobility Fellowship P2BEP2 165343 and partly by NSF Grant No. 2007974.

References

1. Frames-per-second benchmarking tool (FRAPS) v3.5.99 (2021). <http://www.fraps.com/>. Accessed 14 Aug 2021
2. Akenine-Möller, T., Haines, E., Hoffman, N., Pesce, A., Iwanicki, M., Hillaire, S.: Real-Time Rendering, 4th edn. A.K. Peters Ltd, USA (2018)
3. Belcour, L., Barla, P.: A practical extension to microfacet theory for the modeling of varying iridescence. *ACM Trans. Graph. (TOG)* **36**(4), 1–14 (2017)
4. Dhillon, D.S., Teyssier, J., Single, M., Gaponenko, I., Milinkovitch, M.C., Zwicker, M.: Interactive diffraction from biological nanostructures. *Comput. Graph. Forum* **33**(8), 177–188 (2014)
5. He, X.D., Torrance, K.E., Sillion, F.X., Greenberg, D.P.: A comprehensive physical model for light reflection. ACM, New York (1991)
6. Holzschuch, N., Pacanowski, R.: A two-scale microfacet reflectance model combining reflection and diffraction. *ACM Trans. Graph. (TOG)* **36**(4), 1–12 (2017)
7. Levin, A., et al.: Fabricating BRDFS at high spatial resolution using wave optics. *ACM Trans. Graph. (TOG)* **32**(4), 1–14 (2013)
8. Löw, J., Kronander, J., Ynnerman, A., Unger, J.: BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Trans. Graph.* **31**(1), 1–14 (2012)
9. Mashaal, H., Goldstein, A., Feuermann, D., Gordon, J.M.: First direct measurement of the spatial coherence of sunlight. *Opt. Lett.* **37**(17), 3516–3518 (2012)
10. Musbach, A., Meyer, G., Reitich, F., Oh, S.H.: Full wave modelling of light propagation and reflection. In: *CGF*, vol. 32, pp. 24–37 (2013)
11. Oh, S.B., Kashyap, S., Garg, R., Chandran, S., Raskar, R.: Rendering wave effects with augmented light field. In: *CGF*, vol. 29, pp. 507–516 (2010)
12. Stam, J.: Diffraction shaders. In: *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1999*, pp. 101–110. ACM Press/Addison-Wesley Publishing Co., USA (1999)
13. Steinberg, S., Yan, L.Q.: A generic framework for physical light transport. *ACM Trans. Graph. (TOG)* **40**(4), 1–20 (2021)
14. Toisoul, A., Ghosh, A.: Practical acquisition and rendering of diffraction effects in surface reflectance. *ACM Trans. Graph.* **36**(5), 1–16 (2017)
15. Werner, S., Velinov, Z., Jakob, W., Hullin, M.B.: Scratch iridescence: wave-optical rendering of diffractive surface structure. *ACM Trans. Graph. (TOG)* **36**(6), 1–14 (2017)
16. Zabuga, A.V., et al.: Translucent in air and iridescent in water: structural analysis of a salamander egg sac. *Soft Matter* **16**(7), 1714–1721 (2020)