

# Appearance Modeling of Iridescent Feathers with Diverse Nanostructures

YUNCHEN YU, Cornell University, USA

ANDREA WEIDLICH, NVIDIA, Canada

BRUCE WALTER, Cornell University, USA

EUGENE D'EON, NVIDIA, New Zealand

STEVE MARSCHNER, Cornell University, USA

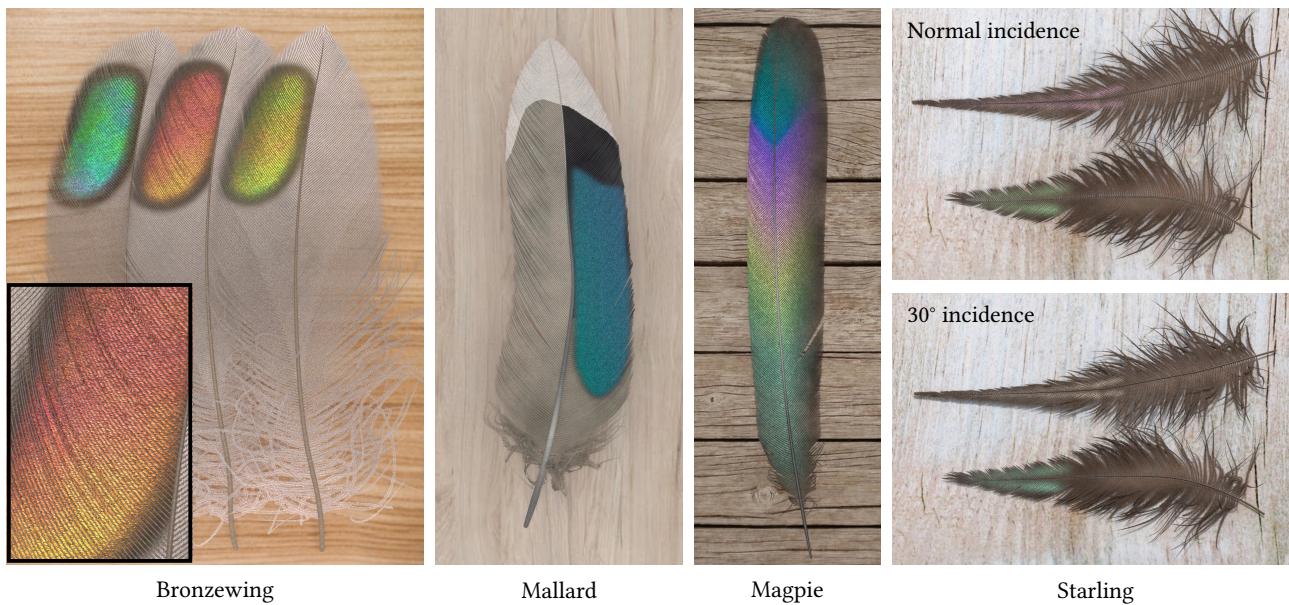


Fig. 1. Rendered feathers of four common bird species, exhibiting iridescent structural color arising from a variety of nanoscale structures within the barbules. Our method starts with randomized procedural models of these structures, uses efficient wave simulations to predict BRDFs and statistical analysis to model spatial and angular irregularity, and results in efficient BRDF models that can be used in standard path tracers.

Many animals exhibit structural colors, which are often iridescent, meaning that the perceived colors change with illumination conditions and viewing perspectives. Biological iridescence is usually caused by multilayers or other periodic structures in animal tissues, which selectively reflect light of certain wavelengths and often result in a shiny appearance—which almost always comes with spatially varying highlights, thanks to randomness and irregularities in the structures. Previous models for biological iridescence tend to each target one specific structure, and most models only compute large-area averages, overlooking spatial variation in iridescent appearance.

Authors' Contact Information: Yunchen Yu, yy735@cornell.edu, Cornell University, USA; Andrea Weidlich, awei@nvidia.com, NVIDIA, Canada; Bruce Walter, bruce.walter@cornell.edu, Cornell University, USA; Eugene d'Eon, ejdeon@gmail.com, NVIDIA, New Zealand; Steve Marschner, srm@cs.cornell.edu, Cornell University, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM 1557-7368/2024/12-ART276  
<https://doi.org/10.1145/3687983>

In this work, we build appearance models for biological iridescence using bird feathers as our case study, investigating different types of feathers with a variety of structural coloration mechanisms. We propose an approximate wave simulation method that takes advantage of quasi-regular structures while efficiently modeling the effects of natural structural irregularities. We further propose a method to distill our simulation results into distributions of BRDFs, generated using noise functions, that preserve relevant statistical properties of the simulated BRDFs. This allows us to model the spatially varying, glittery appearance commonly seen on feathers. Our BRDFs are practical and efficient, and we present renderings of multiple types of iridescent feathers with comparisons to photographic images.

CCS Concepts: • Computing methodologies → Reflectance modeling.

Additional Key Words and Phrases: wave optics, iridescence, material appearance, scattering, feather

## ACM Reference Format:

Yunchen Yu, Andrea Weidlich, Bruce Walter, Eugene d'Eon, and Steve Marschner. 2024. Appearance Modeling of Iridescent Feathers with Diverse Nanostructures. *ACM Trans. Graph.* 43, 6, Article 276 (December 2024), 18 pages. <https://doi.org/10.1145/3687983>

## 1 Introduction

Iridescent structural coloration is a widespread wave optics phenomenon in nature that results from organized nanostructures in biological tissues. Appearance modeling of biological iridescence is an important topic, as it helps us reproduce the striking beauty of nature and also finds applications in biomimicry and meta-material design [Han et al. 2017]. Previous works in graphics have studied structural colors, with methods based on analytical modeling [Sun 2006] or wave simulations [Okada et al. 2013], though most focused on very specific structures, e.g. in one animal species. Moreover, most previous models for iridescent appearance are in the form of large-area averages [Belcour and Barla 2017] and do not sufficiently address the effects of randomness and irregularities in structures—while it is the irregularities that lead to the spatially varying highlights commonly seen on iridescent objects. In this work, we introduce a novel appearance model for biological iridescence arising from multiple types of structures. Iridescence commonly originates from thin films, multilayers, and photonic crystals, and all of these structures have been found in iridescent feathers. Thus, we focus on studying iridescent feathers, using methods that can be generalized toward modeling a wider range of biological iridescence.

Feather iridescence is mainly produced by micro-sized substructures called barbules, which contain quasi-regular nanostructures that selectively reflect light of certain wavelengths. Thus, our appearance modeling relies on computing reflection from barbules, which in turn requires accurate wave simulations. We present a wave optics simulator that computes scattering from barbules with complex internal structures, using different types of approximations for barbules with different nanostructures. By assuming translational symmetry in barbule structures, our simulations only involve 2D calculations and are fast and scalable. Having validated the accuracy of our simulator against our 2D full-wave solver, we use our approximate simulations to efficiently characterize the reflectance distributions of many types of barbules.

More importantly, we want to simulate the spatially varying, glistening appearance of iridescent feathers, caused by small variations in structure from one barbule to the next. The barbules in a feather can be seen as samples from a distribution of similar, but not identical, structures, resulting in a distribution of different BRDFs (Bidirectional Reflectance Distribution Functions). We model not just the mean of this distribution—the average BRDF of the feather—but also the distribution itself. For each type of feather, we perform wave simulations on many randomly generated instances of barbule structures, and then analyze the statistics of the simulation results to derive a mean BRDF and characteristic noise functions, which together define a distribution of BRDFs, preserving important statistical properties of the simulations. Our designed BRDFs can be efficiently sampled and evaluated at render time.

Our contribution is therefore two-fold. We present a fast and accurate wave simulator that computes scattering from small objects with many types of internal structures, and also provide a reference simulator for validation. Furthermore, we introduce a pipeline for distilling wave simulation data into practical appearance models, in the form of BRDF distributions that explicitly model spatially varying appearance. Using our BRDFs, we render different types

of feathers at multiple scales and compare with reference photos. Our barbule geometric models, approximate and reference wave simulators, BRDF generation pipeline, and BRDF implementation can be found at <https://github.com/blaire9989/FeatherLab>.

## 2 Related Work

In this section, we discuss previous work related to feather iridescence, including structural color modeling, wave optics based appearance modeling, and feather rendering and imaging.

### 2.1 Iridescent Reflectance Models

Many models have been proposed to render iridescent surfaces. The earliest [Dias 1991; Gondek et al. 1994; Smits and Meyer 1992] modeled the reflection from a perfect multilayered surface, predicting a mirror-like specular reflection with strong wavelength dependence [Hirayama et al. 2001]; we use a similar computation as a component within our simulation of imperfect multilayer structures. Other methods have modeled diffraction from periodic horizontal structures using grating equations or Fourier optics methods [Egholm and Christensen 2006; Stam 1999; Sun et al. 2000]. For color accuracy, either spectral rendering is used, or careful conversions to RGB reflectance must be made [Belcour and Barla 2017; Sun 2006]. To build more complete and expressive models, thin-film reflection has been combined with microfacet [Belcour and Barla 2017; Kneiphof et al. 2019] or Kirchhoff [Dong et al. 2016; Icart and Arquès 1999, 2000] rough-surface reflection models to incorporate the effects of irregularity on the smooth BRDF. Our method is related to Icart and Arquès [2000] in that we model the effects of irregularity in a physical optics framework, but we model spatial variation rather than only the average BRDF and consider a broader range of structures.

More recent works have applied advanced analytical models to estimate the average BRDFs for specific types of iridescent materials such as pearlescent materials (thin film platelets dispersed in a dielectric) [Guillén et al. 2020] and Bragg mirrors (1D photonic crystals or deeply-layered thin films) [Fourneau et al. 2024].

### 2.2 Wave Optics for Appearance Modeling

An alternative approach to the aforementioned analytic BRDFs is to run more general simulations as a preprocess, using methods such as the finite-difference time-domain (FDTD) method and the boundary element method (BEM). FDTD simulations rely on grid-based spatial discretization, and solve the differential form of Maxwell's equations using time stepping [Oskooi et al. 2010]. Notably, FDTD supports periodic domains and is a good fit for characterizing periodic structures such as photonic crystals, as we do in this work. BEM formulates Maxwell's equations into integral equations on the surface of the scattering object, lowering the dimension of spatial discretization. BEM was previously applied to derive fiber BCSDFs (Bidirectional Curve Scattering Distribution Functions) [Xia et al. 2020] and complicated surface BRDFs [Yu et al. 2023], and in this work, we use BEM to simulate scattering from the internal structures of barbules as a reference for validation. As in [Xia et al. 2020] and [Raymond et al. 2016], we assume translational symmetry in barbule structures and reduce our BEM computations to 2D.

Apart from FDTD and BEM, which rigorously solve Maxwell's equations, approximate wave simulations are also useful for appearance modeling. Recently, [Xia et al. 2023] proposed a fast method for computing fiber scattering, based on pointwise approximations of the reflected field, from which they derived noise-based models for fiber rendering. We also present a fast simulator for computing barbule reflectance based on pointwise approximations, and our feather BRDFs similarly have a noise function component. Still, our work is different from [Xia et al. 2023] in a few aspects. For one, their approximate simulations assume uniform index of refraction (IOR), while our simulations handle a range of internal structures. Moreover, the noise model in [Xia et al. 2023] builds on an existing fiber BCSD, whereas ours is self-contained, with the entire model derived from our simulations. Our parameterized BRDFs can characterize many iridescent objects, modeling their dramatic color changes with respect to illumination and viewing angles.

Moreover, tabulation is frequently used [Cuypers et al. 2012; Guo et al. 2021] to capture the complex behavior of wave optics. We precompute and tabulate certain parameters in our model, and later use them to synthesize our distributions of iridescent feather BRDFs.

### 2.3 Feather Rendering and Iridescence

Previous works have proposed methods for rendering feathers at multiple scales. At the coarsest scale, feathers can be rendered using bidirectional texture functions [Chen et al. 2002; Franco and Walter 2007]. At finer scales, barbs in the feathers, which branch off from the main shafts, are explicitly modeled as curve primitives [Baron and Patterson 2019; Baron 2018; Franco and Walter 2002; Seddon et al. 2008] and are usually rendered using hair shading models [d'Eon et al. 2011; Marschner et al. 2003; Moon et al. 2008; Zinke and Weber 2007]. Only a few works consider barbules, which are microscopic structures that branch off from the barbs—recently, [Baron et al. 2021, 2022] proposed a substructure-based shading technique that simulates the appearance of barbules, and [Huang et al. 2022] implicitly modeled barbules by adjusting the shading frames on the barbs. Among these works, [Huang et al. 2022], which rendered iridescent rock dove feathers using a parameterized BSDF, is the most relevant to our work. Compared to this previous work, our modeling of iridescence is more simulation-driven, but eventually our BRDFs are also parameterized and easy to evaluate.

Moreover, many works in ornithology study the nanostructures in iridescent barbules using scanning electron microscope (SEM) and/or transmission electron microscopy (TEM) imaging, providing valuable data on barbule anatomy [Freyer and Stavenga 2020; Freyer et al. 2021; Giraldo et al. 2018; Han et al. 2017; Nakamura et al. 2008; Stavenga et al. 2017; Xiao et al. 2014]. We rely on these sources for the geometric data that underlie our simulations.

## 3 Overview

As shown in Fig. 2, a feather consists of a central shaft—the rachis—along with a collection of barbs that branch off from the rachis. The barbs contain further branches—the barbules—which are typically microscopic and consist primarily of keratin, melanin, and sometimes air channels [Prum 1999]. These tiny but numerous barbules constitute the majority of the feather's surface area and appearance.

Table 1. List of commonly used symbols.

|  |   |
|--|---|
| $\lambda$  | Wavelength of the light (in air)                |
| $\mathbf{E}, \mathbf{H}$                               | Electric and magnetic fields                    |
| $\mathbf{J}, \mathbf{M}$                               | Electric and magnetic current densities         |
| $\gamma_s, \gamma_p$                                   | Composite reflectivity of a structure           |
| $\omega_i, \omega_o$                                   | Incident and outgoing directions                |
| $\theta_i, \phi_i$                                     | Longitudinal and azimuthal angles of $\omega_i$ |
| $\hat{f}_r(\lambda, \omega_i, \omega_o)$               | BRDF representing average reflectance           |
| $f_r^{(n)}(\lambda, \omega_i, \omega_o)$               | Individual BRDF instances                       |
| $m_{\text{eff}}, \phi_{\text{eff}}, F_{\text{eff}}$    | BRDF parameters: average reflectance            |
| $m_{\text{sta}}, \phi_{\text{sta}}, F_{\text{sta}}, l$ | BRDF parameters: statistics                     |
| $q_r(x), q_i(x)$                                       | Noise functions for BRDF synthesis              |

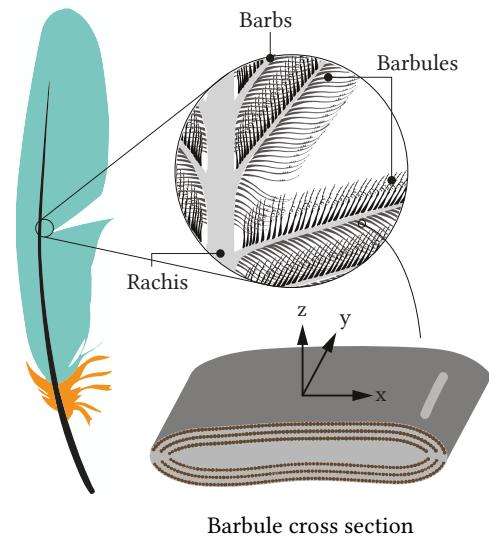


Fig. 2. A typical feather structure with multiple levels of branching. The structures responsible for iridescent colors are the barbules. A barbule is characterized by a non-circular cross section and often exhibits translational symmetry along its longitudinal axis ( $y$ -axis in this context). The feather diagram is adapted from [Thompson 2014].

In barbules from iridescent feathers, different materials are organized into quasi-regular structures that selectively reflect light of certain wavelengths. We study different types of iridescent barbules and build realistic appearance models for iridescent feathers.

### 3.1 Coloration Mechanisms

Many types of birds have iridescent feathers thanks to specially arranged nanostructures in their barbules, and the exact coloration mechanisms vary among bird species. In this work, we study seven bird species, as listed in Table 2. Each species has one unique type of barbule structure, and our featured species cover a representative range of structures, from thin films [Freyer et al. 2021; Nakamura et al. 2008] and multilayers [Giraldo et al. 2018; Xiao et al. 2014] to photonic crystals of varied types [Freyer and Stavenga 2020; Han et al. 2017; Stavenga et al. 2017]. Unlike simpler pigment-based

colors, these iridescent colors are caused by subwavelength-scale structures inside the barbules, and thus the wave nature of light must be considered in our simulations.

### 3.2 2.5D Wave Simulations

The most precise way to compute barbule scattering is to perform 3D wave simulations, which would be very expensive and impractical for rendering. Fortunately, many barbules exhibit translational symmetry along a longitudinal axis, with their cross section profiles approximately invariant, as illustrated in Fig. 2. Therefore, we compute scattering from barbules using “2.5D” wave simulations, as in recent work on fiber scattering [Xia et al. 2020]. That is, we represent each barbule with its cross section profile, and our wave simulations compute electromagnetic fields in 3D, while only using 2D calculations that involve a cross-sectional slice. We introduce our 2.5D simulations in Section 4, and compare BRDFs derived from 2.5D and (very expensive) 3D simulations in our result sections.

We must note that although we study barbules using 2.5D simulations, which were previously applied on fibers, we describe barbule appearance using formalisms for surfaces rather than fibers (i.e. BSDFs rather than BCSDFs). This is because unlike fibers, barbules have distinguishable top and bottom surfaces, and their top surfaces are often nearly flat or mildly curved, allowing us to compute an average, macro-scale surface normal.

### 3.3 BRDF Distributions

Based on simulating barbule scattering, we hope to build realistic appearance models for iridescent feathers. Importantly, each feather consists of many barbules that are structurally similar, while the exact nanostructures vary among barbules, resulting in spatially varying highlights on feathers. Thus, our models need to compute the overall reflected colors from feathers and also incorporate differences in the reflectance distributions of individual barbules.

In our framework, we define a distribution of barbule structures corresponding to each type of feather. Barbules in each distribution are described by similar 2D cross sections that we procedurally generate based on geometric data from the ornithology literature. We perform 2.5D simulations on a number of barbule samples in the distribution, and first construct a smooth BRDF, fitted to an analytical model, that represents the average reflection from the distribution of barbules. We then analyze the reflectance functions of individual barbules to study their single-point statistics (mean and variance) as well as autocorrelations. Based on these statistical properties, we design 1D noise functions and use them to construct a distribution of BRDFs, which statistically resemble the BRDFs computed from simulations. We discuss our BRDF generation pipeline in Section 5.

## 4 Wave Optics Simulations

In this section, we discuss our 2.5D wave simulations that compute reflection from barbules with a diverse range of internal nanostructures. Each wave simulation is provided with a barbule cross section, along with a wavelength and an incident direction. An incident field is propagated to the barbule, and the simulation computes the reflected field from the barbule, which we use to derive the barbule's reflectance distribution.

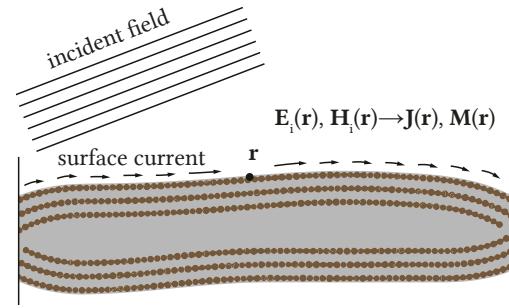


Fig. 3. Illustration of our simulation, where a barbule segment is illuminated by an incident field. The core of the simulation involves computing the surface current densities at each point along the barbule's top surface, based on the local incident field and barbule nanostructure.

One rigorous method to compute the reflected field is to solve Maxwell's equations in the entire domain using the boundary element method (BEM) [Gibson 2021], which formulates the scattering problem on the boundary of the scattering object. Given an incident field and a scattering object, BEM introduces fictitious currents on the object surface, and solves for the current densities from relationships between fields and currents, as dictated by Maxwell's equations. The surface currents are then used to compute the scattered field that propagates outward from the object, which determines the object's scattering distribution. BEM's use of fictitious surface currents lowers the dimensionality of computations, but these simulations can still be very expensive, as the surface currents are computed by solving large systems of equations, assuming all the surface points are interdependent [Xia et al. 2020; Yu et al. 2023].

Here we borrow the notion of fictitious surface currents from BEM, but approximate the current densities on the barbules point by point based on local nanostructures, thereby avoiding large systems of equations. We introduce our simulation setup in 4.1, and in Sections 4.2 and 4.3, we present two methods to approximate surface currents on different types of barbules, which are used as “plug-ins” to our overall simulation pipeline. We discuss how to compute the barbules' reflectance distributions in Section 4.4.

### 4.1 Simulation Setup

We start by introducing our simulation pipeline and important quantities. As shown in Fig. 3, each simulation is provided with a 2D cross section model of a barbule segment, which in general represents a section from a full barbule (since many barbules are larger than needed). As opposed to a closed contour (Fig. 2), our model of a barbule segment has a top and a bottom surface, and the lengths of these surfaces limit the size of the simulation domain along the  $x$ -axis, as shown in Fig. 3. The interior of a barbule segment is modeled as regions that contain keratin, melanin, or air, and keratin and melanin are materials with wavelength-dependent IORs. We adopt the IORs provided in [Stavenga et al. 2015]:

$$\begin{aligned}\eta_k(\lambda) &= 1.532 + 5890/\lambda^2 \\ \eta_m(\lambda) &= 1.648 + 23700/\lambda^2 + j \cdot 0.56e^{-\lambda/270}\end{aligned}\quad (1)$$

Table 2. An overview of the different types of feather barbules studied in this work. In these descriptions, a melanosome refers to a nano-sized structure that contains melanin. Melanosomes are considered rod-shaped, except in hummingbird feathers, where they are described as pancake-shaped.

| Type | Species  | Major coloration scheme   |
|------|--|---|
| 1    | Rock dove ( <i>columba livia</i> )             | Thin film interference from a keratin cortex                            |
| 2    | European starling ( <i>sturnus vulgaris</i> )  | A keratin thin film on top of a basal melanin layer                     |
| 3    | Common bronzewing ( <i>phaps chalcoptera</i> ) | Multiple layers of melanosomes with regular spacing                     |
| 4    | Anna's hummingbird ( <i>calypte anna</i> )     | Layers of pancake-shaped, air-filled melanosomes                        |
| 5    | Common mallard ( <i>anas platyrhynchos</i> )   | 2D photonic crystal: solid melanosomes in a hexagonal grid              |
| 6    | Black-billed magpie ( <i>pica hudsonia</i> )   | 2D photonic crystal: air-filled melanosomes in a hexagonal grid         |
| 7    | Indian peafowl ( <i>pavo cristatus</i> )       | 2D photonic crystal: melanosomes and air channels in a rectangular grid |

where  $\lambda$  is the wavelength (in nm) and  $j$  is the imaginary unit of complex numbers. Note that melanin is an absorbing material, especially at short wavelengths.

Given a barbule segment model, our simulation densely samples a set of points along the 1D top surface. At each point  $\mathbf{r}$ , we evaluate our incident field—a Gaussian-windowed plane wave that propagates along the incident direction—and denote the electric and magnetic fields as  $\mathbf{E}_i(\mathbf{r}), \mathbf{H}_i(\mathbf{r})$ . Based on the local incident field, we estimate the local reflected field  $\mathbf{E}_r(\mathbf{r}), \mathbf{H}_r(\mathbf{r})$  using the methods from Sections 4.2 and 4.3. Knowing the incident and reflected fields at  $\mathbf{r}$ , we can compute the fictitious surface current densities, which are the electric current  $\mathbf{J}(\mathbf{r})$  and the magnetic current  $\mathbf{M}(\mathbf{r})$ :

$$\begin{aligned} \mathbf{J}(\mathbf{r}) &= \mathbf{n}(\mathbf{r}) \times [\mathbf{H}_i(\mathbf{r}) + \mathbf{H}_r(\mathbf{r})] \\ \mathbf{M}(\mathbf{r}) &= -\mathbf{n}(\mathbf{r}) \times [\mathbf{E}_i(\mathbf{r}) + \mathbf{E}_r(\mathbf{r})] \end{aligned} \quad (2)$$

where  $\mathbf{n}(\mathbf{r})$  is the local surface normal at  $\mathbf{r}$ .

With the surface currents, we can evaluate the scattered field from the barbule using the source–field relationships, which provide formulas for the scattered field at any point above the top surface, in terms of the current densities:

$$\begin{aligned} \mathbf{E}_s(\mathbf{p}) &= j\omega\mu(\mathcal{L}\mathbf{J})(\mathbf{p}) - (\mathcal{K}\mathbf{M})(\mathbf{p}) \\ \mathbf{H}_s(\mathbf{p}) &= j\omega\epsilon(\mathcal{L}\mathbf{M})(\mathbf{p}) + (\mathcal{K}\mathbf{J})(\mathbf{p}) \end{aligned} \quad (3)$$

where  $\omega$  is the angular frequency of the fields, and  $\mu, \epsilon$  are the permeability and permittivity of the medium. The operators  $\mathcal{L}$  and  $\mathcal{K}$  compute weighted integrals of  $\mathbf{J}$  and  $\mathbf{M}$  along the barbule's top surface. Eq. 2, Eq. 3 and the expressions for  $\mathcal{L}, \mathcal{K}$  follow from Maxwell's equations and associated boundary conditions, and more detailed formulations can be found in [Xia et al. 2020] and [Yu et al. 2023].

Eq. 2 and Eq. 3 show that our simulation boils down to accurately approximating the reflected field  $\mathbf{E}_r, \mathbf{H}_r$  along the barbule top surface. The next two sections discuss “plug-ins” to our simulation that approximate the reflected field from barbules with different internal structures. Regardless of which “plug-in” we use, our simulation unifies at the last step, where we compute the barbules' reflectance distributions with respect to the outgoing directions.

#### 4.2 Imperfect Multilayer Structure: Barbule Types 1–4

We first introduce a technique for computing reflected fields from barbules with structural colors caused by multilayer arrangements, namely Type 1–Type 4 barbules in Table 2, whose cross sections are shown in Fig. 4(a). We approximate each structure as a spatially

varying multilayer stack, and estimate the reflected field at each point based on the layers below that point.

Notably, we observe that many melanosomes and air bubbles in barbules are closely packed, allowing us to model them as combined layers, as in Fig. 4(b). Moreover, the interior regions of rock dove and starling barbules, with randomly scattered melanosomes, are approximated as layers with IORs computed from estimated percentages of different materials. This way, we turn a barbule that contains many small particles into a layered structure with non-flat layer boundaries—which we call an “imperfect multilayer” structure, as illustrated in Fig. 4(c). These layer structures for modeling each type of barbule are procedurally generated (source code released) based on measurement data from the literature. In our supplemental document, we compare TEM images from related works with our procedural cross section models.

To compute reflection at a point  $\mathbf{r}$  along the barbule top surface, we assume the surface is locally flat and compute a normal vector  $\mathbf{n}(\mathbf{r})$ . We then construct a line  $\mathbf{l}(\mathbf{r})$  that goes through  $\mathbf{r}$  in the direction of  $\mathbf{n}(\mathbf{r})$ . As shown in Fig. 4(c),  $\mathbf{l}(\mathbf{r})$  cuts through the barbule and intersects with the layer boundaries, creating a layer profile at  $\mathbf{r}$  that contains a sequence of layer heights and IORs. At  $\mathbf{r}$ , we treat the incident field as an ideal plane wave with field vectors  $\mathbf{E}_i(\mathbf{r}), \mathbf{H}_i(\mathbf{r})$ , and approximate  $\mathbf{E}_r(\mathbf{r}), \mathbf{H}_r(\mathbf{r})$  by locally reflecting the plane wave off a perfect multilayer stack defined by this layer profile.

Specifically, approximating the field values  $\mathbf{E}_r(\mathbf{r}), \mathbf{H}_r(\mathbf{r})$  requires computing the composite reflection coefficient of the local layer profile, given a wavelength  $\lambda$  and a local incident angle  $\theta(\mathbf{r})$ . With the layer thicknesses  $\{h_k\}$  and the wavelength dependent layer IORs  $\{\eta_k\}$ , the composite reflection coefficients  $\gamma_s(\mathbf{r}), \gamma_p(\mathbf{r})$  for  $s$  and  $p$  polarizations<sup>1</sup> can be computed using the transfer matrix method for multilayer stacks [Born and Wolf 2013; Guillén et al. 2020; Mackay and Lakhtakia 2022].

Given the composite reflection coefficients  $\gamma_s(\mathbf{r}), \gamma_p(\mathbf{r})$ , we compute  $\mathbf{E}_r(\mathbf{r}), \mathbf{H}_r(\mathbf{r})$  by locally decomposing  $\mathbf{E}_i(\mathbf{r})$  into two differently polarized components

$$\mathbf{E}_i(\mathbf{r}) = \mathbf{E}_i^s(\mathbf{r}) + \mathbf{E}_i^p(\mathbf{r}) \quad (4)$$

with respect to the plane of incidence. We then have

$$\mathbf{E}_r(\mathbf{r}) = \gamma_s(\mathbf{r})\mathbf{E}_i^s(\mathbf{r}) + \gamma_p(\mathbf{r})\mathbf{E}_i^p(\mathbf{r}) \quad (5)$$

<sup>1</sup>In the context of plane wave reflection from flat surfaces, the electric field vector is normal to the plane of incidence in the  $s$  polarization, whereas the magnetic field vector is normal to the plane of incidence in the  $p$  polarization.

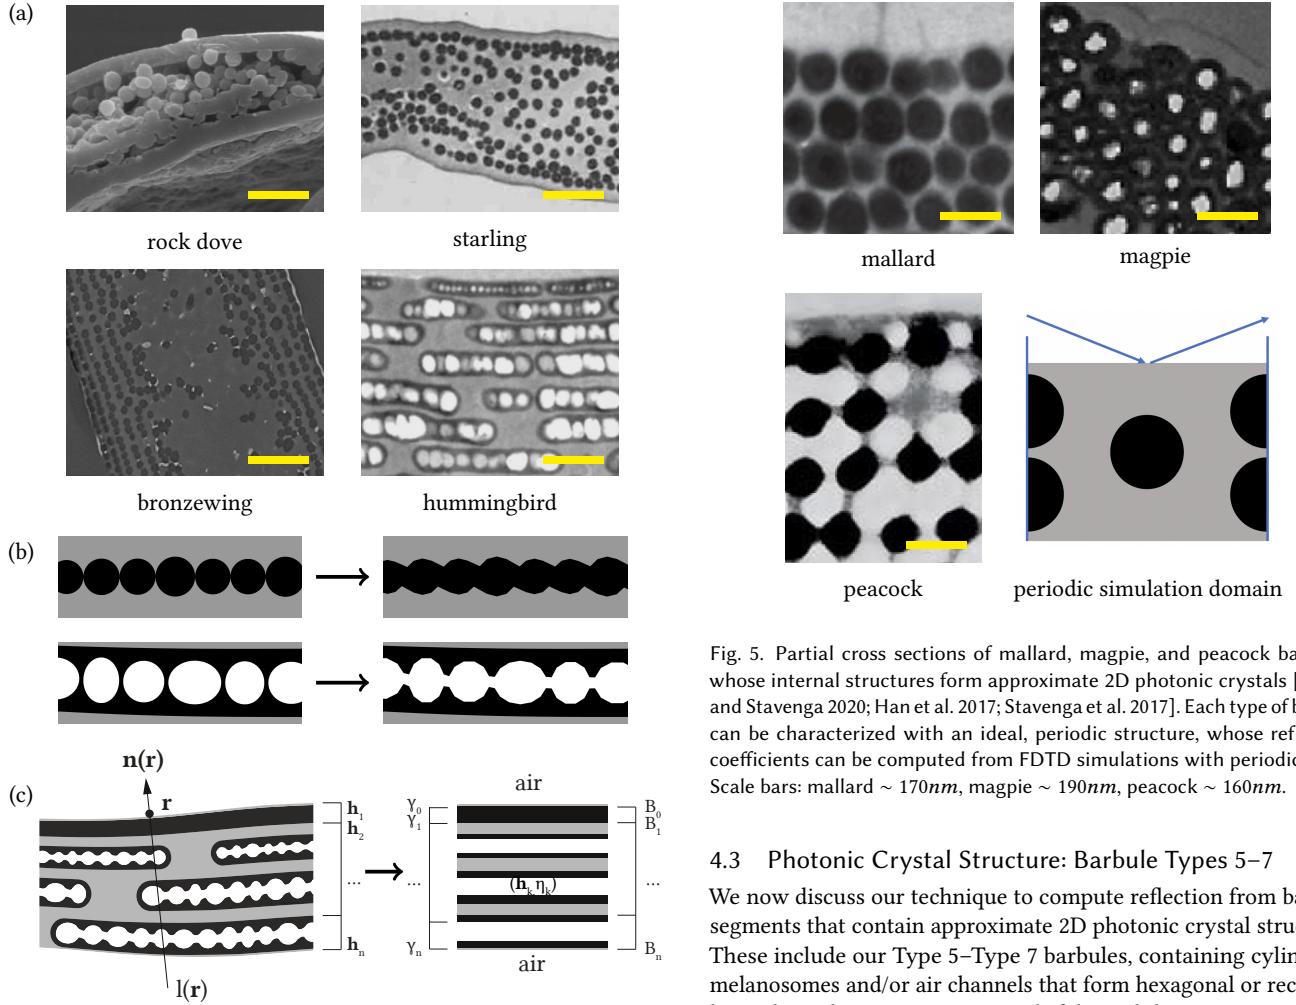


Fig. 4. Barbules that we model as imperfect multilayer structures. (a) Cross sections of rock dove, starling, bronzewing, and hummingbird barbules, from TEM measurements [Freyer et al. 2021; Giraldo et al. 2018; Nakamura et al. 2008; Xiao et al. 2014]. Here and in Fig. 5, white, gray, and black colors refer to air, keratin, and melanin, respectively. Even though rock dove and starling barbules do not resemble multilayer stacks, their structural colors come from the air-keratin-air and air-keratin-melanin layer arrangements at the top of their structures. Scale bars: rock dove  $\sim 1.6\mu\text{m}$ , starling  $\sim 1.5\mu\text{m}$ , bronzewing  $\sim 1.1\mu\text{m}$ , hummingbird  $\sim 0.4\mu\text{m}$ . (b) We model closely packed melanosomes (e.g. from starling, bronzewing) or air bubbles (e.g. from hummingbird) in barbules as connected layers. (c) Our geometric approximations reduce each barbule segment into a spatially varying multilayer structure, and a local layer profile can be obtained at each point on the top surface.

and  $\mathbf{H}_r(\mathbf{r})$  can be derived from  $\mathbf{E}_r(\mathbf{r})$ .  $\mathbf{E}_r, \mathbf{H}_r$  can then be applied in Eq. 2 to compute current densities on barbules, which can in turn be used to compute the scattered field and reflectance distribution of the barbule. As shown in Section 6.1, our multilayer approximation works well on simulating barbules, yielding reflection patterns very similar to those from full-wave simulations.

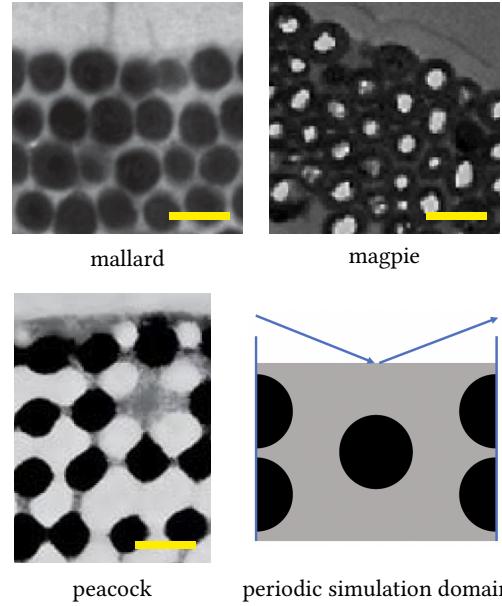


Fig. 5. Partial cross sections of mallard, magpie, and peacock barbules, whose internal structures form approximate 2D photonic crystals [Freyer and Stavenga 2020; Han et al. 2017; Stavenga et al. 2017]. Each type of barbule can be characterized with an ideal, periodic structure, whose reflection coefficients can be computed from FDTD simulations with periodic setup. Scale bars: mallard  $\sim 170\text{nm}$ , magpie  $\sim 190\text{nm}$ , peacock  $\sim 160\text{nm}$ .

#### 4.3 Photonic Crystal Structure: Barbule Types 5–7

We now discuss our technique to compute reflection from barbule segments that contain approximate 2D photonic crystal structures. These include our Type 5–Type 7 barbules, containing cylindrical melanosomes and/or air channels that form hexagonal or rectangular grids, as shown in Fig. 5. Instead of the multilayer approximations from the prior section, we approximate these barbules’ interiors using idealized periodic structures, like the example in Fig. 5 (lower right). Thus, these barbule structures are approximately periodic in one dimension and maintain our assumed translational symmetry along their longitudinal axes. We precompute the reflection coefficients of each 2D photonic crystal at different wavelengths and incident directions using FDTD simulations with periodic domains, and then estimate the reflected fields from a barbule by assuming the periodic structure is locally aligned with its surface at each point.

We use 2D FDTD in the MEEP software [Oskooi et al. 2010] to simulate 2D photonic crystals. In each simulation with a wavelength  $\lambda$  and an incident direction  $\omega_i$ , we define an incident plane wave and use an FDTD simulation to solve for the reflected field in the whole domain. As our photonic crystal structures have very short periods, we can then characterize their reflection using scalar coefficients that depend on  $\lambda$  and  $\omega_i$ , just as in the multilayer case.

In fact, the reflected field from a photonic crystal, along with all other field quantities in the domain, satisfies a Bloch boundary condition along the  $x$  direction:

$$\mathbf{E}_r(\mathbf{r} + \Delta_x \hat{\mathbf{x}}) = e^{jk_x \Delta_x} \mathbf{E}_r(\mathbf{r}) \quad k_x = \hat{\mathbf{x}} \cdot \mathbf{k}_i \quad (6)$$

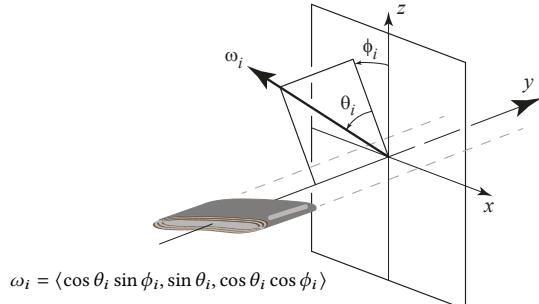


Fig. 6. Illustration of our barbule coordinate system and angles  $\theta_i$  and  $\phi_i$  for an incident direction  $\omega_i$ . Note in this example both angles are negative. The  $y$ -axis is the barbule’s longitudinal axis, and the barbule top surface is approximately parallel to the  $xy$ -plane.

where  $\mathbf{k}_i$  is the incident plane wave’s wavevector, and the same constraint in Eq. 6 applies to the magnetic field. Expressing the reflected field as a sum of plane waves, only those wavevectors that satisfy Eq. 6 can occur; they are the diffraction modes of the crystal.

In our context, the approximate photonic crystal structures in mallard, magpie, and peacock barbules have short periods, with  $\Delta_x < 200\text{nm}$  [Freyer and Stavenga 2020; Han et al. 2017; Stavenga et al. 2017]—less than half of our minimum simulated wavelength. As a result, only a single diffraction mode, corresponding to reflection into the specular direction, is present (see supplemental document for details). Since we know that the reflected field is a plane wave propagating in the specular direction, we can easily compute the reflection coefficient by evaluating the reflected field at a few points.

To fully characterize a photonic crystal structure, we build a reflection coefficient table by running FDTD simulations with many wavelengths and incident directions. For each pair of  $(\lambda, \omega_i)$ , we perform two simulations with  $s$ - and  $p$ -polarized incident plane waves. Reflected fields are computed correspondingly, and reflection coefficients are derived and stored in our table. Many simulations need to be run, while each simulation only takes a few seconds and many simulations can run in parallel. With a precomputed reflection coefficient table of an ideal crystal structure, we can compute the reflected field on any barbule characterized by this type of photonic crystal—using a similar method as in Section 4.2, with  $\gamma_s, \gamma_p$  replaced by our tabulated reflection coefficients.

We have applied substantial approximations in our modeling of crystal-type barbules, since in reality, the photonic crystal arrangements are far from perfect. Nevertheless, our validations in Section 6.1 show that this method is sufficiently accurate when applied to barbules with low surface roughness, which are exactly our targets.

#### 4.4 Far Field Scattering

Given an incident field and a barbule segment, we use the method in Section 4.2 or 4.3 to compute the reflected field from the barbule along its top surface. We can then compute the surface currents using Eq. 2, and further apply Eq. 3 to compute the barbule’s scattered field in the far field (i.e. at points very far from the barbule itself), which determines its reflectance distribution.

Thanks to the assumed translational symmetry, the barbule’s reflectance distribution is a 1D function of the outgoing direction, constrained in a “specular cone” corresponding to the incident direction  $\omega_i$ . To quantify this specular cone, we describe  $\omega_i$  in terms of a longitudinal angle  $\theta_i$  and an azimuthal angle  $\phi_i$ . In our framework, each barbule’s longitudinal axis aligns with the  $y$ -axis, and  $\theta_i$  is the angle formed by  $\omega_i$  and the  $xz$ -plane, while  $\phi_i$  is defined in the  $xz$ -plane with respect to the  $z$ -axis, as shown in Fig. 6. The outgoing direction  $\omega_o$  can be described similarly. Thus, with  $\omega_i \sim (\theta_i, \phi_i)$ , the barbule’s reflectance distribution is defined at outgoing directions satisfying  $\omega_o \sim (-\theta_i, \phi_o)$ , and varies as a function of  $\phi_o$ .

To compute the reflectance distribution, we evaluate the barbule’s scattered field on a semicircle with radius  $\rho$ . In the far field limit where  $\rho$  is much larger than the barbule’s width, the dependence of the scattered field on the radius  $\rho$  is known, and the field corresponding to each outgoing azimuthal angle can be written as

$$\mathbf{E}_s(\lambda, \omega_i; \rho, \phi_o) = e^{jk\rho} \sqrt{1/\rho} \cdot \vec{\mathcal{E}}_s(\lambda, \omega_i; \phi_o) \quad (7)$$

where  $k = 2\pi/\lambda$  and  $\vec{\mathcal{E}}_s(\lambda, \omega_i; \phi_o)$  is a far field analog of the electric field that is independent of  $\rho$  [Gibson 2021; Xia et al. 2020].

Importantly, the barbule’s reflection into  $\omega_o \sim (-\theta_i, \phi_o)$  is associated with the scattered field intensity  $|\vec{\mathcal{E}}_s(\lambda, \omega_i; \phi_o)|^2$ . Moreover, as derived in our supplemental document, the barbule’s total reflected power per unit length is<sup>2</sup>

$$\Phi_r(\lambda, \omega_i) = c_0 \cos \theta_i \int_{-\pi/2}^{\pi/2} |\vec{\mathcal{E}}_s(\lambda, \omega_i; \phi_o)|^2 d\phi_o \quad (8)$$

where  $c_0 = 0.5 \cdot \sqrt{\varepsilon_0/\mu_0}$ , and  $\varepsilon_0$  and  $\mu_0$  are the permittivity and permeability of air. Lastly, the incident power per unit length is computed by integrating the irradiance over the 1D top surface:

$$\Phi_i(\lambda, \omega_i) = \frac{1}{2} \int_S |[\mathbf{E}_i(\mathbf{r}) \times \mathbf{H}_i^*(\mathbf{r})] \cdot \mathbf{n}(\mathbf{r})| d\mathbf{r} \quad (9)$$

where  $\mathbf{n}(\mathbf{r})$  is the local surface normal at the point  $\mathbf{r}$ .

In Section 5, we will use the aforementioned quantities from our 2.5D simulations to derive BRDFs for rendering iridescent feathers. We also note that while the computations are much more expensive, our techniques for approximating the fictitious current densities on barbule surfaces can also be applied to 3D barbule structures that do not possess translational symmetry. These surface currents can be used to compute the 3D barbules’ BRDFs—defined in the entire hemisphere rather than a specular cone (see Section 6.2).

## 5 Iridescent Feather BRDF

Having discussed our 2.5D simulations for computing reflectance functions of individual barbules, we now describe how we distill our simulation results into appearance models for feathers containing many iridescent barbules with structural variations. In Sections 5.1 and 5.2, we will introduce a smooth, microfacet-style BRDF  $\hat{f}_r$ , to describe the average reflection from a given type of barbule. In Section 5.3, we will define our distribution of BRDF instances,  $f_r^{(n)}$ , to model the variations among individual barbules.

<sup>2</sup>The “length” is with respect to the  $y$ -axis.

## 5.1 Average Reflectance Distribution

To characterize the average reflection from a class of barbules, we hope to derive a reflectance function  $\hat{f}_r$  that represents the mean BRDF of the individual barbules in a given distribution. For this purpose, we simulate a number of individual cross sections sampled from a distribution, and based on the 1D reflectance distributions discussed in Section 4.4, we derive an average BRDF  $\hat{f}_r(\lambda, \omega_i, \omega_o)$  that is defined for any pair of directions in the upper hemisphere.

Specifically, for each wavelength  $\lambda$  and incident direction  $\omega_i \sim (\theta_i, \phi_i)$ , we simulate a total of  $M$  barbules, and denote the far field scattered field from the barbule instance  $m$  as  $\vec{\mathcal{E}}_s^{(m)}$ . The incident power  $\Phi_i(\lambda, \omega_i)$  does not vary appreciably across barbule instances and is therefore considered constant for fixed  $\lambda$  and  $\omega_i$ . From these quantities, we define  $U(\lambda, \omega_i; \phi_o)$  in  $\omega_i$ 's specular cone, representing the mean scattering intensity distribution from all the barbules, normalized by the incident power:

$$U(\lambda, \omega_i; \phi_o) = \frac{c_0}{M} \sum_{m=1}^M \frac{|\vec{\mathcal{E}}_s^{(m)}(\lambda, \omega_i; \phi_o)|^2}{\Phi_i(\lambda, \omega_i)} \quad (10)$$

where  $c_0$  was defined in Eq. 8. For outgoing directions  $\omega_o$  in  $\omega_i$ 's specular cone, we want  $\hat{f}_r$  to be approximately proportional to  $U$ :

$$\hat{f}_r(\lambda, \omega_i, \omega_o)(\hat{\mathbf{z}} \cdot \omega_o) \propto U(\lambda, \omega_i; \phi_o) \quad \text{for } \omega_o \sim (-\theta_i, \phi_o) \quad (11)$$

This requirement on  $\hat{f}_r$  is in accordance with the BRDF formulation in [Yan et al. 2018] and [Yu et al. 2023], where  $\hat{\mathbf{z}} \cdot \omega_o$  represents the foreshortening of the barbule as seen by a sensor along  $\omega_o$ .

Moreover, to conserve the average total reflected power, we need

$$\int_{H^2} \hat{f}_r(\lambda, \omega_i, \omega_o)(\hat{\mathbf{z}} \cdot \omega_o) d\omega_o = \cos \theta_i \int_{-\pi/2}^{\pi/2} U(\lambda, \omega_i; \phi_o) d\phi_o \quad (12)$$

where Eq. 12 is derived from Eq. 8 and Eq. 10, as detailed in the supplemental document.

In Section 5.2, we define  $\hat{f}_r$  such that it satisfies Eq. 11 and Eq. 12. We use an analytical, easy-to-sample formulation and make  $\hat{f}_r(\lambda, \omega_i, \omega_o)$  a microfacet-style BRDF, modeling the mean reflectance from single barbule segments in a distribution.

## 5.2 A Microfacet-Style BRDF

We now define the mean reflectance function  $\hat{f}_r$  introduced in Section 5.1. In our context, we construct  $\hat{f}_r(\lambda, \omega_i, \omega_o)$  independently for each pair of incident parameters  $(\lambda, \omega_i)$ , obtaining a microfacet-style formulation through fitting to Trowbridge-Reitz (GGX) models. We use a microfacet-like model because prior works showed that microfacet theories can model some wave optics phenomena despite their geometric optics origins [Dong et al. 2016], and empirically, GGX functions fit well to our simulation data.

For each pair of  $(\lambda, \omega_i)$ , we simulate  $M$  barbule instances and compute their mean scattering distribution  $U(\lambda, \omega_i; \phi_o)$ , using Eq. 10. With many barbule instances,  $U$  becomes a relatively smooth function of  $\phi_o$ . Then for the parameters  $(\lambda, \omega_i)$  where  $\omega_i \sim (\theta_i, \phi_i)$ , we fit  $U(\lambda, \omega_i; \phi_o)$  to a GGX model, such that

$$U(\lambda, \omega_i; \phi_o) \approx s_0 \cdot g(m_{\text{eff}}, \theta_i, \phi_{\text{eff}}, \phi_o) \quad (13)$$

where  $m_{\text{eff}} \equiv m_{\text{eff}}(\lambda, \omega_i)$  and  $\phi_{\text{eff}} \equiv \phi_{\text{eff}}(\lambda, \omega_i)$  are model parameters to be obtained from the fitting and  $s_0$  is a temporary parameter.

The function  $g(m_{\text{eff}}, \theta_i, \phi_{\text{eff}}, \phi_o)$  has the form

$$g(\cdot) = \frac{D'(m_{\text{eff}}, \hat{\mathbf{x}} \cdot \omega_h) G'(m_{\text{eff}}, \hat{\mathbf{z}} \cdot \omega_{\text{eff}}) G'(m_{\text{eff}}, \hat{\mathbf{z}} \cdot \omega_o)}{4(\hat{\mathbf{z}} \cdot \omega_{\text{eff}})}$$

$$\text{where } D'(m_{\text{eff}}, x) = \frac{m_{\text{eff}}^2}{\pi[m_{\text{eff}}^2 + (1 - m_{\text{eff}}^2)x^2]^2} \quad (14)$$

$$G'(m_{\text{eff}}, z) = \frac{2}{1 + \sqrt{1 + m_{\text{eff}}^2 \frac{1-z^2}{z^2}}}$$

where  $\omega_{\text{eff}} \sim (\theta_i, \phi_{\text{eff}})$ ,  $\omega_o \sim (-\theta_i, \phi_o)$ , and  $\omega_h = (\omega_{\text{eff}} + \omega_o) / \|\omega_{\text{eff}} + \omega_o\|$ . Here,  $\omega_{\text{eff}}$  is an “effective” incident direction that depends on the model parameter  $\phi_{\text{eff}}$ , which we find necessary for accurate fitting of  $U$ , especially when the barbules are curved.

Other than missing the Fresnel term and the cosine factor  $(\hat{\mathbf{z}} \cdot \omega_o)$ , the fitted function  $g(m_{\text{eff}}, \theta_i, \phi_{\text{eff}}, \phi_o)$  is quite similar to a classic GGX function [Walter et al. 2007]. The term  $D'(m_{\text{eff}}, \hat{\mathbf{x}} \cdot \omega_h)$  comes from rewriting the classic GGX microfacet distribution function  $D(m, \mathbf{n} \cdot \omega_h)$  as a function of  $\hat{\mathbf{x}} \cdot \omega_h$ . The fitting can be done by finding the parameters  $(s_0, m_{\text{eff}}, \phi_{\text{eff}})$  that minimize

$$\sum_{\phi_o} [U(\lambda, \omega_i; \phi_o) - s_0 \cdot g(m_{\text{eff}}, \theta_i, \phi_{\text{eff}}, \phi_o)]^2 \quad (15)$$

where  $s_0 \geq 0$ ,  $m_{\text{eff}} \in [0, 1]$ ,  $\phi_{\text{eff}} \in [-\pi/2, \pi/2]$ , and the summation is over densely sampled  $\phi_o$  values.  $m_{\text{eff}}(\lambda, \omega_i)$  is the effective roughness parameter, and  $\phi_{\text{eff}}(\lambda, \omega_i)$  defines the effective incident direction. As shown in the supplemental document, for nearly flat barbules, we have  $\phi_{\text{eff}} \approx \phi_i$  for all incident conditions, while for curved barbules (e.g. rock dove, peacock),  $\phi_{\text{eff}}$  varies slightly differently. Finally,  $s_0$  is not used as a model parameter and merely serves to scale  $g$  to the proper magnitude, in order to optimize  $m_{\text{eff}}, \phi_{\text{eff}}$ .

Our GGX fitting determines the shape of  $\hat{f}_r$  in the specular cone of  $\omega_i$ , and we need to extend  $\hat{f}_r$  to the full hemisphere. We first note that  $\hat{f}_r(\lambda, \omega_i, \omega_1)$  and  $\hat{f}_r(\lambda, \omega_i, \omega_2)$  should be highly correlated when  $\hat{\mathbf{y}} \cdot \omega_1 = \hat{\mathbf{y}} \cdot \omega_2$ , as becomes clear after using the 3D version of our simulation to directly compute full hemisphere BRDFs for a few 3D barbules with approximate translational symmetry (see Section 6.2). Thus, with Eq. 11, 13, and 14, we define our full BRDF as

$$\hat{f}_r(\lambda, \omega_i, \omega_o) = F_{\text{eff}} \frac{D(m_{\text{eff}}, \sigma; \omega_h) G(m_{\text{eff}}, \omega_{\text{eff}}, \omega_o)}{4(\hat{\mathbf{z}} \cdot \omega_{\text{eff}})(\hat{\mathbf{z}} \cdot \omega_o)} \quad (16)$$

where  $D(m_{\text{eff}}, \sigma; \omega_h)$ , the normal distribution function (NDF), is

$$D(m_{\text{eff}}, \sigma; \omega_h) = D'(m_{\text{eff}}, \hat{\mathbf{x}} \cdot \omega_h) \cdot e^{-|\hat{\mathbf{y}} \cdot \omega_h|^2 / \sigma^2} \quad (17)$$

Here,  $\sigma$  is a constant parameter that models the spread of the BRDF along the  $y$ -axis. Currently, we manually choose  $\sigma$  for each type of barbule, and acknowledge that Eq. 16 and Eq. 17 only suggest one possible method to define a full hemisphere BRDF.

Moreover, the shadowing-masking term is given by

$$G(m_{\text{eff}}, \omega_{\text{eff}}, \omega_o) = G'(m_{\text{eff}}, \hat{\mathbf{z}} \cdot \omega_{\text{eff}}) G'(m_{\text{eff}}, \hat{\mathbf{z}} \cdot \omega_o) \quad (18)$$

and models the self-occlusion within individual barbule instances.

Lastly,  $F_{\text{eff}}$  in Eq. 16 is chosen to ensure that  $\hat{f}_r$  satisfies Eq. 12. Formally written as  $F_{\text{eff}}(\lambda, \omega_i)$ , this model parameter is an analog to the Fresnel factor in a classic microfacet BRDF, and its dependence on wavelength and angle is the key to modeling iridescence.

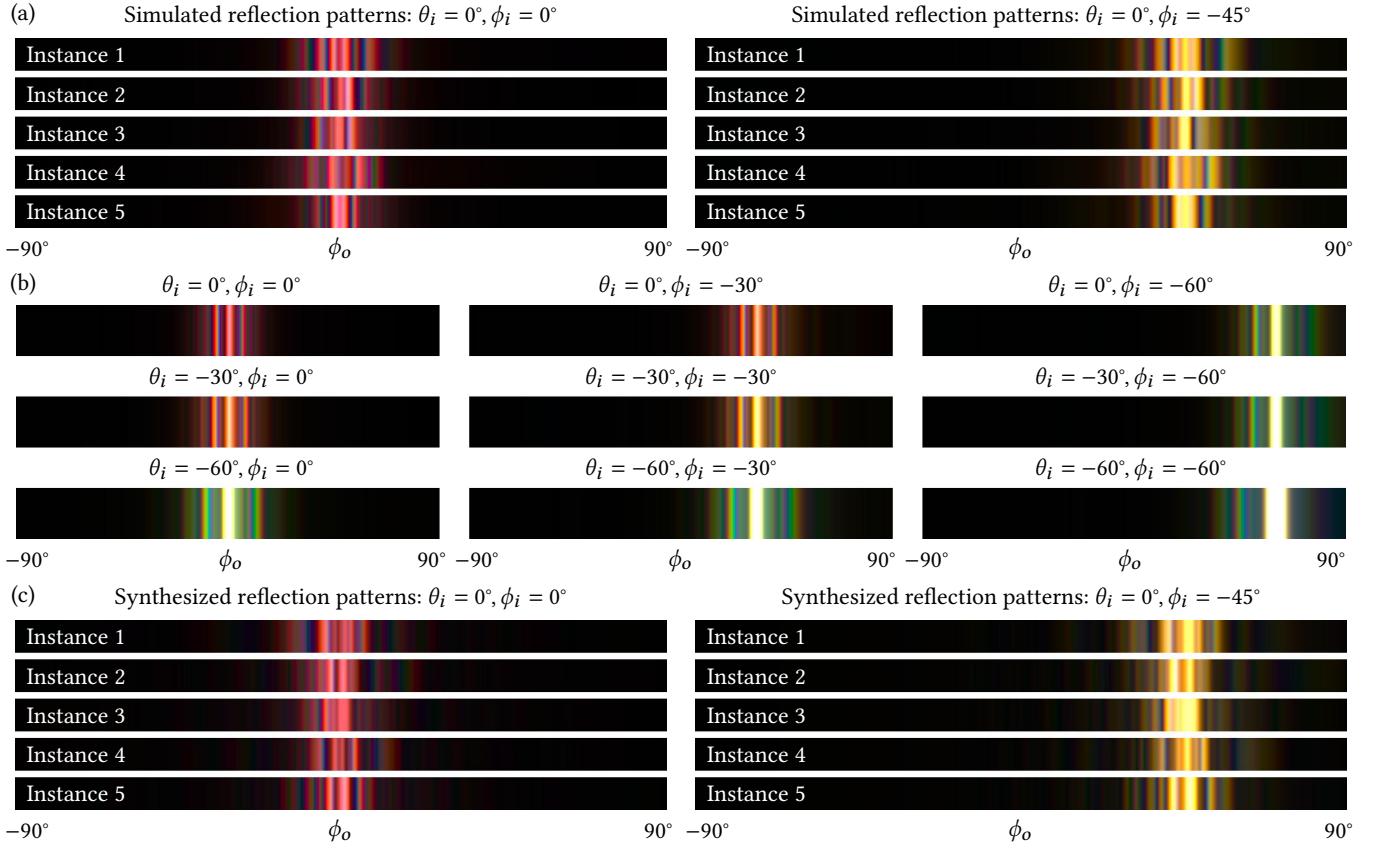


Fig. 7. Demonstration of BRDF instances. (a) Reflection patterns obtained from simulating five procedurally generated red bronzewing barbules, shown for two incident directions in their respective specular cones. (b) Reflection patterns from one bronzewing barbule, shown for nine incident directions. The color patterns are visualized in each  $\omega_i$ 's specular cone, and when  $|\theta_i|$  is large, different  $\phi_o$ 's correspond to scattering directions that are close by in the hemisphere, making the reflection patterns appear "stretched out." Still, one can observe that as  $\omega_i$  changes, the overall reflection pattern is mostly preserved. (c) Some synthesized BRDF instances. As in (a), for each BRDF instance, there is a clear correspondence between the reflection patterns at the two incident directions.

### 5.3 A Distribution of BRDFs

So far we have fitted an analytical BRDF  $\hat{f}_r$  to model the average reflection from barbules coming from a distribution of slightly different structures.  $\hat{f}_r$  could be directly used to render birds from a large distance where no irregularity across a feather is visible, but for closer views it would result in an unrealistically uniform appearance, with all the barbules reflecting in the same way. Thus, we need an efficient process to statistically model the variations in color and brightness from barbule to barbule, which are important to the overall feather appearance. We achieve this by synthesizing a distribution of BRDFs based on  $\hat{f}_r$ , where the BRDF instances  $f_r^{(n)}$  preserve important statistics of the simulated BRDFs.

Fig. 7(a) shows the reflection patterns from a few simulated barbules, which are overall similar but entirely different in details. We find that just like BRDFs of random rough surfaces, these barbule BRDFs statistically resemble partially developed optical speckle [Goodman 2007]. This means that for a distribution of barbules, under each incident condition described by  $(\lambda, \omega_i)$ , in the far field the barbules' scattered field values at each  $\omega_o$  roughly follow a

Gaussian distribution [Steinberg and Yan 2022; Xia et al. 2023]. For one barbule, its BRDF value at each  $(\lambda, \omega_i, \omega_o)$  is produced from the squared magnitude of its scattered field at this point. Thus, the variations in BRDF instances corresponding to different barbules are dictated by the mean and standard deviation of these barbules' scattered field values. We therefore construct  $f_r^{(n)}$  using the single-point statistics in the simulated barbules' scattered fields, while obeying the constraint that these BRDF instances average to  $\hat{f}_r$ .

In addition, it is crucial to quantify the correlations in the reflection from the same barbule at different wavelengths and directions. For instance, each color pattern in Fig. 7(a) demonstrates the correlations in one barbule's scattering into different outgoing directions. Moreover, Fig. 7(b) visualizes one barbule's reflection patterns corresponding to a few  $\omega_i$ 's. The nine subplots indicate that despite the shifting, stretching<sup>3</sup> and color changes, for the same barbule, the details (e.g. fringes) in the simulated reflection patterns are mostly

<sup>3</sup>When  $|\theta_i|$  becomes large and approaches  $\pi/2$ ,  $\omega_i$ 's specular cone "shrinks" and different values of  $\phi_o$  correspond to scattering directions that are close by in the upper hemisphere. This is why the reflection patterns at  $\theta_i = 60^\circ$  look a bit stretched out.

preserved as they vary with the incident direction. These behaviors of the simulated BRDFs should be reproduced in synthesized BRDFs. More exactly, each synthesized  $f_r^{(n)}$  must maintain the correct two-point statistics with respect to different wavelengths and directions, thereby resembling a plausible BRDF that describes a barbule.

As our full pipeline for synthesizing a distribution of BRDFs is complicated, we leave it in the supplemental document and only include the important notations and formulations here. To start with, the aforementioned Gaussian distributions underlying the scattered field values from the barbules allow us to construct each BRDF value in  $f_r^{(n)}$  as the sum of squares of two Gaussian variables:

$$f_r^{(n)}(\lambda, \omega_i, \omega_o) = [r^{(n)}(\lambda, \omega_i, \omega_o)]^2 + [i^{(n)}(\lambda, \omega_i, \omega_o)]^2 \quad (19)$$

where  $r^{(n)}, i^{(n)}$  originate from the real and imaginary parts of the scattered field values. To estimate the mean and standard deviation of these Gaussian variables, we rely on a statistics function  $s(\lambda, \omega_i, \omega_o)$ , which measures the variation in the barbules' scattered field values at each point. Like  $\hat{f}_r$ ,  $s(\lambda, \omega_i, \omega_o)$  is obtained from the simulation data and fitted to a microfacet-like model:

$$s(\lambda, \omega_i, \omega_o) = F_{\text{sta}} \frac{D(m_{\text{sta}}, \sigma; \omega_h)G(m_{\text{sta}}; \omega_{\text{sta}}, \omega_o)}{4(\hat{\mathbf{z}} \cdot \omega_{\text{sta}})(\hat{\mathbf{z}} \cdot \omega_o)} \quad (20)$$

where the model parameters  $m_{\text{sta}}(\lambda, \omega_i)$ ,  $\phi_{\text{sta}}(\lambda, \omega_i)$  (defining  $\omega_{\text{sta}}$ ), and  $F_{\text{sta}}(\lambda, \omega_i)$  are computed in the same way as the analogous model parameters in Section 5.2. At each  $(\lambda, \omega_i, \omega_o)$ ,  $s$  is smaller than or equal to  $\hat{f}_r$  (see supplemental document), and small values of  $s$  indicate high variations in the barbules' scattered field values. More specifically, with  $\hat{f}_r$  and  $s$ , the mean and standard deviation of  $r^{(n)}$  and  $i^{(n)}$  at each  $(\lambda, \omega_i, \omega_o)$  are chosen as

$$\mu_r = \sqrt{s}; \quad \mu_i = 0; \quad \sigma_r = \sigma_i = \sqrt{\hat{f}_r - s} \quad (21)$$

Apart from the single-point statistics, within each BRDF instance, we also need a formulation for modeling the correlations between BRDF values at different wavelengths and directions. Furthermore, we hope to synthesize a large number of BRDF instances efficiently, through a random process. These considerations motivate us to invent a noise process, which generates 1D noise functions  $q$  that we use for synthesizing random BRDF instances. Our complete model for a BRDF instance is given by

$$f_r^{(n)}(\lambda, \omega_i, \omega_o) = [\sigma_r(\lambda, \omega_i, \omega_o) \cdot q_r^{(n)}(x) + \mu_r(\lambda, \omega_i, \omega_o)]^2 + [\sigma_i(\lambda, \omega_i, \omega_o) \cdot q_i^{(n)}(x) + \mu_i(\lambda, \omega_i, \omega_o)]^2 \quad (22)$$

where  $q_r^{(n)}$  and  $q_i^{(n)}$  are two noise functions generated by our 1D noise process. These noise functions satisfy

$$E[q_{r,i}^{(n)}(x)] = 0; \quad \text{Var}[q_{r,i}^{(n)}(x)] = 1 \quad (23)$$

Moreover, the quantity  $x = x(\lambda, \omega_i, \omega_o)$  is given by

$$x(\lambda, \omega_i, \omega_o) = \frac{\hat{\mathbf{x}} \cdot \omega_i + \hat{\mathbf{x}} \cdot \omega_o}{l(\lambda, \omega_i)} \quad (24)$$

where the model parameter  $l(\lambda, \omega_i)$  modulates the autocorrelation in  $f_r^{(n)}(\lambda, \omega_i, \omega_o)$  with respect to the outgoing direction. The parameter  $l$  and the noise function generation process are fundamentally

related to the autocorrelation functions of the simulated barbules' scattered fields, as detailed in the supplemental document.

Fig. 7(c) contains a few synthesized BRDF instances, visualized for two incident directions in their respective specular cones. We note the similarities between the simulated and synthesized BRDFs, as in Fig. 7(a) and Fig. 7(c). We further demonstrate synthesized BRDF instances for multiple types of barbules in Section 6.2.

#### 5.4 Summarizing the BRDF

In previous sections, we discussed our pipeline for distilling our simulation data into iridescent feather BRDFs—from computing an average, full hemisphere BRDF  $\hat{f}_r$ , through estimating the single-point statistics in BRDFs of the simulated barbules, to ultimately synthesizing a distribution of BRDF instances  $f_r^{(n)}$  using a noise generation process. We defined 7 BRDF parameters— $m_{\text{eff}}$ ,  $\phi_{\text{eff}}$ ,  $F_{\text{eff}}$ ,  $m_{\text{sta}}$ ,  $\phi_{\text{sta}}$ ,  $F_{\text{sta}}$ , and  $l$ —associated with the average barbule reflectance as well as the single-point and two-point statistics in the BRDFs of individual barbules. These parameters are computed and stored for each wavelength  $\lambda$  and incident direction  $\omega_i$ . The noise function generation process mentioned in Section 5.3 is fast, with its critical step accelerated using the inverse fast Fourier transform (IFFT). The 1D noise functions  $q_{r,i}^{(n)}$  take up very little storage, so we precompute a large collection of noise functions to use for rendering.

In Section 6, we demonstrate our smooth average BRDF and synthesized BRDF instances for all types of barbules discussed in this work. These BRDFs constructed from 2.5D simulations are further compared to BRDFs directly computed from expensive 3D simulations on 3D barbules. In Section 7, we present our implementation details on BRDF evaluation and importance sampling.

## 6 Results and Validation

In this section, we validate our 2.5D approximate wave simulations against full-wave simulations, and present BRDFs of various types of barbules computed using our methods.

### 6.1 Validations for Barbule Simulations

We demonstrate that our wave simulation method discussed in Section 4 is sufficiently accurate, by validating our simulation results against full-wave BEM simulations. Our BEM simulator generalizes the versions in previous works [Xia et al. 2020; Yu et al. 2023], in that it can simulate objects with arbitrary internal structures. Important mathematical principles behind the full-wave simulations are included in our supplemental document.

We first validate our simulation plug-in from Section 4.2, which applies to our imperfect multilayer barbules. We use a hummingbird barbule as our simulation target, as hummingbird barbules contain the most complicated layer structures (see Fig. 4). As shown in Fig. 8, we feature two incident directions,  $\omega_1 \sim (0^\circ, 0^\circ)$  and  $\omega_2 \sim (0^\circ, -30^\circ)$ , and simulate the barbule in three different ways:

- Full-wave simulations on the unsimplified barbule, where the melanosomes and air bubbles are precisely modeled
- Full-wave simulations on the simplified barbule, where we combine small air bubbles into layers
- Approximate simulations on the simplified barbule, using the technique in Section 4.2

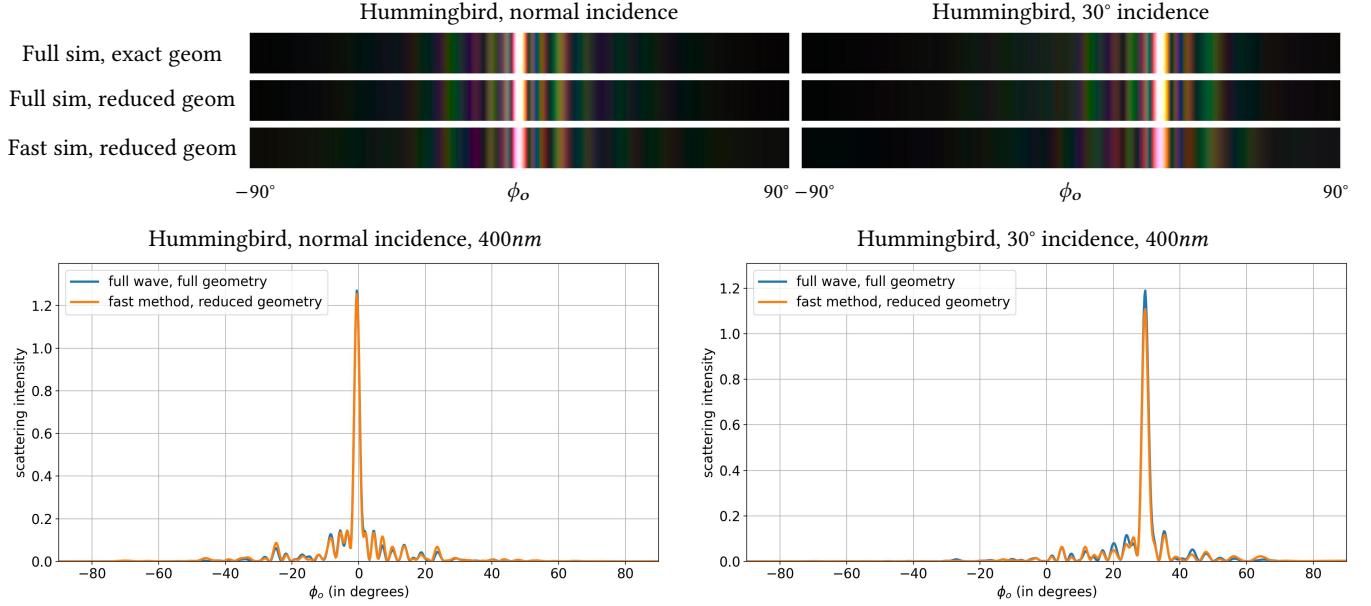


Fig. 8. A three-way comparison of different simulation methods on an “imperfect multilayer” barbule, featuring a pink hummingbird barbule. For each incident direction, we visualize the barbule’s reflection patterns in the specular cone. Featuring  $\lambda = 400\text{nm}$ , we also compare the barbule’s reflectance distributions from the full-wave, exact-geometry and the approximate wave optics, reduced geometry simulations—comparing our practical method to the most accurate setup.

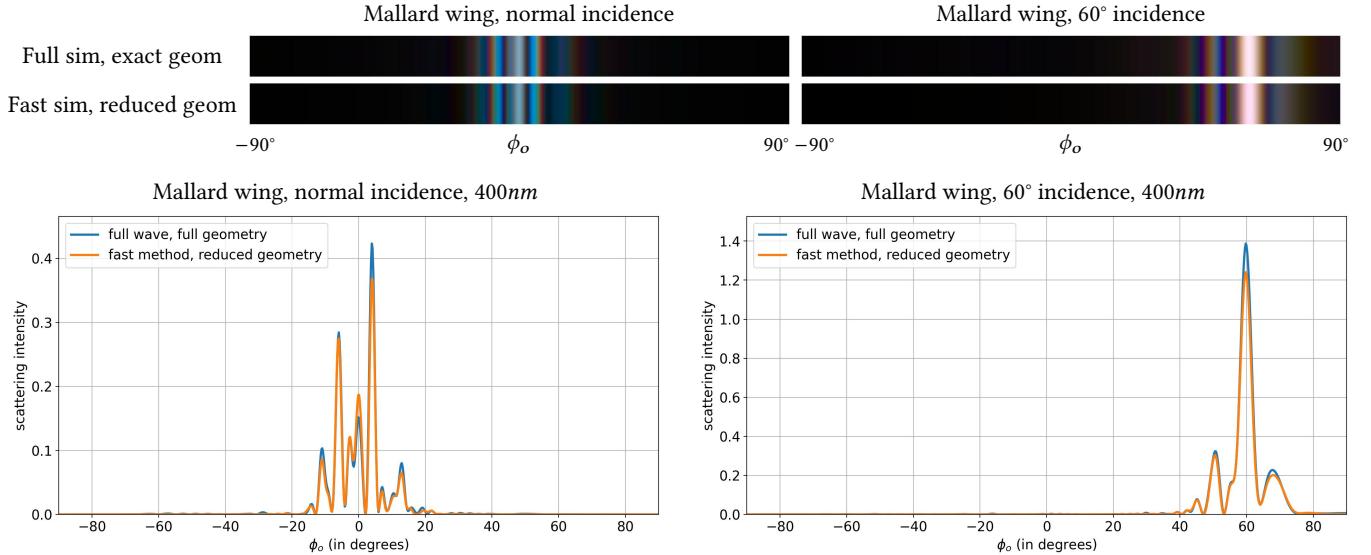


Fig. 9. A two-way comparison of different simulation methods on a blue mallard barbule that contains approximate photonic crystal structures. While our FDTD-based simulations rely on modeling ideal, periodic structures and only explicitly model each barbule’s top surface, our comparisons show that these simulations can in general accurately predict the scattering from photonic crystal type barbules.

The reflection patterns from the three types of simulations, shown in the specular cones with respect to  $\omega_1, \omega_2$ , all appear very similar. Comparing the full-wave simulation results on the full and reduced barbule models shows that our geometric approximations have very little effect on barbule scattering, and comparing the full-wave

and approximate simulation results shows that our fast simulation method is highly accurate, except that it seems to underestimate the reflected intensities along some directions for some wavelengths.

We next validate our simulation plug-in from Section 4.3, which applies to barbules with approximate photonic crystal structures. We

feature a mallard wing feather barbule, characterized by melanosomes arranged in a hexagonal grid. We feature two incident directions and simulate the mallard barbule in two ways:

- Full-wave simulations on the full barbule, where melanosomes are individually modeled. These melanosomes have fluctuating diameters and are arranged in a quasi-regular grid
- Reduced simulations in Section 4.3, based on FDTD precomputations. Only the barbule top surface is explicitly modeled

Despite our substantial approximations, our fast simulations work surprisingly well on the mallard barbule, as shown in Fig. 9. Our supplemental document contains results from more incident conditions, including a case with more significant mismatch between the full-wave and reduced simulation results. Still, our approximate simulations, which are 400–1000 times faster, generally work well on barbules with moderately rough top surfaces. This suggests that ideal photonic crystals are very representative of fully modeled barbules, even when the actual melanosomes are not constant-sized and the underlying grids are not perfectly regular.

## 6.2 Demonstrations of BRDFs

We now present BRDFs generated using our pipeline discussed in Section 5 and compare them to BRDFs computed from simulating 3D barbules with no translational symmetry assumed.

For each type of iridescent feather considered in our work, we procedurally generated 50 barbule cross sections (see our supplemental document and code for geometric models, barbule sizes, etc.) and performed 2.5D simulations on these barbules. We further computed BRDF model parameters for each type of barbule, which allow us to evaluate the average BRDFs and synthesized BRDF instances. The first row of Fig. 10 shows the analytical, smooth BRDF  $\hat{f}_r$  and a synthesized BRDF instance  $f_r^{(n)}$  for a rock dove barbule, featuring three incident directions. The colored BRDFs are shown in the form of “projected hemisphere plots”—each point in the plots corresponds to an outgoing direction in the projected upper hemisphere. We can see that the overall reflected color from the barbule changes with the incident direction, and in each plot, the BRDF instance appears 1D-like, with angular variations along one direction.

To assess the impact of our translational symmetry assumption on BRDFs and renderings, we also developed 3D geometric models for barbules and simulated them using the 3D version of our techniques in Section 4. Specifically, starling, bronzewing, mallard, magpie, and peacock barbules have long, cylindrical interior structures (melanosomes and air channels), so our 3D models of these barbules still possess approximate, but not perfect, translational symmetry. Rock dove and hummingbird barbules contain spherical granules and pancaked-shaped melanosomes, respectively, and we incorporated these features in their 3D models. The second row of Fig. 10 presents BRDFs of 3D rock dove barbules, showing a relatively smooth BRDF (simply averaged across 50 simulated barbules, since we do not have an analytical fitting pipeline for 3D BRDFs) and an individual, simulated instance. The smooth BRDFs are very similar to their counterparts from 2.5D simulations, but the 3D BRDF instance plots contain speckle-like features outside of the specular reflection lobes, revealing some differences between 2.5D and 3D BRDFs. However, the specular reflection lobes themselves, in BRDFs

of our simplified and full 3D barbules, are still qualitatively similar, as barbules indeed have more structural variations in their cross sections than along their longitudinal axes [Nakamura et al. 2008].

BRDFs are presented for other types of barbules in the same format, as shown in the composite hemisphere plots that compare 2.5D and 3D BRDFs, in the rest of Fig. 10 (larger, full images are in the supplemental document). Overall, the unique features in 3D BRDFs are not dominant, especially for barbules that well satisfy the translational symmetry assumption. This suggests that 2.5D simulations are sufficient for deriving BRDFs for realistically rendering iridescent feathers, as we will further show in Section 7.

## 7 Rendering

In this section, we discuss our implementation of the iridescent feather BRDFs in a modern path tracer (Mitsuba3, scalar\_spectral variant), and present feather renderings at multiple scales. Rendering an iridescent feather starts with constructing a distribution of BRDFs for modeling all the barbules in the scene. This distribution of BRDFs is constructed from 2.5D simulation data and represented with seven BRDF parameters and a large collection of pregenerated noise functions. At render time, each barbule primitive or barb segment is assigned a BRDF instance, whose BRDF values depend on relevant model parameters and a unique pair of noise functions in our pool. This way, the reflection from each barbule is modeled by some BRDF instance, and multiple scattering among the barbules is handled by the path tracer, through geometric light transport.

### 7.1 Working with BRDF Parameters

For each type of barbule, our BRDF model parameters— $m_{\text{eff}}$ ,  $\phi_{\text{eff}}$ ,  $F_{\text{eff}}$ ,  $m_{\text{sta}}$ ,  $\phi_{\text{sta}}$ ,  $F_{\text{sta}}$ , and  $l$ —are functions of the simulated wavelength  $\lambda$  and incident direction  $\omega_i$ . We find that with a fixed  $\lambda$ , our model parameters vary smoothly with  $\omega_i \sim (\theta_i, \phi_i)$  and can be considered smooth 2D functions of  $\cos \theta_i$  and  $\cos \phi_i$ .

In our framework, we perform wave simulations at 50 discrete wavelength samples between 400nm and 700nm, and for each wavelength, we consider a collection of  $20 \times 20 = 400$  incident directions, given by  $(\cos \theta_i, \cos \phi_i) = (0.05a, 0.05b)$ , for  $a, b = 1, 2, \dots, 20$ . From simulation results, we use our pipeline in Section 5 (and the supplemental material) to compute the seven BRDF parameters at each  $\lambda$  and  $\omega_i$ . These parameter values are also extrapolated to cover exact grazing angles, where  $\cos \theta_i$  or  $\cos \phi_i$  evaluates to 0.

Upon tabulating our BRDF parameters (< 5MB in total for each type of barbule), at render time, we can estimate the parameter values corresponding to an arbitrary wavelength and incident direction on the fly. Specifically, in spectral rendering each ray is associated with a given wavelength, and we cast this wavelength to the nearest simulated wavelength sample  $\lambda_n$ . Then given a queried direction  $\omega_n \sim (\theta_n, \phi_n)$ , we perform bicubic interpolation on the parameter values at  $\lambda_n$  and estimate  $m_{\text{eff}}$ ,  $\phi_{\text{eff}}$ ,  $F_{\text{eff}}$ ,  $m_{\text{sta}}$ ,  $\phi_{\text{sta}}$ ,  $F_{\text{sta}}$ , and  $l$  at  $(\cos \theta_n, \cos \phi_n)$ . Note that this approach implies that our BRDF parameters are even functions of  $\theta_i$  and  $\phi_i$ , a design choice that reduces the number of required wave simulations.

After estimating the BRDF parameters at queried wavelengths and incident directions, we can importance sample and evaluate our BRDFs, which we discuss in the next section.

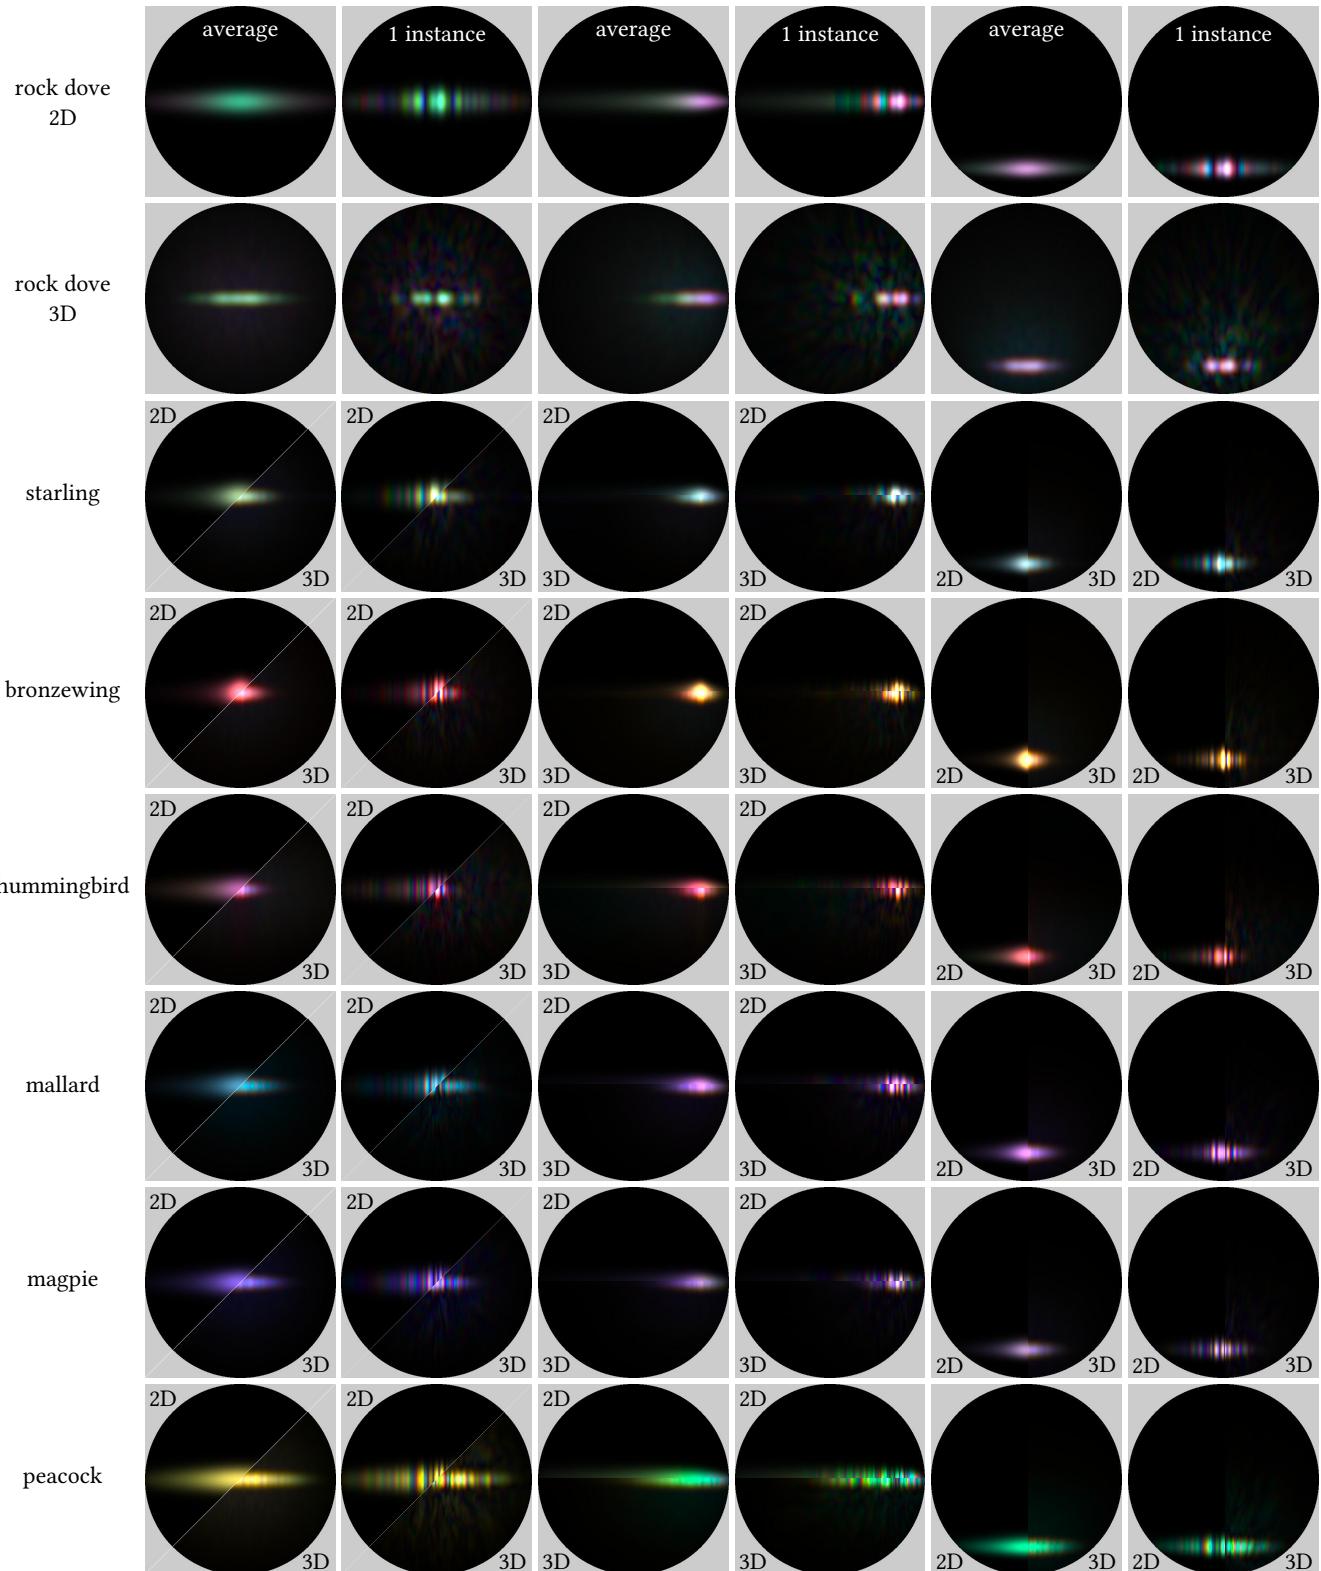


Fig. 10. 2.5D and 3D BRDFs for seven types of barbules, shown as hemisphere plots, featuring incident directions  $(\theta_i, \phi_i) \sim (0^\circ, 0^\circ), (0^\circ, -45^\circ), (45^\circ, 0^\circ)$ .

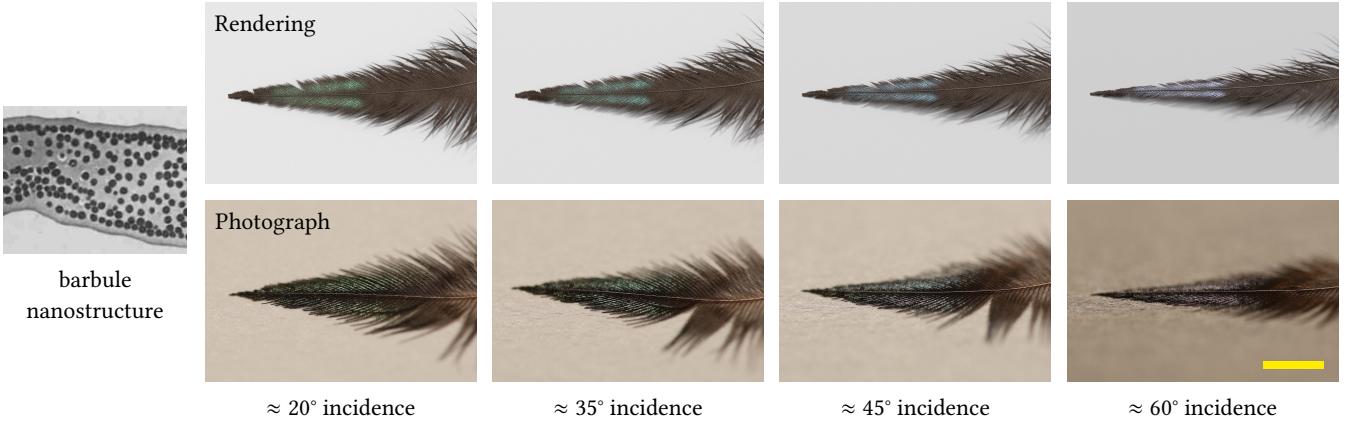


Fig. 11. A European starling feather, rendered (top) and photographed (bottom) at four directions. The feather exhibits weak iridescence and the perceived colors shift from green to violet with increasing illumination and viewing angles. The barbule TEM is adopted from [Freyer et al. 2021]. Scale bar: ~ 4.0mm.

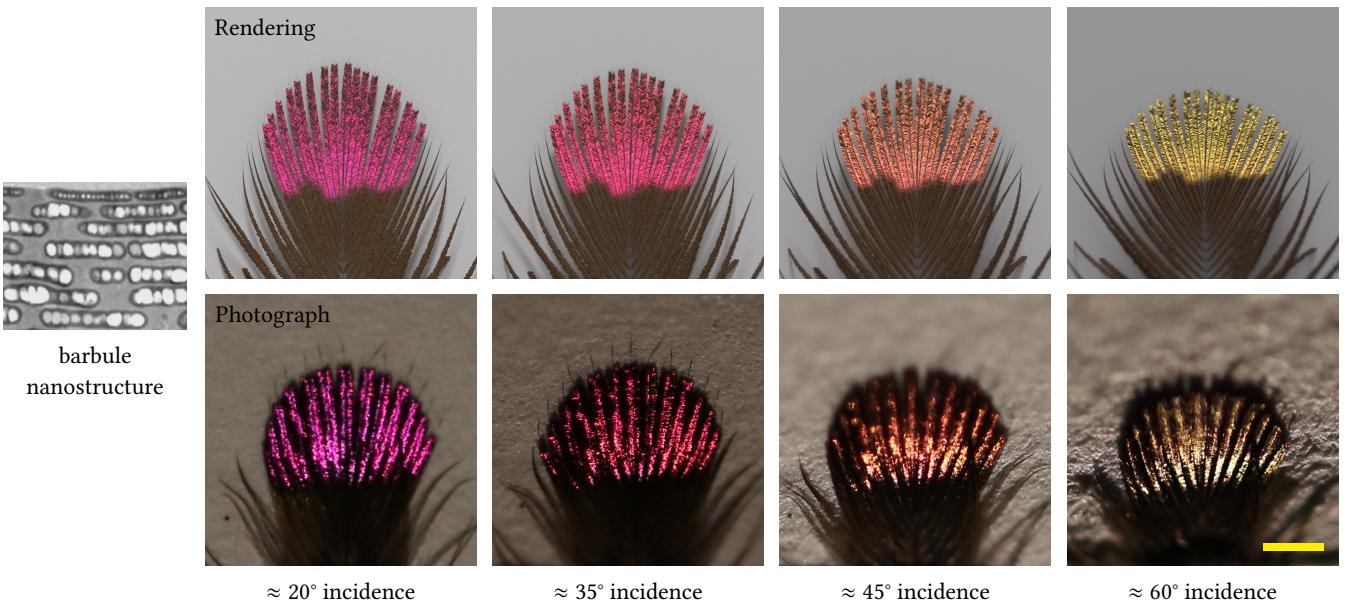


Fig. 12. An Anna's hummingbird feather, rendered (top) and photographed (bottom) at four directions. Hummingbird feather barbules are tilted away from the feather plane by approximately 20° [Giraldo et al. 2018], and the labeled incident angles are with respect to the barbules. The feather exhibits bright structural colors, ranging from pink to yellow when illuminated and viewed from different angles. The TEM is adopted from [Giraldo et al. 2018]. Scale bar: ~ 1.5mm.

## 7.2 BRDF Importance Sampling and Evaluation

Since our synthesized BRDF instances can be drastically different in their details, we perform BRDF importance sampling according to the average BRDF. That is, given a wavelength  $\lambda$  and an incident direction  $\omega_i$ , we importance sample an outgoing direction  $\omega_o$  according to  $\hat{f}_r(\lambda, \omega_i; \omega_o)$ . This requires us to estimate  $m_{\text{eff}}(\lambda, \omega_i)$  and pick a  $\sigma$  value, which together determine the NDF associated with  $\hat{f}_r(\lambda, \omega_i; \omega_o)$ , given in Eq. 17. As for a classic microfacet BRDF, we first sample a half vector  $\omega_h$  according to  $D(m_{\text{eff}}, \sigma; \omega_h)$ .

For convenience, given each pair of  $(m_{\text{eff}}, \sigma)$ , we use a prebuilt sampling table to select  $\omega_h$  based on  $D(m_{\text{eff}}, \sigma; \omega_h)$ . In practice, in our framework  $\sigma$  is manually chosen for each type of barbule and only takes on a few possible values. With  $m_{\text{eff}}$  constrained between 0 and 1, we precompute sampling tables for a finite collection of  $(m_{\text{eff}}, \sigma)$  pairs. Our sampling tables take up less than 1GB combined, and the same tables are used for sampling the BRDFs for all types of barbules. Once we sample a half vector  $\omega_h$  according to the NDF, we can determine  $\omega_o$  from  $\omega_h$  and  $\omega_{\text{eff}}$ , which depends on the estimated parameter  $\phi_{\text{eff}}$ , as introduced in Section 5.2.

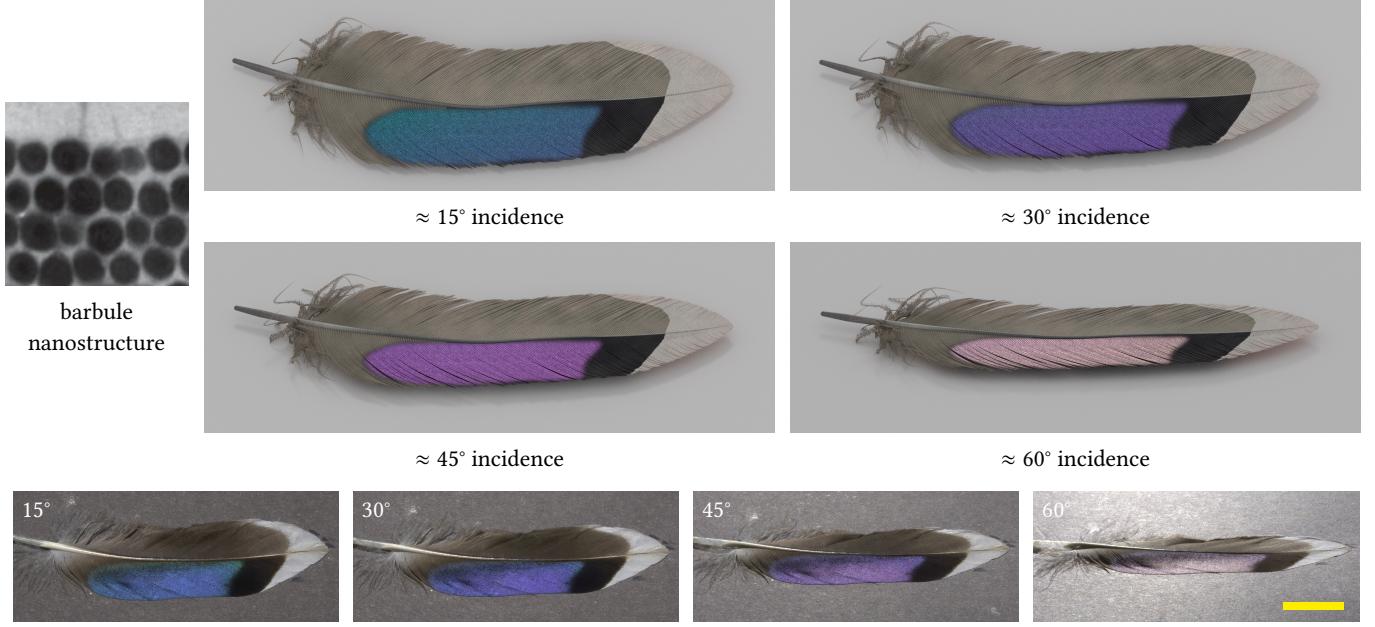


Fig. 13. A mallard wing feather, rendered (top) and photographed (bottom) at four directions. The feather color changes from blue to purple and eventually becomes desaturated as the incident and viewing angles increase. The barbule TEM is adopted from [Stavenga et al. 2017]. Scale bar:  $\sim 25.0\text{mm}$ .

BRDF evaluation relies on the estimated BRDF parameters as well as our precomputed noise functions. When shading a specific barbule, we determine which BRDF instance  $f_r^{(n)}$  it is mapped to and access the associated pair of noise functions  $q_r^{(n)}, q_i^{(n)}$ . We evaluate  $f_r^{(n)}$  at each  $(\lambda, \omega_i, \omega_o)$  according to Eq. 22—after computing the single-point statistics  $\mu_r, \sigma_r, \mu_i, \sigma_i$  using Eq. 16, 20, and 21 and determining the noise function values from  $q_r^{(n)}, q_i^{(n)}$  and Eq. 24.

Lastly, we supplied a diffuse BTDF component to our feather BSDF, and the transmission lobe is chosen small enough such that the BSDF does not violate energy conservation. Accurately computing the transmission through barbules is left for future work.

### 7.3 Single Feathers

We first present our renderings of single feathers. Feather models were made in Houdini 20 and exported as USD files before imported into Mitsuba3. Unlike previous works, barbules are explicitly modeled (as polygon strips) in scenes featuring single feathers. When rendering each barbule, we compute a local shading frame such that the  $t$ -vector aligns with the barbule’s longitudinal axis and the  $n$ -vector aligns with the barbule’s average surface normal. Each barbule is associated with a “barbule ID” that maps to a BRDF instance.

As shown in Fig. 11, 12, and 13, we render single feathers of the European starling, Anna’s hummingbird, and the common mallard, and compare with captured images. Feathers of these birds were obtained from museums and photographed in our laboratory under controlled viewing and lighting conditions. The single feather images were captured using a Canon EOS M6 Mark II camera with a 100mm macro lens, placed on a tripod manually positioned from 15

to 80cm away, depending on the feather size. The light source was a 2cm diameter circular LED mounted on a robotic arm 60cm away.

At small incident angles, the iridescent feathers exhibit their most often perceived colors—weak green, bright pink, and blue. The feather colors gradually blue-shift as the incident and viewing angles increase, and the feathers also appear significantly brighter when illuminated and viewed at larger angles, sometimes exhibiting less saturated colors, as shown in both our renderings and reference images. In all our renderings (also see Fig. 1), we were able to reproduce the spatially varying, glittery appearance on feathers.

### 7.4 Large Feather Assemblies

We also demonstrate larger scenes that contain assemblies of feathers. To save memory, we remove the non-iridescent plumes from the feathers when grooming, and only model the feather down to the barb level. When rendering each barb, we assume that each point on the barb is associated with a barbule, which tilts away from the barb by some angle. Our local shading frame is then computed according to the barb’s normal and the orientation of the underlying barbule, as in [Huang et al. 2022]. Our rock dove scene and mallard head scene consists of 800 and 2200 feathers, respectively, mounting to a total of over 100K bars per scene. Nevertheless, our rendering time was relatively short, as reported in Table 3.

Fig. 14 presents some renderings of the green and purple rock dove neck feathers. Our scene contains a large number of feathers arranged on the surface of a cylinder to approximate the shape of the neck, and our reference images were captured from an intact rock dove specimen. We were not able to bring the bird specimens back to our laboratory, so they were photographed during museum visits

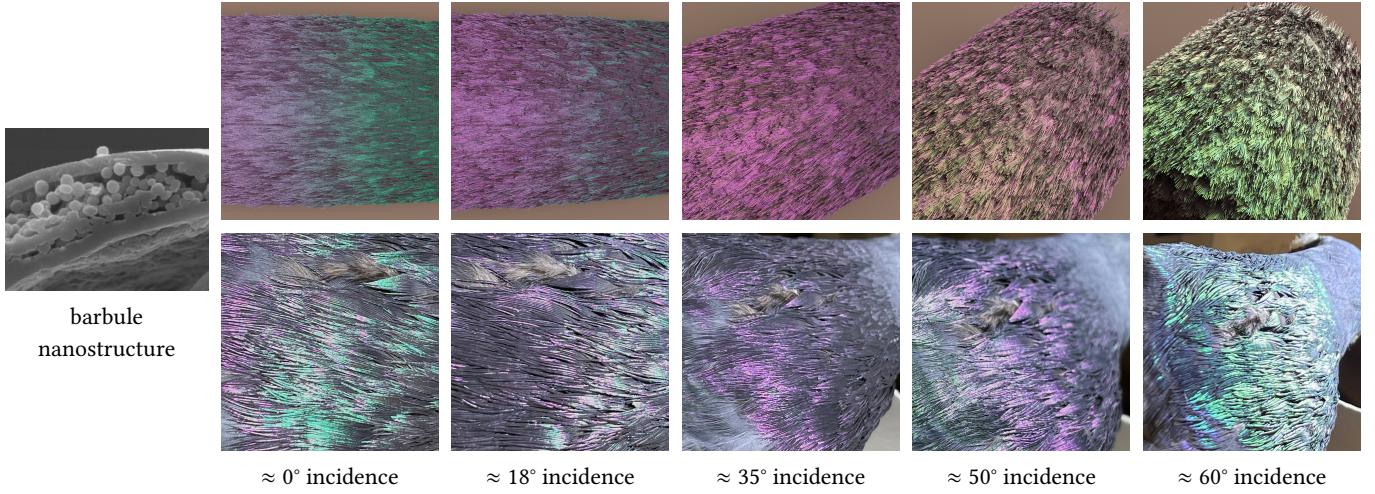


Fig. 14. Rock dove neck feathers, rendered (top) and photographed (bottom) at five directions that give rise to a range of colors. The green and purple colors on the feather come from keratin thin films of different thicknesses. The barbule TEM is adopted from [Nakamura et al. 2008].

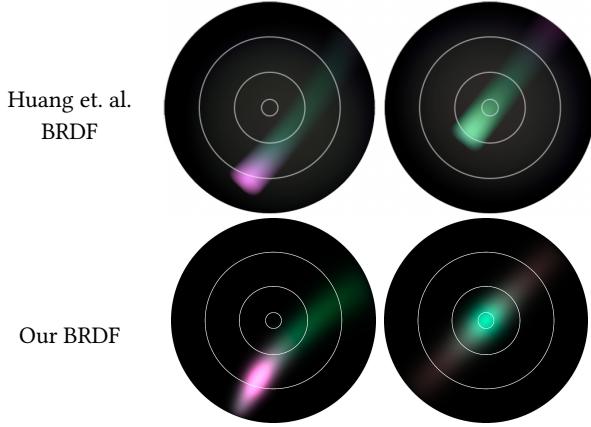


Fig. 15. We compare our rock dove BRDF to those in [Huang et al. 2022]. We use the same visualization scheme as in [Huang et al. 2022], where the white circles indicate viewing angles of 7°, 30°, and 60°. We use an incident angle of 55° for the first configuration and normal incidence for the second.

Table 3. Rendering time corresponding to each presented scene, on a MacBook Pro with Apple M1 chip. For scenes rendered from multiple perspectives, the reported rendering time is averaged across the perspectives.

| Scene                  | Resolution  | Time   |
|------------------------|-------------|--------|
| Fig. 1 (bronzewing)    | 1800 × 2400 | 7.04m  |
| Fig. 1 (mallard)       | 1200 × 2700 | 2.72m  |
| Fig. 1 (magpie)        | 800 × 2400  | 1.60m  |
| Fig. 1 (starling)      | 2400 × 1600 | 2.73m  |
| Fig. 11 (starling)     | 1800 × 1200 | 1.37m  |
| Fig. 12 (hummingbird)  | 1200 × 1200 | 1.47m  |
| Fig. 13 (mallard wing) | 2400 × 800  | 1.56m  |
| Fig. 14 (rock dove)    | 1600 × 1600 | 11.68m |
| Fig. 16 (mallard)      | 2000 × 1500 | 13.83m |

using an iPhone 13 mini. We captured images for a range of viewing and lighting directions but the conditions were less controlled.

The five configurations we present correspond to a range of feather colors that can be observed from the rock dove, and our renderings exhibit color transitions similar to those found in reference photos. Moreover, we compare our smooth mean BRDF  $\hat{f}_r$  for green colored rock dove barbules to the results in [Huang et al. 2022]. We evaluate  $\hat{f}_r$  for two incident directions that roughly match the incident directions used in [Huang et al. 2022], and generate hemisphere plots to visualize the BRDFs. As shown in Fig. 15, our  $\hat{f}_r$  highly resembles the rock dove BRDF from the previous work.

We also rendered some green head feathers from the common mallard, as shown in Fig. 16. Mallard feathers are arranged on the surface of an ellipsoid and the reference images were similarly captured from a bird specimen. At small angles, the rendered feather colors mostly agree with those in captured images, while at our largest angle of around 45°, our rendering and reference image exhibit different shades of blue. This difference is most likely due to the discrepancy between the nanostructures we adopted [Stavenga et al. 2017] and the actual barbule nanostructures in the specimen.

Color differences between rendered and captured images are also present for hummingbird, mallard wing and rock dove feathers. In each case, differences between modeled and real nanostructures could account for some discrepancies, though it is difficult to measure and estimate nanoscale parameters that describe the barbules in each feather sample. The lighting in the renderings and photos may also cause some differences, especially for the rock dove and mallard head scenes, where we had less control over the lighting in the photos. Moreover, the distributions of dark and colored areas on some single feathers are different in rendered and captured images. We believe that the dark areas (e.g. in Fig 12 and 13) are mostly attributed to barbule normals facing varying directions and/or occlusion of some barbules by others, and carefully tuning the feather geometry may lead to better match between renderings and photos.

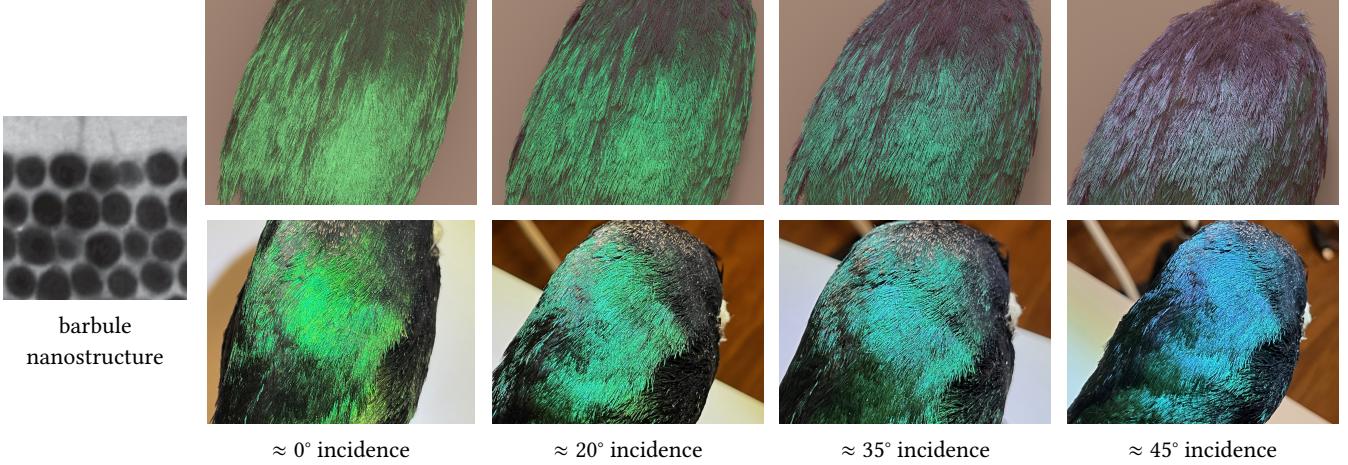


Fig. 16. Mallard head feathers, rendered (top) and photographed (bottom) at four directions that give rise to a color shift from green to blue. The barbule TEM from [Stavenga et al. 2017] is the same one as in Fig. 13 for illustration purposes, while the lattice spacing is larger in head barbules.

## 8 Conclusion and Future Work

We explored an application where the wave nature of light plays a crucial role in appearance, creating dramatic and beautiful iridescence in birds. We showed how simplified—yet accurate—wave simulations can answer questions not only about how reflectance arises from structures but also about how fine-scale variations in reflectance arise from irregularities on structures. For many natural objects, modeling irregularities is as important to realism as modeling the average reflection. Our method also has some limitations, which may nevertheless inspire exciting future works.

**Limitations.** One limitation in our method is that our models are fundamentally based on 2D calculations. Even though our comparisons with 3D simulation results show that 2D calculations seem to be sufficient for modeling barbules, the assumption of translational symmetry makes it difficult to apply our method to more complicated structures. Our simulations are easy to implement in 3D but are expensive, even with GPU accelerations. Also, our method for extending BRDFs from specular cones to the entire hemisphere relies on simple Gaussian functions and manual choice of Gaussian standard deviations, which could be improved in the future.

Moreover, our parameterized BRDFs rely on analytical fitting, and our fitted formulas are limited to single-lobe GGX models. As a result, our method only works well for objects with nearly flat or mildly curved top surfaces, and iridescent objects with more particular structures may require different BRDF models. As an example, the breast feather barbules in a bird of paradise (Lawes' parotia) have unique, boomerang-shaped cross sections [Stavenga et al. 2011] and may need to be described by multi-lobe BRDFs.

Some other limitations lie within our wave simulation method. For one, our simulations rely on pointwise approximation of the reflected field from the simulated object and assumes single reflection from each point on the object surface. This means that multiple reflections, bouncing off different regions on the object surface, are not modeled by our simulations. Such multiple reflections may

rarely occur on natural objects, but could be important to model if we were to apply our method on artificial materials. For another, our simulations are based on leveraging analytic solutions for ideal structures and therefore might become less accurate when applied to objects with less organized structures than barbules, such as structures in fish, cephalopods, or certain plants (e.g. *pollia condensata*). Lastly, our simulations currently do not support accurately computing transmission through the barbules.

**Future Directions.** In the future, we hope to improve and further optimize our wave simulations, and develop more generalized pipelines for converting simulation data into expressive BRDFs. We are also excited to explore ways to develop artist-friendly tools based on our appearance models. For instance, we can precompute a library of BRDF parameters and noise functions for feathers of different overall colors and color changing behaviors. These pregenerated data, combined with an interface for tweaking parameters and controlling spatial variations, can enable artists to render iridescent feathers with desired appearance.

We look forward to the future stages of our work and hope to push toward synthesizing aesthetic images using our appearance models. As always, we will strive to discover and study more applications of wave optics in realistic appearance modeling.

## Acknowledgments

We thank Dr. Mary Margaret Laura Ferraro for giving us access to and helping us handle bird specimens at the Cornell Lab of Ornithology. We also thank Dr. Kevin Epperly from the Burke Museum of Natural History and Dr. Chad Eliason from the Field Museum of Natural History for offering us Anna's hummingbird feathers. This work is supported by the National Science Foundation under grant IIS-2212084 and by a gift funding from NVIDIA Corporation.

## References

Jessica Baron, Daljit Singh Dhillon, and Eric Patterson. 2021. Procedural shading for rendering the appearance of feathers. In *ACM SIGGRAPH 2021 Posters*. 1–2.

Jessica Baron, Daljit Singh Dhillon, N Adam Smith, and Eric Patterson. 2022. Microstructure-based appearance rendering for feathers. *Computers & Graphics* 102 (2022), 452–459.

Jessica Baron and Eric Patterson. 2019. Procedurally generating biologically driven feathers. In *Advances in Computer Graphics: 36th Computer Graphics International Conference, CGI 2019, Calgary, AB, Canada, June 17–20, 2019, Proceedings* 36. Springer, 342–348.

Jessica Renee Baron. 2018. *Procedurally Generating Biologically Driven Bird and Non-Avian Dinosaur Feathers*. Master's thesis. Clemson University.

Laurent Belcour and Pascal Barla. 2017. A practical extension to microfacet theory for the modeling of varying iridescence. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–14.

Max Born and Emil Wolf. 2013. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier.

Yanyun Chen, Yingqing Xu, Baining Guo, and Heung-Yeung Shum. 2002. Modeling and rendering of realistic feathers. *ACM Transactions on Graphics (TOG)* 21, 3 (2002), 630–636.

Tom Cuypers, Tom Haber, Philippe Bekaert, Se Baek Oh, and Ramesh Raskar. 2012. Reflectance model for diffraction. 31, 5, Article 122 (sep 2012), 11 pages. <https://doi.org/10.1145/2231816.2231820>

Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. 2011. An energy-conserving hair reflectance model. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 1181–1187.

Maria Lurdes Dias. 1991. Ray tracing interference color. *IEEE Computer Graphics and Applications* 11, 02 (1991), 54–60.

Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. 2016. Predicting appearance from measured microgeometry of metal surfaces. 35, 1, Article 9 (dec 2016), 13 pages. <https://doi.org/10.1145/2815618>

Janus Egholm and Niels J Christensen. 2006. Rendering compact discs and other diffractive surfaces illuminated by linear light sources. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia*. 329–332.

Gary Fournier, Romain Pacanowski, and Pascal Barla. 2024. Interactive Exploration of Vivid Material Iridescence using Bragg Mirrors. In *Computer Graphics Forum*. Wiley Online Library, e15017. <https://doi.org/10.1111/cgf.15017>

Cristiano G Franco and Marcelo Walter. 2002. Modeling and rendering of individual feathers. In *Proceedings. XV Brazilian Symposium on Computer Graphics and Image Processing*. IEEE, 293–299.

Cristiano G Franco and Marcelo Walter. 2007. Direct texture synthesis of feather pigmentation patterns. In *International Conference on Computer Graphics Theory and Applications*, Vol. 2. SCITEPRESS, 277–283.

Pascal Freyer and Doekele G Stavenga. 2020. Biophotonics of diversely coloured peacock tail feathers. *Faraday Discussions* 223 (2020), 49–62.

Pascal Freyer, Bodo D Wilts, and Doekele G Stavenga. 2021. Cortex thickness is key for the colors of iridescent starling feather barbules with a single, organized melanosome layer. *Frontiers in Ecology and Evolution* 9 (2021), 746254.

Walton C Gibson. 2021. *The method of moments in electromagnetics*. Chapman and Hall/CRC.

Marco A Giraldo, Juan L Parra, and Doekele G Stavenga. 2018. Iridescent colouration of male Anna's hummingbird (*Calypte anna*) caused by multilayered barbules. *Journal of Comparative Physiology A* 204 (2018), 965–975.

Jay S Gondek, Gary W Meyer, and Jonathan G Newman. 1994. Wavelength dependent reflectance functions. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*. 213–220.

Joseph W Goodman. 2007. *Speckle phenomena in optics: theory and applications*. Roberts and Company Publishers.

Íbón Guillén, Julio Marco, Diego Gutierrez, Wenzel Jakob, and Adrian Jarabo. 2020. A general framework for pearlescent materials. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–15. <https://doi.org/10.1145/3414685.3417782>

Yu Guo, Adrian Jarabo, and Shuang Zhao. 2021. Beyond mie theory: systematic computation of bulk scattering parameters based on microphysical wave optics. *ACM Trans. Graph.* 40, 6, Article 285 (dec 2021), 12 pages. <https://doi.org/10.1145/3478513.3480543>

Changhyun Han, Hanbit Kim, Hyunho Jung, Sang-Im Lee, Piotr G Jablonski, and Heonsu Jeon. 2017. Origin and biomimicry of weak iridescence in black-billed magpie feathers. *Optica* 4, 4 (2017), 464–467.

Hideki Hirayama, Kazufumi Kaneda, Hideo Yamashita, and Yoshimi Monden. 2001. An accurate illumination model for objects coated with multilayer films. *Computers & Graphics* 25, 3 (2001), 391–400.

Weizhen Huang, Sebastian Merzbach, Clara Callenberg, Doekele Stavenga, and Matthias Hullin. 2022. Rendering iridescent rock dove neck feathers. In *ACM SIGGRAPH 2022 Conference Proceedings*. 1–8.

Isabelle Icart and Didier Arquès. 1999. An illumination model for a system of isotropic substrate-isotropic thin film with identical rough boundaries. In *Rendering Techniques '99: Proceedings of the Eurographics Workshop in Granada, Spain, June 21–23, 1999* 10. Springer, 261–272.

Isabelle Icart and Didier Arquès. 2000. A physically-based BRDF model for multilayer systems with uncorrelated rough boundaries. In *Rendering Techniques 2000: Proceedings of the Eurographics Workshop in Brno, Czech Republic, June 26–28, 2000* 11. Springer, 353–364.

Tom Kneiphof, Tim Golla, and Reinhard Klein. 2019. Real-time Image-based Lighting of Microfacet BRDFs with Varying Iridescence. *Computer Graphics Forum* 38, 4 (2019), 77–85. <https://doi.org/10.1111/cgf.13772> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13772>

Tom G Mackay and Akhlesh Lakhtakia. 2022. *The transfer-matrix method in electromagnetics and optics*. Springer Nature.

Stephen R Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. 2003. Light scattering from human hair fibers. *ACM Transactions on Graphics (TOG)* 22, 3 (2003), 780–791.

Jonathan T Moon, Bruce Walter, and Steve Marschner. 2008. Efficient multiple scattering in hair using spherical harmonics. In *ACM SIGGRAPH 2008 papers*. 1–7.

Eri Nakamura, Shinya Yoshioka, and Shuichi Kinoshita. 2008. Structural color of rock dove's neck feather. *Journal of the Physical Society of Japan* 77, 12 (2008), 124801.

Naoki Okada, Dong Zhu, Dongsheng Cai, James B Cole, Makoto Kambe, and Shuichi Kinoshita. 2013. Rendering Morpho butterflies based on high accuracy nano-optical simulation. *Journal of Optics* 42 (2013), 25–36.

Ardavan F Oskooi, David Roundy, Mihai Ibanescu, Peter Bermel, John D Joannopoulos, and Steven G Johnson. 2010. MEEP: A flexible free-software package for electromagnetic simulations by the FDTD method. *Computer Physics Communications* 181, 3 (2010), 687–702.

Richard O Prum. 1999. Development and evolutionary origin of feathers. *Journal of Experimental Zoology* 285, 4 (1999), 291–306.

Boris Raymond, Gaël Guennebaud, and Pascal Barla. 2016. Multi-scale rendering of scratched materials using a structured SV-BRDF model. *ACM Trans. Graph.* 35, 4, Article 57 (jul 2016), 11 pages. <https://doi.org/10.1145/2897824.2925945>

Daniel Seddon, Martin Auflinger, and David Mellor. 2008. Rendertime procedural feathers through blended guide meshes. In *ACM SIGGRAPH 2008 talks*. 1–1.

Brian E Smits and Gary W Meyer. 1992. Newton's colors: simulating interference phenomena in realistic image synthesis. In *Photorealism in Computer Graphics*. Springer, 185–194.

Jos Stam. 1999. Diffraction shaders. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. 101–110.

Doekele G Stavenga, Hein L Leertouwer, N Justin Marshall, and Daniel Osorio. 2011. Dramatic colour changes in a bird of paradise caused by uniquely structured breast feather barbules. *Proceedings of the Royal Society B: Biological Sciences* 278, 1715 (2011), 2098–2104.

Doekele G Stavenga, Hein L Leertouwer, Daniel C Osorio, and Bodo D Wilts. 2015. High refractive index of melanin in shiny occipital feathers of a bird of paradise. *Light: Science & Applications* 4, 1 (2015), e243–e243.

Doekele G Stavenga, Casper J Van Der Kooi, and Bodo D Wilts. 2017. Structural coloured feathers of mallards act by simple multilayer photonics. *Journal of The Royal Society Interface* 14, 133 (2017), 20170407.

Shlomi Steinberg and Ling-Qi Yan. 2022. Rendering of subjective speckle formed by rough statistical surfaces. *ACM Transactions on Graphics (TOG)* 41, 1 (2022), 1–23.

Yinlong Sun. 2006. Rendering biological iridescences with RGB-based renderers. *ACM Transactions on Graphics (TOG)* 25, 1 (2006), 100–129.

Yinlong Sun, F David Fracchia, Mark S Drew, and Thomas W Calvert. 2000. Rendering iridescent colors of optical disks. In *Rendering Techniques 2000: Proceedings of the Eurographics Workshop in Brno, Czech Republic, June 26–28, 2000* 11. Springer, 341–352.

Mya Thompson. 2014. Everything You Need To Know About Feathers. <https://academy.allaboutbirds.org/feathers-article/>

Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*. 195–206.

Mengqi Xia, Bruce Walter, Christophe Hery, Olivier Maury, Eric Michielssen, and Steve Marschner. 2023. A Practical Wave Optics Reflection Model for Hair and Fur. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–15.

Mengqi Xia, Bruce Walter, Eric Michielssen, David Bindel, and Steve Marschner. 2020. A wave optics based fiber scattering model. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.

Ming Xiao, Ali Dhinojwala, and Matthew Shawkey. 2014. Nanostructural basis of rainbow-like iridescence in common bronzewing Phaps chalcoptera feathers. *Optics express* 22, 12 (2014), 14625–14636.

Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering specular microgeometry with wave optics. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–10.

Yunchen Yu, Mengqi Xia, Bruce Walter, Eric Michielssen, and Steve Marschner. 2023. A Full-Wave Reference Simulator for Computing Surface Reflectance. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–17.

Arno Zinke and Andreas Weber. 2007. Light scattering from filaments. *IEEE Transactions on Visualization and Computer Graphics* 13, 2 (2007), 342–356.