Three-dimensional modeling of water use and photosynthesis in almond orchards

Project No.: PREC1.Bailey

Project Leader:

Brian Bailey
Dept. of Plant Sciences, UC Davis, One Shields Ave., Davis, CA 95616
Phone: (530) 752-7478
Email: bnbailey@ucdavis.edu

Project Cooperators and Personnel:

- Ken Shackel, Dept. of Plant Sciences, UC Davis
- Bruce Lampinen, Dept. of Plant Sciences, UC Davis
- Ted DeJong, Dept. of Plant Sciences, UC Davis
- Matthew Gilbert, Dept. of Plant Sciences, UC Davis
- Eric Kent, Dept. of Plant Sciences, UC Davis

A. Summary

A computer model that is based on three-dimensional plant geometry was adapted for simulating water use of almond orchards. This model allows light interception and water use to be calculated for each leaf. LiDAR measurements that represent the three-dimensional geometry of the orchard were used to reconstruct trees of three varieties (Nonpareil, Monterey, Wood Colony) in the model. The model was tested using lysimeter measurements of tree water use taken at Kearney Agricultural Research and Extension Center. The three varieties had differing spatial footprints and leaf area densities which affected modeled water use. Typical Wood Colony trees used roughly half the water as Nonpareil trees, with Monterey lying in the middle. Results suggested that irrigating based on the water needs of the largest trees results in substantial over-watering of other varieties, and that for the orchard examined in this work, irrigating based on the actual needs of the variety could save around 20% of the water. The relationship between fraction of shaded ground area and modeled whole-tree transpiration matched well with a previously reported relationship based on measurements, suggesting that fractional ground cover (or percent light interception) was a robust predictor of water use even at the individual tree level. However, results also indicated that certain growth patterns could cause deviation from this relationship, suggesting the need for future work examining how tree architecture impacts water-use efficiency with respect to yield.

B. Objectives

The overall objective of this project was to develop a simulation platform that serves as a tool to perform virtual experiments in almond orchards in order to better understand the effects of various orchard designs and management decisions on water use. Specific objectives are listed below:

1. Adapt the Helios simulation system to almond orchards in order to develop an efficient, three-dimensional modeling system that can accurately predict water transport, microclimate, radiation interception, and photosynthesis.
2. Perform model calibration, verification, and validation.
3. Use the model to examine a range of canopy architectures and quantify trade-offs between water usage and photosynthesis

C. Results and Discussion

Objective 1.

Several important developments were incorporated into the Helios modeling framework (Bailey, 2019) as a part of this project including:

1. A reverse ray-tracing module for accurate radiation interception simulations (Bailey, 2018),
2. A method for 3D leaf-level reconstruction of tree architecture from terrestrial LiDAR scanning data (Bailey and Ochoa, 2018),
3. Implementation and testing of sub-divided leaf patches which can be necessary for simulations of non-linear processes like photosynthesis and transpiration in dense canopies (Kent and Bailey, in review),
4. Implementation and testing of several methods for processing LiDAR point cloud data to more accurately estimate leaf area density using full-waveform LiDAR scanners (Kent and Bailey, in preparation).

The output of this objective was an improved version of the Helios three-dimensional modeling framework capable of modeling almond canopies at scales ranging from sub-leaf to orchard (Fig. 1). It serves as the basis for the simulations conducted for this project and for future model development into areas such as growth and yield. More information on Helios and the model code is available online (https://baileylab.ucdavis.edu/software/index.html).

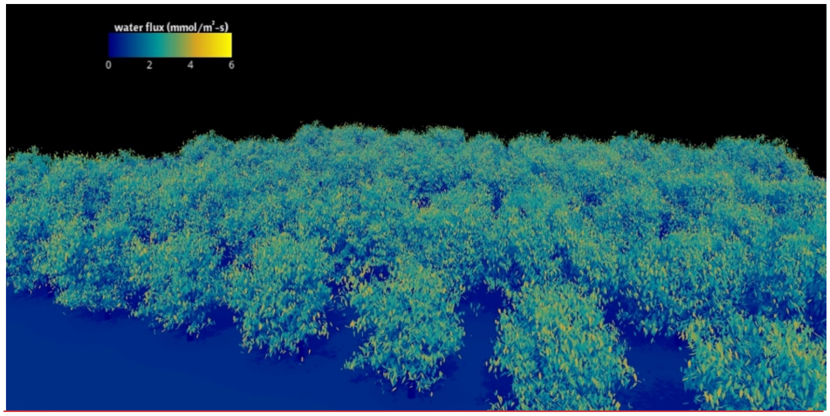

Figure 1. Visualization of model almond orchard canopy reconstructed from LiDAR data, with leaves colored according to simulated transpiration.

Objective 2.

Calibrating and testing of the model was performed based on measurements from the Kearney Agricultural Research and Extension Center, which utilized a weighing lysimeter to provide accurate tree-level water use. Experimental campaigns were conducted during the 2018 and 2019 growing seasons. Initial comparisons showed simulations that overestimated total ET and shifted peak ET later in the afternoon compared with measurements. It was hypothesized that model averaging scale, unrepresentative meteorological input data, overestimation of LiDAR derived reconstruction leaf area, and/or unrepresentativeness of stomatal conductance model parameters could be contributing factors to the observed bias. Each of these factors were investigated, and the total leaf area derived from LiDAR reconstruction as well as the stomatal conductance parameters were found to have the most impact on modeled transpiration. After improving LiDAR reconstructions based on full-waveform LiDAR (Kent and Bailey; in preparation), better agreement was obtained between simulated hourly transpiration rates and measurements from the lysimeter (Fig. 2), although some uncertainty remains. In the future, comparing LiDAR reconstructed geometries with data based on destructive samples of almond trees, as well as exploring other methods of canopy reconstruction would be useful in narrowing potential uncertainty in the canopy representation.

In terms of stomatal conductance parameters, a long sampling time was needed to obtain measurements based on only a limited number of leaves. There is a question of how representative model parameters derived from this limited data are for an entire tree or orchard and how much these parameters vary over time. Applying parameters derived from measurements in one period to other periods did not necessarily result in simulated results in good agreement with the lysimeter data. Our tests and conclusions are therefore somewhat limited in temporal extent.

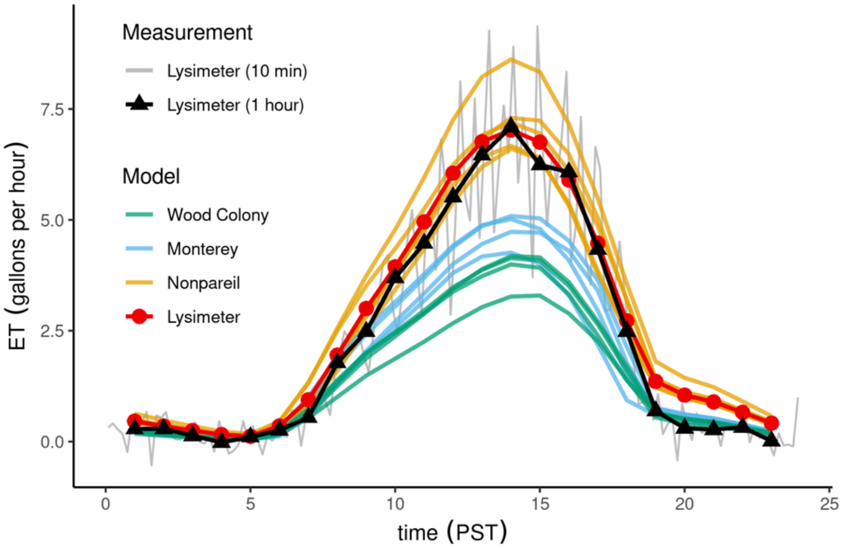

Figure 2. Measured and modeled transpiration over the course of a representative single day (2019-08-02) for multiple trees of the three different varieties at the Kearney lysimeter orchard. Index of agreement between measurements and model results for the lysimeter tree is 0.93.

Objective 3.

The three varieties scanned with the LiDAR at Kearney tended to have different architectures and consequently differences in light interception and water-use. Nonpareil tended to have the largest canopy volume and highest leaf area density, while Monterey had intermediate volume and Wood Colony had the smallest canopy volume. These differences in canopy architecture created significant differences in tree water use between cultivars (Fig. 2). A typical Nonpareil tree used roughly double the amount of water as a typical Wood Colony tree, with Monterey trees lying in the middle. There was some variability in water use within each variety, but these differences were small on average compared to differences between cultivars. According to model predictions, the lysimeter tree ET was representative of the ET of other surrounding Nonpareil trees.

The fraction of ground shaded at mid-day has previously been shown to be correlated with yield (Lampinen et al., 2014) and has also been used to estimate water use of young orchards relative to mature orchards (Schwankl and Prichard, n.d.). A similar approach was tested for this project in which the fraction of shaded ground area per tree in the model was compared with the simulated daily ET relative to trees with 80% fractional ground cover. Simulated results (Fig. 3) were in good agreement with previously published relationships. The linear regression slope for model predictions was 1.28, while a slope of 1.26 was reported by Schwankl and Prichard. The rate of ET for Monterey trees tended to be slightly lower than the expected value (i.e., lower ET per fraction shaded ground area. This may have been due to the slightly different growth pattern of the Monterey trees. This variety tended to grow more horizontally compared to the relatively spherical shape of Wood Colony and Nonpareil varieties, which could increase the ground cover fraction relative to ET.

The above results have several important practical implications. Irrigation decisions are typically made based on the largest trees in the orchard, but results suggest that this leads to substantial over-irrigation of the smaller varieties. For this experimental orchard, about 25% of the trees were Wood Colony, 25% were Monterey, and 50% were Nonpareil. Results suggest that irrigating based on actual tree water demand rather than the demand of the Nonpareil trees would result in a water savings of roughly 20%. Additionally, results suggest that fractional ground cover was a relatively robust indicator of water demand. However, the Monterey variety suggested that certain tree architectures could deviate from that trend, ultimately leading to less ET per fractional ground cover. Whether that reduction in ET results in a corresponding decrease in yield requires further investigation.

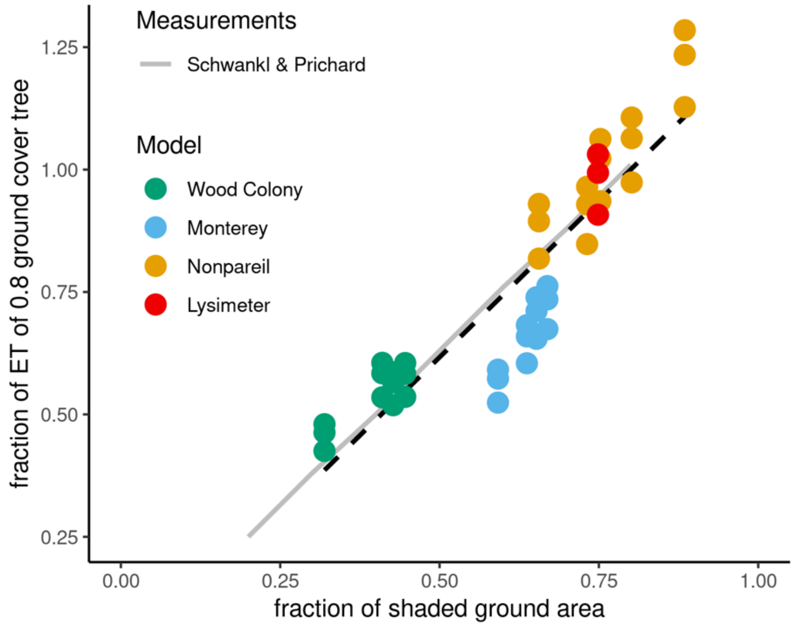

Figure 3. Fraction of shaded ground area plotted against simulated daily total ET as a fraction of ET simulated for tree with 0.8 ground cover fraction. Individual points represent results for individual trees on individual days during a simulation spanning 7/3 to 8/2 2019. The dashed black line shows the linear regression line fitting modeled data. The solid gray line represents the relationship presented by Schwankl and Prichard (n.d.).

D. Outreach Activities

- An oral presentation based on work conducted for this project was given at the American Geophysical Union Fall Meeting on 2020-12-12 in San Francisco. The presentation was part of the session, “Advances in Remote Sensing Theory on Vegetation: Radiative Transfer Modeling, Inversion, and Applications II” and the audience consisted of approximately 50 researchers. Citation: Kent, E. R. and Bailey, B. N. American Geophysical Union Fall Meeting (2020). Remote virtual meeting. “Can estimation of leaf area density from terrestrial LiDAR scanning data be further improved by using additional information provided by full-waveform instruments?” (oral presentation).

- A poster describing this project was presented at the Almond Board of California Conference on 2019-12-10 in Sacramento. The audience included interested growers and researchers.

- An oral presentation based on work conducted for this project was given at the American Geophysical Union Fall Meeting on 2019-12-12 in San Francisco. The presentation was part of the session, “Trends in Digital Agriculture: Crop Sensing Technologies Toward Climate-Resilient Agricultural Systems” and the audience consisted of approximately 50-100 researchers. Citation: Kent, E. R. and Bailey, B. N. American Geophysical Union Fall Meeting (2019). San Francisco, CA. “The Role of Averaging Scale in Modeling Non-linear Leaf Level Processes for Canopies using a Three-dimensional Leaf-resolving Model” (oral presentation).

E. Materials and Methods:

Overall Approach

This work adapted the Helios simulation system to almond orchards. Helios integrates state-of-the-art models for radiation, microclimate, evapotranspiration, and photosynthesis to represent transport processes of interest at scales ranging from leaves to canopies. By explicitly representing every leaf in the canopy using mechanistic physics-based models, the modeling system is robust to changes in architecture.

Study Site and Collection of Field Data

An experimental orchard at the Kearney Agricultural Research and Extension Center in Parlier, CA was used to collect measurements for use as model input and to test the simulation results. The orchard was planted in 2015, had a row spacing of 21 ft, and a tree spacing of about 13 ft, with trees offset from each other in neighboring rows. Nonpareil trees were planted in every other row, with alternating rows of Wood Colony and Monterey. Lysimeter data, which is a direct measure of evapotranspiration by a single almond tree in the orchard, was provided by Ken Schackel and was used to test the model simulations. Measurements were collected as 10 minute averages and a rolling mean with an hour long window was applied to smooth out short-time scale variability (Fig. 2).

While lysimeter data was collected continuously throughout the season, other measurements were collected on several intensive measurement days in 2018 and 2019. A terrestrial full-waveform LiDAR instrument (VZ-1000, RIEGL Inc. Vienna, Austria) was mounted on a wagon and was used to scan a sub-section of the Kearney orchard. Each tree was scanned from four different positions. One challenge encountered was accurately aligning all scans to the same coordinate system. The automatic algorithm typically employed often had trouble aligning scans taken in different rows given the distance and density of vegetation between scan positions. Several approaches were applied to improve alignment, including increasing the number of scans and distributing reflectors throughout the orchard. Manual alignment was performed when necessary, using a meteorological tower in the orchard as a reference. Most simulations run used LIDAR data from the last day of intensity measurements (2019-09-13) when our approach was more refined. Each tree was divided into a series of approximately cubic volumes (about 0.5 m in each direction) and leaf area density (LAD) was estimated for each volume using LiDAR measurements processed with a method published by Grau et al. (2017) and investigated by Kent and Bailey (in preparation). Each tree was then reconstructed using the method of Bailey and Ochoa (2018).

Leaf gas exchange was measured using portable photosynthesis systems (LI-6800, LiCOR Lincoln, NE, USA). Light, temperature, and humidity levels were varied and the leaf stomatal conductance was measured. These data were used to determine coefficients for the stomatal conductance model (Buckley et al. 2012). Simulations reported here all used the same unified set of parameters for all three tree varieties, based on measurements mostly taken on leaves from the Nonpareil lysimeter tree. The long time needed to take light response measurements from a few leaves on this single tree precluded enough measurements from Monterey and Wood Colony trees to evaluate any difference in response between the three varieties. Thus, we note that differences in ET modeled for trees of each variety reported here are due solely to the canopy structure in the model. In the orchard, it is possible the different varieties could have different stomatal responses that could alter their relative ET.

Other measurements collected on intensive measurement days included manual measurements of leaf size, radiative leaf properties (i.e. reflectivity and transmissivity), stem and leaf water potential, and meteorological measurements within the canopy.

Model Simulations

Simulations were run for a sub-region of the Kearney orchard spanning four rows (two Nonpareil, one Wood Colony, and one Monterey) and 4-5 trees within each row (including the lysimeter tree). The percentage of shaded ground area per tree was estimated by sending a known flux of radiation from directly overhead in the model and determining the percentage of that radiation flux absorbed by the ground patch below each tree (with dimensions of the tree and row spacing).

Although we found that modeling non-linear processes at smaller sub-leaf scales can improve accuracy of total canopy fluxes in some cases (Kent and Bailey, in review), initial tests on the model lysimeter tree showed differences in ET occurring primarily near solar noon that were relatively small compared with the uncertainty associated with the tree reconstructions and stomatal conductance model parameters. Consequently, the simulations reported here were carried out at the leaf scale instead of the sub-leaf scale to reduce model run time and enable more trees to be simulated.

The model was primarily run for the intensive measurement days, but days before or after each measurement day when no irrigation occurred were also simulated using the same canopy geometry and stomatal conductance model parameters as on the measurement days. Hourly averaged air temperature, relative humidity, and wind speed from the nearby Parlier CIMIS station were used as model input for each simulation.

F. Publications that emerged from this work

Kent, E. R. and Bailey, B. N. (In preparation). Impact of idealized intensity-based weighting of Multi-echo Terrestrial LiDAR and Partial Misses on Leaf Area Density Estimation.

Kent, E. R. and Bailey, B. N. (In preparation). The role of averaging scale in modeling non-linear leaf level processes for heterogeneous canopies.

Kent, E. R. and Bailey, B. N. (In review). The role of spatial averaging scale in leaf-to-canopy scaling of non-linear processes in homogeneous canopies. Journal of Geophysical Research: Biogeosciences.

Bailey, B.N. (2019). Helios: a scalable 3D plant and environmental biophysical modelling framework. Frontiers in Plant Science 10:1185

Bailey, B.N. (2018). A reverse ray-tracing method for modelling the net radiative flux in leaf-resolving plant canopy simulations. Ecological Modelling 368:233-245

Bailey, B.N. and Ochoa, M.H. (2018). Semi-direct tree reconstruction using terrestrial LiDAR point cloud data. Remote Sensing of Environment 208:133-144.

Other References Cited:

Bailey, B.N. and Mahaffee, W.F. (2017b). Rapid, high-resolution measurement of leaf area and leaf orientation using terrestrial LiDAR scanning data. Measurement Science and Technology 28: 064006.

Bailey, B.N. and Mahaffee, W.F. (2017a). Rapid measurement of the three-dimensional distribution of leaf orientation and the leaf angle probability density function using terrestrial LiDAR scanning. Remote Sensing of Environment 193: 63-76.

Buckley, T. N., Turnbull, T. L., and Adams T. L. (2012). Simple models for stomatal conductance derived from a process model: Cross-validation against sap flux data. Plant Cell Environment., 35: 1647-62.

Grau, E., Durrieu, S., Fournier, R., Gastellu-Etchegorry, J.-P., Yin, T. (2017). Estimation of 3D vegetation density with terrestrial laser scanning data using voxels. a sensitivity analysis of influencing parameters. Remote Sensing of Environment, 191: 373 – 388.

Lampinen, B. D. et al. (2014). A Second Generation Mobile Platform for Assessing Midday Canopy Photo-synthetically Active Radiation Interception in Orchard Systems. Acta Hort. 1058: 105–112.

Schwankl, L., Prichard, T. (n.d.) Almond Irrigation Improvement Continuum. Almond Board of California <https://www.almonds.com/sites/default/files/almond_irrigation_improvement_continuum_1.0[1].pdf>

frontiers
in Plant Science

ORIGINAL RESEARCH
published: 18 October 2019
doi: 10.3389/fpls.2019.01185

Helios: A Scalable 3D Plant and Environmental Biophysical Modeling Framework

Brian N. Bailey*

Department of Plant Sciences, University of California, Davis, Davis, CA, United States

This article presents an overview of Helios, a new three-dimensional (3D) plant and environmental modeling framework. Helios is a model coupling framework designed to provide maximum flexibility in integrating and running arbitrary 3D environmental system models. Users interact with Helios through a well-documented open-source C++ API. Version 1.0 comes with model plug-ins for radiation transport, the surface energy balance, stomatal conductance, photosynthesis, solar position, and procedural tree generation. Additional plug-ins are also available for visualizing model geometry and data and for processing and integrating LiDAR scanning data. Many of the plug-ins perform calculations on the graphics processing unit, which allows for efficient simulation of very large domains with high detail. An example modeling study is presented in which leaf-level heterogeneity in water usage and photosynthesis of an orchard is examined to understand how this leaf-scale variability contributes to whole-tree and -canopy fluxes.

Keywords: biophysical model, functional-structural plant model, software architecture, terrestrial LiDAR, three-dimensional model

INTRODUCTION

Biophysical processes in plant and environmental systems traverse an extraordinary range of spatial and temporal scales, with high heterogeneity commonly present across these scales. In plant ecosystems, this is particularly true, as important effects of heterogeneity have been frequently reported across the full range of scales from cells up through the globe (e.g., Mott and Buckley, 2000; Valladares, 2003). Often, it is convenient to study plant systems at scales most relevant to humans—leaves to canopies in space and seconds to months in time. Obtaining observations beyond these scales often requires high effort that may yield little additional useful information. However, it is clear that heterogeneity across scales can have significant impacts on exchanges of mass, momentum, and energy, and understanding how heterogeneity augments transport processes is key in understanding links between plant structure and function.

To circumvent limitations in our ability to observe plant systems across the entire range of relevant scales, it is common to use mathematical models to translate information obtained at one scale to another scale of interest where data are lacking. In order to do so, assumptions of homogeneity are typically made over a certain range of scales. The earliest, and still most frequently used, class of models of plant systems assumes homogeneity in horizontal directions, thus effectively treating a plant canopy as a "big leaf" (Sinclair et al., 1976; Raupach and Finnigan, 1988; Amthor, 1994; Friend, 2001). In some cases, homogeneity is assumed in all

Edited by:
Hartmut Stützel,
Leibniz University Hannover, Germany

Reviewed by:
Gerhard Buck-Sorlin,
Agrocampus Ouest, France
Jari Perttunen,
Natural Resources Institute Finland (Luke), Finland

***Correspondence:**
Brian N. Bailey
bnbailey@ucdavis.edu

Specialty section:
This article was submitted to Plant Biophysics and Modeling, a section of the journal Frontiers in Plant Science

Received: *03 May 2019*
Accepted: *29 August 2019*
Published: *18 October 2019*

Citation:
Bailey BN (2019) Helios: A Scalable 3D Plant and Environmental Biophysical Modeling Framework. Front. Plant Sci. 10:1185. doi: 10.3389/fpls.2019.01185

directions including the vertical, which is convenient because it means that a measurement or model prediction at any point in space can be considered representative of the entire plant system. A class of big-leaf models called "multilayer models" accounts for vertical heterogeneity by limiting assumptions of homogeneity to a discrete vertical level of vegetation (Meyers and Paw U, 1987; Baldocchi and Harley, 1995). As a compromise between the single big-leaf and multilayer approaches, two-leaf models have also been developed that assume that leaves are either sunlit or shaded, thus effectively limiting model calculations to two-leaf layers (DePury and Farquhar, 1997; Wang and Leuning, 1998).

Although assumptions of large-scale homogeneity are convenient in translating observations and understanding across scales, in nature these assumptions are frequently violated. The current generation of plant system models has tended toward a high-resolution, three-dimensional (3D) approach that explicitly resolves heterogeneity in plant structure at scales of individual plants or smaller (Wang and Jarvis, 1990; Pearcy and Yang, 1996; Sinoquet et al., 2001; Allen et al., 2005; Dauzat et al., 2007; Hemmerling et al., 2008; Pradal et al., 2008; Evers et al., 2018). Early 3D models began by discretizing canopies into individual plants, which allows for the representation of heterogeneity in plant shape, size, and arrangement (e.g., Wang and Jarvis, 1990; Cescatti, 1997). Advances in computational power have enabled more detailed models that discretize plants into homogeneous volumes at submeter scales (e.g., Sinoquet et al., 2001; Bailey et al., 2014; Bailey et al., 2016) or models that resolve individual leaves (e.g., Pearcy and Yang, 1996; Allen et al., 2005; Dauzat et al., 2007; Hemmerling et al., 2008; Pradal et al., 2008).

There is no model that is ideally suited for all applications, and each of the models introduced above makes many trade-offs that are suitable for the particular system and phenomena of interest. A few important trade-offs in plant systems models are discussed below:

Model complexity vs. computational expense. Increases in model complexity generally incur corresponding increases in computational expense. Simple models like the "big-leaf" approach described above are very computationally efficient, and thus they can be used to simulate extremely large problems such as global ecosystem fluxes (Churkina et al., 2005; Reichstein et al., 2005; Lawrence et al., 2019). However, errors and biases can be sizable if subcanopy heterogeneity plays a significant role in the biophysical processes of interest (Friend, 2001; Ponce de León and Bailey, 2019). Models that resolve plant-level heterogeneity often incur a significant computational cost, but simulations are usually limited to domain sizes with a few dozen large plants (Duursma and Medlyn, 2012; Vezy et al., 2018). Leaf-resolving models incur yet another step increase in cost and usually limit the maximum domain size to a dozen or fewer plants depending on plant size and overall model complexity (Hemmerling et al., 2008; Pradal et al., 2008; Kahlen and Stützel, 2011; Woods et al., 2018).

Ease of use vs. flexibility. Providing users with more control over software configuration and execution typically comes with the trade-off of decreasing ease of use (Holzinger, 2005). By automating many tedious or technical tasks, developers can design software that can be readily utilized by inexperienced users. However, for more advanced users who may wish to use the software in ways not originally envisioned by the developers, this can create severe limitations. In the context of plant models, model coupling and execution are often not sequential. For example, if one wishes to simulate photosynthesis of a leaf, this process is cyclically dependent on a number of other processes; photosynthetic rates are dependent on leaf temperature, which is dependent on latent cooling as mediated by stomatal conductance as well as longwave emission, which is also dependent on the leaf temperature. Coupling of the above processes in a model often requires iteration, which can require flexibility if incorporated within a generalized modeling framework. This issue was discussed by Pradal et al. (2008) in the context of the development of the OpenAlea plant modeling framework, which increases ease of use by compromising some flexibility in terms of its execution model. Most 3D plant growth modeling frameworks use a linear work flow in which the execution of various submodules is predefined in order to produce a standard set of outputs (Hemmerling et al., 2008; Pradal et al., 2008; Henke et al., 2016).

Choice of programming language also has important implications in terms of this trade-off. In order to improve ease of use, many modeling frameworks choose to utilize simpler yet less efficient languages such as Python or Java that may not require an explicit compilation step or memory management (Hemmerling et al., 2008; Pradal et al., 2008; Boudon et al., 2012; Henke et al., 2016). Other frameworks have transitioned toward more efficient and flexible languages such as C++ at a sacrifice in usability (Karwowski and Prusinkiewicz, 2003).

Model complexity vs. availability of input data. Increasingly complex models require increasingly complex inputs, and often progress in model development outpaces the development of methods for specifying detailed model inputs. In some cases, models originally built on a solid mechanistic foundation can essentially become overfitted empirical models when inputs turn into free parameters that cannot be measured (Ginzburg and Jensen, 2004). Thus, the development of detailed predictive models is frequently limited by the ability to provide them with realistic inputs, and the argument could therefore be made that in some cases simpler models may be more practical (Raupach and Finnigan, 1988).

This work introduces the new 3D plant and environmental modeling framework "Helios," which is differentiated from other available frameworks in terms of the way in which the above trade-offs are prioritized. First, Helios is a flexible modeling framework that allows for efficient and extensible coupling between arbitrary submodels called plug-ins. Unlike most previous models, it is formulated to allow for maximum control by the user over submodel coupling, execution, and data flow, enabling models with complex feedbacks. However, this comes with a sacrifice in ease of use, as the user often must decide the order and timing of submodel execution. Helios is intended to utilize state-of-the-art biophysical models with high complexity in order to maximize physical realism. In order to afford this high model complexity, many Helios plug-ins perform calculations using graphics processing unit (GPU) hardware, which enables a unique combination of model complexity and range of scales that can be feasibly represented. Finally, Helios includes a plug-in that allows for automatic generation of architectural inputs based on terrestrial LiDAR data.

The goal of this work is to provide a high-level overview of Helios. For specific details regarding implementation and usage, readers are referred to the extensive documentation included with the software.

CORE ENGINE

Model Geometry and Data

At the core of Helios is the Context class, which manages model geometry and data (**Figure 1**). Model geometry is formed using three types of primitive elements: triangles, patches (rectangles), and voxels (parallelepiped) (**Figure 2**). Triangles and patches can be masked using the transparency channel of a PNG image file to create planar elements with arbitrary shapes, which is a common approach in both computer graphics applications (Suffern, 2007) and other plant modeling software (Hemmerling et al., 2008). This often allows for a significant reduction in the number of elements needed to represent complex 3D geometries. For example, a complex leaf shape can be represented by one or a few primitive elements rather than a triangular mesh consisting of dozens of elements (**Figure 2B**).

Upon the creation of each element, a minimal set of data is generated that defines the primitive, such as the coordinates of vertices, surface area, color, and so on. As is common in object-oriented programming, each element is assigned a unique universal identifier (UUID), which can be stored and later used to reference the element. This UUID can be passed to functions that, for example, apply a transformation to the element's position, change an attribute of the element, or be passed to a model plug-in to indicate that model calculations should only be performed for that particular element. This allows for dynamic modification of geometry at any point within the program.

Primitive elements are the basis for most model data structures (**Figure 1**). Scalar or vector data of various types can

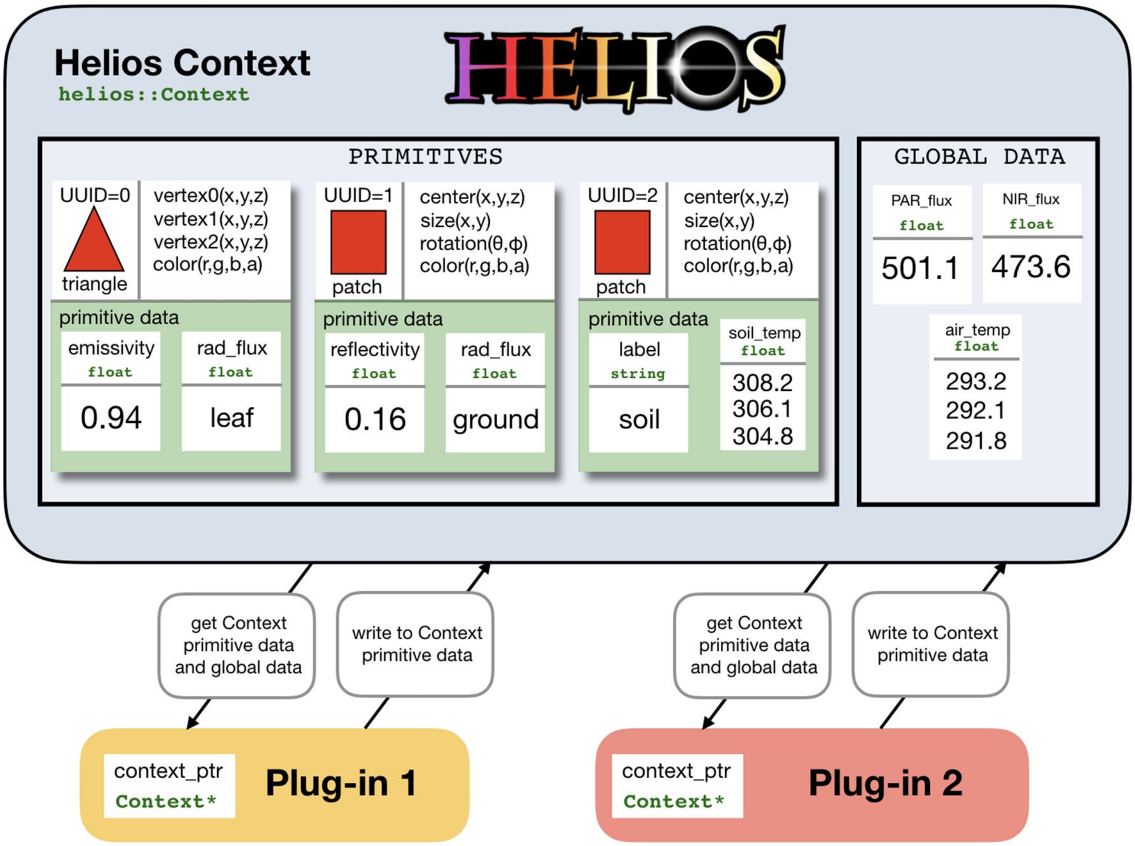

FIGURE 1 | Schematic of data representation and model coupling in Helios. The Context manages model geometry and associated data structures. Model geometry consists of elemental polygons: either triangles, patches (rectangles), or voxels (parallelepiped). Each element may have varying types of associated data called "primitive data," and there may be additional data structures not mapped to individual elements called "global data." Model plug-ins are coupled by operating on common data structures within the context.

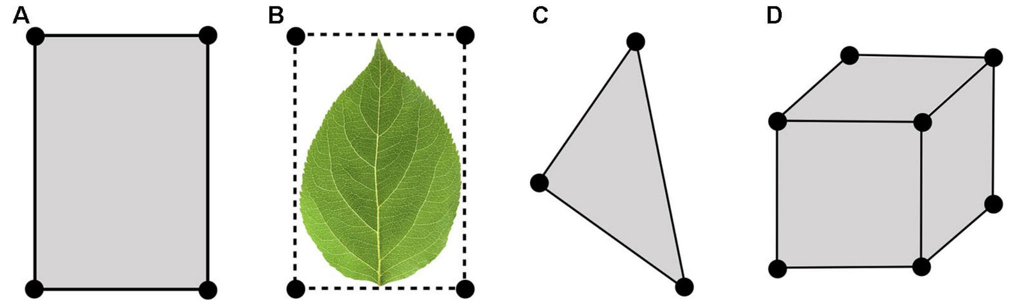

FIGURE 2 | Types of primitive geometric elements available in Helios: **(A)** patch, **(B)** patch masked by an image transparency channel **(C)** triangle, **(D)** voxel.

be associated with each element called "primitive data" (e.g., temperature). These data can be used to specify unique model parameters for each element, or it can provide a container for values computed by a model for a particular primitive. These data structures are also how models are typically coupled. For example, one could create primitive data for each element that specifies its reflectivity, which would be read by the radiation model, which would then write another piece of primitive data that give the value of the computed radiation flux. Another model such as a photosynthesis model could then read this radiation flux and write additional primitive data that give the value of the computed photosynthetic flux. Primitive data values can be set or retrieved using the appropriate setter or getter function (see the section *C++ Application Programming Interface*), which normally takes the UUID of the associated element(s) as an argument.

There is also a more generic data container called "global data," which are not associated with any single geometric element. Global data can be scalar or vector valued and can have a number of different data types (e.g., double, float, integer, string). Global data are set or retrieved using the appropriate setter or getter function, but do not require the UUID of a primitive element because they are independent of any single element.

The data structure formulation used in Helios allows for maximum flexibility in model coupling, but comes with the trade-off of decreased ease of use. The Context itself simply provides a flexible central repository for model geometry and associated data and can also handle file I/O if needed. For this reason, it is very general and allows for arbitrary model coupling and workflows. Plug-ins only need to know the name of the data objects it should read from and write to. Thus, plug-ins can be executed in any order and can share arbitrary data structures.

Time-Series Data

Environmental models are commonly driven by time-series data provided by one or more sensors. The Helios Context includes tools to readily load and access these time-series data. Each data point is associated with some date and time and can either be read automatically from an XML file or added to the Context manually. By setting the date and time in the Context using the appropriate functions, the time-series data will be automatically interpolated to that instant in time and can be queried and used in model calculations.

C++ Application Programming Interface

Users interact with Helios through a C++ application programming interface (API), which means that users write their own program that utilizes the Helios library (see **Listing 1**). As mentioned above, this offers high flexibility but decreases ease of use because users must write their own main function that declares and runs plug-ins. Many tutorials and examples are included within the Helios documentation that illustrate how to utilize the various data structures and functions to perform common modeling tasks.

The Helios Context is a C++ class with many public member functions that are used to access model geometry and data. **Listing 1** provides example code for declaring the Context, adding a triangular element, and then setting primitive data for that element. In this example, the geometry is added through the Context member function "addTriangle()," which takes the Cartesian coordinates of each triangle vertex as arguments. There are a number of additional overloaded versions of the addTriangle() function, which can be used to explicitly set the triangle color, set a texture map, and so on.

The API has several functions that can read/write from/to standard file formats, namely, XML, PLY, and OBJ formats. XML files are used to read and write simulation data and are based on a convention specific to Helios, which is detailed in the documentation. PLY (Stanford polygon) and OBJ (Wavefront) files are standard formats for storing geometric information and are read and written by most 3D computer graphics or computer-aided design software programs. This allows Helios to easily read 3D models generated by other software or write geometry created within Helios to formats that can be read by other software for further analysis. This enables a means by which geometry could be coupled or transferred between other plant modeling platforms that can handle these formats such as OpenAlea or GroIMP.

Listing 1. Example C++ code illustrating the procedure for using the Helios API to add geometry and set associated primitive data.

```
#include “Context.h”

using namespace helios;

int main(void) {

       //Declare the Context class
       Context context;

       //Declare three 3D coordinates defining the triangle vertices
       vec3 vertex0(0, 0, 0);
       vec3 vertex1(0, 1, 0);
       vec3 vertex2(0, 1, 1);

       //Variable to contain unique universal identifier
       unit UUID;

       //Add the triangle, which returns the UUID for the element
       UUID = context.addTriangle(vertex0, vertex1, vertex2);

       //Set the primitive data called “label” to have a value of “mytriangle”
       context.setPrimitiveData(UUID, “label”, “mytriangle”);

}
```

API Documentation

Helios uses Doxygen (www.doxygen.org) to automatically generate documentation for the API and to create a user guide and tutorials with embedded hot-links to associated function documentation. Each plug-in has a documentation page with a consistent structure that defines several key aspects needed to work with the plug-in. This includes, but is not limited to, required dependencies, necessary header files, and any primitive or global data read or written by the plug-in. All API functions and data structures are searchable in order to quickly locate information regarding their purpose and function arguments.

PLUG-INS

Helios plug-ins are implemented as C++ classes with a number of member functions that allow users to set up and run the models. The plug-in classes are typically passed a pointer to the Context class when they are declared, which gives them the ability to access data structures that define model geometry (e.g., vertex positions, surface area, normal vector) and read or write data structures in the Context. A brief description of plug-ins available in version 1.0 is given below, with corresponding software and hardware requirements given in **Table 1**.

Visualizer

The visualizer plug-in creates 3D renderings of model geometry and data based on standard approaches used in computer graphics (Marschner and Shirley, 2015). Utilizing a pointer to the Context, the visualizer parses all geometric elements in the Context and renders them to the screen using OpenGL. There are several means by which elements may be colored. The user can specify a color for the element or provide the path to an image to be used for texture mapping (cf. Marschner and Shirley, 2015). In either case, the Phong lighting model can be optionally used to shade elements, with an additional option to use a model for shadow rendering (**Figure 3**). Alternatively, the user can specify that elements should be colored using a pseudocolor mapping based on primitive data stored in the Context (**Figure 3**).

While the visualizer plug-in provides a seamless means of quickly visualizing model outputs, it is also possible to output geometry and data to file using the standard formats introduced previously, which allows for the use of more sophisticated rendering tools such as Blender. The drawback of this approach is that it adds an additional step to the workflow.

Radiation Transport Model

A GPU-accelerated model for radiation transfer is included as a plug-in to Helios, which is described in detail in Bailey

TABLE 1 | Summary of Helios plug-ins in version 1.0 and their respective software or hardware requirements.

Plug-in	Software/hardware requirements
Visualizer	X11/xorg packages
Radiation model	NVIDIA GPU, CUDA
Energy balance model	NVIDIA GPU, CUDA
Solar model	None
Stomatal conductance model	None
Photosynthesis model	None
Voxel intersection	NVIDIA GPU, CUDA
Procedural tree generation	None
LiDAR data processing	NVIDIA GPU, CUDA

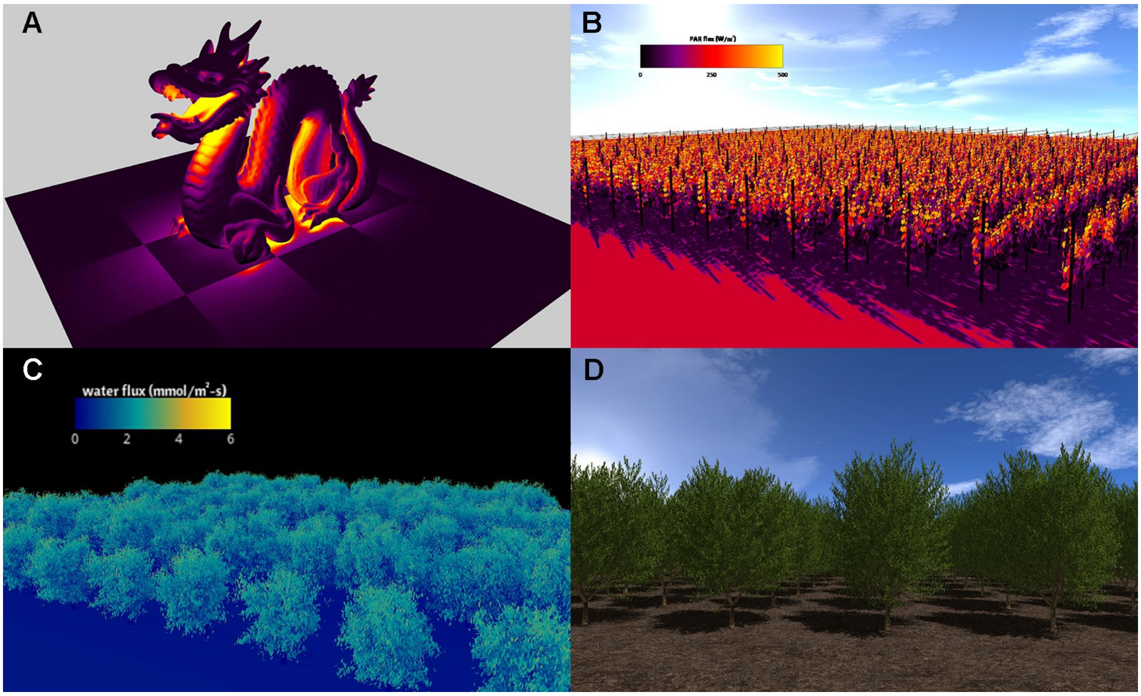

FIGURE 3 | Visualizations of Helios simulation results. **(A)** simulation of radiation emission for model geometry read from a standard polygon file (the so-called "stanford dragon"); **(B)** simulation of absorbed photosynthetically active radiation in a vineyard; **(C)** simulation of transpiration flux in an almond orchard generated from terrestrial LiDAR reconstruction; **(D)** almond tree canopy geometry created using the procedural tree generation plug-in.

(2018). The model uses a novel reverse ray-tracing approach for both solar radiation and terrestrial emission. While reverse ray-tracing approaches have been commonly used in previous models to provide more robust sampling of radiation sources (e.g., Lewis, 1999; Gastellu-Etchegorry et al., 2012; Henke and Buck-Sorlin, 2018), the model of Bailey (2018) presents a new reverse approach for modeling terrestrial emission that ensures that the model satisfies the second law of thermodynamics regardless of the number of rays used. The reduction in the number of rays required, along with the substantial acceleration achieved by utilizing a GPU-based parallelization, means that domains with hundreds of trees and tens of millions of fully resolved leaves can be simulated on a desktop workstation. Using simplified geometries with assumed radiative properties, Bailey (2018) showed that the model converges exponentially toward the exact analytical solution as the number of rays is increased. Currently, the model implementation does not support voxels, but a future release will include the ability to have a mixture of both planar primitive elements and voxels within the domain.

The model can be run over arbitrary wavebands, which are specified in the model by primitive data corresponding to surface radiative properties (i.e., absorptivity, reflectivity, emissivity) integrated over the particular waveband. External radiation sources can be represented by 1) a sphere, 2) collimated radiation propagating in a particular direction, or 3) diffuse ambient radiation, each of which also requires the specification of its position or direction as well as its emissive flux for each radiative band, which can be calculated using the solar model plug-in (see the section *Solar Position and Energy Model*).

Surface Energy Balance Model

A surface energy balance can be solved for each primitive to calculate surface temperature and energy fluxes. The energy balance equation for a surface can be written as

$$R - \varepsilon\sigma T_s^4 = c_p g_H (T_s - T_a) + \lambda g_M \left(\frac{e_s(T_s) f_s - e_s(T_a) h}{p_{atm}} \right) + Q_{other}, \quad (1)$$

where R is the absorbed all-wave radiation flux, ε is the surface emissivity, $\sigma = 5.67 \times 10^{-8}$ W m^{-2} K^{-4} is the Stefan-Boltzmann constant, T_s is the surface temperature in absolute units, $c_p = 29.25$ J mol^{-1} K^{-1} is the heat capacity of air, g_H is the surface boundary-layer conductance to heat, T_a is the air temperature in absolute units, $\lambda = 44{,}000$ J mol^{-1} is the latent heat of vaporization of water, g_M is the overall conductance to water vapor from the surface to air outside the surface boundary layer, $e_s(T)$ is the saturation vapor pressure at temperature T and is computed using the Tetens Equation (Campbell and Norman, 1998), f_s is the fraction of water vapor saturation for the air immediately adjacent to the surface (by default $f_s = 1$ for leaves assuming air in the substomatal cavity is saturated), h is the relative humidity of air outside the boundary layer, and p_{atm} is the ambient air pressure. The flux Q_{other} represents any additional energy fluxes that may be present at the primitive surface (e.g., storage). For the purposes of the case study presented below in the section *Case Study: Quantifying Leaf-Level Variability of Transpiration and Photosynthesis in Whole-Canopies*, it is noted that the rate of water loss E from the surface can be readily calculated from Eq. 1 by isolating the term $g_M\,(e_s\,(T_s)f_s - e_s\,(T_a)h)\ /\ p_{atm} = E$.

All parameters in Eq. 1 can be either specified directly by the user, computed from another plug-in, or otherwise assumed to take the default value given in the documentation. The absorbed radiation flux R can be computed using the radiation transport model plug-in (Section 3.2), and in the case where the primitive corresponds to a leaf, the conductance to moisture g_M can be computed using the stomatal conductance plug-in (see the section *Stomatal Conductance Model*). The energy balance equation is iteratively solved for each primitive in parallel on the GPU using the secant method (Press et al., 2007).

Solar Position and Energy Model

A plug-in is available to estimate the position of the sun, as well as downwelling shortwave and longwave radiative fluxes. The solar position is calculated using standard astronomical relationships as described in Iqbal (2012). In the absence of direct measurements, the clear-sky solar radiative flux incident on a surface normal to the sun can be calculated using the REST2 model of Gueymard (2003). The REST2 model accounts for the effects of Rayleigh scattering and absorption due to water vapor, nitrogen dioxide, ozone, and aerosols. The model also provides an estimation of direct-diffuse partitioning of the incoming solar flux.

If measurements are not available, the downwelling longwave diffuse radiative flux can be calculated using this plug-in, which is based on the model of Prata (1996). The REST2 and longwave models both require specification of the precipitable water in the atmosphere, which is estimated using the model of Viswanadham (1981).

Stomatal Conductance Model

Stomata are typically important active regulators of water vapor transport between the inside of leaves and the atmosphere (Jarvis and McNaughton, 1986). This regulatory effect is represented by specifying a stomatal conductance, which is modeled in Helios using the semimechanistic model of Buckley et al. (2012). This model represents stomatal conductance g_s as a hyperbolic function of photosynthetically active photon flux density and local vapor pressure deficit, which is given by the equation

$$g_s = \frac{E_m(Q+i_0)}{k+bQ+(Q+i_0)D}, \tag{2}$$

where Q is the photosynthetically active photon flux density, and D is the vapor pressure deficit between the intercellular leaf air spaces and the air outside of the leaf boundary layer. Note that the photon flux density is obtained from the energy flux in the PAR band using the factor 4.57 μmol m^{-2} s^{-1}/(W m^{-2}). E_m, i_0, k, and b are treated as empirical model coefficients.

Photosynthesis Model

Two leaf photosynthesis models are available in Helios: an empirical model based on the description of Johnson (2010), and the mechanistic biochemical model of Farquhar et al. (1980) for C_3 photosynthesis. For completeness, the current implementation of the Farquhar et al. (1980) model is described in Appendix 1 because it is the model used in the case study presented in the section *Case Study: Quantifying Leaf-Level Variability of Transpiration and Photosynthesis in Whole-Canopies*. The empirical model is also fully described in the Helios documentation.

Primitive Subvolume Grouping

One important motivation for using a detailed, leaf-resolving plant model is to understand the impacts of aggregation of leaf-level heterogeneity over multiple scales. In order to help facilitate this aggregation, a plug-in is available to rapidly group or bin primitives into arbitrary subvolumes. Users can define arbitrary voxels, and this plug-in will identify any planar primitive elements that are contained within each voxel. This is useful, for example, in calculating leaf area density/index or calculating aggregated attenuation coefficients for comparison with simple models. The primitive binning calculations are performed on the GPU to significantly reduce execution times.

Terrestrial LiDAR Data Processing

Terrestrial LiDAR scanning is a powerful tool for 3D measurement of plant architecture, which has gained popularity in plant modeling applications. While the raw LiDAR point clouds provide a wealth of data that yield an incredibly detailed mapping of the canopy, processing this data into information that is usable in the context of modeling has proven to be a challenge. Raw LiDAR data provide millions of 3D Cartesian coordinates in space. However, models generally cannot use points directly, but rather need information such as surfaces, areas, and so on.

The terrestrial LiDAR plug-in integrates a number of data processing algorithms, along with GPU acceleration, to provide the ability to translate LiDAR point clouds into leaf-by-leaf reconstructions that can be fed directly into the Helios Context. The workflow starts by using the triangulation algorithm of Bailey and Mahaffee (2017b) to calculate the leaf angle distribution, which is used to calculate the leaf area projection function G (Ross, 1981). The G-function is then used to generate estimates of leaf area density for arbitrary volumes of leaves (voxels) following the approach of Bailey and Mahaffee (2017a). To reconstruct individual leaves, the triangulated leaf hit points are segmented to estimate the position and area of individual leaves that are in direct view of the LiDAR scanner (Bailey and Ochoa, 2018). Because a significant fraction of leaves may be occluded from view of the scanner, a statistical backfilling approach is used to ensure that the reconstructed tree leaf orientation and area distributions match the voxel-based measurements described above (see Bailey and Ochoa, 2018).

Each individual LiDAR scan typically consists of tens of millions of points, and grids for calculating leaf area density may consist of thousands of voxels. These dimensions compound to make data processing computationally expensive, and thus several of the LiDAR processing routines are performed in parallel on the GPU. Point-based calculations lend themselves well to parallelization because each laser pulse can be analyzed independently from another.

Procedural Tree Generation

While the LiDAR plug-in provides a powerful means of incorporating measured tree architectures within Helios, certain types of modeling studies may require the ability to simulate a wide range of geometries that cannot be directly measured. The creation of semirandom tree geometries is made possible in Helios through the use of the procedural tree generation algorithm of Weber and Penn (1995). This framework describes the woody architecture of trees as a recursive set of branching levels, each described by their own set of parameters that provide rules for how branching structure should occur. A random perturbation of user-defined magnitude can be added to each parameter to reduce geometric uniformity in order to produce more realistic-looking trees. In the original formulation described by Weber and Penn (1995), leaf orientations are determined through an axial rotation about the branch from which they originate, which may create unphysical leaf orientation distributions. Additional functionality has been added to allow users to specify a custom leaf inclination angle distribution, perhaps that provided by LiDAR measurements (section *Terrestrial LiDAR Data Processing*).

The procedural tree generation plug-in comes with nine predefined tree geometries, which are shown in **Figure 4**. Arbitrary trees can be created by modifying the tree geometric parameters, which are commonly specified in an XML file.

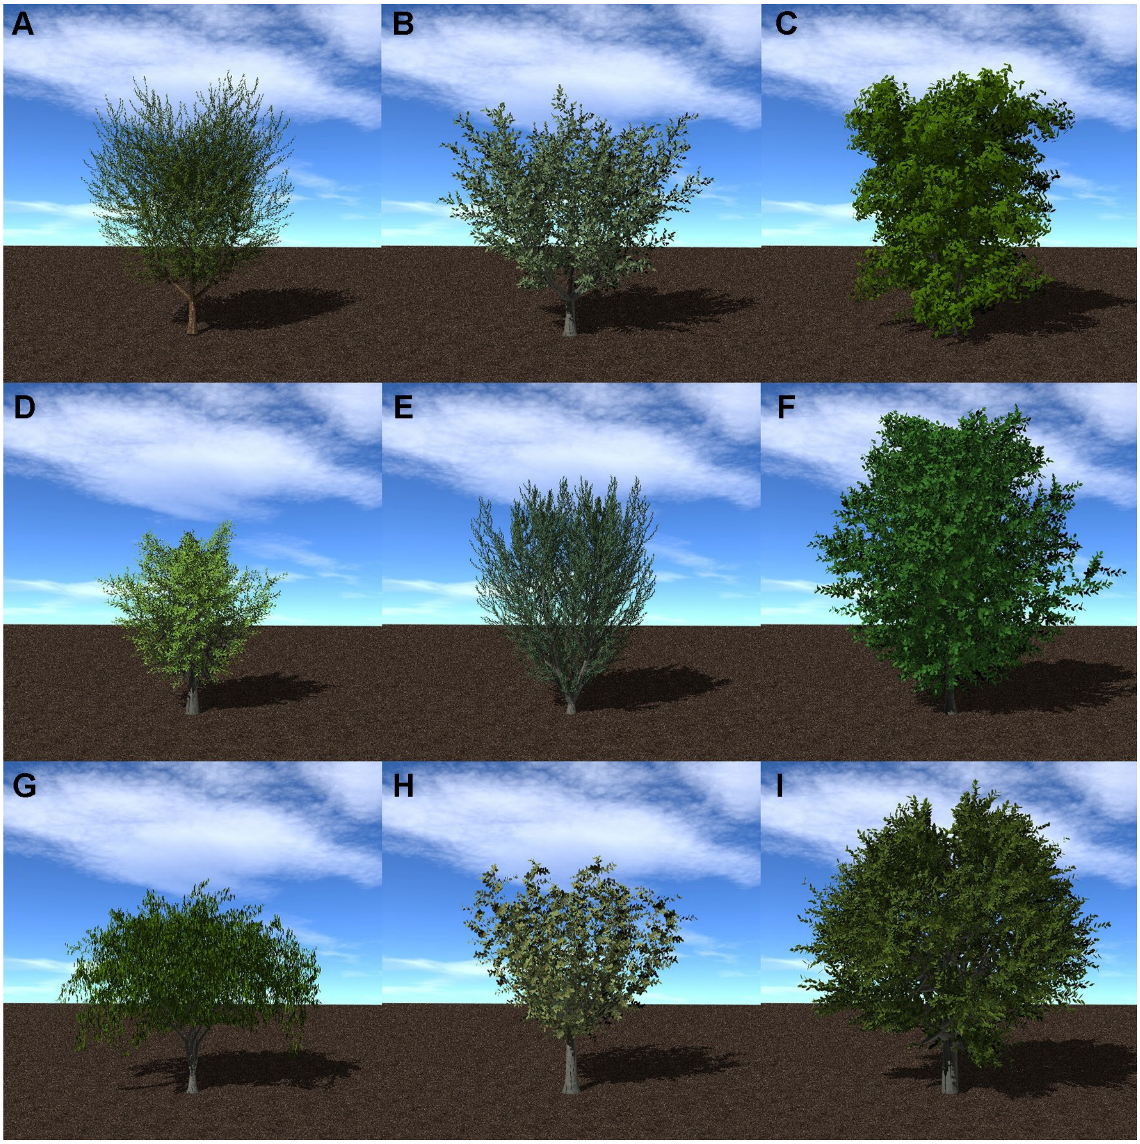

FIGURE 4 | Procedural tree model generation for nine tree species: **(A)** almond; **(B)** apple; **(C)** avocado; **(D)** lemon; **(E)** olive; **(F)** orange; **(G)** peach; **(H)** pistachio; **(I)** walnut.

Parameters include quantities such as the number of recursive branching levels, average angle of branches with respect to their parent branch for each level, and so on. The end geometry produced by the Weber and Penn (1995) and the parameters used to specify the geometry are fairly similar to those produced by the commonly used L-systems approach (Prusinkiewicz and Runions, 2012). L-systems is more elegant in its notational and mathematical representation of the branching structures (it uses a string of characters to encode the structure), but the end result is similar to that used in Helios.

CASE STUDY: QUANTIFYING LEAF-LEVEL VARIABILITY OF TRANSPIRATION AND PHOTOSYNTHESIS IN WHOLE-CANOPIES

Background

While our collective understanding of plant biophysical processes for individual leaves has progressed rapidly over the past several decades, our understanding of canopy-level processes is limited by the need to aggregate highly heterogeneous processes over a wide range of scales. When measurements are performed at the leaf scale, it is often unclear how representative such measurements are of the canopy as a whole. On the other hand, when measurements are performed at large scales that aggregate many smaller scales, it is often unclear how different members of the community (i.e., leaves) contribute to aggregate behavior. In this brief case study, a tree canopy will be examined using Helios to visualize and quantify leaf-level variability in transpiration and photosynthesis in order to understand how individual elements contribute to system-level behavior.

Case Set-up

Canopies of *Prunus dulcis* were simulated to assess the impact of canopy architecture on light interception, microclimate, transpiration, and photosynthesis. Two canopy architectures were considered: an isolated tree and a relatively dense canopy of 100 trees (1 tree per 36 m^2). Individual trees were created using the procedural tree generation plug-in with the same parameters that were used to create the tree shown in **Figure 4A**. In order to maintain consistency within the test case, the isolated tree had the same architecture (within random variation) of each of the trees in the dense canopy, although in reality the architecture of the isolated tree would likely be different. The simulated trees had leaves with constant (one-sided) area of 60 cm^2.

Field data collected in a canopy of *P. dulcis* were used to specify model parameters. The canopy was located in the California Central Valley (36.599°N 119.515°W), and consisted of 4-year-old trees that were approximately 7 m tall. Trees were spaced at 4 m in the East-West direction and 6.4 m in the North-South direction. The ambient air temperature and humidity were assumed to be spatially homogeneous and were specified using data collected from a nearby weather station (**Figure 5**). These data were also fed into the model that predicts the downwelling longwave radiation flux. Incoming direct and diffuse radiation fluxes were estimated using the REST2 model as described above (**Figure 5**), which were equally partitioned into PAR (assumed to be wavelengths <700 nm) and NIR (assumed to be wavelengths >700 nm) bands. A single diurnal cycle was simulated at time step of 15 min. Leaf reflectivity and transmissivity were assumed to be 0.05 for the PAR band and 0.4 for the NIR band.

Leaf angles were specified by randomly drawing from the leaf angle distribution measured in the experimental canopy described above. The average leaf angle distribution was measured by scanning trees using a Riegl VZ-1000 terrestrial LiDAR scanner (Horn, Austria). The scan resolution in the zenithal direction was 0.04° across a range of 100° and 0.08° in the azimuth across a full 360° rotation. Four scans per tree were performed from the northwest, northeast, southwest, and southeast of each trunk at a distance of about 7.5 m. The raw LiDAR data were processed to determine the leaf inclination distribution as described by Bailey and Mahaffee (2017b) using the LiDAR data processing plug-in (**Figure 6**).

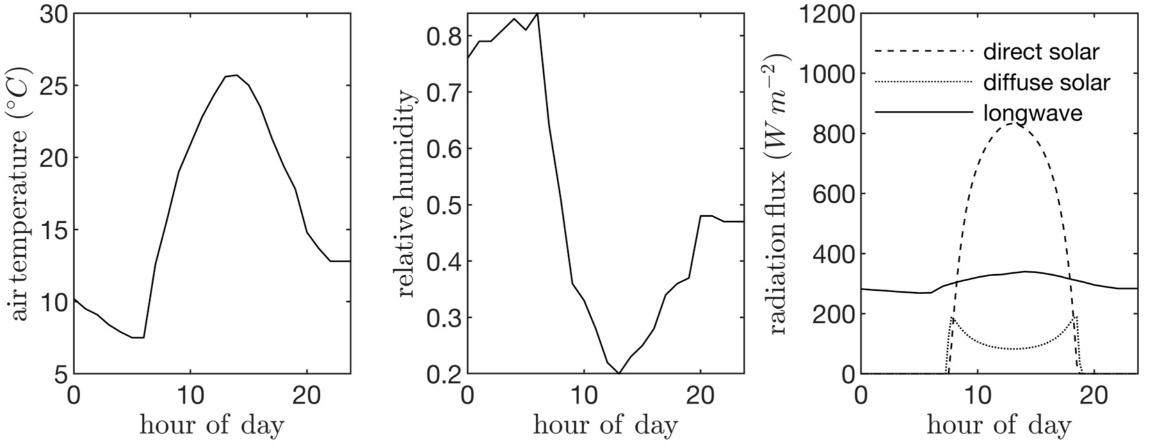

FIGURE 5 | Input time-series data used to drive case study simulations: ambient air temperature (left), air relative humidity (center), and incoming direct solar radiation flux normal to the sum direction, incoming diffuse radiation flux on a horizontal surface, and incoming longwave radiation flux on a horizontal surface (right).

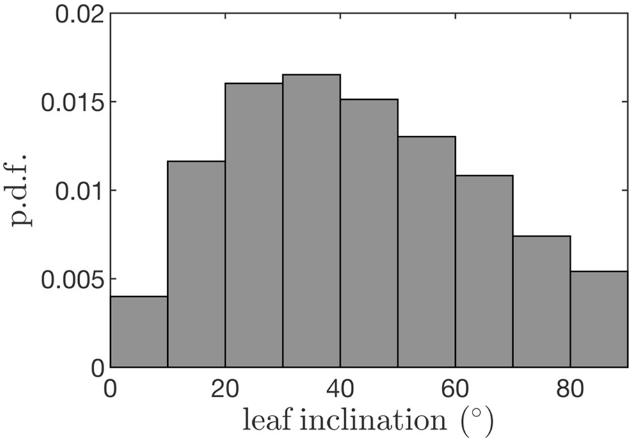

FIGURE 6 | Leaf inclination probability density function (p.d.f.) used to generate tree geometrics for case study simulations.

Leaf-level gas exchange measurements were collected using the LI-6800 portable photosynthesis system (LICOR, Lincoln, NE, USA). All measurements were performed at an ambient CO_2 concentration of 390 µmol/mol, but at varying light, temperature, and humidity levels. The response of photosynthesis and stomatal conductance to light was determined by varying the photosynthetically active photon flux density between levels of 0, 50, 200, 400, 800, 1,200, and 2,000 µmol m^{-2} s^{-1} at a leaf temperature of 25°C and 60% relative humidity. Importantly, a large amount of time was spent at each light level to ensure that stomata had time to fully equilibrate, which took around 1 h per light level. Control measurements with constant conditions were performed to verify that changes in whole-plant water status over this very long period did not significantly affect the response curves. At a saturating light level of 2,000 µmol m^{-2} s^{-1}, leaf temperature varied between 25°C and 35°C, and relative humidity in the chamber varied between 30% and 60% for each leaf temperature. This procedure produced measurements at 10 different combinations of light, temperature, and ambient humidity, which were used to determine model coefficients for the stomatal conductance and photosynthesis models (**Figure 7**; **Table 2**).

In order to refine the initial representation of canopy architecture, a precursor simulation was performed to remove unrealistic leaves. The daily net CO_2 assimilation rate was determined for every leaf within the precursor simulation. Leaves that had negative net CO_2 assimilation over the day (i.e., daily respiration was larger than assimilation) were removed. This resulted in a final canopy LAI of 2.9, which had very few leaves with negative net daily CO_2 assimilation (**Figure 9**). This LAI value is on the high end of what might be observed in real canopies but is reasonable given that nearly all leaves had positive net daily assimilation.

Results

Leaf Probability Distributions

Probability distribution functions (p.d.f.s) of net photosynthesis, transpiration rate, absorbed radiation, and temperature were calculated for all leaves in the tree or canopy. The distributions were formed across leaves for given instants throughout the day (**Figure 8**) or as an integration in time of values for each leaf over the entire day for daylight hours only (**Figure 9**).

Radiation flux. The distribution of absorbed radiation was highly heterogeneous, and followed a nearly exponential distribution, with most leaves absorbing relatively low amounts of radiation. This exponential distribution was an expected result based on Beer's law (Ross, 1981), as the p.d.f. of absorbed flux over all leaves serves to approximate the probability of flux interception along the path of radiation propagation. The distribution is not perfectly exponential due to the presence

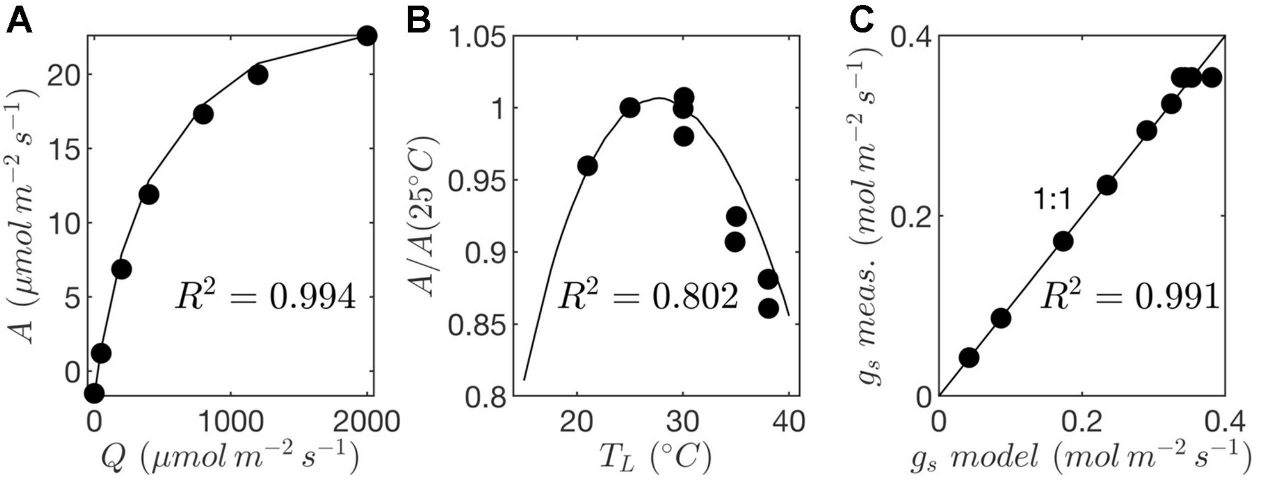

FIGURE 7 | Parameterization of photosynthesis and stomatal conductance models. **(A)** response of net photosynthesis to varying photon flux density (*Q*) at a reference leaf temperature of T_L = 25°C given by gas exchange measurements (circles) and photosynthesis model (solid line); **(B)** response of net photosynthesis to temperature (T_L) given by gas exchange measurements (circles) and photosynthesis model (solid line); and **(C)** comparison of stomatal conductance (g_s) given by gas exchange measurements and predictions by the model.

TABLE 2 | Fitted parameter values for photosynthesis and stomatal conductance models.

Parameter	Description	Value	Units
		Photosynthesis	
$R_{d,25}$	Respiration rate at 25ºC	1.491	μmol m^{-2} s^{-1}
$V_{cmax,25}$	Maximum carboxylation rate at 25ºC	99.5	μmol m^{-2} s^{-1}
$J_{max,25}$	Maximum electron transport rate at 25ºC	185.0	μmol m^{-2} s^{-1}
C_{Jmax}	J_{max} temperature response parameter	17.57*	Unitless
$\Delta H_{a,\,Jmax}$	J_{max} temperature response parameter	43.54*	kJ mol^{-1}
α	Light response parameter	0.41	Unitless
		Stomatal Conductance	
E_m	Maximum transpiration rate	20.43	mmol m^{-2} s^{-1}
i_0	PPFD offset for dark transpiration	38.48	μmol m^{-2} s^{-1}
k	Bulk stomatal parameter	18,383	μmol m^{-2} s^{-1} mmol mol^{-1}
b	Bulk stomatal parameter	49.68	mmol mol^{-1}

**Assumed based on no growth temperature acclimation (Bernacchi et al., 2003).*

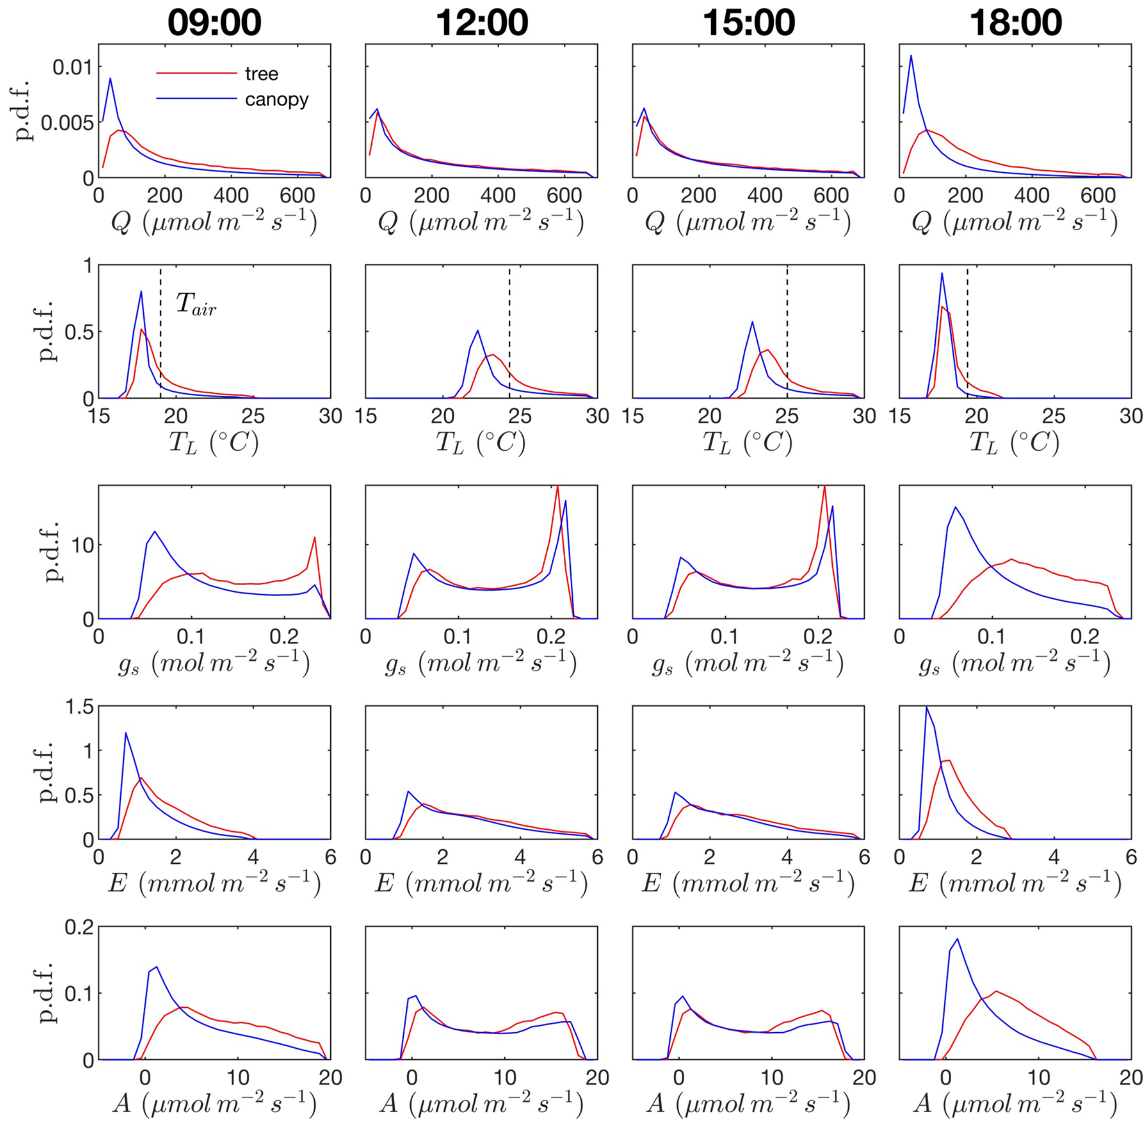

FIGURE 8 | Probability density function (p.d.f.) across all leaves in the canopy for photosynthetic photon flux density *Q*, leaf temperature T_L (ambient temperature given by dashed vertical line), stomatal conductance to water vapor g_s, transpiration rate *E*, and net rate of photosynthesis A. Columns correspond to different times of the day. Red lines correspond to the isolated tree case; blue lines correspond to the dense canopy case

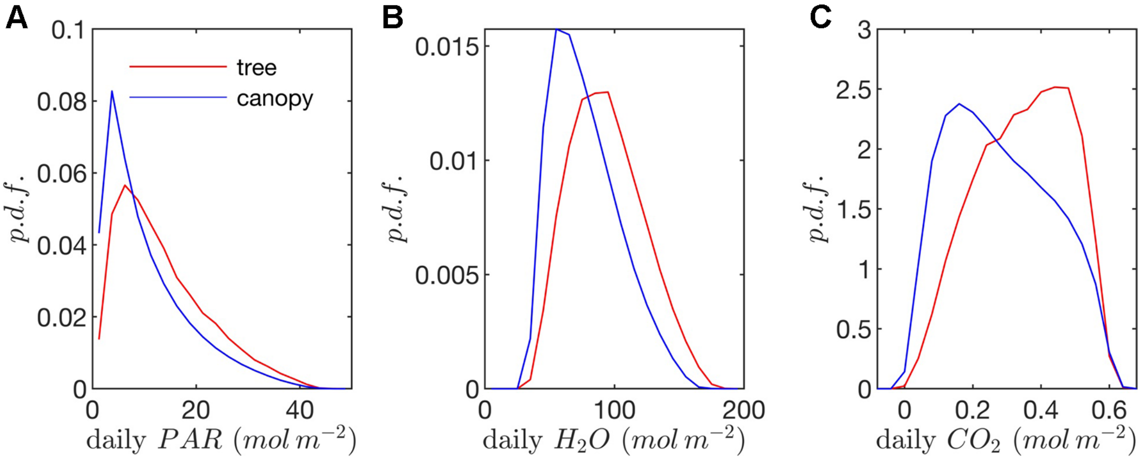

FIGURE 9 | Probability density functions (p.d.f.s) of daily **(A)** absorbed photosynthetic photons, **(B)** transpired water, and **(C)** net CO_2 exchange (all per unit leaf area).

of diffuse ambient radiation and the fact that the tree/canopy is not optically thick, and therefore the ground absorbs some radiation. The instantaneous and daily integrated p.d.f.s showed a similar trend, except that the daily p.d.f. had a shorter tail. The highly skewed distribution meant that a relatively small number of leaves absorbed a large fraction of radiation at any instant of the day. Around the middle portion of the day, leaves in the top 10% in terms of absorbed radiation flux absorbed roughly 65% of the total radiation absorbed by the entire tree or canopy (**Figure 10A**). As the sun angle decreased, this fraction tended to decline, where near dawn and dusk the top 10% of leaves absorbed between 40% and 50% of the total absorbed radiation (**Figure 10A**), which is likely due to the increased diffuse fraction. When integrated over the entire day, 10% of leaves were responsible for absorbing about 48% of the total daily absorbed radiation for the isolated tree and 53% for the dense canopy (**Table 3**).

The probability distributions of absorbed radiation for the isolated tree and dense canopy cases were very similar when the sun was high, and decreasing sun angle tended to smooth the distribution slightly for the isolated tree (**Figure 8**). The distribution for the isolated tree was shifted slightly to higher radiation values, likely due to relatively high fraction of surface area in view of the sun at low sun angles. When integrated over an entire day, the discrepancies between the distributions for the isolated tree and dense canopy were relatively minimal, with the peak in the distribution smoothed slightly for the isolated tree (**Figure 9A**).

Leaf temperature. The distributions of leaf temperature were closer to Gaussian than the distributions of radiation absorption, although the temperature distributions were still positively skewed (**Figure 8**). Most leaves were below the ambient air temperature, with the peak occurring several degrees below the ambient air temperature. There was a significant difference between the leaf

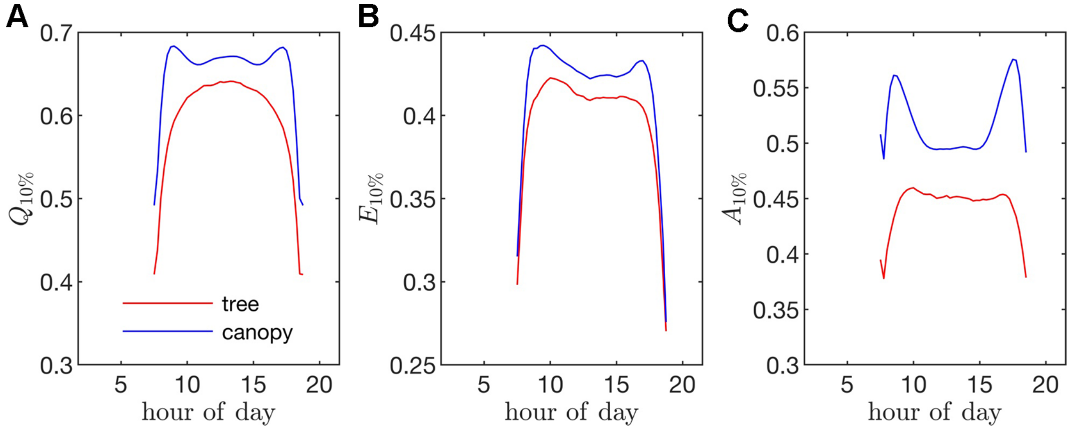

FIGURE 10 | Fraction of the instantaneous flux due the top 10% of all leaves (i.e., fraction of flux contained in the 90th percentile of leaves) for fluxes of **(A)** absorbed photon flux $Q_{10\%}$, **(B)** transpiration rate $E_{10\%}$, and **(C)** net photosynthesis $A_{10\%}$. (For example, a $Q_{10\%}$ value of 0.75 would mean that only 10% of the leaves were responsible for 75% of the total absorbed PAR for the whole tree/canopy.

TABLE 3 | Fraction of the total daily flux due the top 10% of all leaves (i.e., fraction of total flux contained in the 90th percentile of leaves) for fluxes of absorbed photon flux $Q_{10\%}$, transpiration rate $E_{10\%}$, and net photosynthesis $A_{10\%}$.

Variable	Tree	Canopy
$Q_{10\%}$	0.480	0.533
$E_{10\%}$	0.351	0.367
$A_{10\%}$	0.370	0.432

temperature distributions for the isolated tree and dense canopy cases, particularly throughout the middle portion of the day. The lower end of the leaf temperature distributions for the tree was shifted upward by about a degree when compared with the dense canopy. The upper end of the leaf temperature distributions is similar between the tree and canopy cases, indicating that the highest temperature leaves, likely near the tops of the trees, are not significantly affected by the presence of neighboring trees. However, it should be noted that this includes only the radiative effect on temperature because the air temperature was held constant between the isolated tree and canopy cases in order to isolate the effects of geometry.

Stomatal conductance. During the middle portion of the day, the distribution of leaf stomatal conductance followed an interesting bimodal distribution with sharp peaks at either end of the distribution, which also exhibited minimal differences between the isolated tree and dense canopy cases. The lower peak results from the fact that a large portion of leaves are in shade, resulting in a large number of leaves with low stomatal conductance. The upper peak is perhaps more surprising and results from the nonlinearity of the stomatal response to light. For high light levels, stomatal conductance saturates and is relatively insensitive to changes in light, which thus results in a large cluster of leaves with stomatal conductances near the saturating value. Late in the day when sun angles are low, there is a significant positive shift in stomatal conductances in the isolated tree as compared with the dense canopy, which seemingly corresponds with the positive shift in radiation absorption between these two cases.

Transpiration rate. Unlike the distribution of stomatal conductance, the distribution of transpiration flux did not follow a bimodal distribution, but rather had a single sharp peak and large positive skewness. Since the transpiration flux is the product of the stomatal conductance and vapor pressure deficit, this means that the vapor pressure deficit increase at high temperature and light values was sharp enough to dominate the transpiration flux, although stomatal conductance becomes saturated. Overall, discrepancies between the highest and lowest transpiring leaves were smaller than those of absorbed radiation. During much of the day, the top 10% of leaves transpired between 40% and 45% of the total tree/canopy transpiration (**Figure 10B**), and when integrated over the day, the top 10% transpired roughly 35% of the total for both the tree and canopy cases.

The peak in the distribution of transpiration flux was shifted upward in the isolated tree case, which was presumably due to the corresponding upward shift in the leaf temperature and stomatal conductance (**Figure 8**). When the transpiration flux was integrated over the entire day, a similar pattern emerged, except that the positive tail of the distribution was shortened (**Figure 9**).

Net photosynthesis. During the middle portion of the day, the distribution of net leaf CO_2 flux exhibits a peak near a value of zero and a secondary peak near the saturating value (**Figure 8**), with the overall distribution being fairly uniform. At low light levels, photosynthesis is primarily limited by the amount of available photosynthetically active light, and thus it is expected that for low light values the distribution of photosynthesis should be closely related to the distribution of absorbed light, which is evident from **Figure 8**. At high light levels, photosynthesis is relatively insensitive to light (**Figure 7**) but highly sensitive to the CO_2 concentration within the leaf, which is tightly regulated by stomatal conductance. When integrated over an entire day, a strong peak in the photosynthesis distribution at low CO_2 exchange values still exists but is shifted upward, and the region of nearly constant CO_2 exchange values does not exist (**Figure 9**). For the middle of the day, the top 10% of leaves assimilated about 45% of the CO_2 for the isolated tree and 50% for the canopy (**Figure 10C**). When integrated over the day, the top 10% of leaves assimilated 37% of the CO_2 in the isolated tree and 43% in the dense canopy.

Leaf Trajectories

Visualization of time series or "trajectories" of individual leaf exchange rates provides an interesting perspective into how the behavior of individual leaves compares to that of the entire tree or canopy throughout the day. Trajectories are shown in **Figure 11** for absorbed photosynthetically active radiation, leaf temperature, stomatal conductance, transpiration rate, and net rate of photosynthesis for 10 randomly chosen leaves. We can observe the wide range of scenarios encountered by different leaves. Some leaves remain in highly shaded conditions for most of the day except for a brief sunfleck, which allows them to assimilate enough CO_2 to offset daily respiration. Other leaves are "lucky" in that they encounter extended periods of high light conditions. Examples can be observed in which the leaf radiation, temperature, and transpiration rate all increase substantially and in tandem for an extended period, whereas stomatal conductance and photosynthesis reach a maximum value and begin to decline as vapor pressure deficit climbs and stomata start to close.

DISCUSSION

The goal of most modeling efforts is to reduce complex processes to a tractable form that can mathematically represent interrelationships between quantities of interest. Here, our goal was to use a complex model that represents in detail individual members of a complex system (i.e., leaves in a tree/canopy) to help identify emergent behavior that is largely representative of the bulk response of the system, which can provide insight into how simplified experimental and modeling approaches can be formulated and interpreted. In this brief case study, Helios and

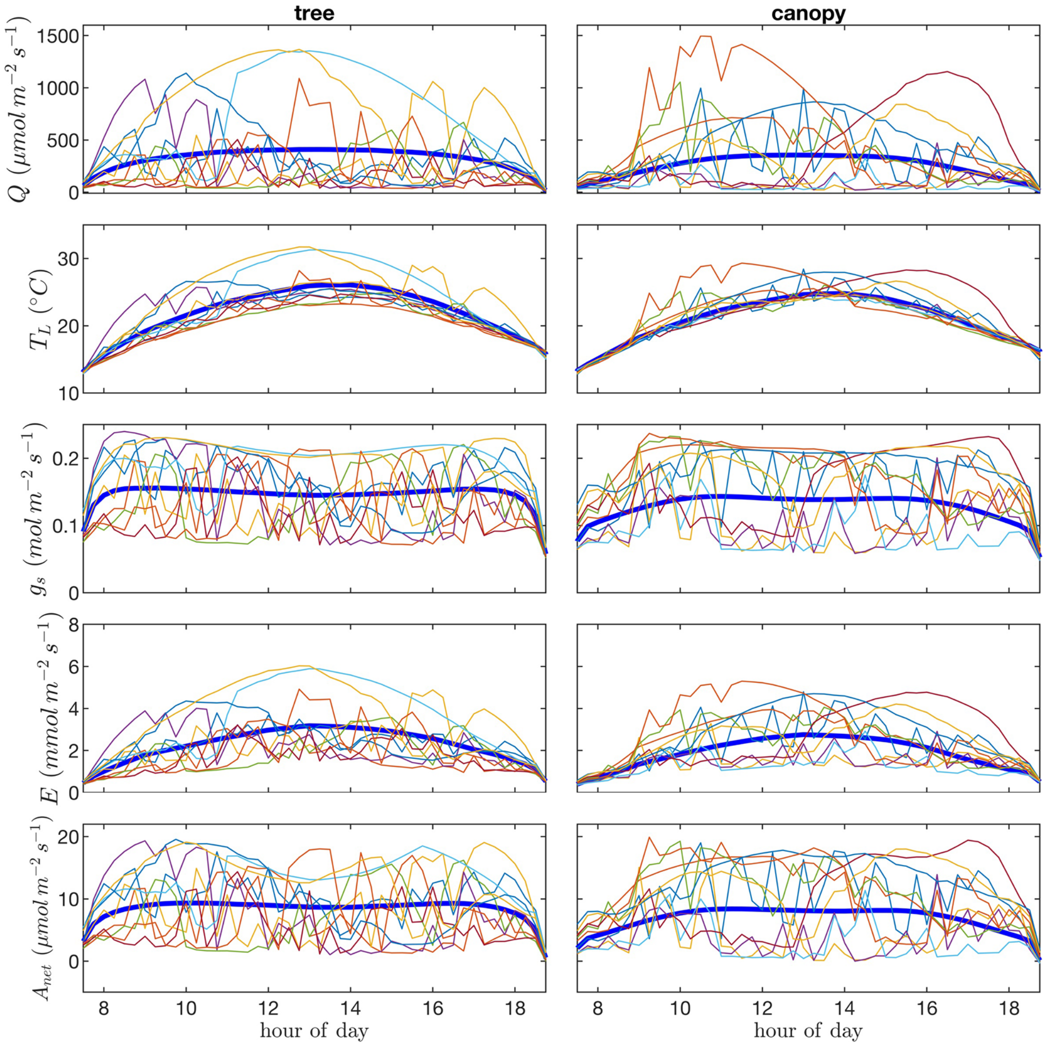

FIGURE 11 | Time-series or "trajectories" for 10 randomly selected leaves of photosynthetic photon flux density Q, leaf temperature T_L, stomatal conductance g_s, transpiration flux E, and net photosynthetic flux A. Each line/color represents the time series of a single leaf. Thick blue lines give the average over all leaves in the tree canopy.

its submodels for radiation transport, leaf temperature, stomatal conductance, and photosynthesis were used to examine leaf-level variability in these processes and how this variability contributes to whole-tree and -canopy behavior.

The results of this case study provide an interesting depiction of the extreme heterogeneity that exists within vegetation for important biophysical processes. Probability distributions across leaves are highly heterogeneous and skewed, and because of inherent nonlinearities in the biophysical processes examined, the general shape of distributions is not consistent across even tightly related processes. At any instance, the whole-tree/canopy behavior in terms of radiation interception and photosynthesis is dominated by a relatively small fraction of the leaf population. When integrated over an entire day, this effect is somewhat reduced, but it was still observed that a small fraction of leaves was responsible for a disproportionate amount of the daily CO_2 assimilation.

A wide range of representations of the above biophysical processes are used in models. So-called "big-leaf" models consider the behavior of only one average leaf assumed to be representative of the entire plant system (e.g., Sinclair et al., 1976; Sellers et al., 1986; Amthor, 1994). For the tree systems examined here, **Figures 8** and **9** illustrate the difficulties in utilizing this approach, given the high variability and skewness of the distributions across leaves, which has also been highlighted in more recent works (De Pury and Farquhar, 1997; Wang and Leuning, 1998; Friend, 2001). As an improved, yet still simple approximation, authors have suggested choosing two representative sets of leaves: sunlit and shaded (De Pury and Farquhar, 1997; Wang and Leuning, 1998). Examination of the distributions of absorbed radiation in **Figures 8** and **9** would call this intuitive approximation into question. Although the naked eye may view two distinct radiation regimes within a tree, this can be deceiving given that leaves are at a variety of orientations with respect to incoming radiation. No clear separation of regimes is evident for absorbed radiation, temperature, and transpiration, although stomatal conductance and photosynthesis had two distinct peaks in the distribution for the middle portion of the day. The degree of separation between sunlit and shaded regimes is expected to vary based on the shape of the light response curve (**Figure 7A**) and the density of vegetation.

More complicated "multilayer models" (e.g., Meyers and Paw U, 1987; Bonan et al., 2012) appear suitable for representing the within-vegetation heterogeneity, provided that enough vertical layers are used. Subdividing the canopy into discrete zones effectively averages across all values within the zone. It is possible that using a zone that is too large can introduce problems due to the fact that the distribution within that zone can have fat tails that give large contributions to overall behavior.

The comparisons between the isolated tree and dense canopy in this study showed surprisingly small differences in the distributions of radiation absorption, transpiration, and photosynthesis through most of the day and in daily integrated distributions, which raises some interesting questions regarding the representation of isolated or sparse vegetation in simplified biophysical models. Because both the isolated tree and canopy cases showed a nearly identical exponential distribution in absorbed radiation, a simple homogeneous Beer's law model could conceivably be used to predict total absorbed radiation per unit total leaf area for the isolated tree. However, the complication arises that we must know the total leaf area and representative ground area for the isolated tree to get an absorbed flux per unit ground area. Models that aggregate trees into homogeneous subvolumes (e.g., see Wang and Jarvis, 1990; Cescatti, 1997; Duursma and Medlyn, 2012) correctly represent tree-scale heterogeneity in absorption, but filter out subtree variability including the tails of the distributions, which were shown to have important contributions to whole-canopy behavior. On the other hand, multilayer models can represent this subtree variability but are not able to represent tree-level heterogeneity in sparse canopies (Ponce de León and Bailey, 2019).

DATA AVAILABILITY STATEMENT

Helios is an open-source software licensed under the GNU GPLv3 license. It can be downloaded from the GitHub repository located at https://github.com/PlantSimulationLab/Helios.

AUTHOR CONTRIBUTIONS

The author confirms being the sole contributor of this work and has approved it for publication.

FUNDING

Financial support of this work by the American Vineyard Foundation grants 2015-1825/2016-1825/2017-1825, Almond Board of California grants 17.PREC1/18.PREC1, and the USDA National Institute of Food and Agriculture Hatch project 1013396.

ACKNOWLEDGMENTS

The author wishes to acknowledge contributions to this work by R. Stoll, E.R. Pardyjak, P. Willemsen, M. Overby, and W.F. Mahaffee.

REFERENCES

Allen, M. T., Prusinkiewicz, P., and DeJong, T. M. (2005). Using L-systems for modeling source-sink interaction, architecture and physiology of growing trees: the L-PEACH model. *New Phytol.* 166, 869–880. doi: 10.1111/j.1469-8137.2005.01348.x

Amthor, J. S. (1994). Scaling CO_2-photosynthesis relationships from the leaf to the canopy. *Photosyn. Res.* 39, 321–350. doi: 10.1007/BF00014590

Bailey, B. N. (2018). A reverse ray-tracing method for modelling the net radiative flux in leaf-resolving plant canopy simulations. *Ecol. Model.* 398, 233–245. doi: 10.1016/j.ecolmodel.2017.11.022

Bailey, B. N., and Mahaffee, W. F. (2017a). Rapid, high-resolution measurement of leaf area and leaf orientation using terrestrial LiDAR scanning data. *Meas. Sci. Technol.* 28, 064006. doi: 10.1088/1361-6501/aa5cfd

Bailey, B. N., and Mahaffee, W. F. (2017b). Rapid measurement of the three-dimensional distribution of leaf orientation and the leaf angle probability density function using terrestrial LiDAR scanning. *Remote Sens. Environ.* 193, 63–76. doi: 10.1016/j.rse.2017.03.011

Bailey, B. N., and Ochoa, M. H. (2018). Semi-direct tree reconstruction using terrestrial LiDAR point cloud data. *Remote Sens. Environ.* 208, 133–144. doi: 10.1016/j.rse.2018.02.013

Bailey, B. N., Overby, M., Willemsen, P., Pardyjak, E. R., Mahaffee, W. F., and Stoll, R. (2014). A scalable plant-resolving radiative transfer model based on optimized GPU ray tracing. *Agric. For. Meteorol.* 198-199, 192–208. doi: 10.1016/j.agrformet.2014.08.012

Bailey, B. N., Stoll, R., Pardyjak, E. R., and Miller, N. E. (2016). A new three-dimensional energy balance model for complex plant canopy geometries: model development and improved validation strategies. *Agric. For. Meteorol.* 218-219, 146–160. doi: 10.1016/j.agrformet.2015.11.021

Baldocchi, D. D., and Harley, P. C. (1995). Scaling carbon dioxide and water vapour exchange from leaf to canopy in a deciduous forest. II. model testing and application. *Plant Cell Environ.* 18, 1157–1173. doi: 10.1111/j.1365-3040.1995.tb00626.x

Bernacchi, C. J., Pimentel, C., and Long, S. P. (2003). *In vivo* temperature response functions of parameters required to model RuBP-limited photosynthesis. *Plant Cell Environ.* 26, 1419–1430. doi: 10.1046/j.0016-8025.2003.01050.x

Bernacchi, C. J., Singsaas, E. L., Pimentel, C., JR, A. R. P., and Long, S. P. (2001). Improved temperature response functions for models of rubisco-limited photosynthesis. *Plant Cell Environ.* 24, 253–259. doi: 10.1111/j.1365-3040.2001.00668.x

Bonan, G. B., Oleson, K. W., Fisher, R. A., Lasslop, G., and Reichstein, M. (2012). Reconciling leaf physiological traits and canopy flux data: use of the TRY and FLUXNET databases in the Community Land Model version 4. *J. Geophys. Res. Biogeosci.* 117, G02026. doi: 10.1029/2011JG001913

Boudon, F., Pradal, C., Cokelaer, T., Prusinkiewicz, P., and Godin, C. (2012). L-Py: an L-system simulation framework for modeling plant architecture development based on a dynamic language. *Front. Plant Sci.* 3 (76), 1–20. doi: 10.3389/fpls.2012.00076

Buckley, T. N., Turnbull, T. L., and Adams, M. A. (2012). Simple models for stomatal conductance derived from a process model: cross-validation against sap flux data. *Plant Cell Environ.* 35, 1647–1662. doi: 10.1111/j.1365-3040.2012.02515.x

Campbell, G. S., and Norman, J. M., (1998). *An introduction to environmental biophysics.* 2nd edition. New York: Springer-Verlag, 286. doi: 10.1007/978-1-4612-1626-1

Cescatti, A. (1997). Modelling radiative transfer in discontinuous canopies of asymmetric crowns. I. Model structure and algorithms. *Ecol. Model.* 101, 263–274. doi: 10.1016/S0304-3800(97)00050-1

Churkina, G., Schimel, D., Braswell, B. H., and Xiao, X. (2005). Spatial analysis of growing season length control over net ecosystem exchange. *Global Change Biol.* 11 (10), 1777–1787. doi: 10.1111/j.1365-2486.2005.001012.x

Dauzat, J., Franck, N., Rapidel, B., Luquet, D., and Vaast, P. (2007). "Simulation of echophysiological processes on 3D virtual stands with the ARCHIMED simulation platform," in *Second international symposium on plant growth modeling, simulation, visualization and applications.* (Beijing, China: IEEE), 101–108. doi: 10.1109/PMA.2006.52

De Pury, D. G. G., and Farquhar, G. D. (1997). Simple scaling of photosynthesis from leaves to canopies without the errors of big-leaf models. *Plant Cell Environ.* 20, 537–557. doi: 10.1111/j.1365-3040.1997.00094.x

Duursma, R. A., and Medlyn, B. E. (2012). MAESPA: a model to study interactions between water limitation, environmental drivers and vegetation function at tree and stand levels, with an example application to $[CO_2]\times$ drought interactions. *Geosci. Model Dev.* 5, 919–940. doi: 10.5194/gmd-5-919-2012

Evers, J. B., Letort, V., Renton, M., and Kang, M. (2018). Computational botany: advancing plant science through functional-structural plant modelling. *Ann. Bot.* 121, 767–772. doi: 10.1093/aob/mcy050

Farquhar, G. D., von Caemmerer, S., and Berry, J. A. (1980). A biochemical model of photosynthetic CO_2 assimilation in leaves of C_3 species. *Planta* 149, 78–90. doi: 10.1007/BF00386231

Friend, A. (2001). Modelling canopy CO_2 fluxes: are 'big-leaf' simplifications justified? *Global Ecol. Biogeo.* 10, 603–619. doi: 10.1046/j.1466-822x.2001.00268.x

Gastellu-Etchegorry, J.-P., Grau, E., and Lauret, N. (2012). DART: a 3D model for remote sensing images and radiative budget of earth images. In *Modeling and simulation in engineering*. Ed. C. Alexandru (Rijeka, Croatia: InTech), 29–68.

Ginzburg, L. R., and Jensen, C. X. J. (2004). Rules of thumb for judging ecological theories. *Trends Ecol. Evol.* 19, 121–126. doi: 10.1016/j.tree.2003.11.004

Gueymard, C. A. (2003). Direct solar transmittance and irradiance predictions with broadband models. Part I: detailed theoretical performance assessment. *Solar Energy* 74, 355–379. doi: 10.1016/S0038-092X(03)00195-6

Hemmerling, R., Kniemeyer, O., Lanwert, D., Kurth, W., and Buck-Sorlin, G. (2008). The rule-based language XL and the modelling environment GroIMP illustrated with simulated tree competition. *Funct. Plant Biol.* 35, 739–750. doi: 10.1071/FP08052

Henke, M., and Buck-Sorlin, G. H. (2018). Using a full spectral raytracer for calculating light microclimate in functional-structural plant modelling. *Comput. Inf.* 36, 1492–1522. doi: 10.4149/cai_2017_6_1492

Henke, M., Kurth, W., and Buck-Sorlin, G. H. (2016). FSPM-P: towards a general functional-structural plant model for robust and comprehensive model development. *Front. Comp. Sci.* 10, 1103–1117. doi: 10.1007/s11704-015-4472-8

Holzinger, A. (2005). Usability engineering methods for software developers. *Commun. ACM* 48, 71–74. doi: 10.1145/1039539.1039541

Iqbal, M. (2012). *An introduction to solar radiation*. Burlington: Elsevier Science.

Jarvis, P. G., and McNaughton, K. (1986). "Stomatal control of transpiration: scaling up from leaf to region," in *Advances in ecological research,* vol. 15. (London, U.K.: Academic Press), 1–49. doi: 10.1016/S0065-2504(08)60119-1

Johnson, I. R. (2010). *PlantMod: exploring the physiology of plant canopies.* Dorrigo, NSW, Australia: Tech. rep., IMJ Software. URL www.imj.com.au/software/plantmod.

Kahlen, K., and Stützel, H. (2011). Modelling photo-modulated internode elongation in growing glasshouse cucumber canopies. *New Phytol.* 190, 697–708. doi: 10.1111/j.1469-8137.2010.03617.x

Karwowski, R., and Prusinkiewicz, P. (2003). Design and implementation of the L+C modeling language. *Electron. Notes Theor. Comput. Sci.* 86, 134–152. doi: 10.1016/S1571-0661(04)80680-7

Lawrence, D., Fisher, R., Koven, C., Oleson, K., Swenson, S., and Vertenstein, M. (2019). *CLM5 documentation*. Tech. rep., Boulder, CO: National Center for Atmospheric Research.

Lewis, P. (1999). Three-dimensional plant modelling for remote sensing simulation studies using the Botanical Plant Modelling System. *Agronomie* 19, 185–210. doi: 10.1051/agro:19990302

Marschner, S., and Shirley, P. (2015). *Fundamentals of computer graphics.* (Boca Raton, FL: A K Peters/CRC Press), 748.

Meyers, T. P., and Paw U, K. T. (1987). Modelling the plant canopy micrometeorology with higher-order closure principles. *Agric. For. Meteorol.* 41, 143–163. doi: 10.1016/0168-1923(87)90075-X

Mott, K. A., and Buckley, T. N. (2000). Patchy stomatal conductance: emergent collective behaviour of stomata. *Trends Plant Sci.* 5 (6), 258–262. doi: 10.1016/S1360-1385(00)01648-4

Pearcy, R. W., and Yang, W. (1996). A three-dimensional crown architecture model for assessment of light capture and carbon gain by understory plants. *Oecologia* 108, 1–12. doi: 10.1007/BF00333208

Ponce de León, M. A., and Bailey, B. N. (2019). Evaluating the use of Beer's law for estimating light interception in canopy architectures with varying heterogeneity and anisotropy. *Ecol. Model.* 406, 133–143. doi: 10.1016/j.ecolmodel.2019.04.010

Pradal, C., Dufour-Kowalski, S., Boudon, F., Fournier, C., and Godin, C. (2008). OpenAlea: a visual programming and component-based software platform for plant modelling. *Funct. Plant Biol.* 35, 751–760. doi: 10.1071/FP08084

Prata, A. J. (1996). A new long-wave formula for estimating downward clear-sky radiation at the surface. *Q.J.R. Meteorol. Soc.* 122, 1127–1151. doi: 10.1002/qj.49712253306

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical recipes: the art of scientific computing.* (Cambridge, U.K.: Cambridge University Press), 1256.

Prusinkiewicz, P., and Runions, A. (2012). Computational models of plant development and form. *New Phytol.* 193, 549–569. doi: 10.1111/j.1469-8137.2011.04009.x

Raupach, M., and Finnigan, J. (1988). 'Single-layer models of evaporation from plant canopies are incorrect but useful, whereas multilayer models are correct but useless': discuss. *Aust. J. Plant Physiol.* 15, 705–716. doi: 10.1071/PP9880705

Reichstein, M., Falge, E., Baldocchi, D., Papale, D., Aubinet, M., Berbigier, P., et al. (2005). On the separation of net ecosystem exchange into assimilation and ecosystem respiration: review and improved algorithm. *Global Change Biol.* 11 (9), 1424–1439. doi: 10.1111/j.1365-2486.2005.001002.x

Ross, J. (1981). *The radiation regime and architecture of plant stands*. The Hague, The Netherlands: Dr. W. Junk Publishers, 424. doi: 10.1007/978-94-009-8647-3

Sellers, P., Mintz, Y., Sud, Y., and Dalcher, A. (1986). A simple biosphere model (SiB) for use within general circulation models. *J. Atmos. Sci.* 43, 505–531. doi: 10.1175/1520-0469(1986)043<0505:ASBMFU>2.0.CO;2

Sinclair, T. R., Murphy, C. E., and Knoerr, K. R. (1976). Development and evaluation of simplified models for simulating canopy photosynthesis and transpiration. *Brit. Ecol. Soc.* 13, 813–829. doi: 10.2307/2402257

Sinoquet, H., Le Roux, X., Adam, B., Ameglio, T., and Daudet, F. A. (2001). RATP: a model for simulating the spatial distribution of radiation absorption, transpiration and photosynthesis within canopies: application to an isolated tree crown. *Plant Cell Environ.* 24, 395–406. doi: 10.1046/j.1365-3040.2001.00694.x

Suffern, K. G. (2007). *Ray tracing from the ground up*. (Boca Raton, FL: A K Peters/CRC Press), 784.

Valladares, F. (2003). "Light heterogeneity and plants: from ecophysiology to species coexistence and biodiversity," in *Progress in botany* (Berlin Heidelberg: Springer-Verlag), 439–471. doi: 10.1007/978-3-642-55819-1_17

Vezy, R., Christina, M., Rouspard, O., Nouvellon, Y., Duursma, R., Medlyn, B., et al. (2018). Measuring and modelling energy partitioning in canopies of varying complexity using MAESPA model. *Agric. For. Meteorol.* 253-254, 203–217. doi: 10.1016/j.agrformet.2018.02.005

Viswanadham, Y. (1981). The relationship between total precipitable water and surface dew point. *J. Appl. Meteorol.* 20, 3–8. doi: 10.1175/1520-0450(1981)020<0003:TRBTPW>2.0.CO;2

Wang, Y. P., and Jarvis, P. G. (1990). Description and validation of an array model—MAESTRO. *Agric. For. Meteorol.* 51, 257–280. doi: 10.1016/0168-1923(90)90112-J

Wang, Y. P., and Leuning, R. (1998). A two-leaf model for canopy conductance, photosynthesis and partitioning of available energy I: model description and comparison with a multi-layered model. *Agric. For. Meteorol.* 91, 89–111. doi: 10.1016/S0168-1923(98)00061-6

Weber, J., and Penn, J. (1995). "Creation and rendering of realistic trees," in *SIGGRAPH '95 Proceedings of the 22nd annual conference on computer graphics and interactive techniques* (New York: ACM), 119–128. doi: 10.1145/218380.218427

Woods, H. A., Saudreau, M., and Princebourde, S. (2018). Structure is more important than physiology for estimating intracanopy distributions of leaf temperatures. *Ecol. Evol.* 8, 5206–5218. doi: 10.1002/ece3.4046

Conflict of Interest: The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

APPENDIX 1. PHOTOSYNTHESIS MODEL EQUATIONS

Diffusion of CO_2 into and out of the leaf is modeled using the following resistance analogy:

$$A + R_d = 0.75 g_M (C_a - C_i), \tag{3}$$

where is the CO_2 assimilation rate, R_d is the respiration rate, and $0.75g_M$ is the conductance to CO_2 transport between the mesophyll and ambient air, which is assumed to differ from water vapor diffusion only based on its lower diffusivity in air (and hence the 0.75 factor). Diffusion is driven by a difference between the intercellular CO_2 concentration C_i and the air CO_2 concentration C_a.

The CO_2 assimilation rate is calculated following Farquhar et al. (1980) as

$$A = \left(1 - \frac{\Gamma^\star}{C_i}\right)\min\{W_c, W_j\} - R_d, \tag{4}$$

where $\Gamma^\star$ is the chloroplastic CO_2 compensation point. A is limited either by the Rubisco-limited carboxylation rate W_c, or by the rate of RuBP regeneration W_j (we neglect the TPU limitation state for conditions typical of the natural environment). W_c is calculated according to

$$W_c = \frac{V_{cmax} C_i}{C_i + K_c(1 + O/K_O)}, \tag{5}$$

where $V_{c,max}$ is the maximum carboxylation rate, K_c is the Michaelis-Menten constant for RuBP carboxylation, K_o is the Michaelis-Menten constant for oxygenation, and O is the partial pressure of oxygen in the air. W_j is calculated according to

$$W_c = \frac{JC_i}{4C_i + 8\Gamma^\star}, \tag{6}$$

where the potential electron transport rate J is modeled using the hyperbolic relationship

$$J = \frac{J_{max} Q\alpha}{Q\alpha + J_{max}}, \tag{7}$$

where J_{max} is the value of J at saturating Q, and α describes the rate at which J reaches J_{max} with increasing Q.

In addition to the original formulation proposed by Farquhar et al. (1980), the temperature dependence of model parameters has also been included following the description given by Bernacchi et al. (2001) and Bernacchi et al. (2003), which are given by the following equations:

$$\Gamma^* = \exp(19.02 - 37.83/(RT_L)), \tag{8a}$$

$$K_c = \exp(38.05 - 79.43/(RT_L)), \tag{8b}$$

$$K_O = \exp(20.30 - 36.38/(RT_L)), \tag{8c}$$

$$R_d = R_{d,25}\exp(18.72 - 46.39/(RT_L)), \tag{8d}$$

$$V_{cmax} = V_{cmax,25}\exp(26.35 - 65.33/(RT_L)), \tag{8e}$$

$$J_{cmax} = J_{max,25}\exp(C_{Jmax} - \Delta H_{a,Jmax}/(RT_L)), \tag{8f}$$

where R is the universal gas constant, T_L is the leaf temperature in absolute units, and the subscript 25 indicates the evaluation of the parameter at a temperature of 25°C. This leaves the following free parameters to be specified in the photosynthesis model: $R_{d,25}$, $V_{cmax,25}$, $J_{max,25}$, α, C_{Jmax}, and $\Delta H_{a,Jmax}$.

Remote Sensing of Environment 208 (2018) 133–144

Contents lists available at ScienceDirect

Remote Sensing of Environment

journal homepage: www.elsevier.com/locate/rse

Semi-direct tree reconstruction using terrestrial LiDAR point cloud data

Brian N. Bailey*, Miguel H. Ochoa

Department of Plant Sciences, University of California, Davis, Davis, CA, USA

ARTICLE INFO

Keywords:
Leaf angle distribution function
Plant architecture
Plant reconstruction
Terrestrial LiDAR

ABSTRACT

A new method was developed for reconstructing the geometric structure of large plants such as trees at the leaf-scale by utilizing terrestrial LiDAR data. The primary goal of the work was to develop a feasible means for accurately and rapidly reconstructing or "digitizing" entire trees in order to specify the position, orientation, and size of every leaf in digital tree models that provide geometric inputs for high-resolution biophysical models or analyses. As with any optical measurement technique, a primary challenge is accurately accounting for plant matter that is occluded from view of the sensor. The present method is termed "semi-direct" because it uses a triangulation procedure to approximately directly reconstruct as many leaves as possible that are in view of the scanner. For plant matter obstructed from view, a statistical backfilling procedure was used to add additional leaves such that the three-dimensional distribution of leaf area and orientation of the reconstructed plant matched that of the actual plant on average. In a best case scenario such as when leaf density is low, nearly all leaf area is directly reconstructed from the scan and the branch and clumping structure is preserved within the reconstruction. In the worst case scenario such as when the leaf density is very high and nearly all leaves are occluded from view of the scanner, only a small fraction of leaves can be directly reconstructed, but at a minimum the distribution of leaf area and the leaf angle distribution across the reconstructed plant will be consistent with that of the actual plant. Unlike many other approaches, the present method does not rely on the woody matter of the plant to provide a skeleton for reconstruction, and can be used in dense plants where little woody matter is visible from the scanner.

1. Introduction

Leaf-level measurements of many biophysical processes (e.g., exchange of water vapor, CO_2, and heat) have become routine, yet scaling these processes up to entire plants and canopies remains a considerable challenge, as performing direct measurements of biophysical processes at these scales if often not possible (Amthor, 1994; Ehleringer, 2000). Instead, our understanding of whole-plant and -canopy biophysical processes typically relies on models that attempt to aggregate information originating at the leaf scale into plant communities. Such models make simplifying assumptions that focus on bulk canopy behavior, such as "big leaf" or "multilayer" models (Sinclair et al., 1976; Amthor, 1994; DePury and Farquhar, 1997). Given the scale of canopy representation in these models, inputs are also typically bulk values specified at or near the canopy scale.

With the continued exponential increase in computational performance (Moore, 1965), we are now in a position where direct scaling from leaves to canopies (i.e., representing every leaf in a canopy) is within reach. High-resolution, three-dimensional models are becoming increasingly common, and are able to represent an incredibly wide range of scales (e.g., Bailey et al., 2014, 2016; Bailey, 2018). The next generation of biophysical models are likely to shed new light on how processes at various scales interact to determine plant behavior over communities.

A considerable challenge in the utilization of such models is the accurate specification of geometric inputs. As the goal of these models is to explicitly represent heterogeneity at various scales and its impact on canopy-level processes, we must be able to accurately measure and input this geometry into the models (Vos et al., 2010; Sarlikioti et al., 2011). Manual measurement of canopy geometry is far too time consuming to be useful at providing canopy-level inputs at the leaf scale.

Remote sensing techniques have provided a means for rapidly measuring and recording the full three-dimensional geometry of plants for use in computer models (i.e., "digitizing"). These techniques make a compromise between level of detail and the size of system that can be represented. Various methods are available to extract plant-scale structural parameters such as crown diameter and height from remote measurements (e.g., Morsdorf et al., 2004; Henning and Radtke, 2006; Rosell et al., 2009; Yang et al., 2013). The clear advantage of these approaches is that they can be used to rapidly measure large spatial

---

* Corresponding author.
E-mail address: bnbailey@ucdavis.edu (B.N. Bailey).

https://doi.org/10.1016/j.rse.2018.02.013
Received 8 September 2017; Received in revised form 28 December 2017; Accepted 11 February 2018

scales, but they do not provide detailed information at the sub-plant scale that may be needed for high-resolution modeling. At the opposite end of the spectrum, methods are also available to measure the full plant structure at the leaf scale. Early work by Sinoquet et al. (1998) used an electromagnetic instrument to directly record the position and orientation of individual foliage elements, which is limited by the need to manually place the instrument next to each leaf. Previous workers have also been relatively successful in using photographic methods to directly reconstruct small plants where nearly all foliage is in direct view of cameras placed on the perimeter of the plant (e.g., Delagrange and Rochon, 2011; Li et al., 2013; Pound et al., 2014). However, these methods cannot be used directly with large plants where a significant portion of plant area is occluded from view.

For large plants such as trees, the problem of measuring the full vegetative structure is complicated by the sheer size of the plants, number of leaves, and potentially large fraction of leaves occluded from view of a remote sensor. If only the woody structure of the tree is of interest, the occlusion problem becomes much less substantial. Numerous methods have been developed based on laser scanning that use the woody structure of the plant as a road map through laser scanning point clouds (e.g., Binney and Sukhatme, 2009; Xu et al., 2007; Côté et al., 2009; Raumonen et al., 2013; Hackenberg et al., 2015; Mèndez et al., 2016). Starting at the trunk, branches can be traced throughout the tree using point connectivity information, which can then be used to generate a reconstruction of the woody tree structure.

If reconstructions of trees at the leaf scale are desired, the occlusion problem must be somehow confronted. Often this involves measurement of the overall tree structure and making reasonable guesses as to where individual leaves should be placed. For example, Shlyakhter et al. (2001) used an aggregate approach based on photographic methods to determine the general shape of tree crowns, and then used a structural model to create a simulated tree that fit within the measured crown shape. In cases where vegetation is sparse or leaf-off measurements are available, a reconstruction of the woody structure can be used as a "skeleton" to guide the placement of individual leaves (e.g., Xu et al., 2007; Côté et al., 2009, 2011). Delagrange and Rochon (2011) demonstrated the possibility of adding leaves to the branch skeleton using allometric relations, but this method relies on empirical relations that may or may not be generally applicable.

Evaluations of plant reconstruction methods are most commonly performed using visual comparisons, as it is difficult to quantitatively evaluate their accuracy given that measurements of the true plant structure is typically not available. While many reconstruction methods produce tree models that appear visually reasonable, it is unclear whether the reconstructions are accurate enough for use in detailed model simulations. Côté et al. (2009) noted that reconstructed plants should be "radiatively consistent" with the actual plants, meaning that radiative transport through the reconstructed plants should be approximately equivalent to that of the actual plants. Côté et al. (2009) were able to produce tree reconstructions for *Pinus* species that demonstrated radiative consistency based on measurements of radiation reflection and transmission.

In this work, we develop a "semi-direct" method that uses terrestrial LiDAR data to reconstruct large plants such as trees that match the three-dimensional leaf area and angle distribution of the actual plant being reconstructed. The method is semi-direct in that it directly reconstructs the majority of leaves that are in direct view of the LiDAR scanner. The method then uses a statistical backfilling approach to recreate occluded leaves in a manner that ensures the overall leaf area and angle distribution matches that of the actual plant. Since the reconstructed leaf area and angle distributions are consistent with the actual trees, the reconstructions are applicable for use in model simulations of processes such as light interception.

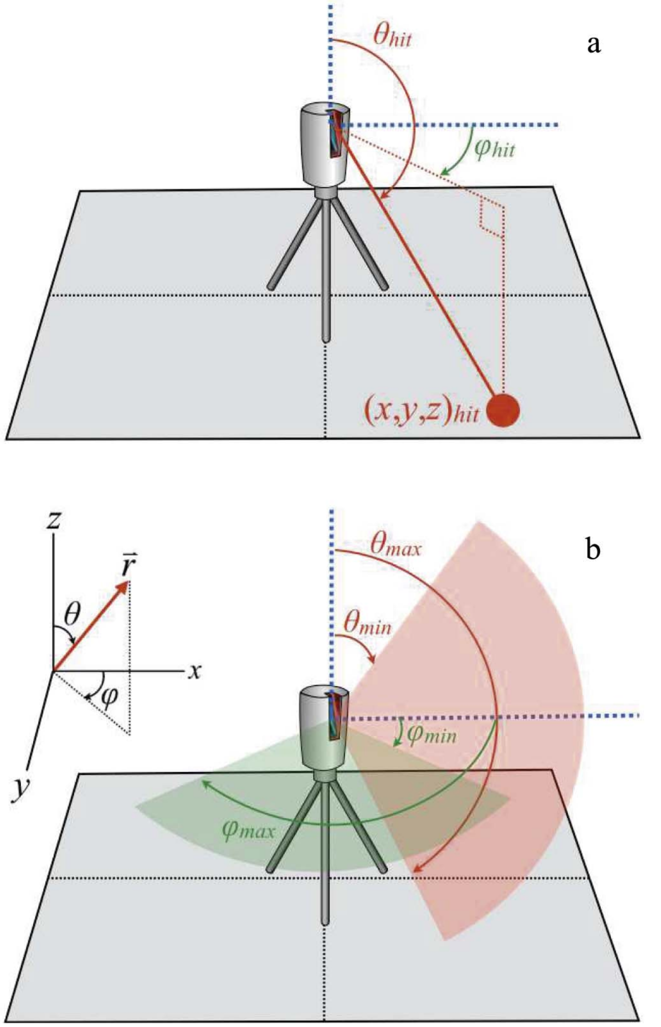

Fig. 1. Schematic depiction of terrestrial LiDAR scanning. (a) Cartesian coordinate $(x,y,z)_{hit}$ of hit point, and corresponding spherical coordinate (θ,φ). (b) Scanning pattern in spherical coordinates, illustrating the range of scan zenithal angles (θ_{min} through θ_{max}) and azimuthal angles (φ_{min} through φ_{max}).

2. Method description

2.1. Terrestrial LiDAR scanning

Typical terrestrial LiDAR scanning instruments are compact units that can be mounted on a tripod, and are used to measure the distance to surrounding objects. The instrument emits a large number of concentrated pulses or beams of radiation into the surrounding spherical space. In the event that a beam intersects solid matter, some fraction of the radiation beam is scattered back to the instrument. Using various methods such as time of flight, the instrument can calculate and record the distance to beam-object intersection points. The direction in which the pulse was sent is also known by the instrument, which allows calculation of the Cartesian (x,y,z) position of beam-object intersection points (Fig. 1a). By emitting millions of beams into the surrounding space, the instrument effectively maps the three-dimensional geometry surrounding the scan location.

Terrestrial LiDAR instruments generally do not emit beams at random, rather they perform a systematic scan of the surrounding spherical space. Most commonly, instruments discretely scan a certain range of zenithal angles while continuously rotating between a range of discrete azimuthal angles (Fig. 1b). This creates an approximately uniform two-dimensional grid of points in spherical space. The scan

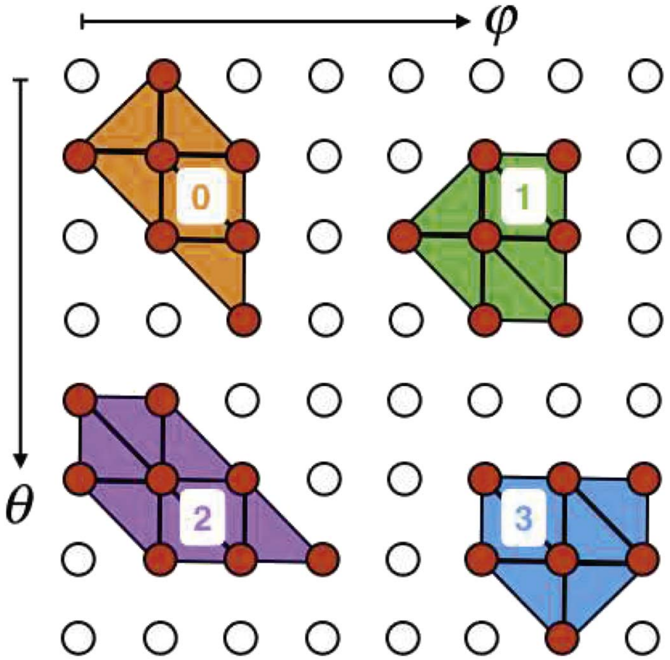

Fig. 2. "Flood-fill" grouping of triangles. A two-dimensional grid of scan points in θ–φ space is shown, with "misses" denoted by open circles and "hits" denoted by filled circles. Connected triangle groups are identified and assigned a group identifier. In the example shown, four continuous triangle groups are formed, which are given identifiers of 0, 1, 2, and 3.

resolution is given by the number of discrete scan zenithal directions N_θ (# rows), and the number of discrete scan azimuthal directions N_φ (# columns), with $N_\theta \times N_\varphi$ being the total number of points in the scan.

2.2. Scan point triangulation

The basic idea behind the plant reconstruction methodology presented in this work is to connect adjacent scan hit points to form triangles, then identify continuous triangle groups that reconstruct individual leaves. The triangulation methodology is described in detail by Bailey and Mahaffee (2017b), and a brief description is repeated below.

The triangulation algorithm first seeks to construct a two-dimensional grid of scan points in spherical space. This grid consists of a (θ,φ) coordinate for each ray sent by the scanner (Fig. 1). This creates a two-dimensional plane of points that can be triangulated (Fig. 2). Bailey and Mahaffee (2017b) suggested an efficient triangulation algorithm that can be used when the indices of the scan points in the 2D spherical grid are recorded by the scanner. This allows for the construction of a "scan table" in which rows correspond to each scan zenithal angle, and columns correspond to each scan azimuthal angle. Given this table, it is relatively straightforward to form triangles between adjacent points in the uniform grid since scan point connectivity is already known. For instruments that do not directly record this information (such as the instrument used in this work), standard 2D Delaunay triangulation can be used (Press et al., 2007), which has the trade-off that it requires more computational effort since point connectivity is not initially known. Triangles exceeding a size or aspect ratio threshold are rejected to prevent erroneous triangles from being formed, such as triangles that connect adjacent leaves. Since each triangle vertex corresponds to a laser hit point, the (x,y,z) coordinates of the vertices are also known. The resulting triangulation gives a set of triangles that follow the surfaces of individual leaves that are in view of the scanner.

2.3. Direct leaf surface reconstruction

Neighboring triangles are connected to form continuous groups, where each group presumably corresponds to all or a portion of an individual leaf's surface. To accomplish this, an algorithm is applied that is similar to a traditional "flood-fill" algorithm (e.g., Lee, 1987), except that it connects adjacent triangles instead of adjacent pixels (Fig. 2). For each triangle, any neighboring connected triangles are identified, where a "connected" triangle is defined as a triangle that shares two vertices with the current triangle being examined. By requiring that two vertices are shared rather than one, this reduces the likelihood that adjacent leaves or branches will inadvertently be merged into a common group. The algorithm begins by iterating over each triangle in the triangulated set. The first triangle is assigned a fill group identifier of "0". For each triangle, any neighboring connected triangles are determined. If any connected triangles exist, each connected triangle is added to the current fill group by assigning it the current group identifier, and the neighbors of each connected triangle are examined in a recursive manner. The recursion halts when there are no connected triangles that have not yet been added to the current fill group. In this case, the current fill group has been completed, and the fill group identifier is increased by one. The original iteration over triangles proceeds, where triangles that have already been assigned to a fill group are skipped. Once the iteration is completed, all possible triangle groups have been formed (Fig. 2).

Triangle groups are filtered by their area to exclude very small or large groups. If only one to a few small triangles are identified in a single group, it is typically not desirable to allocate an entire leaf to this group. These small groups are filtered by specifying a threshold value for the minimum group surface area, below which groups are rejected. Similarly a threshold value is specified for the maximum group surface area, which is typically set to be much larger than the expected area of a single leaf. The primary purpose of filtering large leaf groups is to remove outliers when calculating the characteristic leaf dimension (see below).

Each continuous fill group is then replaced by a "prototype" leaf. Although there are many ways a prototype leaf could be specified (e.g., a rectangle, a triangular mesh), this work used a PNG image to define the leaf shape (Fig. 3). A leaf is specified by a planar rectangle, but a portion of that rectangle is removed according to the transparency channel of the PNG image (Bailey, 2018). The length and width of the prototype are denoted by l and w, and the fraction of the total rectangular area that is not transparent is the solid fraction s (Fig. 3).

There are three quantities that must be specified for each leaf: its (x,y,z) position, size, and orientation. The position and average orientation are readily available from the triangulation; the leaf is placed at the location of the triangle group centroid and oriented in the direction of the average triangle group normal. However, the size is more difficult to determine, because only a relatively few number of leaves on the outside of the plant in full view of the scanner will be completely reconstructed by the triangulation. Most of the leaves are occluded to some degree and will only be partially triangulated, and thus the area of the fill groups will be less than the actual leaf area. One could perform manual measurements of leaf size using a ruler to obtain representative values for leaf sizes. The drawback of this method, aside from having to perform manual measurements, is that leaf size can change with position in the plant and thus specifying a single size value may not be representative. The method used here involved considering only the largest triangulated groups (e.g., 10 largest groups), and taking the characteristic leaf length L to be the average of the square root of the group areas. The spatial distribution of leaf size can be approximately represented by dividing the plant into sub-volumes, and the largest triangulation groups in each volume can be used to determine the representative leaf size for that particular volume. In order to specify the dimension of a leaf from the characteristic leaf size L, we must specify a leaf aspect ratio, which is the ratio r of the length of the leaf parallel (l) to perpendicular (w) to the midrib. Given that $L \equiv \sqrt{a} = \sqrt{wls}$ and $r \equiv l/w$, the leaf length l is equal to $L\sqrt{r/s}$, and $w = l/r$.

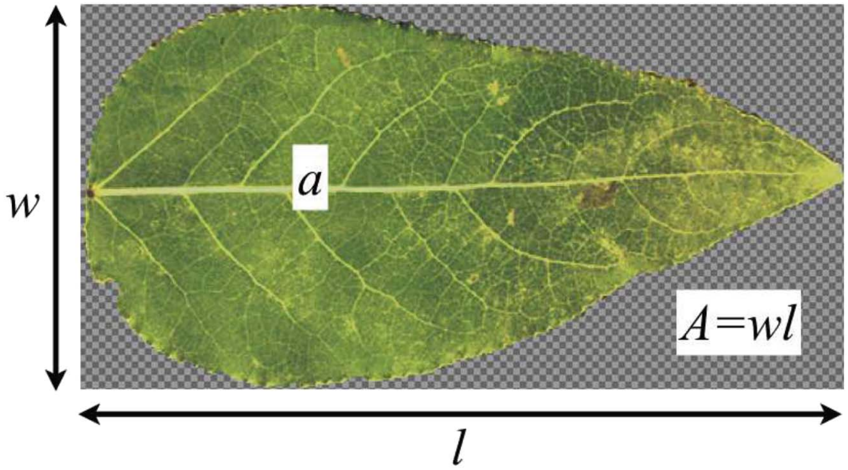

Fig. 3. Example of leaf prototype image. The solid portion of the image is colored, while the checkered portion of the image is considered transparent. The area of the solid portion is a, and the area of the total image is $A = wl$, where w and l are respectively the width and length of the prototype. The fraction of the total image that is solid is $s = a/A$.

2.4. *Backfilling occluded leaves*

Direct leaf reconstruction based on the triangulation only represents a subset of the total leaf area. The leaf area that is not triangulated because it is occluded or because the triangulation failed must be represented through other means. In the present method, the remaining leaf area is reconstructed by backfilling leaves until the leaf area density of the reconstructed plant matches that of the actual plant. The plant is discretized into a grid of rectangular sub-volumes called voxels (see Bailey and Mahaffee, 2017a), and LiDAR points are grouped by the voxel in which they reside. The method described in detail by Bailey and Mahaffee (2017a) can then be used to calculate the leaf area density and leaf angle distribution of the actual plant for each voxel using the LiDAR scan data. More precisely, it should be noted that the method actually measures the area density of all plant matter including branches. This method gives a relatively accurate measure of the total surface area of plant matter within each voxel for the actual plant. It is also straightforward to use the directly reconstructed leaves from Section 2.3 to determine the amount of leaf area in each voxel resulting from the direct reconstruction, as the area of each reconstructed leaf is known. The difference between the total and directly reconstructed area is the amount of leaf area that remains to be added through backfilling.

The backfilling process begins by randomly choosing a directly reconstructed leaf within a given voxel, which is duplicated and placed at a random, uniformly distributed position within the voxel. This process continues for each voxel until the reconstructed leaf area in the voxel matches the “actual” leaf area. It is possible that too much leaf area could have been added during the direct reconstruction, in which case leaf area can be removed by randomly deleting leaves which we term “thinning”. Based on this process, the resulting reconstructed leaf area and leaf angle distribution should be consistent with that of the actual plant for each voxel. This method is dispersive in that it tends to spread out leaves in space. The larger the fraction of leaves that are directly reconstructed, the less dispersive the reconstruction method becomes, and the better the reconstructed tree will match the structure of the actual tree.

2.5. *Woody plant material*

Several methods have been suggested by previous authors for the reconstruction of woody plant material (e.g., Xu et al., 2007; Binney and Sukhatme, 2009; Mèndez et al., 2016; Li et al., 2016). In this work, we focus only on reconstructing leaves within the crown volume, and present a simple method for reconstructing the main trunk similar to that of Xu et al. (2007). The primary purpose of representing the main trunk is simply to provide a visual reference for qualitative evaluation of the reconstruction. A voxel is specified that contains the portion of trunk to be reconstructed. Hit points within this voxel are triangulated, and the flood-fill algorithm of Section 2.3 is applied. The largest fill group is identified, which is assumed to correspond to the trunk. This produces a triangular mesh that approximately reconstructs the portion of the trunk visible from the scanner.

It should also be noted that it is possible that the reconstruction algorithm for leaves could inadvertently identify branches as a leaf group. Rather than attempting to filter out these relatively rare instances, the algorithm is simply applied in the same way as for leaves, and it is assumed that a reconstructed branch is a reasonable location to place a leaf. This work focuses on trees in which the (visible) leaf area is much larger than the woody area. For trees where the woody area is substantial compared to the leaf area, LiDAR hit points corresponding to woody material could be separated within the scan (Béland et al., 2014), and a branch reconstruction algorithm could be applied separate from the leaf reconstruction method presented in this work.

2.6. *Multiple scan positions*

To reconstruct an entire tree, scans from multiple locations surrounding the tree are typically required and must be combined. Generally, the instrumentation on-board the scanner for measuring geographic position is not accurate enough to be used to merge multiple scans (it provides only an estimate). Standard methods are available to register multiple scans to a common global coordinate system, such as the iterative closest point (ICP) method (Zhang, 1994), or methods that use reflectors, checkerboards, spheres, or other common targets placed within the scan. Many instruments also come with software developed by the manufacturer that use proprietary algorithms.

The method for calculating the leaf area contained within each voxel (Bailey and Mahaffee, 2017a) does not distinguish between different scan positions, thus aggregating multiple scans is straightforward. For any given ray direction, the probability that a ray intersects vegetation, the leaf normal vector, and path length through the voxel are simply added to running totals for all scans. The totals for all scan points from all scan locations are used along with Beer's law to solve for leaf area density within the voxel (Bailey and Mahaffee, 2017a). For the leaf reconstruction procedure, the algorithm is applied on a scan-by-scan basis, and reconstructed leaves from each scan are simply aggregated together to form the reconstructed plant.

3. Evaluation of method

3.1. Data collection and processing details

Scanning data was collected for a 5 m tall Emerald Sunshine Elm (*Ulmus propinqua*) located in Davis, California USA to demonstrate application of the method and evaluate its performance. The tree was scanned using a full-waveform Riegl VZ-1000 terrestrial LiDAR scanner (RIEGL Laser Measurement Systems GmbH; Horn, Austria). The scanner sends concentrated beams of radiation with a wavelength of 1550 nm in a uniformly gridded pattern in spherical space, covering a range from 30° to 130° in the zenithal direction and 0 to 360° in the azimuthal direction. The maximum scan resolution is about 41,000 × 150,000 points in the zenithal × azimuthal directions. The beam diameter as it leaves the instrument is approximately 7 mm, which diverges at an angle of approximately 0.3 mrad, meaning that at 10 m range the beam diameter is roughly 8.5 mm. The instrument can scan up to 122,000 points per second, with a range from 2.5 m up to approximately 350–450 m at this scanning rate. The full-waveform LiDAR instrument used can record multiple hit points per pulse, but the point cloud was filtered to consider only the closest hit per pulse. The instrument was equipped with an on-board digital camera (Nikon D810 36 Mega Pixel) that was used to assign RGB color values to each scan point and obtain images for visual comparison with reconstructions.

Four scans were performed at equally spaced intervals surrounding the tree, which were automatically registered to a common coordinate system using Riegl's proprietary RiSCAN Pro software. The scanner was positioned on a tripod approximately 1.25 m above the ground, and approximately 5.5 m from the trunk of the tree. This distance was chosen because it was as close as possible to the tree such that the entire tree was in view of the scanner and digital camera. A modest scan resolution of 2500 × 4500 points (zenith × azimuth) was chosen. At 10 m range, this meant that adjacent points on a surface orthogonal to the beam direction were separated by roughly 3.5–7 mm and 7–14 mm in the zenithal and azimuthal directions, respectively, depending on beam zenithal angle. Given the chosen resolution, the scans took roughly 2 min to complete, with an additional 2–3 min for GPS location and collection of digital photographs. Scans were performed under very low wind speed conditions to minimize leaf disturbances. The above scanning configuration worked well for the particular application of interest, but in general configurations are expected to be application-dependent. Since point density effectively decreases with distance, trees that are larger or further away will require a higher scanning density. Additionally, very large or dense trees could require more scans, potentially at multiple heights to ensure that all portions of the tree are in view of the scanner.

Additionally, the size of 40 random leaves were measured to evaluate the performance of the method for determining the leaf dimensions from the LiDAR data. The lengths of the leaves parallel and perpendicular to the midrib were measured and recorded for each of the 40 leaves. Admittedly, a robust sampling strategy was not used, and only leaves within reach of the ground were measured. This is because only a rough estimate of leaf size was desired in order to assess whether results of the LiDAR method were at least reasonable. Alternatively, a more robust quantification of errors in leaf dimension is presented in Section 4.3 using synthetic data.

For processing the data, a uniformly spaced 3D grid of voxels was overlaid on the tree, within which leaf area was calculated using the method described above and by Bailey and Mahaffee (2017a). The tree crown was divided into a 10 × 10 × 10 grid of rectangular voxels, each of size 0.5 × 0.5 × 0.4 m^3. In the triangulation methodology, triangles were rejected if the length of any of their sides exceeded 5 cm, or if their aspect ratio was greater than 10. In the flood-fill algorithm, triangle groups were rejected if their total area was less than 1 cm^2 or greater than 200 cm^2, which were chosen because they are much smaller or larger than the expected area of a leaf. The maximum leaf area threshold is relatively easy to specify since it is straightforward to estimate the maximum expected leaf area. Understanding the minimum leaf area threshold is slightly less straightforward. It may be undesirable to specify a minimum area threshold that is too small because we typically want at least a few connected triangles for each leaf in order to have confidence that the triangle group uniquely corresponds to a leaf. We recommend a minimum threshold that is roughly an order of magnitude smaller than the maximum area threshold. However, we varied the minimum area threshold between 0.1 and 50 cm^2 and found very little impact on the resulting tree reconstructions. Using tighter area thresholds generally results in slightly less directly reconstructed leaf area, but the overall distribution of leaf area and orientation remains the same.

3.2. Generation of synthetic scanning data

Quantitative evaluation of LiDAR data processing methods is extremely difficult when applied to large, dense trees, since there is typically no “gold standard” measurement against which to compare. Before proceeding to the application of the method under field conditions, an alternative approach is presented that uses simulated or “synthetic” LiDAR data in which the exact vegetation structure is known (see also Côté et al., 2009; Mèndez et al., 2013; Raumonen et al., 2013; Bailey and Mahaffee, 2017a,b). This approach was adopted to test the plant reconstruction method's ability to reproduce the distribution of leaf area, orientation, and characteristic size. Admittedly, this method also has its drawbacks, namely that it is for an idealized case. Thus, it clearly does not replace the need to perform some type of field validation, but represents a powerful tool for algorithm testing and evaluation.

The synthetic LiDAR data was produced by performing a ray-tracing simulation that mimics the actual LiDAR scanning procedure described above in Section 3.1. In short, a model or “reference” tree was created based on the architectural model of Weber and Penn (1995), which specifies the position of the trunk, branches, and leaves. The trunk and branches were made up of a mesh of triangular elements, and the leaves were rectangular transparency masks with zero thickness (see Fig. 3) of size 6 × 20 cm^2 and a solid fraction $s = 0.62$. The overall tree was roughly 7.5 m tall with a crown diameter of about 5.5 m, and had branches with a diameter ranging from 0.36 m at the trunk base to zero at the branch tips. The woody structure of the tree was made up of about 77,000 triangles, and the tree had about 30,000 leaves. Leaf orientations were specified as described in Weber and Penn (1995), where leaves tend to rotate around the axial direction of the branches, which leads to interesting non-uniform angle distributions (see Figs. 8 and 9). Rays were launched from each of the four simulated scanner locations in a spherical pattern approximately matching that of an actual LiDAR scan. Ray-object intersection tests were performed to determine the (x,y,z) location of the closest intersection point (Suffern, 2007). Note that for simplicity it was assumed that a ray had an infinitely small diameter that maintains 100% of the emitted intensity, which is not true for an actual LiDAR beam. The resulting field of (x,y,z) intersection points was taken to be an approximation of an actual LiDAR scan, and was used to run the reconstruction methodology. For the simulated tree case, the voxel grid size was slightly different than that of the real tree because the tree crowns were slightly different sizes (but still consisted of 10 × 10 × 10 total voxels). For this case, the voxels had a size of 0.55 × 0.55 × 0.65 m^3. On average, each voxel contained about 30 leaves.

3.3. Error quantification

Errors between exact and simulated data were quantified using three standard metrics: the index of agreement (Willmott, 1981, 1982), root-mean-squared error (RMSE), and mean bias. The index of agreement is defined as

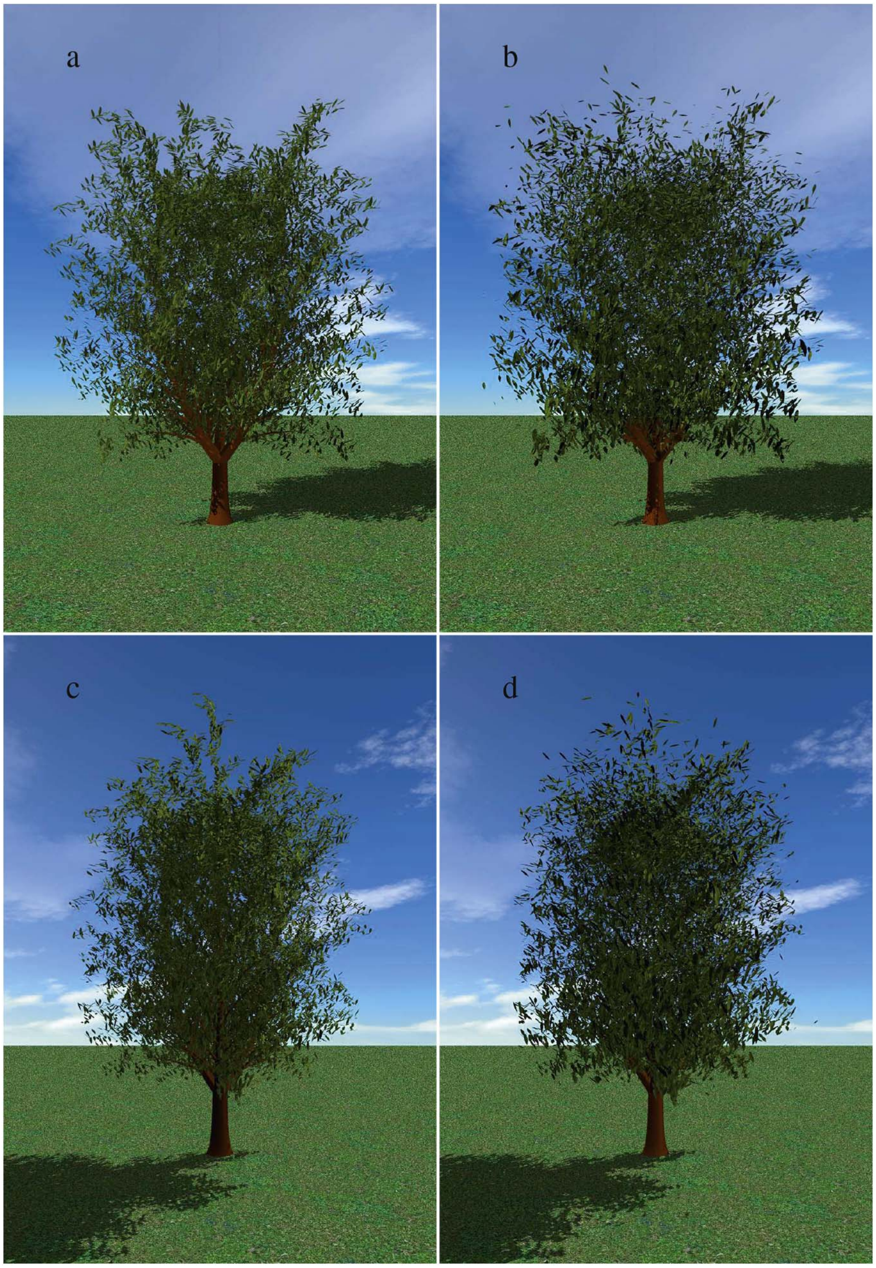

Fig. 4. Visualization of (a, c) computer-generated or "reference" tree, and (b, d) reconstruction of the reference tree based on simulated LiDAR scanning data for two opposing viewing angles.

$$\mathrm{d} = 1 - \frac{\sum_{i=1}^{N}(M_i - L_i)^2}{\sum_{i=1}^{N}(|M_i - \overline{M}| + |L_i - \overline{M}|)^2}, \tag{1}$$

where M_i and L_i are respectively the ith exact and estimated values for each voxel, with N total values, and an overbar denotes an average over all voxels. The RMSE is defined as

$$\mathrm{RMSE} = \left(\frac{1}{N}\sum_{i=1}^{N}(L_i - M_i)^2\right)^{1/2}, \tag{2}$$

and the mean bias is defined as

$$\mathrm{bias} = \frac{1}{N}\sum_{i=1}^{N}(L_i - M_i). \tag{3}$$

4. Evaluation using synthetic scanning data

4.1. Visualization

The visualizations shown in Fig. 4 provide a means for performing a qualitative evaluation of the reconstruction method using the synthetic scanning data. Overall, the reconstruction (Fig. 4b,d) appears visually reasonable in comparison with the reference tree (Fig. 4a,c), and reproduces the general tree structure. Clearly, the reconstruction does not produce an exact replica of the reference tree nor is it intended to do so. As mentioned previously, the reconstruction method is dispersive, meaning that it tends to spread out leaves and diminish structure. As a result, the reconstructed tree has lost some branch and clumping structure compared to the reference tree. The sub-voxel-scale structure that is present is primarily due to directly reconstructed leaves, which are shown in Fig. 5.

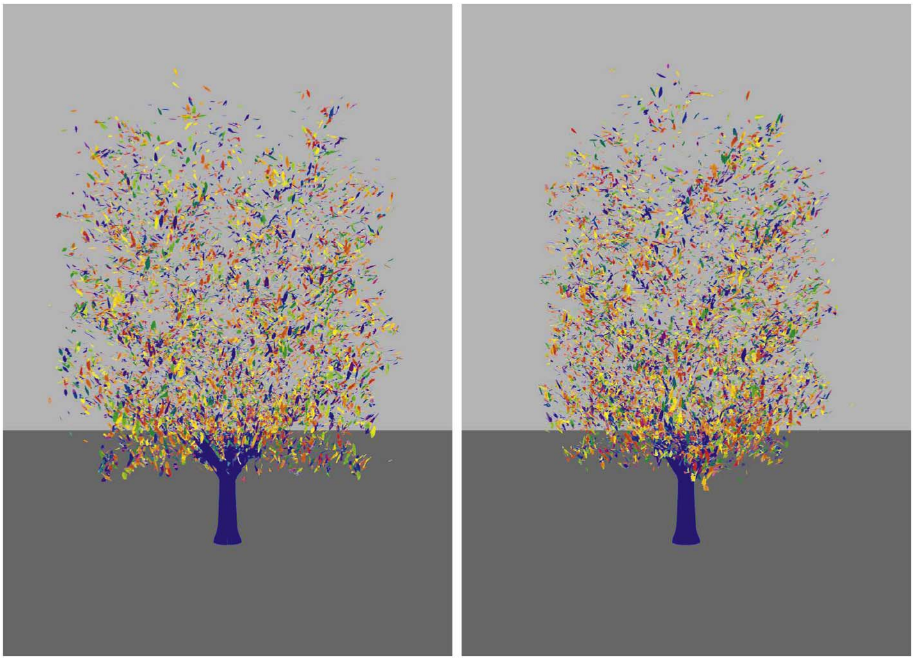

Fig. 5. Visualization of the triangulated leaf groups used to determine the locations of directly reconstructed leaves in the reconstruction shown in Fig. 4b, d for two opposing viewing angles. Each independent fill group is given a unique color. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.2. Leaf area

A more quantitative evaluation of the reconstruction methodology can be conducted by performing a voxel-by-voxel comparison of leaf area between the reconstructed and reference trees (Fig. 6a). Since the exact amount of leaf area in each voxel is known from the reference tree, this provides a means for quantifying the error in measured leaf area. It should be noted that this exercise is primarily a test of the leaf area measurement method of Bailey and Mahaffee (2017a), as this is what determines how much total leaf area should be produced within each voxel.

The index of agreement between the reference and reconstructed total leaf area within the 1000 voxels was 94.2%, and the RMSE was 0.174 m^2 (Fig. 6a), indicating reasonably good overall agreement. There is a notable amount of scatter in the LiDAR measurements, particularly as leaf area density becomes large. There is a small overall negative bias in the estimated leaf area (-0.058m^2), meaning that the LiDAR methodology tended to slightly underestimate the actual amount of total leaf area. Above roughly 1 m^2 of leaves per voxel the scatter becomes increasingly apparent and there is more consistent under prediction. This is likely because the LiDAR inversion methodology used to measure leaf area loses sensitivity as leaf area index along the beam path becomes large (which occurs when either leaf area density or voxel size becomes large). The inversion for leaf area is based on the LiDAR's measurement of the probability that a beam is intercepted by leaves within a given voxel, and as leaf area index along the beam's path becomes large there is little difference in this probability as leaf area varies. There was no clear location in the tree where the relative error in leaf area tended to be largest, but the absolute error was largest wherever leaf area happened to be largest.

Fig. 7 shows the relative amount of leaf area that was directly reconstructed on average. The majority of voxels required backfilling to reach the measured leaf area. Some voxels required that more than 100% of the directly reconstructed leaf area be removed via thinning to match the measured leaf area.

4.3. Characteristic leaf dimension

The ability of the reconstruction method to determine the characteristic leaf size within a given voxel was evaluated in Fig. 6b. The leaf dimension in the reference tree was constant at 8.7 cm. The reconstruction method slightly skews to the left of the actual leaf dimension, which is expected since the leaf is rarely 100% triangulated. However, the majority of the reconstructed leaves are near the actual

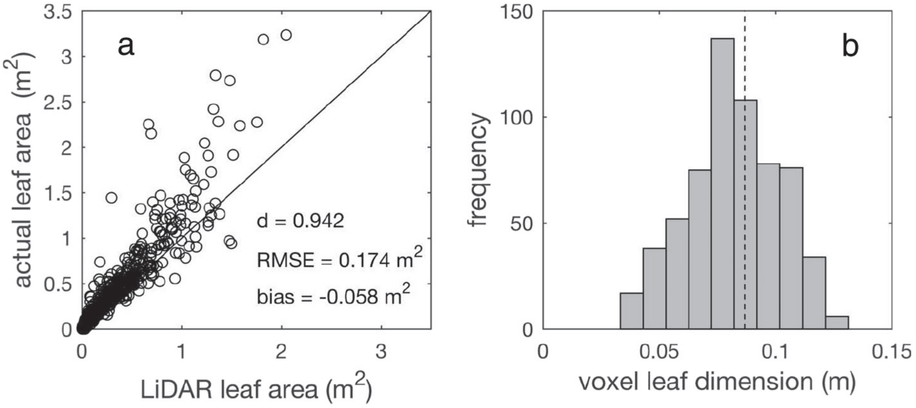

Fig. 6. Comparison of exact values of leaf area (a) and leaf dimension (b) with values obtained from the synthetic LiDAR reconstruction for each voxel. In (a), the diagonal line denotes perfect agreement, and overall agreement is quantified by the index of agreement d, the root-mean-squared error (RMSE), and the mean bias. In (b), the dashed vertical line denotes the (constant) exact value, and bars give a histogram of predicted values over all voxels. Note that the characteristic leaf dimension L was defined as $\sqrt{a}$, where a is the leaf surface area.

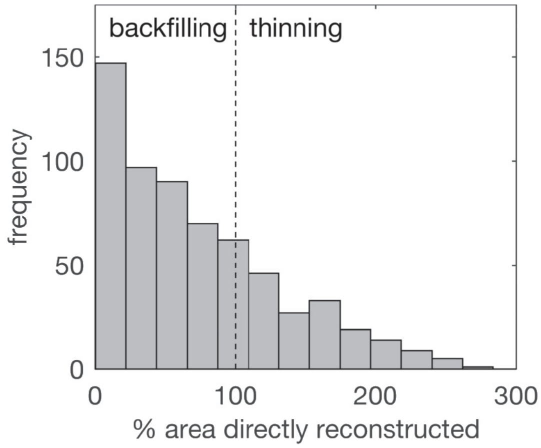

Fig. 7. Histogram of the fraction of leaf area within each voxel that was directly reconstructed. Bars to the left of the vertical dotted line correspond to voxels that had less reconstructed leaf area than actual leaf area, and thus required backfilling. Bars to the right of the vertical dotted line correspond to voxels that had more reconstructed leaf area than actual leaf area, and thus required thinning.

leaf dimension, and the actual mean bias is small at −4.8 mm. The overall RMSE for all reconstructed leaves was 2.0 cm.

4.4. Leaf orientation

To make it feasible to plot voxel leaf angle probability density functions (PDFs), the 10 × 10 × 10 voxel grid was downsampled to a 2 × 2 × 2 grid by simply aggregating neighboring voxels together. Probability density functions are plotted for the leaf inclination (Fig. 8) and azimuthal (Fig. 9) angles within each of these 8 total grid voxels. The exact PDFs from the reference tree are compared against PDFs for the reconstructed tree. PDFs were calculated following the procedure used in Bailey and Mahaffee (2017b), which can be consulted for further details. Overall, the reconstruction is able to qualitatively reproduce the general trends in the inclination and azimuthal angle PDFs. There are some deviations between the reference and reconstructed PDFs due to inadequate sampling of the true PDF, but overall agreement appears visually reasonable. A two sample Kolmogorov-Smirnov test was performed to quantitatively compare the exact and reconstructed leaf angle distributions for each voxel. The distributions for every voxel passed the Kolmogorov-Smirnov test at a 5% confidence interval for both the leaf inclination and azimuthal angle PDFs.

5. Evaluation using field data

5.1. Visualization

Unfortunately, the type of data used above to perform quantitative evaluation of the method is not readily available in the field. Therefore, agreement between the actual (field) and reconstructed trees was assessed based on visual comparisons. In order to do so, the reconstructed trees must be visualized in a manner that is consistent with the way in which the scanner's digital camera perceives the actual tree, which was not an issue in the previous section since identical visualization techniques could be applied for the actual and reconstructed trees. In plotting geometric elements associated with the reconstructed trees, a standard rectangular perspective transformation was applied to the geometry that approximately matched that of the camera lens (Shirley and Morley, 2003). The appropriate field of view for the camera lens was determined through trial-and-error by comparing visualizations of the LiDAR point cloud and photographs. As a result, there is some error in the visualization comparisons due to the camera model used to visualize the reconstructed trees.

Fig. 10 shows a visualization of the tree triangulation, with each fill group given a unique color. Based on visual inspection, the method appears to perform reasonably well in terms of identifying individual leaves. Because of the limited number of distinct colors in the pseudocolor mapping, it can be difficult in some instances to determine

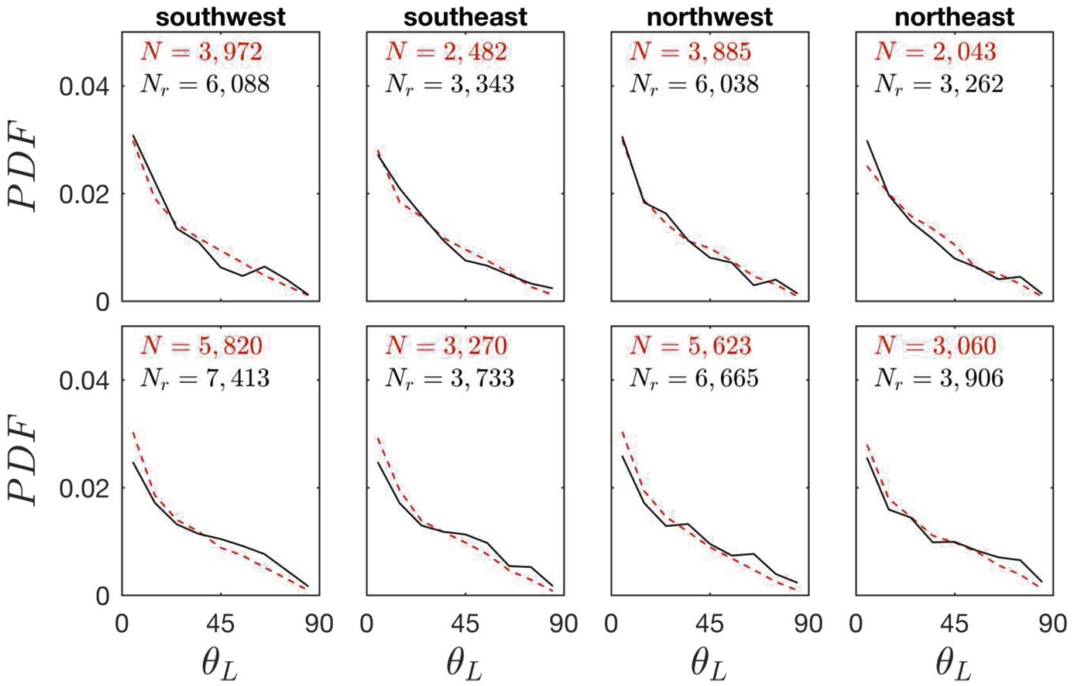

Fig. 8. Probability density functions (PDFs) of leaf inclination angle (θ_L) with a discrete bin size of 10° for eight different leaf zones. The solid black lines correspond to the inclination angle of N total leaves from the tree reconstruction, and the dashed red lines correspond to the inclination angle of N_r total leaves from the reference tree (exact). The leaf zones were determined by downsampling the 10 × 10 × 10 voxel grid to a grid of 2 × 2 × 2 voxels. The top and bottom rows of plots correspond to the top and bottom half of the tree crown, respectively, and each column of plots corresponds to a different azimuthal zone of the tree. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

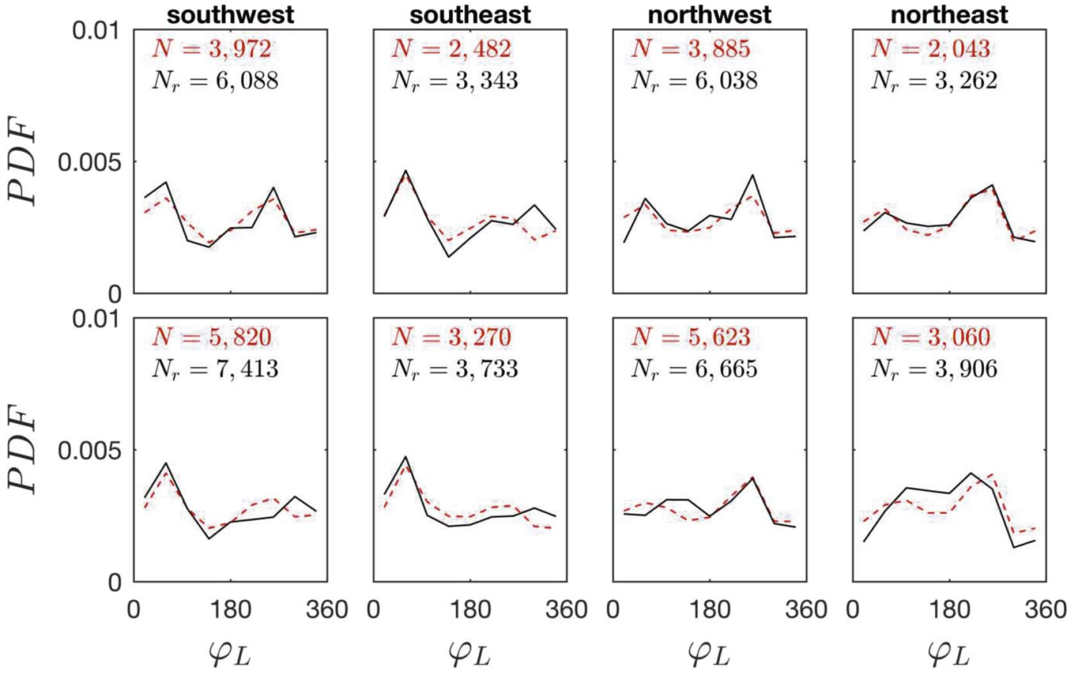

Fig. 9. Probability density functions (PDFs) of leaf azimuthal angle (φ_L) with a discrete bin size of 40° for eight different leaf zones. The solid black lines correspond to the azimuthal angle of N total leaves from the tree reconstruction, and the dashed red lines correspond to the azimuthal angle of N_r total leaves from the reference tree (exact). The leaf zones were determined by downsampling the $10 \times 10 \times 10$ voxel grid to a grid of $2 \times 2 \times 2$ voxels. The top and bottom rows of plots correspond to the top and bottom half of the tree crown, respectively, and each column of plots corresponds to a different azimuthal zone of the tree. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

whether neighboring leaves are in the same fill group or are actually slightly different colors. There appear to be instances in which neighboring leaves that are very close together are inadvertently placed into the same triangle group. However, these occurrences seem to be relatively minimal and still offer reasonable guesses as to where leaves should be placed.

A visualization of the resulting reconstruction as compared with actual photograph and point cloud data is shown in Fig. 11. Qualitative comparison between the actual and reconstructed trees shows close agreement. Individual shoot structures are clearly replicated by the reconstruction. Many individual leaves are closely represented by the reconstructed leaves. Fig. 10 shows which leaves were a result of the direct reconstruction, and indicates that the algorithm is able to identify a large number of individual leaves. The majority of the grid voxels had less than 50% of the leaf area directly reconstructed, and very few required thinning (Fig. 12b). Leaf size prediction seemed to be reasonable (Fig. 12a) and resulted in a visually consistent tree reconstruction.

6. Discussion and conclusions

A semi-direct method was developed and tested that uses terrestrial LiDAR scanning data to reconstruct the architecture of large plants such as trees. A summary of the overall reconstruction algorithm is presented in Fig. 13. The method is termed semi-direct because it seeks to directly reconstruct as many leaves as possible that are in view of the scanner. The resulting direct reconstruction typically represents only a fraction of the total leaf area of the plant. To reconstruct hidden or occluded leaf area, a statistical backfilling procedure was employed in which leaves were added (or removed) such that the overall leaf area and leaf

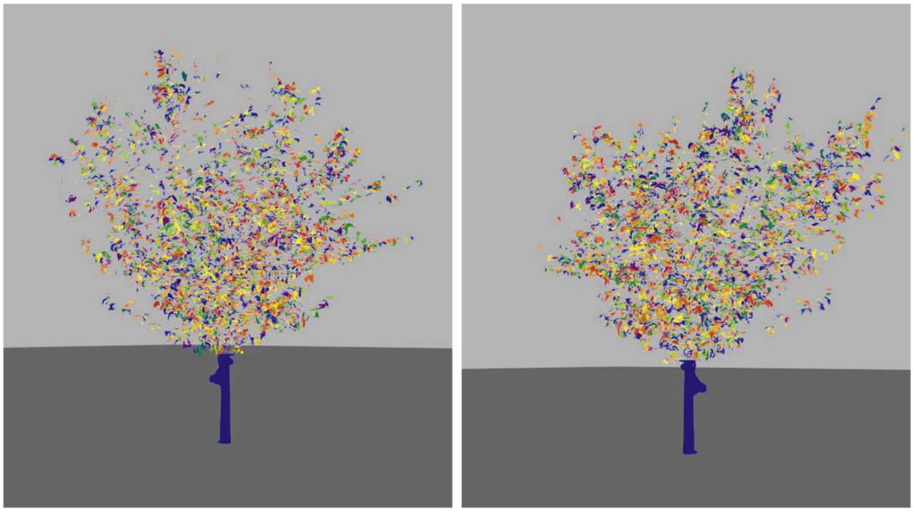

Fig. 10. Visualization of the triangulated leaf groups used to determine the locations of directly reconstructed leaves in the reconstruction shown in Fig. 11b, d (actual elm tree) for two opposing viewpoints. Each independent leaf fill group is given a unique color. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

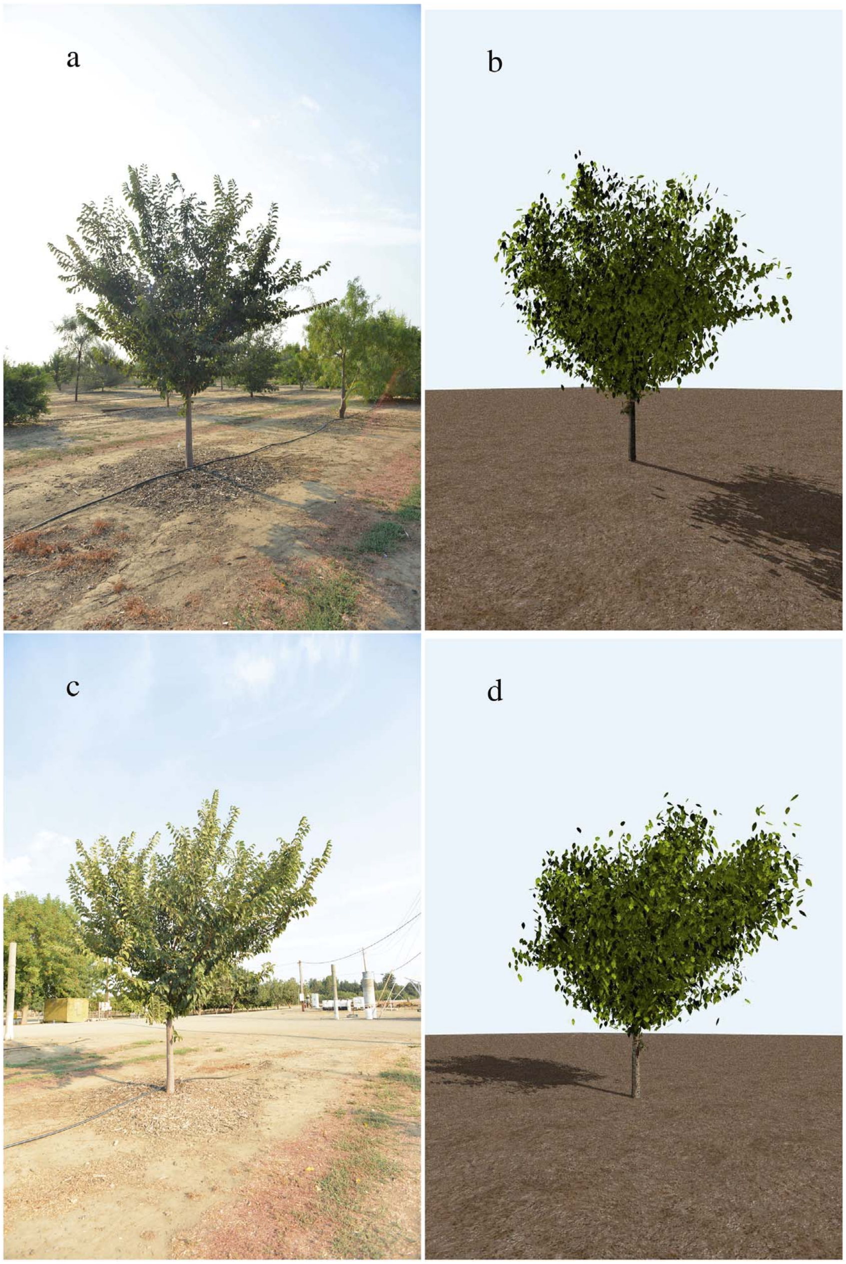

Fig. 11. Visual comparison of actual elm tree photograph (a, c), and reconstructed elm tree (b, d) for two opposing viewpoints.

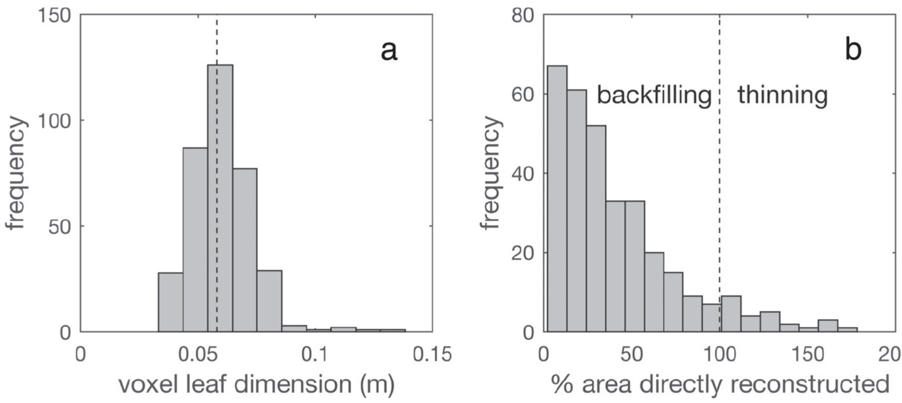

Fig. 12. Histogram of characteristic leaf dimension in each grid voxel for the reconstructed tree in Fig. 11b, d (a), and histogram of the fraction of directly reconstructed leaf area within each grid voxel for the reconstructed tree in Fig. 11b, d (b).

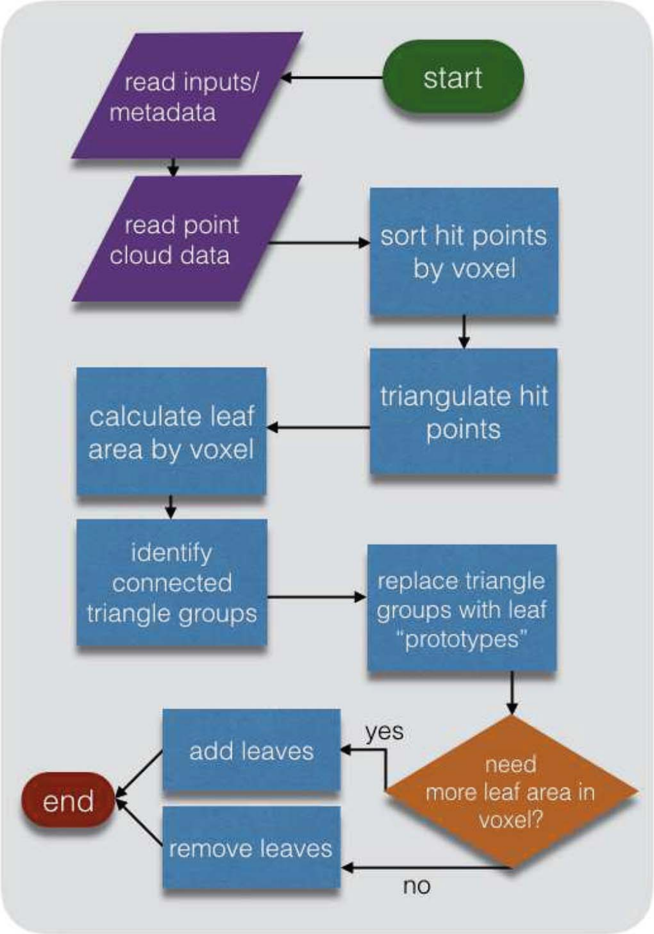

Fig. 13. Flow chart illustrating plant reconstruction methodology.

orientation distributions matched that of the actual plant. This was accomplished by using the methods developed by Bailey and Mahaffee (2017a,b) to measure the leaf angle and leaf area distributions within a user-defined grid of voxels, then adding leaves such that they are consistent with these measured distributions. Thus, the resulting reconstruction is not an exact replica of the plant, rather it is a statistical reconstruction that is consistent with the actual tree at the scale of the voxel grid at that particular instant in time.

In contrast with other methods that rely on the tree branch structure as a skeleton for reconstruction (e.g., Xu et al., 2007; Côté et al., 2009), the present method does not utilize branch structure in the reconstruction of leaves. As a result, the method is applicable to dense plants where little to no wood area is visible from the scanner. The leaf density does, however, affect the quality of the reconstruction. For relatively sparse plants, a larger fraction of leaves are visible to the scanner, and thus the direct portion of the reconstruction represents a larger fraction of the total reconstructed area, which preserves more of the vegetation structure. For dense plants, much of the leaf area is occluded from view of the scanner, and therefore less leaf area is directly reconstructed. Regardless, the reconstructed leaf area and orientation is still consistent with the actual plant at the voxel scale to within the accuracy that the instrument can measure leaf area and orientation for each voxel. A drawback of the present method is that it is dispersive, meaning that it tends to diminish plant structure by spreading out leaves.

Dense vegetation or large voxel sizes have the effect of diminishing the accuracy of the measurement of leaf area. This work suggested that voxels with denser leaves tended to have higher errors in predicted leaf area (Section 4.2). Although not explored in detail, it appeared that for the case examined in this work, errors started to become significant when the voxels contained greater than about 1 m^2 of leaves (note that these values may be case-specific). Future work is needed to more thoroughly examine how various factors affect errors in the leaf area measurement method, as such an exercise was beyond the scope of this work which focused primarily on the reconstruction technique. Small voxels have an additional advantage that they reduce the tendency of the method to disperse or spread out leaves. However, using too small of voxels could become problematic if there are not enough ray samples per voxel.

Aside from the voxel size, there are relatively few tunable parameters in the reconstruction methodology itself. To utilize the triangulation algorithm, the user must specify the maximum allowable triangle dimension. This value is typically easy to specify because results have shown little sensitivity over a wide range, as long as this dimension is much larger than the distance between adjacent hit points and much smaller than the typical distance between adjacent leaves (Bailey and Mahaffee, 2017b). The reconstruction algorithm requires the specification of threshold values for the minimum and maximum allowable surface area of a triangulated leaf "group". Regardless of how these threshold values are specified, the reconstructed tree will still be consistent with the actual tree at the voxel scale in terms of the leaf area and orientation distributions.

The results of this work have important implications in terms of the ability to provide accurate inputs to detailed biophysical models and analyses. Models are now able to represent plant-related processes at the leaf scale (e.g., Vos et al., 2010; Sarlikioti et al., 2011; Bailey, 2018), and combining such models with consistent, leaf-level plant reconstructions provides a means by which these processes can be scaled from leaf-to-tree-to-canopy without the need for often questionable assumptions of homogeneity. In addition to modeling-related efforts, reconstruction data can aid in studies seeking to understand relations between plant structure and function (Meinzer et al., 2011). In order to perform terrestrial scans of entire canopies, scanning throughput needs to be increased. Scanners can be placed on easily movable or autonomous platforms to increase throughput (e.g., Kukko et al., 2012). However, it is important to note that the data processing methods utilized in this work require a stationary sensing platform for the duration of the scan. This also makes utilization of aerial platforms a challenge. At the scan resolution used in this work, scans take only a couple of minutes each (if color photographs are not also collected) and can potentially scan several surrounding trees simultaneously. Canopy-scale reconstruction of very large trees (> 10 m) is likely to introduce additional challenges such as requiring higher scan resolution and high occlusion toward the top of the canopy.

Acknowledgments

Financial support of this work by the American Vineyard Foundation grants 2015-1825/2016-1825/2017-1825, U.S. National Science Foundation grants AGS PREEVENTS 1664175 and DGE-1650042 (Graduate Research Fellowship), and the USDA National Institute of Food and Agriculture Hatch project number CA-D-PLS-2401-H.

References

Amthor, J.S., 1994. Scaling CO_2-photosynthesis relationships from the leaf to the canopy. Photosynth. Res. 39, 321–350.
Bailey, B.N., 2018. Efficient ray-tracing methods for modeling radiation transfer in leaf-resolving plant canopy simulations. Ecol. Model. 398, 233–245.
Bailey, B.N., Mahaffee, W.F., 2017a. Rapid, high-resolution measurement of leaf area and leaf orientation using terrestrial LiDAR scanning data. Meas. Sci. Technol. 28, 064006.
Bailey, B.N., Mahaffee, W.F., 2017b. Rapid measurement of the three-dimensional distribution of leaf orientation and the leaf angle probability density function using terrestrial LiDAR scanning. Remote Sens. Environ. 193, 63–76.
Bailey, B.N., Overby, M., Willemsen, P., Pardyjak, E.R., Mahaffee, W.F., Stoll, R., 2014. A

scalable plant-resolving radiative transfer model based on optimized GPU ray tracing. Agric. For. Meteorol. 198–199, 192–208.

Bailey, B.N., Stoll, R., Pardyjak, E.R., Miller, N.E., 2016. A new three-dimensional energy balance model for complex plant canopy geometries: model development and improved validation strategies. Agric. For. Meteorol. 218–219, 146–160.

Béland, M., Baldocchi, D.D., Widlowski, J.L., Fournier, R.A., Verstraete, M.M., 2014. On seeing the wood from the leaves and the role of voxel size in determining leaf area distribution of forests with terrestrial LiDAR. Agric. For. Meteorol. 184, 82–97.

Binney, J., Sukhatme, G.S., 2009. 3D tree reconstruction from laser range data. In: IEEE International Conference on Robotics and Automation, pp. 1321–1326.

Côté, J.F., Fournier, R.A., Egli, R., 2011. An architectural model of trees to estimate forest structural attributes using terrestrial LiDAR. Environ. Model. Softw. 26, 761–777.

Côté, J.F., Widlowski, J.L., Fournier, R.A., Verstraete, M.M., 2009. The structural and radiative consistency of three-dimensional tree reconstructions from terrestrial LiDAR. Remote Sens. Environ. 113, 1067–1081.

Delagrange, S., Rochon, P., 2011. Reconstruction and analysis of a deciduous sapling using digital photographs or terrestrial-LiDAR technology. Ann. Bot. 108, 991–1000.

DePury, D.G.G., Farquhar, G.D., 1997. Simple scaling of photosynthesis from leaves to canopies without the errors of big-leaf models. Plant Cell Environ. 20, 537–557.

Ehleringer, J.R., 2000. Temperature and energy budgets. In: Pearcy, R.W., Ehleringer, J.R., Mooney, H., Rundel, P.W. (Eds.), Plant Physiological Ecology: Field Methods and Instrumentation. Kluwer Academic Publishers, Dordrecht, The Netherlands, pp. 117–135.

Hackenberg, J., Spiecker, H., Calders, K., Disney, M., Raumonen, P., 2015. SimpleTree - an efficient open source tool to build tree models from TLS clouds. Forests 6, 4245–4294.

Henning, J.G., Radtke, P.J., 2006. Ground-based laser imaging for assessing three-dimensional forest canopy structure. Photogramm. Eng. Remote Sens. 72, 1349–1358.

Kukko, A., Kaartinen, H., Hyppä, J., Chen, Y., 2012. Multiplatform mobile laser scanning: usability and performance. Sensors 12, 11712–11733.

Lee, E., 1987. Region filling using two dimensional grammars. In: Proceedings. 1987 IEEE International Conference on Robotics and Automation, pp. 1475–1478.

Li, F., Chattopadhyay, S., Akbar, S.A., Elfiky, N.M., Kak, A., 2016. A novel visualization tool for evaluating the accuracy of 3D sensing and reconstruction algorithms for automatic dormant pruning applications. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 338–346.

Li, Y., Fan, X., Mitra, N.J., Chamovitz, D., Cohen-Or, D., Chen, B., 2013. Analyzing growing plants from 4D point cloud data. ACM Trans. Graph. 32, 157.

Meinzer, F.C., Lachenbruch, B., Dawson, T.E. (Eds.), 2011. Size- and Age-related Changes in Tree Structure and Function. Springer (514 pp).

Mèndez, V., Catalàn, H., Rosell-Polo, J.R., Arnó, J., Sanz, R., 2013. LiDAR simulation in modelled orchards to optimise the use of terrestrial laser scanners and derived vegetative measures. Biosyst. Eng. 115, 7–19.

Mèndez, V., Rosell-Polo, J.R., Pascual, M., Escolà, A., 2016. Multi-tree woody structure reconstruction from mobile terrestrial laser scanner point clouds based on a dual neighbourhood connectivity graph algorithm. Biosyst. Eng. 148, 34–47.

Moore, G.E., April 1965. Cramming more components onto integrated circuits. Electronics 114–117.

Morsdorf, F., Meier, E., Kötz, B., Itten, K.I., Dobbertin, M., Allgöwer, B., 2004. LiDAR-based geometric reconstruction of boreal type forest stands at single tree level for forest and wildland fire management. Remote Sens. Environ. 92, 353–363.

Pound, M.P., French, A.P., Murchie, E.H., Pridmore, T.P., 2014. Automated recovery of three-dimensional models of plant shoots from multiple color images. Plant Physiol. 166, 1688–1698.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 2007. Numerical Recipes: The Art of Scientific Computing. Cambridge University Press, Cambridge, U.K. 1256 pp.

Raumonen, P., Kaasalainen, M., Åkerblom, M., Kaasalainen, S., Kaartinen, H., Vastaranta, M., Holopainen, M., Disney, M., Lewis, P., 2013. Fast automatic precision tree models from terrestrial laser scanning data. Remote Sens. 5, 491–520.

Rosell, J.R., Llorens, J., Sanz, R., Arnó, J., Ribes-Dasi, M., Masip, J., Escolà, A., Camp, F., Solanelles, F., Gràcia, F., Gil, E., Val, L., Planas, S., Palacín, J., 2009. Obtaining the three-dimensional structure of tree orchards from remote 2D terrestrial LiDAR scanning. Agric. For. Meteorol. 149, 1505–1515.

Sarlikioti, V., de Visser, P.H.B., Marcelis, L.F.M., 2011. Exploring the spatial distribution of light interception and photosynthesis of canopies by means of a functional-structural plant model. Ann. Bot. 107, 875–883.

Shirley, P., Morley, R.K., 2003. Realistic Ray-tracing, Second edition. A. K. Peters, Natick, MA (225 pp.).

Shlyakhter, I., Rozenoer, M., Dorsey, J., Teller, S., 2001. Reconstructing 3D tree models from instrumented photographs. IEEE Comput. Graph. Appl. 53–61.

Sinclair, T.R., Murphy, C.E., Knoerr, K.R., 1976. Development and evaluation of simplified models for simulating canopy photosynthesis and transpiration. Br. Ecol. Soc. 13, 813–829.

Sinoquet, H., Thanisawanyangkura, S., Mabrouk, H., Kasemsap, P., 1998. Characterization of the light environment in canopies using 3D digitising and image processing. Ann. Bot. 82, 203–212.

Suffern, K.G., 2007. Ray Tracing From the Ground Up. A K Peters/CRC Press, Boca Raton, FL (784 pp.).

Vos, J., Evers, J.B., Buck-Sorlin, G.H., Andrieu, B., Chelle, M., de Visser, P.H.B., 2010. Functional-structural plant modelling: a new versatile tool in crop science. J. Exp. Bot. 61, 2101–2115.

Weber, J., Penn, J., 1995. Creation and rendering of realistic trees. In: SIGGRAPH '95 Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques. ACM, pp. 119–128.

Willmott, C.J., 1981. On the validation of models. Phys. Geogr. 2, 184–194.

Willmott, C.J., 1982. Some comments on the evaluation of model performance. Bull. Am. Meteorol. Soc. 63, 1309–1313.

Xu, H., Gossett, N., Chen, B., 2007. Knowledge and heuristic based modeling of laser-scanned trees. ACM Trans. Graph. 26, 19.

Yang, X., Strahler, A.H., Schaaf, C.B., Jupp, D.L.B., Yao, T., Zhao, F., Wang, Z., Culvenor, D.S., Newnham, G.J., Lovell, J.L., Dubayah, R.O., Woodcock, C.E., Ni-Meister, W., 2013. Three-dimensional forest reconstruction and structural parameter retrievals using a terrestrial full-waveform LiDAR instrument (Echidna®). Remote Sens. Environ. 135, 36–51.

Zhang, Z., 1994. Iterative point matching for registration of free-form curves and surfaces. Int. J. Comput. Vis. 13, 119–152.

Ecological Modelling 368 (2018) 233–245

Contents lists available at ScienceDirect

Ecological Modelling

journal homepage: www.elsevier.com/locate/ecolmodel

A reverse ray-tracing method for modelling the net radiative flux in leaf-resolving plant canopy simulations

Brian N. Bailey

Department of Plant Sciences, University of California, Davis, Davis, CA 95616, USA

ARTICLE INFO

Article history:
Received 24 July 2017
Received in revised form
26 November 2017
Accepted 27 November 2017
Available online 12 December 2017

Keywords:
Functional-structural plant model
Graphics processing units
Radiation model
Ray-tracing

ABSTRACT

Radiation is a direct or indirect driver of essentially all biophysical processes in plant systems, and is commonly described through the use of models because of its complex distributions in time and space. Detailed radiation transfer models that represent plant-scale heterogeneity have high computational resource requirements, thus severely limiting the size of problems that can be feasibly considered, while simplified models that can represent entire canopies usually neglect heterogeneity across a wide range of scales. This work develops new methods for computing radiation absorption, transmission, scattering, and emission using ray-tracing approaches that can explicitly represent scales ranging from leaves to canopies. This work focuses on developing a new "reverse" ray-tracing method for describing radiation emission and scattering that ensures all geometric elements (e.g., leaves, branches) are adequately sampled, which guarantees that modelled radiative fluxes are bounded within a reasonable range of values regardless of the number of rays used. This is a critical property when complex model geometries are used, which can be subject to severe sampling errors even when very large ray counts are used. The presented model uses graphics processing units (GPUs) along with highly optimized software to efficiently perform ray-object intersection tests in parallel. This allowed for the simulation of >500 fully resolved trees on a desktop computer in under five minutes.

1. Introduction

A fundamental challenge in studying plant systems is understanding how processes of interest translate across the wide range of relevant scales (Ehleringer and Field, 1993). Plant biophysical processes are often studied locally at the organ level, which are coupled with other plants by environmental processes that traverse the range of scales from leaf to canopy or beyond. Directly measuring physical processes across this wide range of scales is typically not feasible, and generally requires the use of a model at some level. However, representing this range of scales in models is also a considerable challenge, and requires significant simplifications in order to make problems tractable. This means that models usually seek to represent average or representative behaviour and cannot directly resolve plant-scale heterogeneity (e.g., Sinclair et al., 1976; Harley and Baldocchi, 1995; DePury and Farquhar, 1997).

Functional-structural plant models (FSPMs) are a relatively new tool in modelling biophysical processes in plant systems, and seek to describe the three-dimensional development of plant structure over time as influenced by their local environment and physiological function (Vos et al., 2010). These models consist of a coupled set of sub-models that describe various processes involved in plant development such as photosynthesis, nutrient/water transport, carbon allocation, and plant architecture.

FSPM development has progressed rapidly, and holds great potential to aid in our understanding of complex plant system topologies across scales otherwise inaccessible through traditional experimentation. However, despite the continued increase in computing power, FSPMs are often limited in terms of the range of scales they can feasibly represent. Most FSPMs represent the plant at the leaf- and branch-scales, but are typically only able to represent one to a few plants (depending on plant size and model complexity) before computational cost becomes prohibitively expensive (e.g., Allen et al., 2005; Pearcy et al., 2005; Ma et al., 2008; Vos et al., 2010; Sarlikioti et al., 2011). In many cases, this can limit their application in studying plant-to-plant interactions and competition at the field or ecosystem level.

A bottleneck in FSPM computations is the calculation of radiation fluxes absorbed by plant tissues, which directly or indirectly drives nearly all sub-models of physiological processes. Faithfully modelling the transport of radiative energy is complex, particularly when accounting for scattering by millions of elements (e.g., leaves,

---

E-mail address: bnbailey@ucdavis.edu

https://doi.org/10.1016/j.ecolmodel.2017.11.022
0304-3800/

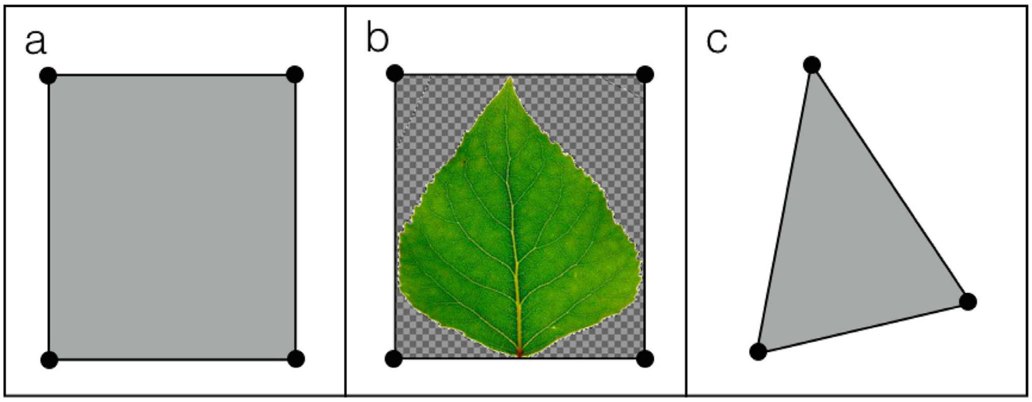

Fig. 1. Geometric element types: (a) patch, (b) alpha mask, and (c) triangle. In (b), the checkered area denotes the portion of the element where material is removed based on the alpha (transparency) channel value.

branches). Most physiological processes have a strong temperature dependence (Johnson and Thornley, 1984). Thus, if temperature is to be included in the model, radiative emission typically must be considered, which adds considerable complexity as each individual element in the domain of interest interacts directly through emission rather than indirectly through scattering. Because of these challenges, models must make compromises in terms of complexity and scale of representation: the more complex the radiation model, the smaller the problem size can be considered.

A very wide range of three-dimensional methods are available to model the transport of sunlight in plant systems (e.g., Ross, 1981; Myneni, 1991; Chelle and Andrieu, 1998; Widlowski et al., 2013), but relatively few are able to model emission of terrestrial radiation. The radiosity method is the standard approach for modelling radiative emission between surfaces or “elements”, and solves a coupled set of equations that represents radiation exchange due to emission and reflection by every element in the domain of interest (Goel et al., 1991; Modest, 2003). Although the radiosity approach is robust, it can quickly become prohibitively expensive as the problem size is increased. This is due to the fact that the radiosity approach involves solving an $N \times N$ system of equations (N being the total number of elements in the domain), which has computational expense that scales as N^3 (cf. Press et al., 2007). Other radiosity-based methods have been developed to improve this scaling by using a multi-scale approach, such as the “nested radiosity” approach of Chelle and Andrieu (1998), which simplifies the contributions due to distant sources of radiation.

When small elements are present in ray-tracing simulations, so-called “reverse” ray-tracing can be used in which radiation is traced backwards by launching rays from elements toward sources, which is common in both computer graphics applications (e.g., Shirley and Morley, 2003) and canopy modelling applications (e.g., Lewis and Muller, 1992; North, 1996; Lewis, 1999; Cieslak et al., 2008). This method has the advantage that every element is guaranteed to be sampled, and thus are more robust when computational cost limits the number of rays that can be afforded. The disadvantage of reverse methods is typically most apparent when scattering of radiation is considered. When a reflection or transmission event occurs along a ray path, the ray traversal becomes irreversible, meaning that tracing of the ray in forward or backward directions is no longer equivalent. This means that more complicated methods must be devised to deal with scattering when the reverse approach is used, which typically leads to increased requirements for memory and run-time.

Currently, methods are not available to effectively model emission of terrestrial radiation using a reverse ray-tracing approach. Using forward methods to model emission in the case of small geometric elements presents similar problems when modelling shortwave radiation, although they may be more severe in the case of emission. Bailey et al. (2016) noted that using a forward tracing method with very small elements resulted in significant sampling errors unless a very large number of rays were used, which could result in substantial errors in the modelled net radiative flux. When coupled with the energy balance equation, it was also noted that very large errors in temperature could result, which led to problematic violations of the second law of thermodynamics.

This paper presents a consistent approach for computing radiative emission in fully-resolved plants using what can be considered a reverse ray-tracing approach adapted to emission. It was hypothesized that using a reverse ray-tracing approach for emission would reduce errors in the modeled net longwave radiation flux. The method is generalized to devise a means for modelling radiation scattering, which is used to produce a complete reverse ray-tracing model of radiation transport due to collimated, point, or terrestrial sources of radiation. The ultimate goal was to develop a model for computing the three-dimensional net radiative flux distribution in fully-resolved canopies that is efficient enough to simulate canopy-scale problems over seasonal time scales in a feasible amount of time.

2. Model description

2.1. Element geometry

It is assumed that the environment of interest is populated by a large number of discrete planar objects, which are termed ‘elements’. In this work, three possible element types will be considered (Fig. 1), which can be combined to form any arbitrary geometry:

- *Patch:* A patch is a planar rectangle defined by four vertices. The patch normal vector $\vec{n}_e$ is defined using a right-hand rule, i.e., following vertices in an anti-clockwise pattern yields an upward pointing normal.
- *Alpha Mask:* An alpha mask (also known as a ‘transparency mask’) is the same as a patch, except that a portion of the patch is removed by specifying a two-dimensional grid of pixels that determines whether or not material is present. The pixel grid is specified using the alpha or transparency channel of a PNG image file.
- *Triangle:* A polygon defined by three arbitrary vertices. The triangle normal vector $\vec{n}_e$ is also defined using a right-hand rule.

2.2. Radiative bands

Arbitrary radiation wavelength bands can be simulated in the model by assigning the appropriate radiative properties to elements. For example, the photosynthetically active radiation (PAR) band (400–700 nm wavelengths) is important when considering sunlight interception by leaves, but typically unimportant when considering radiative emission by leaves since they emit essentially no radiation in this wave band.

Wavelength bands are represented by defining the appropriate radiative properties and emission sources for the band of interest. Each element is assigned total hemispherical radiative properties of ε (emissivity), ρ (reflectivity), τ (transmissivity), and α (absorptivity), where $\rho+\tau+\alpha=1$ (Modest, 2003). Here, the term 'total' indicates a property that is integrated over the spectral band of interest, which is formally defined below (e.g., total reflectivity)

$$\rho = \frac{\int_{\lambda_1}^{\lambda_2} \rho_\lambda I_\lambda \, d\lambda}{\int_{\lambda_1}^{\lambda_2} I_\lambda \, d\lambda}, \tag{1}$$

where ρ_λ is the spectral reflectivity as a function of wavelength λ, I_λ is the spectral intensity of the radiation source, and λ_1 and λ_2 are respectively the lower and upper wavelength limits of the band. Spectral radiative properties can be obtained by direct measurement using a spectroradiometer, or by consulting available spectral databases (e.g., Hosgood et al., 2005; Kotthaus et al., 2014). For simplicity, the spectral intensity and radiative properties are often assumed to be constant over a certain radiative bands such as the PAR band.

2.3. Radiative emission

If the emissivity and absolute temperature of an element is nonzero, the element acts as a source of radiation through emission. The emissive flux of an element can be calculated through the Stefan-Boltzmann law as $\varepsilon\sigma T^4$, where $\sigma = 5.67 \times 10^{-8}$ W m^{-2} K^{-4} is the Stefan-Boltzmann constant, and T is the absolute temperature of the element. Ambient emission from the surroundings may also be present, which is not directly resolved by elements. Given these emitted fluxes, the following procedure describes the method for calculating the absorbed flux due to emission by other elements and the surroundings (illustrated graphically in Fig. 2).

1. Sampling element surfaces: N_{rays} discrete (x, y, z) points are randomly sampled on the surface of each element. In the present implementation, N_{rays} is specified as a constant for all primitives. The sampling of patch surfaces is accomplished by choosing N_{rays} sets of two random numbers from a uniform distribution, and mapping them to the surface of the rectangle (cf. Suffern, 2007). The same procedure is used for alpha masks, except that if points lie in a transparent region they are re-sampled until they are no longer in the transparent region. There are several possible methods for sampling points on an arbitrary triangle, the simplest of which is the rejection method which involves sampling points on a unit rectangle (as in patches) and discarding points that lie outside of the triangle. However, the method used here was to uniformly sample points on a right triangle (i.e., half of a patch), then apply an affine transform to map them onto the arbitrary triangular element.

2. Sampling ray directions: Ray directions are sampled according to a standard cosine-weighted distribution:

$$\theta'_r = \sin^{-1}\sqrt{N_t}, \tag{2}$$

$$\phi'_r = 2\pi N_p, \tag{3}$$

where N_t and N_p are random numbers drawn from a uniform distribution, and θ'_r and ϕ'_r are respectively the ray zenithal and azimuthal directions for a horizontally oriented element. The actual ray directions for an arbitrarily oriented element are found by converting θ'_r and ϕ'_r into a Cartesian unit vector and rotating it into the direction of the element normal. In practice, 'jittered' sampling was used for the calculation of N_t and N_p to improve sampling convergence. This involves generating a uniform grid of N_t and N_p between zero and one, then drawing a random number that gives a displacement that lies between the next adjacent point (see Suffern, 2007, for further details).

A leaf is a thin surface, and could be represented by placing two elements back-to-back. To save computational time and memory,

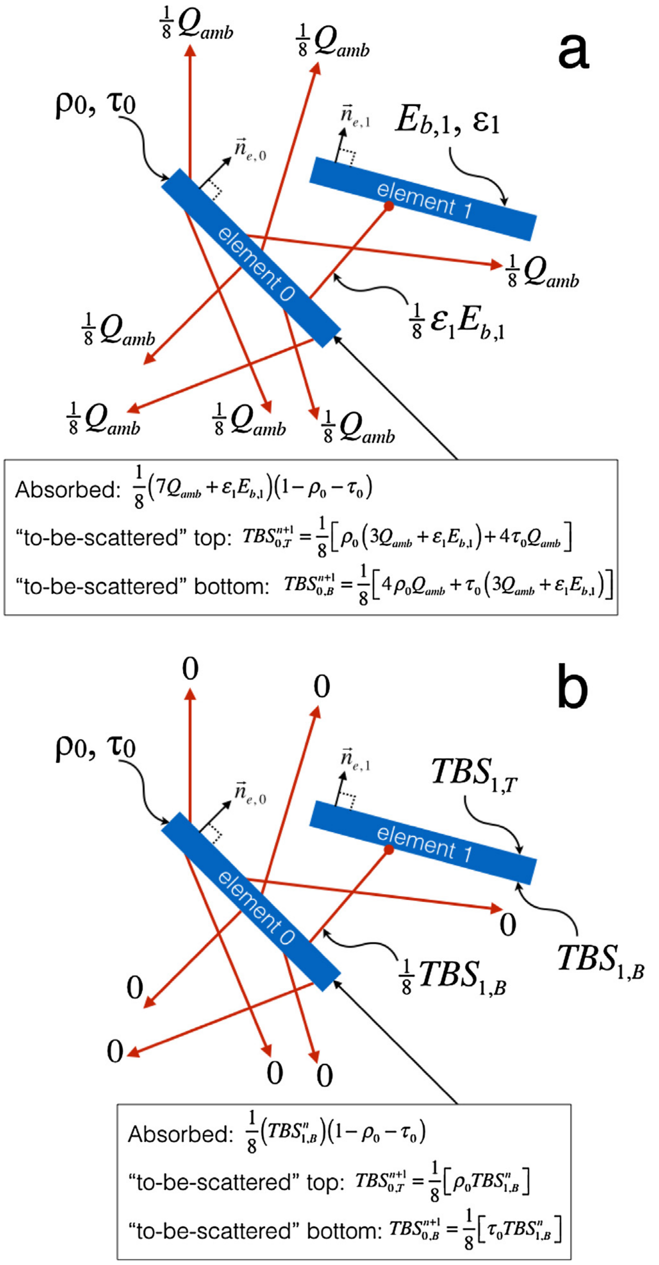

Fig. 2. Sample calculation of absorbed radiation flux by element 0 due to: (a) emission by element 1 and the ambient environment, and (b) the subsequent scattering iteration (denoted as superscript $n = 0, 1, \ldots$). Ray directions in the figure correspond to that of the ray tracing scheme, and not to the direction of radiation propagation. Sample calculations are given for the absorbed flux by element 0, the flux "to-be-scattered" by the top (subscript "T") and bottom (subscript "B") surfaces of the element (all units of W m^{-2}), for the case of 8 rays per element ($N_{rays} = 8$). Note that the top is the side pointing in the direction of the normal vector $\vec{n}_e$. $E_{b,i} = \sigma T_i^4$, ρ_i, τ_i, and ε_i refer respectively to the total hemispherical emissive power, reflectivity, transmissivity, and emissivity of the ith element.

an option can be enabled to emit rays from both sides of an element. For example, an element making up the ground or trunk only needs to emit radiation from one side, while a leaf emits from both sides.

3. Path tracing: The rays are then "launched" from the sampled points away from the element, and ray-object intersection tests are performed (Suffern, 2007). There are one of two possible outcomes for a given ray:

1. *The ray hits another element.* In this case, the emitted flux of the intersected element is queried ($Q = \varepsilon\sigma T^4$) and used to calculate the weighted flux value associated with that ray: Q/N_{rays}. For alpha masks, the alpha map must be queried to determine if the ray intersection occurred on a transparent portion of the element. If so, the ray effectively continues propagating as if no intersection had occurred.
2. *The ray does not intersect any other elements.* In this case, the weighted flux value of the ray is that of the ambient flux Q'_{amb} ($\mathrm{W\,m^{-2}\,sr^{-1}}$), which is the directional-hemispherical flux of an unobstructed surface for the wave band of interest. If the ambient radiation is isotropic, the ambient flux can be assigned a constant hemispherical value Q_{amb} ($\mathrm{W\,m^{-2}}$) and the weighted ray flux is simply Q_{amb}/N_{rays}. If the ambient flux is anisotropic, a directionally-dependent value could be used that varies based on the direction of ray propagation. The ambient flux could be measured using a radiometer equipped with a shadowband (e.g., Horowitz, 1969) or modeled (e.g., Harrison and Coombes, 1988). Since the entire hemisphere around the element is sampled, there is no explicit need for the treatment of boundary conditions. Elements near the edge of the domain may intercept some ambient radiation propagating upward, which replaces the radiation that would have been intercepted from ground elements were the element near the middle of the domain.

The weighted flux of the ray is partitioned into absorbed, reflected, and transmitted components. The absorbed component, $(1-\rho-\tau)Q/N_{rays}$, is accumulated (summed) into an "absorbed" buffer corresponding to the element from which the ray was launched. The reflected component $\rho Q/N_{rays}$, is accumulated in a "to-be-scattered" buffer. There is one "to-be-scattered" for each side of the element, and the reflected flux is accumulated in the buffer corresponding to the same side that was hit (i.e., if the ray hit the top side, the energy is accumulated in the buffer corresponding to the top side). The transmitted component $\tau Q/N_{rays}$ is accumulated in a "to-be-scattered" buffer in a similar manner, except that the transmitted flux is accumulated in the buffer corresponding to the opposite side than was hit (i.e., if the ray hit the top side, the energy is accumulated in the buffer corresponding to the bottom side).

4. Reflection iterations: Energy in the elements' "to-be-scattered" buffers needs to be traced as it leaves the elements. In the present model implementation, it was assumed that reflection and transmission is Lambertian. Measurements indicate that, strictly speaking, this assumption is nearly always false, with specular reflection and scattering of transmitted radiation playing a clear role in creating anisotropy in reflected and transmitted radiation (Woolley, 1971; Bousquet et al., 2005; Combes et al., 2007). While this is usually an incorrect assumption in a strict sense, Chelle (2006) demonstrated that non-Lambertian effects play a minimal role in overall radiation absorption distributions when considering the incident flux due to scattering by many leaves.

Employing the above assumptions, scattering calculations proceed in essentially the same manner as for primary emission. The only differences are that the ray strength is assigned based on the "to-be-scattered" flux rather than the emissive flux, and that the ambient flux is set to zero for scattering iterations. The user specifies the maximum allowable "to-be-scattered" energy (e.g., 0.1 $\mathrm{W\,m^{-2}}$) and a maximum allowable number of scattering iterations (e.g., 10). Scattering iterations proceed until the maximum "to-be-scattered" energy among all elements is below the user-defined threshold, or the number of scattering iterations has exceeded the user-defined maximum value. In either case, any remaining energy in the scattering buffers is absorbed by the element to ensure conservation of energy. In cases where ρ and τ are small, such as in the PAR band, few iterations are required. If ρ or τ are large, such as in the near infrared band, several iterations may be required. However, since reflected and transmitted energy decreases exponentially with each iteration, the number of required scattering iterations is still usually low (e.g., 5–10, depending on user-defined thresholds). The required number of scattering iterations will be further explored in Section 3.1.1.

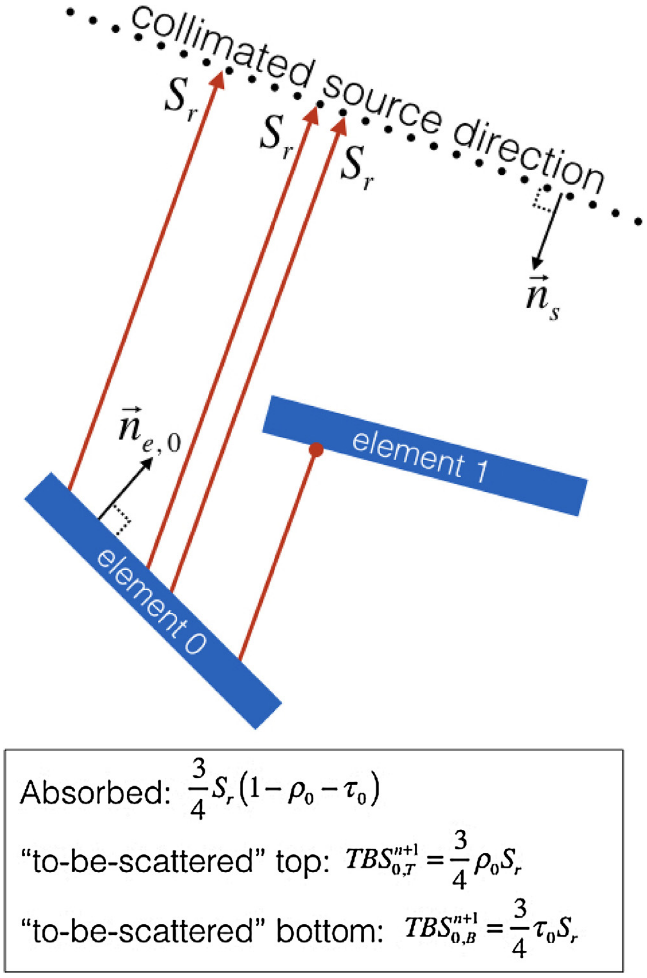

Fig. 3. Ray-tracing schematic for the absorbed flux by element 0 due to a collimated radiation source such as the sun, and amount of energy "to-be-scattered", denoted by *TBS* with subscript "T" or "B" to denote scattering from the top and bottom surfaces, and superscript $n = 0, 1, \ldots$ to denote the scattering iteration. In the example show, the source is sampled using 4 rays per element (i.e., $N_{rays} = 4$). The total hemispherical reflectivity and transmissivity of element 0 is denoted respectively by ρ_0 and τ_0, the hemispherical flux emitted in the direction of the collimated source $\vec{n}_s$ is $S_r N_{rays}$, and $\vec{n}_{e,0}$ is a unit vector normal to the surface of element 0.

2.4. External radiation sources

2.4.1. Collimated radiation sources

Standard approaches are used to calculate the primary incident flux (i.e., before scattering) due to a collimated source of radiation such as the sun (Fig. 3). Collimated radiation is defined as radiation

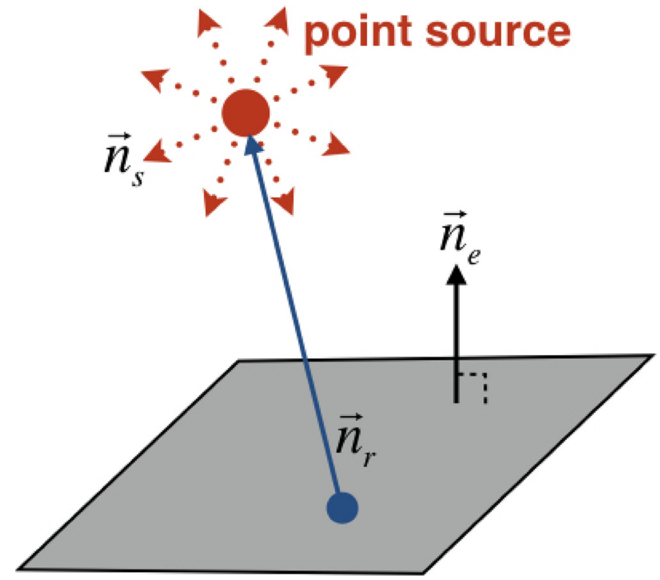

Fig. 4. Ray-tracing schematic for direct radiation calculations with a point radiation source such as a lamp. Points on the element surface are randomly sampled, and $\vec{n}_r$ is a unit vector originating at the sampled point and pointing toward the source. $\vec{n}_e$ is a unit vector normal to the element surface. $\vec{n}_s(\theta_i, \phi_j)$ is a unit vector originating at the location of the point source oriented in the spherical direction (θ_i, ϕ_j).

in which all “beams” of radiation are parallel. Although, emission by the sun is not collimated, the large distance between the sun and Earth means that radiation emitted by the sun that is intercepted by Earth is approximately collimated.

For collimated radiation, points are randomly sampled on the surface of elements in the same way as for emission. The ray direction is simply defined by a line connecting the sampled point and the radiation source. The rays are then traced from the element toward the source, and intersection tests are performed in the same was as for emission. Rays that do not intersect any other elements are assigned a value of

$$S_r = \frac{Q_s|\vec{n}_e \cdot \vec{n}_s|}{N_{rays}}, \tag{4}$$

where Q_s is the source's total hemispherical radiative flux (W m^{-2}) over the band of interest just above the domain in the direction of $\vec{n}_s$. Q_s could come from a sensor measurement (for example, if the measurement of a horizontal sensor was 500 W/m^2 for the direct solar flux, $Q_s = 500/\cos\theta_s$ W/m^2, where θ_s is the solar zenith angle), or from a model of solar transmission (e.g., Liu and Jordan, 1960; Gueymard, 2008).

Absorption and scattering then proceed in the same fashion as with emission. If there is a diffuse flux associated with a given source, the ambient space can be sampled during the first scattering iteration.

2.4.2. Point sources

If the radiation source can be approximated as emenating from a point (e.g., a lamp in the distance), the ray direction is a line that connects the radiation source and a point sampled on the element surface (Fig. 4). The strength of rays that do not intersect any other elements is calculated by integrating incident radiation over all possible directions emitted by the source

$$S_r = \frac{Q_p}{N_{rays}N_\theta N_\phi}\sum_{j=1}^{N_\phi}\sum_{i=1}^{N_\theta}\frac{\left(\vec{n}_e \cdot \vec{n}_r\right)|\vec{n}_r \cdot \vec{n}_s(\theta_i, \phi_j)|}{\pi S^2}\delta, \tag{5}$$

where Q_p is the total emitted radiation flux at the source's surface (W m^{-2}) over the band of interest, $\vec{n}_r$ is a unit vector originating at the point sampled on the element and pointed toward the radiation source, S is the distance from the point sampled on the element to the source. $\vec{n}_s(\theta_i, \phi_j)$ is a unit vector originating at the location of the point source and oriented in the spherical direction (θ_i, ϕ_j), where there are $N_\theta \times N_\phi$ total discrete spherical directions. θ_i is calculated as $\cos^{-1}\left[1-2(i-0.5)/N_\theta\right]$, and ϕ_j is calculated as $2\pi\,(j-0.5)\,/N_\phi$. δ is a function that is equal to 1 m^2 if $\left(\vec{n}_r \cdot \vec{n}_s\right) < 0$ and equal to 0 otherwise, whose purpose is to discard radiation emitted from the back side of the source.

2.5. GPU acceleration

In general, ray tracing calculations can become costly when performed in serial due to the large number of rays required. The computer graphics community has used graphics processing units (GPUs) for decades to accelerate ray-tracing calculations. A GPU is a specialized piece of hardware, with a large number of streaming multiprocessors designed to “crunch numbers” in parallel. Several frameworks have been developed to simplify the task of GPU programming, such as NVIDIA's CUDA framework and OpenCL.

In this work, the above ray-tracing algorithms were implemented using NVIDIA's OptiX ray-tracing library, which is a highly optimized and generalized GPU ray-tracing framework. OptiX does not actually perform any ray-tracing *per se*, rather it is a flexible framework for managing data and tasks associated with ray-tracing. For example, the user writes functions that describes how rays are launched, what constitutes a ray-object intersection, and what should happen in the event of an object hit or miss.

3. Results

3.1. Accuracy and consistency tests

Typically, radiation models are tested by comparing results against real-world data collected in the field (i.e., validation). Performing rigorous validation for detailed radiation models presents several considerable challenges. First, extremely detailed measurements are required to specify geometry within the model. Since individual leaves and branches are represented in the model, it would be necessary to accurately specify the geometry of every leaf and branch in the model if leaf-level measurements are to be reproduced. Another option is to validate using aggregate measurements, which have looser requirements for geometry specification. However, this generally aggregates many different processes and measurement errors, making it difficult to separate natural environmental variability, measurement errors, and model errors.

Another approach that is perhaps of greater value when evaluating these types of detailed models is to verify model predictions for relatively simple cases where the exact answer is known. Since the exact answer is known, all errors in the model predictions must be due to model errors or errors in implementation. This verification over validation approach has been favored by some in the remote sensing community, and has been used for model intercomparisons (Pinty et al., 2001; Widlowski et al., 2013). In the majority of their tests, the ‘exact’ answer was not known, but rather the aggregate of predictions by different models was used as a relative measure of model performance. It should be noted that these tools developed by the remote sensing community focus on the distribution of reflected canopy-level radiation, and not on absorption by individual elements which is the focus of this work.

Here, the model was assessed by ensuring that basic consistency and accuracy conditions were satisfied using problems where the exact answer was known. Accuracy was verified not only by rel-

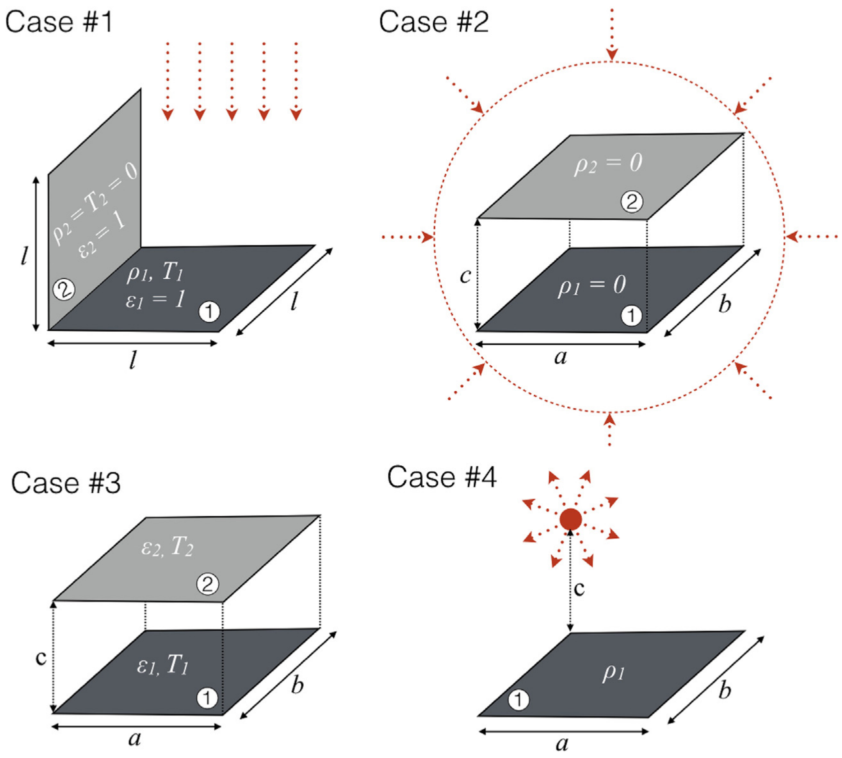

Fig. 5. Schematics of model verification test cases. Case #1 examines both radiative emission and collimated radiation (dotted arrows), Case #2 examines diffuse radiation only (dotted arrows), Case #3 examines radiation emission only, and Case #4 examines a single point radiation source. ρ is the element reflectivity in the case of an external source of radiation, ε is the element emissivity ($1 - \varepsilon$ is the reflectivity due to radiation emitted by elements), and T is the absolute temperature of the elements. Cases and their exact solutions are described in detail in Appendix A.

ative error measures, but importantly by ensuring that the error continually decreases as more rays are used (i.e., the solution converges to the exact answer). Future work will focus on creating new datasets that can be used to more accurately assess real-world errors in model predictions.

3.1.1. Accuracy and convergence

To verify the model and make assessments regarding accuracy, several sample problems were solved that are simple enough that the exact answer can be calculated (Fig. 5). The test cases were designed to progressively evaluate different aspects of the model. Test case parameters such as dimensions, radiative properties, and temperatures were chosen arbitrarily, but with the general goal in mind of providing tests of individual model components and demonstrating the associated error convergence as the number of rays was varied. Each of the test cases are detailed and solved analytically in Appendix A. These details have been included to provide standard cases against which future models can be verified. The performance of the proposed model is contrasted against that of a traditional forward ray-tracing model. In the forward model, rays are launched from the source from which they originated, which may be an element, a collimated source, a point source, or an ambient source (in which case the rays originate from the surface of a sphere that bounds all elements in the domain). A detailed description of the implementation of the forward tracing method is given in Appendix B.

An ensemble of 500 simulations were performed for each case, and the average error was quantified using the normalized root-mean-square error, which can be defined mathematically as

$$nRMSE = \frac{1}{R_{ex}} \left(\sum_i (R_i - R_{ex})^2 \right)^{1/2}, \tag{6}$$

where R_i is the modeled value of the absorbed radiation flux for simulation i in the ensemble, and R_{ex} is the exact value of the absorbed flux.

Case #1 (Fig. 5) was designed to test collimated radiation absorption of a patch, and subsequent reflection to an adjacent patch, as well as patch emission. The reflectivity for radiation originating from the collimated source chosen to be $\rho_1 = 0.3$ and $\rho_2 = 0$. For emission, parameters were chosen to be $\varepsilon_1 = \varepsilon_2 = 1, T_1 = 300$, and $T_2 = 0$. The patch size l is arbitrary and was chosen to be equal to 1.

Fig. 6 reports the normalized error in the absorbed flux for patches #1 and #2, and absorbed flux by patch #2 due to emission by patch #1. Using the reverse method, the error in absorbed radiation for patch #1 is always zero, regardless of the number of direct rays. This is a consequence of the fact that an unobstructed element always perfectly samples radiation sources, a desirable quality of the model. Note that the reverse method is not shown in Fig. 6a because a vale of zero cannot be represented in a logarithmic plot. For the forward method, the absorbed flux for patch #1 is strongly dependent on the chosen number of rays, with the error decreasing exponentially with ray count. This exponential decrease in error with ray count is a standard result for forward ray-tracing simulations (Widlowski et al., 2013). The radiation flux absorbed by patch #2 was strongly dependent on the chosen number of rays for both the forward and reverse methods, and converged to the exact solution at approximately the same exponential rate (Fig. 6b). However, the error in the forward method flux was consistently about four times larger than that of the reverse method. It appears that this is due to the fact that the flux scattered to patch #2 is dependent on proper sampling of the direct radiation by patch #1. Since the reverse method is much more accurate than the forward method at predicting the potential amount of radiation available to be scattered from patch #1, it is consistently more accurate in predicting the absorbed flux by patch #2 due to scattering by patch #1. How-

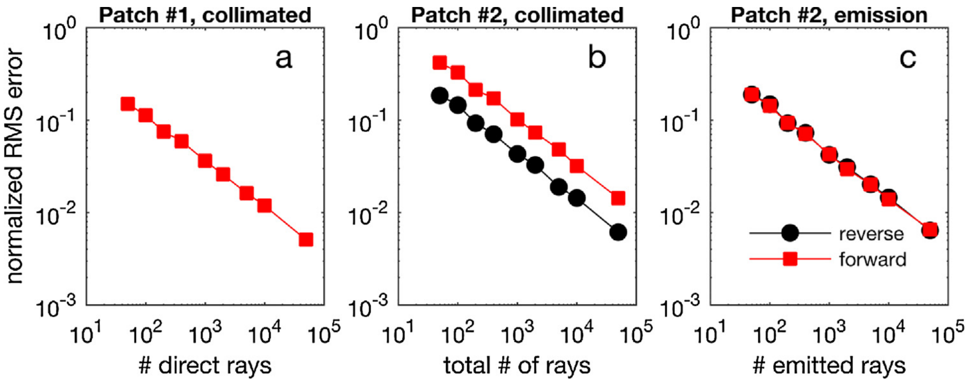

Fig. 6. Convergence of the normalized RMS error (Eq. (6)) in modelled absorbed radiation flux for test Case #1 (see Fig. 5) as the number of rays are varied. Results are reported for the forward and reverse ray-tracing methods. (a) Error in absorbed radiation flux by patch #1 due to a collimated radiation source as the number of collimated rays are varied (note that the reverse method is not shown because the error was always zero); (b) error in the absorbed radiation flux by patch #2 due to scattering of collimated radiation by patch #1 as the total number of rays (collimated + diffuse) are varied; (c) error in the absorbed flux by patch #2 due to radiation emission by patch #1 as the number of emitted rays are varied.

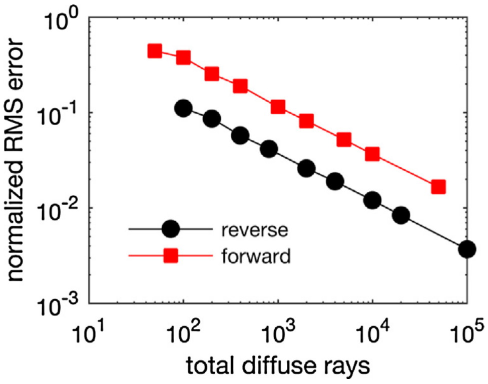

Fig. 7. Convergence of the normalized RMS error (Eq. (6)) in modelled absorbed radiation flux for test Case #2 (see Fig. 5) as the number of rays are varied. Results are reported for the forward and reverse ray-tracing methods. The error is plotted for the absorbed radiation flux by patch #1 due to an isotropic ambient source of radiation as the number of ambient rays are varied.

ever, the diffuse sampling of patch #2 by patch #1 is approximately the same for both methods, as indicated by the similar performance by each method in predicting the absorbed radiation flux by patch #2 due to emission by patch #1 (Fig. 6c). Due to the symmetry of this particular problem, using the forward or reverse methods are essentially equivalent in the simulation of radiative emission by the elements.

Case #2 (Fig. 5) was designed to test sampling of an isotropic ambient radiation source. Radiative properties were chosen to be $\rho_1 = \rho_2 = 0$. The geometry was chosen such that $a = 1$, $b = 2$, and $c = 0.5$. For both the forward and reverse methods, the error in the modelled absorbed flux decreased at approximately the same exponential rate (Fig. 7). However, the error in the forward method was consistently about four times larger than that of the reverse method. For either method, the surface of patch #2 must be adequately sampled by patch #1 in order to calculate the amount of ambient radiation patch #1 intercepts. For the reverse method, this sampling is being performed from the point-of-view of patch #1, whereas for the forward method the sampling is being performed from the point-of-view of the surface of the domain bounding sphere which is on average further away. Thus, the reverse method is more accurate because it better samples the radiation occluded from view of patch #1 by patch #2, but the overall convergence rate is the same.

Case #3 (Fig. 5) was designed to test emission exchange between gray patches and subsequent “infinite” reflections. Here, both the number of diffuse rays and the maximum number of scattering iterations was varied, and only patch emission was considered. Radiative properties were chosen to be $\varepsilon_1 = \varepsilon_2 = 0.6$, and the geometry was specified the same as in Case #2.

When no scattering computations are allowed, both the forward and reverse methods are essentially equivalent because of geometric symmetry as in Case #1. In this case, there is a small amount of initial decrease in the model error as the number of rays is increased, but eventually no improvement is achieved by using more rays given that scattering must be included to correctly predict the absorbed flux (Fig. 8).

As the maximum number of scattering events is increased, the error decreases at an equivalent exponential rate for both methods. Note that this is dependent on the particular method used for scattering in the forward ray-tracing approach. Since, for every ray-object intersection, energy is both absorbed and reflected, the forward and reverse methods are equivalent. Many implementations of forward ray-tracing use a stochastic approach to determine whether absorption or reflection occurs (see Appendix B). In this stochastic implementation, the number of rays available after each scattering event decreases exponentially, thus leading to decreased accuracy and a slower rate of convergence when compared to the reverse ray-tracing approach (not shown).

Case #4 was designed to test emission from a single point radiation source, and subsequent absorption by a patch. The patch reflectivity was $\rho_1 = 0.3$, and the geometry was specified as $a = 1$, $b = 2$, and $c = 1$. N_θ and N_ϕ were set to a constant value of 30 in the case of the reverse tracing method. For both the forward and reverse methods, the error decreased approximately exponentially as the number of direct rays was increased (Fig. 9). However, the error for the reverse method declined at a faster rate than the forward method as the number of rays was increased. On average, the error in the forward method was about fifty times larger than in the reverse method. The reason for this is similar as in Case #1, which is due to the fact that in the reverse method, the source is perfectly sampled by the patch, whereas in the forward method the source must accurately sample a relatively small fraction of the surrounding spherical space. The difference between Case #1 is that the source rays have angular dependence, and thus even in the reverse method, the surface of the patch must be sampled in order to properly calculate the solid angle subtended by the patch with respect to the source. Therefore, the error still has an exponential dependence on ray count even for the reverse method.

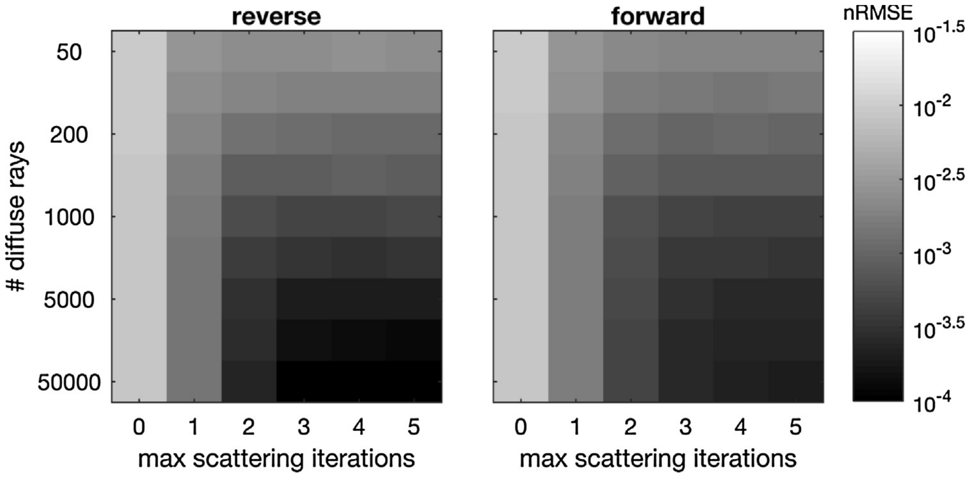

Fig. 8. Convergence of the normalized RMS error (Eq. (6)) for patch #1 in modelled absorbed radiation flux for test Case #3 (see Fig. 5) as the number of rays and number of scattering iterations are varied. Results are reported for the reverse (left pane) and forward (right pane) ray-tracing methods.

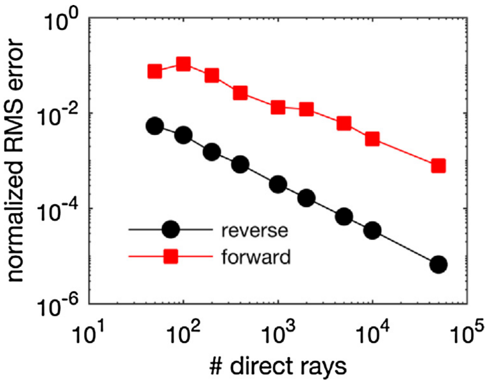

Fig. 9. Convergence of the normalized RMS error (Eq. (6)) in modelled absorbed radiation flux for test Case #4 (see Fig. 5) as the number of rays are varied. Results are reported for the forward and reverse ray-tracing methods. The error is plotted for the absorbed radiation flux by patch #1 due to an isotropic point source of radiation as the number of direct rays are varied.

3.1.2. Conservation of energy (1st law of thermodynamics)

Perhaps the most fundamental consistency requirement for a radiation model is that energy is conserved. Essentially, this means that no energy is gained or lost due to transport; energy can only be relocated. For collimated radiation, the total energy rate incident on the domain is equal to $Q_s \cos\theta_s A_t$, where A_t is the total footprint area of the domain, θ_s is the solar zenith angle, and we note that units are in energy per time. For diffuse radiation, the total incident energy rate is $Q_d A_t$, where Q_d is the diffuse radiation flux on a horizontal surface. For emission, the total energy rate emitted by all elements is $\sum_i^N \varepsilon_i \sigma T_i^4 A_i$. In order to evaluate conservation of energy, we simply calculate the radiation energy rate absorbed by all elements (including energy emitted/reflected to the sky), and compare this value with the total energy rate we started with. If energy is conserved, the two should be equal.

In theory, any consistent radiation model (e.g., ray-tracing, radiosity) should be able to satisfy energy conservation. As mentioned previously, the method described in this work only approximately satisfies conservation of energy. This is because the amount of energy that actually leaves an element is dependent on the element being properly sampled by other elements. To test conservation of energy, a simple domain was created with 140 × 140 patches of size $0.3 \times 0.3\,\mathrm{m}^2$ tiled to create a ground surface, with four trees spaced a 7 m that were 6 m tall on average (Fig. 11). Three-dimensional almond tree geometries were created using the model of Weber and Penn (1995), and placed into the domain at uniform spacing. Each tree consisted of about 6000 triangles making up the trunk and branches, and about 30,000 alpha masks making up the leaves. As detailed in Weber and Penn (1995), each geometric parameter is randomly perturbed to make each tree geometry unique. The shortwave reflectivity and transmissivity of the leaves were arbitrarily set to 0.2 for testing purposes, and the emissivity was set to 0.9. All elements were assigned temperatures of 300 K. For this case, it was found that the error in conservation of energy was about 0.01% of the source flux for a collimated source, and 0.1% of the ambient flux for emission, even if as few as 50 rays per element were used. This indicated that conservation of energy was well satisfied, with errors only slightly larger than machine precision (single precision).

3.1.3. Equilibrium distribution (2nd law of thermodynamics)

A requirement of equal importance as conservation of energy is that the radiation distribution must tend toward a uniform equilibrium state, or that entropy cannot decrease in the absence of external sources or sinks. This is a critically important requirement, since, if the solution does not revert to the correct equilibrium state, we can have little confidence that any observed features of the calculated radiation distribution under non-equilibrium have physical meaning. Testing that this requirement is enforced is relatively simple, and can be performed as a precursor check before every simulation. The temperature of every element is set to some value T, and the ambient flux emanating from the sky is set to σT^4. Regardless of geometry and values of ε, the net radiative flux for all elements should be exactly zero.

The orchard test case with four total trees was used to test model adherence to the 2nd law, and the total number or rays were varied to illustrate its effect on errors in the net radiative flux (Fig. 10). For the forward approach, the total number of rays consisted of both ambient diffuse rays and emission rays (emission and ambient rays are the same in the reverse approach). For each of the four forward simulations, the number of ambient rays was 10^6, 10^7, 10^8, and 10^9, and the corresponding number of emission rays per element was 3000, 5000, 6000, and 10,000 (note that some elements are two-sided and use twice the number of emission rays).

The reverse ray-tracing approach has the very desirable property that the proper equilibrium distribution will always be reached to within machine precision, even with only one ray sample per element. This means that, in the case where the emitted flux from all elements is equal to the diffuse ambient flux, the net modelled

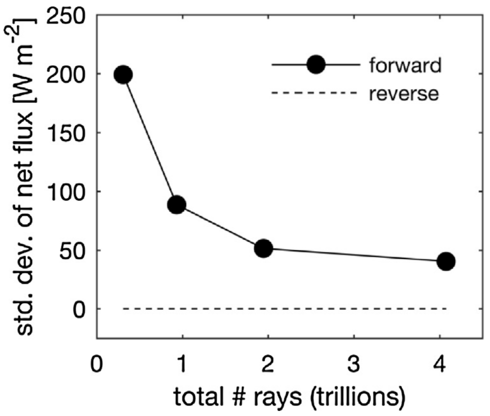

Fig. 10. Standard deviation of the net flux for all elements in the test case shown in this figure as the total number of rays was varied. The reverse tracing method predicts a net flux near zero regardless of the number of rays used.

radiative flux for all elements is essentially zero (about 10^{-7} on average for this case).

When the standard forward ray-tracing method is used, adherence to the 2nd law is strongly dependent on the number of rays in relation to simulated geometry (Fig. 10). For a modest number of rays (i.e., 200 million) the standard error was extremely large at about 150 W m^{-2}, which is about a third of the ambient/emitted flux value σT^4. As the number of rays is increased, the error in the 2nd law steadily decreases. However, even when a very large number of rays was used (i.e., 4 trillion) the error in the net radiation flux was still significant with a standard error of about 40 W m^{-2}. The spatial distribution of this error is visualized in Fig. 11 for this ray count. The range of errors was still very large, with the minimum net flux having a value of −459 W m^{-2} (i.e., no rays sampled the element) and the maximum net flux having a value of about 9000 W m^{-2}. These extreme values corresponded to triangular elements comprising the woody structure of the trees, and occurred at locations of branching/splitting or where there were highly skewed elements. Extreme values could certainly be filtered, or the quality of the triangular mesh could be improved, however this illustrates a problem with using arbitrary meshes with the forward ray-tracing approach that the reverse approach is able to deal with in a straightforward manner.

The reason that the reverse ray-tracing method significantly outperforms the forward method is because it samples the absorbed flux due to both ambient sources and other elements in the exact same way. Therefore, any error in the absorbed ambient flux is exactly offset by the absorbed flux due to other elements. For the forward ray-tracing method, the absorbed ambient flux is sampled from the point-of-view of the domain bounding sphere, and absorbed emission due to other elements is sampled from the point-of-view of the other elements. Each of these samplings can contain substantial and unrelated errors, which can lead to extremely large errors in the net radiative flux.

3.2. Visualizations

Sample visualizations of model outputs are given in Fig. 12, and are simply meant to provide a visual representation of model capabilities. In Fig. 12, each element is coloured based on a mapping of its absorbed radiation flux to the color scale shown in the figure. A flat ground surface was also created by tiling patches of size 0.1×0.1 m^2, which resulted in roughly 1.1 million patches. There were 121 total trees in the orchard, which were 6 m tall and spaced at 7 m. In total, the scene depicted in Fig. 12a contains roughly 5.5 million elements.

Fig. 12a shows a pseudocolor plot for the total absorbed direct and diffuse solar radiation flux in the simulated orchard. The reflectivity and transmissivity of the leaves were set to 0.05, which is characteristic of the photosynthetically active radiation (PAR) band (Jones and Vaughan, 2010), and the rest of the surfaces were assumed to be black. The collimated radiation flux was set to 800 W m^{-2}, and the ratio of the direct to (isotropic) ambient radiation flux (Q_s/Q_{amb}) was 4. The collimated radiation flux zenithal angle was 45°. The number of direct and diffuse rays were set to 100 and 200, respectively.

Fig. 12b depicts direct radiation transport from several point radiation sources corresponding to artificial lamps in a greenhouse. The three-dimensional geometries were read into the code using a standard polygon file that can be generated using nearly any 3D modelling software (e.g., AutoCAD (Autodesk, Inc.), Blender (www.blender.org)). The plants and pots consisted of about 6000 triangles each, each of which were identical except that they were given a random azimuthal rotation. Patches were tiled to make the greenhouse walls, and triangles were tessellated to make lamp housings. In total, the simulation consisted of about 2 million elements.

3.3. Scaling of model run-time

A test case was created to demonstrate model scaling up to very large domain sizes. The tests were performed on an NVIDIA GTX Titan card, which is a consumer-grade gaming GPU. The same general orchard geometry was used as described in Section 3.2. The number of trees was incrementally increased, and the execution time required to simulate both collimated and diffuse radiation for

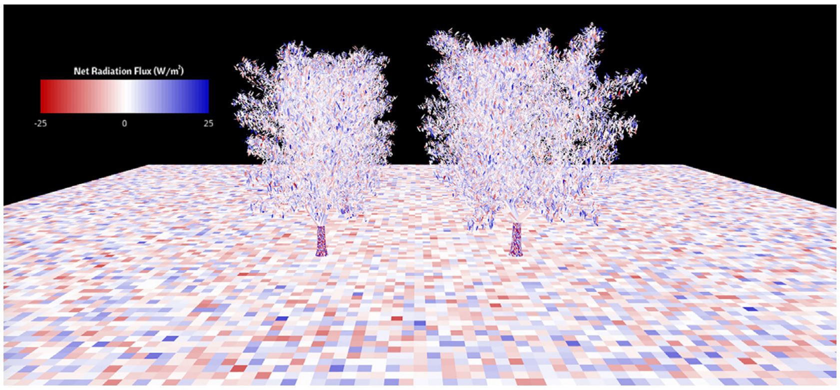

Fig. 11. Visualization of the error in net radiation flux for each element in the 4 tree orchard test case when the forward ray-tracing approach was used. 10^9 ambient and 10^4 emission rays per element were used to yield a total of about 4 trillion total rays. In the test case, the emitted flux from all elements and the ambient flux is σT^4, and thus the model should predict a net flux of zero for all elements.

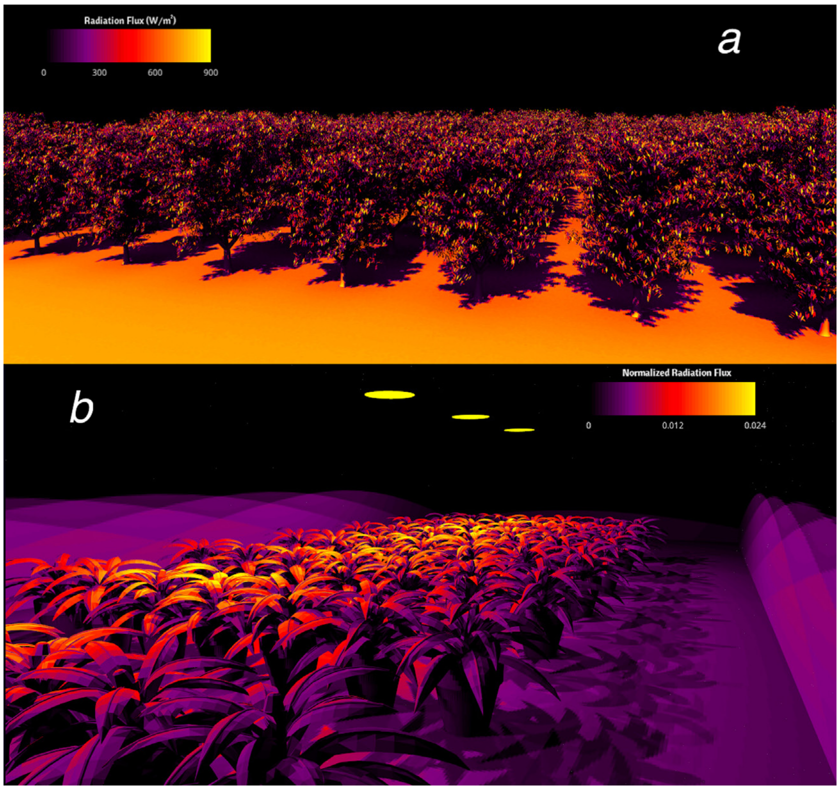

Fig. 12. Sample visualizations of absorbed radiation flux by each element. Top: absorbed direct and diffuse solar radiation in an almond orchard, and bottom: absorbed direct radiation from several point radiation sources corresponding to artificial lighting in a greenhouse normalized by the radiative flux emitted by the sources.

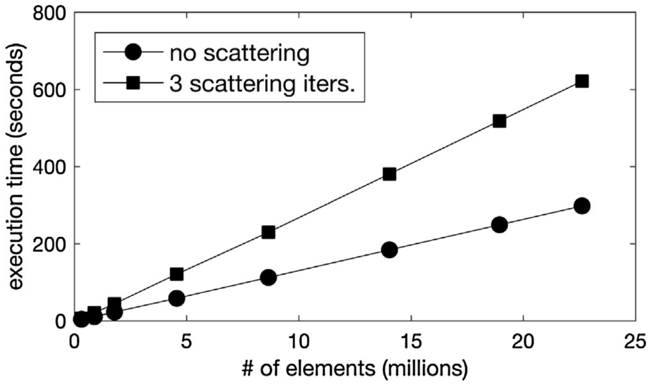

Fig. 13. Scaling of simulation execution time with increasing problem size (almond orchard test case) using an NVIDIA GTX Titan graphics card.

a single band at a single time of day was recorded. Since the purpose of this exercise was simply to demonstrate scaling of run-time, the reflectivity and transmissivity of the leaves were arbitrarily set to 0.1, and the rest of the surfaces were assumed to be black. The number of direct and diffuse rays were chosen to be 100 and 200, respectively. The number of scattering iterations was fixed at 0 and 3 to test its effect on execution time and scaling. By fixing the number of scattering iterations, this removes any effect of the leaf radiative properties on the execution time because scattering will continue in the same manner regardless of the amount of energy to be scattered.

The execution time scaled linearly with number of elements (Fig. 13) for the range of domain sizes examined. Above about 22.5 million elements, the GPU ran out of memory. At this point, larger problems could be simulated by utilizing additional RAM on the host, which comes with an increase in run-time due to the low bandwidth between the GPU and host memories. Once the host runs out of memory, larger problems could still be simulated by utilizing swap space on the host hard drive, which comes with an even larger increase in run-time.

Scattering iterations are expected to incur a similar expense as performing the diffuse radiation calculation without scattering. However, the first scattering iteration is essentially free when performing diffuse calculations, because diffuse calculations and the first scattering iteration can be performed at the same time. Thus, adding three scattering iterations resulted in an increase in execution time of almost exactly 100%. One downside with the method used in this work is that the run-time is virtually independent of the radiative properties of the elements, whereas in traditional forward ray-tracing reflection/transmission rays are not launched if the surface is black since the ray energy becomes zero (see Appendix B). In the reverse method, each element samples the space around it, and thus it cannot assume that any neighboring elements are black until it has actually sampled them.

Results showed that, for this specific geometry and computational hardware, the run-time was about 12.5 seconds per million elements with no scattering. For additional scattering iterations, this run-time can be multiplied by n, where n is the number of scattering iterations. If we were to consider simulating a single tree at hourly timesteps for an entire growing season (assume an average of 12 h of sunlight per day, 200 growing days, and no scattering), we would find the simulation time to be roughly 15 min, which is quite feasible. If the number of trees were to be increased by 100, the execution time would be roughly one day, which is still feasible considering the complexity represented. There are of course opportunities to reduce this execution time by not performing calculations every day, upgrading the GPU, or by using multiple GPUs.

4. Discussion and conclusions

Complex model geometries with small elements typically cause problems for traditional “forward” ray-tracing approaches (e.g., Kimes and Kirchner, 1982; North, 1996; Gastellu-Etchegorry et al., 2004) that trace radiation away from its source. This is because they require an incredibly large number of rays to adequately sample small elements. This forward tracing approach was used to perform simulations of four fully resolved trees to determine whether the model satisfied an equilibrium condition (i.e., if a constant radiative flux is emitted from all objects, the net flux should be zero for all objects). Results showed that there was significant error in the net radiative flux (standard error of about 10%) even when trillions of rays were used. Perhaps more concerning was that there were many outlier elements that had extremely large errors (>100%).

It is fairly well-known that, when elements are small and sampling errors become problematic, “reverse” ray-tracing approaches offer a viable solution (North, 1996; Lewis, 1999; Disney et al., 2000; Cieslak et al., 2008). Here, radiation is traced from the object back toward the source, and therefore all objects are adequately sampled. The trade-off is that, when scattering is considered, the methods become more complicated because radiation propagation is not reversible when scattering occurs. Furthermore, reverse tracing models are not readily available to incorporate radiative emission, which is needed to calculate the net all-wave flux.

In this work, a model was developed that utilizes a reverse tracing approach to calculate the net radiative flux for every element in the simulated domain. The forward and reverse tracing methods were shown to be equivalent in the case of symmetric radiative exchange between two objects. The difference between the methods becomes evident when there is high asymmetry, such as in the case of the sky emitting radiation towards small elements. In the case of forward tracing, the sky is well sampled by an unobstructed element because the sky occupies a large fraction of the spherical space surrounding the element. This same element is relatively poorly sampled by the sky because the element occupies only a small fraction of the total space viewed by the sky. For reverse tracing, the opposite is true. Thus, if we are concerned with absorption by the elements, the reverse method typically offers better sampling of external radiation sources. In the case of emission, either the forward or reverse method may offer better sampling of some elements while giving worse sampling of other elements depending on the relative element sizes. The advantage of the reverse method in all cases is that it guarantees the absorbed flux by any element is bounded within the range of fluxes emitted by all radiation sources, even if only one ray sample per element is used. This is a considerable advantage when only a limited number rays can be afforded because it guarantees that results are at least reasonable, whereas the forward tracing method is likely to produce unreasonable results.

When modelling the net radiative flux at an element, it is advantageous to use a reverse tracing approach both when sampling external sources and emission by elements. It was demonstrated that the reverse tracing method used in this work is at the very least consistent with the second law of thermodynamics or equilibrium condition. The reason that this method satisfies this consistency condition is because any errors in sampling the ambient radiation flux are exactly offset by errors in sampling incoming emission by other elements, which is due to the fact that they are both sampled using the same reverse method.

The reverse ray-tracing method developed in this work bears many similarities with the commonly used radiosity approach (e.g., Goral et al., 1984; Borel et al., 1991; Chelle and Andrieu, 1998), both of which satisfy the consistency condition described above. In theory, the methods are nearly the same, with the major difference being in the treatment of scattering. The reverse tracing approach presented in this work does not require the solution of a linear system of equations that can become costly to solve when the number of elements is large. The present method was shown to scale linearly with the number of geometric elements, while the computational effort associated with solving a linear system of equations increases exponentially with the number of elements (Press et al., 2007).

Results of this work suggest that it is typically preferable to sample the domain from the point-of-view of the entity of interest that is absorbing radiation. If that entity of interest is a leaf or sensor within the canopy, it is typically advantageous to use a reverse ray-tracing approach that ensures the entity is well sampled. In remote sensing applications, the entity of interest is often the distribution of radiation scattered to the sky. Thus, it is more common to find forward tracing methods used in these applications because the domain is sampled from the point-of-view of the sky (Disney et al., 2000; Widlowski et al., 2006, 2013).

In addition to performing detailed studies of radiation transport processes independently, an important future application of the model is coupling with other biophysical models of plant systems. Since the model can represent longwave radiation exchange, it can be coupled with an energy balance model to predict surface temperature throughout canopies. Nearly all physiological processes are strongly dependent on temperature (e.g., photosynthesis, respiration, growth, pest/pathogen development), and such information is needed by physiological models at the organ level. This could be applied not only to understand impacts of heterogeneity in natural ecosystems, but also help to inform or optimize the management of forestry or agricultural systems.

Acknowledgements

Financial support of this work by the American Vineyard Foundation grant 2016-1825, and the USDA National Institute of Food and Agriculture Hatch project number CA-D-PLS-2401-H.

Appendix A. Verification test case descriptions and exact solutions

For completeness, the exact solutions to each of the verification test cases presented in Section 3.1 are given below.

A.1 Case #1

This case consists of two orthogonal square Lambertian patches that share a common edge (Fig. 5). A collimated radiation source is located directly overhead. Patch #1 has a nonzero reflectivity, which reflects some radiation to patch #2 which receives only this reflected radiation. Patch #2 is “cold” and patch #1 is “black”, so the only emission exchange is direct emission from patch #1 to #2.

The fraction of energy leaving patch #1 that is intercepted by patch #2 (i.e., viewfactor F_{1-2} Modest, 2003) is $F_{1-2} = 0.2$ (Modest, 2003, Appendix D39). By symmetry, $F_{2-1} = 0.2$.

Since patch #1 is oriented normal to the radiation source, the absorbed radiation flux is simply

$$Q_1 = (1 - \rho_1) Q_s, \tag{A.1}$$

where Q_s is the collimated radiation flux normal to the radiation source direction. The flux of reflected radiation leaving patch #1 is $\rho_1 Q_s$, and thus the reflected flux absorbed by patch #2 is

$$Q_2 = \rho_1 F_{1-2} Q_s, \tag{A.2}$$

which is only valid for $\rho_2 = 0$. Note that since patch #2 is parallel to the source direction, it receives no direct radiation.

The radiative flux absorbed by patch #2 due to emission by patch #1 is simply

$$E_1 = F_{2-1}\varepsilon_2\sigma T_2^4. \tag{A.3}$$

A.2 Case #2

This case consists of two aligned, parallel, Lambertian patches (Fig. 5). There is an isotropic ambient radiation source. Radiation absorption is ignored on the lower side of patch #1 and on the upper side of patch #2. The reflectivity of both patches is zero.

The fraction of energy leaving patch #1 that is intercepted by patch #2 is given by (Modest, 2003, Appendix D38)

$$F_{1-2} = \frac{2}{\pi XY}\left\{\ln\left[\frac{(1+X^2)(1+Y^2)}{1+X^2+Y^2}\right]^{1/2} + X\sqrt{1+Y^2}\tan^{-1}\frac{X}{\sqrt{1+Y^2}} + Y\sqrt{1+X^2}\tan^{-1}\frac{Y}{\sqrt{1+X^2}} - X\tan^{-1}X - Y\tan^{-1}Y\right\}, \tag{A.4}$$

where $X = a/c$ and $Y = b/c$. By conservation of energy, the fraction of diffuse radiation incident on patch #1 is $(1 - F_{1-2})$. It is also noted that $F_{1-2} = F_{2-1}$ by symmetry. Thus, the absorbed ambient radiation flux for patch #1 is

$$Q_1 = (1 - F_{1-2})Q_s. \tag{A.5}$$

By symmetry, the absorbed ambient radiation flux for patch #2 is the same as that for patch #1.

A.3 Case #3

This case is the same as case #2, except that only emission is considered, and both patches have emissivities less than unity (Fig. 5). As a result, an infinite series of reflections occurs between the two patches.

The viewfactors F_{1-2} and F_{2-1} are given by Eq. (A.4). Since there is no restriction imposed that ε_1 or ε_2 be equal to unity, a radiosity approach can be used to obtain analytical expressions for energy transfer between the two patches. Using conservation of energy, a system of equations describing radiation exchange between the two patches can be written, which yields the standard radiosity equations (Modest, 2003)

$$\frac{e_1}{\varepsilon_1} - \left(\frac{1}{\varepsilon_2} - 1\right)F_{1-2}e_2 = \sigma T_1^4 - F_{1-2}\sigma T_2^4, \tag{A.6a}$$

$$\frac{e_2}{\varepsilon_2} - \left(\frac{1}{\varepsilon_1} - 1\right)F_{2-1}e_1 = \sigma T_2^4 - F_{2-1}\sigma T_1^4, \tag{A.6b}$$

where e_i is the net surface flux for patch #i. By solving for e_2 in Eq. (A.6a), and back-substituting into Eq.(A.6b), an explicit solution can be obtained for e_1 and e_2

$$e_1 = \frac{\sigma\left\{(1-\varepsilon_1)F_{1-2}\left(T_1^4 - F_{2-1}T_1^4\right) + T_1^4 - T_2^4\right\}}{\frac{1}{\varepsilon_1} - \varepsilon_2\left(\frac{1}{\varepsilon_1} - 1\right)\left(\frac{1}{\varepsilon_2} - 1\right)F_{1-2}F_{2-1}} \tag{A.7a}$$

$$e_2 = \varepsilon_2\left\{\left(\frac{1}{\varepsilon_1} - 1\right)F_{2-1}e_1 + \sigma T_2^4 - F_{2-1}\sigma T_1^4\right\}. \tag{A.7b}$$

Finally, subtracting the emitted flux gives the absorbed energy flux

$$E_1 = |e_1 - \varepsilon_1\sigma T_1^4|, \tag{A.8a}$$

$$E_2 = |e_2 - \varepsilon_2\sigma T_2^4|. \tag{A.8b}$$

A.4 Case #4

This case consists of a single point source of radiation incident on a horizontal patch, which has one vertex positioned directly below the source (Fig. 5). The patch may have nonzero reflectivity, which reflects some energy from the source to the surroundings. The fraction of energy leaving the point source that is absorbed by the patch is given by (Modest, 2003, Appendix D47)

$$F_{p-1} = \frac{1}{4\pi}\tan^{-1}\sqrt{\frac{1}{D_1^2 + D_2^2 + D_1^2D_2^2}}, \tag{A.9}$$

where $D_1 = c/a$ and $D_2 = c/b$. Therefore, the flux absorbed by the patch is simply

$$E_1 = (1 - \rho_1)F_{p-1}Q_p, \tag{A.10}$$

where Q_p is the hemispherical flux emitted by the source.

Appendix B. Standard “forward” ray-tracing methodology

The following describes the methodology used to perform forward ray tracing calculations. Most of what is described below can be found in Chapter 11 of Mahan (2002) or other standard radiation textbooks (Modest, 2003; Howell et al., 2010), although some modifications have been made for the treatment of ambient diffuse radiation.

In traditional Monte Carlo ray-tracing approaches, individual ‘bundles’ or ‘rays’ of radiation are launched from their sources and intersected with various elements as they travel through the domain of interest. Ray origins, directions, and associated energy are chosen differently depending on the type of radiation source from which they originate.

Collimated radiation sources: Collimated solar radiation is typically modelled by launching a large number of rays from above the domain and toward the domain in the direction of the source. In order to ensure that rays cover the entire footprint of the domain, a “bounding sphere” can be constructed that encompasses all modelled objects. Rays are then cast from a disk with the same diameter as the bounding sphere, that is positioned tangentially to the bounding sphere (see Overby et al., 2016). The energy assigned to each ray is $Q_s(\pi R_s^2)/N_{rays}$, where Q_s is the radiative flux normal to the source direction, R_s is the radius of the domain bounding sphere, and N_{rays} is the number of direct rays launched toward the domain.

Point radiation sources: Emission from a point radiation source is represented by launching rays uniformly in all directions from the location of the source. The zenithal direction of each ray is chosen as $\theta = \cos^{-1}(1 - R_t)$, where R_t is a number randomly drawn from a uniform distribution. The azimuthal direction of the ray can be similarly specified by $\phi = 2\pi R_p$, where R_p is an additional uniform random number. The energy assigned to each ray is $2Q_p/N_{rays}$, where Q_p is the energy flux emitted by the source, and N_{rays} is the total number of rays launched from each source.

Ambient diffuse radiation sources: Ambient diffuse radiation can be modelled by casting rays in random directions toward the domain. To define a given ray, we must specify both the origin of the ray, and a point through which the ray passes (or equivalent). In the case of an isotropic ambient radiation flux, rays can be defined by randomly choosing two points on the surface of the domain bounding sphere through which the ray will pass. In a spherical coordinate system, the zenithal angle of each point with respect to the bounding sphere center can be given by $\theta = \cos^{-1}(1 - R_t)$, where R_t is a number randomly drawn from a uniform distribution. The azimuthal angle of the points can be similarly specified by $\phi = 2\pi R_p$, where R_p is an additional uniform random number. The energy assigned to each ray is $Q_{amb}(4\pi R_s^2)/N_{rays}$, where Q_{amb} is the isotropic ambient radiation flux, and here, N_{rays} is the number of diffuse rays. If the ambient flux is anisotropic, this could be accounted for by varying the ray energy over the bounding sphere.

Terrestrial emission: Emission by terrestrial objects is modelled by casting rays away from the surface of the objects. Uniformly distributed points are randomly chosen on the surface of the objects (methods for doing so described in Suffern, 2007), which correspond to ray origins. Ray directions are also randomly chosen. For a horizontal surface, the spherical coordinates of the ray can be determined as $\theta = \sin^{-1}\sqrt{R_t}$ and $\phi = 2\pi R_p$, where R_t and R_p are uniform random numbers. If the surface is inclined, the chosen ray direction can be rotated into a coordinate system normal to the object surface. In the case of Lambertian emission, the energy assigned to each ray is $\varepsilon\sigma T^4 A/N_{rays}$, where ε is the object emissivity, σ is the Stefan-Boltzmann constant, T is the absolute temperature of the object's surface, A is the surface area of the object, and here, N_{rays} is the number of rays cast from each object.

Ray-object intersections: Rays are intersected with objects to determine the closest ray-object intersection point. In the event of an intersection, some fraction of energy will be absorbed, transmitted, and reflected. One approach to modelling this is to choose a uniformly distributed random number R. If R is less than the object reflectivity ρ, 100% of the ray's energy is reflected by the object, and a new ray is launched originating at the intersection point with a new random direction away from the same side of the object surface the ray intersected. If R is greater than ρ but less than $\rho + \tau$ (τ is the object's transmissivity), a new ray is launched as in reflection except from the opposite side of the object surface the ray intersected. If $R > \rho + \tau$, 100% of the ray's energy is absorbed by the object and the ray is terminated.

To ensure equivalence with the reverse tracing approach, the forward scattering approach used in this work was a deterministic scattering method. In this case, if a ray intersects an object some fraction of the ray's energy ρ is reflected by launching a new ray whose energy is equal to the product of ρ and the ray energy before scattering. Similarly, some fraction τ is transmitted. The final fraction $1 - \rho - \tau$ is absorbed by the object. This process continues until the ray energy falls below some user-defined threshold value.

References

Allen, M.T., Prusinkiewicz, P., DeJong, T.M., 2005. Using L-sytems for modeling source-sink interaction, architecture and physiology of growing trees: the L-PEACH model. New Phytol. 166, 869–880.

Bailey, B.N., Stoll, R., Pardyjak, E.R., Miller, N.E., 2016. A new three-dimensional energy balance model for complex plant canopy geometries: model development and improved validation strategies. Agric. For. Meteorol. 218-219, 146–160.

Borel, C.C., Siegfreid, A.W., Gerstl, A.W., Powers, B.J., 1991. The radiosity method in optical remote sensing of structured 3-D surfaces. Remote Sens. Environ. 36, 13–44.

Bousquet, L., Lachérade, S., Jacquemoud, S., Moya, I., 2005. Leaf BRDF measurements and model for specular and diffuse components differentiation. Remote Sens. Environ. 98, 201–211.

Chelle, M., 2006. Could plant leaves be treated as Lambertian surfaces in dense crop canopies to estimate light absorption? Ecol. Model. 198, 219–228.

Chelle, M., Andrieu, B., 1998. The nested radiosity model for the distribution of light within plant canopies. Ecol. Model. 111, 75–91.

Cieslak, M., Lemieux, C., Hanan, J., Prusinkiewicz, P., 2008. Quasi-Monte Carlo simulation of the light environment of plants. Funct. Plant Biol. 35, 837–849.

Combes, D., Bousquet, L., Jacquemoud, S., Sinoquet, H., Varlet-Grancher, C., Moya, I., 2007. A new spectrogoniophotometer to measure leaf spectral and directional optical properties. Remote Sens. Environ. 109, 107–117.

DePury, D.G.G., Farquhar, G.D., 1997. Simple scaling of photosynthesis from leaves to canopies without the errors of big-leaf models. Plant Cell Environ. 20, 537–557.

Disney, M.I., Lewis, P., North, P.R.J., 2000. Monte Carlo ray tracing in optical canopy reflectance modelling. Remote Sens. Rev. 18, 163–196.

Ehleringer, J.R., Field, C.B., 1993. Scaling Physiological Processes: Leaf to Globe. Physiological Ecology. Academic Press, Waltham, MA.

Gastellu-Etchegorry, J.P., Martin, E., Gasgon, F., 2004. DART: a 3D model for simulating satellite images and studying surface radiation budget. Int. J. Remote Sens. 25, 73–96.

Goel, N.S., Rozehnal, I., Thompson, R.L., 1991. A computer graphics based model for scattering from objects of arbitrary shapes in the optical region. Remote Sens. Environ. 36, 73–104.

Goral, C.M., Torrance, K.E., Greenberg, D.P., Battaile, B., 1984. Modelling the interaction of light between diffuse surfaces. In: SIGGRAPH '84 Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques, ACM, pp. 213–222.

Gueymard, C.A., 2008. REST2: High-performance solar radiation model for cloudless-sky irradiance, illuminance, and photosynthetically active radiation – validation with a benchmark dataset. Solar Energy 82, 272–285.

Harley, P.C., Baldocchi, D.D., 1995. Scaling carbon dioxide and water vapour exchange from leaf to canopy in a deciduous forest, I. Leaf model parametrization. Plant Cell Environ. 18, 1146–1156.

Harrison, A.W., Coombes, C.A., 1988. Angular distribution of clear sky short wavelength radiance. Solar Energy 40, 57–63.

Horowitz, J.L., 1969. An easily constructed shadow-band for separating direct and diffuse solar radiation. Solar Energy 12, 543–545.

Hosgood, G., Jacquemoud, S., Andreoli, G., Verdebout, J., Pedrini, A., Schmuck, G., 2005. Leaf Optical Properties EXperiment 93 (OPEX93), Ispra, Italy.

Howell, J.R., Siegel, R., Mengüç, M.P., 2010. Thermal Radiation Heat Transfer, 5th ed. CRC Press, Boca Raton, FL, pp. 987.

Johnson, I.R., Thornley, J.H.M., 1984. Temperature dependence of plant and crop processes. Ann. Bot. 55, 1–24.

Jones, H.G., Vaughan, R.A., 2010. Remote Sensing of Vegetation: Principles, Techniques, and Applications. Oxford University Press, Oxford, UK, pp. 400.

Kimes, D.S., Kirchner, J.A., 1982. Radiative transfer model for heterogeneous 3-D scenes. Appl. Opt. 21, 4119–4129.

Kotthaus, S., Smith, T.E.L., Wooster, M.J., Grimmond, C.S.B., 2014. Derivation of an urban materials spectral library through emittance and reflectance spectroscopy. ISPRS J. Photogramm. Remote Sens. 94, 194–212.

Lewis, P., 1999. Three-dimensional plant modelling for remote sensing simulation studies using the Botanical Plant Modelling System. Agronomie 19, 185–210.

Lewis, P., Muller, J.P., 1992. The Advanced Radiometric RAy tracer: ARARAT, for plant canopy reflectance simulation. International ARchives of Photogrammetry and Remote Sensing 29 Commission VII, 26–34.

Liu, B.Y.H., Jordan, R.C., 1960. The interrelationship and characteristic distribution of direct, diffuse, and total solar radiation. Solar Energy 4, 1–19.

Ma, Y., Wen, M., Guo, Y., Li, B., Cournede, P.H., Reffye, P.D., 2008. Parameter optimization and field validation of the functional-structural model GREENLAB for maize at different population densities. Ann. Bot. 101, 1185–1194.

Mahan, J.R., 2002. Radiation Heat Transfer: A Statistical Approach. John Wiley & Sons, New York, pp. 482.

Modest, M.M., 2003. Radiative Heat Transfer. Academic Press, pp. 822.

Myneni, R.B., 1991. Modeling radiative transfer and photosynthesis in three-dimensional vegetation canopies. Agric. For. Meteorol. 55, 323–344.

North, P.R.J., 1996. Three-dimensional forest light interaction model using a Monte Carlo method. IEEE Trans. Geosci. Remote Sens. 34, 946–956.

Overby, M., Willemsen, P., Bailey, B.N., Halverson, S., Pardyjak, E.R., 2016. A rapid and scalable radiation transfer model for complex urban domains. Urban Climate 15, 25–44.

Pearcy, R.W., Muraoka, H., Valladares, F., 2005. Crown architecture in sun and shade environments: assessing function and trade-offs with a three-dimensional simulation model. New Phytol. 166, 791–800.

Pinty, B., Gobron, N., Widlowski, J.L., Gerstl, S.A.W., Verstraete, M.M., Antunes, M., Bacour, C., Gascon, F., Gastellu, J.P., Goel, N., Jacquemoud, S., North, P., Quin, W., Thompson, R., 2001. Radiation transfer model intercomparison (RAMI) exercise. J. Geophys. Res. 106, 11,937–11,956.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 2007. Numerical Recipes: The Art of Scientific Computing. Cambridge University Press, Cambridge, UK, pp. 1256.

Ross, J., 1981. The Radiation Regime and Architecture of Plant Stands. Dr. W. Junk Publishers, The Hague, The Netherlands, pp. 424.

Sarlikioti, V., de Visser, P.H.B., Marcelis, L.F.M., 2011. Exploring the spatial distribution of light interception and photosynthesis of canopies by means of a functional-structural plant model. Ann. Bot. 107, 875–883.

Shirley, P., Morley, R.K., 2003. Realistic Ray-Tracing, 2nd ed. A. K. Peters, Natick, MA, pp. 225.

Sinclair, T.R., Murphy, C.E., Knoerr, K.R., 1976. Development and evaluation of simplified models for simulating canopy photosynthesis and transpiration. Brit. Ecol. Soc. 13, 813–829.

Suffern, K.G., 2007. Ray Tracing from the Ground Up. A K Peters/CRC Press, Boca Raton, FL, pp. 784.

Vos, J., Evers, J.B., Buck-Sorlin, G.H., Andrieu, B., Chelle, M., de Visser, P.H.B., 2010. Functional-structural plant modelling: a new versatile tool in crop science. J. Exp. Bot. 61, 2101–2115.

Weber, J., Penn, J., 1995. Creation and rendering of realistic trees. In: SIGGRAPH '95 Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques, ACM, pp. 119–128.

Widlowski, J.L., Lavergne, T., Pinty, B., Verstraete, M., Gobron, N., 2006. Rayspread: a virtual laboratory for rapid BRF simulations over 3-D plant canopies. In: Graziani, F. (Ed.), Computational Methods in Transport. Springer Berlin Heidelberg, pp. 211–231.

Widlowski, J.L., Pinty, B., Lopatka, M., Atzberger, C., Buzica, D., Chelle, M., Disney, M., Gastellu-Etchegorry, J.P., Gerboles, M., Gobron, N., Grau, E., Huang, H., Kallel, A., Kobayashi, H., Lewis, P.E., Qin, W., Schlerf, M., Stuckens, J., Xie, D., 2013. The fourth radiation transfer model intercomparison (RAMI-IV): proficiency testing of canopy reflectance models with ISO-13528. J. Geophys. Res. 118, 6869–6890.

Woolley, J.T., 1971. Reflectance and transmittance of light by leaves. Plant Physiol. 47, 656–662.