Volume xx (200y), Number z, pp. 1–12

Physically Based Real-Time Translucency for Leaves

Ralf Habel Alexander Kusternig Michael Wimmer

Vienna University of Technology

Abstract

This paper presents a new shading model for real-time rendering of plant leaves that reproduces all important attributes of a leaf and allows for a large number of leaves to be shaded. In particular, we use a physically based model for accurate subsurface scattering on the translucent side of directly lit leaves. For real-time rendering of this model, we formulate it as an image convolution process and express the result in an efficient directional basis that is fast to evaluate. We also propose a data acquisition method for leaves that uses off-the-shelf devices.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Realtime Rendering, Natural Phenomena, Natural Scene Rendering, Physically Based Rendering

1. Introduction

The realistic rendering of plants is an important research topic in computer graphics. In addition to the geometric complexity of plants, the complex light interactions especially inside leaves makes a realistic and efficient display of plants very challenging. Recently, Wang et al. [WWD*05] have incorporated a number of illumination effects into a real-time rendering algorithm for leaves. Their algorithm handles reflectance and translucency, and precomputes the radiance transfer in a whole tree both for indirect illumination and direct sunlight illumination. However, their translucency model is basically Lambertian, and ignores all multiple scattering effects that occur due to changing properties along the leaf surface. These effects, which can only be captured using a BSSRDF [NRH*77] model, include self-shadowing on the leaf surface, varying leaf thickness, and varying leaf surface colors.

In this paper, we show that subsurface scattering effects are crucial for faithfully representing the translucent sides of directly lit leaves (see Figure 11 for an illustration). Therefore, we use a physically based subsurface simulation for thin slabs [DJ05] to account for these effects, and derive a per-pixel translucency transport function that varies according to the direction of the incident light. In order to achieve real-time performance, we introduce the recently developed 'Half Life 2 basis' [McT04] as a general directional basis and expand the translucency transport function into this basis.

Figure 1: *A tree featuring physically based translucency.*

The main contribution of this paper is that we are the first to propose a model for accurate real-time translucent subsurface scattering in plant leaves. While our translucency model could be easily integrated into the framework of Wang et al., the fact that it is local to a leaf makes it straightforward to combine with instancing, animation and shadow mapping, which is important for dynamic, complex scenes. We also show an efficient acquisition method that allows capturing the most important leaf properties using off-the-shelf hard-

ware. Figure 1 shows an example of a scene rendered with real-time physically based translucency.

2. Related Work

A number of techniques have been proposed to calculate the light interactions of leaves. Hanrahan and Krueger [HK93] model subsurface scattering in layered surfaces in terms of one-dimensional transport theory, derive explicit formulae for backscattering and transmission and evaluate the BDFs using a Monte Carlo approach. A model based on the available biological information was proposed by Baranoski [BR97]. He later extended this method by precomputing the reflectance and transmittance values and applying a simplified scattering model [BR01]. The LEAFMOD model proposed by Ganapol et al. [GJH*98] solves the one-dimensional radiative transfer equation in a slab with homogeneous optical properties and generates an estimate of leaf reflectance and transmittance. This model is also used by Wang et al. [WWD*05]. All these methods have in common that they do not take the full subsurface scattering rendering equation using a (8D) BSSRDF [NRH*77] into account for translucency, and therefore cannot simulate variations on the leaf surface due to self shadowing, changes in reflectance parameters or thickness. Instead, these models rely on a (4D) BTDF (bidirectional transmission distribution function).

A hypothetical way to incorporate all scattering effects in a leaf is to construct an apparatus for recording a translucent BTF (bidirectional texture function [DvGNK99, MMS*05]) for a leaf. However, current BTF methods still have significant memory and rendering cost and do not allow high resolutions. Another way to accurately simulate light propagation in a leaf, proposed by Govaerts et al. [GVU96], is to represent the internal three-dimensional structure of leaf tissue and their optical properties geometrically and apply ray tracing to this model. Obviously, this is slow. The most practical way to compute subsurface scattering in a homogeneous material was introduced by Jensen [JMLH01] and is based on an analytical diffusion dipole approximation. Franzke et al. [FD03] showed a simplified single scattering algorithm for leaves based on this idea. However, Jensen's dipole approximation only works for semi-infinite slabs, which excludes leaves. Donner and Jensen [DJ05] extended this to accurately and efficiently calculate subsurface scattering in multi-layered thin slabs. This multi-dipole approach is the basis for the efficient precomputation process for leaf translucency shown in this paper.

Subsurface scattering is an active research area with many results, e.g., for more general materials like skin [BL05], more general lighting conditions [WTL05], or deformable models [MKB*03].

3. A Leaf Model for Real-Time Rendering

3.1. Overview

Before going into details on our new translucency model, this section describes the underlying leaf model and the associated acquisition process. Similar to Wang et al. [WWD*05], we model a leaf as a thin slab of homogeneous material with rough front and back surfaces. The spatially varying reflectance is encoded in an albedo map $\alpha(\vec{x})$, and variations in leaf thickness in a thickness map $d(\vec{x})$. To specify the color of the translucency calculated with the new model, we also save an average translucency map $\rho_t(\vec{x})$. Finally, a normal map $\vec{n}(\vec{x})$ gives us the possibility to accurately simulate high-frequency specular reflections. All mentioned maps exist for both the front and the back surface of the leaf.

The light interaction in a leaf is determined by several terms:

$$L = L^{S_d} + L^{S_i} + L^E,$$

i.e., the contribution of direct sunlight, indirect sunlight and environment lighting. The contribution of direct sunlight can be split into

$$L^{S_d} = L_r + L_t,$$

the reflective and translucent components, where only one is non-zero depending on the dot product between leaf normal and light direction. In this paper, we focus on the translucency from direct sunlight illumination, L_t. We do not deal with indirect or environment illumination (L^{S_i}, L^E), which are approximated by an ambient term in our results. These terms can be calculated following the model of Wang et al. [WWD*05], which requires precomputation for a whole tree. This is quite costly in terms of memory and precludes animation and instancing of leaves.

We briefly cover the reflectance term L_r and acquisition in sections 3.2 and 3.3, while we dedicate sections 4 and 5 to the main contribution of this paper, the leaf translucency model for L_t and its real-time evaluation.

3.2. Leaf reflectance

The structure of a leaf is mostly perceived in its specular reflectance properties due to direct sunlight illumination (L_r), which reveals its high frequency structures. There is a huge variety of leaf BRDFs, ranging from velvet-like due to micro-hairs to highly specular caused by a thick waxy layer. In most cases, the front of a leaf has broad specularity whereas the back of a leaf is diffuse. This is mostly true for tree leaves, but there are also plants which have a highly specular back side, such as *Larrea tridentata*.

Following Bousquet [BLJM05] and Wang et al. [WWD*05], we use the Cook-Torrance shading model [CT82] for the front side of the leaf, and diffuse

shading for the back side. We have considered using a BSSRDF approach also for reflective subsurface scattering, but we have found the difference to a standard BRDF model to be negligible because the non-directional contribution due to subsurface scattering is included in the albedo map from to the acquisition process. This is because in contrast to skin rendering, we calculate subsurface scattering in a very thin and highly absorbing medium, keeping the range of light diffusion down to a few millimeters at maximum.

Since we do not measure the leaf BRDFs ourselves, we take the measured and fitted specular coefficients from Bousquet, ranging from $n = 1.2 - 1.7$ for the effective refractive index and $\sigma = 0.078 - 0.5$ for the roughness, covering highly specular leaves (e.g. *Laurel*) to nearly diffuse specular lobes (e.g. *Hazel*) (see Figure 2).

Our acquisition process provides high resolution normal maps. Therefore, in contrast to previous work which also uses analytical expressions for leaf reflectance, we evaluate the BRDF using the much more accurate normal from the normal map instead of the geometric normal. The diffuse term is taken from the albedo map. The reflective contribution from direct sunlight, which is modeled as a directional light $L(\omega) = L_D \delta(\vec{\omega} - \vec{\omega}_D)$ with light intensity L_D, light direction $\vec{\omega}_D$, and Dirac function δ, therefore evaluates to:

$$L_r(\vec{x}, \vec{\omega}_o) = L_D\left(\frac{\alpha(\vec{x})}{\pi} + f_s(\vec{n}(\vec{x}), \vec{\omega}_D, \vec{\omega}_o)(\vec{n}(\vec{x}) \cdot \vec{\omega}_D)\right) \quad (1)$$

where $\vec{x}$ is a surface point, $\vec{\omega}_o$ is the outgoing direction and f_s is the specular BRDF. The complete reflective BRDF is thus $f_r = \alpha + f_s$. Note that to arrive at the diffuse term, $\alpha(\vec{x})$ would theoretically have to be reduced by the albedo of f_s, however this term is negligible at non-grazing angles.

3.3. Data Acquisition

We briefly describe how we create the required maps for albedo, normal, thickness and average translucency. Our acquisition setup allows the generation of high (i.e., <1mm) resolution maps using a simple process and (slightly expensive) off-the-shelf hardware. This is mostly due to the use of an accurate 3D scanner.

The devices used are a 3D scanner operating at an effective resolution of 0.1 mm (Minolta VI-910), a digital camera (Canon EOS 20D) with fixed exposure time, two 1,000 Watt light sources with large box diffusers, and an easy to construct fixing frame for the leaf. The large diffusers allow approximating hemispherical illumination, which is required for capturing the albedo. The fixing frame guarantees that the leaf remains unchanged during the acquisition process, so that the acquired data is consistent. The leaf is fixed in the frame using wire bridges with small clamps to keep the leaf in a straight but natural position. We arrange the 3D scanner, camera and diffusers as shown in Figure 3.

Figure 2: *Quad patches shaded with highly specular (top) and almost diffuse (bottom) reflectance, and with a directional light at steep (left) and grazing angle (right).*

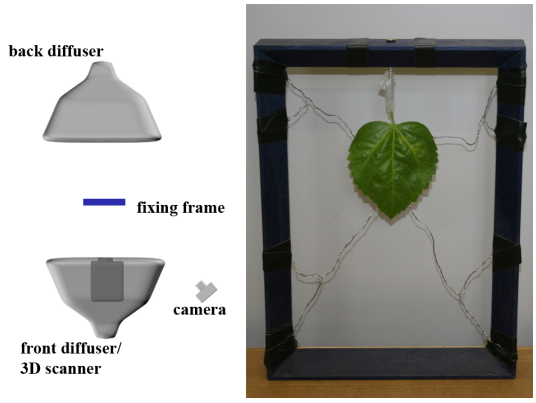

Figure 3: *Schematics of the acquisition setup and a close up of the fixing frame. For the 3D scan, the 3D scanner replaces the diffuser.*

The leaf is sampled by first taking a 3D scan, then replacing the scanner with a diffuser and recording the albedo using the camera, then switching to the back diffuser and recording the translucency. The same steps are repeated for the back side of the leaf after carefully turning the fixing frame.

For postprocessing we use standard tools like Geomagic and Maya. After some filtering, moderate smoothing so as not to impact the high frequency structures of the leaf, and aligning the captured images with the geometry, we geometrically simplify the mesh and encode the high frequency

geometric structures as a normal map [COM98] and a displacement map for each side. The thickness map is generated by subtracting the displacement maps of the front and back side scans, normalized to a user-defined maximum thickness of the leaf. The normal maps generated in this way can be mapped to different geometric levels of detail of the leaf (Figure 4). The generated textures may need highlight removal using standard image processing techniques if the used leaf has strong specular properties. Figure 5 shows a complete data set generated using the proposed method.

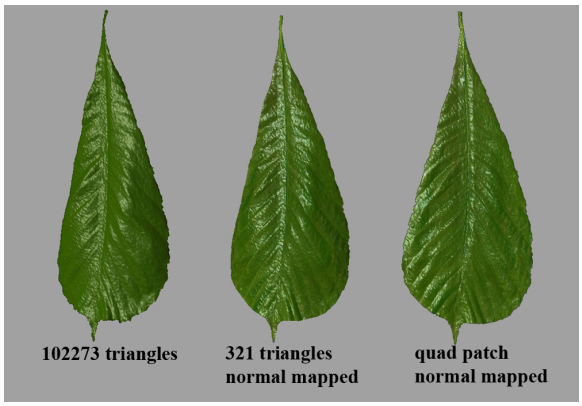

Figure 4: *The scanned geometry, normal-mapped simplified geometry and the normal map on a quad patch. The highlights have been exaggerated for visualization purposes.*

In comparison to the acquisition setup by Wang et al. [WWD*05], we do not capture per-pixel BRDF or BTDF data, which requires a custom-built linear light source device. Therefore we cannot encode spatially varying roughness or specular intensity, but have to rely on one set of parameters for the whole leaf. On the other hand, we create high-resolution normal maps, which causes highlights to be placed more accurately according to the high-frequency structure of the leaf. This is important since specular highlights due to direct illumination are the most prominent feature on the front side of a leaf.

4. Leaf Translucency

One of the main insights in this paper is that while subsurface scattering has only negligible impact on the appearance of the light-facing side of a leaf, it is the dominant factor for the opposite side. Figure 11 demonstrates the difference between a simple, yet state-of-the-art translucency model based on a diffuse BTDF, and the model proposed in this paper based on a BSSRDF. As opposed to the reflective part of the leaf, where high-frequency features are conveniently modeled using a normal map, there is as yet no way to introduce lighting effects due to surface variations in the translucent part, and therefore the simple model is restricted to a simple diffuse shading model using the geometric normal of

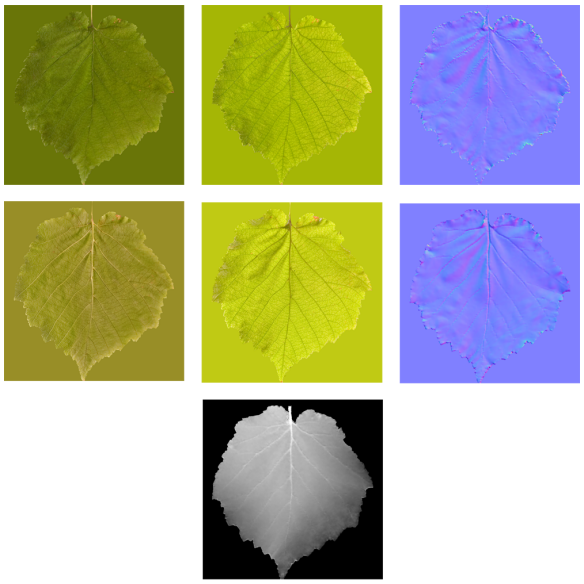

Figure 5: *A complete data set of a leaf, consisting of albedo (left), translucency (middle) and normal map (right) for both sides and a thickness map (bottom).*

the leaf. With our new method, on the other hand, depending on the incident light angle, the leaf appears either smooth (steep angle) or shows high-frequency details (grazing angle) due to leaf structures such as risps or bulges.

The main features taken into account by the BSSRDF model are self shadowing of the leaf before the light penetrates into the leaf interior, variations in leaf thickness, and variations of the reflectance properties over the light-facing leaf surface. These effects lead to variations in the amount of light entering the medium and scattering towards a specific point to be shaded on the opposite leaf surface. Note that our model is local to a leaf, and therefore light variations due to shadows from other leaves or similar only modify the resulting radiance, but do not enter into the subsurface scattering computations (these effects are handled using standard real-time shadow algorithms in our case). However, the influence of such large-scale structures on subsurface scattering in leaves, which happens on a much smaller scale, is negligible.

Recently, Donner and Jensen [DJ05] introduced an efficient approximation for computing subsurface scattering in a thin, homogeneous slab. We introduce the main concepts of this method in Section 4.1, and show an efficient formulation as an image convolution in Section 4.2.

Note that all calculations are carried out in the tangent space of the simplified geometry, scaled so as to preserve physical units. The tangent space can safely be assumed to be locally flat in comparison to the typical length of scattering paths, which is a prerequisite for the BSSRDF model.

4.1. Background: Light Diffusion in Leaves

We consider a leaf which is characterized by the absorption coefficient σ_a, scattering coefficient σ_s and the mean cosine of the scattering phase function g.

Scattering in a translucent material is described by the Bidirectional Scattering Surface Reflectance Distribution Function (BSSRDF) [NRH*77]

$$S(\vec{x}_i,\vec{\omega}_i,\vec{x}_o,\vec{\omega}_o) = \frac{dL_o(\vec{x}_o,\vec{\omega}_o)}{d\Phi(\vec{x}_i,\vec{\omega}_i)}, \tag{2}$$

where L_o is the outgoing radiance on the non-light facing side and Φ is the incident flux on the light-facing side of the leaf. $\vec{x}_i$ and $\vec{\omega}_i$ are the incident position and direction and $\vec{x}_o$, $\vec{\omega}_o$ are the exitant position and direction. Jensen et al. [JMLH01] showed that for a semi-infinite homogeneous slab, the BSSRDF can be approximated by a diffusion dipole, i.e., 2 virtual point light sources, which are derived through a diffusion approximation for the radiative transport. In a thin slab, two boundary conditions have to be taken into account: light that reaches either the front or the back of the leaf never returns. To match these boundary conditions, multiple dipoles are required [DJ05]. The boundary conditions are expressed in terms of the scalar irradiance ϕ, also called *fluence* (for algebraic details see [DJ05]):

$$\phi(r) - 2AD\frac{\delta\phi(r)}{\delta z} = 0 \qquad \text{at } z = 0,\ z = d \tag{3}$$

where $D = \frac{1}{3\sigma_t'}$ with the reduced extinction coefficient $\sigma_t' = \sigma_a + (1-g)\sigma_s$. Expression 3 is applied at the front surface at $z = 0$ and the back surface at a depth of $z = d$. The term A represents the change in fluence due to internal reflection at the surface:

$$A = \frac{1+\rho_d}{1-\rho_d}. \tag{4}$$

Since leaves have a rough reflective surface, we use the average diffuse reflection ρ_d of the reflective BRDF f_r, which can be evaluated using sampling instead of a Fresnel term.

In order to match the boundary condition at $z = 0$ given in 3, a real positive point light is placed under $\vec{x}_i$ at a depth of one mean free path $l = 1/\sigma_t'$ [PCW89]. By placing a negative virtual light source at $(1+4A/3)/\sigma_t' = 2z_b + l$, the net fluence at $-z_b = -2AD$ vanishes and results in a good approximation for the first boundary condition at $z = 0$ [FP92], consisting of a dipole configuration. Both the real and virtual light are treated as being inside the participating medium. To satisfy the second boundary condition at $z = d$, we mirror the existing dipole at the extrapolation distance of the second boundary condition at $z = d + z_b$, which in turn gets mirrored at z_b again to match the first boundary condition. This process converges so that both conditions are matched at the same time [DJ05].

The positions $z_{r,j}$ and $z_{v,j}$ of the positive and negative diffusion dipole poles can be expressed with

$$\begin{aligned} z_{r,j} &= 2j(d+2z_b)+l \\ z_{v,j} &= 2j(d+2z_b)-l-2z_b, \qquad j=-n\ldots n. \end{aligned} \tag{5}$$

The resulting fluence field is defined by

$$\phi(r) = \sum_{j=-n}^{n} \frac{\Phi}{4\pi D}\left(\frac{e^{-\sigma_{tr}d_{r,j}}}{d_{r,j}} - \frac{e^{-\sigma_{tr}d_{v,j}}}{d_{v,j}}\right) \tag{6}$$

where $d_{r,i} = \left|\vec{x}_0 - \vec{x}_{r,j}\right|$ is the distance from $\vec{x}_0$ to the real light sources, $d_{v,j} = \left|\vec{x}_0 - \vec{x}_{v,j}\right|$ the distance to the virtual light sources and $\sigma_{tr} = \sqrt{3\sigma_a\sigma_t'}$ is the effective transport coefficient. Figure 6 shows a slice through a leaf along the z axis of the fluence field for one incident radiance point $\vec{x}_i$.

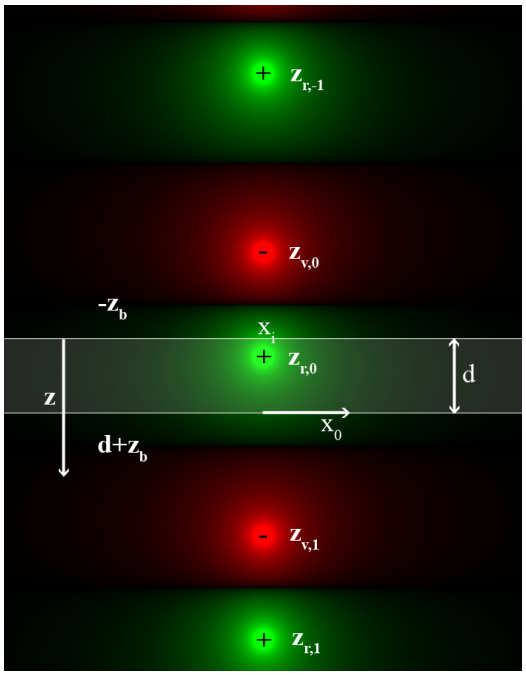

Figure 6: *Fluence field defined by the multi dipole configuration. Green are positive, red are negative values.*

Finally, the diffuse transmittance at the non-light facing surface is equal to the gradient of the fluence, depending on the slab thickness d and the distance $r = |x_o - x_i|$ from the incident point:

$$T(r,d) = -D\frac{\vec{n}\cdot\vec{\nabla}\phi(\vec{x}_0)}{d\Phi(\vec{x}_i,\vec{\omega}_i)} \tag{7}$$

which evaluates to

$$\begin{aligned} T(r,d) = \sum_{j=-n}^{n} \frac{\alpha'}{4\pi}\Bigg(&\frac{(d-z_{r,j})(1+\sigma_{tr}d_{r,j}e^{-\sigma_{tr}d_{r,j}})}{d_{r,j}^3} \\ &-\frac{(d-z_{v,j})(1+\sigma_{tr}d_{v,j}e^{-\sigma_{tr}d_{v,j}})}{d_{v,j}^3}\Bigg), \end{aligned} \tag{8}$$

where $\alpha' = \sigma_s'/\sigma_t'$ is the reduced albedo. Figure 7 shows the transmittance for a fixed set of physical properties, depending on the thickness d and distance r, revealing an exponential falloff as the thickness increases, and a smooth, Gauss curve-like falloff as the distance increases.

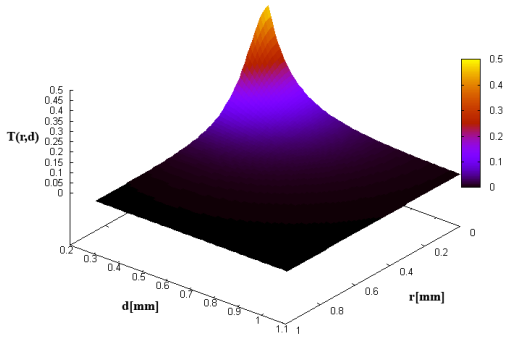

Figure 7: *Transmittance at different thickness d and distance r for a fixed set of physical properties.*

The BSSRDF for translucency is the transmittance within the leaf multiplied by the transmittances at the incident and exiting leaf surfaces:

$$S(\vec{x}_i,\vec{\omega}_i,\vec{x_o},\vec{\omega}_o) = \rho_t(\vec{x}_i,\vec{\omega}_i)T(r,d)\rho_t(\vec{x}_o,\vec{\omega}_o). \quad (9)$$

Again, we use transmittance instead of Fresnel terms because the leaf surfaces are rough. To calculate ρ_t, we assume that all light which is not reflected is transmitted. Therefore, $\rho_t(\vec{x},\vec{\omega}) = 1-\rho_d(\vec{x},\vec{\omega})$, where $\rho_d(\vec{x},\vec{\omega})$ is the reflective albedo. Note that front and back surfaces have different BRDFs and thus different albedos.

4.2. Light diffusion as an image convolution process

In order to achieve a formulation of subsurface scattering that is amenable to real-time rendering, we try to express the BSSRDF evaluation as an image convolution process that operates on the maps available from the acquisition process and allows for fast precomputation.

Reformulating equation (2) gives an expression for the translucent radiance:

$$L_t(\vec{x}_o,\vec{\omega}_o) = \int_\Omega\int_A L(\vec{x}_i,\vec{\omega}_i)S(\vec{x}_o,\vec{\omega}_o,\vec{x}_i,\vec{\omega}_i)(\vec{n}(\vec{x}_i)\cdot\vec{\omega}_i)d\vec{x}_i d\vec{\omega}_i \quad (10)$$

In our case, we assume one directional light of unit intensity ($L_D = 1$) in direction ω_D for the precomputation, represented by a Dirac impulse. We later modulate the result with the actual light intensity to achieve interactive lighting. Also, we take into account self-shadowing from the leaf through a visibility term $V(\vec{x}_i,\vec{\omega}_i)$, so that the radiance arriving at $\vec{x}_i$ is

$$L(\vec{x}_i,\vec{\omega}_i) = \delta(\vec{\omega}_i-\vec{\omega}_D)V(\vec{x}_i,\vec{\omega}_i). \quad (11)$$

Substituting this and (9) into (10) gives

$$L_t(\vec{x}_o,\vec{\omega}_o,\vec{\omega}_D) = \rho_t(\vec{x}_o,\vec{\omega}_o)\int_A T(r,d)E(\vec{x}_i,\vec{\omega}_D)d\vec{x}_i \quad (12)$$

with an irradiance transport function $E(\vec{x},\vec{\omega})$, which describes the light transport from direction $\vec{\omega}$ to the point $\vec{x}$ just below the surface:

$$E(\vec{x},\vec{\omega}) = \rho_t(\vec{x},\vec{\omega})V(\vec{x},\vec{\omega})(\vec{n}(\vec{x})\cdot\vec{\omega}). \quad (13)$$

Equation 12 describes the translucent light transport to $\vec{x}_o$ in a thin slab (leaf) for a given light direction $\vec{\omega}_D$. This equation already has the form of a continuous convolution process of the signal $E(\vec{x},\vec{\omega})$ with kernel $T(r,d)$.

To take the thickness variations into account, we take the measured thickness $d(\vec{x}_i)$ as a local approximation of the boundary conditions that lead to the transmittance term. This makes the convolution kernel $T(r,d(\vec{x}_i))$ non-stationary. This local approximation, which is inherent in the multi-dipole model, matches the actual boundary conditions closely in the case of leaves, except for risp sides, which get slightly softer. Taking into account arbitrary boundary conditions accurately is still the main limitation of dipole approximations in comparion to a full Monte Carlo simulation [DJ05].

Finally, we exploit the fact that all spatially variant variables are given in maps with the same resolution and convert the continuous convolution into a discrete one by discretizing the area integral. As area element, the constant physical area of a texel A_p in the maps according to the scaled tangent plane is used.

$$L_t(\vec{x}_o,\vec{\omega}_o,\vec{\omega}_D) = \rho_t(\vec{x}_o,\vec{\omega}_o)A_p\sum_{\vec{x}_i} T(r,d(\vec{x}_i))E(\vec{x}_i,\vec{\omega}_D) \quad (14)$$

Now the calculation of L_t can be implemented as an image convolution, using albedo, transmission, thickness and normal maps as input. Please note that although equation 14 represents an image convolution, all the variables such as distances still are the physical lengths on the leaf, and apart from discretization, the result is still an exact representation of 10.

This result is not limited to preprocessing for real-time rendering but can also be used in a physically based raytracer, integrating over the light direction $\vec{\omega}$ for global illumination.

5. Real-Time Translucency

In this section, we show how to exploit the image convolution derived in the previous section in order to arrive at a form that can be evaluated in real time. With equation 14, we have the ability to calculate the translucency for a given light direction $\vec{\omega}_D$ for each pixel $\vec{x}_o$ on the leaf surface. We separate the equation to obtain a *directional* part that only depends on $\vec{\omega}_D$:

$$L_t(\vec{x}_o,\vec{\omega}_o,\vec{\omega}_D) = \rho_t(\vec{x}_o,\vec{\omega}_o)L_t^D(\vec{x}_o,\vec{\omega}_D) \quad (15)$$

with

$$L_t^D(\vec{x}_o,\vec{\omega}_D) = A_p \sum_{\vec{x}_i} T(r,d(\vec{x}_i))E(\vec{x}_i,\vec{\omega}_D) \qquad (16)$$

For real-time evaluation, we want to precompute the expensive image convolution, thus requiring a new means to efficiently store and evaluate the resulting hemispherical function $L_t^D(\vec{x}_o,\vec{\omega}_D)$ for every pixel. For this purpose, we introduce the so-called *Half Life 2 basis*.

Furthermore, we make the following important simplification for real-time rendering: in principle, the function L_t is wavelength dependent, and therefore would have to be evaluated for several spectral bands, which could be done by calculating $L_t(\vec{x}_o,\vec{\omega}_D)$ with a trichromatic convolution, using the measured albedo and transmission coefficients of both sides of the leaf, and color-dependent coefficients. However, the result is far more convincing if one uses the measured average translucency, which contains both $\rho_t(\vec{x}_o)$ and the very complex spectral absorbance behavior and microstructures that appear in a backlit leaf. Therefore, we evaluate the directional part L_t^D only for *one dominant wavelength* (510nm for leaves, corresponding to green), basically capturing the effects on *luminance* effected by subsurface scattering. We add the chromatic effect using the final exiting transmittance $\rho_t(\vec{x}_o,\vec{\omega}_o)$ by substituting this quantity with the captured translucency $\rho_t(\vec{x}_o)$. In this step, we also drop the view dependence of subsurface scattering, which can be considered practically diffuse [BR01].

5.1. The Half Life 2 Basis

The Half Life 2 Basis (HL2 basis) is a little documented vector basis that was introduced in the Source engine [McT04] to combine light mapping and normal mapping by expressing the normal map in terms of the HL2 basis and evaluating the light map for the HL2 basis vectors. This works for example for radiosity, which can be interpreted as a linear function of the normal vector. It was also used to achieve normal mapping in combination with spherical harmonics-based precomputed radiance transfer [Slo06], in which case the vectors are interpreted as functions in terms of the dot product.

The HL2 basis is generated by 3 orthogonal vectors, rotated relative to the tangent coordinate system so that the angle between adjacent vectors projected on the tangent plane is isotropic, and the angle between the tangent plane and each vector is identical (Figure 8):

$$\vec{H}_1 = \left(-\tfrac{1}{\sqrt{6}}, -\tfrac{1}{\sqrt{2}}, \tfrac{1}{\sqrt{3}}\right)$$
$$\vec{H}_2 = \left(-\tfrac{1}{\sqrt{6}}, \tfrac{1}{\sqrt{2}}, \tfrac{1}{\sqrt{3}}\right)$$
$$\vec{H}_3 = \left(\sqrt{\tfrac{2}{3}}, 0, \tfrac{1}{\sqrt{3}}\right)$$

These vectors define three cosine basis functions on the hemisphere (Figure 9):

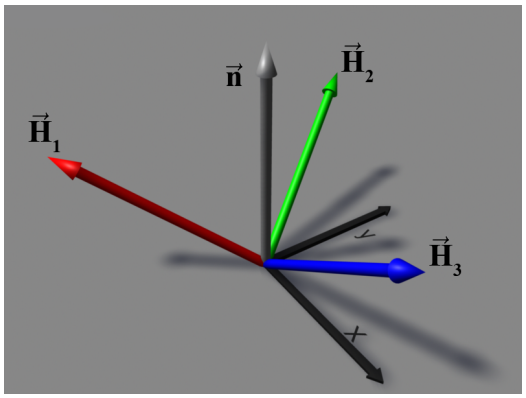

Figure 8: *The 3 vectors that define the Half Life 2 basis. The colors correspond to their coefficient channel.*

$$\mathcal{H}_i(\vec{\omega}) = \sqrt{\frac{3}{2\pi}}\vec{H}_i \cdot \vec{\omega}. \qquad (17)$$

The cosine functions do not get clamped, so all basis functions contribute over the whole hemisphere, including negative values where the dot product is negative. It can be easily verified that the HL2 basis is orthogonal with respect to hemispherical integration:

$$\int_\Omega \mathcal{H}_i(\vec{\omega})\mathcal{H}_j(\vec{\omega})d\vec{\omega} = \delta_{ij}. \qquad (18)$$

and therefore a hemispherical function $f(\vec{\omega})$ can easily be projected into the basis:

$$f(\vec{\omega}) = \sum_{i=1\ldots3} h_i\mathcal{H}_i(\vec{\omega}) \qquad (19)$$

using the basis coefficients

$$h_i = \int_\Omega f(\omega)\mathcal{H}_i(\vec{\omega})d\vec{\omega}. \qquad (20)$$

Note that even if the hemispherical function represented is always positive, one of the coefficients h_i can be negative.

The advantage of the HL2 basis is that it is very cheap to evaluate the represented function in one given hemispherical direction, which is exactly what is needed for our direct sunlight evaluation. This is in contrast to other hemispherical bases (e.g. spherical harmonics [SKS02] or wavelet bases [LSSS04]), which require the light to be transformed into the basis, and are therefore preferable only in situations where the illumination itself is also a hemispherical function and not a single direction. Since the basis has only three terms, only low frequency signals on the hemisphere can be represented, but due to the blurring properties of subsurface scattering, it is sufficient for the proposed method, which we also show numerically in Section 6.

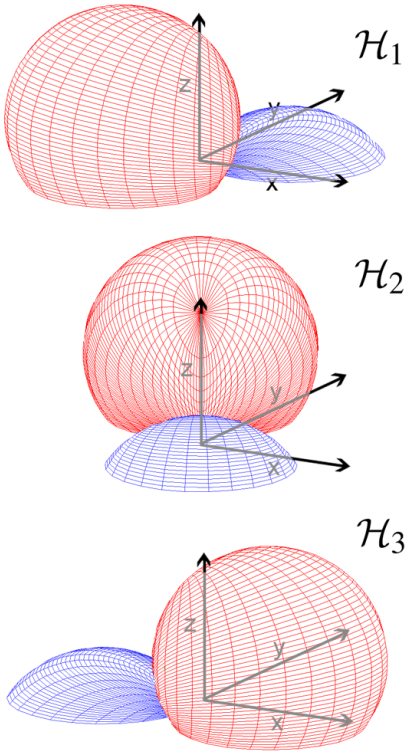

Figure 9: *The Half Life 2 basis functions. Red are positive and blue are negative values.*

5.2. Projecting Translucency into the HL2 basis

For each position $\vec{x}_o$ on the leaf, the translucency function is projected into the HL2 basis by evaluating equation (20) for $L_t^D(\vec{x}_o, \vec{\omega}_D)$. In order to evaluate the hemispherical integral, we sample it with a uniform distribution of N_L light directions over the hemisphere using Shirley's square to hemisphere mapping [SC97] to map stratified random points to a hemisphere:

$$h_i(\vec{x}_o) = \frac{2\pi}{N_L} \sum_{m=1..N_L} \mathcal{H}_i(\vec{\omega}_m) L_t^D(\vec{x}_o, \vec{\omega}_m) \qquad (21)$$

For each light direction $\vec{\omega}_m$, we perform the image convolution (16) using the acquired maps and the convolution kernel $T(r, d(\vec{x}))$. However we do not need to convolve every pixel with every pixel. As a maximum size for the kernel, we evaluate equation 8 at the thinnest part of the leaf and calculate the radius where the transmittance falls under a small threshold.

The result of the projection is a HL2 coefficient map (Figure 10) (note again that we calculate only one wavelength for scattering, see above).

It is interesting to note that the special case of perfectly

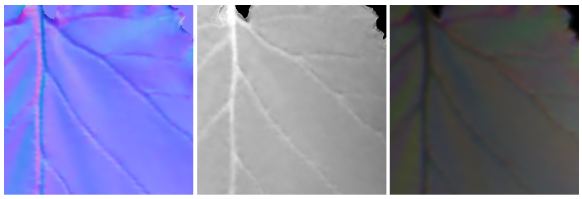

Figure 10: *The normal map, height map, and the resulting HL2 coefficient map.*

diffuse translucency $L_t^D(\vec{x}_o, \vec{\omega}_D) = \vec{n} \cdot \vec{\omega}_d$, which corresponds to the diffuse BTDF used in previous work, can be represented exactly using the coefficients $h_i = \sqrt{2\pi}/3$, since $\vec{n} = (0,0,1)$ in tangent space and therefore

$$\sum_{1..3} \frac{\sqrt{2\pi}}{3} \mathcal{H}_i(\vec{\omega}_d) = \sum_{1..3} \frac{\sqrt{3}}{3} H_i \vec{\omega}_d = (\sum_{1..3} \frac{\sqrt{3}}{3} H_i) \vec{\omega}_d = \vec{n} \cdot \vec{\omega}_d \qquad (22)$$

The coefficients can thus be said to record the deviations of the actual physically based translucency function from this special case. Note that actually any cosine function can be represented exactly in the HL2 basis.

5.3. Precomputing Visibility

The image convolution contains an evaluation of the visibility function $V(\vec{x}_i, \vec{\omega}_k)$ which is expensive even for preprocessing. In order to speed up the projection into the HL2 basis, we turn the visibility function into a lookup by precomputing the visibility term in the form of a horizon map per $\vec{x}_i$ [Max88] to capture the self shadowing of risps and bulges. We divide the horizon circle into 16 slices and record the average horizon height for each slice. At convolution time, i.e., when evaluating (16) in (21), we compare the current light direction $\vec{\omega}_m$ with the linearly interpolated horizon height and return whether $\vec{\omega}_m$ is above or below the horizon.

5.4. Rendering Translucency

At render time, we need to transform the light vector $\vec{\omega}_D$ into tangent space, look up the HL2 basis coefficients and evaluate the translucency with

$$L_{t,rec}(\vec{x}_o, \vec{\omega}_D) = L_D \rho_t(\vec{x}_o) \sum_{i=1..3} h_i(\vec{x}_o) \sqrt{\frac{3}{2\pi}} \vec{H}_i \cdot \vec{\omega}_D, \qquad (23)$$

resulting in only two texture lookups (ρ_t and h_i are each stored in an RGB texture) and 3 added and weighted dot products, which allows calculating accurate and dynamic translucency for a large number of leaves on a per-pixel basis.

This computation is extremely simple and memory efficient, and can be integrated into any modern rendering pipeline using pixel shaders. Since we consider only interactions local to a single leaf, the leaves can be instanced and

also animated, which allows many trees with large numbers of leaves to be animated and rendered interactively.

6. Implementation, Results and Discussion

We have implemented and tested the proposed leaf model and applied it to two leaf datasets acquired with the described acquisition setup. Subsurface light transport L_t^D makes use of physical parameters (averages taken from [MIPK90] and [Woo71]) which are summarized in Table 1, together with parameters for the precomputation. We evaluate L_t^D at a wavelength of 510 nm, so the parameters are chosen for that wavelength.

Mean cosine	g	0.07
Scattering coeff.	σ_s	10.2 $1/mm$
Absorption coeff.	σ_a	0.4 $1/mm$
Refraction index ratio	η	1.33
Multi dipoles	n	3
Number of light directions	N_L	128

Table 1: *Parameters used for the calculation and precomputation of subsurface scattering.*

First we show in Figure 11 the comparison of the proposed translucency method with the current state-of-the-art using a diffuse BTDF (from our measured data) and the geometric normal $\vec{n}_g$, giving the following expression:

$$L_t^{BTDF}(\vec{x}_0, \omega_D) = -L_D \rho_t(\vec{x}_0)(\vec{n}_g \cdot \vec{\omega}_D). \qquad (24)$$

The figure shows the variations in the translucent radiance in dependence of the incoming light direction. For example, at steep angles, the structural features are smudged out, whereas at grazing angles the high frequency structures of leaves are observable. Scattering effects in the leaf due to variations in the light-facing surface, e.g., self shadowing, thickness and reflectance variations, are accurately modeled. In contrast, the standard model appears flat and responds to changes in light direction only through the cosine term.

Second, we analyze the accuracy with which the Half Life 2 Basis is able to approximate an exact evaluation of the multi-dipole model [DJ05]. For this, we calculate the relative signal reconstruction error $L_{t,rec}/L_t - 1$ for all texels in the data set seen in 5 for 3 light angles ($\pi/2$, $\pi/4$, $\pi/8$) relative to the horizon, and plot the resulting error distributions over the leaf as histograms in Figure 12. Figure 13 shows a direct comparison at a light angle of $\pi/8$. The average reconstruction error is only 3%, which shows that the Half Life 2 Basis is well suited to represent translucency. Also note that a leaf only transmits about 5-20% [ML51] of the incident light, an effect which is accurately modeled with the proposed method. If the leaf is rendered without tone mapping, the HL2 basis coefficients can be scaled so that translucent leaves don't appear too dark.

Finally, we discuss the implementation and performance

Figure 11: *Physically based leaf translucency (top) with light at different angles from steep (left) to grazing angles (right) in comparison to the standard diffuse translucency model (bottom).*

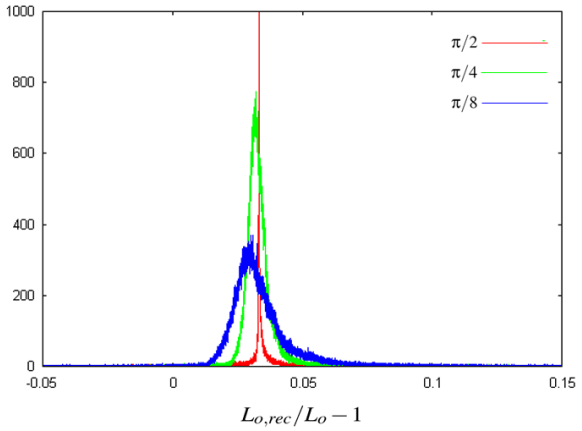

Figure 12: *Error histogram of the relative signal reconstruction error for 3 light directions.*

of the method. All results were generated on a Pentium4 3.2GHz with an NVIDIA 8800 GTX graphics card running DirectX 10. Starting from the acquired datasets, the precomputation to generate the HL2 coefficient map takes about 20s per image convolution, giving about 45 minutes for a complete leaf dataset. At runtime, the evaluation of the translucency term involves two texture lookups into the translucency and HL2 coefficient map, and a few arithmetic instructions. The cost of this is significantly lower than the rest of the shading model, for example the arithmetic required to evaluate the Cook-Torrance shading model for the light-

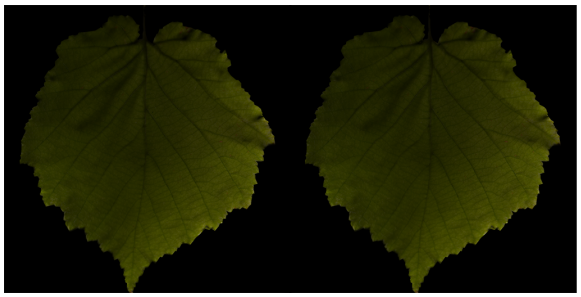

Figure 13: *Leaf shaded with multi-dipole model L_t (left) and the reconstructed translucency $L_{t,rec}$ (right) at a light angle of $\pi/8$.*

facing surfaces. The DirectX 10 shader compiler reports approximately 119 instruction slots for the complete shader, of which the translucency part takes about 15 instructions. An arbitrary number of leaves can be rendered using this dataset. The complete shader currently uses 8 1024x1024 RGB textures: albedo, translucency, normal map and HL2 map, both for front and back sides of the leaf, resulting in 24MB of texture memory for a single leaf dataset.

We have measured the performance in an interactive application showing a tree with 6,392 leaves rendered as quad patches with normal maps. Each leaf is dynamically shaded. The light-facing leaf sides use the Cook-Torrance shading model discussed in Section 3, whereas the opposite sides are shaded using our physically based translucency model. To account for the high dynamic range of the scene, we apply Reinhard's tone mapping algorithm [RSSF02] and simple blooming in a post-process. Shadows are generated using a 2048x2048 shadow map which is evaluated with a percentage closer filter (PCF) [RSC87] with 6 Poisson-distributed samples. Since we do not have animation data for the tree available, we decided to precompute a simple indirect lighting term using ambient occlusion. Figures 1 and 14 (color plate) show screenshots of the resulting rendering quality for the whole tree, however a better impression can be obtained from the accompanying video. The frame rate for a fly through of the tree varies from 66 frames per second (fps) for closeup views where the whole screen is covered with fragments, to 116 fps for more distant views, including all mentioned effects.

The datasets for the leaf models, consisting of albedo, translucency, normal, thickness and HL2 basis coefficient maps and an example shader are available at [Hab07].

7. Conclusions and Future Work

We have presented the first physically based translucency model for the real-time rendering of plant leaves. Our model takes full subsurface scattering and the full structure of a leaf into account, including effects like self shadowing, spatially varying reflectance and thickness into the subsurface scattering simulation. For the light-facing sides of leaves, we use a state-of-the art shading model augmented by high-resolution normal maps, which provide improved specular reflections. All required input maps are created through a novel acquisition process that gives <1mm resolution on a leaf.

Our translucency model is based on the multiple dipole approximation to thin slab subsurface scattering introduced by Donner and Jensen [DJ05]. We reformulated the translucent contribution as an image convolution process and projected it into the Half Life 2 basis for efficient rendering. For the future, we would like to automate the proposed data acquisition method to make it possible to generate a database of leaves suitable for both real-time and physically based rendering.

In the future, it would be interesting to validate the model using a full Monte Carlo solution. This could also be used to compare our one-layer approximation to a model with two or more layers for a leaf, as proposed by Donner and Jensen [DJ05].

8. Acknowledgements

This research was funded by the Austrian Science Fund (FWF) under contract no. P17261-N04. The authors would like to thank Oliver Mattausch and Thomas Gamper for very helpful discussions and comments. The camera has been supplied by the Hochschuljubiläumsstiftung der Stadt Wien, Project H-1085/2004.

References

[BL05] BORSHUKOV G., LEWIS J. P.: Realistic human face rendering for "the matrix reloaded". In *SIGGRAPH '05: ACM SIGGRAPH 2005 Courses* (New York, NY, USA, 2005), ACM Press, p. 13.

[BLJM05] BOUSQUET L., LACHERADE S., JACQUEMOUD S., MOYA I.: Leaf BRDF measurements and model for specular and diffuse components differentiation. *Remote Sensing of Environment 98* (2005), 201Ű–211.

[BR97] BARANOSKI G., ROKNE J.: An Algorithmic Reflectance and Transmittance Model for Plant Tissue. *Computer Graphics Forum 16*, 3 (1997).

[BR01] BARANOSKI G., ROKNE J.: Efficiently simulating scattering of light by leaves. *The Visual Computer 17*, 8 (2001), 491–505.

[COM98] COHEN J., OLANO M., MANOCHA D.: Appearance-preserving simplification. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1998), ACM Press, pp. 115–122.

[CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Trans. Graph. 1*, 1 (1982), 7–24.

[DJ05] DONNER C., JENSEN H. W.: Light diffusion in multi-layered translucent materials. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers* (New York, NY, USA, 2005), ACM Press, pp. 1032–1039.

[DvGNK99] DANA K. J., VAN GINNEKEN B., NAYAR S. K., KOENDERINK J. J.: Reflectance and texture of real-world surfaces. *ACM Trans. Graph. 18*, 1 (1999), 1–34.

[FD03] FRANZKE O., DEUSSEN O.: Rendering plant leaves faithfully. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Sketches & Applications* (New York, NY, USA, 2003), ACM Press, pp. 1–1.

[FP92] FARRELL T. J., PATTERSON M. S.: A diffusion theory model of spatially resolved, steady-state diffuse reflections for the noninvasive determination of tissue optical properties in vivo. *Med. Phys. 19* (1992), 879–888.

[GJH*98] GANAPOL B., JOHNSON L., HAMMER P., HLAVKA C., PETERSON D.: LEAFMOD: A new within-leaf radiative transfer model. *Remote Sensing of Environment 63* (1998), 182–193.

[GVU96] GOVAERTS Y., VERSTRAETE S. J. M., USTIN S.: Threedimensional radiation transfer modeling in a dycotyledon leaf. *Applied Optics 35*, 33 (1996), 6585Ű–6598.

[Hab07] http://www.cg.tuwien.ac.at/research/publications/2007/Habel_2007_RTT/, 2007.

[HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1993), ACM Press, pp. 165–174.

[JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), ACM Press, pp. 511–518.

[LSSS04] LIU X., SLOAN P.-P., SHUM H.-Y., SNYDER J.: All-Frequency Precomputed Radiance Transfer for Glossy Objects. *Proceedings Eurographics Symposium on Rendering 15* (2004), 337–344.

[Max88] MAX N. L.: Horizon mapping: shadows for bump-mapped surfaces. *The Visual Computer 4* (1988), 109Ű–117.

[McT04] MCTAGGERT G.: *Half-Life 2/Valve Source Shading*. Tech. rep., Valve Corporation, 2004.

[MIPK90] MA Q., ISHIMARU A., PHU P., KUGA Y.: Transmission, Reflection, and Depolariastion of an OPtical Wave For a Single Leaf. *IEEE Transactions on Geoscience and Remote Sensing 28*, 5 (1990), 865–872.

[MKB*03] MERTENS T., KAUTZ J., BEKAERT P., REETH F. V., SEIDEL H.-P.: Efficient rendering of local subsurface scattering. In *PG '03: Proceedings of the 11th Pacific Conference on Computer Graphics and Applications* (Washington, DC, USA, 2003), IEEE Computer Society, p. 51.

[ML51] MOSS R. A., LOOMIS W. E.: A Low Distortion Map Between Disk and Square. *Journal of the Iowa Agricultural Experiment Station J*, 2017 (1951), 370–391.

[MMS*05] MÜLLER G., MESETH J., SATTLER M., R. S., KLEIN R.: Acquisition, Synthesis, and Rendering of Bidirectional Texture Functions. *Computer Graphics Forum 24*, 1 (2005), 83Ű–109.

[NRH*77] NICODEMUS F. E., RICHMOND J. C., HSIA J. J., GINSBERG I. W., LIMPERIS T.: *Geometrical considerations and nomenclature for reflectance*. Jones and Bartlett Publishers, Inc., USA, 1977.

[PCW89] PATTERSON M. S., CHANCE B., WILSON B. C.: Time resolved reflectance and transmittance for the noninvasive measurement of tissue optical properties. *Appl. Opt. 28* (June 1989), 2331–+.

[RSC87] REEVES W. T., SALESIN D. H., COOK R. L.: Rendering antialiased shadows with depth maps. In *SIGGRAPH '87 Proceedings* (New York, NY, USA, 1987), ACM Press, pp. 283–291.

[RSSF02] REINHARD E., STARK M., SHIRLEY P., FERWERDA J.: Photographic tone reproduction for digital images. *ACM Trans. Graph. 21*, 3 (2002), 267–276.

[SC97] SHIRLEY P., CHIU K.: A Low Distortion Map Between Disk and Square. *Journal of Graphics Tools 2*, 3 (1997), 45–52.

[SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *SIGGRAPH '02 Proceedings* (New York, NY, USA, 2002), ACM Press, pp. 527–536.

[Slo06] SLOAN P.-P.: Normal mapping for precomputed radiance transfer. In *SI3D '06: Proceedings of the 2006 symposium on Interactive 3D graphics and games* (New York, NY, USA, 2006), ACM Press, pp. 23–26.

[Woo71] WOOLLEY J.: Reflectance and Transmittance of Light by Leaves. *Plant Physiol. 47* (1971), 656–662.

[WTL05] WANG R., TRAN J., LUEBKE D.: All-frequency interactive relighting of translucent objects with single and multiple scattering. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers* (New York, NY, USA, 2005), ACM Press, pp. 1202–1207.

[WWD*05] WANG L., WANG W., DORSEY J., YANG X., GUO B., SHUM H.-Y.: Real-time rendering of plant leaves. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers* (New York, NY, USA, 2005), ACM Press, pp. 712–719.

Figure 14: *A tree featuring physically based translucency (top left). The top right picture shows a leaf with both reflectant and translucent parts. A direct comparison of the standard method (bottom left) to the proposed method (bottom right) shows improved structure due to normal mapping, and the greatly improved appearance of translucency in the leaves.*