Contents lists available at ScienceDirect

Renewable and Sustainable Energy Reviews

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

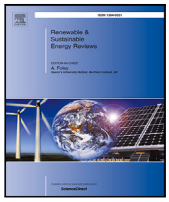

Highlights

Shadow modeling in urban environments for solar harvesting devices with freely defined positions and orientations

Andrés Arias-Rosales, Philip R. LeDuc*

Renewable and Sustainable Energy Reviews xxx (xxxx) xxx

- Four shadow modeling approaches are developed, characterized, and compared.
- Five testing scenarios based on different locations, urban settings, and weather.
- The analytical approach is orders of magnitude faster.
- Focused ray tracing is competitive in reliability, performance, and simplicity.
- Power RMSE decrease from 3.8% to 1.11% with shadow modeling convergence.

Contents lists available at ScienceDirect

Renewable and Sustainable Energy Reviews

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

Shadow modeling in urban environments for solar harvesting devices with freely defined positions and orientations

Andrés Arias-Rosales, Philip R. LeDuc *

Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, United States

ARTICLE INFO

Keywords:
Shadow modeling
Solar energy
Partial shading
Photovoltaic
Ray tracing
Rasterization

ABSTRACT

As a sustainable alternative regarding environmental impact, cost-effectiveness, and social integration, solar energy is expected to become an ever more ubiquitous part of our intricate human world. Dropping prices and growing demand are making it more viable for a variety of solar devices to be implemented in urban and other complex environments. From devices helping people meet their energy needs to solar-powered drones fulfilling urban services like maintenance, security, carrying goods, or even transporting people. These environments involve constrained and dynamic conditions, encouraging the use of solar harvesting devices that can freely adopt tailor-made positioning and tracking strategies to make the most of available resources. A crucial challenge is improving the geometrical flexibility and efficiency of modeling capabilities. In particular, developing practical approaches that account for detailed shadow effects in complex scenarios can be computationally challenging, and it is not clear how different approaches compare face-to-face in urban contexts and with freely defined harvesting surfaces. In this work, four shadow modeling approaches are developed and demonstrated in urban scenes of varying complexity; accuracy and precision are characterized versus computational cost; run-time trends are analyzed as functions of scene complexity, and energy estimation implications are examined. The approaches converge within 1% deviations, and the highest performing approach is three orders of magnitude faster than the most computationally costly. This work supports the selection and development of accurate, efficient, and flexible modeling frameworks that will play a role in enabling a diverse range of solar harvesting devices in challenging urban environments.

1. Introduction

The steadily falling prices of Photovoltaics (PV) (by 90% from 2007–2017 [1]), with collectors that can harvest clean energy in a decentralized, safe, and modular manner are making it more viable and sustainable [2] for many types of solar devices to be implemented in urban [3,4] and other complex settings (Fig. 1). As a large fraction of energy demand comes from cities [5], bringing energy production closer to consumers reduces transmission and distribution losses [6,7] and may make the electrical grid more efficient and adaptable. There has been notable research interest in urban solar harvesting, including solar accessibility in developing cities [4,8], energy implications of shadow distribution on roofs [9], solar potential maps [10] toward more sustainable urban planning [11], energy-efficient architecture [12], solar-powered aerial vehicles [13], fixed building-integrated PV [14–16], and dynamic building envelopes [17]. Other complex environments include heliostats [18], step-like fields [19], and wind–solar dual land use [20]. In this work, "complex" refers to scenarios where the surrounding objects cannot be reduced to simple analytical expressions and their shape, the shape of the available harvesting area, and the positions and orientations of the harvesting surfaces can be freely defined.

There are many types of potentially shadow-casting objects that can be relevant for solar harvesting in intricate settings (Fig. 1(B)). Near objects [21] include buildings [13], façade prominences [22], trees, cables, chimneys, antennas, poles or even the surrounding terrain [23]. Self-shadows [24,25] include near solar devices (Fig. 1(A)) or parts from the same device, such as side-walls in solar cookers [26] or trough fins in solar heaters [27]. Distant objects or the topography at the horizon may also be considered [28,29]. The cast shadows may be "soft" [30], which can reduce the overall intensity of the accessible solar irradiance, or they may be "hard" partial shadows [30,31], with abrupt changes in irradiance distribution. This can lead to PV degradation, bumpy IV curves [32], or dangerous voltage and current mismatches and hot-spots [32]. This work focuses on modeling the "hard" shadows that result from blocking (partially or fully) the incoming Beam radiation. Many studies focus on Beam shadows (as opposed

* Corresponding author.
E-mail addresses: aariasro@andrew.cmu.edu (A. Arias-Rosales), prl@andrew.cmu.edu (P.R. LeDuc).

https://doi.org/10.1016/j.rser.2022.112522
Received 25 March 2022; Received in revised form 19 April 2022; Accepted 26 April 2022

Nomenclature

Acronyms and Names:

Ana	Analytical Areas approach for Beam shadow modeling
B	Simulated cuboid volumes as building objects
CL	Convergence Level
FRT	Focused Ray Tracing approach for Beam shadow modeling
GT	Naive Ground Truths
MAD	Median Absolute Deviations
Med	Median
NSRDB	National Solar Radiation Database
Pix	Rasterization and Pixel-Counting approach for Beam shadow modeling
POV	Point Of View
PV	Photovoltaics in general or simulated harvesting surfaces
Reg	Regression model
RMSE	Root-Mean-Square Errors or deviations
RT	Forward Ray Tracing approach for Beam shadow modeling
SD	Standard Deviation

Variable subscripts:

A	Related to Albedo radiation
a	Related to ambient temperature
B	Related to Beam radiation
cov	Related to the PV covering
DC	Related to Circumsolar radiation
DH	Related to Horizon Brightening radiation
DI	Related to Sky Isotropic radiation
gr	Related to ground
sky	Related to sky

Variables:

α_T	Total absorptance
β	Inclination of the harvesting surface
ϵ	Emissivity
η_{PV}	Effective PV efficiency
η_{ref}	Reference PV efficiency reported by the manufacturer at T_{ref} temperature
σ	Stefan–Boltzmann constant
τ	Transmittance factor
B	PV temperature coefficient
C	Thermal conductivity
h	Convection coefficient
I	Incident Irradiance
$I\tau s$	Total incident irradiance considering transmittance, shading, and electrical mismatch losses
$I\tau sg$	Same as $I\tau s$ but without NSB considerations, i.e., shadow implications are only geometrical
Is	Total incident irradiance unobstructed by shadows
L	PV covering thickness
M/M_{ref}	Spectral response of the PV material
NB	Number of PV blocks in a PV panel
$NRays$	The total number of rays
NSB	Number of shaded PV blocks
P	Power output
R	Unitless radius for relative comparisons between simulated scenes
rT	Relative run-time
T	Temperature
USF_B	Unshaded Factor for Beam radiation

to Diffuse shadows) because they usually account for most of the shading-induced energy losses [11,33], and Beam is usually regarded as the radiation component that is the most affected by shadow-casting structures [9,19,34–36].

Some common simplifying constraints and assumptions are useful in traditional solar harvesting scenarios, e.g., solar farms in open spaces, but may be limiting or lead to considerable inaccuracies in more complex and dynamic environments. Regarding positioning, solar harvesting devices are often arranged in uniform and fixed grids [36,37] along with fixed orientations [34], and when active tracking is used, the trackers are often assumed to be aligned with each other [19,24,35,37]. This simplifies energy estimations, especially concerning potentially intricate shadowing effects [6,11], but it may restrain flexibility in environments with challenging constraints [36]. In complex environments, more dynamic tracking [13] and positioning strategies [28] can improve energy efficiency and versatility. For instance, solar modules may be positioned in 3D arrangements [34], in non-regular grids [24,36], or use individualized [35] and context-aware tracking (e.g., backtracking [22]) to minimize shadows and maximize harvesting.

Shadow modeling can be crucial in the estimation of power output [21,29] and the evaluation and optimization of solar collectors [19]. It can greatly influence design decisions [33] and play a major role in making the most of available resources, such as enabling efficient land use [37]. Developing accurate, efficient, and flexible modeling frameworks will play a role in enabling a diverse range of solar harvesting devices, namely, systems that harvest solar energy to provide a variety of urban services, such as helping people meet their energy needs in their homes (e.g., building-integrated PV) or vehicles (e.g., vehicle-integrated PV [38]), assisting traffic control, public lighting, and drones for surveillance or carrying goods [39,40].

There are a variety of ways to account for Beam shadows toward solar energy estimation. Ray-tracing techniques [10,33] are among the prevailing ways of simulating how solar radiation interacts with an environment. In the context of Beam shadows, solar radiation can be modeled as parallel rays that behave according to the laws of geometrical optics. This is known to be a good approximation of radiation when the scale of the objects involved is orders of magnitude larger than the wavelength [33]. In “forward” ray-tracing, a set of rays are generated in the direction of the light source (i.e., the sun) to sample the space of interest, and then the number of rays that reach the harvesting surfaces are counted. Some computational resources may be wasted calculating ray paths that do not interact with harvesting surfaces, so “backward” ray-tracing [33] generates the rays directly from these surfaces and analyzes whether the rays are obstructed. On the other hand, instead of counting unobstructed rays, rasterization techniques count the number of pixels that correspond to the plotted harvesting surfaces and that are visible from the sun’s point of view [21].

Analytical shadow modeling approaches provide direct results that do not depend on resolution (e.g., the number of rays or pixels used to sample the space). They differ in the geometries that they allow [29], as

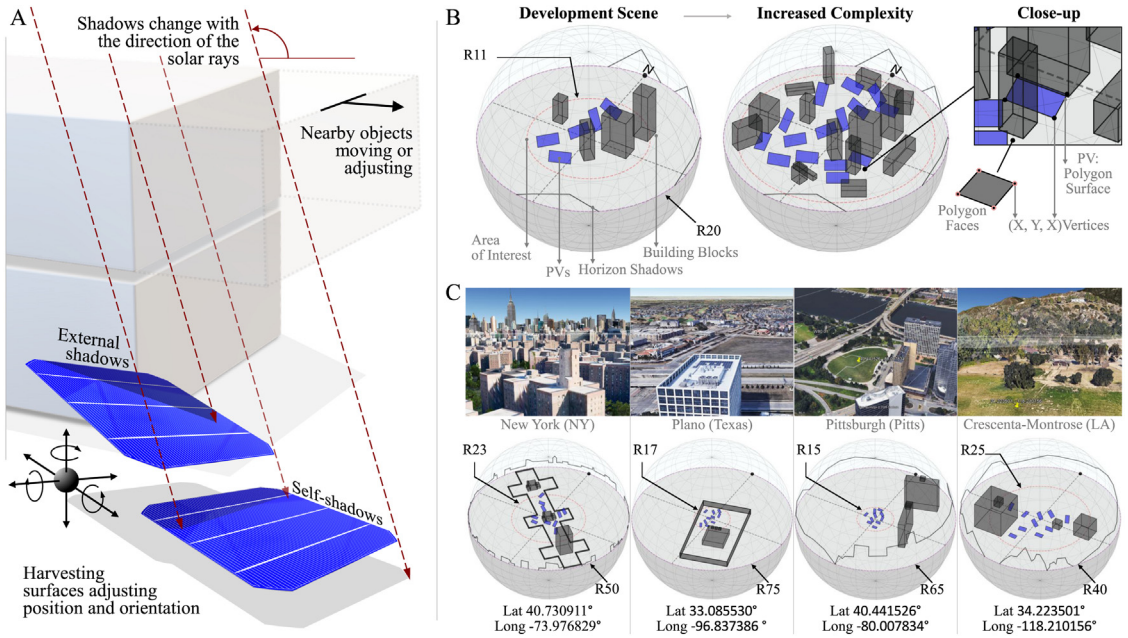

Fig. 1. Shadow modeling for urban solar harvesting. Shadows that hinder the harvesting performance (A) are dependent on solar position, the positions and orientations of the harvesting surfaces, the surrounding objects, and the distant silhouettes at the horizon (B). A virtual environment is developed to test and demonstrate different shadow modeling approaches in a range of urban scenes. The "Development" scene (B) is a generic setting with varying complexity used as the main testing ground. Four additional scenes (C) are generated based on real locations (screenshots from Google Earth®, 2021) to further explore the modeling capabilities and test the main findings. Radii (R) are unitless and only for relative comparisons.

well as in the practical restrictions and simplifications that they entail. Using vector analysis, linear algebra, and trigonometric relationships, they typically project the shadow-casting geometries along the Beam's direction and delimit the shaded area. Some are elegantly simple and effective, but constrained to specific applications, such as shadows cast over windows [41], stratospheric aerostats with solar arrays [42], solar cookers [26], roofs as common 3D shapes [9], wind turbine poles [20], fixed parallel solar arrays [19,43], and solar trackers aligned with each other [24,35,37]. Overall, the more general and detailed these methods are, the more complex and difficult to implement they are [21,25,44]. One main challenge is tracing the relevant isolated contours from arbitrarily intricate overlapping geometries, which may be addressed with convex-hull and polygon-clipping algorithms [6,12,29,36].

Developing practical modeling frameworks that account for detailed shadow effects in complex scenarios can be computationally challenging [21,29,45,46]. Also, computational cost has been reported to increase with scene complexity [11,29], which may be limiting for some applications. From the published literature, it is not well understood how the different Beam shadow modeling approaches compare in terms of accuracy, precision, run-time efficiency, and practicality, especially in intricate urban contexts and with the position and orientation of the harvesting surfaces being freely and dynamically defined. How to make a one-to-one comparison considering that these approaches are based on different practical conceptualizations, imply different underlying assumptions, work with different resolution variables, and their performance is expected to be a function of the specificities of the scene? These are some of the questions motivating this work.

Given the growing interest in solar harvesting in urban and other intricate settings, the push toward more flexible modeling frameworks, and the questions about comparative performance that remain, the novelty and contributions of this work are (I.) Four shadow modeling approaches that allow for flexible positioning and orientation (Figs. 3–6), developed to have control over most constraints and assumptions, and which are demonstrated in five virtual urban scenes with varying complexity (Fig. 1); (II.) characterizing the trade-off between accuracy and computational cost (Figs. 7–8); (III.) directly comparing the approaches (Table 1 and Fig. 9) in terms of computational cost, accuracy, and precision based on discretized convergence levels and significance heat maps; (IV.) characterizing the run-time trends as functions of scene complexity (Fig. 10); and (V.) exploring the implications in energy estimation (Figs. 11–13). These computational capabilities are developed and characterized seeking to find a balance between accuracy, ease of implementation, geometrical flexibility, and computational cost. This work supports the development, selection, and set-up of computationally efficient modeling frameworks, which can enable designs and optimizations that will help to diversify and broaden the adoption of solar harvesting technologies.

In the following sections, the virtual testing scenarios are described (Section 2.1), the shadow modeling approaches are presented (Section 2.2), the comparative strategy and energy model are discussed (Sections 2.3–2.4), the results on convergence, run-time versus scene complexity, and energy implications are described (Sections 3.1, 3.2, and 3.3, respectively), and the main findings and implications are summarized in Section 4. Please refer to the Supplementary Material for relevant data and code and additional details on the shadow modeling approaches and the energy model.

2. Methods

Fig. 2 summarizes the methodological steps and phases of this work, from the development of the code "infrastructure" to the main outputs and results. The steps are discussed in further detail in Sections 2.1–2.4 and in the Supplementary Material.

2.1. Virtual environment and testing scenarios

To test the different shadow modeling approaches, a virtual environment is developed in Python 3.0. The virtual scenes are generated using basic functions from the *Matplotlib* library for visualization and *Numpy* for geometric constructions, projections, and for employing linear algebra in space manipulations. Scenes (Fig. 1(B)–(C)) are represented mainly with PV and Building objects (Bs).

PVs are flat polygon surfaces and Bs are cuboid volumes bounded by polygon faces. Every polygon is characterized by four vertices

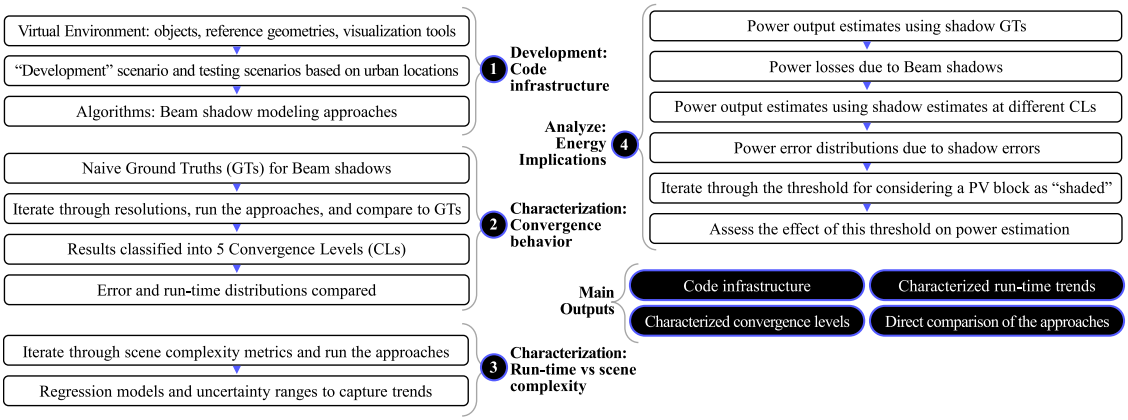

Fig. 2. Methodological Steps.

with (X, Y, Z) coordinates. PVs can have any arbitrary length, width, number of sub-modules (i.e., "blocks", see Section 2.2) represented by the number of block columns and rows, any roll angle (about an axis perpendicular to the PV surface at its center), azimuth angle (North-clockwise direction of the normal vector), and inclination angle from the horizontal. Similarly, Bs can have any width, length, height, and North-clockwise azimuth direction. Large Bs objects can be composed of several adjacent smaller cuboids so that the shadow modeling approaches perceive the relative distances more accurately. Bs and PVs are first generated at the center of the scene, rotated using rotation matrices (using linear algebra), and then displaced to their target (X, Y, Z) positions. All PVs can freely and dynamically adjust their positions and orientations in these recreated virtual worlds.

As an angular reference of the sky and ground domes, a spherical surface is centered ideally with respect to the PV objects. An area of interest can be delimited with a circle to focus the shadow modeling efforts and make better use of the space-sampling resources. Distant objects and surrounding topography can be represented as Horizon Shadows projected onto the spherical sky surface. Similar to [28,29], Horizon Shadows are represented as sets of (Azimuth, Elevation) coordinates and projected onto the spherical sphere. These projected shapes cast shadows at an angular level. Namely, if Beam radiation from a given Point Of View (POV) direction falls within a Horizon Shadow, all radiation from that POV is considered blocked. This is analogous to when the solar disc falls behind the top of a mountain and objects in the local scene can no longer access Beam radiation, regardless of their local positions and orientations. Considering reflections from nearby objects is not part of the scope of this work. Similar studies have described such considerations as impractical, by Cascone et al. [29], or often negligible, by Erdélyi et al. [7]. Nonetheless, that may be an interesting addition to the modeling problem that is worth exploring in future work.

A hypothetical "Development" scene, located in Chicago, is generated as the main environment for debugging, testing, and characterizing the different approaches (Fig. 1(B)). Additionally, to test the approaches in more diverse conditions, four scenes are generated based on real urban locations (explored using Google Earth®) in New York, Plano Texas, Pittsburgh, and near Los Angeles (Fig. 1(C)). These locations are meant to represent a range of different weather and latitude conditions, and the scenes can vary in the number of PV and Building objects to represent different complexities. For power calculations, meteorological data of the 5 locations is obtained through The National Solar Radiation Database (NSRDB) [47] for the year 2019.

2.2. Beam shadow modeling approaches

Four approaches for modeling Beam shadows are developed: Forward Ray Tracing (RT), Focused Ray Tracing (FRT), Rasterization and Pixel-Counting (Pix), and Analytical Areas (Ana). Figs. 3–6 illustrate the main steps of how these approaches work. A more detailed description of these steps, including pseudo-code, is presented in Supplementary Sec. 2. They are programmed in Python 3.0 using functions from the *Numpy* and *Random* libraries for the ray-based and analytical approaches, as well as *Matplotlib* and the *io* module of *Skimage* library for rasterization and pixel-counting. The code for these algorithms is openly available in the Supplementary Material.

The goal of these approaches is to find the Beam Unshaded Factor (USF_B) of the PV objects, namely, the proportion of PV surface that is visible (unobstructed), divided by the total PV surface from Beam's POV. For example, 20% USF_B means that 80% of the solar harvesting surface is blocked by shadows from a particular POV. Some studies refer to the Shadow or Shading Factor instead [11,25], which is $ShadowFactor = 1 - USF$. Shadow calculations are performed at the level of panels (Panel USF_B) and blocks (Block USF_B), i.e., groups of solar cells connected in series and protected by the same bypass diode [25]. The use of blocks and diodes is one of the main electrical hardware solutions to minimize risks in the event of hard partial shadows. In addition to calculating the proportion of a panel that is being visually obstructed (the geometrical aspect of shading), the number of shaded blocks is identified, as this has electrical implications that can be approximated in the modeling framework [48] (Section 2.4).

The four approaches begin by projecting the objects in the scene (PVs and Bs) onto a POV plane, which is normal to the sun vector (solar azimuth and elevation) (see Supplementary Fig. 1). This way, the problem becomes 2-dimensional, and the depth distances from the POV plane to the objects are used to establish what objects may cast shadows over others. For instance, a building block can only cast a shadow over a PV if it is "in front" from the perspective of incoming rays.

As illustrated in Fig. 3(A), the characterization of Beam shadows can be carried out using ray-tracing techniques. RT simulates Beam radiation as a set of parallel rays that stochastically sample the area of interest. Whenever a projected ray is inside a polygon, the unobstructed ray would go through that object. If this object is a PV and it is the first object to make contact with the ray, then the ray is illuminating the harvesting surface (Actual Hit). Conversely, if this is not the first object to make contact with the ray, then the ray would hit that PV if unobstructed (Ref Hit), but it is actually blocked by a shadow. The Unshaded Factors are the total number of Actual Hits over the number of Ref Hits. The resolution variable is the *mean number of rays per PV block area*, which is related to the ray density. To enforce this ratio, the total number of rays ($NRays$) must increase proportionally to the area of interest. To improve efficiency, if a ray is too far from any of the PVs (beyond their radius of influence), the algorithm moves on to the next ray. Similarly, when the intersections of a ray are being analyzed, only the objects that are close enough are considered. The radii of influence of both PVs and Bs are half their maximum diagonals. Also, the ray

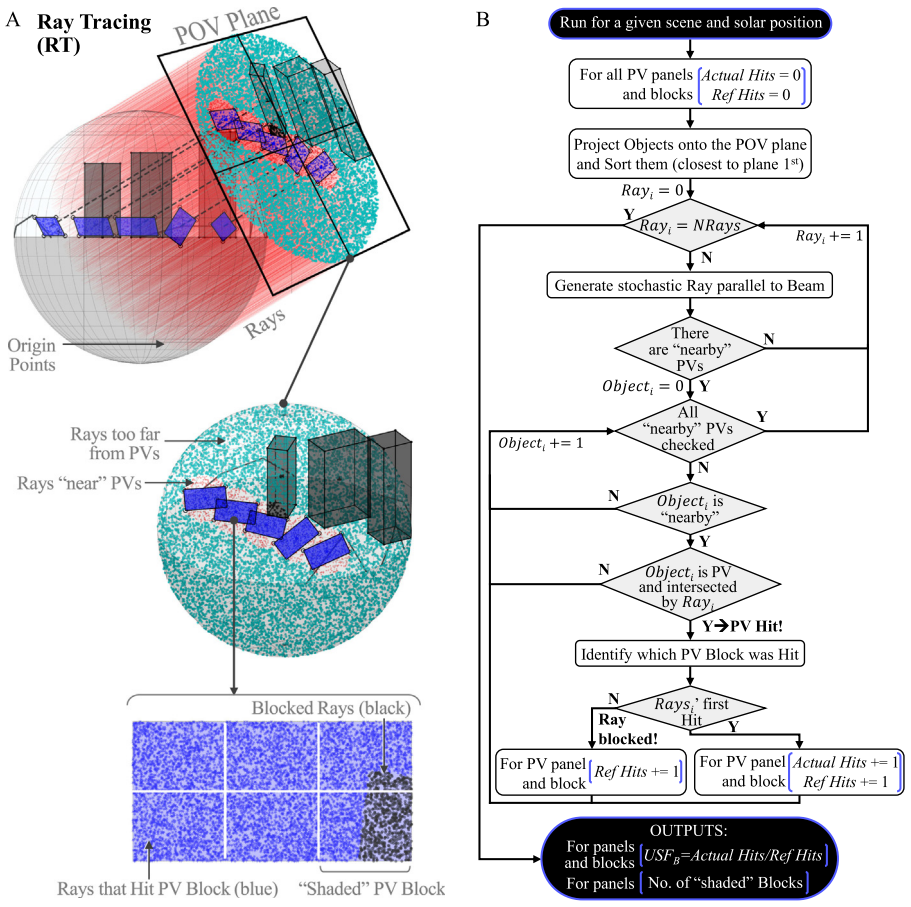

Fig. 3. Forward Ray Tracing approach for modeling Beam shadows. (A) Illustrated example of how the approach works and (B) Flowchart of the main steps.

density (the number of rays per unit volume) is constant regardless of the POV direction because the plane where the rays are originated is normal to the POV.

FRT also uses parallel rays (Fig. 4), but this approach does not sample the whole area of interest. Instead, the rays are generated directly from or toward the PV surfaces (similar to Backward ray-tracing [33]) to avoid processing rays that do not interact with PVs. It is established by design that all unobstructed rays intersect a PV surface (Ref Hits). Thus, it is not necessary to process all intersections between every ray and every projected polygon, so the algorithm stops after the first intersection is found. If this intersection is with the PV object that originated the ray, then this ray is illuminating the surface (Actual Hit). To reduce the computational cost further, FRT also only considers the objects that are close enough to a ray while looking for intersections. As with RT, shadow characterization improves with resolution, i.e., the number of rays per block ($NRaysBlock$) and panel ($NRaysPV$). But unlike RT, the total number of rays does not depend on the area of interest; only on the resolution and number of blocks. Also unlike RT, the ray density in FRT increases with the angle of incidence because a set of rays sample an increasingly narrower cross-sectional area. This may be a convenient aspect of this approach because shadow modeling accuracy is not necessarily hindered by large incidence angles, namely, when geometrical overlaps often appear smaller from POV and are therefore more challenging to characterize. Additionally, the effective ray density can be different for each PV surface, as they may all have different orientations.

Pix (Fig. 5) computes Beam shadows using rasterization and pixel-counting. The vector geometry of the scene is projected onto the POV plane by setting a virtual camera along the Beam rays. Then, every projected polygon is plotted in order (farthest from the sun first); similar to [21]. For this, a raster PNG image is generated using the *Matplotlib* library functionalities (in Python 3.0). Every PV block is assigned a particular color value and everything else (B objects and background) is plotted in black. When polygons are plotted, their vector-like geometric description becomes discretized and rasterized. This procedure is used to simulate the PV areas that are visible from POV, namely, the unshaded areas are approximated by counting the number of pixels in the raster image that correspond to PV block colors (Actual Pixels). The Unshaded Factors are the total number of Actual Pixels, divided by the number of PV block pixels that are visible in a Reference scene without shadows (Ref Pixels). The Reference scene is produced by ignoring the B objects, arranging the PVs in a grid where no shadows can be cast, setting the POV from the Zenith, and rotating the PVs so that the angles of incidence between their surface and the POV are preserved.

The resolution variable in Pix is the *mean number of pixels per PV panel side length*. Similar to RT, the total number of pixels must increase proportionally to the area of interest to enforce this ratio. There is a trade-off between pixel resolution and accuracy, and some studies have suggested hardware considerations to boost similar procedures, such as GPU computing and Z depth buffering [11,44]. The higher the pixel resolution is, the better the accuracy is, but also the more computationally costly the process becomes due to larger output images.

In the case of the Ana approach (Fig. 6), the aim is to fully characterize the projected shadows (from the POV perspective) so that the unshaded PV areas (Actual Areas) can be delimited analytically. A polygon overlap with an object that is "in front" implies that a shadow is cast. To find the "Actual Area" of a PV block, all shadow-casting polygons are identified and carefully merged (whenever possible) into isolated shadow envelopes. The Shoelace Formula [49] is used to calculate the non-overlapping areas of these shadow envelopes. The shadow envelopes are then merged to the block polygon, and the Actual

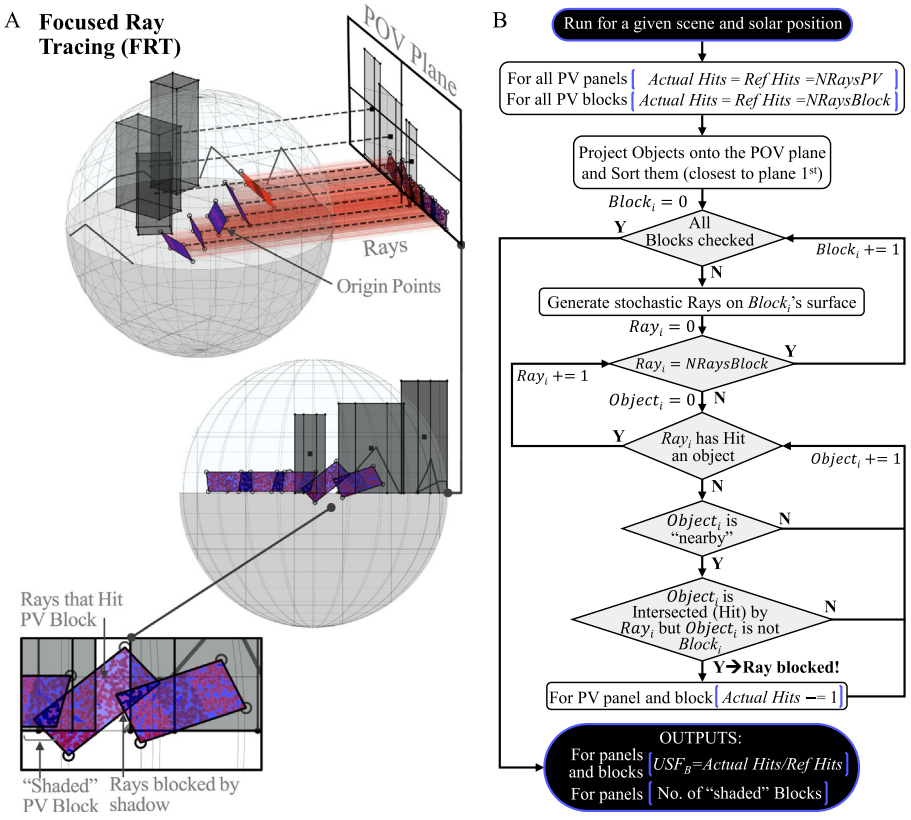

Fig. 4. Focused Ray Tracing approach for modeling Beam shadows. (A) Illustrated example of how the approach works and (B) Flowchart of the main steps.

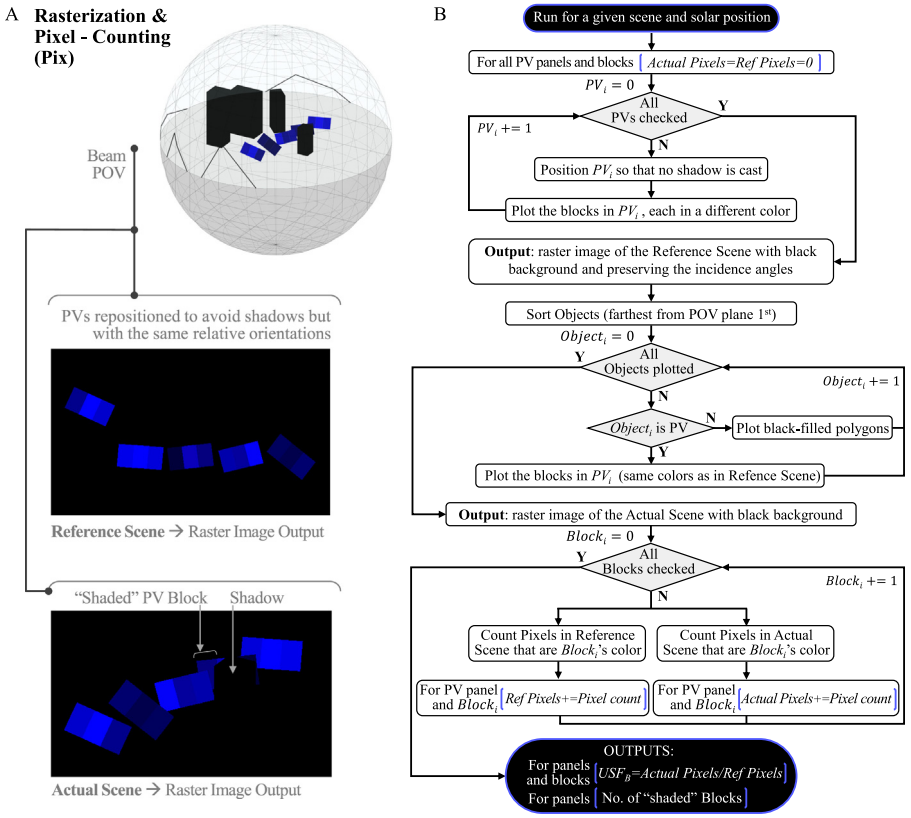

Fig. 5. Rasterization and Pixel-Counting approach for modeling Beam shadows. (A) Illustrated example of how the approach works and (B) Flowchart of the main steps.

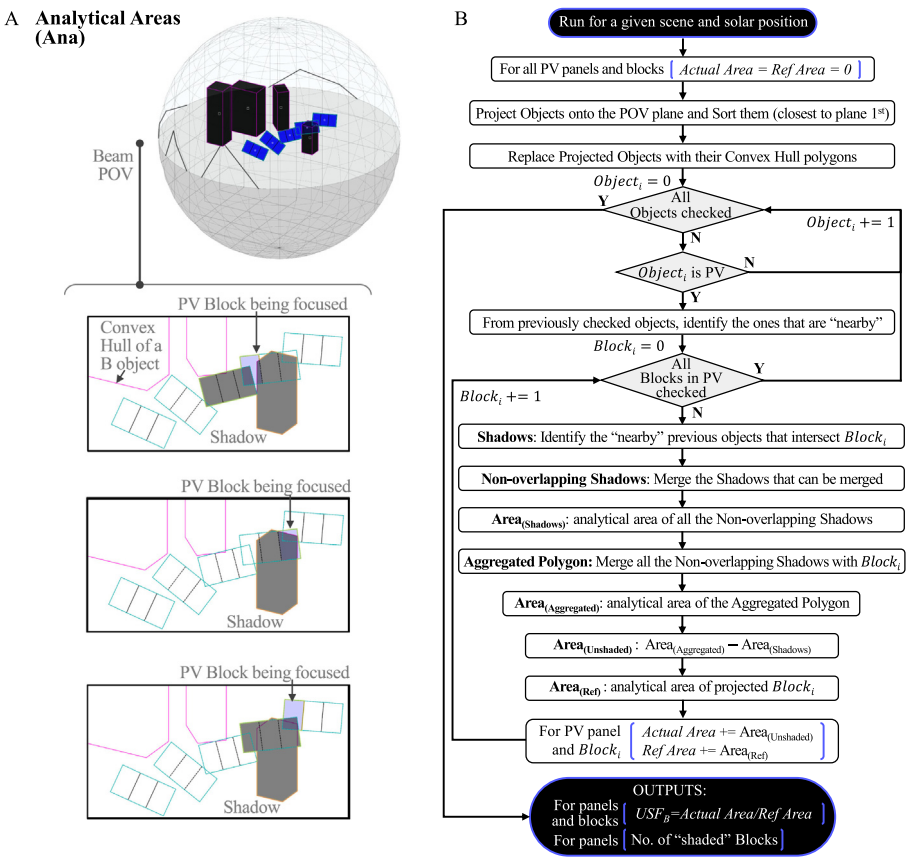

Fig. 6. Analytical Areas approach for modeling Beam shadows. (A) Illustrated example of how the approach works and (B) Flowchart of the main steps.

Area results from this aggregated area minus the area of the shadow envelopes alone. The reference area (Ref Area) results from applying the Shoelace formula to the projected vertices of the block. To limit the impact of potential errors, the Actual Area is constrained to positive values equal to or smaller than Ref Area.

Ana is considerably more complex than the other approaches, and the main challenge is tracing the relevant isolated contours from arbitrarily intricate overlapping geometries. To reduce complexity in Ana, B objects are represented by finding the Convex Hull [50] envelope of their projected volume as opposed to 6 individual polygon faces. The envelopes of B and PV objects can be assumed to be simple and convex because they are cuboids, but the overall shadows cannot be assumed to be convex because they result from interactions between several shadows. Additionally, rounding errors can lead to inconsistencies, so a uniform and fine-tuned rounding procedure is used for all coordinates and angular comparisons (e.g., checking whether two lines are parallel). For instance, if too many decimal places are used, then several definitions of the same point in space may be incorrectly treated as different points. Conversely, too few decimal places may hinder general accuracy or lead to merging points that should be distinct.

Figs. 3–6 describe the basic structure of the Beam shadow modeling approaches. However, additional considerations are required for the appropriate behavior of actual implementations, such as accounting for interactions with Horizon Shadows, dealing with division by zero and PV surfaces that are not facing the POV plane, or constraining USF_B between 0 and 1.

2.3. Comparative strategy

The shadow modeling approaches developed in this work aim to capture the same optical phenomena. However, how to make a one-to-one comparison considering their conceptual and practical differences? To address this, an exploratory and systematic comparison is carried out based on simulations, convergence criteria, and inferential statistics. The results are expected to be generalizable enough to be useful for theoretical and practical purposes in comparable scenarios. Namely, since the modeling performance trends characterized in this work are obtained from exploring different locations, incidences, panel orientations, and surroundings, then similar trends are expected in future applications with similar scene complexity.

A set of USF_B are established as the target values, referred to as "Naive Ground Truths" (GTs), in the five scenes and for two representative dates, i.e., March 20 (Equinox) and December 21 (Winter solstice), 2019. For this, Forward Ray Tracing (RT) is selected as Baseline because it is more directly related to how Beam radiation is conceptualized in Geometrical Optics [33], i.e., as parallel rays in the sun's direction filling a space. RT is run with progressively higher resolutions until the USF_B Standard Deviation (SD) falls well below 1% and performance reaches an apparent point of diminishing returns. The algorithm is run several times (10 in the case of Development scene) and GTs result from the averages among repetitions. These may be considered "naive" because absolute Ground Truths would require extensive experimental validations, yet these values are robust enough to serve as useful practical targets.

Resolutions are sampled in logarithmic progression to assess the convergence behavior of the shadow modeling approaches toward the target values. Approximate resolution ranges are then delimited and classified into Convergence Levels (CL) that allow for more direct comparisons between the shadow modeling approaches. The mean runtime is also analyzed as a function of resolution and scene complexity (i.e., Area of interest and number of PVs, Bs, and blocks per PV). Different aspects of scene complexity are relevant depending on the modeling approach. For instance, Area of interest is not defined for FRT and Ana. Time behavior is characterized in relative terms for the

findings to be potentially more generalizable, and all simulations are run on a Dell XPS 15-7590, 9th Generation Intel® Core™ i7-9750H processor, 32 GB RAM, and NVIDIA® GeForce® GTX 1650 4 GB video card.

Regression models are fitted to the relative results (using the *"scipy.optimize.curve_fit"* function in Python 3.0.) to capture run-time trends. For all linear models, the significance of the regression is tested using a Monte Carlo random permutations test of hypothesis [51, 52]. Monte Carlo methods can be powerful alternatives for testing inferential hypotheses without assuming underlying distributions [51]. Manly [53] suggests generating a minimum of 1,000 Monte Carlo simulations for a 5% significance level. In this work, 100,000 simulations are used in every Monte Carlo analysis to better approximate exact statistical tests of hypothesis [54]. Additionally, Monte Carlo analyses are used to test the significance of the differences in run-time and error distributions comparing the shadow modeling approaches at the same convergence levels.

2.4. Energy model

To have a meaningful interpretation of the results in terms of power output implications [W/m^2], an energy modeling framework is employed based on an expansion of the one described in detail by Arias-Rosales and LeDuc [55]. This energy model accounts for the date, time, geographic location, altitude, measured local weather and available solar resources, anisotropic sky model, transmittance losses, spectral response, heat transfer model, efficiency as a function of temperature, Beam Unshaded Factors, and some additional electrical considerations.

The Unshaded Factors for Beam radiation (USF_B) at the panel level and the Number of Shaded Blocks (NSB) can be obtained with the approaches described in this work. These variables are key components in the "Shading Effect Model", which accounts for the impact of hard partial shadows on the effective irradiance toward estimating power output. The Shading Effect Model is based on the experimentally-adjusted model proposed by Martínez-Moreno et al. [48], which considers the geometrical proportion of shaded area as well as an electrical factor equal to the number of blocks that are shaded (NSB) over the adjusted number of blocks per panel ($1 + NB$). In this work, a block is considered shaded whenever its unshaded area is 97% or less of its Reference area (i.e., the visible area ignoring shadows). As an example, this corresponds to shading half of the area of one solar cell in a panel with 60 cells and 3 blocks.

In this modeling framework, only Beam shadows are considered and USF_B is applied to both Beam and Circumsolar irradiances. Based on these assumptions, Is, $I\tau sg$, and $I\tau s$ (all in [W/m^2]) are defined in Eqs. (1)–(3). Is is the total irradiance that reaches the solar device unobstructed by shadows, while $I\tau s$ is the total effective irradiance that reaches the PV surface considering losses due to transmittance, shadows, and electrical mismatches. $I\tau sg$ is the same as $I\tau s$, but it does not include NSB considerations, i.e., shadow implications are only geometrical. The distinction between these three irradiances is important to differentiate the inputs and outputs in the heat transfer model below. In Eqs. (1)–(3), irradiances are denoted by (I), transmittances by (τ), and the subscripts correspond to the anisotropic components of solar radiation [56], i.e., Beam (B), Sky Isotropic (DI), Albedo (A), Circumsolar (DC), and Horizon Brightening (DH). Transmittances are calculated with the "Default Diffuse Transmittance Models" presented in [55].

$$Is = \left(I_B + I_{DC}\right) USF_B + I_{DI} + I_A + I_{DH} \tag{1}$$

$$I\tau sg = \left(I_B \cdot \tau_B + I_{DC} \cdot \tau_{DC}\right) USF_B + I_{DI} \cdot \tau_{DI} + I_A \cdot \tau_A + I_{DH} \cdot \tau_{DH} \tag{2}$$

$$I\tau s = \left(I_B \cdot \tau_B + I_{DC} \cdot \tau_{DC}\right) USF_B \left(1 - \frac{NSB}{1+NB}\right) + I_{DI} \cdot \tau_{DI} + I_A \cdot \tau_A + I_{DH} \cdot \tau_{DH} \tag{3}$$

As PV efficiency is known to be affected by temperature [57], a heat transfer model modified from [55] is used to estimate the temperature of the harvesting surfaces (T_{PV} in [Kelvin]). Eqs. (4) and (5) describe the energy balance for the PV material and the covering layer, respectively, with the inputs to the left of the equality and the outputs to the right. The input for the PV layer is the effective transmitted irradiance, and the outputs are the power output (P) and the heat conducted to the covering layer, where T_{cov} is the covering temperature (in [Kelvin]), C is the thermal conductivity (taken as 1.8 W/mK [58]), and L is the covering thickness (taken as 0.002 m [55]). The inputs for the covering layer are the absorbed irradiance and the heat conducted from the PV layer, and the outputs are heat losses due to convection and emission toward the sky and the ground. Here, α_T is the total absorptance, h is the convection coefficient, T_a is the ambient temperature, ϵ is the covering emissivity (0.9 for glass [59]), σ is the Stefan–Boltzmann constant, T_{sky} is the sky temperature, β is the inclination of the harvesting surface, and T_{gr} is the ground temperature (taken as $T_{gr} \approx T_a$).

$$I\tau sg = P + \frac{C}{L}\left(T_{PV} - T_{cov}\right) \tag{4}$$

$$Is \cdot \alpha_T + \frac{C}{L}\left(T_{PV} - T_{cov}\right) = h\left(T_{cov} - T_a\right) + \epsilon\sigma\left(T_{cov}^4 - T_{sky}^4\right)\left(\frac{1+\cos\beta}{2}\right) + \epsilon\sigma\left(T_{cov}^4 - T_{gr}^4\right)\left(\frac{1-\cos\beta}{2}\right) \tag{5}$$

The coupled two-equation system has two unknowns (T_{PV} and T_{cov}), which are calculated simultaneously for every time step with a fixed-point iteration method [60] and with the equations re-organized as Eqs. (6)–(7). The stopping criterion for this procedure is a maximum of 250 iterations or an energy balance threshold of 1 W/m^2, whichever occurs first.

$$T_{PV_{i+1}} = \frac{L}{C}\left[I\tau sg - I\tau s \cdot \eta_{PV_i}\frac{M}{M_{ref}}\right] + T_{cov_i} \tag{6}$$

$$T_{cov_{i+1}} = -\frac{L}{C}\left\{h\left(T_{cov_i} - T_a\right) + \epsilon\sigma\left[\left(T_{cov_i}^4 - T_{sky}^4\right)\left(\frac{1+\cos\beta}{2}\right) + \left(T_{cov_i}^4 - T_a^4\right)\left(\frac{1-\cos\beta}{2}\right)\right] - Is \cdot \alpha_T\right\} + T_{PV_{i+1}} \tag{7}$$

Lastly, power output (P in [W/m^2]) is a function of $I\tau s$ as well as the effective PV efficiency (η_{PV}) and the Spectral Response (M/M_{ref}) (Eq. (8)). For calculating η_{PV} (Eq. (9)), η_{ref} is the reference PV efficiency reported by the manufacturer at a temperature T_{ref}, and B is the PV temperature coefficient. Based on commonly used values for silicon crystalline modules, B is assumed to be 0.004, η_{ref} as 0.2, and T_{ref} as 298.15 Kelvin [57].

$$P = I\tau s \cdot \eta_{PV}\frac{M}{M_{ref}} \tag{8}$$

$$\eta_{PV} = \eta_{ref}\left[1 - B\left(T_{PV} - T_{ref}\right)\right] \tag{9}$$

The steps toward calculating power output are outlined in further detail in Supplementary Sec. 1. Aside from Eqs. (1)–(9), all the required equations and procedures are readily available in Section 2.3 of [55]. In particular, those procedures explain how to obtain the incident irradiances and transmittance factors for every anisotropic component of the radiation, as well as α_T, T_{sky}, h, and M/M_{ref}.

3. Results and discussion

3.1. Convergence: Accuracy and computational cost trade-off

The convergence behavior toward the Naive Ground Truths is characterized as a function of resolution for all the Beam shadow modeling approaches (Fig. 7). The main trends (curves) are explored using the Development scene with base complexity (Fig. 1(B)-left) and then compared to the other scenes (markers). As RT and FRT involve stochastic ray instantiations, they are run repeatedly, seeking more robust results (10 repetitions in the case of the Development scene). Root Mean Square Errors (RMSE) are used as proxies for accuracy versus GTs and Standard Deviations (SD) as proxies for precision (variation among repetitions).

There appears to be a distinct pattern of convergence for both RT and FRT, with RMSE and SD consistently improving with resolution

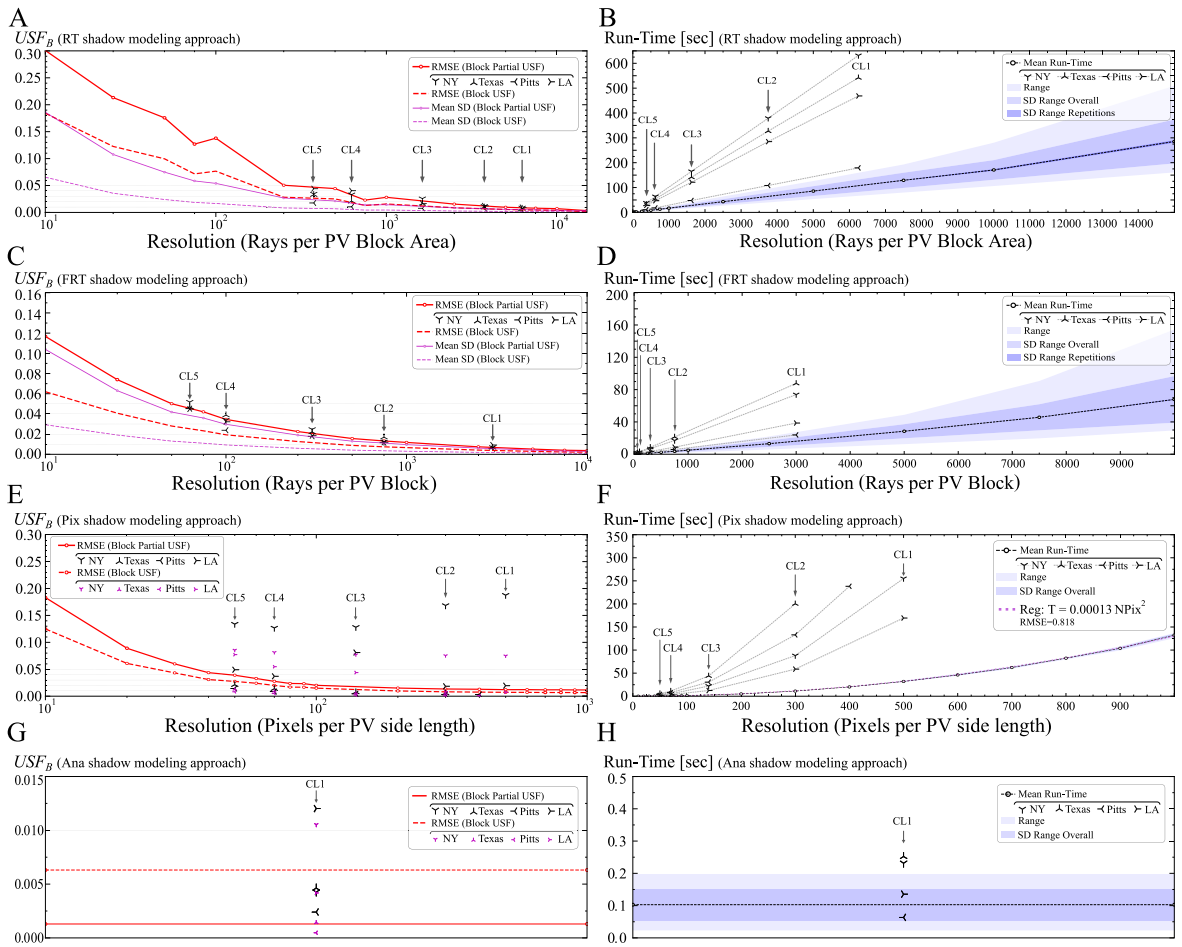

Fig. 7. Characterizing the trade-off between accuracy (RMSE), precision (SD), and computational cost (Run-Time). This figure suggests that the different approaches do converge, mostly agree with each other (within 1% deviations), and their shadow modeling errors tend to decrease with resolution but with increasing computational cost. RT is the Baseline (A–B) and Ana is the most efficient (G–H). Partial USF are between 0 and 1 (non-inclusive); the most challenging cases. Resolution ranges are classified into Convergence Levels (CL): 5%–1%. The Mean Run-Time, used as an experimental proxy for computational cost, is the mean time to run an algorithm on a given scene and time of the day (solar position). "SD Range Overall" accounts for the average variation in run-time caused by both the change in solar position and stochastic repetitions. "SD Range Repetitions" compares among repetitions, so it only accounts for the effect of stochasticity in the generation of rays.

(Fig. 7(A) & (C)). The five scenes exhibit a similar pattern, suggesting that the Convergence Levels (CL) are potentially generalizable to comparable environments. The Mean Run-Times (Fig. 7(B) & (D)) present linear relationships that are directly proportional to resolution and, in the case of RT, the higher the scene complexity is (e.g., in NY), then the steeper the slope is. FRT makes more effective and efficient use of the generated rays compared to RT, so it converges with fewer rays and takes less time to run the algorithm with the same number of rays per block. "SD Range Repetitions" is remarkably narrow in both approaches, suggesting that run-time is not considerably sensitive to how rays are stochastically instantiated.

The Pix approach is not stochastic, so there is no measure of precision with SD (Fig. 7(E)). Although there is an apparent pattern of convergence for the Development, Texas, and Pitts scenes, the results diverge for NY and LA. This seems to be due to a vulnerability that this approach has of plotting polygons in the wrong order when the distance of multiple objects to the POV plane is too similar. The main challenge, however, is representing polygon borders as accurately as possible. As details become smaller and more intricate, such as when a PV looks very narrow from the side, border limitations start dominating how the PV areas are characterized. This is mitigated with better resolutions but only up to a point. Higher resolutions imply larger image outputs, and there is a limit to the image size that is effectively dealt with using the conventional functionalities of *Matplotlib* (for Python 3.0). Because of this, it is not possible to test the Texas and Pitts scenes at CL 1% as those are the scenes with the largest outer areas. The run-times (Fig. 7(F)) follow the decreasing rank order of outer areas, i.e., Texas-Pitts-NY-LA-Development. An exponential regression model is fitted to capture this behavior, in which the pixel resolution has an approximately quadratic effect on mean run-time (i.e., $NPix^2$). Also, the SD range is notably narrow compared to RT and FRT, suggesting that run-time in Pix is barely sensitive to the solar position.

In the case of Ana, which does not depend on any resolution, most error distributions (Fig. 7(G)) are concentrated below the 1% threshold. Ana tends to overestimate the illuminated PV area when the merged shadow has holes (see Supplementary Fig. 5), but this issue appears to be infrequent (once in Development, NY, and LA) and has a marginal effect on performance. Additionally, Ana is sensitive to rounding errors, but uniform and fine-tuned rounding has been carried out to overcome this. The Mean Run-Times (Fig. 7(H)) are orders of magnitude smaller than the other approaches, especially toward CL1%.

Based on the main RMSE and SD trends, approximate resolution ranges are classified into Convergence Levels. Every CL is represented by a characteristic resolution value within the range or at the lower boundary of the range (Table 1). For example, CL3% means that both precision (SD) and accuracy (RMSE) are expected to have average deviations below 3%. This allows comparing the different modeling approaches at the same level in a more direct way despite differences in conceptualization, underlying assumptions, and the type of resolution. From Table 1, Ana is markedly more efficient when compared to the

Table 1
Comparing the Beam shadow modeling approaches at every Convergence Level. General format: [low range boundary–high range boundary] characteristic value. Resolution: [approximate CL thresholds] representative resolution. RMSE: [corresponding to the lower and upper resolution boundaries] mean between boundaries; in units of USF. Run-Times: [from the Min of the lower end to the Max of the higher end] mean between boundaries. Relative Run-Times: characteristic run-times relative to Ana's run-times. For example, in situations comparable to the ones tested, the RT approach is expected to converge to Naive Ground Truth values within 5% regarding precision and accuracy if it is run with a resolution of 250–500 Mean Rays per PV block Area and this takes approximately 2.59-13.39 s to run, which is 62.3 times longer compared to Ana.

CL	Metric	Ray Tracing	Focused Ray Tracing	Pixel-Counting	Analytical Areas
5%	Resolution Unit	Mean Rays/Block A	Rays/Block	Mean Pixels/PV side	–
	Resolution	**[250-500] 375**	**[50-75] 63**	**[40-60] 50**	–
	Mean Run-Time [s]	[2.59-13.39] 6.41	[0.08-0.64] 0.31	[0.28-0.62] 0.44	[0.02-0.2] 0.1
	Relative Run-Time	**62.3**	**3.0**	**4.2**	**1**
	$RMSE_{(Block)}$	[0.025-0.028] 0.0267	[0.023-0.028] 0.0256	[0.024-0.031] 0.0275	0.0063
	$RMSE_{(BlockPartial)}$	[0.044-0.05] 0.0474	[0.042-0.05] 0.0461	[0.033-0.044] 0.0383	0.0013
	$RMSE_{(Panel)}$	[0.011-0.015] 0.0131	[0.013-0.016] 0.0143	[0.018-0.023] 0.0207	0.0036
	$RMSE_{(PanelPartial)}$	[0.017-0.022] 0.0194	[0.02-0.023] 0.0213	[0.022-0.03] 0.0262	0.0054
4%	**Resolution**	**[500-750] 625**	**[75-100] 100**	**[60-80] 70**	–
	Mean Run-Time [s]	[5.32-20.02] 10.7	[0.12-0.91] 0.43	[0.54-0.99] 0.73	[0.02-0.2] 0.1
	Relative Run-Time	**104.0**	**4.2**	**7.1**	**1**
	$RMSE_{(Block)}$	[0.013-0.025] 0.019	[0.019-0.023] 0.0213	[0.017-0.024] 0.0205	0.0063
	$RMSE_{(BlockPartial)}$	[0.023-0.044] 0.0336	[0.035-0.042] 0.0383	[0.024-0.033] 0.0282	0.0013
	$RMSE_{(Panel)}$	[0.009-0.011] 0.0098	[0.011-0.013] 0.0122	[0.014-0.018] 0.016	0.0036
	$RMSE_{(PanelPartial)}$	[0.013-0.017] 0.0146	[0.017-0.02] 0.0182	[0.018-0.022] 0.02	0.0054
3%	**Resolution**	**[750-2500] 1625**	**[100-500] 300**	**[80-200] 140**	–
	Mean Run-Time [s]	[8-63.9] 27.98	[0.17-4.13] 1.47	[0.87-5.17] 2.94	[0.02-0.2] 0.1
	Relative Run-Time	**272.0**	**14.3**	**28.6**	**1**
	$RMSE_{(Block)}$	[0.009-0.016] 0.0122	[0.009-0.019] 0.0141	[0.01-0.017] 0.0133	0.0063
	$RMSE_{(BlockPartial)}$	[0.015-0.028] 0.0218	[0.016-0.035] 0.0251	[0.015-0.024] 0.0194	0.0013
	$RMSE_{(Panel)}$	[0.005-0.011] 0.0084	[0.005-0.011] 0.0082	[0.007-0.014] 0.0104	0.0036
	$RMSE_{(PanelPartial)}$	[0.008-0.017] 0.0124	[0.007-0.017] 0.0121	[0.01-0.018] 0.0137	0.0054
2%	**Resolution**	**[2500-5000] 3750**	**[500-1000] 750**	**[200-400] 300**	–
	Mean Run-Time [s]	[26.99-126.1] 64.42	[0.89-8.15] 3.73	[4.88-21.03] 12.55	[0.02-0.2] 0.1
	Relative Run-Time	**626.3**	**36.2**	**122.0**	**1**
	$RMSE_{(Block)}$	[0.005-0.009] 0.0071	[0.007-0.009] 0.0077	[0.008-0.01] 0.0087	0.0063
	$RMSE_{(BlockPartial)}$	[0.01-0.015] 0.0125	[0.012-0.016] 0.0137	[0.013-0.015] 0.014	0.0013
	$RMSE_{(Panel)}$	[0.003-0.005] 0.0044	[0.004-0.005] 0.0045	[0.005-0.007] 0.0059	0.0036
	$RMSE_{(PanelPartial)}$	[0.005-0.008] 0.0065	[0.006-0.007] 0.0066	[0.007-0.01] 0.0085	0.0054
1%	**Resolution**	**[5000-7500] 6250**	**[1000-5000] 3000**	**[400-600] 500**	–
	Mean Run-Time [s]	[53.59-192.71] 107.44	[1.97-48.77] 16.73	[19.73-48.97] 33.12	[0.02-0.2] 0.1
	Relative Run-Time	**1044.5**	**162.6**	**321.9**	**1**
	$RMSE_{(Block)}$	[0.004-0.005] 0.0049	[0.003-0.007] 0.0048	[0.007-0.008] 0.0073	0.0063
	$RMSE_{(BlockPartial)}$	[0.008-0.01] 0.0087	[0.005-0.012] 0.0085	[0.012-0.013] 0.0125	0.0013
	$RMSE_{(Panel)}$	[0.0033-0.0035] 0.0034	[0.002-0.004] 0.0028	[0.004-0.005] 0.0047	0.0036
	$RMSE_{(PanelPartial)}$	[0.0048-0.0051] 0.005	[0.003-0.006] 0.0042	[0.006-0.007] 0.0068	0.0054

others. FRT at CL1% takes approximately 160 times more time to run on average (two orders of magnitude), Pix takes 320 times more, and RT more than a thousand times (three orders of magnitude). Focusing rays with FRT is also notably effective at improving efficiency, as RT takes from 6.4 more times to run at CL1% to 25 more at CL4%. Also, FRT takes approximately half the time compared to Pix across CLs. The results suggest that FRT is competitive in terms of reliability, performance, and simplicity of implementation, while Ana is exceptionally accurate and the most computationally efficient. The most appropriate approach and CL depend on the particular scene and modeling purpose, so this table may be used as a comparative tool to aid selection, where modeling accuracy can be weighted along with computational cost (Relative Run-Time). For instance, the optimization of the positions of a solar array may begin with fast calculations at CL5% and end with detailed evaluations at CL1%.

The Error distributions are compared in Fig. 8. These distributions show very similar patterns of convergence at the block and panel levels. USF_B deviations tend to be considerably less substantial at the panel level because larger areas are usually easier to characterize and because some of the errors at the block level cancel out at the panel level. From their general shape and medians, RT and FRT deviations are symmetrical, so there does not seem to be systematic under or over-estimation. Ana also appears to be symmetrical at the Box-plot level, although one extreme case visually affects symmetry in the Violin depiction. Conversely, Pix distributions seem to be negatively biased (i.e., under-estimation). The rank order of shadow modeling efficiency is illustrated in terms of run-time distributions (Fig. 8(C)); in decreasing run-time efficiency: Ana (barely visible in comparison), FRT, Pix, and RT.

The distributions and patterns of convergence are compared further regarding the statistical significance of their differences (Fig. 9). By comparing distributions from two groups at a time, these P-values estimate (by randomly reassigning the groups) how likely it is to obtain differences at least as large as the ones observed due to randomness alone. The smaller the P-values are, the more significant the differences observed are. For instance, comparing the Panel RMSE at CL5% (Fig. 9(A)-left), all approaches are significantly different (P-value < 3.2%) except for RT-FRT, while there are no significant differences in the Panel RMSE distributions at CL2% (P-value > 48%). This suggests that, as the approaches become more convergent, the error distributions become more indistinguishable in terms of Panel RMSE. On the other hand, Mean run-time distributions present highly significant differences (most with P-value ≈ 0%) in pairwise comparisons between approaches (Fig. 9(A)) and between Convergence Levels (Fig. 9(B)).

Fig. 9(B) supports the assessment of how significant the RMSE improvements are as a function of Convergence Level. These improvements are all highly significant in the case of Pix, which is expected because the error distributions (Fig. 8) visibly become more compact and less biased (Medians closer to 0). The majority of Panel RMSE improvements are significant for FRT, especially comparing CL2% to worse CLs (P-value < 0.13%). However, FRT at CL1% does not achieve significant RMSE improvements at the panel level (P-value = 13.5%),

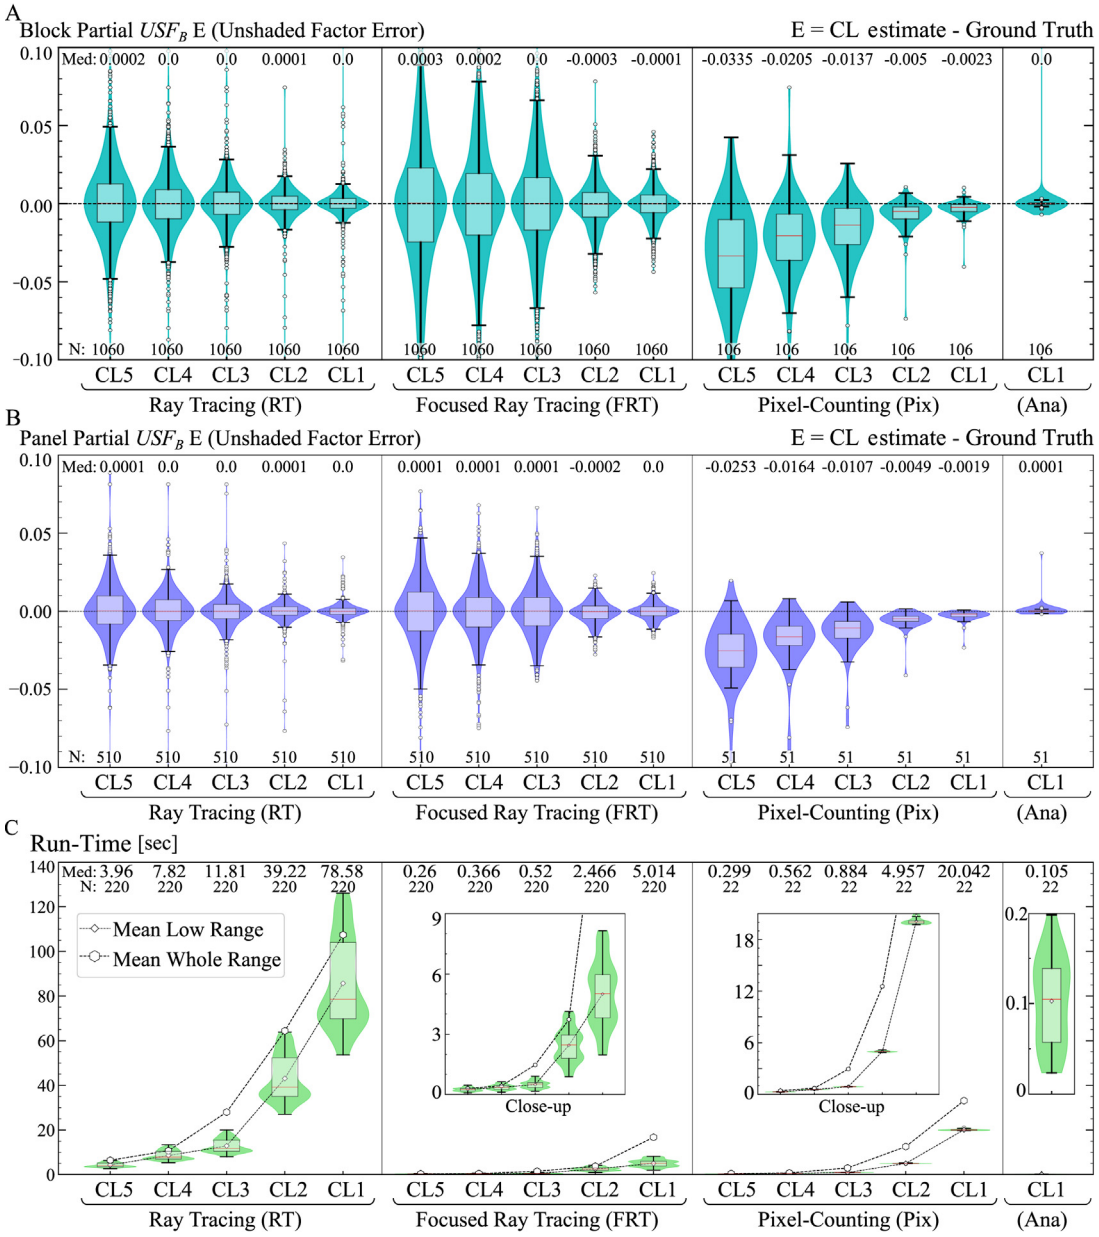

Fig. 8. Error and run-time distributions for the four Beam shadow modeling approaches. Errors with respect to GTs are presented at the different Convergence Levels and at the block level (A), panel level (B), and with the corresponding run-times (C). Only partial shadows are included. Scene: Development; Resolutions: low-range of the CLs; Dates: March 20 (Equinox) and December 21 (Winter solstice), 2019; Repetitions: 10 in the stochastic approaches, i.e., RT and FRT.

which can be useful information when choosing between CL1% or CL2% resolutions. These are examples of how these significance heat maps can inform decision-making in shadow modeling.

3.2. Run-time and scene complexity

Mean run-time is also a function of scene complexity, which should be considered when choosing among approaches for a particular environment. Relative run-time is analyzed in Fig. 10 for the four Beam shadow modeling approaches. Regression models (Reg) are used for comparing the strength of the effect that the different complexity metrics have on run-time. They may also be used for rough run-time estimations for different scenes based on a reference scene. As an example of how these graphs can be used, Fig. 10(B) illustrates how the mean run-time has a linear relationship with the number of PV objects, and the trend is consistent across the five scenes. If there is a known scene available as a reference (e.g., in this case having 5 PVs and taking 51 s to run the RT algorithm), then it is possible to estimate how much more (or less) time it would take to run the algorithm in the same scene (or similar) but with a different number of PVs. This is accomplished by using the regression models displayed in the figure, which are stated in relative terms, and using the SD ranges and RMSE values as an indication of the expected uncertainty or variation from regression estimates. From Fig. 10, the most relevant metrics for each approach are, in decreasing order:

- **For RT**: relative No. of PV panels (rNPV), relative Ray Area (rRA), and relative No. of Building cuboids (rNB); all linear trends. Run-time is expected to increase by 60% if the number of panels is doubled (rNPV = 2), by 40% if rRA = 2, and by 12% if rNB = 2.
- **For FRT**: rNPV, relative No. of PV Blocks per panel (rNPVB), and rNB; all approximately linear trends. Run-time is expected to increase by 121% if rNPV = 2, by 107% if rNPVB = 2, and by 23% if rNB = 2.

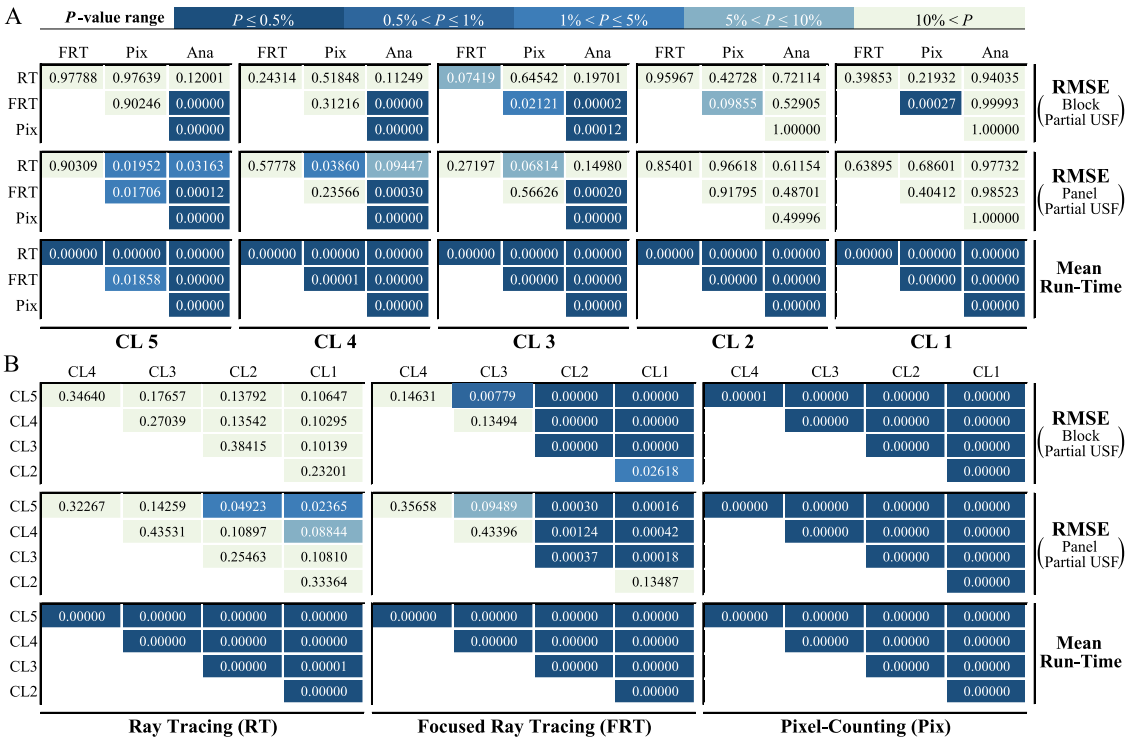

Fig. 9. Testing the significance of the pairwise differences between RMSE and Mean run-times of the Beam shadow modeling approaches. Comparing differences between approaches within Convergence Levels (A), as well as between Convergence Levels within approaches (B). All *P*-values are obtained with Monte Carlo analyses, testing one pair of distributions at a time (same distributions as in Fig. 8) with 100,000 randomized group re-assignments.

The relative run-time responses in FRT are almost twice as strong as in RT.

- **For Pix**: rRA, rNPVB-rNPV, and rNB; all linear trends. Run-time is expected to increase by 95% if rRA = 2 (more than twice as much as in RT), by 27%–44% if rNPVB = 2 (less than half the response in FRT), by 28% if rNPV = 2 (less than half the response in RT and FRT), and by 0.2% if rNB = 2 (no meaningful effect). The response as a function of rNPVB becomes steeper with the number of panels (5PV-10PV).
- **For Ana**: rNPV (almost one-to-one correspondence), rNPVB, and rNB; all linear trends. Run-time is expected to increase by 105% if rNPV = 2, by 72% if rNPVB = 2, and by 58% if rNB = 2.

Relative run-times in the testing cases seem to follow similar trends as the ones characterized in the Development scene, suggesting that these trends may be approximately generalizable to comparable environments. This is also supported by how narrow and stable the "PV iter" and "B iter" ranges are. This implies that the particular PV and B objects that are randomly iterated are not usually as consequential in terms of run-time as the relative number of such objects.

3.3. Energy estimation implications

Many solar harvesting applications that require the use of shadow modeling approaches will ultimately need to perform power or energy estimations, treating shadows as potential energy losses. This section explores some of the energy implications of this challenge by addressing the following questions: How relevant is Beam shadow modeling in terms of power output estimation in the context of intricate environments and PVs with freely defined positions and orientations? (Fig. 11), How do error distributions and convergence in Beam shadow modeling relate to error distributions in power estimates? (Fig. 12), and how does the threshold for considering a PV block as a "shaded Block" (SB) affect power estimation? (Fig. 13).

Generating variation in the simulation conditions is important for improving how generalizable the explored trends are expected to be. Sources of variation here include the use of different locations with distinct environments, as well as harvesting surfaces (PVs) with different positions and orientations. Every simulated time step implies a shift in the angle of incidence between Beam rays and the PVs, which introduces changes in shadow events. As an additional source of variation and to further explore the implications of dynamic and freely defined harvesting surfaces, the PV orientations (azimuth and inclination) are "nudged" every time step, i.e., the initial PV angular coordinates are shifted. Each angular nudge is randomly generated in a −45° to +45° range. This nudging is akin to PV devices dynamically and individually adjusting their orientations throughout the day. The energy model used to simulate the power outputs of these harvesting surfaces is described in Section 2.4 and Supplementary Sec. 1. As part of the modeling framework, Beam shadows are modeled using the FRT approach, which in all tested cases converged (Fig. 7(C)) to approximately the same values as RT (the Baseline approach) but with markedly lower run-time (from 6 to 20 times faster, Table 1). The power "Naive Ground Truths" are calculated using USF_B "Naive Ground Truths" (as in Section 2.3 but using FRT instead of RT). Seeking to test the worst-case scenarios concerning resolution, low-range resolutions (Table 1) are used to get additional USF_B and power distributions for every convergence level.

Fig. 11(A) illustrates how the USF_B weighted averages change throughout the day in the five scenes. Mornings and afternoons tend to be the moments with the most pronounced shadows because Beam rays are closer to the ground and, thus, more likely to be blocked by obstacles before reaching the PVs. The red ranges (Fig. 11(B)) between the transmitted irradiances with and without shadows illustrate the potential energy losses due to Beam shadows. Interestingly, not all the regions where USF_B is low (i.e., considerable shadows) coincide with wide energy loss ranges. This is because some of the cases with pronounced shadows are from PVs that would not be harvesting a lot of solar radiation anyway, so in those cases, the shadows do not have a

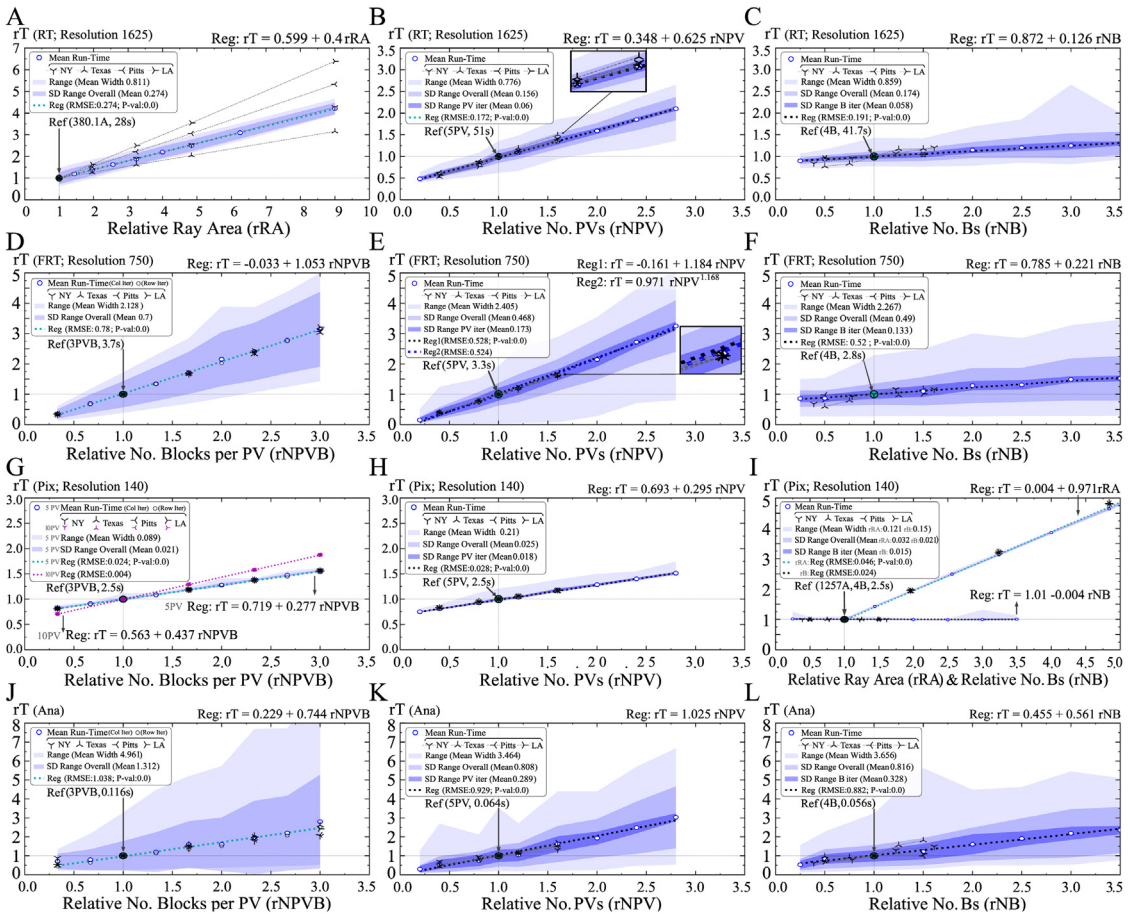

Fig. 10. Relative run-time as a function of scene complexity based on the variables that are relevant to every approach. This figure characterizes the trends of how computational cost increases as the scene complexity increases (Area of interest and the Number of PVs, Bs, and blocks per PV). RT approach in (A–C), FRT in (D–F), Pix in (G–I), and Ana in (J–L). Trends are characterized in the Development scene, captured with regression models, and compared to other scenes. Run-times and complexity metrics are analyzed in relative terms (with reference points shown as Ref) to look for potentially generalizable trends. "PV iter" and "B iter" ranges account for the effect of randomly choosing a number of different objects from a larger set of objects (i.e., PV panels and Building cuboids). "Overall" ranges additionally account for changes in solar position. No. of blocks (D, G & J) is explored by iterating the number of block columns (Col Iter) and rows (Row Iter) per panel.

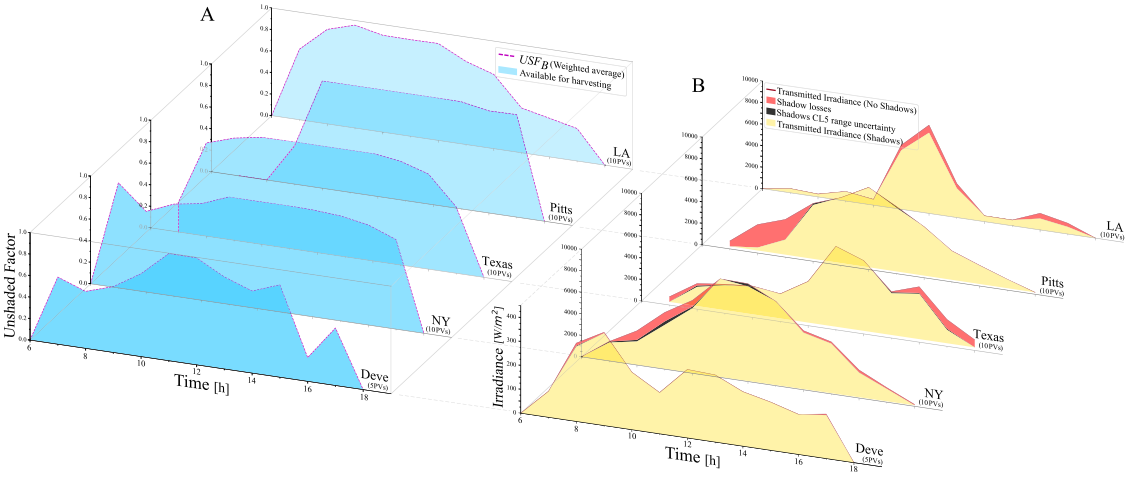

Fig. 11. The effect of Beam shadows on the irradiance that is available for energy harvesting. The PVs in every scene are treated as part of the same system, so the USF_B curves (A) are calculated as the weighted averages among the PVs. The weights are the irradiances (B) that would be transmitted to the PVs without shadows. (B) Illustrates the potential energy losses and uncertainties due to shadow modeling. The variations in these curves (A–B) result from the movement of the sun, the PVs, and their shadow interactions with the environment.

considerable impact on the aggregate. The black ranges (Fig. 11(B)) illustrate the uncertainty in shadow energy losses due to challenges with convergence in shadow modeling. Supplementary Figure 6 explores in further detail how these power estimation uncertainties tend

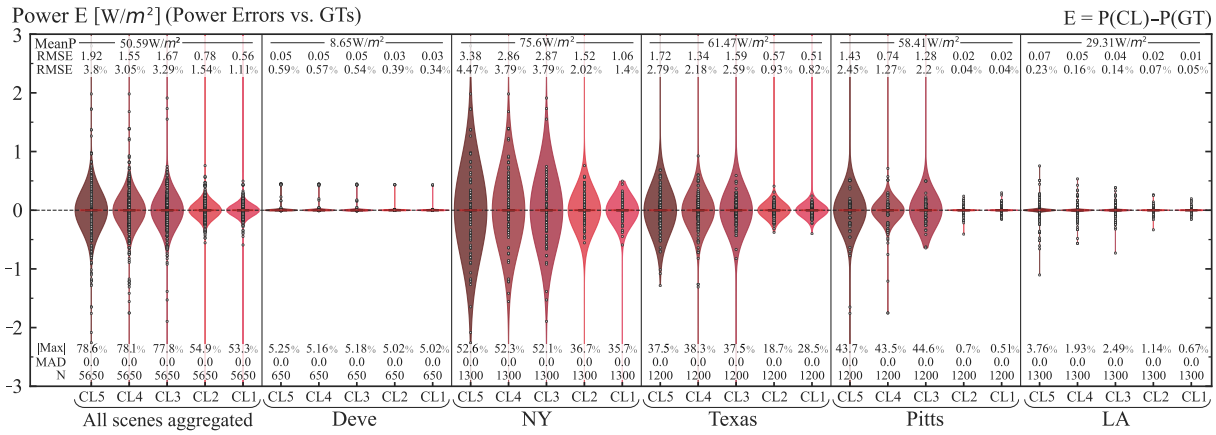

Fig. 12. Violins and box-plots of how power estimation error distributions are considerably reduced as Beam shadow modeling converges from CL5 to CL1. The results are aggregated from all simulated times, PVs, and stochastic repetitions in the 5 scenes. Power E (Errors) compare the power estimates that use USF_B Ground Truths versus power estimates that use USF_B distributions at different convergence levels. Mean Power (MeanP), RMSE (absolute and relative to MeanP), maximum relative error ($|Max|$), MAD (Median Absolute Deviations), and N (No. instances) are displayed.

to occur with FRT at the level of each harvesting surface, which appear as sudden but sparse deviations that, for the most part, disappear or diminish from CL5 to CL1. Nonetheless, the uncertainty ranges at CL5 in Fig. 11(B) do not appear to be very significant compared to overall irradiance aggregates, which suggests that using the approaches with CL5 resolutions (which are faster to run; Fig. 8(C)) may still be useful for practical power estimations.

Fig. 12 illustrates how errors in shadow modeling estimation translate into errors in power output estimates, how these power errors are distributed, and how convergence in shadow modeling can meaningfully reduce uncertainty concerning error distributions. In the aggregate (left), RMSE drop from 3.8% at CL5 to 1.11% at CL1, and the maximum power errors drop from 78.6% at CL5 to 53.3% at CL1. This suggests that convergence in shadow modeling is an important aspect for accurate power estimations at individual cases, even though on average they may not seem very appreciable (Fig. 11(B)). Also, the relevance of shadow modeling convergence greatly varies among scenes, having a greater effect in NY, Texas, and Pitts, and less in Deve and LA. Although there are extreme cases, the bulk of the error distributions is notably compact, which is why the box-plots are barely visible (i.e., collapsed toward 0) while the violin distributions are mostly shaped by outliers (i.e., white dots). This agrees with the Median Absolute Deviations being 0 in all cases, implying that at least half of the shadow calculations fully agree with the GT target values (i.e., Power Error ≈ 0).

Another aspect that is influential in the relationship between shadow and energy errors is the threshold for considering a PV block as a "shaded Block" (SB) (see Section 2.4). This threshold accounts for additional electrical losses due to hard partial shadows that can go beyond the geometrical proportion of PV panel being blocked. For the previous calculations in this work, a block is considered shaded whenever its unshaded area is 97% or less of its reference area. Fig. 13(A) examines the implications of completely ignoring these electrical considerations by setting a 0% threshold (i.e., no SB) and comparing the resulting distributions to the power Ground Truths that use a 97% threshold. Compared to Fig. 12(left), these error distributions are notably larger (RMSE ≈ 8.7%) and more asymmetrical toward (+) errors, which indicates a tendency to over-estimate power.

Fig. 13(B) explores this relationship more by iterating threshold values in the whole range from 0 to 1. Power estimates decrease as the threshold increases, but changes in the low range seem to have negligible effects while they become substantial as the threshold approaches 90%. Some scenes like Deve and LA seem to be impervious to this effect, while NY presents considerable drops at the extreme of the range, i.e., by −28.6%, −28.9%, and −21.2% in terms of Power Mean, Median, and SD, respectively. The most appropriate threshold may depend on the panel-block-cell configurations being used and the electrical specificities of the actual devices, and this becomes more relevant the closer it gets to a 100% threshold. A 100% threshold means that any shadow (regardless of its size) can render the block as electrically "shaded".

4. Conclusions

Enabling solar technologies that are more diverse and dynamic regarding positioning and orientation can lead to better and more fine-tuned integration of solar harvesting in urban environments. A key challenge is improving modeling capabilities in terms of geometrical flexibility and efficiency because some common constraints and assumptions can be limiting when attempting to maximize the available solar resources in such intricate environments. In particular, developing practical modeling frameworks that account for detailed shadow effects in complex scenarios can be computationally challenging. It is not clear how the different available approaches compare directly in terms of accuracy, precision, efficiency, and practicality, especially in the intricate urban context and with the position and orientation of every harvesting surface being freely and dynamically defined. In this work, four Beam shadow modeling approaches are developed and demonstrated in a range of urban landscapes. The novelty consists of the characterized trends and the direct comparison of these approaches, at different convergence levels, in the context of urban environments with varying complexity and freely defined solar harvesting devices. These characterizations involve implications in accuracy, precision, computational cost, ease of implementation, and energy estimation.

The Beam shadow modeling approaches converge below 1% deviations from target values and present considerable differences in computational efficiency. The fully analytical approach is, on average, three orders of magnitude faster than the forward ray-tracing approach, two orders of magnitude faster than the focused ray-tracing approach, and 320 times faster than rasterization and pixel-counting. Focused ray-tracing is competitive in terms of reliability, performance, and simplicity of implementation, while the analytical approach is the most efficient.

Critical trade-offs between accuracy and computational cost are characterized. Mean run-times present linear and superlinear trends that depend on scene complexity, i.e., the area of interest and the number of PV panels, building cuboids, and blocks per PV panel. The strength and relevance of the effect depend on the nature of the modeling approach. These trends underscore the importance of

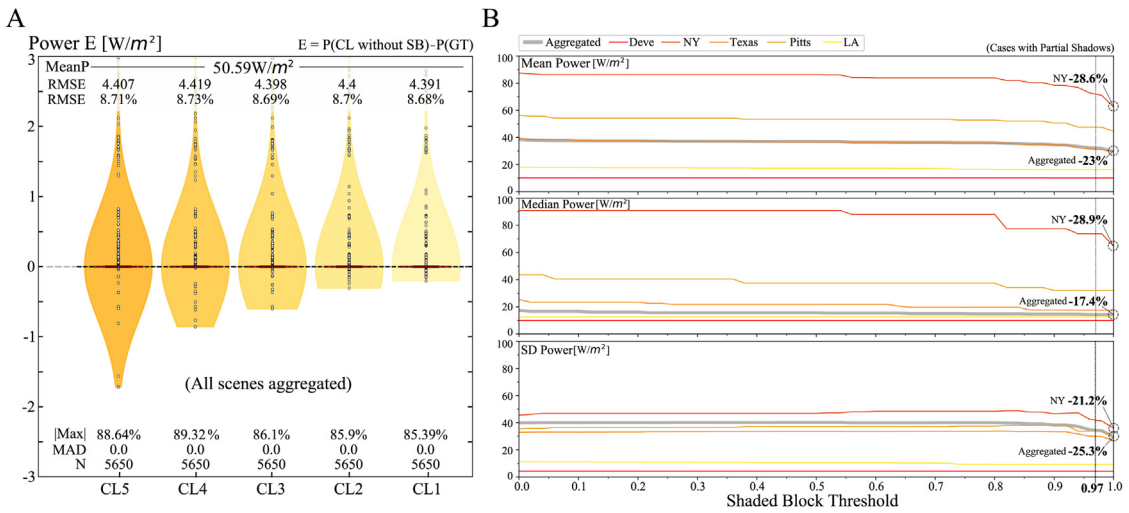

Fig. 13. Exploring the relationship between energy estimation errors and the threshold for considering a PV block as a "shaded Block" (SB). (A) Shows the Power Error distributions when this effect is completely ignored (0% threshold), while (B) iterates the threshold in the whole range and plots the trends in terms of means, medians, and Standard Deviations. Only cases with Partial shadows are included in (B) to focus on the most challenging shadow events.

considering the particular environment and modeling purpose when choosing among approaches and resolutions.

The results demonstrate how errors in shadow modeling estimation translate into errors in power output estimates, how these power errors are distributed, and how convergence in shadow modeling can meaningfully reduce uncertainty. The use of higher shadow modeling resolutions (at convergence levels from 5% to 1%), across all the analyses, leads to a drop in power estimation RMSE from 3.8% to 1.11%, and improvements in maximum power errors from 78.6% to 53.3%. Nonetheless, using the shadow modeling approaches at a 5% convergence level may have efficiency and practicality advantages while still leading to average power estimates that may be sufficiently accurate in the aggregate. The most appropriate approach and resolution depend on the particular scene and modeling purpose, where modeling accuracy and precision can be weighted along with computational cost. For instance, the optimization of the positions of a solar array may begin with fast calculations at a 5% convergence level and end with detailed evaluations at a 1% convergence level. Therefore, the characterizations, direct comparisons, and trends presented in this work can be a valuable guide to aid the selection and development of modeling frameworks.

The modeling capabilities presented are expected to be useful in the design, optimization, control, and forecasting of dynamic solar harvesting applications in intricate settings, such as cities and solar installations with shared land use. Developing accurate, efficient, and flexible modeling frameworks will play a role in enabling a diverse range of solar harvesting devices. From urban farming and devices that help people meet their energy needs, to solar-powered drones that fulfill a wide range of urban services like maintenance, security, carrying goods, or even transporting people.

CRediT authorship contribution statement

Andrés Arias-Rosales: Conceptualization, Methodology, Writing, Supervision, Analysis, Software, Data curation, Validation. **Philip R. LeDuc:** Conceptualization, Methodology, Writing, Supervision.

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.

Data and code availability statement

The main resulting datasets and code are available in the Supplementary Materials and described in Supplementary Secs. 4-5.

Acknowledgments

The authors would like to thank the Mechanical Engineering Department at Carnegie Mellon University, as well as the Colciencias-Fulbright Scholarship for Doctoral studies, cohort 2018 for their support. We would also like to thank the U.S. Department of Energy (DOE)/ NREL/ ALLIANCE for making The National Solar Radiation Database (NSRDB) available.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.rser.2022.112522. Supplementary Materials contain relevant data, code, and details on the shadow modeling approaches and the energy model.

References

[1] Louwen A, van Sark W. Chapter 5 - photovoltaic solar energy. In: Junginger M, Louwen A, editors. Technological learning in the transition to a low-carbon energy system. Academic Press; 2020, p. 65–86. http://dx.doi.org/10.1016/B978-0-12-818762-3.00005-4, URL https://www.sciencedirect.com/science/article/pii/B9780128187623000054.

[2] MIT. The future of solar energy-an interdisciplinary mit study. Cambridge, MA. USA: MIT; 2015.

[3] Yan J, Yang Y, Campana PE, He J. City-level analysis of subsidy-free solar photovoltaic electricity price, profits and grid parity in China. Nat Energy 2019;4(8):709–17.

[4] Zhu R, You L, Santi P, Wong MS, Ratti C. Solar accessibility in developing cities: A case study in Kowloon east, Hong Kong. Sustainable Cities Soc 2019;51:101738. http://dx.doi.org/10.1016/j.scs.2019.101738, URL http://www.sciencedirect.com/science/article/pii/S2210670719313046.

[5] Freitas S, Catita C, Redweik P, Brito MC. Modelling solar potential in the urban environment: State-of-the-art review. Renew Sustain Energy Rev 2015;41:915–31.

[6] Melo EG, Almeida MP, Zilles R, Grimoni JA. Using a shading matrix to estimate the shading factor and the irradiation in a three-dimensional model of a receiving surface in an urban environment. Sol Energy 2013;92:15–25.

[7] Erdélyi R, Wang Y, Guo W, Hanna E, Colantuono G. Three-dimensional solar radiation model (SORAM) and its application to 3-D urban planning. Sol Energy 2014;101:63–73.

[8] Broto VC, Stevens L, Ackom E, Tomei J, Parikh P, Bisaga I, et al. A research agenda for a people-centred approach to energy access in the urbanizing global south. Nat Energy 2017;2(10):776–9.

[9] khasraw Bani R, Jalal SJ. Impact of shadow distribution on optimizing insolation exposure of roofs according to harness or transfer of solar energy in sulaimani city, Iraq. Renew Energy 2019;136:452–62. http://dx.doi.org/10.1016/j.renene.2019.01.002, URL http://www.sciencedirect.com/science/article/pii/S0960148119300023.

[10] Oh M, Park H-D. A new algorithm using a pyramid dataset for calculating shadowing in solar potential mapping. Renew Energy 2018;126:465–74.

[11] Robledo J, Leloux J, Lorenzo E, Gueymard CA. From video games to solar energy: 3D shading simulation for PV using GPU. Sol Energy 2019;193:962–80.

[12] Rodrigues E, Fernandes MS, Gomes A, Gaspar AR, Costa JJ. Performance-based design of multi-story buildings for a sustainable urban environment: A case study. Renew Sustain Energy Rev 2019;113:109243.

[13] Wu J, Wang H, Li N, Yao P, Huang Y, Yang H. Path planning for solar-powered UAV in urban environment. Neurocomputing 2018;275:2055–65. http://dx.doi.org/10.1016/j.neucom.2017.10.037, URL http://www.sciencedirect.com/science/article/pii/S0925231217316880.

[14] Ballif C, Perret-Aebi L-E, Lufkin S, Rey E. Integrated thinking for photovoltaics in buildings. Nat Energy 2018;3(6):438–42.

[15] Yadav S, Panda S, Tripathy M. Performance of building integrated photovoltaic thermal system with PV module installed at optimum tilt angle and influenced by shadow. Renew Energy 2018;127:11–23. http://dx.doi.org/10.1016/j.renene.2018.04.030, URL http://www.sciencedirect.com/science/article/pii/S0960148118304373.

[16] Tripathy M, Sadhu P, Panda S. A critical review on building integrated photovoltaic products and their applications. Renew Sustain Energy Rev 2016;61:451–65.

[17] Svetozarevic B, Begle M, Jayathissa P, Caranovic S, Shepherd RF, Nagy Z, et al. Dynamic photovoltaic building envelopes for adaptive energy and comfort management. Nat Energy 2019;4(8):671–82.

[18] Collado FJ, Guallar J. A review of optimized design layouts for solar power tower plants with campo code. Renew Sustain Energy Rev 2013;20:142–54.

[19] Alsadi SY, Nassar YF. A general expression for the shadow geometry for fixed mode horizontal, step-like structure and inclined solar fields. Sol Energy 2019;181:53–69. http://dx.doi.org/10.1016/j.solener.2019.01.090, URL http://www.sciencedirect.com/science/article/pii/S0038092X19301021.

[20] Mamia I, Appelbaum J. Shadow analysis of wind turbines for dual use of land for combined wind and solar photovoltaic power generation. Renew Sustain Energy Rev 2016;55:713–8.

[21] Niewienda A, Heidt F. SOMBRERO: A PC-tool to calculate shadows on arbitrarily oriented surfaces. Sol Energy 1996;58(4):253–63. http://dx.doi.org/10.1016/S0038-092X(96)00088-6, Selected Proceedings of the ISES 1995: Solar World Congress. Part II, URL http://www.sciencedirect.com/science/article/pii/S0038092X96000886.

[22] Pedro MC-RM. Modelling of shading effects in photovoltaic optimization. (Ph.D. thesis), 2016.

[23] Díaz F, Montero H, Santana D, Montero G, Rodríguez E, Aguiar LM, et al. Improving shadows detection for solar radiation numerical models. Appl Math Comput 2018;319:71–85. http://dx.doi.org/10.1016/j.amc.2017.01.046, Recent Advances in Computing URL http://www.sciencedirect.com/science/article/pii/S0096300317300644.

[24] Fernández-Ahumada L, Ramírez-Faz J, López-Luque R, Varo-Martínez M, Moreno-García I, de la Torre FC. A novel backtracking approach for two-axis solar PV tracking plants. Renew Energy 2020;145:1214–21. http://dx.doi.org/10.1016/j.renene.2019.06.062, URL http://www.sciencedirect.com/science/article/pii/S0960148119308900.

[25] Fartaria TO, Pereira MC. Simulation and computation of shadow losses of direct normal, diffuse solar radiation and albedo in a photovoltaic field with multiple 2-axis trackers using ray tracing methods. Sol Energy 2013;91:93–101. http://dx.doi.org/10.1016/j.solener.2013.02.008, URL http://www.sciencedirect.com/science/article/pii/S0038092X13000650.

[26] Farhadi R, Taki M. The energy gain reduction due to shadow inside a flat-plate solar collector. Renew Energy 2020;147:730–40. http://dx.doi.org/10.1016/j.renene.2019.09.012, URL http://www.sciencedirect.com/science/article/pii/S0960148119313424.

[27] Hosseini SS, Ramiar A, Ranjbar AA. The effect of fins shadow on natural convection solar air heater. Int J Therm Sci 2019;142:280–94. http://dx.doi.org/10.1016/j.ijthermalsci.2019.04.015, URL http://www.sciencedirect.com/science/article/pii/S1290072919300808.

[28] Gardashov R, Eminov M, Kara G, Kara EGE, Mammadov T, Huseynova X. The optimum daily direction of solar panels in the highlands, derived by an analytical method. Renew Sustain Energy Rev 2020;120:109668.

[29] Cascone Y, Corrado V, Serra V. Development of a software tool for the evaluation of the shading factor under complex boundary conditions. In: Proceedings of building simulation 2011: 12th conference of international building performance simulation association, sydney. 2001, p. 14–6.

[30] Maghami MR, Hizam H, Gomes C, Radzi MA, Rezadad MI, Hajighorbani S. Power loss due to soiling on solar panel: A review. Renew Sustain Energy Rev 2016;59:1307–16.

[31] Bayrak F, Oztop HF. Effects of static and dynamic shading on thermodynamic and electrical performance for photovoltaic panels. Appl Therm Eng 2020;169:114900. http://dx.doi.org/10.1016/j.applthermaleng.2020.114900, URL https://www.sciencedirect.com/science/article/pii/S1359431119346095.

[32] Rezk H, Mazen A-O, Gomaa MR, Tolba MA, Fathy A, Abdelkareem MA, et al. A novel statistical performance evaluation of most modern optimization-based global MPPT techniques for partially shaded PV system. Renew Sustain Energy Rev 2019;115:109372.

[33] Jakica N. State-of-the-art review of solar design tools and methods for assessing daylighting and solar potential for building-integrated photovoltaics. Renew Sustain Energy Rev 2018;81:1296–328.

[34] Dey S, Lakshmanan MK, Pesala B. Optimal solar tree design for increased flexibility in seasonal energy extraction. Renew Energy 2018;125:1038–48. http://dx.doi.org/10.1016/j.renene.2018.02.017, URL http://www.sciencedirect.com/science/article/pii/S0960148118301605.

[35] Díaz-Dorado E, Cidrás J, Carrillo C. A method to estimate the energy production of photovoltaic trackers under shading conditions. Energy Convers Manage 2017;150:433–50.

[36] Díaz-Dorado E, Suárez-García A, Carrillo CJ, Cidrás J. Optimal distribution for photovoltaic solar trackers to minimize power losses caused by shadows. Renew Energy 2011;36(6):1826–35.

[37] Perpiñán O. Cost of energy and mutual shadows in a two-axis tracking PV system. Renew Energy 2012;43:331–42. http://dx.doi.org/10.1016/j.renene.2011.12.001, URL http://www.sciencedirect.com/science/article/pii/S0960148111006537.

[38] Commault B, Duigou T, Maneval V, Gaume J, Chabuel F, Voroshazi E. Overview and perspectives for vehicle-integrated photovoltaics. Appl Sci 2021;11(24):11598.

[39] Janszen J, Shahzaad B, Alkouz B, Bouguettaya A. Constraint-aware trajectory for drone delivery services. 2021, arXiv preprint arXiv:2111.00339.

[40] Alkouz B, Bouguettaya A, Mistry S. Swarm-based drone-as-a-service (sdaas) for delivery. In: 2020 IEEE international conference on web services. IEEE; 2020, p. 441–8.

[41] Pongpattana C, Rakkwamsuk P. Efficient algorithm and computing tool for shading calculation. Songklanakarin J Sci Technol 2006;28(2):375–86.

[42] Liao J, Jiang Y, Liao H, e Xiao D, Yuan J, Yang Z, et al. Investigation on the effects of shadow on output performance and thermal characteristic of the solar array on stratospheric aerostat. Energy 2019;182:765–76. http://dx.doi.org/10.1016/j.energy.2019.06.092, URL http://www.sciencedirect.com/science/article/pii/S0360544219312241.

[43] Budin R, Budin L. A mathematical model for shading calculations. Sol Energy 1982;29(4):339–49.

[44] Yezioro A, Shaviv E. Shading: A design tool for analyzing mutual shading between buildings. Sol Energy 1994;52(1):27–37.

[45] Braunholtz-Speight T, Sharmina M, Manderson E, McLachlan C, Hannon M, Hardy J, et al. Business models and financial characteristics of community energy in the UK. Nat Energy 2020;5(2):169–77.

[46] Hadwan M, Alkholidi A. Assessment of factors influencing the sustainable performance of photovoltaic water pumping systems. Renew Sustain Energy Rev 2018;92:307–18.

[47] Sengupta M, Xie Y, Lopez A, Habte A, Maclaurin G, Shelby J. The national solar radiation data base (NSRDB). Renew Sustain Energy Rev 2018;89:51–60.

[48] Martínez-Moreno F, Muñoz J, Lorenzo E. Experimental model to estimate shading losses on PV arrays. Sol Energy Mater Sol Cells 2010;94(12):2298–303.

[49] Anderson SM, Byer OD. Finding polygonal areas with the Corset theorem. College Math J 2017;48(3):171–8.

[50] Jarvis RA. On the identification of the convex hull of a finite set of points in the plane. Inform Process Lett 1973;2(1):18–21.

[51] Gotelli N, Ellison A. A primer of ecological statistics. Macmillan Education; 2013, URL https://books.google.com.co/books?id=zTbjMQEACAAJ.

[52] Ewens WJ. On estimating p values by Monte Carlo methods. Am J Hum Genet 2003;72(2):496.

[53] Manly BF. Randomization, bootstrap and monte carlo methods in biology, Vol. 70. CRC Press; 2006.

[54] Anderson MJ, Robinson J. Permutation tests for linear models. Aust N Z J Stat 2001;43(1):75–88.

[55] Arias-Rosales A, LeDuc PR. Modeling the transmittance of anisotropic diffuse radiation towards estimating energy losses in solar panel coverings. Appl Energy 2020;268:114872. http://dx.doi.org/10.1016/j.apenergy.2020.114872, URL http://www.sciencedirect.com/science/article/pii/S0306261920303846.

[56] Arias-Rosales A, LeDuc PR. Comparing view factor modeling frameworks for the estimation of incident solar energy. Appl Energy 2020;277:115510.

[57] Dubey S, Sarvaiya JN, Seshadri B. Temperature dependent photovoltaic (PV) efficiency and its effect on PV production in the world–a review. Energy Procedia 2013;33:311–21.

[58] Notton G, Cristofari C, Mattei M, Poggi P. Modelling of a double-glass photovoltaic module using finite differences. Appl Therm Eng 2005;25(17):2854–77.

[59] Rosell J, Vallverdu X, Lechon M, Ibanez M. Design and simulation of a low concentrating photovoltaic/thermal system. Energy Convers Manage 2005;46(18):3034–46.

[60] Chapra SC, Canale RP, et al. Numerical methods for engineers. Boston: McGraw-Hill Higher Education; 2010.