4-2020

# Randomizable phenology-dependent corn canopy for simulated remote sensing of agricultural scenes

M Grady Saunders
ms3103@rit.edu

# Randomizable phenology-dependent corn canopy for simulated remote sensing of agricultural scenes

by

M. Grady Saunders

B. S. East Tennessee State University, 2017

A thesis submitted in partial fulfillment of the
requirements for the degree of Master of Science
in the Chester F. Carlson Center for Imaging Science
College of Science
Rochester Institute of Technology

April, 2020

Signature of the Author _____

Accepted by _____
Coordinator, M.S. Degree Program Date

CHESTER F. CARLSON CENTER FOR IMAGING SCIENCE
COLLEGE OF SCIENCE
ROCHESTER INSTITUTE OF TECHNOLOGY
ROCHESTER, NEW YORK

CERTIFICATE OF APPROVAL

---

M.S. DEGREE THESIS

---

The M.S. Degree Thesis of M. Grady Saunders
has been examined and approved by the
thesis committee as satisfactory for the
thesis required for the
M.S. degree in Imaging Science

Dr. Jan van Aardt, Thesis Advisor

Dr. John Kerekes

Dr. Scott Brown

Dr. Quirine Ketterings

Dr. David Ross

Date

# Contents

# Randomizable phenology-dependent corn canopy for simulated remote sensing of agricultural scenes

by

M. Grady Saunders

Submitted to the
Chester F. Carlson Center for Imaging Science
in partial fulfillment of the requirements
for the Master of Science Degree
at the Rochester Institute of Technology

## Abstract

Crop health assessment and yield prediction from multi-spectral remote sensing imagery are ongoing areas of interest in precision agriculture. It is in these contexts that simulation-based techniques are useful to investigate system parameters, perform preliminary experiments, etc., because remote sensing systems can be prohibitively expensive to design, deploy, and operate. However, such techniques require realistic and reliable models of the real world. We thus present a randomizable time-dependent model of corn (*Zea mays L.*) canopy, which is suitable for procedural generation of high-fidelity virtual corn fields at any time in the vegetative growth phase, with application to simulated remote sensing of agricultural scenes. This model unifies a physiological description of corn growth subject to environmental factors with a parametric description of corn canopy geometry, and prioritizes computational efficiency in the context of ray tracing for light transport simulation. We provide a reference implementation in C++, which includes a software plug-in for the 5$^{\text{th}}$ edition of the Digital Imaging and Remote Sensing Image Generation tool (DIRSIG5), in order to make simulation of agricultural scenes more readily accessible. For validation, we use our DIRSIG5 plug-in to simulate multi-spectral images of virtual corn plots that correspond to those of a United States Department of Agriculture (USDA) site at the Beltsville Agricultural Research Center (BARC), where reference data were collected in the summer of 2018. We show in particular that 1) the canopy geometry as a function of time is in agreement with field measurements, and 2) the radiance predicted by a DIRSIG5 simulation of the virtual corn plots is in agreement with

radiance-calibrated imagery collected by a drone-mounted MicaSense RedEdge imaging system. We lastly remark that DIRSIG5 is able to simulate imagery directly as digital counts provided detailed knowledge of the detector array, e.g., quantum efficiency, read noise, and well capacity. That being the case, it is feasible to investigate the parameter space of a remote sensing system via "end-to-end" simulation.

# Chapter 1

# Introduction

## 1.1. Introduction

Crop monitoring via passive remote sensing with multispectral sensors and imaging spectrometers, mounted on Unmanned Aerial Systems (UAS) and Earth-orbiting satellites, is fast becoming a critical part of general crop management and yield prediction [1]. That being the case, an important research objective is to quantify the spectral relationship between the continually changing environment and the phenology of commodity crops, and to improve remote sensing technology in pursuit of more robust systems for crop monitoring and assessment of crop ecology. However, remote sensing systems can be prohibitively expensive to design, deploy, and operate. Simulation-based techniques therefore are useful to investigate system parameters, perform preliminary experiments, and pre-visualize data products, among other things. Of course, such techniques require models of the real world which are realistic and reliable (provide high structural and spectral fidelity), as well as accessible (user-friendly and computationally efficient).

Corn (*Zea mays L.*) is perhaps the most ubiquitous commodity crop. It is the primary ingredient in myriad food and industrial products, as well as the primary feed grain [2], [3]. Approximately 90 million acres of corn are projected to be planted, with 15 billion bushels projected to be harvested, in the United States in 2020 [4]. We thus present a randomizable time-dependent model of corn canopy, which is suitable for procedural generation of high-fidelity virtual corn fields at any time in the vegetative growth phase, with application to simulated remote sensing of agricultural scenes. That is, the model should be sufficient to simulate physically-accurate multi-spectral imagery of

corn fields (from drone-level to satellite-level distances) at arbitrary vegetative growth stages, where quantities of interest (e.g., plant height, leaf area, leaf chlorophyll content, leaf water content, etc.) are known by construction. Physically-accurate simulated images with known ground-truth quantities should be useful to optimize the parameters of a hypothetical imaging system, as well as to provide a source of data for machine learning detection and classification algorithms. To that end, we provide a freely available C++ implementation of the model. This implementation includes a software plug-in to the 5$^{\text{th}}$ edition of the Digital Imaging and Remote Sensing Image Generation tool (DIRSIG5) [5], described in greater detail in Section 1.3, with the hope that this plug-in will make simulated remote sensing of agricultural scenes more accessible to the research community. We used this plug-in to simulate multi-spectral images of virtually reconstructed corn plots corresponding to those from a USDA site at the Beltsville Agricultural Research Center (BARC), where reference data were collected in summer 2018. We subsequently used the real multispectral imagery to validate the efficacy of the virtual simulation-based model.

## 1.2. Related work

Plant growth modeling is varied in the literature, in the sense that motivations, techniques, inputs, and outputs may differ drastically from one model to the next. Many crop growth models in particular are *agronomical*, i.e., motivated by interest in optimizing soil management and crop production. Agronomical models aim to be comprehensive in their consideration of crop ecology, phenology, and physiology. Such models simulate the shoot and root systems of crops through the end of the reproductive stage, accounting for eco-physiological processes via day-to-day integration of abstract "rate" and "state" variables. Modeled processes include photosynthesis, assimilation (the conversion of $CO_2$ to organic material), respiration (the conversion of soil nutrients into energy for growth), and transpiration (the process by which water moves through and evaporates from plant tissue).

van Diepen et al. [6] introduced WOFOST, the WOrld FOod STudies model of annual field crops. WOFOST represents a crop variety as a generic parameter set (including, e.g., optimum day length, maximum rates of respiration and transpiration, maximum nutrient concentrations in leaves). As such, it is able to simulate the growth and development of different types of crops as well as different varieties for each type. WOFOST has been improved, expanded, and used to great effect

in the 25 years since its introduction [7]. Jones et al. [8] introduced CERES-Maize, an agronomical cropping model which predates WOFOST and focuses specifically on the growth and development of corn varieties. CERES-Maize thus exposes more genotype-specific parameters, and features more explicit (though still abstract) process modeling. Yang et al. introduced the Hybrid-Maize model [9]–[11], which combines the corn-specific formulation of CERES-Maize with the mechanistic formulation of WOFOST and other generic cropping models. While useful for soil management analysis and yield forecasting, none of these models prioritize high-fidelity canopy geometry and/or the associated optical properties as outputs, and thus are not immediately applicable to light-transport simulation. Fortunately, there is another category of growth models which we may call *graphical* or *architectural*, i.e., motivated by interest in graphics applications, visualization, or otherwise reproducing physical plant architecture.

There is in fact much interest in the computer graphics research community related to the geometrical structure and optical properties of plants. Of particular interest is the fractal arrangement of many plant varieties. There are two competing viewpoints in terms of understanding the mechanisms that produce such fractal arrangements. On one hand is the notion that plants are *self-similar*, as proposed by Honda [12]. This is to say that plant structure is inherently recursive and replicates with well-defined rules at each level. This has motivated explicitly recursive algorithms [13], [14]. The opposing viewpoint considers that plants are *self-organizing*, as is traceable to Ulam [15]. This is to say that plant structure is owed to competition between branch modules for space and resources, and not necessarily only to inherent recursion rules. Pałubicki et al. [16] presented a method for generating realistic trees with self-similar, but also self-organizing branching structures, remarking that trees do exhibit recursive, self-similar branching, but the genetic mechanism that controls branching is not distinct enough between tree types to be entirely responsible for the emergent structure [17].

The widely adopted approach to actually represent plant structure is to use an *L-system* [18], [19]. An L-system is a formal grammar used to "grow" a hierarchical string of symbols by repeatedly applying a set of rewriting rules in parallel to an initial string of symbols. The first L-system was introduced by Lindenmayer [20] (after whom it is named) to model cellular interaction over discrete time steps, given a set of rules that constrain how cells may interact at each time step. In the context of plant modeling, this idea is adapted such that the string of symbols encodes the

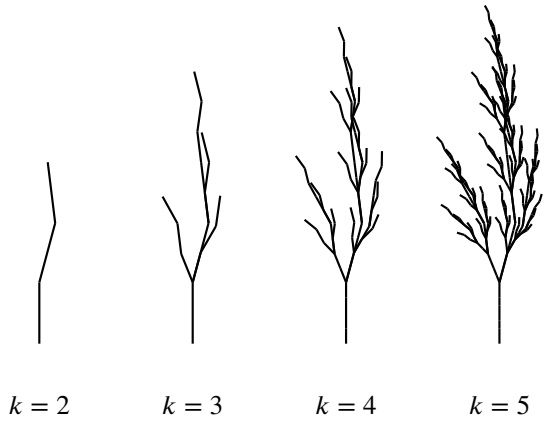


**Figure 1.1:** Iterations $k = 2$ through $k = 5$ of a basic L-system modeling a tree-like structure. To make the structural changes between iterations more visually comparable, each iteration is scaled by $1/2^k$.

hierarchical organization of branching structures in a plant, and the set of rewriting rules encodes how the branching structures may grow at each time step. An example is shown Figure 1.1. Here, the initial string at iteration $k = 1$ is just $X$. Subsequent strings for iterations $k = 2, 3, 4, ...$ are generated by applying two rewriting rules,

$$X \to F+[[X]-X]-F[+FX]+X,$$
$$F \to FF.$$

We use a "turtle algorithm" to generate the drawing for each $k$ in the figure. That is, we consider a turtle starting at $(0, 0)$ facing north. We then look at each symbol in the $k$th string as an instruction for the turtle. If the symbol is $X$, the turtle does nothing. If the symbol is $F$, the turtle draws a line from its current position to a unit ahead in the direction it is facing. If the symbol is $+$, it rotates by $22°$. If the symbol is $-$, it rotates by $-37°$. If the symbol is [, it remembers its current position and direction. If the symbol is ], it goes back to the position and direction it most recently remembered (without drawing a line). This L-system is simple and ad hoc. By attributing quantities of interest to the symbols (length, biomass, growth potential, etc.), it is possible to engineer an L-system which emulates more complicated biophysical interactions.

Fournier and Andrieu [21] modeled the growth and development of corn canopy architecture using L-systems. In their model, the authors use the L-system representation of a corn plant

to produce associated canopy geometry, which is then used to compute the microclimate, i.e., temperature and Photosynthetically Active Radiation (PAR), at each leaf. The microclimate is used in turn to inform how the L-system is updated in the next time step. This cycle repeats to simulate the growth of the plant. Drouet and Pagès [22] presented GRAAL, a model of the GRowth, Architecture, and carbon ALlocation in corn plants which also uses explicit canopy geometry to simulate microclimate. As its name may imply, GRAAL pays even greater attention to agronomical quantities of interest such as carbon allocation and biomass. We may think of such "microclimate models" as being in part agronomical, in part architectural, but not exactly graphical. That is to say that the motivation is predominantly agronomical, such that interest in plant architecture and modeling of canopy geometry serves only to improve modeling of agronomical processes. That being the case, the canopy geometry is not necessarily intended for graphics applications or light transport simulation.

España et al. [23] presented a more extensive parameterization of corn canopy geometry intended for application to reflectance simulation, which built upon the work of Prévot et al. [24]. This parameterization represents the lengths, positions, and orientations of leaves and internodes as a function of three primary variables, being 1) the final number of leaves, 2) the final height, and 3) the cumulative leaf area. This parameterization also includes a plethora of secondary variables, e.g., per-leaf insertion angle, curvature, and undulation, all of which affect the specific appearance of each plant. By randomizing these secondary variables according to empirically deduced statistical distributions, the authors were able to generate an arbitrary number of visually unique plants with known primary characteristics. Such approaches, however, almost exclusively dealt with 3D plant geometry, with little attention paid to spectral or optical properties, especially those beyond the visible electromagnetic spectrum (400–700nm).

Regarding the simulation of leaf optical properties spectra, Jacquemoud and Baret [25] most notably introduced a model called PROSPECT. The model considers the interaction of diffuse light with a stack of $N$ dielectric plates separated by $N-1$ empty spaces to model leaf scattering. While this construction implies $N$ to be any positive integer, the result of the derivation permits $N$ to be any positive real number. The parameter $N$ is known as the *structure parameter*, representing in some capacity the internal organization of the leaf. For a so-called compact leaf, e.g., a corn leaf, $N \approx 1$. The model accepts other notable biophysical parameters, such as the refractive


**Figure 1.2:** A DIRSIG5-simulated image of the Port of Tacoma, as captured by a low-earth orbiting pan-chromatic framing array.

index of leaf material and the concentrations and specific absorptance spectra of pigments and water in the leaf. Through its inversion, PROSPECT is most often used to retrieve biophysical parameters from measured spectra [26]–[29], though it has also been used directly in scattering models for computer graphics [30], [31]. The PROSPECT+SAIL variety of models [29] extends the individual-leaf optical properties spectra predicted by PROSPECT to directional scattering distributions of entire canopies. The objective of this modeling is to deduce biophysical parameters of canopies via inversion/fitting to bidirectional reflectance measurements, such that there is no explicit representation of leaf geometry. As such, the PROSPECT+SAIL modeling approach in particular is not applicable within this work, where explicit geometry is a central component. Next, it is helpful to introduce DIRSIG5 in reasonable detail to establish the overall context of this simulation-based research effort.

## 1.3. The DIRSIG5 simulation model

DIRSIG5 is the $5^{\text{th}}$ and most recent edition of the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model [5], which has been developed and validated by the Digital Imaging and Remote Sensing (DIRS) laboratory at the Rochester Institute of Technology over the past two decades. The DIRSIG model performs physics-based light transport simulation in order to provide radiometrically-accurate solutions to light transport problems of interest to remote sensing

applications. Moreover, the model is able to generate passive broadband, multi-spectral, and hyper-spectral imagery [32]–[36], as well as response signals from Light Detection And Ranging (LiDAR) [37]–[39] and Synthetic Aperture Radar (SAR) [40]. DIRSIG numerically integrates the radiance distribution in a virtual scene over the area of a virtual detector element by explicitly tracing light rays. Therefore, the geometric structure of the scene being simulated, e.g., a landscape, city, or forest canopy, must be modeled explicity by 3D primitive shapes. The most frequently used primitive is the *triangle*. That is to say, a virtual scene is usually facetized as a collection of (potentially millions or more) triangles, though other primitive shapes, e.g., disks, cylinders, and spheres, may also be used. Each primitive shape is associated with a set of material properties. Such properties include, most importantly, the Bidirectional Reflectance Distribution Function (BRDF) and the Bidirectional Transmittance Distribution Function (BTDF) of the material, which characterize how an incident light ray is reflected, transmitted, and absorbed. Material properties may be different for every primitive, and may even vary over a single primitive using a feature called *material mapping*. An example of a high-fidelity DIRSIG5 scene is shown in Figure 1.2. This scene represents the Port of Tacoma in Tacoma, Washington. It contains cars, crates, streetlights, buildings, bar targets, and even water, among other things. Every object in the scene is constructed in the way just described, i.e., as a collection of geometric primitives with different material properties.

While the Tacoma scene shown in Figure 1.2 is an excellent showcase of what DIRSIG5 is capable of, it is imporant to note that *creating a high-fidelity scene is not easy*. Doing so requires a lot of time, effort, and data. Typically, the geometry that represents an object is modeled by hand in a 3D content-authoring application, e.g., Blender (`https://blender.org`), or is obtained from 3D scan data. Proper characterization of an object's material properties requires either measured goniometer data or hand-tuning of a parametric scattering function to obtain the desired scattering shape, and often further requires authoring of *material mixture maps* to model textured surfaces.

## 1.4. Objectives

The objectives of this work are to simulate the growth and development of a virtual corn field, and to assess the fidelity of such a virtual corn field in accurately reproducing spectral response, as measured by a UAS multi-spectral agricultural imaging system, via physics-based light-transport simulation. It is also important to note that an outcome of this work, namely the DIRSIG5

software plug-in, is meant to automate aspects of DIRSIG scene creation, specifically regarding the generation of virtual corn canopy geometry and associated optical properties, in order to make it easier for users of the software to assemble and simulate realistic agricultural scenes. The overview of the simulation approach is described in Section 2.1. The specific simulation methods implemented for the modeling of corn canopy geometry, physiological growth and development, and associated optical properties are described in Sections 2.2, 2.3, and 2.4, respectively. The qualitative and quantitative validation of virtual corn models in reproducing spectral response is described and discussed in Section 2.5. Closing remarks and suggested directions for future work are given in Sections 3.1 and 3.2, respectively.

# Chapter 2

# Structural and spectral modeling of corn canopy

## 2.1. Overview

The primary contribution of this work is, in brief, to develop and validate a high-fidelity, time-dependent geometric parameterization of corn canopy that is suitable for light-transport simulation by ray tracing, where the time dependence is governed by a physiological model of growth and development. That is, we model a corn plant as a collection of parametric surfaces whose primary physical characteristics (e.g., surface area) depend on *physiological age*, and whose secondary physical characteristics (e.g., leaf angle) are randomizable within plausible distributions.

There are four main points that set our model apart from previous models in the literature. 1) The geometric parameterization more rigorously considers, and thus more faithfully reproduces, lengths and areas of leaves. 2) The geometric parameterization introduces additional features, including leaf curl and explicit sheath modeling, that improve visual fidelity. 3) The growth model operates in continuous time. This is to say that plant geometry is efficiently and smoothly evaluable as a function of time. We mean by "efficiently evaluable" that the algorithmic complexity is constant $O(1)$ with respect to time. We mean by "smoothly evaluable" that time is a real number, not an integer corresponding to a discrete time step. As this may suggest, the model does not rely on a typical L-system representation of plant architecture. We note that corn does not exhibit deeply self-similar branching, and therefore we propose that corn does not necessarily warrant representation

as an L-system. 4) The growth model translates expressions from previous models into a unified "physiological age space". This collapses consideration of environmental factors to their impact on a single equation and a handful of parameters. While this is a simplification, it is not without merit, and it permits the model to be more general, extensible, and implementable for the task of generating spectrally and structurally plausible corn canopy geometry.

We have constructed the model to be suitable for simulated remote sensing of agricultural scenes. This requires that the model be able to produce visually high-fidelity geometry, to capture in-field spectral and structural variability with some consideration of environmental conditions, and to operate efficiently for various viewing distances. For example, it should be possible to simulate imagery captured by a drone flying less than 100 meters above the ground, with a ground sample distance (GSD) on the order of centimeters, as well as imagery captured by a satellite in low-Earth orbit, with a GSD on the order of tens of meters or greater. The vast number of corn plants potentially visible at satellite-scale distances is problematic, and coping with this has greatly influenced the design of the model. It is impractical to simulate the nutrient uptake and microclimate for millions of corn plants as part of the growth model, where neighboring plants shadow each other, compete for resources, etc. This level of simulation would further necessitate an impractical amount of input data, consisting of spatial maps of all incorporated environmental variables, some or many of which may not even be relevant to an end user's specific scenario. Even so, the model should be agronomically informed, and it should be possible for an end user to reproduce the effects of biophysical interactions of interest through exposed parameters in the model. Stated differently, the model is designed to be able to reproduce relevant biophysical phenomena, but not necessarily to predict these phenomena from first principles.

We describe the geometric parameterization of the relevant plant surfaces, namely the lamina, sheath, and internode, in section 2.2, building on the parameterization of España et al. [23]. We describe the physiological growth model in section 2.3, where we derive piecewise linear equations for growth in "physiological age space"from the work of Fournier and Andrieu [21] and Drouet and Pagès [22]. In Section 2.4, we describe our approach for modeling leaf optical properties, which includes uniting the ordinary PROPSECT model [25] with modern, physically-based microsurface scattering functions [41], [42]. Finally, in Section 2.5, we validate the efficacy of the virtual corn model in reproducing spectral response by comparing simulated imagery to multi-spectral UAS

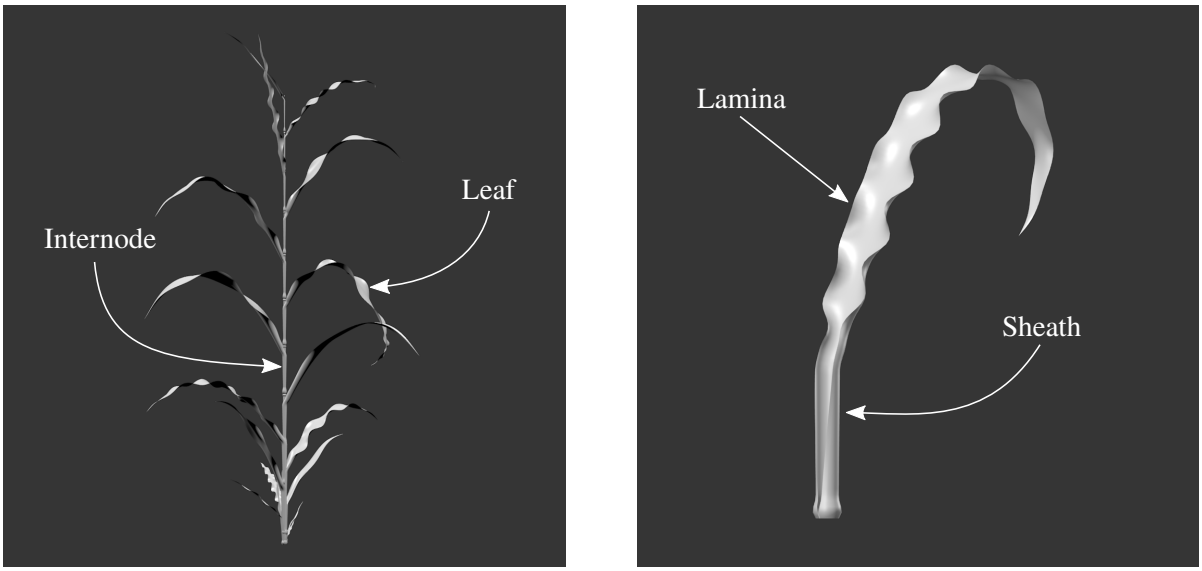


**Figure 2.1:** Two images labeling the relevant parts of the vegetative structure of a corn plant. On the left, the distinction between the internode (stalk) and leaf. On the right, the lamina (leaf blade) and sheath (connective tissue).

imagery captured by the MicaSense RedEdge-3.

## 2.2. Methods: Geometric parameterization

The vegetative structure of a corn plant consists of two major parts: leaves and internodes. Each leaf is logically divisible into its *lamina* and its *sheath*—the lamina is the leaf blade itself, while the sheath is the protective tissue joining the leaf blade to an internode. An internode is a segment of the stalk of the plant which connects a pair of leaves (see Figure 2.1). In this section, we construct a high-fidelity parametric expression for each of these surfaces, such that each has an exact length, and the lamina in particular has a "nearly exact" area, which is properly exact in the absence of curl and undulation deformations. We have referred most extensively to the work of España et al. [23] and Prévot et al. [24]. The following discussion focuses on the aspects of our parameterization that differ in this work. We omit simpler or less important mathematical expressions in favor of descriptions in plain English where it is sensible to do so. Such expressions and other implementation details are documented as part of the implementation itself.

### 2.2.1. Forming the midrib curve

We form the midrib (the middle rib in the center of the lamina) as a piecewise curve in the XY plane, analogous to España et al., with the exception that we explicitly parameterize by arc-length, which requires that we derive and invert the arc-length integral for each of the two piecewise segments. It turns out that neither arc-length integral is invertible in closed form, so we determine an appropriate initial guess for numerical inversion by Newton-Rhaphson iteration. This additional work ultimately guarantees that the midrib has an exact length $L$, and the magnitude of the derivative of the midrib is constant and in fact equivalent to $L$. This information is useful later to show that the lamina has an exact area in the absence of deformations. Let $\boldsymbol{P}_{\mathrm{m}} : [0,1] \to \mathbb{R}^2$ parameterize the midrib curve, again in the XY plane. Then,

$$\boldsymbol{P}_{\mathrm{m}}(t) = \begin{cases} \boldsymbol{P}_{\mathrm{m0}}(t) & t \le t_{\mathrm{ms}}, \\ \boldsymbol{P}_{\mathrm{m1}}(t) & t > t_{\mathrm{ms}}, \end{cases} \tag{2.1}$$

where $\boldsymbol{P}_{\mathrm{m0}}$ is quadratic, $\boldsymbol{P}_{\mathrm{m1}}$ is elliptic, and $t_{\mathrm{ms}} \in (0,1]$ is the parametric split location.

*Quadratic segment.* As a function of $x$ instead of parameter $t$, the quadratic segment is of the form

$$\boldsymbol{P}_{\mathrm{m0}}(x) = \begin{bmatrix} x & ax^2 + bx \end{bmatrix}^\top \tag{2.2}$$

where $a$ is the curvature and $b$ is the slope at the origin. The arc-length integral $L_{\mathrm{m0}}$ is obtainable by change of variables, such that

$$L_{\mathrm{m0}}(x) = \frac{1}{2a}\left[\hat{L}_{\mathrm{m0}}(2ax+b) - \hat{L}_{\mathrm{m0}}(b)\right] \tag{2.3}$$

$$\text{where } \hat{L}_{\mathrm{m0}}(x) = \frac{1}{2}x\sqrt{1+x^2} + \frac{1}{2}\operatorname{arcsinh} x. \tag{2.4}$$

We must evaluate Equation 2.2 at $x = L_{\mathrm{m0}}^{-1}(\ell)$ to parameterize $\boldsymbol{P}_{\mathrm{m0}}$ by arclength $\ell$. By manipulating Equation 2.3, we have

$$L_{\mathrm{m0}}^{-1}(\ell) = \frac{1}{2a}\left[\hat{L}_{\mathrm{m0}}^{-1}(2a\ell + \hat{L}_{\mathrm{m0}}(b)) - b\right] \tag{2.5}$$

which requires only that we invert $\hat{L}_{\mathrm{m0}}$. Although $\hat{L}_{\mathrm{m0}}$ is not invertible in terms of elementary functions, it is nonetheless bijective and differentiable, and thus subject to efficient numerical inversion by Newton-Rhaphson iteration given a sufficiently good initial guess $x_*$. To determine

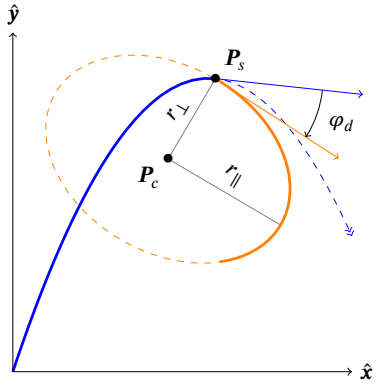


**Figure 2.2:** A possible configuration of the midrib curve, with the quadratic segment $\boldsymbol{P}_{\mathrm{m}0}$ shown in blue and the elliptic segment $\boldsymbol{P}_{\mathrm{m}1}$ shown in orange. The elliptic segment has parallel and perpendicular radii $r_{\parallel}$ and $r_{\perp}$ with respect to its tangent at $\boldsymbol{P}_s$. The elliptic tangent may be angularly displaced from the quadratic tangent by setting $\varphi_d > 0$, as shown.

the initial guess, we notice that the arcsinh term in Equation 2.4 becomes insignificant as $x \to \infty$. That is,

$$\lim_{x\to\infty} \frac{\hat{L}_{\mathrm{m}0}(x)}{\tilde{L}_{\mathrm{m}0}(x)} = 1 \text{ for } \tilde{L}_{\mathrm{m}0}(x) \equiv \frac{1}{2}x\sqrt{1+x^2}.$$

where the simplified function $\tilde{L}_{\mathrm{m}0}$ *is* invertible in terms of elementary functions for $x \geq 0$. We thus select the initial guess as

$$x_* = \tilde{L}_{\mathrm{m}0}^{-1}(\ell) = \sqrt{\frac{1}{2}\sqrt{16\ell^2+1} - \frac{1}{2}}. \tag{2.6}$$

*Elliptic segment.* As a function of ellipse eccentric angle $\varphi$ instead of parameter $t$, the elliptic segment is given by

$$\boldsymbol{P}_{\mathrm{m}1}(\varphi) = \boldsymbol{P}_c + \begin{bmatrix} \cos\varphi_s & \sin\varphi_s \\ -\sin\varphi_s & \cos\varphi_s \end{bmatrix} \begin{bmatrix} -r_{\perp}\cos\varphi \\ r_{\parallel}\sin\varphi \end{bmatrix} \tag{2.7}$$

where $\boldsymbol{P}_c$ is the ellipse center, $\varphi_s$ is the tangent angle at $t_{\mathrm{ms}}$, and $r_{\perp}$ and $r_{\parallel}$ define the perpendicular and parallel radii with respect to the tangent at $t_{\mathrm{ms}}$. Refer to Figure 2.2, which also features the split location $\boldsymbol{P}_s$ and the displacement angle $\varphi_d$ between the quadratric and elliptic tangents at $t_{\mathrm{ms}}$. In accordance with España et al. [23], setting $\varphi_d > 0$ causes the leaf to "break" at $t_{\mathrm{ms}}$.

We accept the parallel radius $r_{\parallel} > 0$ and the signed eccentricity $\varepsilon \in (-1, 1)$ as parameters, then

compute $r_\perp$ as

$$r_\perp = r_\parallel \times \begin{cases} (1-\varepsilon^2)^{+1/2} & \varepsilon \geq 0, \\ (1-\varepsilon^2)^{-1/2} & \varepsilon < 0. \end{cases} \tag{2.8}$$

The arc-length integral $L_{\mathtt{m1}}$ of $\boldsymbol{P}_{\mathtt{m1}}$ is again obtainable by change of variables with careful consideration of the eccentricity $\varepsilon$,

$$L_{\mathtt{m1}}(\varphi) = \begin{cases} r_\parallel E(\varphi, \varepsilon^2) & \varepsilon \geq 0, \\ r_\perp E(\varphi + \frac{\pi}{2}, \varepsilon^2) - r_\perp E(\frac{\pi}{2}, \varepsilon^2) & \varepsilon < 0, \end{cases} \tag{2.9}$$

where $E$ denotes the incomplete elliptic integral of the second kind

$$E(\varphi, m) = \int_0^\varphi \sqrt{1 - m \sin^2 u}\, \mathrm{d}u.$$

As before, we must evaluate Equation 2.7 at $\varphi = L_{\mathtt{m1}}^{-1}(\ell)$ to parameterize by arc-length $\ell$. By manipulating Equation 2.9, we obtain

$$L_{\mathtt{m1}}^{-1}(\ell) = \begin{cases} E^{-1}(\frac{\ell}{r_\parallel}, \varepsilon^2) & \varepsilon \geq 0, \\ E^{-1}(\frac{\ell}{r_\perp} + E(\frac{\pi}{2}, \varepsilon^2), \varepsilon^2) - \frac{\pi}{2} & \varepsilon < 0, \end{cases} \tag{2.10}$$

where $E^{-1}$ denotes the inverse of $E$ with respect to $\varphi$ for fixed $m$. Similarly to the quadratic case, there is no elementary expression for $E^{-1}$, but $E$ is bijective and differentiable with respect to $\varphi$, and thus subject to efficient numerical inversion by Newton-Rhaphson iteration given a sufficiently good initial guess $\varphi_*$. Notice that $E$ is a strictly increasing function of $\varphi$ which oscillates about a line through the origin with slope $(2/\pi)E(\pi/2, m)$. We therefore choose

$$\varphi_* = \frac{\pi}{2} \frac{1}{E(\pi/2, m)} \ell' \tag{2.11}$$

where $\ell'$ denotes the appropriate argument expression from equation 2.10.

### 2.2.2. Forming the lamina surface

We form the lamina surface similarly to España et al. [23], i.e., we use a quadratic width profile $q$ to extrude the lamina in the Z direction at each point in XY on the midrib. However, we substitute the arc-length parameterized midrib developed in the previous section, and subsequently show that the preliminary definition of the lamina surface has an exact area $A$ in addition to an

exact length $L$. We further introduce leaf curl, a new parameter, to work in conjunction with leaf undulation as defined by España et al [23].

*Width profile.* Let $q : [0,1] \to \mathbb{R}$ denote the width profile, which defines the extent of the lamina in the Z direction at each point on the midrib. The width profile is given by

$$q(t) = -q_w(t + q_0)(t - 1) \tag{2.12}$$

where $q_w$ depends on desired area $A$ and length $L$, and $q_0 \in [0,1]$ is a shaping parameter. For $q_0 = 0$, the width is largest at $t = 1/2$. For $q_0 = 1$, the width is largest at $t = 0$. To ensure the area calculation for the preliminary lamina surface works out properly, we define $q_w$ as

$$q_w = \frac{A}{L}\frac{1}{q_0 + 1/3} \implies \int_0^1 q(t)\,\mathrm{d}t = \frac{1}{2}\frac{A}{L}. \tag{2.13}$$

*Preliminary lamina surface.* From equations 2.1 and 2.12, we form a preliminary expression for the lamina surface $\boldsymbol{S}_{\mathfrak{l}0} : [0,1]^2 \to \mathbb{R}^3$ as

$$\boldsymbol{S}_{\mathfrak{l}0}(s,t) = \boldsymbol{P}_{\mathfrak{m}}(t) + (2s - 1)q(t)\hat{\boldsymbol{z}}. \tag{2.14}$$

It is necessary to integrate the magnitude of the cross product of the partial derivatives over the domain of interest, in order to compute the area of a parametric surface in general. This is achievable in closed-form for Equation 2.14, as vector calculus reveals that

$$\left\|\frac{\partial \boldsymbol{S}_{\mathfrak{l}0}}{\partial s} \times \frac{\partial \boldsymbol{S}_{\mathfrak{l}0}}{\partial t}\right\| = 2q\left\|\frac{\partial \boldsymbol{P}_{\mathfrak{m}}}{\partial t}\right\| = 2Lq,$$

where the last equality follows because $\boldsymbol{P}_{\mathfrak{m}}$ is parameterized by arc-length. Then, it is straightforward to show that $\boldsymbol{S}_{\mathfrak{l}0}$ has desired area $A$ by evaluating

$$\iint_{[0,1]^2}\left\|\frac{\partial \boldsymbol{S}_{\mathfrak{l}0}}{\partial s} \times \frac{\partial \boldsymbol{S}_{\mathfrak{l}0}}{\partial t}\right\|\mathrm{d}s\,\mathrm{d}t = 2L\int_0^1 q\,\mathrm{d}t = A.$$

*Incorporating curl.* In order to incorporate leaf curl, we wrap the preliminary lamina surface $\boldsymbol{S}_{\mathfrak{l}0}$ cylindrically around a copy of the midrib curve $\boldsymbol{P}_{\mathfrak{m}}$ displaced in the XY plane along its normal direction with a bilinearly-varying radius. That is, we define

$$r_{\mathfrak{l}}(s,t) = \mathrm{lerp}(s; \mathrm{lerp}(t; r_{\mathfrak{l}00}, r_{\mathfrak{l}01}), \mathrm{lerp}(t; r_{\mathfrak{l}10}, r_{\mathfrak{l}11})) \tag{2.15}$$

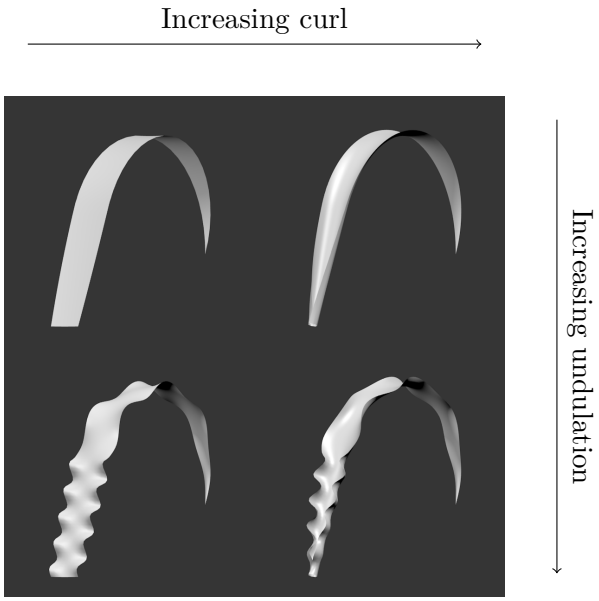


**Figure 2.3:** Lamina surface deformation in the presence of curl and undulation.

for choice of radii $r_{\mathfrak{l}st} > 0$, and where lerp denotes ordinary 1D linear interpolation. Varying the radius with $s$ allows the lamina to spiral inward or outward, while varying the radius with $t$ allows the lamina to taper from one end to the other. Next, we define the cylindrical angle as

$$\theta_{\mathfrak{l}}(s,t) = \frac{q(s,t)}{r_{\mathfrak{l}}(s,t)} \tag{2.16}$$

where $q(s,t) = (2s-1)q(t)$, from Equation 2.14, represents arc-length around the cylinder. Then, we define the curled lamina surface $\boldsymbol{S}_{\mathfrak{l}1}$ as

$$\boldsymbol{S}_{\mathfrak{l}1}(s,t) = \boldsymbol{P}_{\mathfrak{m}}(t) + r(s,t)\boldsymbol{V}(s,t), \tag{2.17}$$

$$\boldsymbol{V}(s,t) = (1-\cos(\theta_{\mathfrak{l}}(s,t)))\hat{\boldsymbol{N}}_{\mathfrak{m}}(t) + \sin(\theta_{\mathfrak{l}}(s,t))\hat{\boldsymbol{z}}, \tag{2.18}$$

where $\hat{\boldsymbol{N}}_{\mathfrak{m}}$ is the right-hand normal direction with respect to the tangent of $\boldsymbol{P}_{\mathfrak{m}}$,

$$\hat{\boldsymbol{N}}_{\mathfrak{m}} = \hat{\boldsymbol{z}} \times \frac{1}{L}\frac{\partial \boldsymbol{P}_{\mathfrak{m}}}{\partial t}.$$

The reader is referred to the top right of Figure 2.3 to see the effect of bilinearly-varying curl on the lamina surface.

It is important to note that Equation 2.17 is a nearly, but not perfectly, area-preserving rendition of Equation 2.14 (though it is perfectly area-preserving in two limiting cases, being in the limit

as $r \to \infty$, and when the $\boldsymbol{P}_{\mathfrak{m}}$ is perfectly straight and $r$ is constant). Calculating the differential area element for $\boldsymbol{S}_{l,1}$ is cumbersome, but is nonetheless achievable in closed-form. We find that the deformed area is typically within 1% of the defined area, even if the radius is small and/or bilinearly-varying.

*Incorporating undulation.* We displace the lamina along its unit-length surface normal to incorporate undulation. We thus define the curled, undulated lamina surface $\boldsymbol{S}_{\mathfrak{l}2}$ as

$$\boldsymbol{S}_{\mathfrak{l}2}(s,t) = \boldsymbol{S}_{\mathfrak{l}1}(s,t) + \gamma(s,t)\hat{\boldsymbol{N}}_{\mathfrak{l}1}(s,t) \tag{2.19}$$

where $\boldsymbol{S}_{\mathfrak{l}1}$ is the curled lamina surface, $\hat{\boldsymbol{N}}_{\mathfrak{l}1}$ is its unit-length surface normal, and $\gamma$ is the magnitude of displacement. Following España et al. [23], we parameterize $\gamma$ by the number of cycles $n_t$ and the angular amplitude $\Omega_t$, being the maximum angular deviation in the surface derivative. Then,

$$\begin{aligned}\gamma(s,t) &= (2s-1)^2 q(1,t)\tan(\Omega_t) \\ &\times \begin{cases} \sin(+\omega_0 t) & t < t_{\mathfrak{ms}} \\ \sin(-\omega_1(t-t_{\mathfrak{ms}})) & t \geq t_{\mathfrak{ms}} \end{cases}\end{aligned} \tag{2.20}$$

where $\omega_0 = 2\pi(n_t + 1/2)/t_{\mathfrak{ms}}$ and $\omega_1 = 2\pi n_t/(1-t_{\mathfrak{ms}})$.

We further allow the number of undulations to be different on either side of the lamina, such that we choose between two values $n_{t0}$ and $n_{t1}$ for $s$ less-than/greater-than 1/2. Note that, just as in the curl deformation, the undulation is approximately but not strictly area-preserving. Refer to the bottom left and bottom right of Figure 2.3 for the effect on lamina geometry.

### 2.2.3. Forming the leaf and internode surfaces

We introduce an explicit parameterization of the sheath, then connect the sheath to the lamina with a cubic patch in order to form the leaf surface. We consider the sheath as a modified cylinder, where the radius and angular extent vary over the domain, such that the sheath surface $\boldsymbol{S}_{\mathfrak{s}} : [0,1]^2 \to \mathbb{R}^3$ is given by

$$\begin{aligned}\boldsymbol{S}_{\mathfrak{s}}(s,t) =& \rho(s,t) r_{\mathfrak{s}}(t)[\cos(\theta_{\mathfrak{s}}(s,t))\hat{\boldsymbol{x}} + \sin(\theta_{\mathfrak{s}}(s,t))\hat{\boldsymbol{z}}] \\ &+ (th + h_0)\hat{\boldsymbol{y}}.\end{aligned} \tag{2.21}$$

**Figure 2.4:** A leaf, formed as a first-order continuous union of the lamina and sheath surfaces.

This indeed looks like a parameterization of a cylinder about the Y axis with height $h$ and vertical offset $h_0$, where the effective radius is the product of an envelope $\rho$ and a profile curve $r_{\mathfrak{s}}$. We define the angular extent $\theta_{\mathfrak{s}}$ to vary bilinearly, and to be symmetric in $s$ given principal angular extents $\theta_{\mathfrak{s}t}$ at $t = 0$ and $t = 1$ respectively, such that

$$\theta_{\mathfrak{s}}(s,t) = (s - 1/2)\,\mathrm{lerp}\,(t; \theta_{\mathfrak{s}0}, \theta_{\mathfrak{s}1}). \tag{2.22}$$

The radius envelope $\rho$ is a bilinear function as well, defined for constants $\rho_{st}$ analogous to Equation 2.15. The radius profile $r_{\mathfrak{s}}$ is a piecewise curve uniting a cubic Hermite segment, which characterizes a bulge at the base of the cylinder achieving a maximum displacment $d$, with a linear segment at a given split point $t_{\mathfrak{s}\mathfrak{s}} \in [0, 1]$.

We connect the sheath and lamina surfaces with a cubic patch to form the final leaf surface. In particular, at each $s$, we construct a cubic Hermite spline connecting $\boldsymbol{S}_{\mathfrak{s}}(s, 1)$ to $\boldsymbol{S}_{\mathfrak{l}2}(s, 0)$ with first-order continuity. We choose the top-edge parameters of the sheath to match the bottom edge of the lamina in terms of radius and angular extent, then position the lamina a distance $c$ from the sheath in the direction of the leaf insertion angle. A realization of the final leaf surface is shown in Figure 2.4. We form the internode similarly to the sheath, i.e., as a cylinder with a profile curve modulating the radius, though without varying angular extent. The internode radius profile curve is piecewise, uniting an interior linear segment with two cubic Hermite segments on the ends. The Hermite segments are formed identically to those in the sheath, also achieving a maximum

displacement $d$.

### 2.2.4. Tessellation

To actually populate a three-dimensional scene with corn plants, we must discretize the geometric parameterization of each plant in order to aggregate a set of facets suitable for ray tracing. We thus obtain a set of vertices by evaluating parametric surfaces on a rectangular grid, then we obtain a set of facets by forming edges between adjacent vertices. This process is known as tessellation. Suppose that we 1) use 32-bit floating point numbers to represent vertex components, 2) use 16-bit unsigned integers to represent the vertex indices for each facet, and 3) tessellate at a low resolution, e.g., 16 vertices and 32 triangular facets per leaf. The facets for a plant having 10 to 20 leaves then will occupy ∼1 KB. This is problematic, as a typical corn field may occupy ≈ 300 acres with ≈ 30 thousand plants per acre (or ≈ 8 plants per square meter). Even for the limited tessellation resolution given above, the facets for this many corn plants would occupy 20–70 GB. To circumvent this issue, we tessellate geometry as needed while performing ray tracing, such that we need not store facets explicitly on disk or in memory. We compute a bounding box for each plant, which bounds the region its facets may occupy, and we tessellate its geometry only if a ray enters the box. For efficiency, we maintain a Least-Recently Used (LRU) cache of recently tessellated geometries. The maximum size of the cache is configurable, and when this size is exceeded, we remove the least-recently used geometry.

We further allow the end user to specify the desired tessellation density of plant geometries. This is given as a pair of floating point values that describe the desired number of subdivisions per centimeter in $s$ and $t$ respectively. Using a higher tessellation density yields higher fidelity geometry suitable for near viewing distances, where each plant may occupy many pixels in the output image. Using a lower tessellation density produces lower fidelity geometry suitable for far viewing distances, where each pixel in the output image may contain many plants.

## 2.3. Methods: Physiological growth modeling

In this section, we discuss the physiological model of growth and development that determines the primary physical characteristics of the geometric parameterization as a function of time. The model does not depend on time explicitly, but instead on *physiological age*. Physiological age is

an abstracted measure of the age of a plant which quantifies its stage of development independent of its chronological age. Physiological age is itself a function of time, temperature, and other environmental factors. We have referred largely to the partly architectural, partly agronomical models presented by Fournier and Andrieu [21] and Drouet and Pagès [22]. Both cited models operate in discrete, fixed time steps, and explicitly simulate things like microclimate and resource allocation within the plant at each time step. From these more rigorous discrete-time models we have distilled a more efficient and more elegant, albeit less rigorous, continuous-time model. That is to say that our model depends continuously on physiological age, which in turn depends continuously on time-varying quantities.

It is worth stating at the outset that we measure physiological age in *plastochrons*, also called *plastochron numbers* or *plastochronic units*; Fournier and Andrieu measured physiological age in so-called growing degree units (GDUs), though the authors made use of plastochrons in certain aspects of their model as well. As we will examine further in the next section, GDUs and plastochrons happen to be redundant measuring systems, and for this reason we adopt plastochrons as the canonical units of physiological age.

### 2.3.1. Growth concepts

A corn plant develops rather predictably, in the sense that there is no deeply recursive, self-similar branching as there is in a common oak tree, for example. In describing the geometric parameterization, we have already introduced the two major components of a corn plant, being leaves and internodes. Leaves and internodes in fact come in pairs, such that there is a leaf for each internode and vice versa. Moreover, each pair forms a functional unit of growth and development known as a *phytomer*. While the plant is actively growing, there is a region of actively subdividing tissue at the apex known as the *apical meristem*, which is responsible for sequentially initiating phytomers.

The plastochron is defined as the length of time between the initiation events of successive phytomers [43]. The rate at which the apical meristem initiates phytomers, which we denote by $R_p$, therefore is of particular interest. Warrington and Kanemasu [44] found that this initiation rate, having units of phytomers per day, depends only on temperature under conditions where other environmental factors are not prohibitive. In particular, $R_p$ is well modeled as a cubic function of

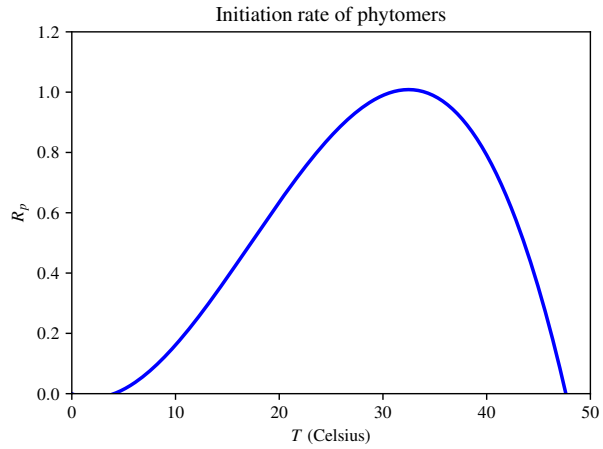


**Figure 2.5:** Initiation rate of phytomers $R_p$ given by Warrington and Kanemasu [44] as a cubic function of temperature $T$ in degrees celsius. $R_p$ is understood to be zero to left and right of the minimum and maximum $T$-intercepts respectively.

temperature $T$ in degrees celsius as

$$\begin{aligned} R_p(T) = &-0.00065 - 0.0138T \\ &+0.00372T^2 - 0.000072T^3, \end{aligned} \tag{2.23}$$

which achieves a maximum value of $R_p \approx 1$ phytomer per day at the optimal growing temperature $T = T_{\text{opt}} \approx 30\,^{\circ}\text{C}$ as shown in Figure 2.5. We may thus obtain the plastochrons $p$ elapsed over a time interval $[\tau_0, \tau_0 + \Delta\tau]$ as

$$p = \int_{\tau_0}^{\tau_0+\Delta\tau} R_p \, \mathrm{d}\tau \tag{2.24}$$

where $R_p$ depends on temperature $T$, which may depend on time $\tau$ in turn. Because $R_p$ is positive for plausible temperatures, the plastochron number $p$ is strictly increasing, and therefore invertible, with respect to the duration $\Delta\tau$.

A unit similar to the plastochron is the *growing degree unit* (GDU), also known as the *growing degree day* (GDD), which is more widely-used in agricultural practice. A measurement in GDUs is obtained by integrating the difference of the meristem temperature $T$ from the base temperature $T_{\text{base}}$ over time. The base temperature is the minimum temperature necessary for growth to occur. We thus can think of GDUs as an approximate, scaled version of plastochrons where $R_p$ is replaced by the simple linear equation $T - T_{\text{base}}$. For typicial mid-latitude summer temperatures, GDUs $G$ for $T_{\text{base}} = 9.8\,^{\circ}\text{C}$ and plastochrons $p$ are in fact linearly related by $G \approx p/0.06$. This is shown in

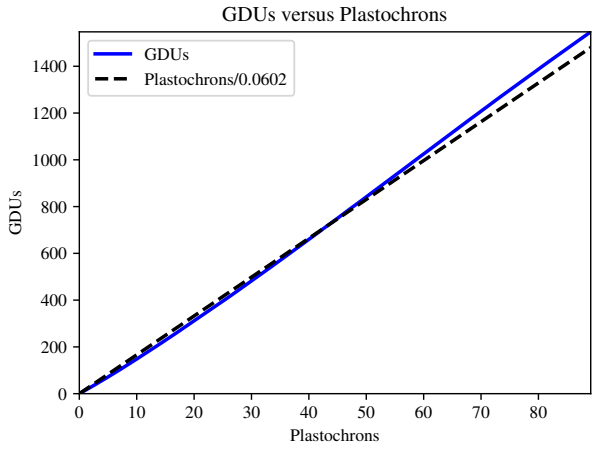


**Figure 2.6:** GDUs versus plastochrons from May 1 (DOY 121) to August 31 (DOY 243) for typical mid-latitude summer temperatures. As shown, the relationship is approximately linear, such that $G \approx p/0.0602$.

figure 2.6. We compute "typical" mid-latitude summer temperatures as discussed in appendix A.

We next provide expressions for the growth and development of leaves and internodes that depend on physiological age in plastochrons $p$. We thus far suppose $R_p$ depends only on temperature, which in turn depends on time, as given by Equation 2.23. However, the plastochron number $p$ as given by Equation 2.24 does not inherently assume anything about the definition of $R_p$. That is to say that the resulting model is generalizable to consider more than temperature and time by substituting an alternative expression for $R_p$ which considers the impact of more, and perhaps more extreme, environmental factors. This is what we mean when we say that the model is defined in a generic "physiological age space". The built-in assumption is that the primary physical characteristics of all corn plants develop identically under identical conditions, and $R_p$ is the fundamental equation that defines how environmental variables affect development.

As a first attempt to incorporate environmental stress, we rewrite $R_p$ as

$$R_p(T, S_w, S_n) = (1 - S_w)e^{-\frac{1}{2}S_n} R_p(T) \tag{2.25}$$

where $S_w \in [0, 1], S_n \in [0, 1]$ are normalized water and nutrient stress parameters, such that a value of zero indicates no stress and a value of one indicates maximal stress. When both parameters are zero, $R_p$ is unchanged. Under maximal water stress $S_w = 1$, $R_p = 0$ and growth ceases. Under maximal nutrient stress $S_n = 1$, $R_p$ is reduced to $\approx 60\%$ of its temperature-limited value. The

form of Equation 2.25 is consistent with the implementation of WOFOST 7.1 in de Wit's Python Crop Simulation Environment (PCSE). The water stress parameter $S_w$ is the complement of the Reduction factor due to excessive transpiration, which is the soil moisture fraction between the critical soil moisture and the wilting point. The nutrient stress parameter $S_n$ is the complement of the least nutrient index in Nitrogen/Phosphorous/Potassium (NPK) simulation.

### 2.3.2. Growth expressions

We characterize a corn plant by four primary physiological parameters, consisting of the total number of leaves $N_{\mathrm{T}}$, the maximum length of the largest lamina $L_{\mathfrak{l}\mathrm{M}}$, the maximum length of the largest internode $L_{\mathfrak{i}\mathrm{M}}$, and the rank of the first elongating internode $n_0$. We use piecewise linear functions of plastochron number $p$, which we derive from discrete-time expressions present in the literature, to model the lengths of laminae, sheaths, and internodes.

*Laminae.* Let $L_{\mathfrak{l}}(k;p)$ denote the length of lamina $k$ as a function of plastochron number $p$. At a high level, we express this as

$$L_{\mathfrak{l}}(k;p) = L_{\mathfrak{l}\mathrm{m}}(k) \cdot \operatorname{clamp}\left[\frac{p - p_{\mathfrak{l}0}(k)}{D_{\mathfrak{l}}(k)} \in [0,1]\right] \tag{2.26}$$

where $L_{\mathfrak{l}\mathrm{m}}(k)$ is the maximum length of lamina $k$. which multiplies a normalized "growth factor" clamped to $[0,1]$. The growth factor is a unitless linear expression of the plastochron number $p$, where $p_{\mathfrak{l}0}(k)$ is the age in plastochrons at which growth begins, and where $D_{\mathfrak{l}}(k)$ is the growth duration in plastochrons. Following Drouet and Pagès [22], we use the allometric length/width relationship of Bonhomme and Varlet-Grancher [45] to determine lamina area $A$ and shaping $q_0$ as a function of lamina length $L$. The width profile associated with this relationship is directly convertible to the width profile $q(t)$ as given by Equation 2.12, constraining $A = 0.0795L^2$ and $q_0 = 0.381$.

Following Fournier and Andrieu [21], the maximum length of lamina $k$ is given by

$$L_{\mathfrak{l}\mathrm{m}}(k) = L_{\mathfrak{l}\mathrm{M}} \mathrm{e}^{\frac{1}{2}a_{\mathrm{M}}(k/n_{\mathrm{M}}-1)^2 + \frac{1}{2}b_{\mathrm{M}}(k/n_{\mathrm{M}}-1)^3} \tag{2.27}$$

where $L_{\mathfrak{l}\mathrm{M}}$ represents the length of largest lamina. Fournier and Andrieu [21] used published data from 22 genotypes to linearly regress the remaining curve parameters against the total number of

leaves $N_{\mathrm{T}}$ that ultimately develop in the plant, thereby obtaining

$$a_{\mathrm{M}} = 0.25N_{\mathrm{T}} - 10.61,$$

$$b_{\mathrm{M}} = 0.27N_{\mathrm{T}} - 5.99,$$

$$n_{\mathrm{M}} = 0.33N_{\mathrm{T}} + 5.93.$$

The duration $D_{\mathrm{l}}(k)$ is determined by dividing the terminal length $L_{\mathrm{lm}}(k)$ by a growth rate $v_{\mathrm{l}}^*$ in units of length per plastochron. Fournier and Andrieu [21] used the growth rate parameter $v_l = 0.564(T - 9.8)$ proposed by Salah [46], which has units of centimeters per day and is a function of meristem temperature $T$ in celsius. This parameter is constructed to be effective in the region where growth rate increases quasi-linearly with temperature. This region is between the base temperature of $T_{\mathrm{base}} = 9.8\,^{\circ}\mathrm{C}$, below which it is understood that growth does not occur, and the optimal growing temperature $T_{\mathrm{opt}} \approx 30\,^{\circ}\mathrm{C}$, above which it is understood that growth rate decreases with temperature. However, $v_l$ is not directly compatible with our approach because it explicitly considers temperature and time, whereas we would like to express growth rate independently in units of length per plastochron. We consider that the derivative of $v_l$ with respect to temperature is a constant $dv_l/dT = 0.564\,\mathrm{cm}/^{\circ}\mathrm{C}/\mathrm{d}$ having units of length per degree-day. As we have previously established, degree-days or GDUs are approximately linearly related to plastochrons within typical mid-latitude summer temperatures (refer back to Figure 2.6). This range of temperatures is between the base and optimal growing temperatures, and so is consistent with the region $v_l$ intends to model. We thus convert $dv_l/dT$ to the desired units of length per plastchron using the relationship $G \approx p/0.06$. This yields

$$v_{\mathrm{l}}^* = \frac{0.564\,\mathrm{cm}/^{\circ}\mathrm{C}/\mathrm{d}}{0.06\,\mathrm{pl}/^{\circ}\mathrm{C}/\mathrm{d}} = 9.4\,\mathrm{cm}/\mathrm{pl}.$$

The age in plastochrons at which growth begins $p_{\mathrm{l0}}(k)$ is given by the age at which the phytomer is initiated plus the delay between initiation and linear elongation. The age in plastochrons at which phytomer $k$ is initiated by the apical meristem is, by definition, $p = k$. The delay between initiation and linear elongation in plastochrons was fit by Fournier and Andrieu [21] to data published by Cao et al. [47] and Zur et al. [48]. This fit is part linear and part cubic. The linear part is $1.94k - 5.16$ for $k < n_{\mathrm{M}}$, and the cubic part is constrained to be first-order continuous at $k = n_{\mathrm{M}}$, to have zero slope at $N_{\mathrm{T}}$, and to increase by 3.65 from $k = n_{\mathrm{M}}$ to $k = N_{\mathrm{T}}$.

*Sheaths and internodes.* We express the lengths of sheaths and internodes in exact analogy to Equation 2.26. For brevity, we do not include the expressions here. We calculate the growth duration for sheaths and internodes analogously, by dividing terminal length by growth rate. The growth rate of sheaths $v_{\mathfrak{s}}^*$ is equivalent to the growth rate of laminae $v_{\mathfrak{l}}^*$, and the growth rate of internodes $v_{\mathfrak{i}}^*$ is set to $0.18v_{\mathfrak{l}}^*$. Note that the growth rates of sheaths and internodes are also defined in units of length per plastochron, but are otherwise analogous to the values used by Fournier and Andrieu.[1] Elongation of the sheath is 30% complete when lamina elongation ceases, and elongation of the internode begins when the elongation of the sheath is 60% complete. These heuristics are reported, but not used, by Fournier and Andrieu [21] based on data by Sharman [49], Hesketh et al. [50], Grant and Hesketh [51], and Robertson [52].

*Randomized parameters.* The primary physical characteristics of each plant are determined by the four parameters $N_{\mathrm{T}}$, $L_{\mathfrak{l}\mathrm{M}}$, $L_{\mathfrak{i}\mathrm{M}}$, and $n_0$ stated previously. We further store the location and azimuthal angle, as well as a 32-bit integer identifier or "seed value". The seed value is used to initialize a Pseudo-Random Number Generator (PRNG), which is used in turn to randomize all remaining secondary physical characteristics, such as leaf insertion angles and curvatures, within plausible distributions given by España et al. [23] and Drouet and Pagès [22]. Regarding the choice of PRNG, we use and recommend the 32-bit Permuted-Congruential Generator (PCG) [53], which is a principled modification of the Linear-Congruential Generator (LCG) that is comparably efficient and produces statistically better output.

It is important to emphasize that, while these parameters are randomized, they are not stochastic. In other words, each seed value uniquely and deterministically identifies a configuration of secondary physical characteristics. Furthermore, we linearly interpolate many randomized parameters from sensible initial values to the randomized terminal values using the appropriate growth factors. For example, each leaf insertion angle is always initially vertical, growing upward from the apex, and interpolates to its randomized terminal value via the lamina growth factor. This results in each plant being truly physically unique and growing continously over time, as shown in Figure 2.7. By modeling the growth and development of plant geometry explicitly, known phenomenology such as shadow-hiding and opposition surge (an extreme case of shadow-hiding) is inherently reproducable,

[1] Also, Drouet and Pagès report a value of $0.41v_{\mathfrak{l}}^*$ for the growth rate of internodes, yet cite Fournier and Andrieu. This is a misprint, and should be instead $0.14v_{\mathfrak{l}}^*$.

**Figure 2.7:** Plant geometry shown at 15-day increments as output by the growth model under temperature-limited (no stress) conditions.

as shown in Figure 2.8. This is scalable to scenes containing millions of unique plants, using the runtime tessellation technique described in Section 2.2.4, since each plant is only represented by a handful of physiological parameters, its location and azimuthal angle, and an integer seed value.

## 2.4. Methods: Optical properties modeling

We implement a PROSPECT-inspired model to obtain leaf optical properties. In particular, we use PROSPECT [25] to simulate leaf hemispherical reflectance and transmittance given the concentrations of chlorophyll $a + b$ and water. We further simulate a perfectly energy conserving Bidirectional Scattering Distribution Function (BSDF) for a rough dielectric plate, in order to capture directional variation. This is consistent with the theoretical understanding of a compact leaf in the PROSPECT model for $N \approx 1$. We re-weight the reflective hemisphere of this BSDF by the hemispherical reflectance and the transmissive hemisphere of this BSDF by the hemispherical transmittance.

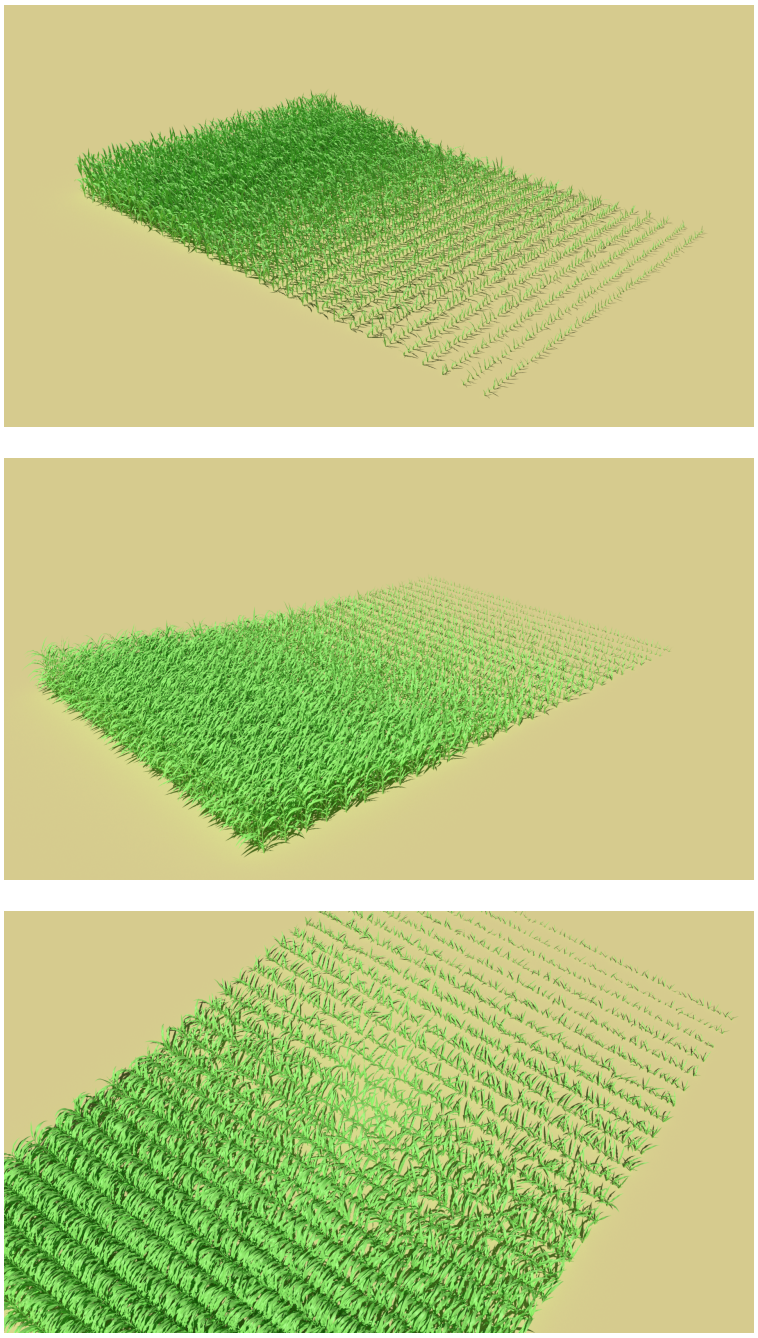

**Figure 2.8:** An example corn plot, where the stage of development is incremented uniformly across rows. These images were rendered with Blender's RGB-centric Cycles path tracer (not DIRSIG) with ad hoc optical properties. The intent of this figure is to demonstrate observable phenomenology due to explicit plant geometry, namely the effects of shadow-hiding (top two images) as well as the extreme case of opposition surge (bottom image).

*Spherical Quad Tree BSDFs.* DIRSIG5 features a numerical BSDF format called Spherical Quad Tree (SQT) [5]. This format constructs a quad-tree-like data structure on the surface of a sphere from given BSDF measurements for efficient and smooth interpolation and sampling of an arbitrary scattering function at runtime. An SQT is most often used to incorporate real-world Bidirectional Reflectance Distribution Function (BRDF) data measured by a goniometer into a DIRSIG scene. It is equally valid however to use an SQT to incorporate a simulated BSDF into a DIRSIG scene, such as that of a rough dielectric plate.

### 2.4.1. Layered-SQT

We developed a separate tool, called `layered-sqt`, for simulating the emergent BSDF of a layered assembly in SQT format for use with DIRSIG scenes. A *layered assembly* is a theoretical construction consisting of $M$ layers, separated by $M+1$ participating media for $M \geq 1$. A layer is an infinite plane which is offset along (and normal to) the $z$-axis, and which is associated with a constituent BSDF that describes how light scatters upon intersection. A medium occupies the space between adjacent layers or, in the boundary cases, the spaces above and below the top and bottom layers respectively. See Figure 2.9 for clarification. The *emergent BSDF* is the BSDF observed in the limit as one backs infinitely far away from the assembly or, identically, as the assembly shrinks to an infinitesmial point. To ensure that the emergent BSDF is well-defined, `layered-sqt` requires that the properties of surfaces and participating media be homogeneous, i.e., independent of spatial location. Furthermore, `layered-sqt` does not account for wavelength-dependence for simplicity/tractability.

*Basics of program usage.* The `layered-sqt` program is run from the command-line. It scans command-line arguments for (optional) configuration flags and a (required) input filename appearing somewhere as a positional argument, i.e., an argument not consumed by a flag. It then parses the input file, simulates the emergent BSDF, and writes the results to a RAW-format file ready for conversion to SQT-format via the existing `raw2sqt` program. For instance,

```
$ ./layered-sqt example.lsqt -p 50000 -o example.raw
```

simulates a layered assembly described in `example.lsqt` with 50,000 paths, and writes the emergent BSDF in plain-text RAW format to `example.raw`. Then, run

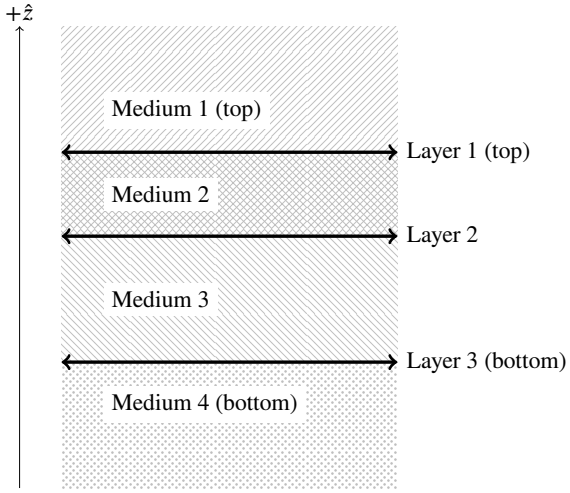


**Figure 2.9:** An example diagram of a layered assembly with three layers and four participating media.

```
$ ./raw2sqt example.raw
```

to write a new SQT-format file `example.sqt` for use with DIRSIG.

Two of the most common flags appear above, being `-p` (or `--path-count`) to specify the number of light paths sampled in the simulation and `-o` (or `--output`) to specify the output filename. To see a list of all acceptable flags with brief descriptions and default values, pass `-h` (or `--help`), or simply run `layered-sqt` with no filename. As an aside, `layered-sqt` verifies parameters specified in command-line flags, as well as keyword arguments specified in the input LSQT file. In the event that something has an unreasonable value, the program issues an error and fails in a controlled manner. Lastly, `layered-sqt` recognizes the single dash filename "-" as standard input. This makes it possible to pipe the (presumably LSQT format) output of a script into `layered-sqt` directly, if this happes to be convenient. As a trivial example,

```
$ cat example.lsqt | \
  ./layered-sqt -p 50000 -o example.raw -
```

is equivalent to just passing `example.lsqt`.

*Basics of LSQT format.* The structure of a layered assembly is easy enough to convey in plain-text with rudimentary syntax, so this is the format `layered-sqt` accepts as input. We refer to this as "LSQT format", and we suffix associated filenames with the extension `.lsqt` (though this suffix

is not strictly required for the program to run). An LSQT file is therefore a line-by-line plain-text description of a layered assembly from top to bottom. So, the first line describes the top medium, the second line describes the top layer beneath the top medium, the third line describes the medium beneath the top layer, and so on until the bottom layer and bottom medium. That being the case, odd-numbered lines describe media and even-numbered lines describe layers. The syntax is

```
Name key1=val1 key2=val2
```

where the identifier `Name` is described by keyword arguments `key1` and `key2`. Importantly, this syntax is whitespace-delimited, so keyword arguments of the form `key=val` must not contain whitespace. It is also worth mentioning that the keyword arguments may appear in any order.

### 2.4.2. Layered-SQT tutorial

The most intuitive way to introduce the functionality of `layered-sqt` is by a brief tutorial. This section is not intended to provide a comprehensive treatment of the structure and usage of `layered-sqt`, but instead to provide an idea of how the program is used and what the program simulates. We will review two examples, the first of which is trivial, and the second of which is more interesting.

We first consider simulating a trivial layered assembly which represents a 60% reflective Lambertian surface. To do so, we create a plain-text file `Lambertian.lsqt` with three lines.

```
1 Medium
2 Layer z=0 Lambertian fR=0.6
3 Medium
```

Above, lines #1 and #3 specify the top and bottom media as vacuum, i.e., writing `Medium` with no keyword arguments indicates vacuum. Line #2 specifies a layer at $z = 0$, with a Lambertian BSDF with reflectance coefficient $f_R = 0.6$. We next run `layered-sqt` on `Lambertian.lsqt`, without passing any additional command-line flags. Doing so creates two new files: `Lambertian.raw`[2] and `Lambertian.lsqt.lss`. As stated in the previous section, we may convert the plain-text RAW file to a binary SQT file for use with DIRSIG by running `raw2sqt`. The LSS file is a binary cache `layered-sqt` uses to store simulation data. LSS is an initialism for LSQT-Slice. It is useful for 1)

[2]By default, `layered-sqt` replaces the extension `.lsqt` with `.raw` to form the output filename.

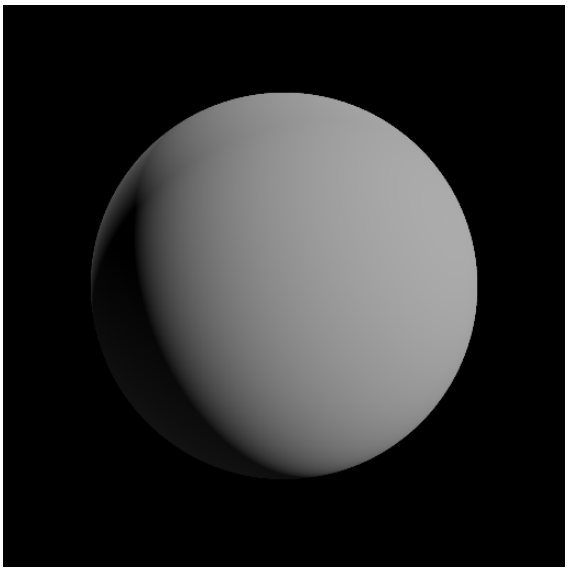

**Figure 2.10:** A preview of a 60% reflective Lambertian BRDF output by `layered-sqt-lssview`.

previewing the emergent BSDF *without setting up and rendering a DIRSIG scene* and 2) simulating the emergent BSDF progressively, with multiple runs of `layered-sqt`.

We may run `layered-sqt-lssview`, an accompanying program packaged with `layered-sqt`, to preview the emergent BSDF in an LSS file. For example,

```
$ ./layered-sqt-lssview Lambertian.lsqt.lss
```

writes a new file `Lambertian.lsqt.png`, which is a $512 \times 512$ rendering of the BSDF applied to a ball, shown in Figure 2.10. It is important to note that `layered-sqt-lssview` is *not* a full-blown path-tracer, and may not perfectly represent how the BSDF will appear in DIRSIG. It only accounts for the direct (first bounce) contributions of a few directional light sources, and it further uses tone-mapping and sRGB correction, such that the output image is not suitable for any radiometric analysis. The intended use of the preview image is to determine if a simulated BSDF is suitably convergent/noise-free.

We next consider simulating a slightly more interesting layered assembly which represents a 20% reflective Lambertian surface with a dielectric coating and a layer of back-scattering dust. We create a plain-text file `Dusty.lsqt` with seven lines.

```
1 Medium
2 Layer z=2 Null
```

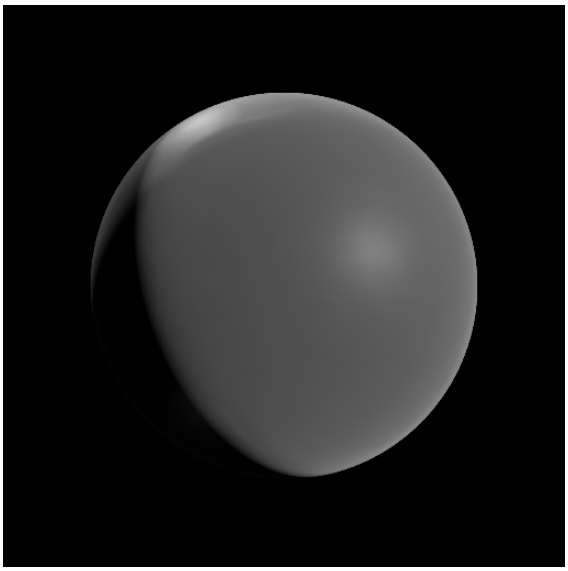

**Figure 2.11:** A preview of a 20% reflective Lambertian BRDF with a dielectric coating and a layer of backscattering dust output by `layered-sqt-lssview`.

```
3 Medium mus=0.2 HenyeyGreenstein g=-0.4
4 Layer z=1 MicrosurfaceDielectric alpha=0.2
5 Medium eta=1.4
6 Layer z=0 Lambertian fR=0.2
7 Medium
```

Similarly to the previous example, lines #1 and #7 specify the top and bottom media as vacuum, and line #6 specifies a layer at $z = 0$ with a 20% reflective Lambertian BRDF. There is much to unpack, however, in lines #2, #3, #4, and #5. Line #2 specifies a null layer at $z = 2$. A *null layer* is a layer with no associated BSDF, used only to separate media with different properties. Line #3 specifies a scattering medium with coefficient $\mu_s = 0.2$, characterized by the Henyey-Greenstein phase function (the volume-scattering analog of a surface BSDF) with shape parameter $g = -0.4$. Line #4 specifies a layer at $z = 1$ with a dielectric microsurface BSDF with roughness coefficient $\alpha = 0.2$. Line #5 specifies a medium with refractive index $\eta = 1.4$.

Regarding the new keyword arguments that appear for media—every medium in an LSQT file has a refractive index $\eta$ (`eta`), which defaults to one, and two volume scattering parameters, which default to zero. The volume scattering parameters consist of the scattering coefficient $\mu_s$ (`mus`) and

the absorption coefficient $\mu_a$ (`mua`). Scattering and absorption events happen according to Beer's Law in a homogeneous medium. That is, the probability of scattering within a particular distance $d$ is given by an exponential distribution $1 - \exp(-\mu_s d)$, where $\mu_s$ is the distribution parameter. This applies analogously to the probability of absorption and $\mu_a$. We thus identify the units of $\mu_s$ and $\mu_a$ as inverse distance, and we interpret the reciprocals $1/\mu_s$ and $1/\mu_a$ as the mean distances between scattering and absorption events, respectively. Note that in the limit $\mu_s = \mu_a = 0$ (the default values), the mean distances tend to infinity, such that scattering and absorption events cease to exist. This is why `Medium` with no keyword arguments, having $\eta = 1$ and $\mu_s = \mu_a = 0$, is equivalent to vacuum.

To simulate a thin layer of back-scattering dust, line #3 specifies a medium with a scattering coefficient $\mu_s = 0.2$ and a Henyey-Greenstein phase function with shape parameter $g = -0.4$. We choose $\mu_s$ to be 0.2 so that the mean distance between scattering events is $1/\mu_s = 5$ units. This is somewhat large in comparison to the thickness of the dust layer, which is 1 unit, so that the dust appears suitably thin. The Henyey-Greenstein phase function is widely used in computer graphics due to its simplicity and intuitive parameterization (although it was originally intended to model scattering of interstellar dust clouds). This phase function is equipped with a single shape parameter $g \in (-1, +1)$, which is identical to the mean scattering cosine. That is, it tends to forward scatter as $g \to +1$ and back scatter as $g \to -1$, and is isotropic at $g = 0$.

Line #4 specifies a layer with a dielectric microsurface BSDF. A *microsurface* is thought to be an infinitesimally thin cloud of microscopic facets, where each facet is thought to scatter light according to another simpler BSDF [42]. The cloud is characterized geometrically by a distribution of the slopes of the facets and a distribution of the heights of the facets. The distribution of slopes is parameterized by its so-called roughess $\alpha$ (more generally, anisotropic roughness $\alpha_x$, $\alpha_y$). As $\alpha \to \infty$, the distribution of slopes widens and the microsurface BSDF appears rougher. As $\alpha \to 0$, the distribution of slopes collapses. This recovers the initial, simpler BSDF in the limiting case where $\alpha = 0$. A so-called *dielectric* microsurface applies the dielectric Fresnel mirror delta BSDF to the microfacets. So, line #4 using $\alpha = 0.2$ corresponds to a somewhat-rough Fresnel mirror, where the relative refractive index is formed as 1/1.4 from the media above and below on lines #3 and #5.

To simulate this layered assembly effectively, we run

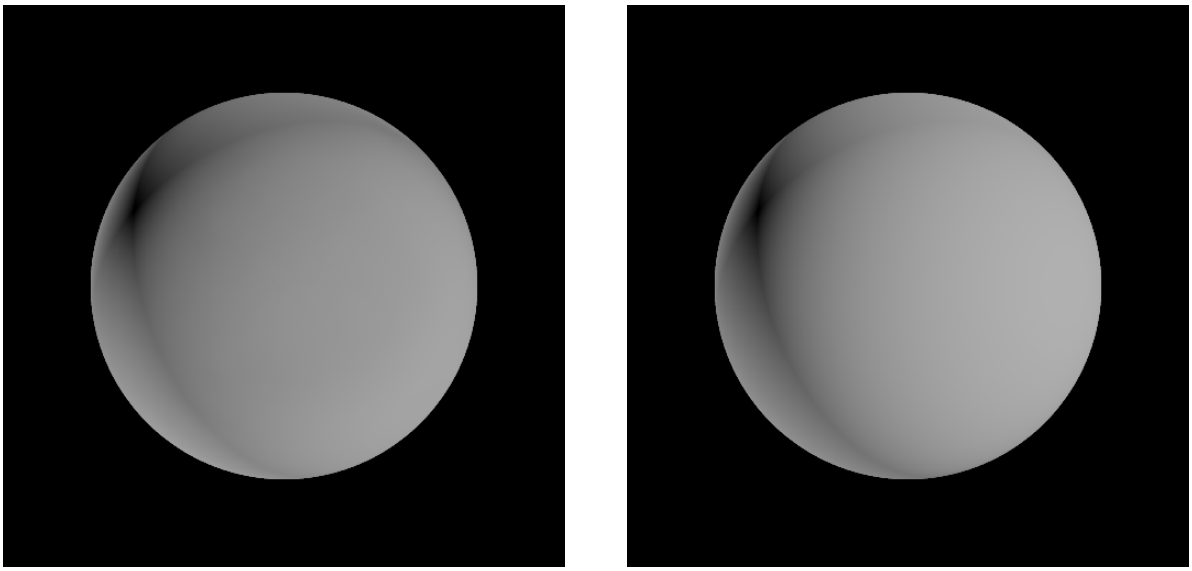

**Figure 2.12:** On the left, a PROSPECT-inspired rough dielectric plate BSDF with refractive index 1.4. On the right, a standard 50% reflective, 50% transmissive Lambertian BSDF for comparison. Both are spectrally unweighted and perfectly energy conserving.

```
$ ./layered-sqt -wi=300 -p=100000 Dusty.lsqt
```

where `-wi=300` increases the number of incident direction samples $\omega_i$ from 80 (the default) to 300 and `-p=100000` increases the number of simulated light paths from 10,000 (the default) to 100,000. The simulation takes a few minutes to complete. The preview produced by `layered-sqt-lssview` is shown in Figure 2.11.

### 2.4.3. PROSPECT-inspired plate BSDF

A leaf is represented as a stack of $N$ dielectric plates in the PROSPECT framework. A corn leaf is a *compact leaf*, meaning that its internal structure is relatively simple and uniform. As such, it is in fact modeled as a single plate with $N \approx 1$. Though this representation is well-suited to predicting net hemispherical optical properties, it is not ideal for predicting the directional variation of these properties. That is to say that a smooth dielectric plate exhibits directional scattering similar to a Fresnel mirror. Corn leaves in the real world do not appear mirror-like but instead appear much more diffuse. PROSPECT hemispherical spectra thus are typically applied in conjunction with the Lambertian assumption, i.e., that scattering is uniform in all directions. This is more realistic but still far from accurate. In order to unite the ordinary PROSPECT framework with the observation that leaves appear diffuse, we consider a corn leaf to scatter directionally as a dielectric plate

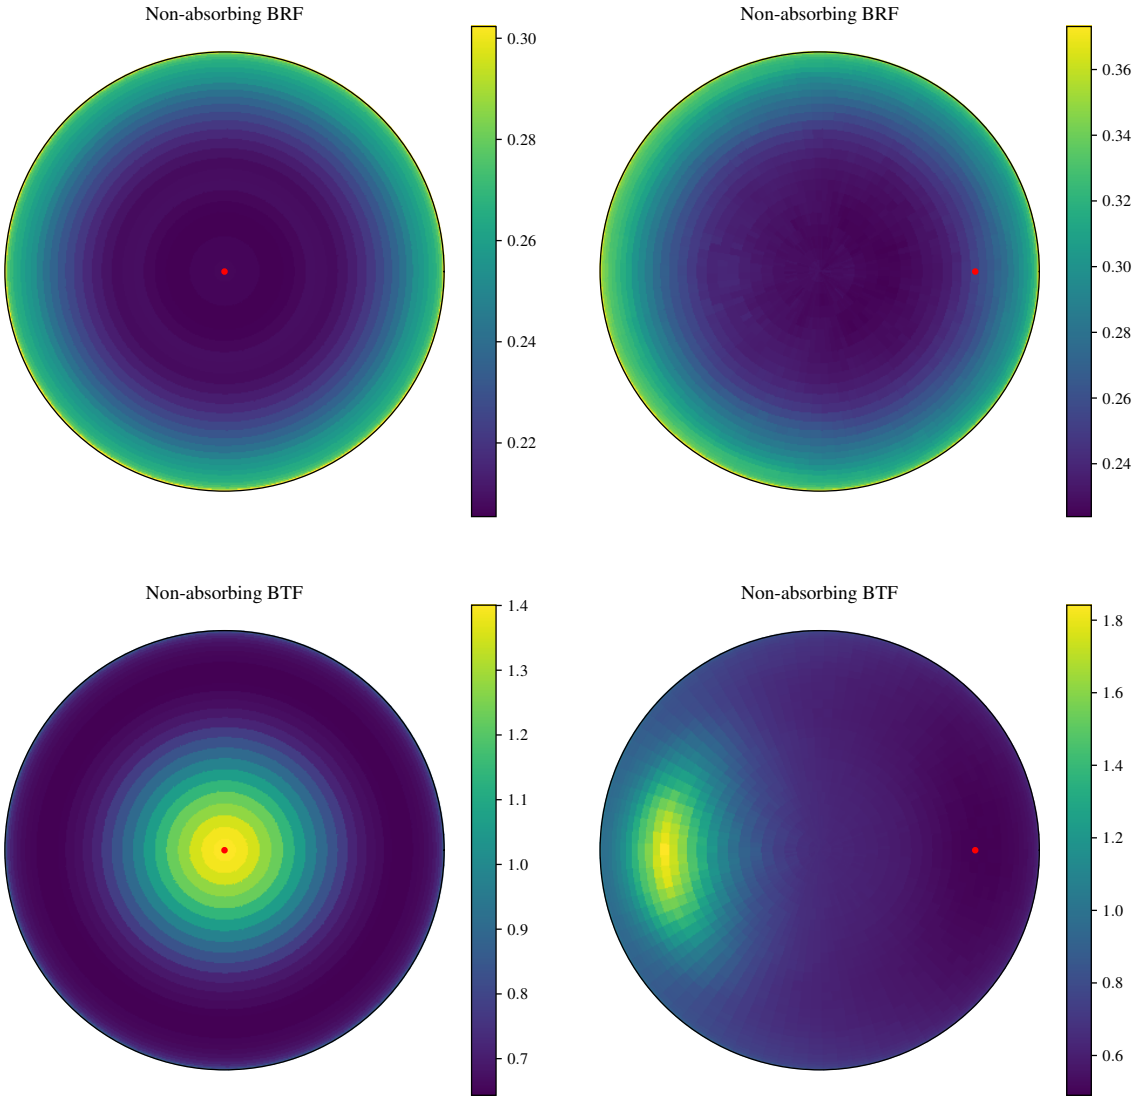


**Figure 2.13:** Hemispherical polar plots of the BRF (top row) and BTF (bottom row) of the non-absorbing PROSPECT-inspired BSDF, shown for viewing zenith angles of 0° (left column) and 45° (right column). The red dot in each plot corresponds to the viewing direction.

with rough top and bottom surfaces instead of smooth ones. For the purpose of simulating the scattering distribution of the BSDF, we further assume that the plate is perfectly energy conserving (no energy is absorbed). We reincorporate absorption by subsequently re-weighting the reflective and transmissive hemispheres of the BSDF to reproduce the Directional Hemispherical Reflectance (DHR) and Directional Hemispherical Transmittance (DHT) at nadir, as predicted by the standard PROSPECT model. This separates the directional behavior of the BSDF into a more tractable subproblem.

We use `layered-sqt` to simulate the emergent BSDF of a rough dielectric non-absorbing plate as a layered assembly with two layers, both characterized by a rough, isotropic microsurface BSDF. Using the microsurface formalism to model reflectance by rough surfaces was first popularized in the computer graphics community by Cook and Torrance [54]. The generalization of reflective microsurface BRDFs to both reflective and transmissive BSDFs was given by Walter et al. [55]. Until recently however, microsurface BSDFs suffered from energy loss at high roughness values due to over-estimation of shadowing and disregard of multiple-scattering. This was resolved by rigorous analysis of microsurface shadowing properties and the introduction of multiple-scattering microsurface models by Heitz et al. [41], [42]. We attribute the latter, perfectly energy-conserving BSDF model of Heitz et al. [42] to the top and bottom surfaces of the plate with a roughness of 2 and a refractive index of 1.4. In a more comprehensive simulation, roughness and refractive index would be parameters that potentially vary with wavelength. We note, however, that the emergent BSDF is not particularly sensitive to refractive index within a sensible range, say $[1.2, 1.8]$, and is decreasingly sensitive to increasing roughness $>> 2$.

The emergent BSDF is shown in Figure 2.12 next to a 50% reflective, 50% transmissive Lambertian BSDF for reference. Hemispherical polar plots of the Bidirectional Reflectance Factor (BRF) and Bidirectional Transmittance Factor (BTF) for viewing angles at nadir and 45 degrees off axis are shown in figure 2.13. Notice that the plate BSDF reflects proportionally more energy at grazing angles and transmits proportionally more energy parallel to the incident direction.

## 2.5. Validation of the corn models

We validated our model against UAS data collected in the summer of 2018 at the United States Department of Agriculture (USDA) Beltsville Agricultural Research Center (BARC) in Beltsville,

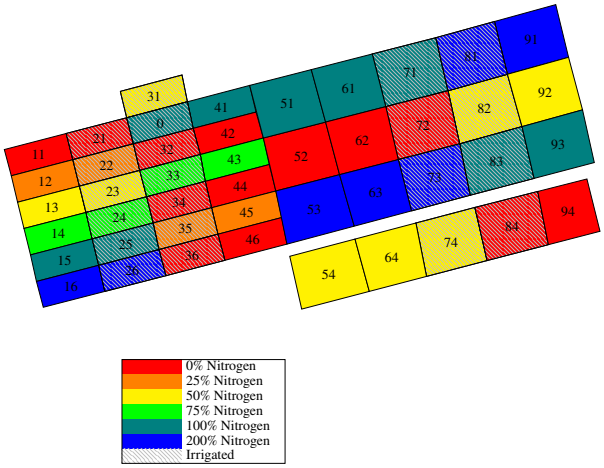


**Figure 2.14:** A top-down diagram of the corn plots at the BARC site, with 25 smaller plots (appearing to the left/west) and 20 larger plots (appearing to the right/east). Each is associated with an integer ID, displayed in the center. Each is filled to indicate how much nitrogen it received, and hatched if it received irrigation.

| | Non-irrigated | Irrigated |
|---|---|---|
| 0% N | 7 | 6 |
| 25% N | 2 | 2 |
| 50% N | 4 | 4 |
| 75% N | 1 | 2 |
| 100% N | 5 | 3 |
| 200% N | 5 | 3 |

**Table 2.1:** The distribution of the 44 plants for which measurements were collected over nitrogen and irrigation treatments.

Maryland. That is, we verified that the radiance predicted by a DIRSIG5 simulation of virtual corn canopy is suitably close to the radiance captured by a real multispectral agricultural imaging system, namely the MicaSense RedEdge-3, flown by the RIT UAS team on the morning of June 18, 2018. The BARC site is located at approximately 76.8° West and 39.0° North and occupies approximately 200 meters by 100 meters. The site consisted of 25 smaller corn plots and 20 larger corn plots, as depicted in Figure 2.14, planted on May 9. The plots were variously subjected to six different nitrogen treatments, being 0%, 25%, 50%, 75%, 100%, and 200% of the recommended rate of 130 pounds per acre. Around half of the plots were irrigated. The USDA collected measurements from 44 different plants on nine occassions in the months of June and July. The measurements include plant height and chlorophyll content, quantities of particular interest to this work. As may be evident, examining 44 plants across six different nitrogen treatments and two different irrigation treatments left little headroom to draw conclusions confidently from statistical analyses of specific treatment combinations. The effective sample sizes were single digit numbers, and the 44 plants were not evenly stratified across treatment categories as shown in table 2.1. We nonetheless considered general trends in plant height and chlorophyll content in our virtual reconstruction of the site.

Plant height seemed to depend more on nitrogen treatment than on whether or not irrigation was present, as shown in Figure 2.15. The USDA remarks widespread drought-stress occurring only from July 3 until July 20, sufficiently after the early stages of development. There was notable variability in plant heights, even within specific treatment combinations, from $\pm 6\,\text{cm}$ on June 12 to $\pm 40\,\text{cm}$ by July 16. A handheld SPAD-502 meter was used to measure leaf chlorophyll content non-destructively. SPAD really measures "greenness" via a spectral band index which correlates with chlorophyll content, such that an empirical equation is required to obtain absolute concentrations [56]. We used the so-called literature consensus equation for converting SPAD readings to chlorophyll concentrations given by Cerovic et al. [57], which is $C_{a+b} = 99S/(144 - S)$ micrograms per square centimeter. Leaf chlorophyll content was highly variable, even more so than plant height, such that it is not sensible to speculate about its relationship to the applied nitrogen and irrigation treatments. Overall, the mean chlorophyll content increased from $\approx 25\,\mu\text{g/cm}^2$ on June 12 and plateaued at $\approx 40\,\mu\text{g/cm}^2$ in early July.

We virtually reconstructed the BARC site as a DIRSIG scene to obtain simulated imagery that

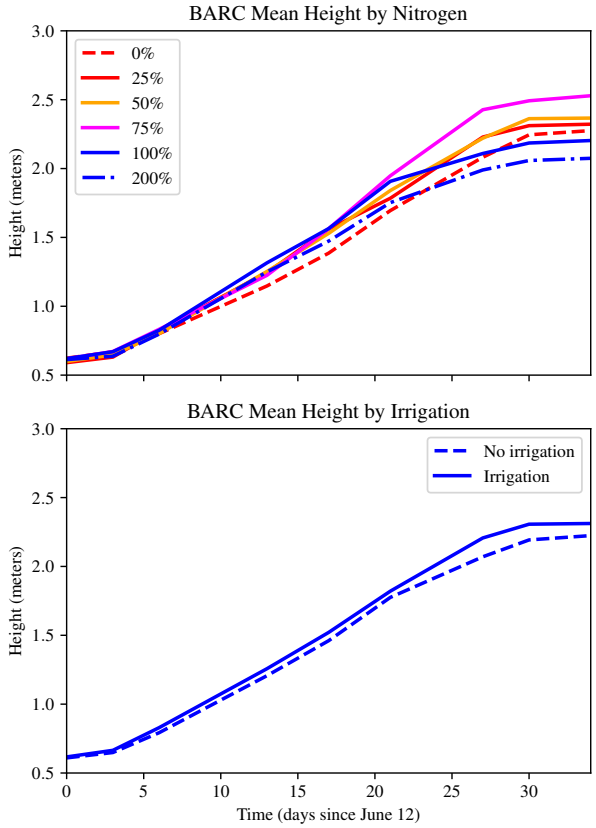


**Figure 2.15:** The mean heights of plants measured at the BARC site over time, separated by nitrogen and irrigation treatment.

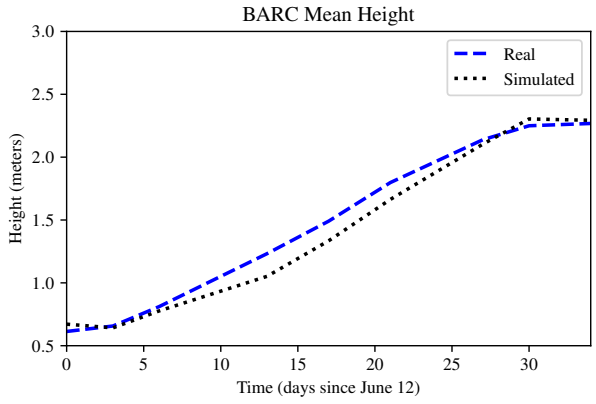


**Figure 2.16:** The mean height of corn plants measured at the BARC site over time versus mean height of simulated plants. Note that the height of a corn plant is measured from the ground to the highest point on the top leaf which has begun to curl downward. We calculated the height of simulated plants according to this principle.

| Band name | Band center | FWHM |
|---|---|---|
| Blue | 0.475 µm | 0.02 µm |
| Green | 0.560 µm | 0.02 µm |
| Red | 0.668 µm | 0.01 µm |
| Red Edge | 0.717 µm | 0.01 µm |
| Near IR | 0.840 µm | 0.04 µm |

**Table 2.2:** MicaSense RedEdge band information used to simulate imagery with DIRSIG5.

is qualitatively and quantitatively comparable to MicaSense imagery collected by the RIT UAS team. The scene was defined relative to $(341821, 4321180, 30.2031)$ in the EPSG:26918 coordinate system for UTM zone 18N. The underlying terrain was based on tessellation of the Digital Elevation Map (DEM) of the surrounding area, acquired through Maryland's iMAP data portal, at one-meter resolution. The USDA provided ESRI Shapefiles describing the plot regions shown in Figure 2.14. We interpolated plant locations within these regions to have row and column spacing roughly equal to that at the BARC site, i.e., $\approx$ 36 inches between rows and $\approx$ 8 inches between columns. We modeled plant growth via equation 2.25 with daily average temperatures measured by the Baltimore Washington International Airport weather station. We applied up to 40% nutrient stress (ad hoc) for plots without any applied nitrogen, with spatially-correlated noise (specifically, three layers of 2D simplex noise [58]) to introduce lower-frequency variability. This produced $\approx 70,000$ unique plants at varying stages of development. Figure 2.16 shows the mean height of measured plants versus the mean height of simulated plants. A handheld Spectra Vista Corporation (SVC) spectrometer was used to collect reflectance measurements at the site. We fit a PROSPECT spectrum to SVC measurements of leaf reflectances in order to determine the Directional Hemispherical Reflectance (DHR) and Transmittance (DHT) at nadir, to use in combination with our PROSPECT-inspired BSDF described in Section 2.4. We varied the chlorophyll concentration of the fit within the range indicated by SPAD readings to introduce optical properties variability. We further used SVC measurements to characterize the reflectance of the ground directly via the Lambertian assumption. We simulated imagery captured by the MicaSense RedEdge-3 at 10:50AM on June 18, flying at $\approx$ 60 meters above the ground, moving at $\approx$ 2 meters per second. The RedEdge has five spectral bands in the visible through the near infrared, shown in table 2.2, and a $1280 \times 960$ focal plane with a 5.5 µm focal length and $3.75 \times 3.75\,\mu m^2$ pixels. We used these quantities to characterize the virtual sensor in DIRSIG. Running the simulation took around five minutes for decent convergence,

and between 20 and 30 minutes for nearly noiseless convergence. Figure 2.17 shows the blue, green, and red bands of real and simulated images.

### 2.5.1. Results

We statistically analyzed sets of *image tiles* extracted from real and simulated images in order to quantitatively compare the real and simulated imagery. That is, we partitioned visually similar real and simulated $1280 \times 960$ images into $8 \times 6$ grids of $160 \times 160$ square tiles, thus forming a set of 48 real tiles and a set of 48 simulated tiles for analysis. The purpose of partitioning the larger images into sets of smaller image tiles was to preserve the spatially-local statistical properties of the imagery in the analysis. We first calculated normal summary statistics $(\mu, \sigma)$, being the mean radiance $\mu$ and standard deviation $\sigma$, for each band in each individual tile. We then calculated the average $(\mu, \sigma)$ statistics across all tiles in each of the real and simulated tile sets. The average $(\mu, \sigma)$ statistics provide a broad measure of the "center" and "spread" of the radiance distributions by which the general agreement of the real and simulated tile sets may be assessed. The average $(\mu, \sigma)$ statistics do not provide a proper measure of "shape" however, as we have no reason to believe that the radiances are normally distributed. In order to assess the agreement of the "shape" of radiance distributions between real and simulated tile sets, we further examined the $\chi^2$ histogram distance between tiles. For each band of each tile, we formed a radiance histogram with $N$ bins over the range $[0, L_{\max}]$, where $L_{\max}$ is the maximum recorded radiance in the real imagery. We selected $N = 30$ to be roughly double the number of bins given by Sturges's rule for $160^2$ observations, i.e., $1 + \log_2 160^2 \approx 15.6$, to avoid overly-smoothing trends in the data. Examples of real and simulated tiles are shown in figure 2.18, and the radiance histograms of these tiles are shown in Figure 2.19. We used the $\chi^2$ distance metric to quantitatively compare radiance histograms in each band. The $\chi^2$ distance between normalized histograms $H_1$ and $H_2$ is given by

$$\chi^2(H_1, H_2) = \frac{1}{2} \sum_{k=1}^{N} \frac{(H_1(k) - H_2(k))^2}{H_1(k) + H_2(k)}.$$

Note that this gives a proper sense of "distance", i.e., it is symmetric $\chi^2(H_1, H_2) = \chi^2(H_2, H_1)$ and non-negative $\chi^2(H_1, H_2) \geq 0$ for all histograms $H_1, H_2$. Further, $\chi^2(H_1, H_2) = 0$ only if there is an exact match $H_1 = H_2$. To provide a baseline distance to assess the relative closeness of radiance distributions of corn-filled image tiles, we further selected a $160 \times 160$ tile containing only

**Figure 2.17:** The blue, green, and red bands from a real radiance-converted image collected by the MicaSense RedEdge (top) and from a DIRSIG5 simulation (bottom). The simulated date is June 18, 2018, with the plants in growth stages V6–V8.

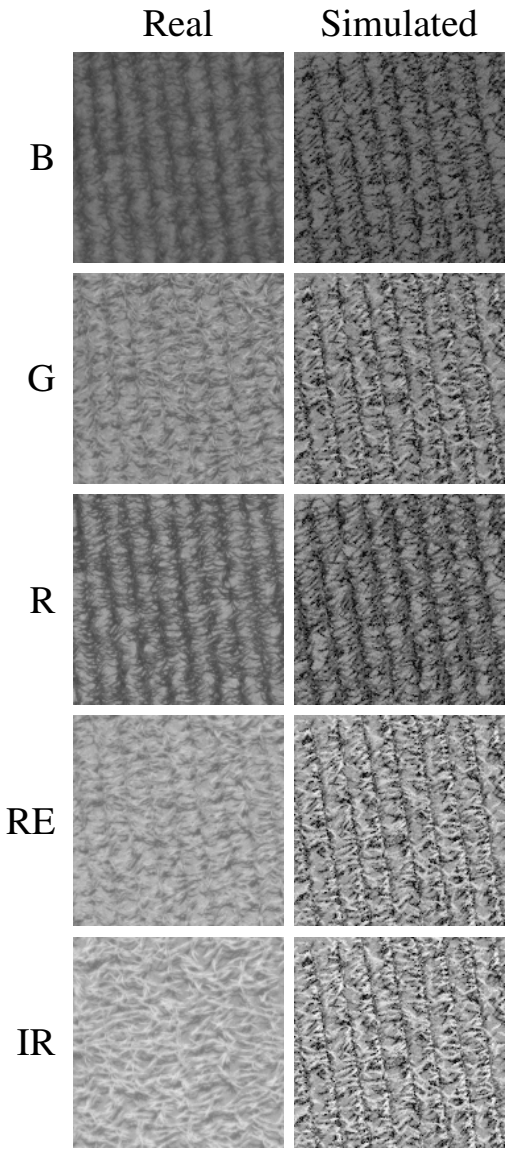


**Figure 2.18:** Per-band comparison of $160 \times 160$ tiles of real MicaSense imagery versus simulated DIRSIG5 imagery. Note that all images have been contrast-stretched with respect to the same min/max radiances so that they are properly visually comparable.

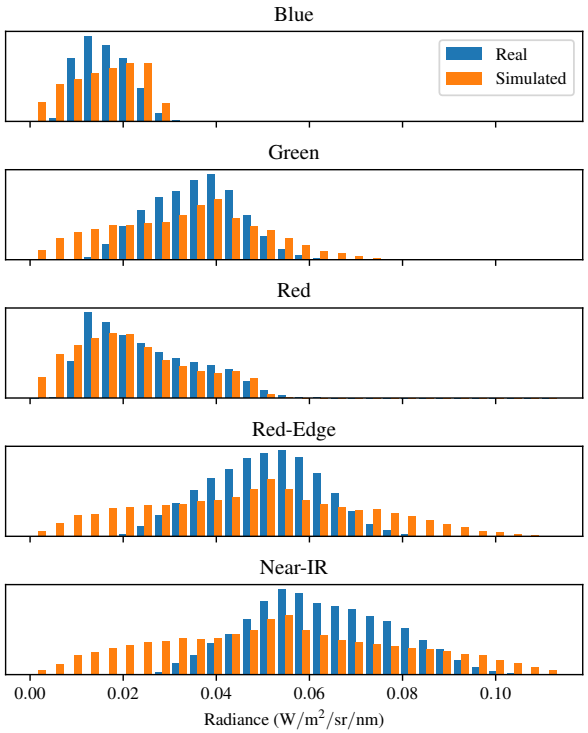


**Figure 2.19:** Per-band radiance histograms for $160 \times 160$ tiles of real MicaSense imagery versus simulated DIRSIG imagery.

.

**Per-band average $(\mu, \sigma)$ statistics**

| | Real | Simulated |
|---|---|---|
| Blue | $(0.0152, 0.0052)$ | $(0.0158, 0.0074)$ |
| Green | $(0.0335, 0.0101)$ | $(0.0316, 0.0153)$ |
| Red | $(0.0226, 0.0101)$ | $(0.0223, 0.0126)$ |
| Red Edge | $(0.0476, 0.0125)$ | $(0.0454, 0.0221)$ |
| Near IR | $(0.0610, 0.0159)$ | $(0.0510, 0.0250)$ |

**Table 2.3:** Per-band average $(\mu, \sigma)$ statistics, being mean $\mu$ and standard deviation $\sigma$ in units of Watts per square meter per steradian per micron, for radiance distributions of $160 \times 160$ image tiles from real and simulated images.

the ground, without any corn plants, from another real image (captured as the UAS took off). Then, for each band, we calculated the average $\chi^2$ distance between 1) all pairings of real tiles with simulated tiles, 2) all pairings of real tiles with other (non-identical) real tiles, 3) all pairings of simulated tiles with other (non-identical) simulated tiles, 4) all pairings of real tiles with the reference ground tile, and 5) all pairings of simulated tiles with the reference ground tile.

The average $(\mu, \sigma)$ statistics are given in table 2.3. The average $\chi^2$ histogram distance metrics are given in table 2.4. Regarding the interpretation of table 2.4, the Real/Sim column measures the shape variability of radiance distributions between the real and simulated tile sets. The Real/Real and Sim/Sim columns measure within-population shape variability of radiance distributions in the real and simulated tile sets respectively. The Real/Ground and Sim/Ground columns provide the baseline $\chi^2$ distances between the real and simulated image tiles and a real image tile containing only the ground. That being the case, if the simulated imagery perfectly mimicked the real imagery, then the Real/Sim column would contain all zeros, the Sim/Sim column would be identical to the Real/Real column, and the Sim/Ground column would be identical to the Real/Ground column.

### 2.5.2. Discussion

In examining table 2.3, we see that there is reasonable concensus in the mean radiance between the real and simulated tiles, though the standard deviation in simulated tiles is wider. The percent error of the mean $\mu$ is within $\pm 6\%$ in all bands except the near IR, where it is $-16.4\%$. The percent error of the standard deviation $\sigma$ is 42.3% in the blue, 51.5% in the green, 24.8% in the red, 76.8% in the red edge, and 57.2% in the near IR. In examining table 2.4, we see that

**Per-band average $\chi^2$ histogram distances**

| | Real/Sim | Real/Real | Sim/Sim | Real/Ground | Sim/Ground |
|---|---|---|---|---|---|
| Blue | 0.133 | 0.054 | 0.024 | 0.993 | 0.995 |
| Green | 0.118 | 0.046 | 0.022 | 0.829 | 0.746 |
| Red | 0.114 | 0.068 | 0.030 | 0.868 | 0.835 |
| Red Edge | 0.162 | 0.042 | 0.023 | 0.234 | 0.367 |
| Near IR | 0.178 | 0.023 | 0.023 | 0.655 | 0.588 |

**Table 2.4:** Per-band average $\chi^2$ histogram distances between $160 \times 160$ image tiles from real and simulated images. Each column "A/B" contains the average distance between unique histogram pairs in groups A and B for each band.

the Real/Sim $\chi^2$ distances are seven to ten times smaller than the baseline Real/Ground and Sim/Ground $\chi^2$ distances in the blue, green, and red bands—thus indicating that the shapes of the radiance distributions of these bands are reasonably close. The Real/Sim $\chi^2$ distances are only two to five times smaller than the baseline distances in the red edge and near IR, which is perhaps consistent with the increased percent error in the standard deviations of these bands. The Sim/Sim $\chi^2$ distance is on the same order of magnitude as the Real/Real $\chi^2$ distance, though it is consistently about half of the Real/Real $\chi^2$ distance in all bands except the near IR, where it is approximately equal. This suggests that there is greater shape variability in the real radiance distributions. The fact that there is simultaneously greater shape variability and lesser standard deviation in the real radiance distributions may indicate that the variability is more spatially dependent. In any case, we found that the simulated imagery is quantitatively plausible, given that the real and simulated scenes are not expected to be completely equivalent, and that the real and simulated images are not expected to be entirely free of confounding sources of error.

*Discrepancies.* The simulated imagery agreed reasonably well with the real imagery, but it is not perfect. It is therefore worth discussing the potential sources of discrepancies. It is always possible that an aspect of the simulation was too simplistic or otherwise flawed. We characterized the ground as a flat Lambertian surface whose reflectance is given by a single SVC measurement. In reality, the ground is neither perfectly Lambertian nor a flat, singular surface. That is to say that portions of the ground may be covered with grass, shrubs, and pebbles, as well as remnants of other plants from no-till agricultural practices. Such textural variability in the ground was not modeled, nor was

textural variability within individual plant leaves. The optical system of the RedEdge was modeled as a pinhole with a rudimentary Point Spread Function (PSF) based on aperture diameter. The real RedEdge optical system is certainly not as ideal. In Figure 2.18, the apparent row structure is nearly lost in the red-edge and near IR in the real image tiles shown. These particular tiles were located toward the edge of the image, where vignetting is the strongest, and the system PSF seems to introduce the most blur. Moreover, we have performed our analysis on radiance images, where the values at each pixel represent the average band-integrated radiances over the corresponding detector element in physical units of power per unit area per unit solid angle. Real imaging systems do not inherently capture radiance, but instead record arbitrary *digital counts* which must be converted to radiance via calibrated equations. Radiometric-correction using such equations introduces another potential source of error. For the RedEdge in particular, there is a standarized radiometric-correction process (used in this analysis) which aims to correct for vignetting and to account for exposure time and sensor gain. This involves scaling each digital count by an expression which depends on metadata recorded in the raw TIFF files and on the coordinate of each pixel. This process is not entirely reliable however. Mamaghani and Salvaggio [59] found that the average percent error in radiance imagery using the "factory default" radiometric correction for a particular RedEdge-3 sensor was $-10.98\%$ in the blue, $-0.43\%$ in the green, $3.59\%$ in the red, $32.81\%$ in the red edge, and $-17.08\%$ in the near IR. The authors further note that this type of calibration error is specific to each sensor, and is subject to change over time as a sensor ages.

*PROSPECT-BSDF versus Lambertian assumption.* As previously stated, we used the PROSPECT-inspired BSDF given in section 2.4 to characterize the directional scattering of leaves in the DIRSIG simulation. While this BSDF is based on widely accepted theoretical footing, it is reasonable to wonder whether the use of this BSDF, instead of a Lambertian BSDF, actually improved the simulation results per tables 2.3 and 2.4. We therefore ran an alternate DIRSIG simulation, which was identical in all aspects except that the PROSPECT-inspired rough plate BSDF was replaced by a simple Lambertian BSDF (still weighted to reproduce standard PROSPECT hemispherical optical properties, in order to preserve the total proportions of energy reflected, transmitted, and absorbed). Using a Lambertian BSDF, we found that the percent errors and Real/Sim $\chi^2$ distances were relatively unchanged in the blue, green, and red bands, but significantly increased in the red

edge and near IR—where the leaves were the most reflective and transmissive. The percent error in the mean radiance $\mu$ was 0.7% in the blue, 3.6% in the green, $-4.4\%$ in the red, 18.7% in the red edge, and 21.0% in the near IR. The percent error in the standard deviation $\sigma$ was 42.3% in the blue, 62.5% in the green, 22.8% in the red, 116.0% in the red edge, and 132.1% in the near IR. Further, the Real/Sim $\chi^2$ distance increased by 0.1% in the blue, 0.1% in the green, 2.5% in the red, 23.8% in the red edge, and 28.1% in the near IR. This seems to suggest that the PROSPECT-inspired BSDF is closer in shape to the true BSDF (at least in the red edge and near IR) than an ordinary Lambertian BSDF. This may also suggest that there is further room for improvement in the leaf scattering model.

We next will conclude with brief remarks and suggested future research efforts.

# Chapter 3

# Conclusions

## 3.1. Closing remarks

We presented a randomizable time-dependent model of corn canopy with application to simulated remote sensing of agricultural scenes. The model is designed to be suitable to generate high fidelity geometry at very near ($< 100\,\text{m}$ above ground) viewing distances, but also to generate lower fidelity geometry efficiently enough to render millions of unique corn plants at very far, satellite-level viewing distances. In doing so, we provided novel innovations to existing modeling approaches in the literature. We parameterized the midrib curve of plant leaves by arclength, such that the basic leaf shape has a well-defined area. We introduced a leaf curl deformation to allow leaves to unfurl from base to tip. We further introduced explicit, visually interesting parameterizations of the sheath and internode surfaces, and we continuously bridged the sheath and lamina surfaces in order to form more visually compelling canopy geometry. We based the time-dependency of the geometry on a physiological model of growth and development, which distills more rigorous models in the literature to an efficient, elegant, and time-continuous framework. We united the PROSPECT model of leaf optical properties spectra with a numerically simulated bidirectional scattering distribution function of a rough dielectric plate in order to characterize leaf scattering. We implemented the model as a C++ library, and we have used this library to implement a DIRSIG5 software plug-in. This plug-in automates the geometric and optical properties modeling of virtual corn plants, in order to make it easier for users to generate and simulate realistic, high-fidelity agricultural scenes. We have further used DIRSIG5 to simulate multispectral radiance imagery of corn plants mimick-

ing real radiometrically-corrected imagery collected by a UAS-mounted MicaSense RedEdge sensor at a USDA site at the Beltsville Agricultural Research Center. We quantitatively compared the simulated imagery to the real imagery by examining normal statistics and $\chi^2$ histogram distances, and found that the spatially-local radiance distributions within the simulated imagery were plausible. In particular, we found that the percent error of the mean radiance was within $\pm 6\%$ for all bands except the near IR, where it was $-16.4\%$. The percent error of the standard deviation varied between 24% (red) and 77% (red edge). The $\chi^2$ distances between the real and simulated radiance histograms, which compare the shapes of the distributions, were seven to ten times smaller than the baseline $\chi^2$ distances (measured to the radiance histogram of the ground, containing no corn plants) in the blue, green, and red bands, and two to five times smaller in the red edge and near IR. We further found that the percent errors and $\chi^2$ distances were reduced (improved) by using the PROSPECT-inspired plate BSDF, instead of the typical Lambertian assumption, in the red edge and near IR bands where leaves are the most reflective and transmissive. There is still room for improvement, however, both regarding the fidelity of the virtual scene, as well as the fidelity of the virtual imaging system.

## 3.2. Future work

There are many potential avenues for the continuation of this research. First, there is much room for improvement regarding the modeling itself. This research has focused only the vegetative stage of development—that is, the period of growth from emergence to the beginning of the reproductive stage. Thus, we have not modeled the development of the tassel (which is initiated in the vegetative stage, but is not visually striking until the reproductive stage), nor have we modeled the development of actual ears of corn, nor have we modeled senescence (plant death following the reproductive stage). Future work may consider extending the modeling approach to account for later growth stages. Future work may also consider investigating sensible relationships between the material concentrations determining leaf optical properties spectra (i.e., chlorophyll $a + b$ and water) and the proposed growth model, in order to form a more unified framework. It would also be worthwhile to incorporate dirurnal cycles, i.e., changes in canopy structure and arrangment over a day in response to light, temperature, as moisture, as an additional layer of complexity to the model. We further recommend more thorough investigation of the Bidirectional Scattering Dis-

tribution Function (BSDF) of corn leaves. We developed the `layered-sqt` BSDF simulation tool to calculate the emergent BSDF of a general layered structure, containing many layers separated by participating media with different optical properties. However, we restricted the theoretical description of the layered structure of a corn leaf to be consistent with the fundamental assumptions of the PROSPECT model. That is, we made as few additional assumptions as possible—the leaf was modeled as a dielectric plate with rough top and bottom surfaces, without any scattering or absorption by a participating medium in between. Future work may consider proposing a more comprehensive description which accounts for, e.g., epicuticular wax and sub-surface scattering by leaf material. Future work may also consider investigating the wavelength-dependence of the parameters of such a description. Regarding the representation of the virtual imaging system, we have not accounted for the system PSF beyond a first approximation from the aperture diameter, nor have we explicitly modeled optical elements in any fashion. We have further limited our quantitative analysis to radiance images, which introduced the radiometric calibration of the real imaging system as another potential source of error. It would be advisable for future work to apply a more robust radiometric correction process, such as that proposed by Mamaghani and Salvaggio [59], to improve the reliability of quantitative analysis. We also note that it is possible to simulate digital counts, the fundamental quantity collected by real imaging systems, directly with DIRSIG5. Alternatively, it is possible to convert DIRSIG5-simulated radiance images to digital counts with a separate program or script. This type of comprehensive "end-to-end" imaging simulation requires detailed information of the detector array however, e.g., quantum efficiency, read noise, dark noise, photo-response non-uniformity (PRNU), and well capacity. We did not possess, and did not attempt to measure, such information for the RedEdge-3. Future work may consider accounting for such information in order to comprehensively simulate, analyze, and compare digital counts instead of radiances.

**Implementation**

We make the C++ implementation freely available to the research community via a GitHub repository:

`github.com/mgradysaunders/labyrinth-v2`

`github.com/mgradysaunders/layered-sqt`

# Appendix A

# Random air temperature

For convenience (or for instances where densely-measured data are not available), we developed a continuous-time model of air temperature. We represent air temperature in degrees Celsius as a function of time $t_d$ as fractional Day-Of-Year (DOY). We mean by "fractional DOY" that $t_d = 1.0$ is 12AM on January 1, $t_d = 1.5$ is 12PM on January 1, $t_d = 2.0$ is 12AM on January 2, and so on. The top-level function, $T_{\text{air}}$, is broken down as the sum of a mean and an offset,

$$T_{\text{air}}(t_d) = \bar{T}_{\text{air}}(t_d) + \Delta T_{\text{air}}(t_d).$$

The mean is thought to oscillate sinusoidally between yearly extrema, denoted by $\bar{T}_{\text{air,min}}$ and $\bar{T}_{\text{air,max}}$ respectively, such that

$$\bar{T}_{\text{air}}(t_d) = (1-\mu)\bar{T}_{\text{air,min}} + \mu\bar{T}_{\text{air,max}}$$

where

$$\begin{aligned}\mu &\leftarrow \frac{1}{2}\sin 2\pi\nu + \frac{1}{2},\\ \nu &\leftarrow 0.0038872 t_d - 0.5411472.\end{aligned}$$

The offset is modeled as a skewed sinusoid with amplitude $\Delta T_{\text{air,amp}}$,

$$\begin{aligned}\Delta T_{\text{air}}(t_d) = \Delta T_{\text{air,amp}}&\\ \times 4.368815 \tan^{-1} &\frac{\sin 2\pi\alpha}{4.368815 - \cos 2\pi\alpha}\end{aligned}$$

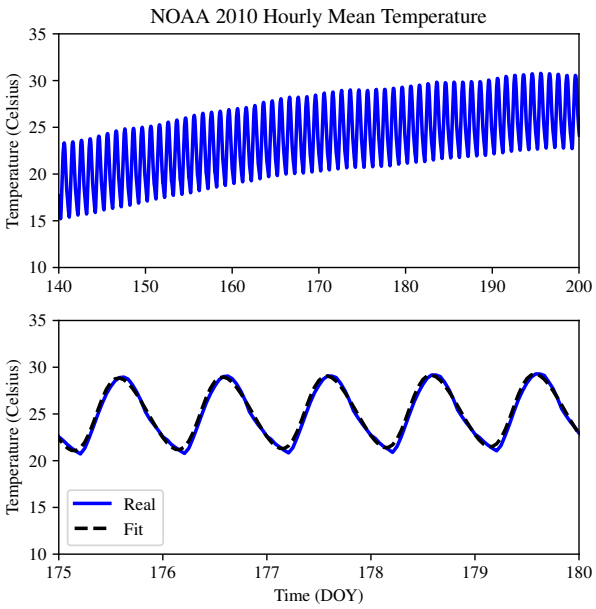


**Figure A.1:** NOAA 2010 hourly mean temperature data from the Washington-Reagan international airport. In the lower plot, our fit with optimized parameters of $\bar{T}_{\text{air,min}} \approx 12.1\,^{\circ}\text{C}$, $\bar{T}_{\text{air,max}} \approx 26.6\,^{\circ}\text{C}$, $\Delta T_{\text{air,amp}} \approx 3.8\,^{\circ}\text{C}$ is shown for a five day period.

where

$$\alpha \leftarrow t_d - 0.360914.$$

We fit the various constants to 2010 NOAA hourly mean temperatures recorded at the Washington Reagan International Airport station over the growing season from May 1 to August 31. As a point of reference, this predicts the average maximum temperature to occur in the early afternoon on July 21. Figure A.1 plots the NOAA data alongside our parametric fit with optimized parameters $\bar{T}_{\text{air,min}} \approx 12.1\,^{\circ}\text{C}$, $\bar{T}_{\text{air,max}} \approx 26.6\,^{\circ}\text{C}$, $\Delta T_{\text{air,amp}} \approx 3.8\,^{\circ}\text{C}$ achieving an MSE of approximately $0.2\,^{\circ}\text{C}$.

# Bibliography


[1] N. J. Stehr, "Drones: The newest technology for precision agriculture," *Journal of natural resources and life sciences education*, vol. 44, no. 1, pp. 89–91, 2015, ISSN: 1059-9053. DOI: 10.4195/nse2015.04.0772.

[2] L. Hoffman, A. Baker, L. Foreman, and E. Young, *Feed grains backgrounder*, Feed Outlook No. (FDS-07C01), U. S. Department of Agriculture, 2007.

[3] T. Capehart, O. Liefert, and D. W. Olson, *Feed outlook: April 2020*, Feed Outlook No. (FDS-20d), U. S. Department of Agriculture, 2020.

[4] *World Agricultural Supply and Demand Estimates*, U. S. Department of Agriculture, Apr. 2020.

[5] A. A. Goodenough and S. D. Brown, "DIRSIG5: Next-generation remote sensing data and image simulation framework," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 10, no. 11, pp. 4818–4833, Nov. 2017.

[6] C. A. van Diepen, J. Wolf, H. van Keulen, and C. Rappoldt, "WOFOST: A simulation model of crop production," *Soil Use and Management*, vol. 5, no. 1, pp. 16–24, 1989. DOI: 10.1111/j.1475-2743.1989.tb00755.x.

[7] A. de Wit, H. Boogaard, D. Fumagalli, S. Janssen, R. Knapen, D. van Kraalingen, I. Supit, R. van der Wijngaart, and K. van Diepen, "25 years of the wofost cropping systems model," *Agricultural Systems*, vol. 168, pp. 154–167, 2019, ISSN: 0308-521X. DOI: 10.1016/j.agsy.2018.06.018. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0308521X17310107.

[8] C. A. Jones, J. R. Kiniry, and P. T. Dyke, *CERES-Maize: A Simulation Model of Maize Growth and Development*. Texas A&M University Press, 1986, ISBN: 9780890962695.

[9] H. Yang, A. Dobermann, J. Lindquist, D. T. Walters, T. J. Arkebauer, and K. Cassman, "Hybrid-maize - a maize simulation model that combines two crop modeling approaches," *Field Crops Research*, vol. 87, pp. 131–154, May 2004. DOI: 10.1016/j.fcr.2003.10.003.

[10] H. Yang, A. Dobermann, K. Cassman, and D. Walters, "Features, applications, and limitations of the hybrid-maize simulation model," *Agronomy Journal - AGRON J*, vol. 98, May 2006. DOI: 10.2134/agronj2005.0162.

[11] H. S. Yang, P. Grassini, K. G. Cassman, R. M. Aiken, and P. I. Coyne, "Improvements to the hybrid-maize model for simulating maize yields in harsh rainfed environments," *Field Crops Research*, vol. 204, pp. 180–190, 2017, ISSN: 0378-4290. DOI: 10.1016/j.fcr.2017.01.019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378429017301727.

[12] H. Honda, "Description of the form of trees by the parameters of the tree-like body: Effects of the branching angle and the branch length on the shape of the tree-like body," *Journal of Theoretical Biology*, vol. 31, pp. 331–338, 1971. DOI: 10.1016/0022-5193(71)90191-3.

[13] P. E. Oppenheimer, "Real time design and animation of fractal plants and trees," in *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '86, New York, NY, USA: Association for Computing Machinery, 1986, pp. 55–64, ISBN: 0897911962. DOI: 10.1145/15922.15892.

[14] P. Prusinkiewicz, L. Mündermann, R. Karwowski, and B. Lane, "The use of positional information in the modeling of plants," in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '01, New York, NY, USA: Association for Computing Machinery, 2001, pp. 289–300, ISBN: 158113374X. DOI: 10.1145/383259.383291.

[15] S. M. Ulam, "On some mathematical properties connected with patterns of growth of figures," in *Proceedings of Symposia on Applied Mathematics 14*, 1962, pp. 215–244.

[16] W. Pałubicki, K. Horel, S. Longay, A. Runions, B. Lane, R. Měch, and P. Prusinkiewicz, "Self-organizing tree models for image synthesis," *ACM Trans. Graph.*, vol. 28, no. 3, Jul. 2009, ISSN: 0730-0301. DOI: 10.1145/1531326.1531364.

[17] T. Sachs and A. Novoplansky, "Tree form: Architectural models do not suffice," *Israel Journal of Plant Sciences*, vol. 43, no. 3, pp. 203–212, 1995. DOI: 10.1080/07929978.1995.10676605.

[18] P. Prusinkiewicz and J. Hanan, *Lindenmayer Systems, Fractals and Plants*. Berlin, Heidelberg: Springer-Verlag, 1989, ISBN: 3540970924.

[19] P. Prusinkiewicz and A. Lindenmayer, *The Algorithmic Beauty of Plants*. Berlin, Heidelberg: Springer-Verlag, 1990, ISBN: 0387972978.

[20] A. Lindenmayer, "Mathematical models for cellular interaction in development, i and ii," *Journal of Theoretical Biology*, vol. 18, pp. 280–315, 1968. DOI: 10.1016/0022-5193(68)90079-9.

[21] C. Fournier and B. Andrieu, "A 3d architectural and process-based model of maize development," *Annals of Botany*, vol. 81, pp. 233–250, 1998. DOI: 10.1006/anbo.1997.0549.

[22] J.-L. Drouet and L. Pagès, "GRAAL: A model of GRowth, Architecture and carbon ALlocation during the vegetative phase of the whole maize plant," *Ecological Modelling*, vol. 165, pp. 147–173, Jul. 2003. DOI: 10.1016/S0304-3800(03)00072-3.

[23] M. L. España, F. Baret, F. Aries, M. Chelle, B. Andrieu, and L. Prévot, "Modeling maize canopy 3D architecture: Application to reflectance simulation," *Ecological Modelling*, vol. 122, no. 1, pp. 25–43, 1999, ISSN: 0304-3800. DOI: 10.1016/S0304-3800(99)00070-8.

[24] L. Prévot, F. Aries, and P. Monestiez, "Modélisation de la structure géométrique du maïs," *Agronomie*, vol. 11, no. 6, pp. 491–503, 1991. [Online]. Available: https://hal.archives-ouvertes.fr/hal-00885391.

[25] S. Jacquemoud and F. Baret, "Prospect: A model of leaf optical properties spectra," *Remote Sensing of Environment*, vol. 34, no. 2, pp. 75–91, 1990, ISSN: 0034-4257. DOI: https://doi.org/10.1016/0034-4257(90)90100-Z.

[26] Y. Zhang, J. M. Chen, J. R. Miller, and T. L. Noland, "Leaf chlorophyll content retrieval from airborne hyperspectral remote sensing imagery," *Remote Sensing of Environment*, vol. 112, pp. 3234–3247, 2008. DOI: 10.1016/j.rse.2008.04.005.

[27] S. Jacquemoud, S. L. Ustin, J. Verdebout, G. Schmuck, G. Andreoli, and B. Hosgood, "Estimating leaf biochemistry using the prospect leaf optical properties model," *Remote Sensing of Environment*, vol. 56, no. 3, pp. 194–202, 2000. DOI: 10.1016/0034-4257(95)00238-3.

[28] J. B. Feret, C. François, G. P. Asner, A. A. Gitelson, R. E. Martin, L. P. R. Bidel, S. L. Ustin, G. le Maire, and S. Jacquemoud, "Prospect-4 and 5: Advances in the leaf optical properties model separating photosynthetic pigments," *Remote Sensing of Environment*, vol. 112, no. 6, pp. 3030–3043, 2008. DOI: 10.1016/j.rse.2008.02.012.

[29] S. Jacquemoud, W. Verhoef, F. Baret, C. Bacour, P. Zarco-Tejada, G. P. Asner, C. François, and S. L. Ustin, "Prospect + sail models: A review of use for vegetation characterization," *Remote Sensing of Environment*, vol. 113, S56–S66, 2009. DOI: 10.1016/j.rse.2008.01.026.

[30] G. V. Baranoski and J. G. Rokne, *Light Interaction with Plants: A Computer Graphics Perspective*. Woodhead Publishing, 2004, ISBN: 1904275117.

[31] T. Miao, C. Zhao, X. Guo, and S. Lu, "A framework for plant leaf modeling and shading," *Mathematical and Computer Modeling*, vol. 58, pp. 710–718, 2011. DOI: 10.1016/j.mcm.2011.10.027.

[32] E. J. Ientilucci, S. D. Brown, J. R. Schott, and R. V. Raqueno, "Multispectral simulation environment for modeling low-light-level sensor systems," in *Image Intensifiers and Applications; and Characteristics and Consequences of Space Debris and Near-Earth Objects*, C. B. Johnson, T. D. Maclay, and F. A. Allahdadi, Eds., International Society for Optics and Photonics, vol. 3434, SPIE, 1998, pp. 10–19. DOI: 10.1117/12.331229.

[33] J. Schott, S. Brown, R. Raqueño, H. Gross, and G. Robinson, "An advanced synthetic image generation model and its application to multi/hyperspectral algorithm development," *Canadian Journal of Remote Sensing*, vol. 25, no. 2, pp. 99–111, 1999. DOI: 10.1080/07038992.1999.10874709.

[34] E. J. Ientilucci and S. D. Brown, "Advances in wide-area hyperspectral image simulation," in *Targets and Backgrounds IX: Characterization and Representation*, W. R. Watkins, D. Clement, and W. R. Reynolds, Eds., International Society for Optics and Photonics, vol. 5075, SPIE, 2003, pp. 110–121. DOI: 10.1117/12.488706.

[35] W. Yao, D. Kelbe, M. van Leeuwen, P. Romanczyk, and J. van Aardt, "Towards an improved LAI collection protocol via simulated and field-based PAR sensing," *Sensors*, vol. 7, p. 1092, Jul. 2016. DOI: 10.3390/s16071092.

[36] W. Yao, J. van Aardt, M. van Leeuwen, D. Kelbe, and P. Romanczyk, "A simulation-based approach to assess subpixel vegetation structural variation impacts on global imaging spectroscopy," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 7, pp. 4149–4164, 2018.

[37] S. D. Brown, D. D. Blevins, and J. R. Schott, "Time-gated topographic LIDAR scene simulation," in *Laser Radar Technology and Applications X*, G. W. Kamerman, Ed., International Society for Optics and Photonics, vol. 5791, SPIE, 2005, pp. 342–353. DOI: 10.1117/12.604326.

[38] D. D. Blevins, S. D. Brown, and J. R. Schott, "First-principles based LIDAR simulation environment for scenes with participating mediums," in *Laser Radar Technology and Applications XI*, G. W. Kamerman and M. D. Turner, Eds., International Society for Optics and Photonics, vol. 6214, SPIE, 2006, pp. 141–152. DOI: 10.1117/12.665958.

[39] P. Romanczyk, J. van Aardt, K. Cawse-Nicholson, D. Kelbe, J. McGlinchy, and K. Krause, "Assessing the impact of broadleaf tree structure on airborne full-waveform small-footprint LiDAR signals through simulation," *Canadian Journal of Remote Sensing*, vol. 39, no. sup1, S60–S72, 2013. DOI: 10.5589/m13-015.

[40] M. Gartley, A. Goodenough, S. Brown, and R. Kauffman, "A comparison of spatial sampling techniques enabling first principles modeling of a synthetic aperture radar imaging platform," *Proc. SPIE*, vol. 7699, Apr. 2010. DOI: 10.1117/12.849552.

[41] E. Heitz, "Understanding the masking-shadowing function in microfacet-based brdfs," *Journal of Computer Graphics Techniques (JCGT)*, vol. 3, no. 2, pp. 48–107, Jun. 2014, ISSN: 2331-7418. [Online]. Available: http://jcgt.org/published/0003/02/03/paper.pdf.

[42] E. Heitz, J. Hanika, E. d'Eon, and C. Dachsbacher, "Multiple-scattering microfacet bsdfs with the smith model," *ACM Transactions on Graphics*, vol. 35, no. 4, pp. 1–58, Jul. 2016, ISSN: 0730-0301. DOI: 10.1145/2897824.2925943. [Online]. Available: https://doi.org/10.1145/2897824.2925943.

[43] W. K. Silk and R. O. Erickson, "Kinematics of plant growth," *Journal of Theoretical Biology*, vol. 76, no. 4, pp. 481–501, 1979. DOI: 10.1016/0022-5193(79)90014-6.

[44] I. J. Warrington and E. T. Kanemasu, "Corn growth response to temperature and photoperiod, iii. leaf number," *Agronomy Journal*, vol. 75, pp. 762–766, 1983. DOI: 10.2134/agronj1983.00021962007500050010x.

[45] R. Bonhomme and C. Varlet-Grancher, "Estimation of the gramineous crop geometry by plant profiles including leaf width variations," *Photosynthetica*, vol. 12, pp. 193–196, 1978.

[46] H. B. H. Salah, "Analyse de la vitesse de croissance foliaire du maïs (zea mays l.) en conditions climatiques fluctantes. conséquences de deficits hydriques dans le sol.," PhD thesis, Institut National Agronomique Paris Grignon, 1996.

[47] J. Cao, J. D. Hekseth, B. Zur, and J. F. Reid, "Leaf area development in maize and soybean plants," *Biotronics*, vol. 17, pp. 9–15,

[48] B. Zur, J. F. Reid, and J. D. Hekseth, "The dynamics and mazie canopy development, 1. leaf ontogeny," *Biotronics*, vol. 18, pp. 55–66,

[49] B. C. Sharman, "Developmental anatomy of the shoot of zea mays l," *Annals of Botany*, vol. 6, pp. 246–282, 1942.

[50] J. D. Hesketh, I. J. Warrington, J. F. Reid, and B. Zur, "The dynamics of corn canopy development: Phytomer ontogeny," *Biotronics*, vol. 17, pp. 69–77, 1988.

[51] R. F. Grant and J. D. Hesketh, "Canopy structure of maize (zea mays l.) at different populations: Simulation and experimental verification," *Biotronics*, vol. 21, pp. 11–24, 1992.

[52] M. J. Robertson, "Relationships between internode elongation, plant height, and leaf appearance in maize," *Field Crops Research*, vol. 38, pp. 135–145, 1994.

[53] M. E. O'Neill, "Pcg: A family of simple fast space-efficient statistically good algorithms for random number generation," Harvey Mudd College, Claremont, CA, Tech. Rep. HMC-CS-2014-0905, Sep. 2014.

[54] R. L. Cook and K. E. Torrance, "A reflectance model for computer graphics," *ACM Trans. Graph.*, vol. 1, no. 1, pp. 7–24, Jan. 1982, ISSN: 0730-0301. DOI: 10.1145/357290.357293. [Online]. Available: https://doi.org/10.1145/357290.357293.

[55] B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, "Microfacet models for refraction through rough surfaces," in *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, ser. EGSR'07, Grenoble, France: Eurographics Association, 2007, pp. 195–206, ISBN: 9783905673524.

[56] Q. Ling, W. Huang, and P. Jarvis, "Use of a SPAD-502 meter to measure leaf chlorophyll concentration in arabidopsis thaliana," *Photosynthesis Research*, vol. 107, no. 2, pp. 209–14, 2011. DOI: 10.1007/s11120-010-9606-0.

[57] Z. G. Cerovic, G. Masdoumier, N. B. Ghozlen, and G. Latouche, "A new optical leaf-clip meter for simultaneousnon-destructive assessment of leaf chlorophylland epidermal flavonoids," *Physiologia plantarum*, vol. 146, no. 3, pp. 251–260, 2012, ISSN: 0031-9317. DOI: 10.1111/j.1399-3054.2012.01639.x.

[58] S. Gustavson, "Simplex noise demystified," Jan. 2005. [Online]. Available: http://staffwww.itn.liu.se/~stegu/simplexnoise/simplexnoise.pdf.

[59] B. Mamaghani and C. Salvaggio, "Multispectral sensor calibration and characterization for sUAS remote sensing," *Sensors*, vol. 19, 20 2019. DOI: 10.3390/s19204453.

[60] E. Myers, J. Kerekes, C. Daughtry, and A. Russ, "Assessing the impact of satellite revisit rate on estimation of corn phenological transition timing through shape model fitting," *Remote Sensing*, vol. 11, 21 2019. DOI: 10.3390/rs11212558.