

Real-Time Rendering of Plant Leaves

Lifeng Wang Wenle Wang* Julie Dorsey[†] Xu Yang[‡] Baining Guo Heung-Yeung Shum
Microsoft Research Asia *Tsinghua University [†]Yale University [‡]Nankai University

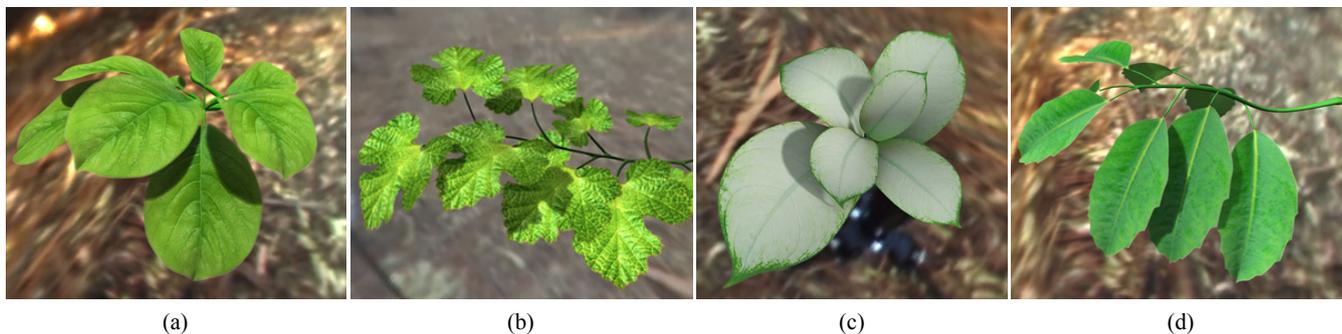


Figure 1: Leaves rendered with our approach. (a)-(d) are balata, pelargonium, omoto and prunus leaves respectively.

Abstract

This paper presents a framework for the real-time rendering of plant leaves with global illumination effects. Realistic rendering of leaves requires a sophisticated appearance model and accurate lighting computation. For leaf appearance we introduce a parametric model that describes leaves in terms of spatially-variant BRDFs and BTDFs. These BRDFs and BTDFs, incorporating analysis of subsurface scattering inside leaf tissues and rough surface scattering on leaf surfaces, can be measured from real leaves. More importantly, this description is compact and can be loaded into graphics hardware for fast run-time shading calculations, which are essential for achieving high frame rates. For lighting computation, we present an algorithm that extends the Precomputed Radiance Transfer (PRT) approach to all-frequency lighting for leaves. In particular, we handle the combined illumination effects due to low-frequency environment light and high-frequency sunlight. This is done by decomposing the local incident radiance of sunlight into direct and indirect components. The direct component, which contains most of the high frequencies, is not pre-computed with spherical harmonics as in PRT; instead it is evaluated on-the-fly using pre-computed light-visibility convolution data. We demonstrate our framework by the rendering of a variety of leaves and assemblies thereof.

Keywords: real-time rendering, appearance modeling, reflectance and shading models, natural phenomena

*This work was done while Wenle Wang was a visiting student and Xu Yang was a visiting researcher at Microsoft Research Asia.

1 Introduction

Realistic rendering of botanical structures, such as trees, is essential to portraying plants and landscape scenery, but achieving realism means confronting many challenges inherent in the composition and appearance of these structures. Leaves are particularly difficult to simulate due to their intricate underlying structure and their complex and subtle interaction with light. The difficulties are exacerbated in scenes comprised of entire tree models illuminated by daylight, where accounting for the shadowing effects, in particular, demands a sophisticated and time-consuming ray tracing approach. However, there are variety of applications, such as environmental assessment and games, for which the interactive rendering of plant models with global illumination effects is desirable.

In this paper, we present a framework for rendering of plant leaves in real-time with global illumination effects. The basis of our framework is a realistic leaf appearance model that is amenable to real-time rendering. This model describes leaf appearance in terms of a few parametric bidirectional reflectance distribution functions (BRDF) and bidirectional transmittance distribution functions (BTDF). These spatially-variant BRDFs and BTDFs are compactly stored in a set of parameter maps, which can be loaded into graphics hardware for fast on-the-fly shading calculations. This is critical for real-time rendering. In our system, a BRDF-BTDF pair is stored as two 720×540 RGBA textures. In general a spatially-variant BRDF is a 6D function which, at this resolution (720×540), could easily consume many gigabytes of memory with a brute-force tabular representation.

For realistic rendering, we derive BRDFs and BTDFs by taking into account the main scattering behaviors of leaves—i.e., the rough surface scattering over leaf surfaces and the subsurface scattering inside leaf tissues. More importantly, we formulate these BRDFs and BTDFs such that they can be measured from real leaves. The subsurface scattering inside leaf tissues determines the BTDF and the BRDF diffuse term [Hanrahan and Krueger 1993]. Our subsurface scattering analysis is based on LEAFMOD, a radiative transfer model for a slab of homogeneous material, which has been experimentally validated with measured data from real leaves [Ganapol et al. 1998]. Using this model we derive parametric forms for the BTDF and the BRDF diffuse term, with parameters including the leaf thickness, as well as the scattering and absorption coefficients of leaf tissues. We also show that these parametric forms can be

fit to reflectance and transmittance data measured by a linear light source (LLS) device [Gardner et al. 2003]. As this model incorporates the key aspects of leaf appearance, it obviates the need for the complex, three-layer model commonly used in plant rendering [Baranoski and Rokne 2002] and supports intuitive editing of leaf appearance. Moreover, the compactness of this representation enables real-time, realistic rendering.

Rough surface scattering is responsible for the glossy reflection over leaf surfaces. Graphics researchers have developed a number of models for the glossy term in the BRDF. However, these models are primarily based on experimental data of inorganic materials; it is not clear which, if any, of these models is appropriate for leaf rendering. We propose the use of the Torrance-Sparrow model [1967] for rendering the glossy reflection of leaves. Our proposal is based on the work of Ma et al. [1990], who conducted extensive experiments to establish that Stogryn’s formula for the normalized scattering cross section per unit area [Stogryn 1967] is well-suited for leaves. We show that the Torrance-Sparrow model is in fact equivalent to Stogryn’s formula and thus suitable for leaf rendering. Like the BRDF diffuse term, the glossy term formulated this way can also be measured using an LLS device.

With the above appearance model, we render leaves using a novel two-pass algorithm built upon the Precomputed Radiance Transfer (PRT) approach [Sloan et al. 2002]. Unlike PRT, which is intended for low-frequency lighting, our algorithm can capture high-frequency lighting effects including soft shadows cast by the sun. We achieve this by decomposing the incident sunlight radiance at each surface point into direct and indirect components and processing them separately in two passes. In the first pass, the indirect component, along with the low-frequency environment light, is efficiently handled by PRT. In the second pass, we use pre-computed light-visibility convolution data to enable quick evaluation of the contribution of the direct sunlight. This avoids the loss of high-frequency details by not using the low-order spherical harmonics basis usually required by PRT. The final rendering result is the sum of the outputs of the two passes.

The remainder of the paper is organized as follows. The following section reviews existing techniques and compares them with ours. Section 3 discusses our leaf model and how to fit this model to reflectance and transmittance data acquired from real leaves. In Section 4, we describe our two-pass rendering algorithm. Section 5 presents some of our results, and Section 6 discusses areas for future work.

2 Related Work

Leaf Models: A variety of techniques exist for creating leaf geometry, which can be modeled as hinged polygons [Bloomenthal 1985], fractal sets [Demko et al. 1985], or L-systems [Prusinkiewicz et al. 1988; Prusinkiewicz et al. 2001]. Our work focuses on the texture and appearance of leaves; the leaf geometry may be obtained with any technique.

Several graphics researchers have studied subsurface scattering in leaves. Hanrahan and Krueger [1993] modeled leaves as layered surfaces and used Monte Carlo ray tracing to evaluate the BRDF and BTDF. Baranoski and Rokne [1997] proposed the algorithmic BDF model (ABM) which accounts for biological factors that affect light propagation and absorption in leaves. Baranoski and Rokne [2001] later introduced the foliar scattering model (FSM), which gains efficiency over ABM by pre-computing reflectance and transmittance values and applying a simplified scattering model. Both the ABM and FSM models are based on Monte Carlo ray tracing. Recently, Franzke and Deussen [2003] reported good rendering speeds (several minutes per frame) with a ray tracer based on a simplified subsurface scattering model. A difficulty with all of these ray-tracing-based models is that they cannot support the fast

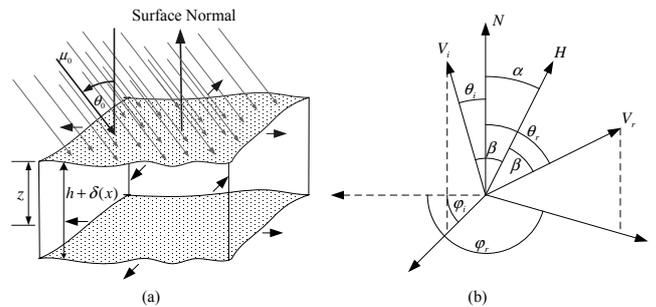


Figure 2: Our leaf model. (a) The plane-parallel leaf geometry (after Ganapol et al.), where μ_0 is the source direction, $h + \delta(\mathbf{x})$ is the leaf thickness, and z is the coordinate measured from the top surface. (b) The angles and vectors for the glossy reflectance of the BRDF.

run-time shading calculations required by real-time rendering. Another problem is that the parameters of these models are usually set by hand rather than measured from real leaves.

Researchers have also developed leaf scattering models for botany and remote sensing applications [Vogelmann 1993; Jacquemoud and Ustin 2001; Baranoski and Rokne 2002]. These models typically make heavy use of biological information of plant tissues. For example, [Govaerts et al. 1996] explicitly modeled the 3D geometry of internal cellular structure of leaf tissues (epidermis, elongated palisade cells, and spongy cells) and used Monte Carlo ray tracing to simulate the propagation of light. Our leaf model does not depend on detailed knowledge about leaf internal structure; instead we rely on measured reflectance and transmittance data for realistic rendering. Essentially, our model is designed for rendering leaves using *measured* data, whereas the biologically-based models are intended for *predicting* measured data through within-leaf light transport simulation. Adapting biologically-based models for real-time rendering is challenging because the light transport simulation is fairly slow.

There exist many analytical BRDF models, which can be isotropic [Torrance and Sparrow 1967; Cook and Torrance 1982; Oren and Nayar 1994] or anisotropic [Kajiya 1985; Ward 1992; Poulin and Fournier 1990; Ashikhmin et al. 2000]. While these models are compact and fast to evaluate, most are designed based on experimental data for inorganic, rather than organic, materials. In particular, none of the models takes into account subsurface scattering, which is important for plant tissues.

Leaf Rendering: The realistic rendering of plant and tree models has a long history in computer graphics (e.g., [de Reffye et al. 1988; Weber and Penn 1995; Max 1996; Deussen et al. 1998; Meyer et al. 2001; Qin et al. 2003; Reche et al. 2004]). Recent techniques, including PRT [Sloan et al. 2002; Sloan et al. 2003] and the all-frequency approach [Ng et al. 2003; Ng et al. 2004], precompute global transport effects in a way that can be exploited by graphics hardware for real-time rendering. Our rendering algorithm resembles these recent techniques in that it pre-computes light transport information to facilitate run-time rendering. It is worthwhile to note that the all-frequency approach, while effective for general environment lighting, is not ideal for our scenario. For our case the all-frequency approach would have to use a very high-dimensional signal and have it sampled very densely over all surfaces, making it impractical for processing a large leaf assembly in real-time.

3 Parametric Leaf Model

We model a leaf as a slab with rough surfaces as illustrated in Fig. 2. The slab interior is assumed to be homogeneous. The slab surface is textured with an albedo map $\gamma(\mathbf{x})$, which accounts for spatially varying reflectance properties. The slab thickness is written

as $h + \delta(\mathbf{x})$, where h is a positive constant for user control of the overall leaf thickness, and $\delta(\mathbf{x})$ is a function that describes local thickness variations in different parts of the leaf slab. $\gamma(\mathbf{x})$ and $\delta(\mathbf{x})$ are computed from reflectance and transmission data of real leaves. In the following we first present our leaf model and then show how each term of the model is derived.

The reflectance and transmittance properties of each surface (top or bottom) of the leaf slab are described by a 6D spatially-variant BRDF $f_r(\mathbf{x}, \theta_i, \phi_i; \theta_r, \phi_r)$ and BTDF $f_t(\mathbf{x}, \theta_i, \phi_i; \theta_r, \phi_r)$

$$\begin{aligned} f_i(\mathbf{x}, \theta_i, \phi_i; \theta_t, \phi_t) &= \frac{1}{\pi} e^{-(\sigma_a + \sigma_s)(h + \delta(\mathbf{x}))} + \frac{B}{2\pi} \frac{\sigma_s}{\sigma_a + \sigma_s} \\ f_r(\mathbf{x}, \theta_i, \phi_i; \theta_r, \phi_r) &= \frac{A}{2\pi} \frac{\sigma_s}{\sigma_a + \sigma_s} \gamma(\mathbf{x}) \\ &+ \frac{\rho_s(\mathbf{x})}{\cos \theta_i \cos \theta_r \cos^4 \alpha} \cdot \frac{\exp(-\frac{\tan^2 \alpha}{m(\mathbf{x})^2})}{4\pi m(\mathbf{x})^2}, \end{aligned} \quad (1)$$

where A and B are constants given below, $\mathbf{x} = (x, y)$ is the position on the leaf surface, (θ_i, ϕ_i) and (θ_r, ϕ_r) describe the incident and reflected directions as Fig. 2 illustrates. There are also three parameters related to subsurface light transport: the absorption coefficient σ_a and scattering coefficient σ_s of the material inside the leaf slab, and the leaf thickness h .

The BRDF $f_r(\mathbf{x}, \theta_i, \phi_i; \theta_r, \phi_r)$ consists of a diffuse term and a glossy term. The diffuse term accounts for the diffuse reflection due to subsurface scattering [Hanrahan and Krueger 1993]. The diffuse term is independent of the incident and reflected directions. Thus we have

$$f_r(\mathbf{x}, \theta_i, \phi_i; \theta_r, \phi_r) = \frac{1}{\pi} \rho_d(\mathbf{x}) + f_s(\mathbf{x}, \theta_i, \phi_i; \theta_r, \phi_r),$$

where $f_s(\mathbf{x}, \theta_i, \phi_i; \theta_r, \phi_r)$ is the glossy term that describes the glossy reflection due to surface roughness. Later we shall derive a formula for the diffuse reflectance $\rho_d(\mathbf{x})$, based on analysis of subsurface light transport. The BTDF has only a diffuse term and can be written as

$$f_t(\mathbf{x}, \theta_i, \phi_i; \theta_t, \phi_t) = \frac{1}{\pi} \rho_t(\mathbf{x}),$$

because light transmitted through materials becomes diffuse.

Subsurface Scattering: Now we derive formulae for the diffuse reflectance $\rho_d(\mathbf{x})$ and transmittance $\rho_t(\mathbf{x})$. Explicit expressions of $\rho_d(\mathbf{x})$ and $\rho_t(\mathbf{x})$ will allow us to answer important questions such as how the leaf BRDF and BTDF vary as the leaf thickness h changes. Another benefit of the explicit expressions is that they allow us to derive a more compact leaf model from the measured data, as we shall see. Our subsurface scattering analysis is based on a within-leaf radiative transfer model called LEAFMOD, which has been experimentally validated with reflectance and transmittance data measured from real leaves [Ganapol et al. 1998].

In general the radiative transfer equation can be written as

$$\Omega \cdot \nabla I(\mathbf{r}, \Omega) + \sigma_t I(\mathbf{r}, \Omega) = \sigma_s \int_{4\pi} d\Omega' p(\Omega', \Omega) I(\mathbf{r}, \Omega'),$$

where $\sigma_t = \sigma_a + \sigma_s$ and $\mathbf{r} = (x, y, z)$. At a given point $\mathbf{x} = (x, y)$ on a homogeneous slab, we can rewrite this equation in 1D form as follows

$$\left[\mu \frac{\partial}{\partial z} + 1 \right] I(z, \mu) = \sigma_s \int_{-1}^1 f(\mu', \mu) I(z, \mu') d\mu',$$

where $I(z, \mu)$ is the radiance at z in direction $\mu = \cos \theta$, $\mu' = \cos \theta'$, and $f(\mu', \mu)$ is the azimuthal average of the general phase function $p(\Omega', \Omega)$. Let $\tau = \sigma_s z$ be the optical path length. In LEAFMOD, the leaf interior is assumed to be filled with isotropic material based on

biological considerations [Ganapol et al. 1998]. Thus $f(\mu', \mu) = \frac{1}{2}$ and we have

$$\left[\mu \frac{\partial}{\partial \tau} + 1 \right] I(\tau, \mu) = \frac{\omega}{2} \int_{-1}^1 I(\tau, \mu') d\mu', \quad (2)$$

where $\omega = \frac{\sigma_s}{\sigma_t}$.

To obtain the diffuse reflectance ρ_d and transmittance ρ_t we follow [Ganapol et al. 1998; Siewert 1978]. Let $\mu_0 = \cos \theta_0$ with $\theta_0 = 0$. The adaxial (front surface) and abaxial (back surface) boundary conditions are

$$I(0, \mu) = \delta(\mu - \mu_0), \quad I(\Delta, -\mu) = 0, \quad (3)$$

where $\mu > 0$, Δ is the optical thickness defined as $\sigma_t h_0$ for a physical thickness h_0 . The boundary condition $I(\Delta, -\mu)$ depends on the Lambertian reflectance r_s of the surface adjacent to the back leaf surface and in our case $r_s = 0$.

The diffuse reflectance $\rho_d = \int \mu I(0, -\mu) d\mu$ and the transmittance $\rho_t = \int \mu I(\Delta, \mu) d\mu$ are obtained by solving $I(0, -\mu)$ and $I(\Delta, \mu)$ from Eq. (2) with boundary conditions Eq. (3). Expanding the exit radiances in a set of shifted Legendre polynomials $\psi_n(\mu)$, we solve Eq. (2) and get

$$\rho_d = \int_0^1 \frac{\omega}{2} \sum_{n=0}^{N-1} a_n \psi_n(\mu) \mu d\mu, \quad \rho_t = \int_0^1 (e^{-\Delta} + \frac{\omega}{2} \sum_{n=0}^{N-1} b_n \psi_n(\mu)) \mu d\mu, \quad (4)$$

where a_n and b_n are constants and N is chosen such that two solutions of consecutive orders are within relative error of 10^{-3} . Finally,

$$\rho_d = \frac{A}{2} \frac{\sigma_s}{\sigma_a + \sigma_s}, \quad \rho_t = e^{-(\sigma_a + \sigma_s)h_0} + \frac{B}{2} \frac{\sigma_s}{\sigma_a + \sigma_s},$$

where A and B are constants determined by a_n and b_n from Eq. (4)

$$A = \sum_{n=0}^{N-1} a_n \int_0^1 \mu \psi_n(\mu) d\mu, \quad B = \sum_{n=0}^{N-1} b_n \int_0^1 \mu \psi_n(\mu) d\mu.$$

The above equations hold at every point $\mathbf{x} = (x, y)$. By adding $\gamma(\mathbf{x})$ and $\delta(\mathbf{x})$ to account for albedo variations and local thickness details respectively, we can extend the above analysis to the whole leaf and obtain

$$\rho_d(\mathbf{x}) = \frac{A}{2} \frac{\sigma_s}{\sigma_a + \sigma_s} \gamma(\mathbf{x}), \quad \rho_t(\mathbf{x}) = e^{-(\sigma_a + \sigma_s)(h + \delta(\mathbf{x}))} + \frac{B}{2} \frac{\sigma_s}{\sigma_a + \sigma_s} \quad (5)$$

In the above equation, h is a constant for user control of the leaf thickness whereas the local thickness variation function $\delta(\mathbf{x})$ is obtained by fitting it to the transmittance data measured from real leaves. Note that $\rho_d(\mathbf{x})$ is not affected by the leaf thickness h .

Note that $r_s = 0$ comes from the fact that our leaf reflectance and transmittance are measured with the LLS device, in which the surface adjacent to the abaxial leaf surface is made nearly non-reflective by covering the light box with a diffuse dark gel [Gardner et al. 2003].

Rough Surface Scattering: The glossy term $f_s(\mathbf{x}, \theta_i, \phi_i; \theta_r, \phi_r)$ describes the glossy reflection due to the rough surface of the slab model. On a rough surface light is scattered in various directions. Because surface roughness of a leaf is large compared to the wavelength of the incident light and undulates at a large scale, we can apply Kirchhoff rough surface scattering theory [Beckmann and Spizzichino 1963]. Stogryn [1967] has derived the following formula for the normalized scattering cross section per unit area for isotropic rough surfaces using Kirchhoff rough surface scattering

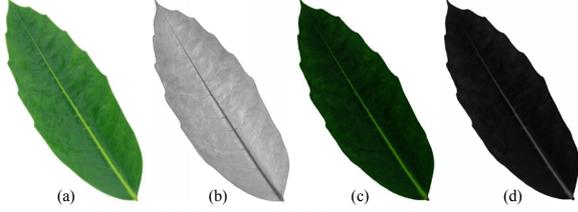


Figure 3: A set of BRDF and BTDF parameter maps for the front surface of a leaf slab. (a) Albedo map $\gamma(\mathbf{x})$. (b) The local thickness variation function $\delta(\mathbf{x})$. (c) The specular intensity map $\rho(\mathbf{x})$. (d) The specular roughness map $m(\mathbf{x})$.

theory:

$$\sigma_k^0 = \frac{l^2(1 + \cos \theta_i \cos \theta_r + \sin \theta_i \sin \theta_r \cos(\phi_r - \phi_i))^2}{(\cos \theta_i + \cos \theta_r)^4 \sigma^2} \cdot \exp\left(-\frac{l^2(\sin^2 \theta_i + \sin^2 \theta_r + 2 \sin \theta_i \sin \theta_r \cos(\phi_r - \phi_i))}{4\sigma^2(\cos \theta_i + \cos \theta_r)^2}\right),$$

where l is the correlation length and σ is the RMS height and the angles θ_i , θ_r , ϕ_i , and ϕ_r are illustrated in Fig. 2. Ma et al. experimentally established that Stogryn's formula are suitable for rough surface scattering on leaf surfaces [1990].

Let $m(\mathbf{x}) = \frac{2\sigma}{\cos \alpha}$ be the root mean square slope of the microfacets at point \mathbf{x} . We have

$$\sigma_k^0 = \frac{1}{\cos^4 \alpha} \cdot \frac{\exp(-\tan^2 \alpha / m(\mathbf{x})^2)}{m(\mathbf{x})^2}, \quad (6)$$

where α is the angle between the surface normal N and the half vector H between the incident and reflected light directions as illustrated in Fig. 2. From this we can express the glossy reflection as follows

$$f_s(\mathbf{x}, \theta_i, \phi_i; \theta_r, \phi_r) = \frac{\rho_s(\mathbf{x})}{\cos \theta_i \cos \theta_r \cos^4 \alpha} \cdot \frac{e^{-\tan^2 \alpha / m(\mathbf{x})^2}}{4\pi m(\mathbf{x})^2}, \quad (7)$$

where $\rho_s(\mathbf{x})$ is called the *specular intensity map*, and $m(\mathbf{x})$ the *specular roughness map*. Eq. (7) is Cook-Torrance model with the geometrical attenuation factor and the Fresnel term merged into $\rho_s(\mathbf{x})$. See Appendix A in the conference DVD for derivation details.

Fitting BRDF and BTDF: We obtain our final BRDFs and BTDFs by fitting these parametric models in Eq. (1) to reflectance and transmittance data measured from real leaves. We acquire this data using an LLS device we built following [Gardner et al. 2003]. Two BRDF-BTDF pairs are acquired, one for each of the top and bottom surfaces of the leaf slab. For each surface, we fit a diffuse lobe and a specular lobe to the reflectance data acquired by the LLS device and thus obtain the diffuse reflectance $\rho_d(\mathbf{x})$, specular intensity map $\rho_s(\mathbf{x})$ and specular roughness map $m(\mathbf{x})$. We also measure the transmittance $\rho_t(\mathbf{x})$ for each surface.

From an estimated leaf thickness h and measured $\rho_d(\mathbf{x})$ and $\rho_t(\mathbf{x})$, we can compute σ_a , σ_s , $\gamma(\mathbf{x})$ and $\delta(\mathbf{x})$ as follows. We first solve for σ_a and σ_s values at every point \mathbf{x} using the following equations:

$$\rho_d(\mathbf{x}) = \frac{A}{2} \frac{\sigma_s}{\sigma_a + \sigma_s}, \quad \rho_t(\mathbf{x}) = e^{-(\sigma_a + \sigma_s)h} + \frac{B}{2} \frac{\sigma_s}{\sigma_a + \sigma_s}.$$

Then we average these values over the leaf surface to get two scalar constants σ_a and σ_s . Once σ_a and σ_s are known, it is straightforward to get $\gamma(\mathbf{x})$ and $\delta(\mathbf{x})$ using Eq. (5). In practice we have to iterate through the above equations multiple times because σ_a , σ_s , $\rho_d(\mathbf{x})$, and $\rho_t(\mathbf{x})$ all have RGB channels.

Fig. 3 exhibits a set of parameter maps for a surface of the leaf slab. Recovering the parameters $\gamma(\mathbf{x})$, $\delta(\mathbf{x})$, σ_a , and σ_s through

the fitting process has two advantages. First, it removes the redundancy in the measured diffuse transmittance $\rho_t(\mathbf{x})$ and makes the leaf model more compact. The measured $\rho_t(\mathbf{x})$ has RGB channels. In contrast $\delta(\mathbf{x})$ is only a greyscale map that can be stored in the alpha-channel of one of the texture maps needed for the BRDF parameters. As a result, no separate texture map is needed to store the BTDF. The other advantage is that we can now perform meaningful editing of leaf appearance by perturbing parameters such as σ_a , σ_s , and the leaf thickness h from their estimated values. See Fig. 7 for editing examples.

4 Lighting Computation

In this section we present a two-pass algorithm for real-time rendering of plant leaves with global illumination. Our approach builds upon the PRT framework [Sloan et al. 2002]. Unlike PRT, which is intended for low-frequency lighting, our algorithm is designed for illumination that includes both a low-frequency environment map and the sun, an all-frequency source.

Our goal is efficient global illumination, including the important high-frequency lighting and detailed, soft shadowing effects due to the sun. To achieve this goal, our algorithm decomposes the sunlight illumination at each surface point into direct and indirect components and processes them separately in two rendering passes. In the first pass, the indirect component, along with the low-frequency environment light, is efficiently handled by PRT. In the second pass, we quickly evaluate the contribution of the direct component using pre-computed light-visibility convolution data at all vertices in the scene. The second pass does not use a low-order spherical harmonics basis and thus avoids the loss of high-frequency details. The final rendering result is the sum of the outputs of the two passes.

Sunlight Decomposition: According to the formulation of [Sloan et al. 2002; Kautz et al. 2002], PRT pre-computes and stores a linear operator M_p at every surface point \mathbf{p} in the scene. M_p transforms the source lighting vector l into a transferred incident radiance vector $l^T(\mathbf{p}) = M_p l$ with $l^T(\mathbf{p})$ representing the local incident radiance at \mathbf{p} . M_p attenuates the source lighting by shadowing and increases it through inter-reflections.

We pre-compute M_p with a ray-tracer, in which both BRDF and BTDF are evaluated at each surface point to account for the fact that leaves are translucent. The exit radiance at a surface point \mathbf{p} for the given view direction \mathbf{v}_p , $e(\mathbf{p}, \mathbf{v}_p)$, is computed as a dot product $e(\mathbf{p}, \mathbf{v}_p) = b(\mathbf{v}_p) l^T(\mathbf{p})$, where $b(\mathbf{v}_p)$ is a view-dependent BRDF-BTDF vector.

Our source lighting vector $l = S + E$, where S is the sunlight and E is the low-frequency environment light. Since M_p is a linear operator, the exit radiance is

$$e(\mathbf{p}, \mathbf{v}_p) = b(\mathbf{v}_p) S^T(\mathbf{p}) + b(\mathbf{v}_p) E^T(\mathbf{p}), \quad (8)$$

where $S^T(\mathbf{p}) = M_p S$ and $E^T(\mathbf{p}) = M_p E$. In order to capture high-frequency details of the soft shadows cast by the sun, we decompose the transferred sunlight radiance $S^T(\mathbf{p})$ into a direct component $S_d^T(\mathbf{p})$ and an indirect component $S_i^T(\mathbf{p})$. $S_d^T(\mathbf{p})$ consists of all sunlight illumination at \mathbf{p} that comes directly from the sun. $S_i^T(\mathbf{p})$ includes of all indirect sunlight illumination at \mathbf{p} through transmissions and inter-reflections. Thus Eq. (8) becomes

$$\begin{aligned} e(\mathbf{p}, \mathbf{v}_p) &= b(\mathbf{v}_p) S_d^T(\mathbf{p}) + b(\mathbf{v}_p) S_i^T(\mathbf{p}) + b(\mathbf{v}_p) E^T(\mathbf{p}) \\ &= b(\mathbf{v}_p) S_d^T(\mathbf{p}) + b(\mathbf{v}_p) (S_i^T(\mathbf{p}) + E^T(\mathbf{p})). \end{aligned} \quad (9)$$

We shall discuss the first term shortly. For the second term, we obtain $S_i^T(\mathbf{p}) + E^T(\mathbf{p})$ using PRT, with the modification that we only record the indirect component of the transferred sunlight radiance. This modification is straightforward because the transfer operator

M_p is pre-computed with a ray tracer, and the direct illumination is simply the first light bounce. Since PRT projects M_p , S , and E onto a low-order spherical harmonics basis, we only get low-frequency visual effects for the second term in Eq. (9). This is not an issue for E which is assumed to be of low-frequency. For the sunlight S , the limitation of PRT implies that inter-reflections involving S are captured only at low-frequencies.

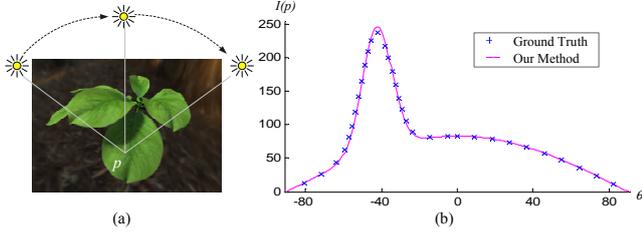


Figure 4: BRDF approximation inside the solid angle subtended by the sun. (a) At point P , we compare the rendering results with and without the approximation for a variety of sunlight directions. (b) The intensity comparison at P . The ground truth is generated by ray tracing.

Now we examine the first term in Eq. (9). We wish to compute $b(\mathbf{v}_p)S_d^T(\mathbf{p})$ directly without involving low-order spherical harmonics basis and thus avoid the loss of high-frequency visual effects. By definition

$$b(\mathbf{v}_p)S_d^T(\mathbf{p}) = \int_{\Omega} f_r(\mathbf{s}, \mathbf{v}_p)S_d(\mathbf{s})V(\mathbf{p}, \mathbf{s})s_z ds,$$

where $f_r(\mathbf{s}, \mathbf{v}_p)$ is the BRDF, $S_d(\mathbf{s})$ is the sunlight as a function of the light direction \mathbf{s} , $V(\mathbf{p}, \mathbf{s})$ is the visibility function of the sun at \mathbf{p} , s_z is the ‘‘cosine factor’’ (z -component of \mathbf{s}), and Ω is the hemisphere of light directions. For general incident lighting, the integral in the above expression is quite expensive to evaluate. For the special case of sunlight, we can quickly calculate this integral by pre-computing the light-visibility convolution at all vertices in the scene.

Light-Visibility Convolution: We model the sun as an area light source of the shape of a circular disk. Let Ω_0 be the solid angle extended by the sun disk and \mathbf{s}_0 be the sunlight direction. $S_d(\mathbf{s})$ is non-zero only inside Ω_0 . Since the sun is far away, Ω_0 is very small and we have

$$\begin{aligned} b(\mathbf{v}_p)S_d^T(\mathbf{p}) &= \int_{\Omega_0} f_r(\mathbf{s}, \mathbf{v}_p)S_d(\mathbf{s})V(\mathbf{p}, \mathbf{s})s_z ds \\ &\approx f_r(\mathbf{s}_0, \mathbf{v}_p)V_{s_0}(\mathbf{p}), \end{aligned} \quad (10)$$

where

$$V_{s_0}(\mathbf{p}) = \int_{\Omega_0} S_d(\mathbf{s})V(\mathbf{p}, \mathbf{s})s_z ds \quad (11)$$

is called the *light-visibility convolution* (LVC) at \mathbf{p} . $V_{s_0}(\mathbf{p})$ is essentially a shadow factor that accounts for the illumination at \mathbf{p} by an area source. In Eq. (10) we approximately regard the BRDF as constant inside the solid angle Ω_0 extended by the sun disk. We found that this is a fairly accurate approximation, as Fig. 4 demonstrates.

For a given sunlight direction \mathbf{s}_0 , the *light visibility map* V_{s_0} consists of the LVC values $V_{s_0}(\mathbf{p})$ of all vertices \mathbf{p} in the scene. Fig. 5 shows computation of the light visibility map. The most important fact about the light visibility maps is that they can be pre-computed. With $V_{s_0}(\mathbf{p})$ available, we can quickly evaluate $b(\mathbf{v}_p)S_d^T(\mathbf{p})$ at runtime accordingly to Eq. (8). Here we take advantage of our parametric BRDF model, which is compact and can be loaded into the GPU for fast calculations.

To pre-compute all light visibility maps, we first calculate the LVC value $V_{s_0}(\mathbf{p})$ at all vertices \mathbf{p} for all sunlight directions \mathbf{s}_0 . Then we rebin the LVC data for each sunlight direction \mathbf{s}_0 to obtain the corresponding light visibility map V_{s_0} . For a given vertex \mathbf{p}

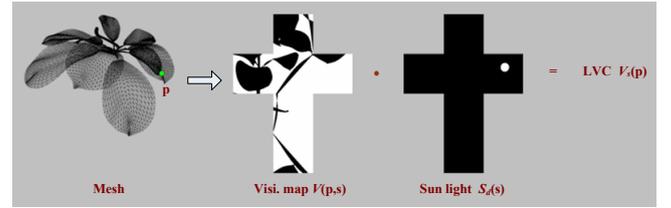


Figure 5: Computing the light-visibility convolution at point \mathbf{p} with the sun visibility $V(\mathbf{p}, \mathbf{s})$ and the sun mask $S_d(\mathbf{s})$.

and sunlight direction \mathbf{s}_0 , the evaluation of $V_{s_0}(\mathbf{p})$ may be thought of as a ray casting process: a set of rays are cast from \mathbf{p} to the sun disk and for each ray, the contribution to the light-visibility convolution integral is calculated using $S_d(\mathbf{s})$ and the sun visibility at \mathbf{p} . The final value of the integral is the sum of contributions of all cast rays.

Because the evaluation of $V_{s_0}(\mathbf{p})$ is part of the pre-processing step, we can afford to use more expensive techniques without worrying about our system’s run-time performance. Nevertheless, there is an efficient way to compute $V_{s_0}(\mathbf{p})$ for all sunlight directions \mathbf{s}_0 . With a cube map placed around \mathbf{p} , we can render the scene onto the cube map using graphics hardware, producing in effect the values of $S_d(\mathbf{s})V(\mathbf{p}, \mathbf{s})s_z$ at all cube map pixels. Then for every sunlight direction \mathbf{s}_0 , we obtain the light-visibility convolution integral as the pixelwise dot-product of the cube map with the sun mask corresponding to \mathbf{s}_0 .

Compression: The collection of light visibility maps of all sunlight directions is fairly large and needs compression for efficient processing. For a $32 \times 32 \times 6$ environment map and a scene with 100k vertices, the collection of all light visibility maps takes about 600 MB (each pixel of a light visibility map takes a byte). Fortunately, to render a given frame we only need to uncompress a single light visibility map, since the sunlight direction is fixed per frame. A light visibility map is small (100 KB in the above example) and can be decoded quickly. We uncompress the light visibility map on the CPU and upload the result onto the GPU as vertex attributes. For compression, we apply run-length encoding (RLE) to each light visibility map. RLE is a lossless scheme that preserves image quality and supports real-time rendering. For the above example, RLE compresses the 600 MB of light visibility data down to 100 MB. Other compression schemes of course could be used to improve the compression ratio.

Rebinning the LVC values for each sunlight direction is important for compression. If the LVC values were rebinned for every vertex, we would have to randomly access these data when rendering each frame. In that case, data coding and decoding becomes difficult due to the random access.

Level of Detail: To accelerate our PRT rendering pass, we construct a discrete geometry LOD for each leaf mesh and derive radiance transfer matrices for all LOD meshes. Specifically, we first pre-compute the radiance transfer matrices at the finest-level mesh vertices. Then we derive the radiance transfer matrices at coarse-level mesh vertices using a simple averaging scheme with Gaussian weights. Let \mathbf{p} be a vertex on the coarse mesh. The transfer matrix at \mathbf{p} is a weighted sum of the transfer matrices of vertices of the finest mesh (that is, only those vertices within a given radius r from \mathbf{p} within the surface, rather than the spherical neighborhood.) Once this LOD hierarchy is constructed for PRT, rendering is straightforward: we need only to determine the current mesh LOD level for each vertex, and then compute the radiance transfer from the corresponding adjacent PRT LODs.

Discussion: Instead of pre-computing and storing the LVC values at all vertices and for all sunlight directions, we could try to cast soft shadows on-the-fly using a shadow algorithm. However, it is hard to compute soft shadows for a large leaf assembly because it tends to generate complicated self-occlusions. For example, in a leaf assem-

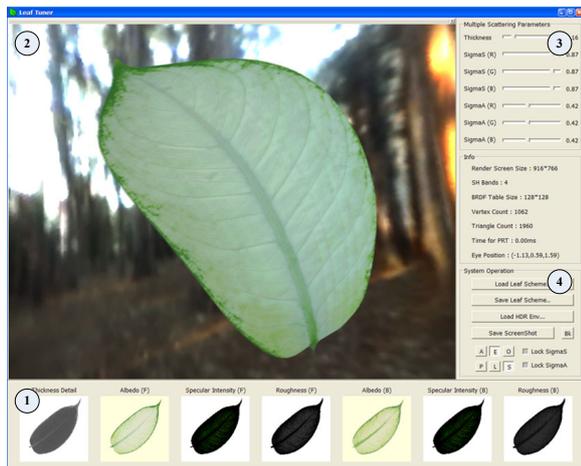


Figure 6: The user interface for editing leaf appearance.

bly, a leaf is often simultaneously a receiver and an occluder, which makes it impossible to use convolution shadow textures [Soler and Sillion 1998]. With a large number of leaves, shadow volume techniques including [Assarsson and Akenine-Möller 2003] would suffer impractically heavy fill rates. “Smoothies” is a fast soft shadow technique that seems applicable to leaves, however, the geometrically approximate shadows could cause disruptive artifacts, particularly in close-up views [Chan and Durand 2003].

Although we developed our two-pass algorithm with sunlight in mind, our approach is applicable to the rendering of other types of scenes illuminated by an environment map as well as several small area light sources. There is no restriction on the shapes of the area light sources, but they must be small enough for the constant BRDF approximation to be sufficiently accurate.

5 Results

We implemented our system in OpenGL on a PC with a 2.8 GHz Pentium IV processor and an ATI Radeon 9800Pro graphics card. We also built an LLS device following [Gardner et al. 2003]. In this section we report rendering results using leaf reflectance and transmittance data acquired with our LLS device.

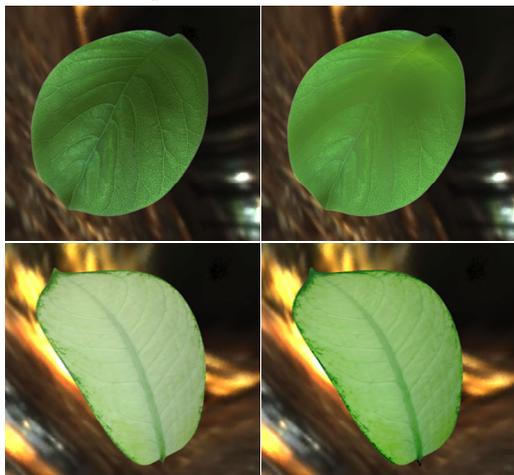


Figure 7: Appearance editing examples. Top row shows the results of changing leaf thickness, and the thicker leaf is shown on the left. The bottom row shows the results of changing the absorption coefficient σ_a .

Leaf Model and Appearance Editing: A leaf model consists of two pairs of BRDFs and BTDFs: one pair for the top surface of the leaf slab and the other for the bottom. For each BTDF, we

store the thickness detail map $\delta(\mathbf{x})$. For each BRDF, we store three maps: the albedo map $\gamma(\mathbf{x})$, the specular intensity map $\rho_s(\mathbf{x})$, and the specular roughness map $m(\mathbf{x})$. These four maps are stored as two RGBA textures of resolution 720×540 . One texture contains $\gamma(\mathbf{x})$ in its RGB channels and $\delta(\mathbf{x})$ in the alpha channel. The other texture contains $\rho_s(\mathbf{x})$ in its RGB channels and $m(\mathbf{x})$ in the alpha channel. In total a leaf model is stored as four RGBA textures.

Fig. 6 shows a simple user interface for editing leaf appearance. The bottom panel (Panel #1) exhibits the parameter maps such as the albedo maps, the specular intensity maps, and the specular roughness maps for the top and bottom leaf surfaces. The top right panel (Panel #3) controls the subsurface scattering parameters including the leaf thickness and the RGB channels of the absorption and scattering coefficients. The editing result is interactively displayed in the middle window (Window #2) to give the user immediate feedback. Fig. 7 shows examples of leaf appearance editing. Appendix B in the conference DVD provides additional editing examples.

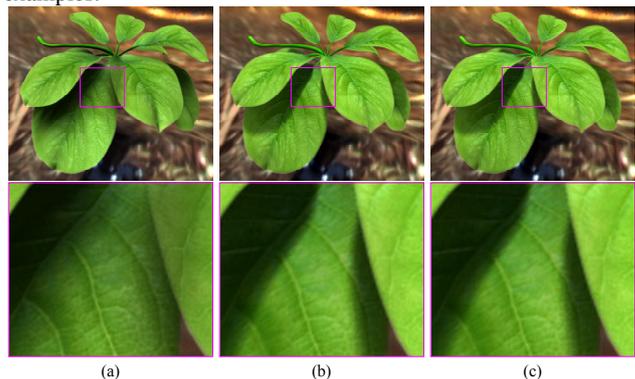


Figure 8: Rendering quality comparison. (a) The result by PRT. (b) Our result. (c) The ground truth as rendered by ray tracing.

Rendering: Fig. 8 compares the rendering quality of our system with PRT and the ground truth generated by ray tracing. Our result compares favorably with the ground truth. As expected, high-frequency details of shadows are lost in the PRT result but they are well captured by our system.

Fig. 1 and Fig. 10 show images with a variety of different leaves that were modeled and rendered with our system. Notice the soft edge of the elm leaves is well captured, which would be difficult to do using previous approaches. Table 1 provides detailed information for Fig. 1 and Fig. 10. The rendering resolution is 800×600 . The “SH data” column provides for each model the number of mesh vertices and the size of spherical harmonics data in megabytes (MB). This data is used for the PRT rendering pass, in which we use fifth order spherical harmonics with 25 coefficients. The “LVC data” column provides for each model the number of mesh vertices and the size of the light-visibility convolution data. Since spherical harmonics data only contains low-frequency information, a relatively sparse sampling is sufficient. For the all-frequency LVC data, a higher sampling density is required.

Model	SH data		LVC data		Fps
	# vertices	size (MB)	# vertices	size (MB)	
Balata	3799	9.3	33767	50.6	41.2
Omoto	2778	7.4	59770	69.6	38.5
Pelargonium	3927	10.4	58169	60.1	40.8
Alpinia	2978	8.0	43154	65.9	45.9
Prunus	2534	6.8	43126	55.8	37.2
Elm	3827	9.4	44687	59.8	41.8

Table 1: Rendering performance statistics for Fig. 1 and Fig. 10. The rendering speeds are reported in the “fps” column. The pre-computation time for each of these examples is about 15 minutes.

Fig. 11 shows the rendering results of a balata tree with over 500k vertices. For such a large model, our system achieves renders at about 10 fps.

Fig. 9 compares our method with the all-frequency shadow technique [Ng et al. 2003]. The image qualities of the two approaches are comparable; the main difference is in rendering speeds. For *dynamic viewpoint* and dynamic lighting, our method achieves over 35 fps. For glossy surfaces such as leaf surfaces, an all-frequency approach with dynamic viewpoint is not currently available. For this reason, the all-frequency approach with *fixed viewpoint* (called “image relighting” in [Ng et al. 2003]) is used and it runs at the speed of about 5 fps. In image relighting, a $32 \times 32 \times 6$ environment map is used and about 1000 wavelet coefficients are retained (16%). The visibility value is quantized to 1 byte. The compression ratio is about 6.25. The image size is 800×600 .

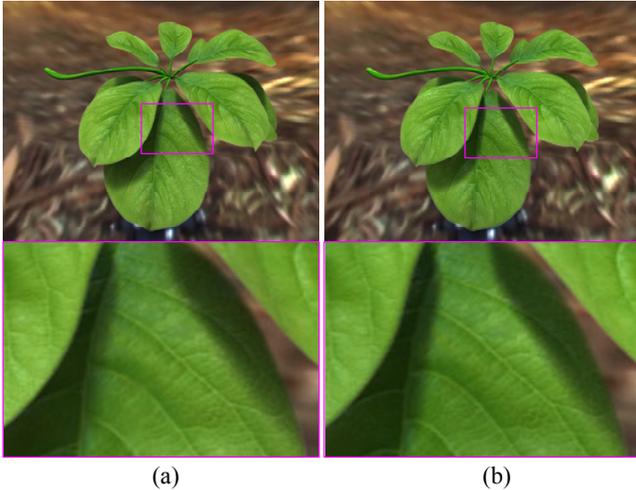


Figure 9: Quality comparison with the all-frequency approach. (a) Result by the all-frequency approach. (b) Our result.

6 Summary and Discussion

We have presented a framework for real-time rendering of plant leaves with global illumination effects. A key component of our framework is a parametric leaf model that is both sophisticated enough to incorporate subsurface scattering and rough surface scattering and compact enough to support real-time rendering. Our leaf model can be captured from real leaves, which makes it easy to create highly realistic leaf appearance models. Another important component of our framework is a two-pass algorithm that renders global illumination effects without losing high-frequency details of shadows. We have demonstrated our framework by rendering a variety of plant leaves.

Our system suggests several interesting areas for future work. The geometric modeling of leaf structures is an active area in computer graphics. An attractive area for future work might involve combining our appearance modeling technique with a more elaborate geometric model. Our leaf model does not consider small features such as hairs on leaves [Führer et al. 2004]. This is a topic that merits much additional research. Finally, our framework is aimed at the rendering of leaves from broadleaf plants. In the future, we would like to develop an approach for rendering other types of leaves, such as those from conifers.

Our work demonstrates that when measured data is available, it can lead to significantly simpler appearance models, without compromising the quality – indeed, in this case, substantially enhancing the speed and quality – of the rendered results. Moreover, there is a difference between the model required to *predict* the details of how light interacts with a material and the model needed to *describe* that interaction. That is, there is no need to model the interior of a material in order to predict an appearance that can be captured via measurement. This is a new way to think about appearance modeling of leaves and other thin objects, which differs from conventional

thinking and practice. However, we demonstrate that this approach is both viable and promising.

While there has been extensive recent work in computer graphics in the areas of capture and real-time rendering, these two areas have to large extent developed in isolation from one another. In developing a complete, end-to-end system, we demonstrate that there is great benefit in establishing a relationship between these two seemingly disparate problems: what is captured and how it is stored can be coupled directly to the rendering algorithm, thereby yielding efficient, high-fidelity rendering of materials with very complex appearances.

Last, in computer graphics, the rendered results, as perceived by the viewer, are ultimately what count. Raw, measured data, on the other hand, can lead to results that are far from what the user desires. One of the key challenges in appearance modeling is determining and providing the appropriate handles, such that a user can achieve the desired results. The ability to edit select aspects of appearance data is crucial in this regard. The leaf model presented in this paper is an example of a model that provides intuitive control parameters. An exciting area for future work is the development of comparable models for a broader classes of materials.



Figure 10: Rendering results of clusters of elm and balata leaves.

Acknowledgement: The authors would like to thank Jiaping Wang and Andrew Gardner for their help in building LLS device. Many thanks to Zhunping Zhang for implementing the initial system for leaf appearance editing and to Steve Lin for his help in video production. We also want to thank Przemyslaw Prusinkiewicz and Xin Tong for useful discussions. The geometry models of leaves and branches were made by Mingdong Xie. We are grateful to the anonymous reviewers for their helpful suggestions and comments.

References

- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *Proceedings of SIGGRAPH '00*, 65–74.
- ASSARSSON, A., AND AKENINE-MÖLLER, T. 2003. A geometry-based soft shadow volume algorithm using graphics hardware. *ACM Transactions on Graphics* 22(3), 511–520.
- BARANOSKI, G. V. G., AND ROKNE, J. G. 1997. An algorithmic reflectance and transmittance model for plant tissue. *Computer Graphics Forum* 16(3), 141–150.
- BARANOSKI, G., AND ROKNE, J. 2001. Efficiently simulating scattering of light by leaves. *The Visual Computer* 17(8), 491–505.
- BARANOSKI, G. V. G., AND ROKNE, J. G. 2002. *Light Interaction with Plants*. SIGGRAPH '02 Course Notes.
- BECKMANN, P., AND SPIZZICHINO, A. 1963. *The Scattering of Electromagnetic Waves from Rough Surfaces*. MacMillan, New York.
- BLOOMENTHAL, J. 1985. Modeling the mighty maple. *Computer Graphics (SIGGRAPH '85 Proceedings)* 19, 305–311.
- CHAN, E., AND DURAND, F. 2003. Rendering fake soft shadows with smoothies. *Proc. of the Eurographics Symposium on Rendering 2003*.
- COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Transactions on Graphics* 1(1) (Jan.), 7–24.



Figure 11: Rendering result of trees.

- COOK, TORRANCE, K. E. AND SPARROW, E. M. 1967. Theory for Off-Specular Reflection from Roughened Surfaces. *Journal of the Optical Society of America* 57(9) (sep.), 1105-1114.
- DE REFFYE, P., EDELIN, C., FRANCON, J., JAEGER, M., AND PUECH, C. 1988. Plant models faithful to botanical structure and development. *Computer Graphics, Proceedings of Siggraph'88* 22(4), 151-158.
- DEMKO, S., HADGES, L., AND NAYLOR, B. 1985. Construction of fractal objects with iterated function system. *Computer Graphics* 19(3), 271-278.
- DEUSSEN, O., HANRAHAN, P. M., LINTERMANN, B., MECH, R., PHARR, M., AND PRUSINKIEWICZ, P. 1998. Realistic modeling and rendering of plant ecosystems. In *Proceedings of SIGGRAPH 98*, 275-286.
- FRANZKE, O., AND DEUSSEN, O. 2003. Rendering plant leaves faithfully. *SIGGRAPH '03 Sketches*.
- FUHRER, M., JENSEN, H. W., AND PRUSINKIEWICZ, P. 2004. Modeling hairy plants. In *Proc. of Pacific Graphics '04*.
- GANAPOL, B., JOHNSON, L., HAMMER, P., HLAVKA, C., AND PETERSON, D. 1998. LEAFMOD: A new within-leaf radiative transfer model. *Remote Sensing of Environment* 63, 182 - 193.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. *ACM Transactions on Graphics* 22(3) (July), 749-758.
- GOLOMB, S. 1966. Run-Length Encodings. *IEEE Transactions on Information Theory* 12, 399-401.
- GOVAERTS, Y., VERSTRAETE, S. J. M., AND USTIN, S. 1996. Three-dimensional radiation transfer modeling in a dicotyledon leaf. *Applied Optics* 35(33), 6585 - 6598.
- HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. *Proceeding of Siggraph '93*, 165-174.
- JACQUEMOUD, S., AND USTIN, S. 2001. Leaf optical properties: A state of the art. In *Proc. 8th Int. Symp. Physical Measurements and Signatures in Remote Sensing*, 223-232.
- KAJIYA, J. T. 1985. Anisotropic reflection models. In *Computer Graphics (Proceedings of SIGGRAPH 85)*, vol. 19, 15-21.
- KAUTZ, J., SLOAN, P.-P., AND SNYDER, J. 2002. Fast, arbitrary BRDF shading for low-frequency lighting using spherical harmonics. *Proceedings of the 12th Eurographics Workshop on Rendering*, 301-308.
- MA, Q., NISHIMURA, A., PHU, P., AND KUGA, Y. 1990. Transmission, reflection and depolarization of an optical wave for a single leaf. *IEEE Trans. on Geoscience and Remote Sensing* 28, 5 (september), 865 - 872.
- MAX, N. 1996. Hierarchical rendering of trees from precomputed multi-layer z-buffers. In *Eurographics Rendering Workshop 1996*, 165-174.
- MEYER, A., NEYRET, F., AND POULIN, P. 2001. Interactive rendering of trees with shading and shadows. *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, 183-196.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2003. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Transaction on Graphics* 22(3) (July), 376-381.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2004. Triple product wavelet integrals for all-frequency relighting. *ACM Transaction on Graphics* 23(3) (August), 477-487.
- OREN, M., AND NAYAR, S. K. 1994. Generalization of Lambert's reflectance model. In *Computer Graphics (Proceedings of SIGGRAPH 94)*, 239-246.
- POULIN, P., AND FOURNIER, A. 1990. A model for anisotropic reflection. In *Computer Graphics (Proceedings of SIGGRAPH 90)*, vol. 24, 273-282.
- PRUSINKIEWICZ, P., LINDENMAYER, A., AND HANAN, J. 1988. Development models of herbaceous plants for computer imagery purposes. In *Computer Graphics (Proceedings of SIGGRAPH 1988)*, 141-150.
- PRUSINKIEWICZ, P., MUENDERMANN, L., KARWOWSKI, R., AND LANE, B. 2001. The use of positional information in the modeling of plants. *Proceedings of Siggraph '01* (August), 289-300.
- QIN, X., NAKAMAE, E., TADAMURA, K., AND NAGAI, Y. 2003. Fast photo-realistic rendering of trees in daylight. *Computer Graphics Forum* 22(3), 243-252.
- RECHE, A., MARTIN, I., AND DRETTAKIS, G. 2004. Volumetric reconstruction and interactive rendering of trees from photographs. *ACM Transactions on Graphics* 23(3) (July), 720-727.
- SIEWERT, C. E. 1978. The fn method for solving radiative-transfer problems in plane geometry. *Astrophysics and Space Science* 58, 131-137.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transaction on Graphics* 21(3), 527-536.
- SLOAN, P.-P., HALL, J., HART, J., AND SNYDER, J. 2003. Clustered principal components for precomputed radiance transfer. *ACM Transaction on Graphics* 22(3) (July), 382-391.
- SOLER, C., AND SILLION, F. 1998. Fast calculation of soft shadow textures using convolution. *Proceeding of SIGGRAPH '98* (July), 321-332.
- STOGRYN, A. 1967. Electromagnetic scattering from rough, finitely conducting surface. *Radio Sciences 2 (New Series)*, 4, 415-428.
- VOGELMANN, T. C. 1993. Plant tissue optics. *Annual Review of Plant Physiology and Plant Molecular Biology* 44, 231-251.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. *Proceedings of SIGGRAPH '92*, 265-272.
- WEBER, J., AND PENN, J. 1995. Creation and rendering of realistic trees. *Computer Graphics (Proceeding of SIGGRAPH 1995)*, 119-128.