

Document Version

Final published version

Licence

CC BY

Citation (APA)

Katsikogiannis, O. A., Isabella, O., Santbergen, R., & Ziar, H. (2026). Tracing rays from leaves to sky: Multispectral, penumbra-aware irradiance modeling for agrivoltaic orchards. *Applied Energy*, 420, Article 128087. <https://doi.org/10.1016/j.apenergy.2026.128087>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.

Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Tracing rays from leaves to sky: Multispectral, penumbra-aware irradiance modeling for agrivoltaic orchards

Odysseas Alexandros Katsikogiannis^{*}, Olindo Isabella, Rudi Santbergen, Hesam Ziar

Photovoltaic Materials and Devices Group, Electrical Sustainable Energy Department, Delft University of Technology, Delft, the Netherlands

HIGHLIGHTS

- Atmosphere-specific sun-sky generation across the solar spectrum.
- Practical solar-disk sampling for stable penumbra rendering.
- Canopy porosity and phenology impact simulated orchard light levels.

ARTICLE INFO

Keywords:
Agrivoltaics
Orchards
Multispectral
Penumbra
Radiance
Sunflecks

ABSTRACT

Light-simulation tools—exemplified by Radiance—are widely used for quantitative daylight studies and are increasingly adopted in agrivoltaics (agri-PV) to handle complex geometry via ray tracing. Yet, beyond typical workflows three practical limitations persist: spectrally resolved skies are restricted to the visible band; soft-shadow (penumbra) rendering relies on runtime-intensive solar-disk sampling; and fast, integrated canopy models remain scarce. We present a Radiance-compatible Python framework that adds: (i) atmosphere-specific sun-sky generation across the solar spectrum; (ii) efficient, equal-area sampling of the solar disk; and (iii) a simple canopy reconstruction tailored to narrow-trained orchards. To improve spectral fidelity, resolution, and range, we couple SMARTS-derived spectra to a Perez-based sky, leveraging Radiance's multispectral rendering. We deterministically sample the sun's finite extent using a Fibonacci lattice, yielding stable penumbra without prohibitive runtimes. The canopy model parameterizes porosity and seasonal development at a daily rate. Canopy representation matters: opaque-static models, common in agri-PV simulations, systematically underestimate light levels and miss spatiotemporal patterns needed to diagnose suboptimal conditions. Comparatively, a porous-dynamic model led to $\approx 26\%$ higher seasonal light levels, with gains attaining $\approx 100\%$ early in the season and converging to $\approx 16\%$ after foliage matured. While penumbra is limited under conventional PV modules, penumbra-capable renderings enable exploration of design pathways—narrower cell layouts (half-cell and beyond) with greater module-canopy separation—that smooth lighting extremes.

1. Introduction

Accurate irradiance simulations support a wide range of energy and environmental applications—from architecture to agriculture and photovoltaics. Ray-tracing methods have become central across these fields for their ability to handle complex 3-D geometry [1–4]. Agrivoltaic (agri-PV) orchards—where tree canopies are sheltered by overhead PV arrays—are no exception. Striking a balance between canopy light interception and shading requires tailoring PV array design to canopy architecture. While experimental fine-tuning is slow and site-dependent, simulations provide early design insight before orchard establishment.

Recent ray-tracing studies on agri-PV orchards have advanced practice, yet several limitations persist. Canopy representation is typically simplified to opaque and static geometry [5,6], neglecting light porosity and seasonal phenology. Spectral effects are usually ignored and a constant broadband-to-PAR (400–700 nm) conversion is applied [7]—common in practice but insensitive to seasonal atmospheric variation. Finally, the sun is approximated as a point source, neglecting *penumbra* (soft shadows)—a justified simplification for conventional opaque modules, but limiting for semi-transparent modules with narrower cells. The impact of these assumptions on irradiance distribution is highly case-specific, underscoring the need for a unified framework

^{*} Corresponding author.

E-mail address: O.A.Katsikogiannis@tudelft.nl (O.A. Katsikogiannis).

<https://doi.org/10.1016/j.apenergy.2026.128087>

Received 19 November 2025; Received in revised form 31 March 2026; Accepted 19 May 2026

Available online 8 June 2026

0306-2619/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

that captures spatial, temporal, and spectral variability with higher resolution.

We focus on three aspects: (i) spectral sun–sky generation, (ii) penumbra-aware renderings, and (iii) porous, phenology-dependent canopy reconstructions. We present an irradiance modeling framework that quantifies how atmospheric variability, solar-disk angular extent, and canopy architecture jointly shape light availability. This enables site-specific diagnosis and opens new pathways for agri-PV orchard design.

We extend Radiance—a physically based suite of light-simulation tools [8]—by integrating spectral, penumbra, and canopy modeling improvements. We use Radiance's multispectral modifiers to characterize sun and sky spectra from an atmospheric radiative-transfer model. These feed into the ray-tracing process—a different approach from [6], where the global spectrum was merged with spectrally flat ray-traced results, neglecting spectral contrast between sunlit and shaded areas. We introduce a deterministic solar-disk sampler that distributes rays uniformly across the disk, improving efficiency and reproducibility relative to stochastic sampling. Finally, because external canopy reconstruction pipelines are impractical during early design phases, we implement a parametric canopy model that represents growth dynamics with a few interpretable parameters. While a similar approach was used in [7], foliage placement was purely stochastic, lacking the natural clustering of leaves around branches—relevant for wall-like, narrow-trained orchard systems (Fig. 1) [9].

Section 2 reviews these aspects. Section 3 details the simulation framework, Section 4 presents a case study on an agri-PV narrow-trained apple orchard in Bolzano, Italy, and Section 5 discusses the results and concludes.

2. Background

2.1. Spectrally resolved skies

We review three methods for spectrally resolved sky generation in Radiance (excluding sky-hemisphere discretization techniques).

LARK (Rhino/Grasshopper) drives the Perez-based sky generator *gendaylit* for time-series simulation and applies a user-specified, spatially uniform sky spectrum (default: CIE D65) to the sky dome [13]. While traditional Radiance renderings use RGB triples to represent spectra, LARK increases spectral resolution through multi-pass rendering—the spectrum is divided into nine wavelength bands, requiring three passes to cover the entire spectrum.

Radiance v6 introduced multispectral modifiers for light sources or materials, overcoming the boundaries of traditional RGB-based workflows [14]. The accompanying sky generator, *genssky*, produces spectrally anisotropic skies using precomputed (or externally generated) scattering look-up tables (LUTs) [15]. However, *genssky*'s LUTs and wavelength grid currently target only the visible range as there is no runtime parameterization of gas absorption.

ALFA (commercial Rhino plugin) provides physics-based spectrally anisotropic skies driven by libRadtran [16]. Annual simulations have

been pursued via temporal downsampling (using 56 representative point-in-time simulations) and interpolation [17]. ALFA has not yet been calibrated for use with standard ground-based weather inputs.

These tools were developed for architectural daylighting (380–780 nm), whereas silicon PV response extends well beyond this range. We couple SMARTS [18] with a spectral extension of *gendaylit* to generate atmosphere-specific solar spectra—eliminating Rhino dependencies and setting up a Python-based workflow.

2.2. Rendering penumbra

In typical Radiance-based simulations, the sun is treated as a point-source, sampled with a single ray, overlooking the solar disk's angular size. A common workaround combines *direct jittering* (randomly distributing rays over the solar disk) with *supersampling* (rendering at high resolution and downscaling by averaging luminance across pixels) [19]. While this method can produce visually realistic shadows, it introduces a heavy computational burden with limited reproducibility.

We render penumbra by distributing a set of proxy suns across the solar disk to mimic its spatial extent. This “many-suns” concept was recently evaluated [20]: compared against high-dynamic-range (HDR) measurements, it achieved low error with superior computational speed and storage efficiency to alternatives—making it a practical approach for estimating penumbra lengths. However, their implementation compiles an octree to render an image of the sun, distributes suns (one per pixel), and overwrites the *gendaylit* sky file. Our method avoids these and achieves more uniform coverage with fewer proxy suns (details in Section 3.1).

2.3. Canopy geometric models

Geometric canopy representations can be imported via three pipelines.

Point-cloud methods. LiDAR offers high geometric precision but is costly, requires expertise, and suffers with occlusion in dense canopies—limiting internal-layer characterization [21,22]. Commercial systems provide limited spectral discrimination, enforcing manual separation of wood and foliage [23]. Photogrammetry provides a lower-cost alternative for reconstructing the visible canopy envelope with semi-automatic component classification. However, high-quality results demand stable irradiation and minimal wind, and image alignment and point-cloud cleaning remain non-trivial [21,24]. While scalability has improved via UAV/mobile LiDAR [22,23] and orchard-specific pipelines [25–27] occlusion and acquisition quality remain fundamental constraints. In addition, point clouds (and their artifacts) must be meshed prior to Radiance import. Recent methods improve robustness to noise and incompleteness [28,29], yet reconstruction remains bounded by acquisition quality and cannot recover unobserved structure [30].

Radiance-field methods. 3D Gaussian Splatting (3DGS) [31] achieves high visual fidelity from posed multiview images but requires dense, well-posed coverage, stable outdoor conditions, and remains sensitive to dense foliage [32]. Mesh extraction via SuGaR [33] can be




Fig. 1. Narrow-trained orchard systems with multiple vertical leaders (e.g., Guyot [10]): (a) canopy in dormancy reproduced from [10]), (b) early development (reproduced from [11]), (c) fully developed (reproduced from [12]).

fast, but it is not yet the standard for tracking canopy growth [34]. Like point-cloud pipelines, they capture “snapshots” of a plant’s architecture, necessitating repeated field measurements to follow phenology.

Functional-Structural Plant Models (FSPMs) address this limitation by integrating physiology with structural modeling but demand extensive data collection and crop-specific parameter tuning [35–37]. These requirements, coupled with challenging validation procedures, limit scalability and constrain their use to scenarios with ample existing datasets.

3. Methods

The workflow comprises three stages (Fig. 2): (i) daylight-source generation, (ii) surface generation, and (iii) ray tracing. Measured atmospheric constituents—most notably aerosol optical depth at 500 nm (AOD_{500}) and precipitable water vapor (PWV)—are supplied directly or retrieved from ground-based networks (e.g., AERONET). These parameterize the spectral composition of direct and diffuse irradiance, applied to the sun and sky via spectral modifiers. When required, penumbra is modeled efficiently by distributing proxy-suns over the solar disk using an equal-area Fibonacci lattice. Training-system and climate-specific data characterize canopy architecture and phenology, updated daily. Each scene element is defined using an object-oriented design inspired by bifacial_radiance [38].

3.1. Daylight-source generation

Daylight-source generation centers on the sky generator that produces the sun and a continuous sky irradiance distribution using the Perez All-Weather sky model [39]. The model requires inputs of direct normal irradiance (DNI) and diffuse horizontal irradiance (DHI); however, typical radiometric stations measure global horizontal irradiance (GHI), necessitating the use of a decomposition model.

Global irradiance decomposition. We use GISPLITv3 for its robustness at unseen sites [40,41]. GHI decomposition is conditioned by CAELUS, a sky classification algorithm that categorizes sky conditions into six classes [42]. Integrating CAELUS and GISPLIT requires high-quality atmospheric inputs (daily or finer) and an advanced clear-sky radiative transfer model.

Atmospheric radiative-transfer model. We integrate SMARTS into the pipeline to provide clear-sky reference irradiances. When driven by reliable atmospheric data, SMARTS matches spectroradiometer observations within overall instrumental uncertainty ($\approx 5\%$) [43]. Furthermore, it excels at predicting DNI and GHI across various arid locations with high aerosol variability [44]. We have translated a subset of SMARTS v2.9.5 [45] from Fortran to Python. Unlike pySMARTS [46], our implementation does not launch external Fortran executables, streamlining the workflow.

Solar disk sampling technique. Efficient penumbra rendering requires a solar-disk sampler that achieves high uniformity with the least number of points. We evaluated four disk samplers (Fig. 3):

- Pixel grid (clipped): Cartesian pixel centers kept if they fall inside the disk (as in [20]).
- Hexagonal (clipped): Triangular (hex) lattice clipped to the disk.
- Poisson-disk: Grid-accelerated dart-throwing [47] that enforces minimum separation.
- Fibonacci: Equal-area radial placement (golden-angle).

Each distribution was assessed on (i) local regularity via the coefficient of variation of nearest-neighbor distances CV_{NN} (Eq.1), and (ii) radial uniformity via a Kolmogorov–Smirnov distance D_{KS} (Eq. 2).

$$d_i = \min_{j \neq i} \|x_i - x_j\|, \quad \bar{d} = \frac{1}{n} \sum_i d_i, \quad s_d = \sqrt{\frac{1}{n} \sum_i (d_i - \bar{d})^2}, \quad CV_{NN} = \frac{s_d}{\bar{d}} \quad (1)$$

where x_i are 2D coordinates of point i (origin at disk center); d_i is the nearest-neighbor spacing for point i ; n is the number of points (proxy suns).

$$\hat{F}(r) = \frac{1}{n} \sum_i \mathbf{1}(\|x_i\| \leq r), \quad D_{KS} = \sup_{0 \leq r \leq R} \left| \hat{F}(r) - \left(\frac{r}{R}\right)^2 \right| \quad (2)$$

where r a test radius; R is the disk radius; $\mathbf{1}(\bullet)$ the indicator function; $\hat{F}(r)$ the empirical radial CDF (fraction of points within r); and $(r/R)^2$ the ideal CDF for perfectly area-uniform points.

Lower is better for both metrics. Pixel-grid and hexagonal samplers are maximally regular locally by construction ($CV_{NN} = 0$), but clipping causes under-sampling near the disk’s circumference, increasing D_{KS} with only a few points. By contrast, the Fibonacci lattice consistently outperforms all samplers on radial uniformity. Boundary-aware adaptations (e.g., CVT/Lloyd relaxation) were not considered. Because render time scales with the number of proxies, the Fibonacci lattice is adopted to achieve uniformity with fewer proxies. Its construction is as follows.

The Fibonacci lattice distributes points by mapping them from a square lattice onto a unit disk via equal-area transformation. Let the golden ratio be $\phi = (1 + \sqrt{5})/2$. For n proxy suns indexed $i = 0, 1, \dots, n - 1$, the polar coordinates (r_i, θ_i) are computed using:

$$y_i = \frac{i + \varepsilon}{n - 1 + 2\varepsilon} \quad (3)$$

$$r_i = R_s \sqrt{y_i} \quad (4)$$

$$\theta_i = 2\pi \left(\frac{i}{\phi} \right) \quad (5)$$

where R_s is the apparent solar-disk radius (in radians) and ε sets the start offset along the y -axis.

Daylight modeling pipeline (Fig. 4). SMARTS reconstructs solar spectra for clear sky (CS) and for a clean-and-dry atmosphere (CDA). After spectral integration, these yield broadband GHI_{CS} and GHI_{CDA} , which—together with measured GHI—drive CAELUS for sky




Fig. 2. Radiance-based irradiance workflow. Site-specific weather and atmospheric inputs set the sun–sky spectra, while structural and phenological inputs define canopy (and orchard) geometry. The scene is assembled and ray traced to sample multispectral irradiance at specified virtual sensors.

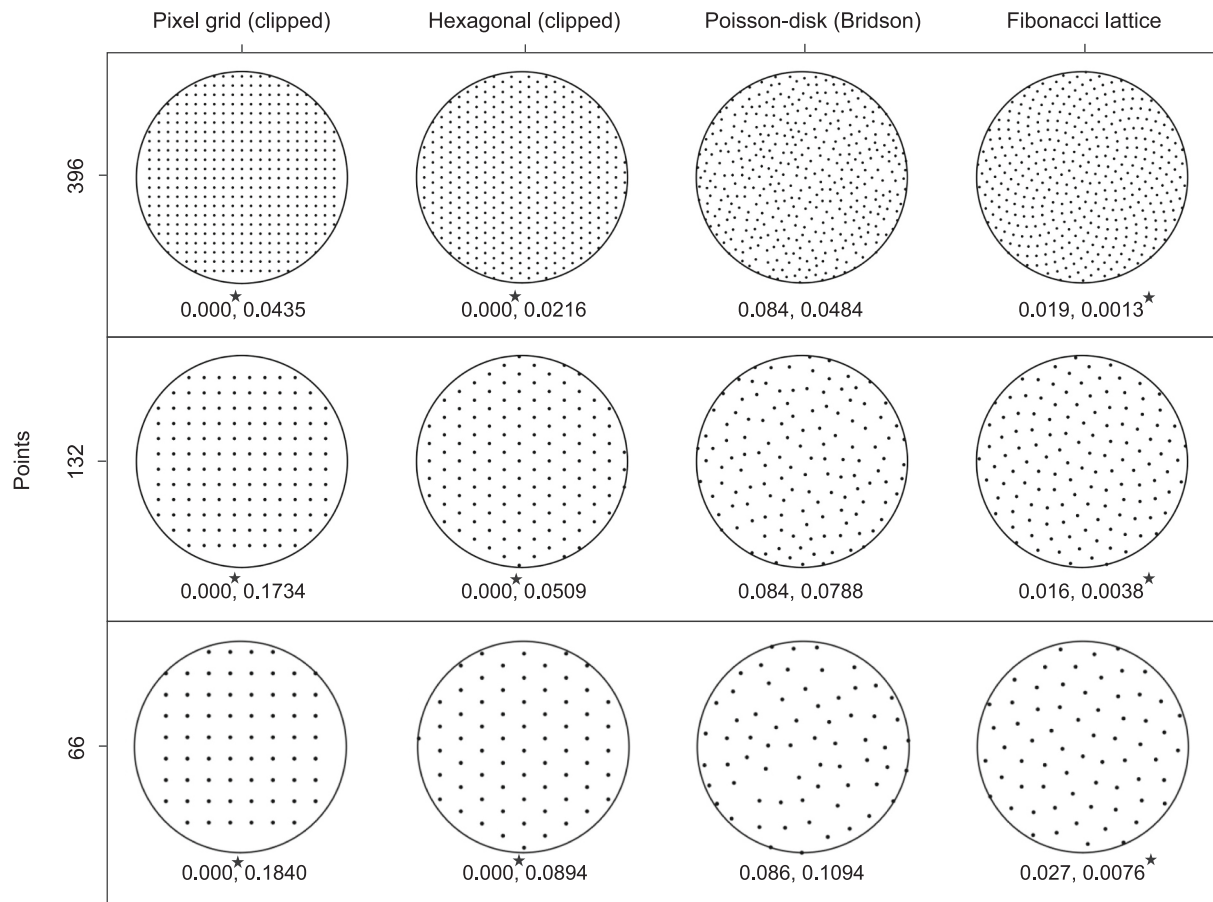


Fig. 3. Comparison of four disk-sampling methods (columns) against total number of points (rows). Numbers under each panel report CV_{NN}, D_{KS} (lower is better). An asterisk marks the best sampler (per metric) at that number of points.




Fig. 4. Daylight modeling pipeline linking SMARTS (clear-sky spectra), CAELUS + GISPLIT (optional sky classification + GHI decomposition), and the sky generator gensdaylit to produce Radiance-compatible sun-sky files. SMARTS-derived spectral shapes are scaled to all-sky DNI/DHI and applied via Radiance's spectral pattern modifiers. Optionally, gensdaylit replaces the point-source sun with a Fibonacci-lattice distribution of proxy suns.

classification. The resulting sky class, measured GHI, and modeled clear-sky inputs (GHI_{CS}, DHI_{CS}) are then passed to GISPLIT for all-sky decomposition.

At each time step, SMARTS-derived direct/diffuse spectral shapes are scaled to match the decomposed all-sky DNI and DHI. This scaling assumes spectrally neutral cloud attenuation and may be biased. After selecting the number of wavebands, both spectra are partitioned and written as per-band fractions to a .dat file.

This wavelength-dependent behavior is applied to broadband DNI and DHI using our Python-based sky generator gensdaylit (the “s” denotes spectral), an extension of gendaylit. Rather than defining the sun and sky with RGB triples, gensdaylit uses Radiance's specfile pattern primitive—a file-based spectral modifier that reads (wavelength, value) samples—to support an increased number of spectral bands. In this configuration, the sky retains the Perez brightness distribution

(anisotropic radiance) while remaining directionally uniform in spectral composition across the hemisphere.

Penumbra is enabled by distributing proxy suns and passing their relative positions to gensdaylit. Each proxy is assigned uniform radiance (no limb darkening), and radiances are scaled so that the sum of proxy contributions equals the target DNI. The polar coordinates from Eqs. (3–5) are converted to Cartesian offsets and stored in a .rad file; at run time these pre-computed angular offsets are applied to the current solar zenith/azimuth via gensdaylit. An implementation and sample I/O for gensdaylit (v1.0) are available in Section Data Availability.

3.2. Surface generation – Orchard reconstruction

The dynamic orchard is built in two stages: (i) reconstruct daily canopy objects (snapshots); (ii) for each day, replicate the corresponding

object across the orchard layout via instancing. Inputs include training-system branch configuration, climate-driven phenological timing, and crop-specific reflectance spectra.

Narrow-trained systems (e.g., Guyot; Fig. 1) exhibit wall-like canopies that permit simplification from full 3D foliage to 2D planar surfaces. Vertical and horizontal foliage planes define the canopy envelope with time-varying porosity (Fig. 5).

Canopy reconstruction. A Boolean grid (initialized to zero) is defined over each foliage plane at a resolution slightly finer than the mean leaf size. Cells corresponding to branches are set to 1 and act as attractors for subsequent foliage placement. Leaf placement is stochastic with distance-based weights,

$$W = (1 + d)^{-CC} \quad (6)$$

where d is the distance from an empty cell to the nearest branch and CC is a clustering coefficient; larger CC yields tighter clusters near branches early in the season; while smaller CC produces more dispersed filling later.

Seasonal development is modeled via cumulative leaf area $A_l(t)$, which often follows a sigmoidal trajectory in fruit-tree canopies [48,49]. We approximate this behavior with a logistic function:

$$A_l(t) = \frac{A_{lfd}}{1 + e^{-k(t-t_{mid})}} \quad (7)$$

where k is the growth-rate parameter, t_{mid} is the inflection time, and A_{lfd} is the leaf area at full development. We set t_{mid} from phenological timing and, given A_{lfd} and one observation $A_l(t)$, solve Eq. (7) for k .

Planar foliage is parameterized by gap fraction GF , a proxy for canopy transmittance [50]. We assume canopy fill $1 - GF$ scales linearly with normalized leaf area $A_l(t)/A_{lfd}$, justified for narrow-trained canopies due to their low internal volume relative to outer surface area. The time-dependent GF that drives progressive filling of the 2D grid is:

$$GF(t) = 1 - \frac{1 - GF_{fd}}{1 + e^{-k(t-t_{mid})}} \quad (8)$$

where GF_{fd} is gap fraction at full development. Foliage is updated daily, while branch structure remains static (consistent with mature canopies maintained by training and pruning).

Trunk and branches define the woody structure. Side and cap foliage (defined as grids using Eqs. 6–8), are assembled into the canopy envelope. Optical properties are assigned per-component and partitioned across the selected wavebands using Radiance's *spectrum* pattern primitive—a basic spectral color modifier. A Radiance *.rad* file is generated (see Fig. 5) and compiled into an *octree*, saved in “frozen”

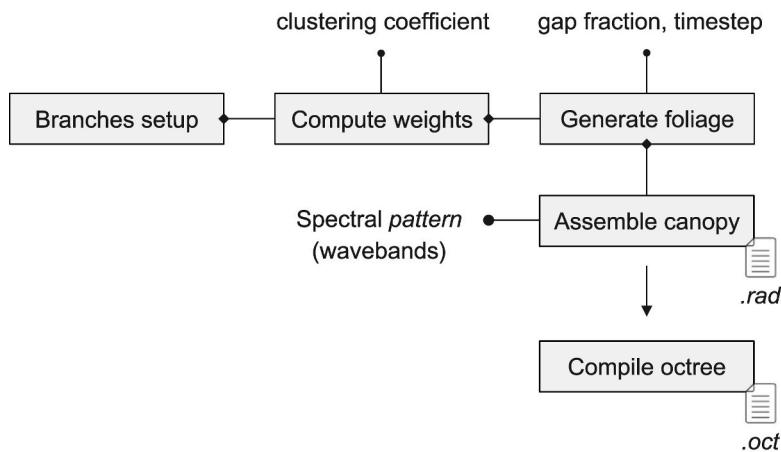


Fig. 5. Canopy reconstruction workflow. Branch attractors are defined on a 2D canopy grid. Distance-based weights, controlled by clustering, guide stochastic foliage placement to achieve a time-dependent gap fraction. Foliage and branches are assembled, assigned spectral optical properties, and compiled into a Radiance octree; the procedure is repeated daily to capture foliage growth.

form (geometry and materials embedded) to enable fast daily instancing.

Canopy replication to orchard. To limit runtime, the orchard is modeled by arraying a single canopy instance according to the layout. Orchard size is chosen so that the central tree experiences minimal light leakage from orchard boundaries at low solar elevations. The overall process of reconstructing and compiling such an orchard incurs $\approx 26\%$ computational overhead relative to an opaque, static-canopy baseline. Nevertheless, this overhead is minor in comparison to *rtrace*.

Steps to calibrate reconstruction: fit growth parameters in Eq. (7) using leaf counts or leaf area index (LAI) during the rapid-growth phase; tune CC (Eq. 6), and GF_{fd} (Eq. 8) against HDR side-view images (April–June) to match observed foliage density and distribution.

3.3. Ray tracing

Daylight sources, geometry, and materials are compiled into the final *octree* and passed to the *rtrace* command, which is well-suited for sampling discrete locations (e.g., individual leaves or PV cells). Irradiance sampling surfaces are placed horizontally (Fig. 6) to align with predominantly planophile leaf-angle tendencies reported for many temperate broadleaf species (e.g., apple, pear) [51]. Simulation accuracy depends on the Radiance parameters (Appendix A), while discretization is governed by three aspects:

Temporal — timestep, controlling generation of dynamic skies and geometry.





Fig. 6. Point-wise irradiance sampling on the canopy cross-section. Virtual leaves are placed adjacent to the canopy envelope. These infinitesimal, horizontal surfaces act as sampling targets for both front- (sky-facing) and back-side (ground-facing) irradiance.



Spectral — number of wavebands, consistent with the spectral pattern modifiers.

Spatial — density and arrangement of sampling points across the scene.

3.4. Simulation validity

The irradiation modeling framework builds on methods with established validation, but the added sub-models introduce uncertainties. This subsection summarizes prior validation efforts, discusses transferability, and identifies components that remain unvalidated.

To isolate the ray tracer from sky-model uncertainty, sky luminance was measured with a scanner and the discrete values were mapped onto the simulated sky hemisphere [52]. Under these conditions, Radiance predicted indoor illuminance with $\approx 5\%$ error. In practice, such datasets are rare; therefore, sky models are used. Beyond the sky input, the largest sensitivities lie in geometric fidelity and surface optics [53,54].

Outdoors, the same principles apply, with added sensitivity to site-specific ground albedo, which varies with solar-zenith angle, season, and weather conditions [55–58]. Spectral-albedo curves from databases often mismatch local conditions [59], limiting their use for validation when ground-reflected irradiance is a major contributor—as in bifacial PV arrays that rely on back-side gains. In [60], on-site spectral albedo was measured and used to compare simulations with pyranometer readings across several tools. We plugged the same inputs into our framework and reproduced back-side irradiance with MBE $\approx 1.5 \text{ W/m}^2$ and RMSE $\approx 5 \text{ W/m}^2$ [61], comparable to [60] and within ISO 9060:2018 Class C pyranometer performance limits [62].

Beyond typical workflows, this framework advances daylight-source generation in two respects. First, spectral skies are driven by SMARTS, which matches high-end spectroradiometer observations when provided with reliable atmospheric data [43]. For this work, quality-assured AERONET (v2.0) AOD and PWV [63] were used. Second, penumbra is rendered via a solar-disk sampling technique that differs from [20]; however, since both approaches sum equal-energy proxy suns, the validation in [20] broadly transfers.

The principal unvalidated component is the canopy's reconstruction, owing to the absence of on-site canopy-level irradiance measurements. Accordingly, results are interpreted comparatively under identical inputs and settings, so shared systematic uncertainties (sky model and surface optics) largely cancel. The canopy model targets orchard-average shading; point-in-time values will differ but mean light levels and spatiotemporal traits are expected to be comparable.

4. Results

Simulation inputs (e.g., sky, geometry, materials) are summarized in Appendix A (Tables A.1–A.3).

4.1. Canopy-model comparison

This section evaluates how canopy representation affects simulated light in an open-field orchard (without PV modules) by comparing the porous–dynamic model with an opaque–static baseline. The opaque–static model replaces the porous canopy envelope with a solid box (no transmission, no phenology). Photosynthetically active radiation (PAR, 400–700 nm, W/m^2) was sampled at virtual leaves on the canopy envelope (Fig. 6) and aggregated over time; we denote cumulative PAR over the specified period as PAR_c (Wh/m^2). Results are reported for both canopy sides (southeast- and northwest-facing, hereafter “east-facing” and “west-facing”).

Seasonal PAR_c vertical profiles (April–September; Fig. 7a) show height gradients: PAR_c decreased from top to bottom due to canopy self-shading and inter-row occlusion. When averaged across the canopy envelope the porous–dynamic model allowed $\approx 26\%$ higher PAR_c than the opaque–static baseline. Regardless of canopy model, back-side contributions (Fig. 7b) enhanced total PAR_c by $\approx 14\%$, with larger gains lower in the canopy where ground view is less obstructed.

Light distributions were evaluated using kernel density estimate (KDE) based violin plots which visualize the probability density of simulated PAR samples; wider sections indicate higher sample density (more frequent values). Seasonal PAR_c distributions (Fig. 8a) show clear peaks corresponding to zones where similar light levels prevail. Although shapes are broadly similar, the porous–dynamic distribution is more compact due to sunflecks—brief, patchy bursts of direct sunlight that pass through small canopy gaps—enhancing light homogeneity. A clear horizontal asymmetry was also observed between east- and west-facing sides, primarily because row orientation deviates from true east–west alignment, increasing irradiation on the east-facing side.

Monthly analysis (Fig. 8b) revealed phenological effects. Early in the season, differences between canopy models were most pronounced, while as the season progressed PAR_c distributions became increasingly similar in shape and magnitude—a convergence largely driven by foliage development. During early canopy development in April, the porous–dynamic model shows relatively uniform PAR_c distributions. By early June, as foliage matures, pronounced vertical and horizontal inhomogeneities emerge, further modulated by sky and solar dynamics.

For both canopy models, inhomogeneity peaked around June, primarily due to the sun's wider, asymmetric azimuthal path relative to the




Fig. 7. Light-penetration profiles for two canopy models: opaque–static, and porous–dynamic. Each point represents seasonal PAR_c at each sampling surface (virtual leaf). (a) Total PAR_c (front + back); (b) back-side PAR_c only. Both plots include east- and west-facing canopy sides.

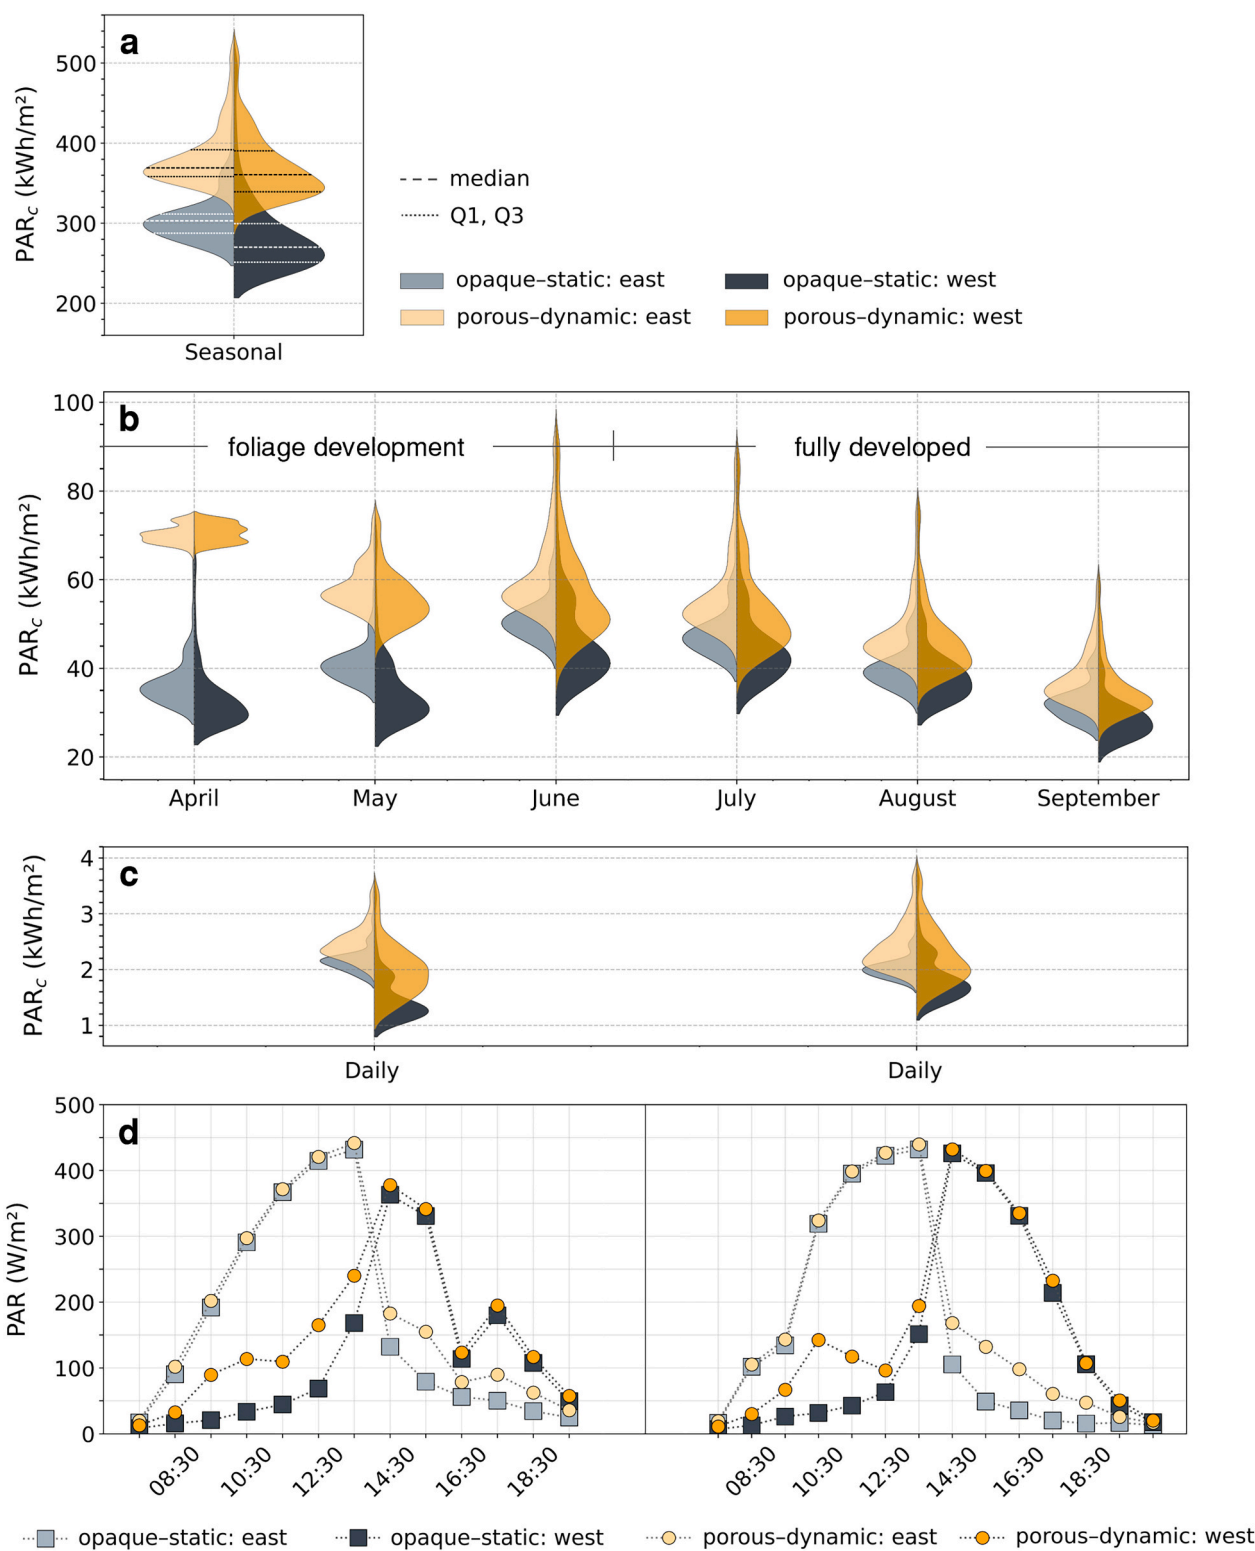


Fig. 8. Total PAR (front + back) at multiple time scales for opaque-static and porous-dynamic canopy models. (a–c) KDE-based violin distributions of PAR_c across the canopy envelope: (a) seasonal, (b) monthly, and (c) daily for two June days with contrasting skies (left: clear-to-cloudy transition; right: predominantly clear). (d) Hourly time series of height-averaged PAR for each canopy side, for the same two June days.

orchard's axis. This wider path intensified shading and increased vertical PAR_c inhomogeneity, by producing elongated distribution tails (Fig. 8b)—noticeably longer on the west-facing side. Conversely, the azimuthal asymmetry amplified horizontal PAR_c inhomogeneity by increasing the mismatch between east- and west-facing canopy sides.

Both effects gradually reduced after the summer solstice.

These trends are quantified in Table 1, where the gain in PAR_c and spatial homogeneity of the porous-dynamic are expressed relative to the opaque-static canopy model.

Daily to hourly analysis. Over the simulated period, clearer

Table 1Comparative summary of PAR_c and spatial homogeneity. Values represent the gain of the porous-dynamic relative to the opaque-static canopy model.

	Season	Apr	May	Jun	Jul	Aug	Sep
Mean PAR_c gain (%) ^a	26.2	100.5	48.7	18.6	15.3	16.0	16.6
Vertical homog. gain (%) ^{a, b}	4.9	15.7	9.1	2.7	2.2	3.3	3.2
Horizontal homog. gain (%) ^{a, b}	2.7	5.6	6.7	2.6	1.2	0.3	2.4
Overall homog. gain (%) ^{a, b}	5.4	16.6	10.9	3.2	2.2	3.0	3.3

^a computed across both leaf sides (front and back).^b homogeneity equals $1 - CV$, where CV is the coefficient of variation; horizontal homogeneity is computed between canopy sides; vertical homogeneity is computed along the height profile.

morning skies increased PAR_c on the east-facing canopy side, while cloudier afternoons produced more uniform conditions across both sides. This diurnal pattern (left panels of Figs. 8c–d) sustained horizontal inhomogeneity up to seasonal scales, and the asymmetry persisted even on predominantly clear-sky days (right panels of Figs. 8c–d).

Sunflecks partially mitigated this mismatch by increasing irradiation on the shaded canopy side. The spatiotemporal distribution of sunflecks depends on canopy architecture, row orientation, and solar position. In our scene, sunfleck intensity peaked in the morning and afternoon (Fig. 8d), when the sun was nearly perpendicular to the canopy surface, maximizing the projected open area. Conversely, as the sun aligned with the orchard's axis (around solar noon), the influence of canopy porosity diminished, as did differences between the two models.

4.2. Significance of penumbra in Agri-PV

We compared renderings using Radiance's default point-source sun (sharp-edged shadows) versus our proxy-sun area-source approach (soft, extended shadows). For the studied geometry (Fig. 9), we found that 66 proxy suns are the minimum required to produce representative penumbra. We do not include a sensitivity analysis because this threshold is geometry-specific (PV cell dimensions and mounting height).

Under clear-skies, shadows cast by upper canopy layers and PV arrays are visibly softer and spatially extended due to the penumbra effect. The effect is redistribution, not gain: daily PAR_c distributions (Fig. 10) became more homogeneous when penumbra was considered.

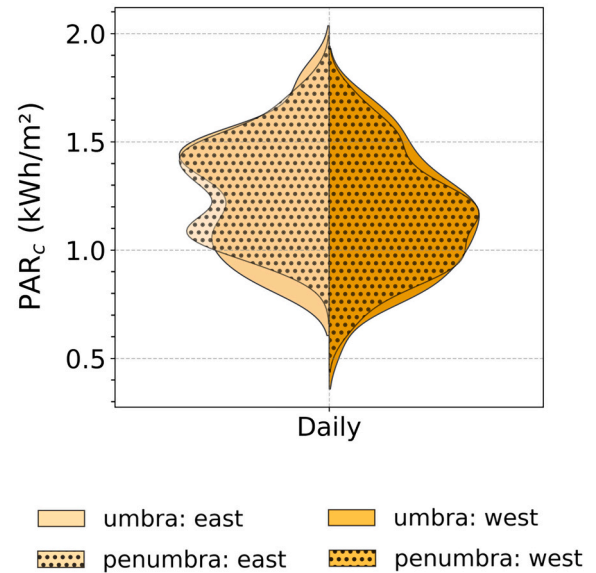


Fig. 10. Daily KDE-based violin distributions of PAR_c across the canopy envelope for the same scene as Fig. 9 under two solar treatments: point-source (umbra) and area-source (penumbra). Only front-side PAR_c is shown.




Fig. 9. Renderings under clear-sky afternoon conditions in June. Both images use the same scene including the porous-dynamic canopy; differing only in the solar representation: (a) point-source sun (umbra); (b) area-source sun (penumbra). Renderings were generated using Radiance's rpict command.

Specifically, they exhibited shorter, thinner tails and increased density near the mean—indicating a reduced likelihood of both excessive shading and overexposure. However, spatially aggregated PAR_c values remained unchanged. Modeling the sun's angular size redistributes irradiation—replacing step-like patterns with gradual PAR transitions. Over longer periods, the influence of penumbra diminishes as diffuse light dominates (predominantly cloudy conditions).

4.3. Limitations

The usefulness of our porous–dynamic canopy model is not strictly tied to how well the canopy is visually replicated. Nevertheless, the resulting architecture differs from real canopies in leaf orientation and surface area, which affects both the magnitude and spatial distribution of light (see Section 5.1 for implications). Furthermore, we did not reconstruct internal canopy layers and instead flattened foliage into 2D planar surfaces. Accounting for internal layers together with per-leaf sampling would increase occlusion and shift the PAR_c distributions (Fig. 8a–c) toward lower values.

Two further scope-related considerations concern spectral skies: (i) narrow spectral features (e.g., O_2/O_4 , CO_2) may be smoothed by Radiance's current maximum number of wavebands ($MAXCSAMP = 24$); truly high-resolution, full-spectrum runs may require multiple passes; (ii) wavelength-dependent cloud attenuation is not explicitly modeled.

We also held solar disk angular size constant at 0.533° (*gendaylit* default) and assumed uniform radiance across the disk (no limb darkening), assigning equal radiance to each proxy. These assumptions are likely negligible at the occluder scales studied here but may require investigation for detailed pergolas [64] and perforated facades [65].

5. Discussion and conclusion

We contribute a Python-based, Radiance-compatible framework that (i) raises solar spectral fidelity driven by site-specific atmospheric inputs, (ii) accelerates penumbra rendering with deterministic solar-disk sampling, and (iii) captures seasonal canopy porosity in narrow-trained orchards.

5.1. Modeling implications for narrow-trained, open-field orchards

Seasonal back-side leaf PAR is significant ($\approx 14\%$ of total) and largely insensitive to canopy model: the porous–dynamic canopy admits additional ground-reflected PAR to the back side but also increases front-side PAR , maintaining the ratio approximately constant. The exact share is dictated by ground albedo and orchard layout. Accounting for internal foliage and actual leaf-angle distributions would likely raise this share by reducing the relative dominance of front-side direct light.

Our porous–dynamic canopy captures sunflecks, yielding substantially higher seasonal light levels than the opaque–static baseline. The monthly gain is phenology-dependent, diminishing with foliage development—from $\approx 100\%$ in April to $\approx 16\%$ in July–September. After foliage maturity, the mean gain remains well below the gap fraction at full development ($GF_{fd} = 40\%$): the canopy consists of multiple planar surfaces (Fig. 5), so their combined effect (compounded by inter-row occlusion) reduces overall canopy transmittance further than a single-surface gap fraction would suggest. Beyond PAR magnitude, sunflecks also improved spatial homogeneity, with an overall seasonal gain of $\approx 5\%$ relative to the baseline; the net effect depends on row orientation, diurnal weather, and solar position dynamics. Given the minimal runtime overhead, porous–dynamic should be the default—particularly for lower-vigor canopies where porosity persists.

Consistent with Beer-Lambert formulations [66–68], direct-light penetration scales with the projected foliage area normal to the beam. For east–west rows, sunflecks peak morning and afternoon—a pattern expected to persist even with explicit internal-layer reconstruction,

because low incidence reduces planophile leaf projected area and effective optical depth.

Future work will reconstruct internal foliage using voxels (rather than 2D grids), validate the voxel-based canopy model, and perform leaf-level sampling to inform agri-PV designs that support canopy-level photosynthesis. Extending the canopy reconstruction to other training systems is non-trivial and will be tackled in future work.

5.2. Design implications for Agri-PV orchards

Light levels in open-field orchards are rarely optimal—either in magnitude or homogeneity. This framework enables site-specific diagnosis for targeted design. Although the impact of penumbra is limited under conventional PV modules, modeling them opens new design pathways in clear-sky regions. Overhead PV arrays using half-cell modules (and, prospectively narrower cells) with appropriate module-canopy separation can damp lighting extremes—mitigating sunburn and lifting minima in shaded regions.

5.3. Broader applicability

Although the canopy reconstruction approach presented here is specific to narrow-trained orchards; the remainder of the framework can be applied to any (agri-)PV scene. In particular, the upgraded sky generator has broader relevance: coupling SMARTS with *gendaylit* yields full-spectrum, atmosphere-specific sun/sky unlike visible-only alternatives—useful for daylighting and thermal studies. In addition, integrating a deterministic and efficient solar-disk sampler into *gendaylit* we support applications that demand high-fidelity penumbra rendering.

CRediT authorship contribution statement

Odysseas Alexandros Katsikogiannis: Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Olindo Isabella:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization. **Rudi Santbergen:** Writing – review & editing, Supervision. **Hesan Ziar:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization.

Funding

This work was funded by the European Union's Horizon Europe Innovation Action programme under grant agreement N° 101096352 (SYMBIOSYST). The views expressed are those of the authors and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

Data from the integrated atmospheric observation system KITcube of the Karlsruhe Institute of Technology (KIT), obtained during the TEAMx Observational Campaign, were used in this study. We gratefully acknowledge the AERONET (Aerosol Robotic NETwork) team for atmospheric data processing and quality assurance. We thank the Radiance community—especially Greg Ward and Taoning Wang—for their openness and willingness to share expertise.

Appendix A. Appendix

Collates major case-study inputs and simulation parameters. Daylight-source inputs: EURAC Meteo Browser (Laimburg station, 2023), and AERONET v2.0 (KITcube_Bolzano; AOD₅₀₀, PWV, 2023). Surfaces: parameters needed for canopy, orchard reconstruction and PV array configuration (Table A.1) and optical properties (Table A.2). Ray-tracing: key Radiance parameters governing accuracy (Table A.3). Simulation discretization: hourly timestep; 18 wavebands spanning 300–1200 nm in 50 nm intervals (PAR results use the 400–700 nm subset); 1 cm spatial sampling resolution. Simulation speed was insensitive to the number of wavebands.

Table A.1
Canopy, orchard, and PV configuration.

Category	Parameter	Value	Units	Note
Canopy setup	Leaf growth rate (k)	0.116	–	Eq. 7
	Inflection time (t_{mid})	135	days	Eq. 7 (day of year)
	Gap fraction (GF_{fd})	0.4	–	At full development Eq. 8
	Clustering coefficient (CC)	1–1.3	–	Linear seasonal variation (Eq. 6)
	Grid unit size	3	cm	2D canopy grid (leaf patch size)
Orchard layout	Canopy dimensions	$0.4 \times 3 \times 2.6$	m	$W \times L \times H$ (including trunk)
	Foliage elevation	0.5	m	Offset from ground to foliage base
	Row spacing	2.5	m	Open-field scenario (no PV)
	Row axis azimuth	110/290	°	Measured clockwise from north
	Cell dimensions	18.2×9.1	cm	$W \times L$ (half-cells)
PV module	Cell spacing	1×6	cm	$W \times L$
	Number of cells	6×12	–	$W \times L$
PV array	Hub height	4.0	m	Ground to torque tube center
	Row spacing	3.25	m	Agri-PV scenario (trackers & orchard)
	Tracker axis azimuth	110/290	°	Measured clockwise from north

Table A.2
Optical properties.

Element	Property	Model	Note
Foliage	$\rho(\lambda)$	Lambertian	Black cherry leaves from [69]
Ground	$\rho(\lambda)$	Lambertian	Measured site albedo [60] (seasonally constant)
PV cells	$\rho(\lambda)$	Mixed	Measured (bare cell)
PV frame	$\rho(\lambda)$	Mixed	Weathered aluminum from [70] (<i>specularity</i> = 0.15)
PV glass	$\tau(\lambda)$	Dielectric	Measured (normal incidence)

$\rho(\lambda)$: spectral hemispherical reflectance; $\tau(\lambda)$: spectral hemispherical transmittance. Spectra resampled to 18 bands and applied via Radiance's spectral modifiers. A mixed optical model combines Lambertian diffuse and specular reflection components.

Table A.3
Radiance parameters (disabling selected built-in algorithms).

Parameter	Value	Note
Ambient bounces (-ab)	2	Include 1st order reflections from sky and 2nd order from sun.
Ambient divisions (-ad)	1024	Set high enough to reduce Monte Carlo sampling error.
Ambient accuracy (-aa)	0	Disable interpolation; enforce hemispherical sampling at each point.
Direct threshold (-dt)	0	Disable selective shadow testing to ensure visibility of all proxy suns.
Direct jittering (-dj)	0	Disable stochastic jitter; penumbra handled via proxy-suns.
Direct sampling (-ds)	0	Disable source subdivision; proxy-sun method takes over.

Generative AI statement

During the preparation of this work, the authors used ChatGPT (OpenAI) to improve readability and for programming support (e.g., refactoring, debugging). The authors reviewed and edited all AI-assisted text and code and take full responsibility for the submitted manuscript.

Data availability

The Python code and sample input/output for the sun–sky generator are made publicly available.

References

- [1] Wong IL. A review of daylighting design and implementation in buildings. *Renew Sust Energ Rev* 2017;74:959–68. <https://doi.org/10.1016/j.rser.2017.03.061>.

- [2] Honningdalsnes EH, Marstein ES, Nygård MM, Wiig MS, Riise HN. Benchmarking irradiation models for photovoltaic applications: a comparative analysis of radiance-based tools. *Sol Energy* 2025;296:113566. <https://doi.org/10.1016/j.solener.2025.113566>.
- [3] Pao Y-C, Kahlen K, Chen T-W, Wiechers D, Stützel H. How does structure matter? Comparison of canopy photosynthesis using one- and three-dimensional light models: a case study using greenhouse cucumber canopies. *In Silico Plants* 2021;3(2):diab031. <https://doi.org/10.1093/inilicoplants/diab031>.
- [4] Oswald D, Pourreza A, Chakraborty M, Khalsa SDS, Brown PH. 3D radiative transfer modeling of almond canopy for nitrogen estimation by hyperspectral imaging. *Precis Agric* 2025;26:12. <https://doi.org/10.1007/s11119-024-10207-z>.
- [5] Asa'a S-N, Kaaya I, Dupon O, de Jong R, van der Heide A, et al. Assessing the energy yield and irradiation distribution in fixed and tracking agrivoltaic orchards. *Proc. 41st Eur. Photovolt. Sol. Energy Conf. Exhib. (EUPVSEC)*. 2024. <https://doi.org/10.4229/EUPVSEC2024/4DV.1.33>.
- [6] Mouhib A, Fernández-Solas Á, Pérez-Higueras PJ, Fernández-Ocaña AM, Micheli L, et al. Enhancing land use: integrating bifacial PV and olive trees in agrivoltaic systems. *Appl Energy* 2024;359:122660. <https://doi.org/10.1016/j.apenergy.2024.122660>.
- [7] Bruno M, Gfüllner LJ, Berwind MF. Enhancing agrivoltaic synergies through optimized tracking strategies. *J Photon Energy* 2025;15(3):032703. <https://doi.org/10.1117/1.JPE.15.032703>.
- [8] Ward GJ. The RADIANCE lighting simulation and rendering system. *Comput Graph* 1994;28(2):459–72. <https://doi.org/10.1145/192161.192286>.
- [9] Scalisi A, O'Connell MG, Stefanelli D, Zhou S, Pitt T, et al. Narrow orchard systems for pome and stone fruit—a review. *Sci Hortic* 2024;338:113815.
- [10] Dorigoni A, Micheli F. Guyot training: a new system for producing apples and pears. *Eur Fruit Mag* 2018;2:18–23. <https://openpub.fmach.it/handle/10449/54108>.
- [11] CherryTimes. Gregory Lang: cherry orchard training system innovations. <https://cherrytimes.it/en/news/gregory-lang-cherry-orchard-training-system-innovations>; 2023. accessed 10 Apr 2025.
- [12] Apple and Pear Australia Ltd. Lessons learnt from Covid-19. <https://apa.org.au/lessons-learnt-from-covid-19/>; 2021. accessed 10 Apr 2025.
- [13] Gkaintatzi Masouti M, Pierson C, van Duijnhoven J, Andersen M, Aarts M. A simulation tool for building and lighting design considering ipRGc-influenced light responses. *E3S Web of Conferences* 2022;362:01001. <https://doi.org/10.1051/e3sconf/202236201001>.
- [14] Lawrence Berkeley National Laboratory, 2025. Radiance, version 6.0. Release notes: <https://radsite.lbl.gov/radiance/dist/ReleaseNotes.html> (accessed 20 Sep 2025). GitHub release: <https://github.com/LBNL-ETA/Radiance/releases/tag/ra6R0P1a>.
- [15] Bruneton E, Neyret F. Precomputed atmospheric scattering. *Comput Graph Forum* 2008;27(4):1079–86. <https://doi.org/10.1111/j.1467-8659.2008.01245.x>.
- [16] Solemma. ALFA (Adaptive Lighting for Alertness). Retrieved 20 June 2025, from <https://www.solemma.com/alfa/>; 2017.
- [17] Alight A, Jakubiec JA. An ipRGc-influenced/non-visual spectral occupant model (INSOM) for lighting design, Part 1: Light simulation method. *Light Res Technol* 2025;1–17. <https://doi.org/10.1177/14771535251368379>.
- [18] Gueymard CA. Parameterized transmittance model for direct beam and circumsolar spectral irradiance. *Sol Energy* 2001;71(5):325–46. [https://doi.org/10.1016/S0038-092X\(01\)00054-8](https://doi.org/10.1016/S0038-092X(01)00054-8).
- [19] Larson GW, Shakespeare RA. Rendering with radiance: The art and science of lighting visualization. San Francisco: Morgan Kaufmann; 1998. <https://radsite.lbl.gov/radiance/book/>.
- [20] de Vries SW, Chamilothoni K, Aarts MPJ. Experimental validation of radiance-based methods for simulating solar penumbra and pinhole projections. *LEUKOS* 2025;21(3):235–56. <https://doi.org/10.1080/15502724.2024.2365691>.
- [21] Bailey BN, Ochoa MH. Semi-direct tree reconstruction using terrestrial LiDAR point cloud data. *Remote Sens Environ* 2018;208:133–44. <https://doi.org/10.1016/j.rse.2018.02.013>.
- [22] Escolà A, Peña JM, López-Granados F, Rosell-Polo JR, de Castro AI, et al. Mobile terrestrial laser scanner vs. UAV photogrammetry to estimate woody crop canopy parameters – part 1: methodology and comparison in vineyards. *Comput Electron Agric* 2023;212:108109. <https://doi.org/10.1016/j.compag.2023.108109>.
- [23] Farhan SM, Yin J, Chen Z, Memon MS. A comprehensive review of LiDAR applications in crop management for precision agriculture. *Sensors* 2024;24(16):5409. <https://doi.org/10.3390/s24165409>.
- [24] Scher CL, Griffoul E, Cannon CH. Drone-based photogrammetry for the construction of high-resolution models of individual trees. *Trees* 2019;33:1385–97. <https://doi.org/10.1007/s00468-019-01866-x>.
- [25] Meyer L, Gilson A, Scholz O, Stammering M. CherryPicker: semantic skeletonization and topological reconstruction of cherry trees. *Proc IEEE/CVF Conf Comput Vis Pattern Recognit Workshops (CVPRW)*. 2023. p. 6244–53. <https://openaccess.thecvf.com/content>.
- [26] Qiu T, Wang T, Han T, Kuehn K, Cheng L, et al. AppleQSM: geometry-based 3D characterization of apple tree architecture in orchards. *Plant Phenomics* 2024;6:0179. <https://doi.org/10.34133/plantphenomics.0179>.
- [27] Xia Y, Li H, Zhang F, Sun G, Qi K, et al. OrchardQuant-3D: combining drone and LiDAR to perform scalable 3D phenotyping for characterising key canopy and floral traits in fruit orchards. *Plant Biotechnol J* 2025. <https://doi.org/10.1111/pbi.70229>.
- [28] Ando R, Okura F, Chegini T, Nagahara H, Mita T. Robust surface reconstruction of plant leaves from 3D point clouds. *Plant Methods* 2021;17:52. <https://pubmed.ncbi.nlm.nih.gov/33860276/>.
- [29] Whebell RM, Moroney TJ, Turner IW, Pethiyagoda R, McCue SW. Implicit reconstructions of thin leaf surfaces from large, noisy point clouds. arXiv:10286. <https://arxiv.org/abs/2009.10286>; 2020.
- [30] Boukhana M, Ravaglia J, Hétyroy-Wheeler F, de Solan B. Geometric models for plant leaf area estimation from 3D point clouds: a comparative study. *Graph Vis Comput* 2022;2:200057. <https://doi.org/10.1016/j.gvc.2022.200057>.
- [31] Kerbl B, Kopanas G, Leimkühler T, Drettakis G. 3D Gaussian splatting for real-time radiance field rendering. *ACM Trans Graph* 2023;42(4). <https://doi.org/10.1145/3592433>. Art. 139, 1–14.
- [32] Chen Y, Xiao K, Gao G, Zhang F. High-fidelity 3D reconstruction of peach orchards using a 3DGS-ag model. *Comput Electron Agric* 2025;234:110225. <https://doi.org/10.1016/j.compag.2025.110225>.
- [33] Guédon A, Lepetit V. SuGaR: surface-aligned Gaussian splatting for efficient 3D mesh reconstruction and high-quality mesh rendering. In: *proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 5354–5363. 2024. <https://doi.org/10.1109/CVPR52733.2024.00512>.
- [34] Shen P, Jing X, Deng W, Jia H, Wu T. PlantGaussian: exploring 3D Gaussian splatting for cross-time, cross-scene, and realistic 3D plant visualization and beyond. *Crop J* 2025;13(2):607–18. <https://doi.org/10.1016/j.cj.2025.01.011>.
- [35] van der Heijden G, de Visser PHB, Heuvelink E. Measurements for functional-structural crop models. In: Vos J, Marcelis LFM, de Visser PHB, Struijk PC, Evers JB, editors. *Functional-Structural Plant Modelling in Crop Production*. Dordrecht: Springer; 2007. p. 13–25.
- [36] Grisafi F, DeJong TM, Tombesi S. Fruit tree crop models: an update. *Tree Physiol* 2022;42(3):441–57. <https://doi.org/10.1093/treephys/tpab126>.
- [37] DeJong TM. Simulating fruit tree growth, structure, and physiology using L-systems. *Crop Sci* 2022;62(6):2091–106. <https://doi.org/10.1002/csc.220847>.
- [38] Ayala Pelaez S, Deline C. Bifacial radiance: a python package for modeling bifacial solar photovoltaic systems. *J Open Source Softw* 2020;5(50):1865. <https://doi.org/10.21105/joss.01865>.
- [39] Perez R, Seals R, Michalsky J. All-weather model for sky luminance distribution—preliminary configuration and validation. *Sol Energy* 1993;50:235–45. [https://doi.org/10.1016/0038-092X\(93\)90017-1](https://doi.org/10.1016/0038-092X(93)90017-1).
- [40] Ruiz-Arias JA, Gueymard CA. Solar irradiance component separation benchmarking: the critical role of dynamically-constrained sky conditions. *Renew Sust Energy Rev* 2024;202:114678. <https://doi.org/10.1016/j.rser.2024.114678>.
- [41] Ruiz-Arias JA, Gueymard CA. GISPLIT: high-performance global solar irradiance component-separation model dynamically constrained by 1-min sky conditions. *Sol Energy* 2024;269:112363. <https://doi.org/10.1016/j.solener.2024.112363>.
- [42] Ruiz-Arias JA, Gueymard CA. CAELUS: classification of sky conditions from 1-min time series of global solar irradiance using variability indices and dynamic thresholds. *Sol Energy* 2023;263:111895. <https://doi.org/10.1016/j.solener.2023.111895>.
- [43] Gueymard CA. Prediction and validation of cloudless shortwave solar spectra incident on horizontal, tilted, or tracking surfaces. *Sol Energy* 2008;82:260–71. <https://www.sciencedirect.com/science/article/pii/S0038092X07001004>.
- [44] Ruiz-Arias JA, Gueymard CA. A multi-model benchmarking of direct and global clear-sky solar irradiance predictions at arid sites using a reference physical radiative transfer model. *Sol Energy* 2018;171:447–65. <https://doi.org/10.1016/j.solener.2018.06.048>.
- [45] Gueymard CA. SMARTS code, version 2.9.5: user's manual. Solar consulting services; 2005. https://www.solarconsultingservices.com/SMARTS295_Users_Manual_PC.pdf.
- [46] National Renewable Energy Laboratory, 2025. pySMARTS: Python wrapper for SMARTS. GitHub repository: <https://github.com/NREL/pySMARTS> (accessed 10 Jan 2025).
- [47] Bridson R. Fast Poisson disk sampling in arbitrary dimensions. *ACM SIGGRAPH 2007 Sketches* 2007. <https://doi.org/10.1145/1278780.1278807>.
- [48] Lakso AN, Corelli-Grappadelli L. Leaf area development patterns in young pruned and unpruned apple trees. *J Amer Soc Hortic Sci* 1994;119:391–6. <https://doi.org/10.21273/JASHS.109.6.861>.
- [49] Forshey CG, Elfving DC. The relationship between vegetative growth and fruiting in apple trees. *Hortic Rev* 1989;11:229–67. <https://doi.org/10.1002/9781118060841.ch7>.
- [50] Balakrishnan P, Jakubiec JA. Measuring light through trees for daylight simulations: a photographic and photometric method. *Proc BSO* 2016;2016:343–52. https://publications.ibpsa.org/conference/paper/?id=bso2016_1152.
- [51] Chianucci F, Pisek J, Raabe K, Marchino L, Ferrara C, Corona P. A dataset of leaf inclination angles for temperate and boreal broadleaf woody species. *Ann For Sci* 2018;75:50. <https://doi.org/10.1007/s13595-018-0730-x>.
- [52] Mardaljevic J. Validation of a lighting simulation program under real sky conditions. *Light Res Technol* 1995;27:79–86. <https://doi.org/10.1177/14771535950270040701>.
- [53] Reinhart CF, Walkenhorst O. Validation of dynamic RADIANCE-based daylight simulations for a test office with external blinds. *Energy Buildings* 2001;33:683–97. [https://doi.org/10.1016/S0378-7788\(01\)00058-5](https://doi.org/10.1016/S0378-7788(01)00058-5).
- [54] Mardaljevic J. Verification of program accuracy for illuminance modelling: assumptions, methodology and an examination of conflicting findings. *Light Res Technol* 2004;36:217–39. <https://doi.org/10.1191/1477153504lii20oa>.
- [55] Dittmann A, Sanchez H, et al. Comparative analysis of albedo measurements (plane-of-array and horizontal) at multiple sites worldwide. In: *Proc. 36th EU PVSEC*; 2019. <https://doi.org/10.4229/EUPVSEC20190219-5DO.1.4>. Marseille, France.
- [56] Marion W. Measured and satellite-derived albedo data for estimating bifacial photovoltaic system performance. *Sol Energy* 2021;215:321–7. <https://doi.org/10.1016/j.solener.2020.12.050>.

- [57] Riedel-Lyngskær N, Ribaconka M, Pó M, Thorseth A, Thorsteinsson S, Dam-Hansen C, et al. The effect of spectral albedo in bifacial photovoltaic performance. *Sol Energy* 2022;231:921–35. <https://doi.org/10.1016/j.solener.2021.12.023>.
- [58] Ziar H, Sönmez FF, Isabella O, Zeman M. A comprehensive albedo model for solar energy applications: geometric spectral albedo. *Appl Energy* 2019;255:113867. <https://doi.org/10.1016/j.apenergy.2019.113867>.
- [59] Riedel-Lyngskær N, Ribaconka M, Pó M, Thorseth A, Thorsteinsson S, Dam-Hansen C, et al. Spectral albedo in bifacial photovoltaic modeling: what can be learned from onsite measurements? In: *proc. 48th IEEE PVSC*, pp. 942–949. 2021. <https://doi.org/10.1109/PVSC43889.2021.9519085>.
- [60] Riedel-Lyngskær N, Berrian D, Alvarez-Mira D, Protti AA, Thorsteinsson S, Poulsen PB, et al. Validation of bifacial PV simulation software against monitoring data from large-scale single-axis trackers and fixed-tilt systems in Denmark. *Appl Sci* 2020;10:8487. <https://doi.org/10.3390/app10238487>.
- [61] Katsikogiannis OA, Isabella O, Ziar H. Modelling, design, and performance evaluation of Agri-PV orchards (session 4DO.3.5). In: *Proc. 41st EU PVSEC*; 2024. Vienna, Austria.
- [62] ISO. ISO 9060:2018—Solar energy—Specification and classification of instruments for measuring hemispherical solar and direct solar radiation. Geneva, Switzerland: International Organization for Standardization; 2018.
- [63] Holben BN, Eck TF, Slutsker I, Sinyuk A, Schafer JS, Giles DM, et al. AERONET'S version 2.0 quality assurance criteria. *Proc. SPIE* 6408, remote sensing of the atmosphere and clouds, 64080Q. 2006. <https://doi.org/10.1117/12.706524>.
- [64] Salazar Trujillo JH. Calculation of the shadow–penumbra relation and its application on efficient architectural design. *Sol Energy* 2014;110:139–50. <https://doi.org/10.1016/j.solener.2014.08.043>.
- [65] Chamilothoni K, Wienold J, Moscoso C, Matusiak B, Andersen M. Subjective and physiological responses towards daylight spaces with contemporary façade patterns in virtual reality: influence of sky type, space function, and latitude. *J Environ Psychol* 2022;82:101839. <https://doi.org/10.1016/j.jenvp.2022.101839>.
- [66] Monsi M, Saeki T. On the factor light in plant communities and its importance for matter production. *Ann Bot* 2005;95:549–67. <https://doi.org/10.1093/aob/mci052>.
- [67] Wilson JW. Inclined point quadrats. *New Phytol* 1960;59:1–8. <https://doi.org/10.1111/j.1469-8137.1960.tb06195.x>.
- [68] Nilson T. A theoretical analysis of the frequency of gaps in plant stands. *Agric Meteorol* 1971;8:25–38. [https://doi.org/10.1016/0002-1571\(71\)90092-6](https://doi.org/10.1016/0002-1571(71)90092-6).
- [69] Hosgood B, Jacquemoud S, Andréoli G, Verdebout J, Pedrini A, Schmuck G. Leaf Optical properties experiment 93 (LOPEX93). European Commission, Joint Research Centre; 1995. EUR 16095 EN.
- [70] Baldrige AM, Hook SJ, Grove CI, Rivera G. The ASTER spectral library version 2.0. *Remote Sens Environ* 2009;113(4):711–5. <https://doi.org/10.1016/j.rse.2008.11.007>.