


# Upgrading the PVMD Toolbox: Selective Skydome Refinement for Improved Shading Analysis

MSc. Thesis

Vasiliki Psaraki

Delft University of Technology



7 July 2025

# Upgrading the PVMD Toolbox: Selective Skydome Refinement for Improved Shading Analysis

MSc. Thesis

by

Vasiliki Psaraki

to obtain the degree of Master of Science  
at the Delft University of Technology  
to be defended publicly on July 17, 2025 at 10:30

**Thesis committee:**

Chair:	Prof.dr.ir. Olindo Isabella
Supervisors:	Dr. ir. Rudi Santbergen ir . Sathya Shanka Vasuki
External examiner:	Dr. Jianning Dong
Place:	Faculty of Faculty Electrical Engineering, Mathematics and Computer Science Delft University of Technology
Project Duration:	November, 2024 - July, 2025
Student number:	6016766

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Faculty Electrical Engineering, Mathematics and Computer Science, Delft University of Technology

# Abstract

Solar energy is predicted to be prevalent in the energy mix in the upcoming years, with its global targeted capacity for 2050 being more than tenfold compared to now. In a country like the Netherlands, where land is primarily devoted to agriculture, there is limited space available for energy production. The solution comes with Hybrid Renewable Energy Systems (HRES), and in the case of solar and wind, combining them eases the inherent variability due to their complementarity. Nevertheless, this integration introduces a possible threat, which is the shade that wind turbines can cast on adjacently placed solar panels, reducing their energy yield.

The PVMD Toolbox is a PV system energy yield prediction tool, developed in the Photovoltaic Materials and Devices (PVMD) group of TU Delft. It takes into account cell optics, system surroundings, and system losses to compute the energy outcome of the simulated PV system. When the goal is to perform simulations with modules in such a hybrid environment, it is critical to have a highly accurate shading analysis on the PV modules. This means ensuring precision in how the module perceives its surroundings. However, doing so comes at the expense of high computational cost and potential memory limitations. The objective of this thesis is to create an updated version of the current part of the toolbox that models how a PV panel processes its surroundings in order to perform highly accurate shading analysis efficiently.

Shading analysis becomes more detailed as the skydome, perceived from the module's point of view, is more finely discretized. The skydome is a hemispherical projection of the sky from the module's position on the ground, divided into triangular patches representing different sky directions. Instead of uniformly refining the entire skydome to increase accuracy, the Selective Skydome Refinement method focuses computational effort only on the regions of the sky affected by shading.

The method begins with a low-resolution simulation to identify which parts of the skydome contribute to shading. Only these regions, along with a surrounding buffer to ensure full coverage, are then refined to a higher resolution. Sensitivity to incident light values are recalculated only for this refined subset of sky triangles using targeted ray tracing, while the rest of the skydome retains coarser values replicated from the initial simulation. This selective approach reduces the number of rays that must be traced to portray the surroundings by more than 70 % in the example case. The model was fully integrated into the PVMD Toolbox with only minor adjustments, enabling faster and more scalable simulations of shading effects in complex environments such as combined wind and solar energy systems.

The most significant achievement of the developed method is that it can reduce the computational time of a single-module simulation of the required electric calculations by around 90 %, at the small cost of less than 1 % relative error in the simulated DC energy yield. Beyond its overall performance, the method was also tested for different starting skydome resolutions to balance speed and accuracy. Comparing two refinement levels, the results showed that starting from a moderately refined skydome improves shading detection while maintaining fast simulations. Both cases resulted in very low energy yield errors, confirming the method's reliability. An optimal starting level was identified, offering accurate results with minimal computation, suitable for future simulations. Finally, the method was applied to a large PV-wind layout, where contour plots of simulated and interpolated results revealed consistent spatial shading trends and highlighted the importance of high-resolution weather data in capturing short-duration shadow effects.

# Preface

*This thesis marks the end of my Master's in Sustainable Energy Technology at TU Delft. What began as a challenging journey turned into an experience that ultimately shaped me and helped me grow both academically and personally.*

*I would like to express my gratitude to my supervisors who made this journey as smooth as possible. To Sathya, my daily supervisor, for his immense help, always present and punctual at weekly meetings, each time offering insights, thinking along with me and suggesting ideas. To Rudi, for sharing his knowledge, being flexible, and always bringing a sense of calm. It meant a lot that you valued my work, your reassurance gave me confidence and motivation. I would also like to thank Youri Blom for his guidance with the toolbox and for being so approachable and eager to help whenever I asked. I'm grateful as well to Professor Olindo Isabella and Jianning Dong for being part of the thesis committee and for their input during the final stage of my project.*

*Finally, I have to thank my parents for making this Master's possible for me, supporting me emotionally and financially, being my biggest fans and most loving parents. I also want to thank the friends I made here for their love and support and for making me feel at home.*

*Vasiliki Psaraki  
Delft, July 2025*



# Contents

<b>Abstract</b>	<b>i</b>
<b>Preface</b>	<b>ii</b>
<b>Nomenclature</b>	<b>viii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Background . . . . .	1
1.2 Hybrid Integration of Wind and Solar . . . . .	2
1.3 Shading in Wind-Solar Systems . . . . .	4
1.4 PVMD Toolbox: Overview and Structure . . . . .	7
1.5 Thesis Objective & Scope . . . . .	8
1.6 Thesis Outline . . . . .	9
<b>2 Case Study for Shading Analysis: Wind Turbine</b>	<b>10</b>
2.1 Building the Simulation Environment . . . . .	10
2.1.1 Fundamentals of 3D Environment Construction in PVMD Toolbox . . . . .	10
2.1.2 Single Wind Turbine Environment . . . . .	11
2.1.3 Adding PV Modules to the Environment . . . . .	15
2.1.4 PV Module Placement Strategy . . . . .	16
2.2 Simulations Setup Assumptions . . . . .	18
2.3 Simulations Setup Case Studies . . . . .	19
2.3.1 Example Case A . . . . .	19
2.3.2 Example Case B . . . . .	20
2.3.3 Example Case C . . . . .	20
2.3.4 Example Case D . . . . .	21
2.4 Conclusions . . . . .	23
<b>3 Shadow Modeling with the Existing PVMD Toolbox</b>	<b>24</b>
3.1 Ray Tracing and Sensitivity Maps . . . . .	24
3.2 Importance of Skydome Discretization . . . . .	26
3.3 Irradiance Distribution . . . . .	28
3.4 Energy Output . . . . .	30
3.5 Tradeoff of Shading Sensitivity Analysis . . . . .	33
3.5.1 Computational Burden Analysis . . . . .	33
3.5.2 DC Energy Yield Analysis . . . . .	34
3.6 Conclusions . . . . .	36
<b>4 Selective Skydome Refinement Method</b>	<b>37</b>
4.1 Conceptual Overview of Selective Refinement . . . . .	37
4.2 Selective Ray Tracing Methodology . . . . .	38
4.2.1 Integration and Key Variables . . . . .	38
4.2.2 Identification of Shaded Sky Triangles in the Initial Refinement . . . . .	39
4.2.3 Buffer Zone and Final Selection . . . . .	41
4.2.4 Increase Mesh to Target Refinement . . . . .	42
4.2.5 Selective Backward Ray Tracer . . . . .	42
4.3 Output . . . . .	44
4.4 Conclusions . . . . .	46
<b>5 Performance Evaluation and Application of the Developed Skydome Refinement Method</b>	<b>47</b>
5.1 Validation of the Selective Refinement Method . . . . .	47

5.1.1	Computational Cost Analysis . . . . .	47
5.1.2	Energy Yield Analysis . . . . .	48
5.2	Optimal Starting Resolution for Selective Refinement . . . . .	49
5.3	Shading Analysis for a larger PV system . . . . .	52
5.3.1	Energy Yield Variability Between Adjacent Modules . . . . .	52
5.3.2	Analysing shading patterns . . . . .	54
5.4	Conclusions . . . . .	59
<b>6</b>	<b>Conclusions and Recommendations</b>	<b>61</b>
6.1	Conclusions . . . . .	61
6.2	Future Work Recommendations . . . . .	62
<b>A</b>		<b>68</b>
A.1	PVMD Toolbox Flowchart . . . . .	68
<b>B</b>		<b>69</b>
B.1	Non-linear Exponential Model Used for Computational Time Fitting . . . . .	69
<b>C</b>		<b>70</b>
C.1	Relative Errors in Energy Yield Between Levels of Full Refinements . . . . .	70
C.2	Comparison of Sensitivity Maps for Example Case A . . . . .	71
C.3	Number of Refined Triangles per Module for $N_{\text{ref}} = 2, 3 \rightarrow 6$ . . . . .	72
<b>D</b>		<b>73</b>
D.1	Example Case C: Grid Analysis Setup . . . . .	73
D.2	Single Grid Square Uncertainty Analysis . . . . .	75
D.3	Analysis of a cluster of 16 Modules . . . . .	77
D.4	Clustering the 184 modules . . . . .	79
D.5	Uncertainty analysis of clusters . . . . .	80
D.6	Contour plot analysis using different weather data resolution . . . . .	82
D.6.1	Solar elevation - Shadow length . . . . .	82
D.6.2	Velocity of the Turbine Tower Shadow Tip . . . . .	83

# List of Figures

1.1	Global transforming energy scenario by 2050. Reproduced from IRENA, "Global Renewables Outlook 2020" [1]. . . . .	2
1.2	Monthly wind, solar, and combined energy production at the East site offshore Asturias, Spain. Solar output exceeds wind in summer, illustrating seasonal complementarity. Reproduced from [26]. . . . .	3
1.3	Monthly wind and solar energy production at the London Array offshore wind farm (UK), illustrating seasonal complementarity. Reproduced from [19]. . . . .	4
1.4	Vattenfall's largest hybrid energy park, Energy Park Haringvliet Zuid, features an efficient combination of wind turbines, solar panels and batteries. Reproduced from [15]. . . . .	5
1.5	Tower shading and shadow components of a vertical pole. Reproduced from [33]. . . . .	5
1.6	Comparison of shading simulation methods: (a) dynamic simulation using Monte Carlo blade sampling and GPU rendering, (b) static simulation using fixed turbine geometry in PVSyst. . . . .	6
2.1	Simulation environment including a ground plane and one wind turbine. . . . .	13
2.2	Close-up views of a wind turbine nacelle and blades from various angles. . . . .	14
2.3	PV system of one wind turbine and 200-module incorporated in the environment on the ground. . . . .	15
2.4	PV system of 200 modules. . . . .	16
2.5	Hybrid solar-wind park layouts in (a) onshore and (b) offshore environments. . . . .	17
2.6	Visualization of the simulation environment of Example Case A . . . . .	19
2.7	Visualization of the simulation environment of Example Case B . . . . .	20
2.8	Visualization of the simulation environment of Example Case C . . . . .	21
2.9	Visualization of the simulation environment for Example Case D. . . . .	22
3.1	Sensitivity map, of high accuracy, of one module of Example Case A. . . . .	25
3.2	Seasonal variation in solar beam concentration due to changing solar angle. Reproduced from D. V. Schroeder, licensed under CC BY-NC 3.0 [45]. . . . .	26
3.3	Sensitivity Maps for increasing skydome refinement levels from 1 to 6. . . . .	27
3.4	Irradiance distribution on the 60-module PV system at 14:00, on June 1 <sup>st</sup> . . . . .	28
3.5	Irradiance distribution on the 60-module PV system at 15:00, on June 1 <sup>st</sup> . . . . .	29
3.6	Three chosen modules with different shadings at two consecutive hours on June 1 <sup>st</sup> . . . . .	29
3.7	Comparison of irradiance and photocurrent for modules 22, 59, and 45. throughout the 1 <sup>st</sup> of June . . . . .	31
3.8	Comparison of absorbed irradiance and DC power output for modules 15, 22, and 46 throughout the 1 <sup>st</sup> of June. . . . .	32
3.9	Increasing computational time with higher skydome discretization levels, fitted with an exponential model. . . . .	33
3.10	DC power output and relative error across skydome discretization levels. Purple data points indicate daily DC energy output (Wh) and gray bars are the errors of each value relative to $N_{\text{ref}} = 6$ . . . . .	35
4.1	Flowchart for selective refinement and sensitivity processing . . . . .	38
4.2	Sensitivity map for skydome discretization level $N_{\text{ref-initial}} = 2$ . . . . .	39
4.3	Visible indexing for the sky facets of discretization level $N_{\text{ref}} = 2$ . . . . .	40
4.4	Initial <i>potentially shaded triangles</i> with buffer area outlined in red, including 2 layers of buffer zone. . . . .	41
4.5	Resulting sensitivity map for selectively discretized skydome, from $N_{\text{ref}} = 2$ to $N_{\text{ref}} = 6$ . . . . .	44
4.6	Evolution of the sensitivity map through the selective refinement process. . . . .	45

5.1	Sensitivity maps of the module of Example Case C, for four different skydome discretization levels: (a) $N_{\text{ref}} = 2$ , (b) $N_{\text{ref}} = 3$ , (c) $N_{\text{ref}} = 2 \rightarrow 6$ , and (d) $N_{\text{ref}} = 3 \rightarrow 6$ . . . . .	51
5.2	The 184-module PV system and 1 wind turbine used in the simulation. . . . .	52
5.3	Comparison of contour plots for illustrating the DC power losses due to a wind turbine shadow over 184 modules, using 1h and 10 minute resolution data. . . . .	56
5.4	Contour plots for DC power output variation within a day for two modules, with 1h and 10min resolution, showing how shading conditions differ with the weather data resolution. . . . .	57
5.5	Contour plots for DC power output variation within a day for a module, where with 1h the yield is overestimated compared to 10min resolution. . . . .	58
5.6	Contour plots for DC power output variation within a day for a module, with 1h and 10min resolution. North of the turbine, the shadow influences less the energy yield. . . . .	59
A.1	Flowchart of PVMD Toolbox describing the simulation methodology. Reproduced from [39]. . . . .	68
C.1	Comparison of sensitivity maps using full versus selective skydome refinement at level $N_{\text{ref}} = 6$ . . . . .	71
D.1	Grid Layout with 576 squares for PV placements . . . . .	73
D.2	The 184 grid squares with simulated PV modules. . . . .	74
D.3	Spatial layout of the 25 PV modules within a $12.5 \text{ m} \times 12.5 \text{ m}$ grid square. . . . .	75
D.4	Spatial layout of the 16 PV modules within a $50 \text{ m} \times 50 \text{ m}$ . . . . .	77
D.5	Clustering of PV modules for simulations. . . . .	79
D.6	Solar position and resulting shadow length throughout 1 <sup>st</sup> June 2005 in Delft. As the elevation angle increases, the turbine shadow shortens, reaching a minimum at solar noon. . . . .	82

# List of Tables

3.1	Computational time in minutes for each level of skydome discretization. . . . .	33
3.2	Computational time in minutes for each level of skydome discretization. . . . .	34
5.1	Comparison of total computational time between the original and selective refinement methods. . . . .	48
5.2	Comparison of DC energy output (in kWh) for unrefined and selectively refined methods. . . . .	48
5.3	Relative error (%) between unrefined and selectively refined methods for selected periods. . . . .	49
5.4	Comparison of DC energy output (in kWh) for unrefined and the two selectively refined methods. . . . .	50
5.5	Relative error (%) between unrefined and the two selectively refined methods for selected periods. . . . .	50
5.6	Statistical overview of center-based sampling accuracy within a 12.5 m × 12.5 m grid square. . . . .	53
5.7	Statistical overview of center-point sampling accuracy within a 50 m × 50 m module cluster. . . . .	54
B.1	Estimated parameters for the exponential model $y = y_0 + A \cdot \exp(R_0 \cdot x)$ . . . . .	69
C.1	Relative error (%) in DC energy yield at each level of skydome discretization, compared to $N_{\text{ref}} = 6$ , for two representative dates. . . . .	70
C.2	Number of refined parent triangles and total child triangles to recalculate per module for 32 modules. . . . .	72
D.1	DC yield and relative error (vs. central position) for all 25 modules in Position 406 D.1 or 110 D.2. . . . .	76
D.2	DC yield and relative error (vs. central position) for all 16 modules in Batch 10. . . . .	78
D.3	DCP simulation metrics for clusters 1–3 and statistical summary. . . . .	80
D.4	DCP simulation metrics for clusters 4–6 and statistical summary. . . . .	80
D.5	DCP simulation metrics for clusters 7–10 and statistical summary. . . . .	81
D.6	Summary statistics of solar elevation angle and turbine shadow length on 1 <sup>st</sup> June 2005 in Delft (threshold = 0°). . . . .	83

# Nomenclature

## Abbreviations

AC	Alternating Current
c-Si	Crystalline Silicon
DC	Direct Current
DHI	Diffuse Horizontal Irradiance
DNI	Direct Normal Irradiance
EQE	External Quantum Efficiency
GHI	Global Horizontal Irradiance
GPV	Ground-mounted Photovoltaics
GUI	Graphical User Interface
HRES	Hybrid Renewable Energy System
II3050	Integrated Infrastructure Outlook 2050
IRENA	International Renewable Energy Agency
IV	Current–Voltage
MCS	Monte Carlo Simulation
MoE	Margin of Error
MPPT	Maximum Power Point Tracking
MRQ	Main Research Question
PERC	Passivated Emitter and Rear Contact
POA	Plane of Array
PS	Power Smoothing
PV	Photovoltaic
RSD	Relative Spectral Distribution
SM	Sensitivity Map
SVF	Sky-View Factor
UV	Ultraviolet

## Symbols

$\Delta t$	Time step duration [h]
$\Delta T_{i \rightarrow j}$	Computational time increase from refinement level $i$ to $j$
$\epsilon_{\text{rel}}(i - j)$	Relative error in DC energy yield between refinement levels $i$ and $j$ [%]
$\sigma$	Standard deviation of energy yield sample [kWh]
$\tau$	Kendall's coefficient of concordance
$\theta_{\text{AOI}}$	Angle of incidence of the incoming ray [°]
$\theta_{\text{AOI, primary}}$	Angle of incidence of the primary ray [°]
$A$	Absorbed irradiance [ $\text{W m}^{-2}$ ]
$\text{DC\_EY}_{\text{no turbine}}$	DC energy yield without wind turbine shading [kWh]
$\text{DC\_EY}_{\text{with turbine}}$	DC energy yield with wind turbine shading [kWh]
$EY$	Energy yield [ $\text{W h}$ ]
$F$	Global face matrix defining triangular surfaces
$I_{\text{mpp}}$	Current at maximum power point [A]
$I_{\text{sc}}$	Short-circuit current [A]
$I_{rr}$	Total irradiance [ $\text{W m}^{-2}$ ]
$J_{\text{ph}}$	Photocurrent density [ $\text{mA cm}^{-2}$ ]
$L_{\text{shading}}$	DC power loss due to shading by wind turbine [%]
$MoE$	Margin of error in estimated energy yield [kWh]
$N_{\text{ref}}$	Skydome refinement level
$P_{\text{mpp}}$	Power at maximum power point [W]
$T$	Computational time [min]
$V$	Global vertex matrix for 3D object geometry
$V_{\text{mpp}}$	Voltage at maximum power point [V]
$V_{\text{oc}}$	Open-circuit voltage [V]
$(x, y, z)$	Spatial coordinates in 3D space [m]
$n$	Number of samples (modules, grid points, or triangles)
$n_{\text{children}}$	Number of child triangles per parent triangle after refinement
$r_l$	Refinement level difference between base and target skydome resolution
$s$	Sample standard deviation of energy yield [kWh]
$t^*$	Critical value from Student's $t$ -distribution
$\bar{x}$	Sample mean energy yield [kWh]
$v_{\text{shadow}}$	Shadow-tip velocity across the PV field [ $\text{m s}^{-1}$ ]
$t_{\text{shadow}}$	Time for shadow to pass over a module [s]
$f(x, y)$	Interpolated energy yield at position $(x, y)$ in local grid
$a, b, c, d, e, f$	Coefficients of 2D quadratic interpolation surface
$f_{00}, f_{10}, f_{01}, f_{11}, f_c$	Known values used to constrain 2D surface interpolation [kWh]
$z$	Z-score for normal distribution confidence intervals

# 1

## Introduction

This chapter introduces the topic of the thesis and establishes essential concepts and background information that serve as a foundation for this work. Section 1.1 explores the current state of the renewable energy sector, with emphasis on the solar industry. Section 1.2 builds on this by examining the hybrid wind–solar integration and discusses the complementarity between wind and solar resources in such systems. The challenge of shading that comes with this synergy is discussed in Section 1.3, where also some of the existing shadow modeling approaches in the literature are presented. Section 1.4 introduces the PVMD Toolbox and its relevant components. Section 1.5 outlines the motivation, the objective, and the scope of this thesis, as well as the research questions it seeks to address. Finally, Section 1.6 provides an overview of the thesis structure.

### 1.1. Background


Decarbonization in the energy sector is a central goal in the energy transformation plans outlined by the International Renewable Energy Agency (IRENA) [1], as it accounts for over 75 % of total GHG emissions in the EU. For 2050, according to IRENA, in order to meet the agreed global climate goals, renewables would need to provide two-thirds of the world's energy supply [1]. Solar PV and wind energy would lead the way in the transformation of the global electricity sector.

As shown in Figure 1.1, it is projected, in the 'Transforming Energy Scenario', that by 2050 the electricity generation coming from renewable energy sources will reach the 86 %. Namely, in this scenario wind power will be supplying one third and then solar, the second biggest contributor, supplying 25 % of total electricity demand. The noteworthy is that the expansion in the installed capacity of solar photovoltaics (PV) would be 8519 GW which entails a more than tenfold increase in its share of the generation mix compared to 2017 [1].

Specifically, the European Union has committed to reduce its greenhouse gas emissions by at least 55 % until 2030 and to go climate neutral by 2050 [2, 3]. In order to achieve that, the renewable energy production needs to increase at least up to 42.5 % renewable energy in the energy mix by 2030, therefore, with an aspirational goal of reaching 45 % [2].

In the Netherlands in particular, solar power is the second major renewable energy source [4]. The country is committed to this goal, targeting to produce 42 TWh of land-based renewable electricity every year by 2030 [5]. The even more longterm goal, according to the Integrated Infrastructure Outlook 2050 (II3050), the expected installed capacity for solar PV in 2050 is projected to range between 100 GW to 180 GW [6, 7]. The government is actively supporting that with several actions and initiatives. Subsidies that promote renewable energy production are also provided. An example is the 'Renewable electricity' SDE++, launched in 2024, a €13 billion subsidy [8].





**Figure 1.1:** Global transforming energy scenario by 2050. Reproduced from IRENA, "Global Renewables Outlook 2020" [1].

These targeted capacities cannot be accommodated by relying solely on large-scale ground-mounted photovoltaic (GPV) farms [9], due to the scarcity of space and land restrictions. In the Netherlands agriculture has a central role in national spatial planning and land management. It is estimated that more than 80 % of the Netherlands' total area is allocated to agriculture, recreation, forests, and natural landscapes, with nearly two-thirds of this area dedicated specifically to agricultural activities in 2015 [10]. Thus, multiple land uses are encouraged, such as adding solar into buildings and infrastructure or vehicles.

Out of the total 42 TW h, 7 TW h is expected to come from small-scale solar installations in the already built environment, like domestic rooftops [5]. In 2019, the Netherlands introduced the "Zonneladder" scheme, prioritizing solar installations on rooftops and built environments over agricultural lands [11]. In 2022, 1.8 GW of rooftop solar capacity was added, which is 38 % more than the amount added in 2021 [12]. Apart from that, there is a growing trend to combine solar and wind energy, primarily due to severe grid congestion in the Netherlands. Co-locating these technologies allows the use of shared substations, cabling, and, most importantly, grid connection infrastructure [13]. To support this, in April 2021, Holland Solar, NWEA, and Energie Samen introduced a model agreement to help large-scale wind and solar asset owners share connection points and ease grid pressure [14]. One example is the 60 MW hybrid energy park, consisting of 124,000 solar panels, six wind turbines and batteries at Haringvliet in the Netherlands [15]. Shell also has announced that plans to construct a hybrid 50 MW plus 50 MW solar-wind farm [16]. More ambitious ideas entail incorporating solar panels jointly with agriculture or even floating on inland water bodies or offshore sites [5, 17].

## 1.2. Hybrid Integration of Wind and Solar

The renewable energy sources are nonnegotiable, but they come with limitations and inherent disadvantages. The biggest is their inherent variability since they are driven by meteorological factors. Existing power grids may be strained by the inconsistent supply of standalone renewable sources, leading to variations in frequency and voltage [18]. This causes fluctuations in power output and challenges grid stability, often requiring costly system oversizing [19, 20].

These can be tackled when combining sources, in other words, forming Hybrid Renewable Energy Systems (HRES). Integrating sources entails many advantages. First and foremost, there is more energy availability and supply stability. Even more, when adding storage into the system, excess energy is not lost during times of high production and used during low production periods. That improves the overall efficiency of the system, makes the grid more stable and can reduce grid congestion [18].

This thesis is specifically focused on the wind and solar synergy. A big strength of the wind and solar mix is that both renewable sources are at a high level of technological maturity [21, 22] and that their intermittency is on a large scale complementary. Adding solar capacity increases cable usage, known


as cable pooling, and improves utilization of cables during periods of suboptimal wind power generation [23]. Golroodbari et al. [23] state that although the cable capacity is exceeded 12.07 % of the time, only 1.72 % of energy is lost because unused wind capacity allows for more solar power to be generated without overloading the system.

The complementarity of solar and wind power helps mitigate fluctuations and intermittency, reducing the need for energy storage while improving supply reliability and grid integration [20, 22, 24]. Vattenfall [25] observed in one of their combined solar and wind farms that only during 2 % of the time do solar and wind farms operate simultaneously at more than 75 % of their full capacity. Also, it is known that the power generation efficiency of solar modules is higher at lower temperatures. Thus, placing them in a space selected for its wind capacity, adjacent or in between a wind park, means that the wind can cool down the panels and improve their efficiency. Reducing power output intermittency enhances energy production, resulting in a smoother and more stable power supply throughout the year [19, 21, 22].

Naturally, many studies discuss the complementarity between wind and solar energy. Various correlations, such as Pearson's, Kendall's, and Spearman's coefficients, can quantify the relationship between wind and solar energy. Kendall's coefficient  $\tau$  is many times preferred because it doesn't require linearity, normal distribution, or homoscedasticity [19]. It ranges from -1 (perfect complementarity) to 1 (perfect synchronicity), with negative values indicating complementarity, which is desirable. Results of [19] and [22] indicate weak complementarity between wind and solar energy on hourly and daily timescales but strong complementarity on weekly and monthly. Similarly, [24] reveals that on a larger time scale deviations in power fluctuations are significantly reduced. In most regions, reverse fluctuations occur more often than synchronous, and times that generation becomes zero are uncommon.

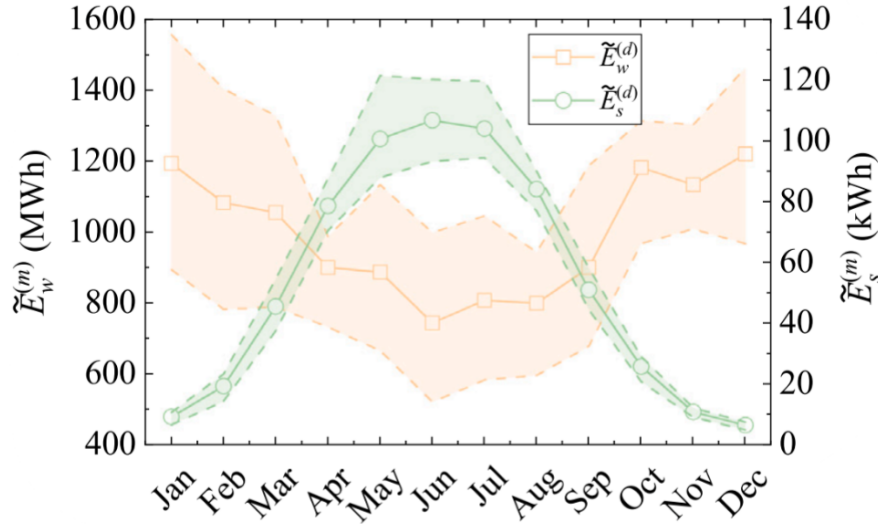
Other studies introduce different methods for measuring complementarity between wind and solar energy. López et al. [26] introduced a novel Power Smoothing (PS) index to quantify the reduction in power output variability in hybrid offshore wind–solar systems. Using this index, they reported up to 63 % smoothing in power output. Ren et al. [24], in an effort to provide a more accurate complementarity measure, introduced a novel index that considers fluctuation states and amplitudes, suggesting that traditional correlation coefficients may overestimate complementarity. They found that this helps to optimize the capacity ratio when selecting components for hybrid systems, accounting for both available and unavailable power periods.

Figures 1.2, 1.3 present two studies [19, 26] that demonstrate satisfactory complementarity over the course of a year. The contrasting patterns of variation between wind and solar energy suggest that they are likely to complement one another.



**Figure 1.2:** Monthly wind, solar, and combined energy production at the East site offshore Asturias, Spain. Solar output exceeds wind in summer, illustrating seasonal complementarity. Reproduced from [26].

Both reveal the interannual seasonal variability, with higher wind energy yield in winter and lower yield in summer. The graphs are based on modeled energy outputs using location-specific meteorological data and commercial specifications for wind turbines and PV panels. In the Asturias case 1.2, wind and solar resources were derived from the SIMAR and NASA POWER datasets. In the London Array case 1.3, ERA5 data and manufacturer power curves were used to assess complementarity and inform the sizing of PV. Overall, solar energy exhibits less year-to-year variability than wind energy [19].



**Figure 1.3:** Monthly wind and solar energy production at the London Array offshore wind farm (UK), illustrating seasonal complementarity. Reproduced from [19].

In the study of Delbeke et al. [22], the behavior of wind-solar complementarity was investigated over a time frame of 70 years under two climate scenarios of global warming. It is observed that complementarity becomes more prominent as the time scale shifts from hourly to monthly, aligning with other studies. According to the study's projections, complementarity may gradually improve in the face of significant global warming. Despite the significant fluctuations, even across decades, the overall systematic pattern remains consistent toward the end of the century.

### 1.3. Shading in Wind-Solar Systems

Despite the benefits of integrating wind and solar PV systems, several challenges must be addressed to optimize their synergy. When considering the combination of solar and wind energy, it is imperative to identify the factors that can impact the power output of PV modules and, by extension, the entire PV system. Specifically, it is important to understand how the presence of nearby wind turbines affects the irradiance on the plane of the array (POA) of the modules and, consequently, their energy production. The POA loss depends on the sun position, the array orientation, and the shading [27].

One major concern that can lead to changes in POA irradiance, and thus in energy yield in such a synergy, is shading, which is the focus of this dissertation. Shading on solar panels leads to uneven exposure across cells and modules, causing differences in operating conditions, reducing output, and leading to power dissipation and heating [27]. In a wind-solar HRES, the shading effect can be partial, as the nacelle rotates with the wind direction while the PV modules stay stationary, resulting in a dynamic shading pattern over time [28, 29]. For modeling purposes, wind turbine shading can be interpreted as comprising two types of shadows [30, 31]: the slowly moving shadow of the entire wind turbine, caused by the sun's movement across the sky, and the dynamic, rapidly moving shadow of the turbine blades, which is driven by wind variation.

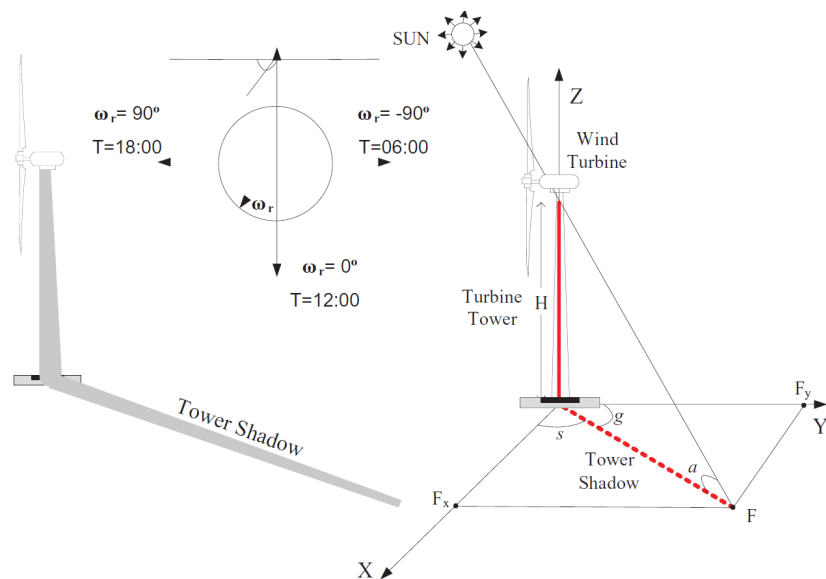
The sun's angles change throughout the day and across seasons, altering where and how shadows from the turbines would fall on the adjacent solar panels. This affects the total shade of the turbine from

the tower, the nacelle and hub, and the blades. The shadows cast by the tower can be perceived as the most significant type of shadow, as they fully obstruct direct sunlight [32]. Alternating shading also affects the inverter performance, causing delays in adjusting to optimal energy output and resulting in further energy losses [30].



**Figure 1.4:** Vattenfall's largest hybrid energy park, Energy Park Haringvliet Zuid, features an efficient combination of wind turbines, solar panels and batteries. Reproduced from [15].

Researchers in the past few years have employed methods to model the shading of a turbine on an adjacent PV module in an effort to quantify the significance of these losses. Some studies model static shading losses in a hybrid wind–solar farm, focusing only on the tower shadow. For instance, Routray et al. [33] that approximate the wind turbine as a vertical pole and use a combination of solar position geometry and graphical modeling. The shadow projection is determined analytically via trigonometric formulations accounting for latitude, solar declination, and hour angle.




**Figure 1.5:** Tower shading and shadow components of a vertical pole. Reproduced from [33].



They employ a Monte Carlo Simulation (MCS) to generate a number of random scenarios from these distributions to simulate hourly power output, including losses from shading or wake effects. Their results show that tower shadows can lead to daily solar losses of up to around 3 %, highlighting the importance of spatial layout in co-located renewable systems.

Taking it one step further, others also attempt to model the blade shadows, following different methods such as approximating their shape as a sphere [28], a disc [28, 34], or an elliptical shape [35]. In some cases, studies also model the actual shape of the wind turbine. Ludwig et al. [36] model the full turbine geometry (tower, nacelle, blades) in PVSyst's 3D environment to estimate string-level shading losses, seen in Figure 1.6a. Their results show that turbine shading accounts for at most 2.3 % yield loss, often less than the self-shading among PV rows (3.1 %). Robledo et al. [28], who developed a dynamic 3D simulation method using GPU rendering to model the movement of turbine blades and their shading impact on PV modules, seen in Figure 1.6b. Using a Monte Carlo sampling of blade positions and GPU-based 3D shading to compute cell-level losses, which are translated to power reduction via a diode model. Results show shading losses stay below 2% when modeled accurately.



**Figure 1.6:** Comparison of shading simulation methods: (a) dynamic simulation using Monte Carlo blade sampling and GPU rendering, (b) static simulation using fixed turbine geometry in PVSyst.

As previously mentioned, as the wind turbine blades rotate, they cast moving shadows that are neither static, uniform, nor predictable. These shadows constantly change in position and size depending on the current position of the wind turbine blades, which is influenced by the wind [37]. Moving shadows cause alternating light intensity changes and create multi-steps in the I–V characteristic with local maximum powers [38], potentially decreasing energy production [37]. Modeling this effect is complex, as the blade movement is dynamic and influenced by factors such as dominant wind directions and the yaw mechanism of the turbine [28].

Views differ on how important it is to model the flickering in a hybrid wind-solar park. Some researchers argue that quantifying the impact of shadow flicker on solar energy production is crucial [37]. Others believe that the effect of flickering is far less when compared to the tower shading. Hence, the losses incurred due to shadow flickering would not undermine the viability of a hybrid wind-solar project [32].

Dekker et al. [25], focusing on the dynamic shading of wind turbine blades, noticed that there was an energy loss of about 6 %, depending on the park configuration, PV modules, and inverter type. Assuming a simplified rotor blade area, the study finds that shading causes less than 1 % land loss for PV deployment. On the contrary, a previous TNO study estimated that dynamic shading can cause greater energy losses than slower-moving shadows, with temporary outliers of up to 50 % energy loss in some areas [30, 31].

## 1.4. PVMD Toolbox: Overview and Structure

The PVMD Toolbox is an in-house MATLAB based modeling software developed by the Photovoltaic Materials & Devices (PVMD) Group at TU Delft. It is a photovoltaic (PV) system yield prediction tool, that can simulate the AC energy yield of a PV system without measurements but based on fundamental material parameters [39]. It is a holistic tool in the sense that it starts from being able to dive in cell level optical properties with a semiconductor simulation to model, but also expands to module and system level. It can simulate real-size industrial PV modules, taking into account the module's mounting conditions and the system's surroundings.

Using spectrally resolved ray tracing, a measure of each module cell's sensitivity to incoming rays is computed, resulting in the so-called Sensitivity Map (SM). One step further, the toolbox incorporates weather data to compute the absorbed irradiance and resulting photocurrent at each cell for any specified moment within the user-defined time period. Through the use of a thermal model, it also provides the individual cell temperatures and then the DC energy yield and each solar cell's IV-curve hourly to then convert into the AC energy yield of the system. Lastly, it gives the possibility of analyzing system losses, evaluating solar cell or PV system degradation over time, calculating the Levelized Cost of Electricity. A detailed workflow chart of the PVMD Toolbox can be found in Appendix A. For further information on the toolbox and its specific uses, see Vogt et al. [39].

In the present thesis all simulations were conducted using the PVMD toolbox. The toolbox includes the following interconnected components: CELL, MODULE, WEATHER, THERMAL, ELECTRIC, DEGRADATION, CONVERSION and FINANCIALS. The functions are performed in the order specified in the aforementioned list whilst running a simulation. For this study, four of these are relevant and will be therefore analysed more:

- **MODULE:** In simulation cases that may include more than one modules and properties as incident energy, operating temperatures, etc. vary between the modules the simulation is *non-periodic*. The sub-function `MODULE_nonperiodic` will be then used. In this part ray tracing is performed, which simulates the PV module mounting conditions and surroundings allowing for shading and albedo effects to be considered. After the ray tracing, based on the wavelength and angle of incidence, material properties such as reflectivity, absorptance, and opacity are assigned to the surfaces. The ultimate output of this part is the sensitivity values of each module and the generation of a Sensitivity Map (SM). The values in the sensitivity map represent how responsive the surface is to incoming rays. It is defined as the ratio between the number of rays absorbed by the module and the total number of incoming rays for each solid angle in the skydome. It's important to note that creating these sensitivity maps is the most computationally intensive part of the toolbox simulations.
- **Weather:** Integrates real-world meteorological data to provide site-specific conditions for simulations. It includes sun azimuth and elevation, Global Horizontal Irradiance (GHI), Direct Normal Irradiance (DNI), Diffuse Horizontal Irradiance (DHI), ambient temperature, wind speed, relative humidity, and UV radiation. If horizon reconstruction is enabled, DNI and DHI values are adjusted based on the sky-view factor (SVF) and terrain shading. A key step in the WEATHER module is adjusting the Relative Spectral Distribution (RSD) based on air mass (AM). Spectral irradiance is computed using either SMARTS (clear-sky conditions, no cloud effects) or SBDarts (includes cloud scattering and absorption, modifying spectral composition). The total irradiance  $I_{rr}$  ( $\text{W m}^{-2}$ ), the absorbed irradiance  $A$  ( $\text{W m}^{-2}$ ), and the generated photocurrent  $J_{ph}$  ( $\text{mA cm}^{-2}$ ) in each absorption layer of every cell are also computed at each time step.
- **Thermal:** Executes calculations for the cell temperatures of all modules of the PV system. The toolbox offers multiple thermal models to simulate temperature dynamics, including the Fluid Dy-

namic Model, Duffie-Beckman Model, Faiman Model, Sandia Model, and Incropera Model. However, in this stage of this study, we are not interested in the thermal effects on energy losses. The goal is to portray the losses coming from shading only. Therefore, the temperature for each PV cell is directly assigned a constant value of  $25^{\circ}\text{C}$  for every module, across all time steps of the simulation. This ensures that the impact of wind turbine shading on PV performance is isolated and that variations in power output are attributed solely to shading effects rather than temperature fluctuations.

- **Electric:** Handles the electrical performance of PV modules based on irradiance and temperature conditions. It generates the current-voltage (IV) characteristics of the modules by utilizing data from the WEATHER, THERMAL, and MODULE modules. The primary goal of this module is to determine the energy output of each panel by computing key electrical parameters. For each module and every simulation hour, key electrical parameters are calculated: short-circuit current  $I_{sc}$  and open-circuit voltage  $V_{oc}$  represent the maximum achievable current and voltage. The maximum power point current  $I_{mpp}$  and voltage  $V_{mpp}$  define the operating conditions for peak power output. The DC power  $P_{mpp}$  (W) is computed at each time step, and summing it over all modules and time steps yields the total energy yield  $EY$  (W h).

While the PVMD Toolbox provides a comprehensive framework for photovoltaic modeling, its current structure is best suited to relatively simple environments. It offers the possibility to perform either optical simulations of a single PV panel in isolation or simulations of full PV systems that include nearby objects representing simplified environments, such as buildings. In the latter case, the inclusion of surrounding elements is handled through a MATLAB-based GUI, which is intended for constructing simple, static environments. Its design and interface are optimized for simulations on a module or small system level, and not for the complexity introduced by large-scale, hybrid configurations such as wind–solar farms. These limitations are further discussed in the next section.

## 1.5. Thesis Objective & Scope

This thesis is motivated by the current trend toward combining renewable energy sources, and in particular the system integration of wind and solar systems. The proven benefits of wind–solar synergy underscore the relevance of this research. To investigate this integration and assess shading effects on photovoltaic (PV) modules, a specific system geometry is used throughout the thesis. A static single wind turbine is chosen in this study as a case study. In this context, the term *static* refers to the use of a fixed wind turbine structure without modeling the movement of the blades. The dynamic shading effects caused by rotating blades are acknowledged but fall outside the scope of this work. In addition, the current implementation models the sun as a point source, assigning all direct irradiance to a single sky triangle. While this introduces some limitations in angular accuracy, improving this aspect falls outside the scope of the present work and is proposed as a future extension.

The PVMD Toolbox is a powerful simulation tool ranging from intricate modeling at the cell level up to the module and system levels. It supports photovoltaic system modeling capable of predicting the energy yield of a PV module placed in different environments. One of its many capabilities is analyzing the surroundings in which the PV module is placed and assigning a value called *sensitivity*, which quantifies how much light from each sky direction contributes to absorption in the PV cell.

Nevertheless, as explained in the previous section, performing the shading sensitivity analysis with high accuracy can become very time-consuming. This becomes even more challenging when the environment includes slender objects, such as wind turbines, because their thin shape requires a much finer resolution to capture their shadow properly. Aside from that, MATLAB's memory issues can also pose a problem, as simulations may encounter "out of memory" errors depending on the computer they are run on. This creates a limitation in the current modeling approach of the PVMD toolbox, especially when evaluating energy yield in complex environments such as hybrid wind–solar systems. Addressing this challenge requires enhancing the toolbox to support more accurate shading analysis within a reasonable computational time. This thesis contributes to this goal by refining the method for calculating shading sensitivity across the skydome.

The primary and secondary questions guiding this research are as follows:

**Main Research Question (MRQ): How can the PVMD Toolbox be upgraded to efficiently and accurately perform shading analysis, including the surroundings?**

**Sub-questions (SQs):**

1. How is a hybrid wind–solar simulation environment, including a detailed wind turbine, constructed and configured in the PVMD Toolbox?
2. How can the current version of the PVMD Toolbox model and analyze shading effects of a wind turbine on a PV module, and how does skydome discretization affect accuracy and performance?
3. What enhancements can be introduced in the PVMD toolbox to improve the accuracy and efficiency of the shading analysis of a PV module with respect to its surroundings?
4. What are the performance benefits of the developed method, and how can it be further optimized? Furthermore, how does wind turbine shading influence spatial energy yield variability across a large-scale PV layout?

## 1.6. Thesis Outline

The structure of this report is organized to address the research sub-questions outlined above. Each chapter is dedicated to answering one of the sub-questions, providing a step-by-step approach to answering the main research question.

Chapter 2, addresses Sub-question 1., introducing the simulation case study that is carried out throughout this thesis. That is a hybrid wind–solar environment with a single wind turbine in the middle. It contains the detailed construction of a wind turbine geometry in the PVMD toolbox, along with key parameters, assumptions, and representative setup cases used in the thesis. Next, Chapter 3 answers Sub-question 2., by examining how the existing version of the PVMD Toolbox models shading using the current sensitivity analysis. A big limitation appears concerning the computational cost of a highly accurate sensitivity analysis. A detailed analysis of this tradeoff is presented, validating that such simulations with complex geometry and high accuracy is not viable with the existing toolbox. These limitations are addressed in the Chapter 4, which responds to Sub-question 3. This chapter demonstrates in detail a newly developed *selective refinement method* to drastically reduce the computational time of the sensitivity calculation while maintaining high accuracy. After this, in Chapter 5 the evaluation of this method, performance gains, further optimization, and an application in a more realistic case. These are addressing Sub-question 4. Finally, Chapter 6 concludes the thesis, summarizing key practices and findings and outlining interesting directions for future work.



# 2

## Case Study for Shading Analysis: Wind Turbine

This chapter addresses Sub-question: **1. How is a hybrid wind–solar simulation environment, including a detailed wind turbine, constructed and configured in the PVMD Toolbox?** It focuses on building a simulation environment that includes the complex geometry of a wind turbine, as it is required to evaluate interactions between wind turbines and photovoltaic (PV) systems. Section 2.1 explains in detail how the environment can be introduced to the simulation and the geometric construction of a sample wind turbine, including tower, nacelle, and blade modeling. It also describes how the PV modules are added in the environment. Then, Section 2.2 showcases all the relevant parameters and assumptions for the simulations. Finally, Section 2.3 presents some example simulation setup cases that are used throughout the thesis.

### 2.1. Building the Simulation Environment

The environment is the first critical step of the simulation when targeting the analysis of shading. Objects in the environment are the reason shade is cast on the PV modules, varying the irradiance conditions and consequently the electrical output. Thus, a detailed analysis of the construction of the simulation environment will follow.

#### 2.1.1. Fundamentals of 3D Environment Construction in PVMD Toolbox

Using the PVMD Toolbox, simulations can be conducted through a MATLAB GUI or a script version. Either way, for every simulation, a specific 3D environment that represents the surroundings in which the PV system will be installed has to be created and loaded. The environment is defined by a structured matrix of 3D coordinates (vertices), where each row represents a single point in space with an (x, y, z) position. From these, the face matrix is created by grouping the vertices in sets of three or four, forming triangular or rectangular faces, respectively. In the face matrix, each row contains the indices of the vertices, referring to the rows in the vertex matrix, that define a particular face.

For every face in this matrix, five other properties have to be defined. The first property is the *material type*, which is selected from a predefined database of materials and their associated properties, known as the Material Library, included with the toolbox. Then the next property is *albedo*, which is set to be equal to the spectral reflectivity in the Spectral Reflectivity Library, inside the toolbox's Material Library. The other three are the *scattering* value, the *color*, and the *opacity* of each face, which are selected by the user. Finally, the complete environment is stored as a `.mat` file containing a set of structured variables: the vertex matrix, face matrix, and five corresponding property arrays for each face: material type, RGB color, opacity, albedo, and scattering. When loaded, these variables are directly available for use in the simulation.

There are two approaches to creating a simulation environment of this type: a graphical user interface (GUI)-based method and a script-based method. If operating with the GUI version, in order to construct

the simulation environment, a custom GUI can be used in the MATLAB App Designer. This interface allows the user to interactively define 3D rectangular and triangular objects. This is done by inputting the coordinates of their vertices ( $x$ ,  $y$ ,  $z$ ). Each object can be assigned a material type (e.g., solar cell, water, support structure) so that relevant optical properties such as scattering are also assigned. The interface includes buttons to draw individual objects, delete the last object, and save the environment configuration. Once the desired layout is complete, the scene can be finalized and exported for use a specific file.

This GUI-based environment-building approach is functional and simple for conventional use cases, such as small-scale systems like residential rooftops. Nevertheless, it becomes cumbersome when trying to model large-scale, slender, or geometrically complex environments like a wind turbine. Since the user must manually determine and input the exact coordinates of each vertex, it becomes increasingly difficult and impractical for complex geometries to predict every. Additionally, focusing on geometric adaptability, it lacks the flexibility to accommodate reconfigurable environments, such as altering the nacelle orientation or adjusting blade positions. Which are essential when modeling multiple geometric scenarios for shading analysis.

To overcome these limitations, instead of using the GUI, the environment has to be constructed programmatically using a MATLAB script. In this approach, each component is built through a sequence of structured commands that generate the geometry procedurally by defining vertices and assembling faces rather than using predefined functions or modules. Optical and physical properties, including material type, RGB color, opacity, albedo, and scattering, are then assigned to each face. Once the full environment is assembled, all relevant variables are saved into a single `.mat` file, which can be loaded directly into the simulation. This script-based method offers greater flexibility and control, particularly for large-scale or geometrically complex scenarios that exceed the practical capabilities of the GUI.

### 2.1.2. Single Wind Turbine Environment

The simulation environment for this thesis includes a horizontal-axis wind turbine (HAWT) with a gear-box, which represents the most common commercial turbine configuration. It consists out of four parts in total. The complex geometry of a wind turbine is portrayed by three components: the wind turbine tower, nacelle (blade support), and the blades. This turbine is placed on a fourth component, representing the ground. The process involves building a global vertex matrix ( $V$ ) and a corresponding face matrix ( $F$ ), followed by assigning material and optical properties to each face, as mentioned before.

#### Tower

The wind turbine tower is modeled as a tapered vertical cylinder, meaning a varying radius decreasing from the base to the top. Its geometry is created by generating the vertices of a wider bottom and a narrower top circle with radii of 5 m and 4 m, respectively, and placing them 165 m apart [40]. This is done by creating the center of the circles at (0,0) and discretizing the perimeters of the base and top circles into a fixed number of segments by creating the corresponding number of angular points. Then, by connecting the centers of the circles with two of the circumference points, a triangular segment/face is created.

Next, the side faces of the tower are generated by connecting corresponding points on the bottom and top circles. For each segment, a quadrilateral patch is defined between two adjacent points on the bottom circle and their corresponding points on the top circle. Since the simulation environment is based on triangular meshes, each quadrilateral is subdivided into two triangular faces by specifying appropriate combinations of three vertex indices. In this way, the cylindrical surface of the tower is constructed as a closed and continuous 3D shape through a series of connected triangular faces, resulting in a faceted approximation of a vertical hollow cylinder. Once all faces of the tower are combined, the material type `concrete` is assigned, an albedo corresponding to concrete, and a scattering coefficient set to 1, along with a gray RGB color, and an opacity value of 1.

#### Nacelle (Blade Support Structure)

The nacelle, which also serves as the blade support structure, is modeled as a short horizontal cylinder positioned atop the tower. Geometrically, the nacelle is centered at the top of the tower and extends along the  $x$ -axis. It has a radius of 3 m and a total length of 8 m [40], giving it a compact, streamlined shape. The front edge of the cylinder is slightly extended to resemble a conical tip, forming a shape that

visually resembles a rectangular body with a conical extension on one side, like a nacelle in a geared wind turbine.

The nacelle geometry is created by discretizing two circular cross-sections, one at the front and one at the back, each defined by 360 angular divisions to ensure a smooth appearance. The vertices of each circular cross-section is formed by generating a set of 360 vertices evenly spaced along its perimeter. Similar to the tower, the circular faces at both ends are constructed by connecting the center point of the circle with each pair of adjacent perimeter vertices, forming a fan of triangular faces. The faces of the body of the cylinder are then created by connecting adjacent vertices from the front and back circles to form rectangular patches, which are again split into triangular facets for compatibility with the triangular mesh system.

An additional feature of the nacelle geometry is the ability to adjust its orientation by rotating it around the  $z$ -axis. It can be manually adjusted in the script by rotating its front and back circular cross-sections, allowing the simulation of different yaw angles for shading analysis. A rotation matrix is applied to the nacelle vertices to allow reorienting the entire structure around the  $z$ -axis. This enables the nacelle to be manually rotated by any desired angle between  $0^\circ$  and  $360^\circ$ , simulating different yaw positions.

Once the nacelle geometry is constructed, it is assigned the material type `aluminum`, along with a dark blue RGB color, an opacity of 1, a scattering coefficient of 1, and an albedo value retrieved from the toolbox's Material Library.

### Blades

The wind turbine blades are modeled as three triangular surfaces, each extending radially outward from the front face of the nacelle. To define their geometry, the perimeter of the nacelle's front circle is divided into three equal sections, corresponding to the number of blades. From these, three pairs of vertices are selected to serve as the base edges of the blades. For each blade, the midpoint between the two base points is first calculated. This midpoint still lies on the circular edge of the nacelle. A unit vector is then constructed, pointing outward and perpendicular to the chord connecting the base points. This is achieved by taking the cross product between the chord direction and the nacelle's local normal vector, ensuring that the blade points away from the nacelle's front face. The tip of each blade is positioned by extending this outward-pointing normal vector from the midpoint by a fixed blade length of 115.5 m [40, 41].

The resulting triangle is formed by the two base points and the computed tip. Together, the three triangles represent the full rotor, with evenly spaced blades projecting forward from the nacelle. An additional feature of the script is the ability to rotate the blades collectively by specifying a starting angle, allowing their orientation to be manually adjusted in a similar manner to the nacelle for simulating different rotor positions. Each blade is assigned the material type `aluminum`, as defined in the toolbox's Material Library, a light gray RGB color, an opacity of 1, an albedo value corresponding to aluminum, and a scattering coefficient of 1. This is a modeling assumption, since in reality wind turbine blades are usually made of composite materials such as carbon and glass fiber [40]. However, the choice of aluminum was made for simplicity and to ensure compatibility with the toolbox's predefined material properties.

### Ground

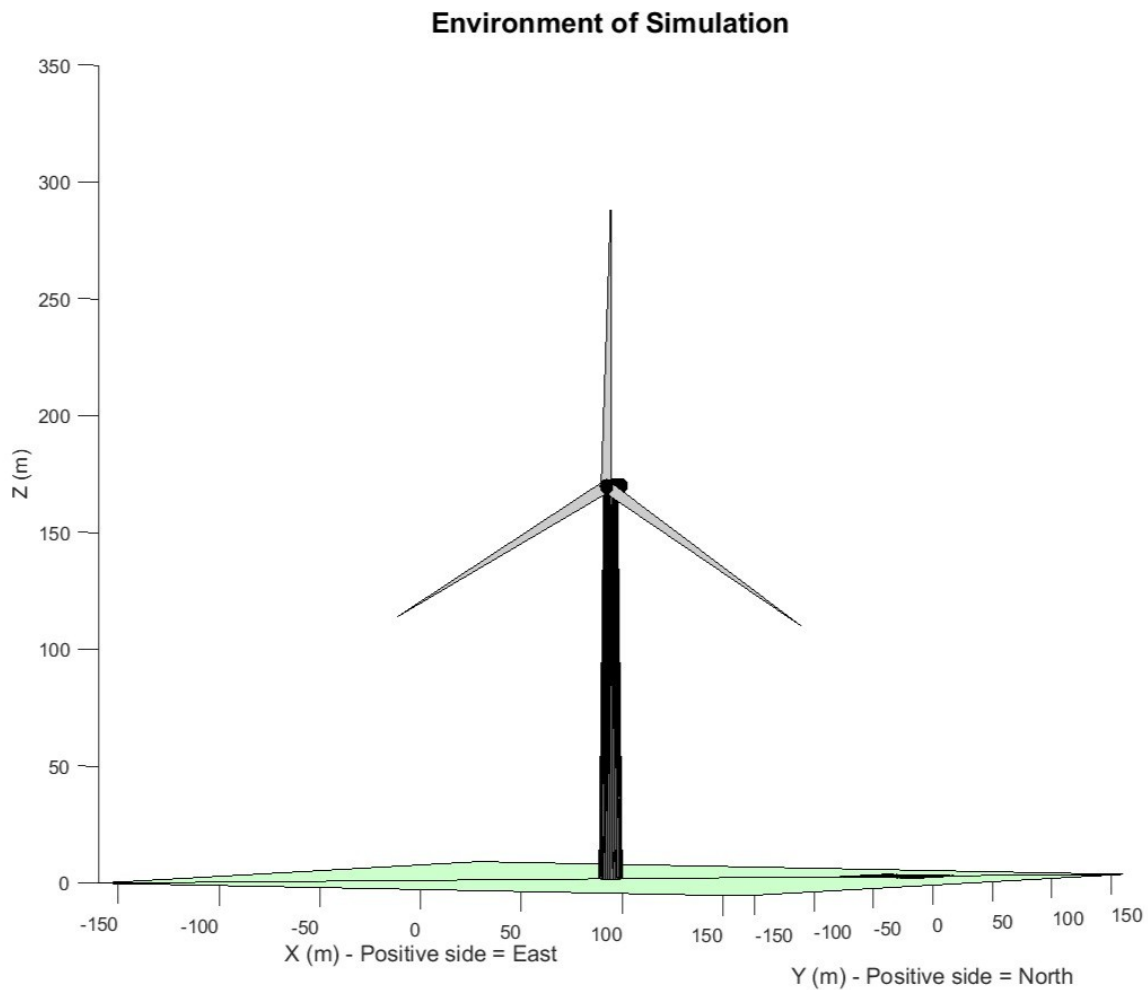
The ground is modeled as a flat rectangular plane centered at the origin in the  $x$ - $y$  plane. It spans  $\pm 160$  m in both directions, resulting in a square area of  $320 \times 320$  m<sup>2</sup>. The geometry is defined by four corner vertices at a height of  $z = 0$ , and the surface is formed using two triangular faces covering the full rectangular area. Once created, the ground surface is assigned the material type `ground`, with a green RGB color, an opacity of 0.2, a scattering coefficient of 1, and an albedo value consistent with that material.

### Visualization


This completes the construction of the wind turbine environment. Finally, all generated vertices and faces from each component are aggregated into the global matrices  $V$  and  $F$ , while the associated material and optical properties are stored in their respective arrays saved to a single `.mat` file that can be directly loaded into the PVMD Toolbox for simulation.

To visualize the constructed environment, the `plot_Environment` function, developed by Youri Blom, is used. It takes as input the global vertex and face matrices, along with visual properties such as RGB color and opacity, and renders the 3D scene using MATLAB's `patch` function. The environment and PV modules are plotted together, but this will be explained in the next section 2.1.3.


Since these simulations are meant to be relevant for offshore installations, the inspiration for all the the wind turbine dimensions is a 15 MW wind turbine from Vestas (V236-15.0 MW) [41]. The blade length is known from Vestas, but the dimensions of the tower and the nacelle design were assumed because they are not published by Vestas.




**Figure 2.1:** Simulation environment including a ground plane and one wind turbine.



(a) Turbine from the back.



(b) Wind turbine close-up for nacelle and blades.



(c) Nacelle front view.

**Figure 2.2:** Close-up views of a wind turbine nacelle and blades from various angles.

### 2.1.3. Adding PV Modules to the Environment

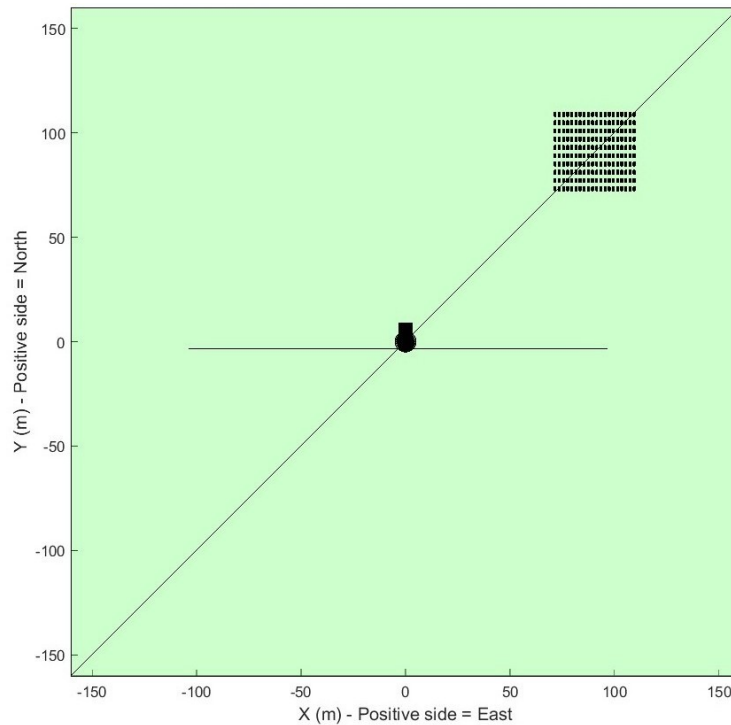
Once the environment is established and loaded in the simulation, PV modules are integrated into the scene. The incorporation of photovoltaic (PV) panels into the simulation environment is performed through a function that generates their full 3D geometry and integrates them with the pre-existing environment. Each PV panel is defined by a set of user-specified geometric and positional parameters that are provided through the simulation input structure.

For each simulation case, the user specifies the number of panels to be included and provides a list of 3D coordinates that determine their placement in the environment. These coordinates correspond to the bottom-left corner of each panel, which serves as the local origin for constructing the full module geometry. Each panel's orientation is defined by its tilt and azimuth angles. Its internal structure is determined by the number of solar cells arranged in rows and columns, along with cell dimensions, inter-cell and edge spacing, and module thickness.

Based on this input, the function constructs each panel individually by creating front (and optionally rear) solar cells and a module bulk (representing the frame). Each component is assigned appropriate optical and material properties. These properties define how each component interacts with light, including its reflectance, transmittance, and absorption characteristics, while material tags (e.g., aluminum) distinguish structural elements from active cell surfaces. The three surface elements of the panel are also geometrically transformed to reflect the specified tilt, azimuth, and location. The geometric transformation involves rotating the panel according to its tilt and azimuth angles to match its orientation in space and then translating the entire structure to the correct position in the environment based on its bottom-left corner coordinates.

This process is repeated iteratively for all specified panels. Each panel's geometry is appended to the existing environment mesh, resulting in a complete 3D scene that contains both static elements (e.g., ground, wind turbines) and dynamic elements (e.g., PV modules). The final combined environment is then used for subsequent ray-tracing simulations to assess irradiance and shading effects under various conditions.

A demonstration example is shown in Figure 2.3 and Figure 2.4. The simulation environment consists of a single wind turbine positioned at the center of the ground area. A total of 200 photovoltaic panels are placed around it with a fixed spacing of 30 cm between adjacent modules.



**Figure 2.3:** PV system of one wind turbine and 200-module incorporated in the environment on the ground.




Figure 2.4: PV system of 200 modules.


#### 2.1.4. PV Module Placement Strategy

In order to have as realistic simulations as possible, it is important to figure out the positioning of the photovoltaic (PV) modules relative to the wind turbines in the hypothetical hybrid farm.


In [36], onsite integrated hybrid PV-wind plants were evaluated. They were testing configurations where the minimum distance between turbines and PV arrays was either 50 m or 100 m. An example layout of their work can be seen below in Figure 2.5a. Similarly, in the work of López *et al.* [26], the co-existence of wind and solar power offshore is again discussed. The exact separation between the PV farm and the wind turbines in their work is not specified, however, the schematic, reproduced in Figure 2.5b, visually suggests a spacing of approximately 0.5 to 1 rotor diameter.

Tripp *et al.* [27] suggest a parameterized method for optimizing the layout of wind and solar hybrid power plants. They introduce a total of eleven parameters, five for turbine placement and six for solar placement. The most relevant parameters are the lateral and southern buffer zones between PV regions and turbines, which are bounded between 1 and 10, to give a lot of flexibility. With the minimum setback (turbine spacing) of 200 m, this would translate to buffer distances ranging from approximately 201 meters to 2010 m. Compared to the other sources, such large buffer distances would be unrealistic for practical hybrid power plant designs.

Based on these sources, in this thesis for the simulations it will be considered an acceptable clearance area between the wind turbine and the PV modules any distance ranging from 0.25 to 1.25 times the rotor diameter. As mentioned above in Section 2.1 with a rotor diameter of 231 m, the simulations include acceptable clearance distances starting from roughly 57.75 m to 288.75 m. This range permits more cautious placements near one full rotor diameter or just beyond in offshore scenarios, while also balancing the results from the literature, which demonstrate that distances as short as 50 meters are feasible.



(a) Layout configuration of an onshore hybrid solar-wind park, where  $d_1$ : the distance between the turbines in the main wind direction,  $d_2$ : distance perpendicular to the main wind direction, and  $d_3$ : the minimum distance between wind turbines and photovoltaic racks. Reproduced from [36].



(b) Combined floating wind and solar energy farm and schematic layout. Reproduced from [26].

**Figure 2.5:** Hybrid solar-wind park layouts in (a) onshore and (b) offshore environments.



## 2.2. Simulations Setup Assumptions

The simulations are a crucial and time-intensive part of this work. Multiple parameters define each simulation, consequently, careful decisions are required to ensure they effectively illustrate the results and provide a clear understanding of shading effects. The key parameters that characterize the simulations include the environment, the number of PV modules, system arrangement, the placement of the PV system in the environment, and the weather data based on the simulation period.

Before presenting the simulation results, it is essential to first outline all the assumptions that define the simulation framework. Once these are established, each aspect of the simulation will be explained in detail, followed by the corresponding results at every stage of the process.

The following assumptions are considered for the simulations:

### General Simulation Setup:

- The simulations take place in an onshore environment. The wind turbine and PV modules are placed on a square area representing ground.
- When building the environment, the scattering values for all faces of the surroundings (turbine and ground) are set equal to 1. This means that it is assumed that light hitting the surface is scattered uniformly in all directions (diffuse scattering) [42].
- When building the environment, the opacity values of the wind turbine faces are set to one and of the ground to 0.2.

### PV Module and System:

- The solar cell type is standard monocrystalline Passivated Emitter and Rear Contact (PERC).
- The PV module used is the RSM72-6-360-380M model [43]: 99.2 cm width, 195.6 cm height, 72 cells (12 rows, 6 columns), and rated power range of 360 W to 380 W, under standard test conditions (STC).
- The PV modules are mounted horizontally, with a tilt angle of  $0^\circ$  relative to the ground.
- Modules have a module mounting height of 0 m, no induced inter-row shading.
- Each module includes three bypass diodes.
- The system consists of non-reconfigurable modules.

### Weather Aspects:

- The meteorological data used correspond to the location of Delft, South Holland, for the year 2005. They include direct and diffuse irradiance components, measured or referenced for a surface tilted at  $31^\circ$ , with values provided on an hourly basis.
- SMARTS is used for spectral data, assuming clear sky conditions.
- The horizon reconstruction feature is not included.

### Electrical Aspects:

- No metallization losses are considered.
- A global maximum power point tracking (MPPT) method is applied.
- No heat dissipation effects are included.
- No reverse bias is applied.

Throughout this thesis, whenever a module is described as being placed at a specific position  $(x, y)$ , the coordinates refer to the distance from the center of the ground plane at  $(0, 0)$ , which is also the position of the wind turbine tower. These coordinates correspond to the bottom-left corner of the module, not its center.

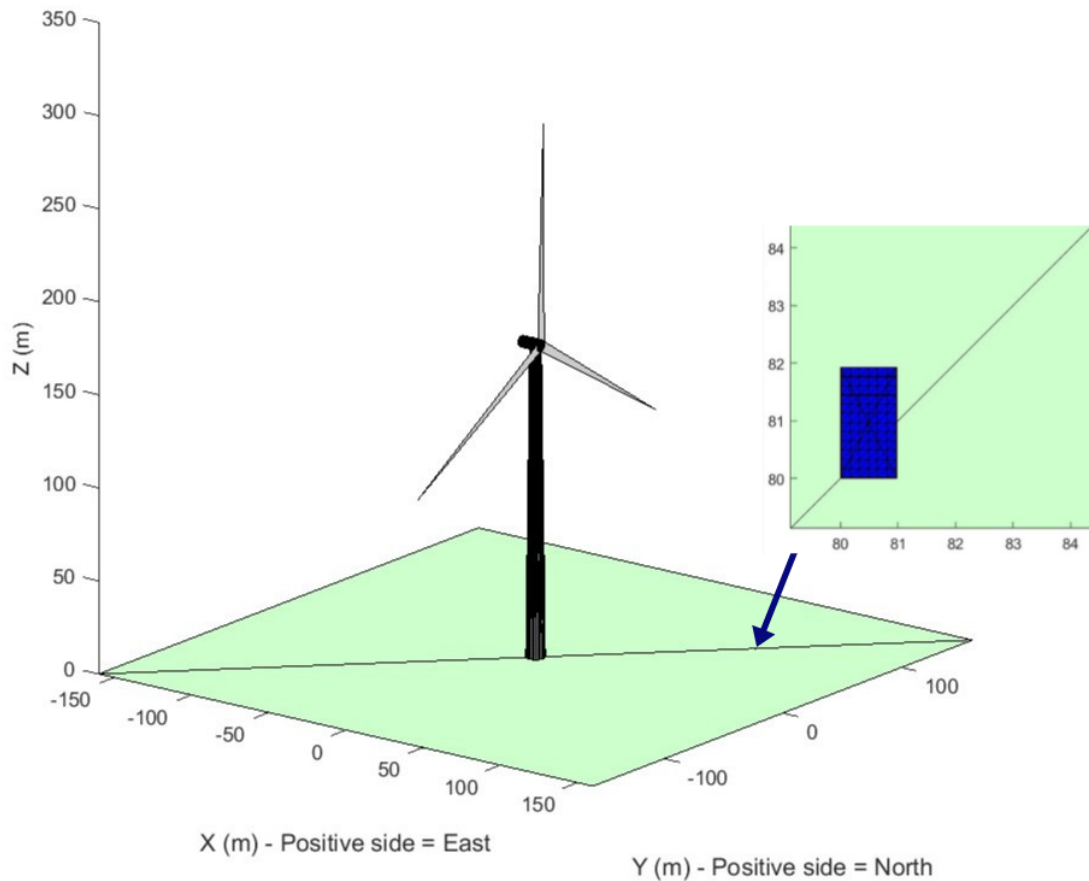
## 2.3. Simulations Setup Case Studies

To enable consistent analysis across different chapters, three reference simulation cases are defined and used throughout the thesis. Each case represents a distinct spatial configuration of PV modules around a wind turbine. These *Example Cases* used throughout the thesis to support different analyses, including shading evaluation, method validation, and result interpretation under various conditions.

### 2.3.1. Example Case A

This case is repeatedly used to validate the selective skydome refinement method and assess its impact on computational time and accuracy.

- The simulation area covers a square of  $300\text{ m} \times 300\text{ m}$ , ranging from  $-150\text{ m}$  to  $150\text{ m}$  in both  $x$  and  $y$  directions.
- A single PV module is placed at coordinates  $(80\text{ m}, 80\text{ m})$ .
- A single wind turbine is positioned at the origin  $(0\text{ m}, 0\text{ m})$ .
- The turbine faces East, rotated by  $0^\circ$  counterclockwise
- The turbine blades are fixed at a rotation angle of  $0^\circ$ .

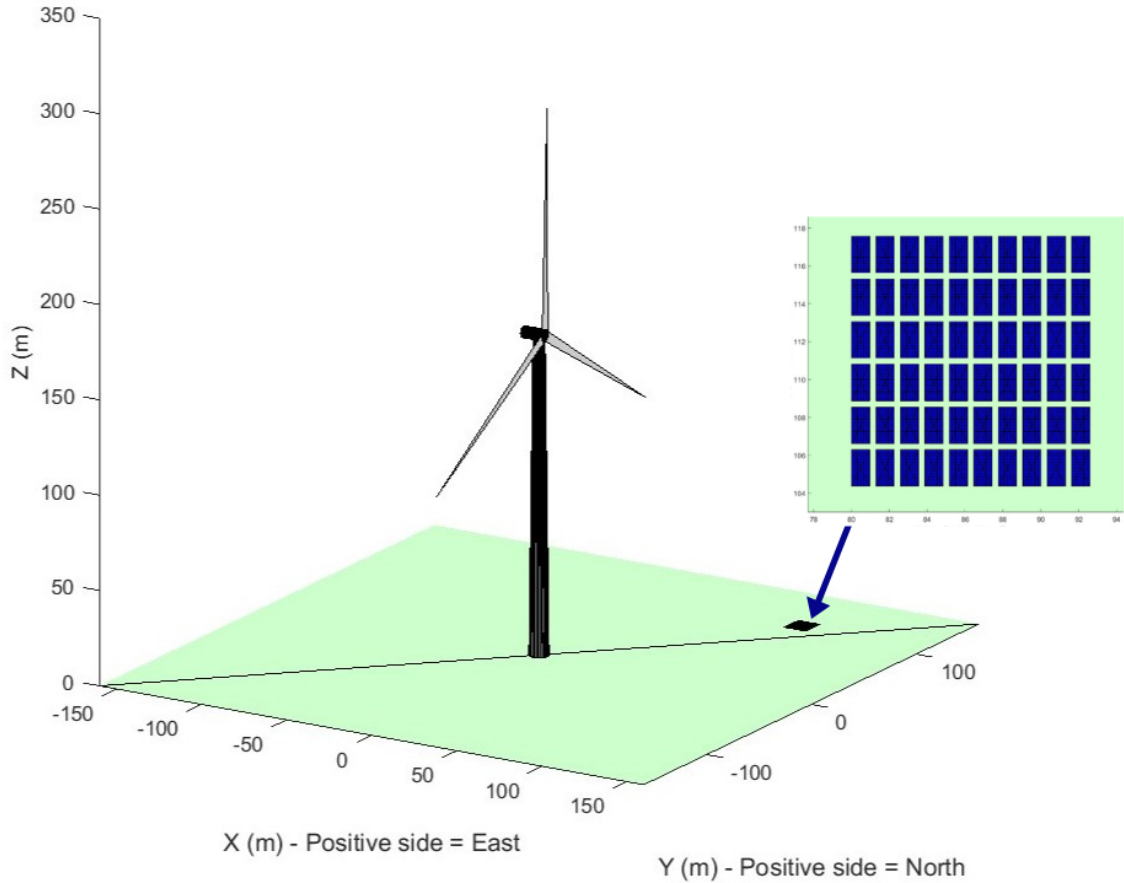


**Figure 2.6:** Visualization of the simulation environment of Example Case A

### 2.3.2. Example Case B

This case is used to analyze the effect of turbine shading on the spatial distribution of irradiance and the resulting energy output across the module array.

- The simulation area covers a square of  $300\text{ m} \times 300\text{ m}$ , ranging from  $-150\text{ m}$  to  $150\text{ m}$  in both  $x$  and  $y$  directions.
- A total of 60 PV modules are simulated, placed North-East of the turbine.
- A single wind turbine is positioned at the origin  $(0\text{ m}, 0\text{ m})$ .
- The turbine faces East, rotated by  $0^\circ$  counterclockwise
- The turbine blades are fixed at a rotation angle of  $0^\circ$ .



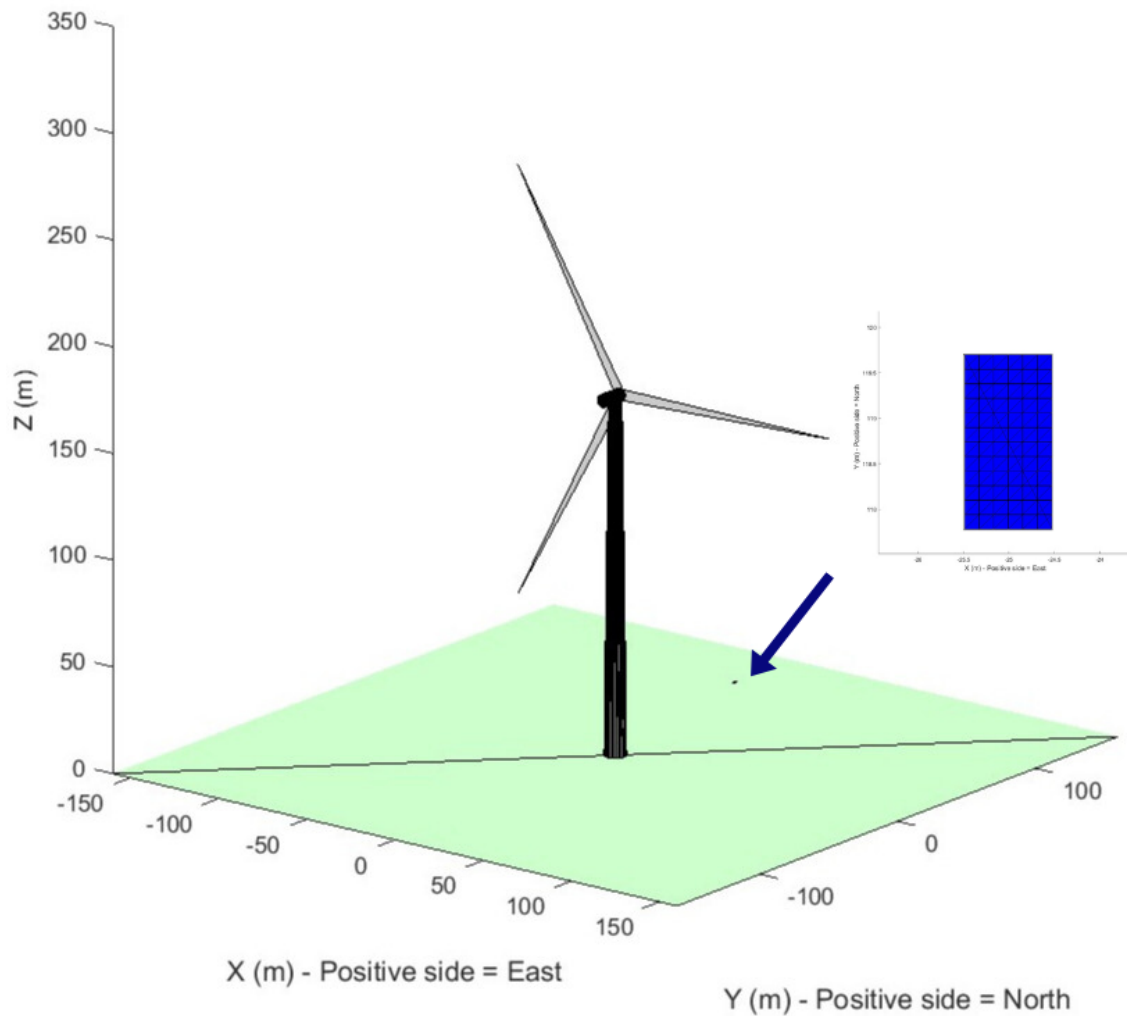
**Figure 2.7:** Visualization of the simulation environment of Example Case B

### 2.3.3. Example Case C

This case is used for assisting the representation of the selective refinement method.

- The simulation area covers a square of  $300\text{ m} \times 300\text{ m}$ , ranging from  $-150\text{ m}$  to  $150\text{ m}$  in both  $x$  and  $y$  directions.
- A single PV module is placed at coordinates  $(-25.496\text{ m}, 117.772\text{ m})$ .
- A single wind turbine is placed at the origin  $(0\text{ m}, 0\text{ m})$ .
- The turbine faces North, rotated by  $90^\circ$  counterclockwise

- The turbine blades are fixed at a rotation angle of  $270^\circ$ .

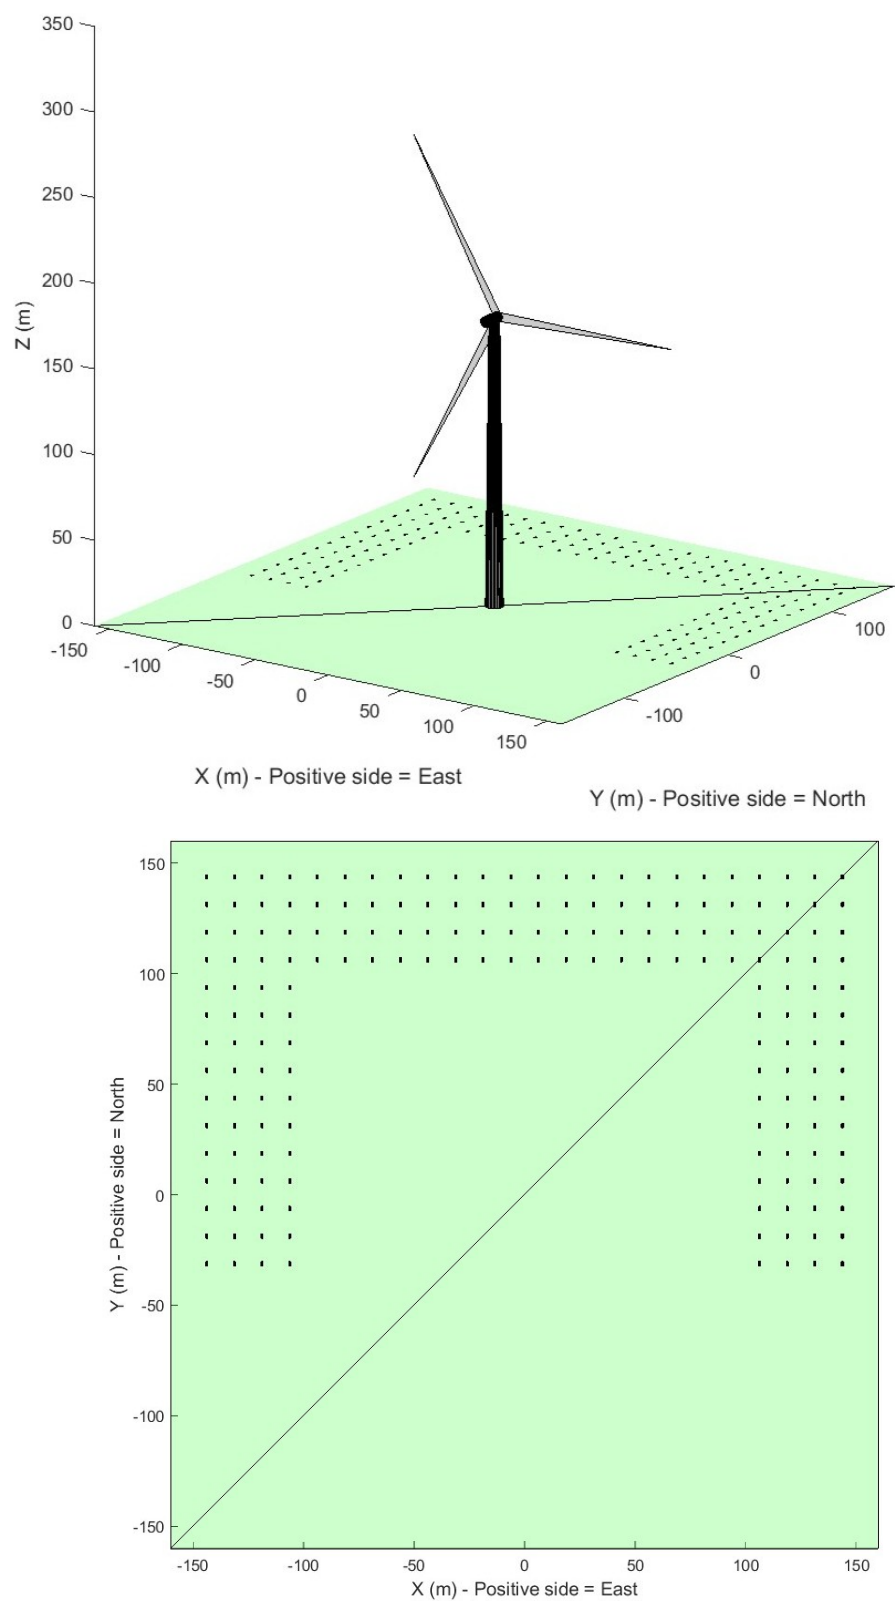


**Figure 2.8:** Visualization of the simulation environment of Example Case C

#### 2.3.4. Example Case D

This setup is mostly used to visualize and analyze shading patterns around a wind turbine by generating contour plots of shading intensity over the area.

- The simulation area covers a square of  $300\text{ m} \times 300\text{ m}$ , ranging from  $-150\text{ m}$  to  $150\text{ m}$  in both  $x$  and  $y$  directions.
- A total of 184 PV modules are simulated.
- A single wind turbine is placed at the origin  $(0\text{ m}, 0\text{ m})$ .
- The turbine faces North, rotated by  $90^\circ$  counterclockwise
- The turbine blades are fixed at a rotation angle of  $270^\circ$ .



**Figure 2.9:** Visualization of the simulation environment for Example Case D.

## 2.4. Conclusions

This chapter addressed Sub-question 1.: *How is a hybrid wind–solar simulation environment, including a detailed wind turbine, constructed and configured in the PVMD Toolbox?*

A comprehensive 3D environment was built in the PVMD Toolbox to replicate shading effects from a wind turbine in a hybrid PV–wind system. The three components of the wind turbine (i.e., tower, nacelle, and three blades), are modelled in the environment using a script-based technique that enables control over geometry, materials, and orientation. This method provides the flexibility required to depict intricate, large-scale objects that would otherwise be impossible to create using the graphical user interface (GUI). The PV modules are then integrated and that defines the final geometry of the whole simulation environment.

A wide range of simulation parameters and assumptions are explicitly documented, covering aspects such as the environment, PV system design, weather conditions and electrical behaviour. Three reference cases (A, B, and C) were introduced, each representing a different layout of PV modules around a wind turbine and are used throughout the thesis. These are established to support the analysis of shading patterns, showcase results, examine the simulation performance, and energy yield. Together, these elements provide a robust and flexible simulation framework for investigating the shading impact of a stationary wind turbine on PV systems in a hybrid energy environment.

# 3

## Shadow Modeling with the Existing PVMD Toolbox

This chapter addresses Sub-question: **2. How can the current version of the PVMD Toolbox model and analyze shading effects of a wind turbine on a PV module, and how does skydome discretization affect accuracy and performance?** It explains how the PVMD toolbox simulates the static shadow cast by a wind turbine or any environmental object on photovoltaic modules. Section 3.1 introduces the ray tracing and sensitivity analysis that the PVMD toolbox perform. This is followed by Section 3.3 with irradiance distribution and Section 3.4 results. The downside of the current sensitivity analysis of the toolbox is the computational burden it causes. This is examined in Section 3.5, and the solution will be addressed in the next chapter.

### 3.1. Ray Tracing and Sensitivity Maps

Once the full environment has been assembled and all PV modules have been added, the `CELL` function is called to compute the performance of the selected solar cell chemistry and its electrical behavior under reference conditions. Then, the simulation proceeds to evaluate the optical performance of the system at the cell level in the `MODULE_nonperiodic` function.

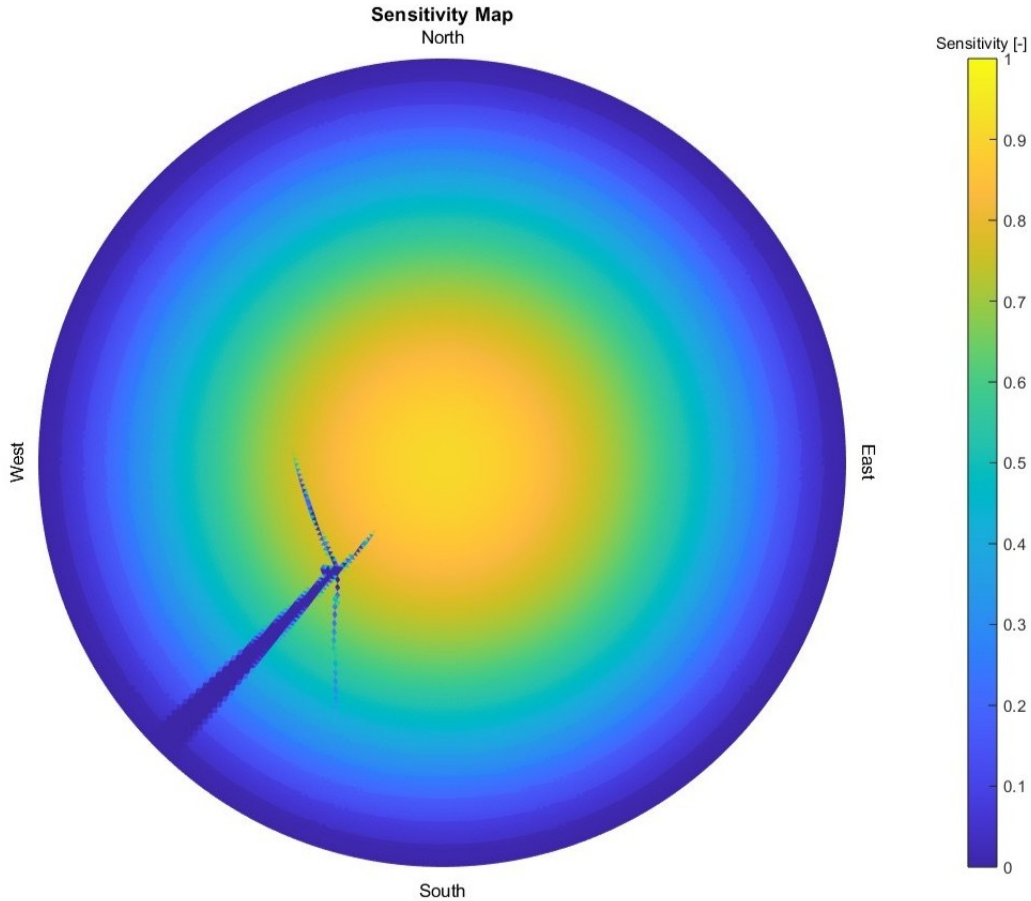
To achieve this, a method known as *backward ray tracing* technique is employed. In this process the entire object (PV) is illuminated from that direction, accounting for all possible reflections [44]. Then from each solar cell rays are traced in reverse: from the PV cell surface back to the sky. By tracing the rays, it is determined how much of the skydome is visible to each PV cell. The process is involved in the `BackwardTracer` function.

It is important to consider how the rays are defined in the context of backward ray tracing. Since the ray tracing is so closely related to the structure of the sky, the representation of the skydome is essential. A hemispherical skydome is constructed, which is then divided into triangular segments, each corresponding to a specific position and angle of the sun. Each triangular segment of the skydome represents a distinct direction in the sky. And furthermore, the number of rays that will be traced for each cell is equal to the number of the triangular sky faces.

During backward ray tracing, rays are cast from each solar cell towards the centroids of these triangular sky elements to simulate incoming sunlight from those directions. The coordinates of each ray correspond to the centroid of a triangular sky face. Each ray that is traced is evaluated to determine whether it is blocked by any object in the 3D environment, like the wind turbine, accounting for both direct and reflected light paths. If not, the ray returns back to the sky. If the ray escapes to the sky unobstructed, the triangle is considered fully visible and assigned a non-zero sensitivity value. On the contrary, if the ray is fully blocked by an object in the environment, the sensitivity for that triangle and cell is set to zero, indicating full shading in that direction. The `BackwardTracer` function also incorporates the optical properties of scene elements, such as reflectivity and scattering, when computing the absorbed

light fraction.

The output of the `BackwardTracer` is the sensitivity array of a surface to incident light as a function of the hemispherical angle of incidence [44]. The *Sensitivity Map (SM)* that results from it is a visual representation of how effectively the PV module captures light arriving from different angles. The SM is illustrated in a graph like Figure 3.1, which is a hemispherical representation of the visible sky to the module, like what a 360° camera would depict if located on a module laying on the ground. It is important to note that the sensitivity array and map are completely independent of time and does not correspond to a specific simulation period and it is provided per PV module. The example Sensitivity Map corresponds to the PV module of Example Case A, presented in 2.3.




**Figure 3.1:** Sensitivity map, of high accuracy, of one module of Example Case A.

The result provides a clear, scalar representation of how visible each sky direction is from the module surface. The visible sensitivity values assigned to each skydome triangle are obtained after averaging. First averaging the first absorption layer across all wavelengths, followed by an additional averaging across all cells in the module. The brighter colors indicate that the module is more sensitive to irradiation coming from that direction, and accordingly, the darker implies that for that particular sun angle, there is a shadow cast on the module.

Practically, the sensitivity output array is a four-dimensional one with dimensions that are:  $(sky\ faces) \times (cells) \times (absorption\ layers) \times (wavelengths)$ . In this array, each value is the fraction of light that is absorbed in the corresponding layer with respect to the total amount of light coming from that vertex. For instance, if the value of the sensitivity map for a certain vertex, a certain wavelength, a certain cell, and a certain layer is 0.5, it means that 50% of the irradiance from that vertex will be absorbed by that layer. It is not necessarily a 0 for shading or 1 for no shading, because it also takes into account the angle of incidence of the sun, like seen in Figure 3.2. If the sun comes in at a large angle, the irradiance is more spread out. The sensitivity map indicates the sensitivity of a surface to light from each direction.





**Figure 3.2:** Seasonal variation in solar beam concentration due to changing solar angle. Reproduced from D. V. Schroeder, licensed under CC BY-NC 3.0 [45].

It is worth explaining the meaning of the layers in the sensitivity array. In the case of the crystalline silicon (c-Si) cell used in this study, the physical layer structure is: *air-c-Si-air*. However in the sensitivity array three irradiance absorption matrices are included :

1. Module absorption: total absorption within the complete module and is calculated as  $1 - R - T$ , which contributes to thermal heating of the module.
2. Absorption in the active-absorption layer: irradiance absorbed specifically by the c-Si absorber layer, which relates to photocurrent generation, and thus in the energy yield.
3. Total incident irradiance: A hypothetical absorptance of 1 is assumed, which allows to find the theoretical upper limit of conversion efficiency.

From all the above, it is evident that the correct calculation of the sensitivity of a module is the most important output when working on shadow modelling. Naturally the accuracy of the skydome and consequently the ray tracing are essential, so it will be explained in detail how this is modeled in the toolbox.

### 3.2. Importance of Skydome Discretization

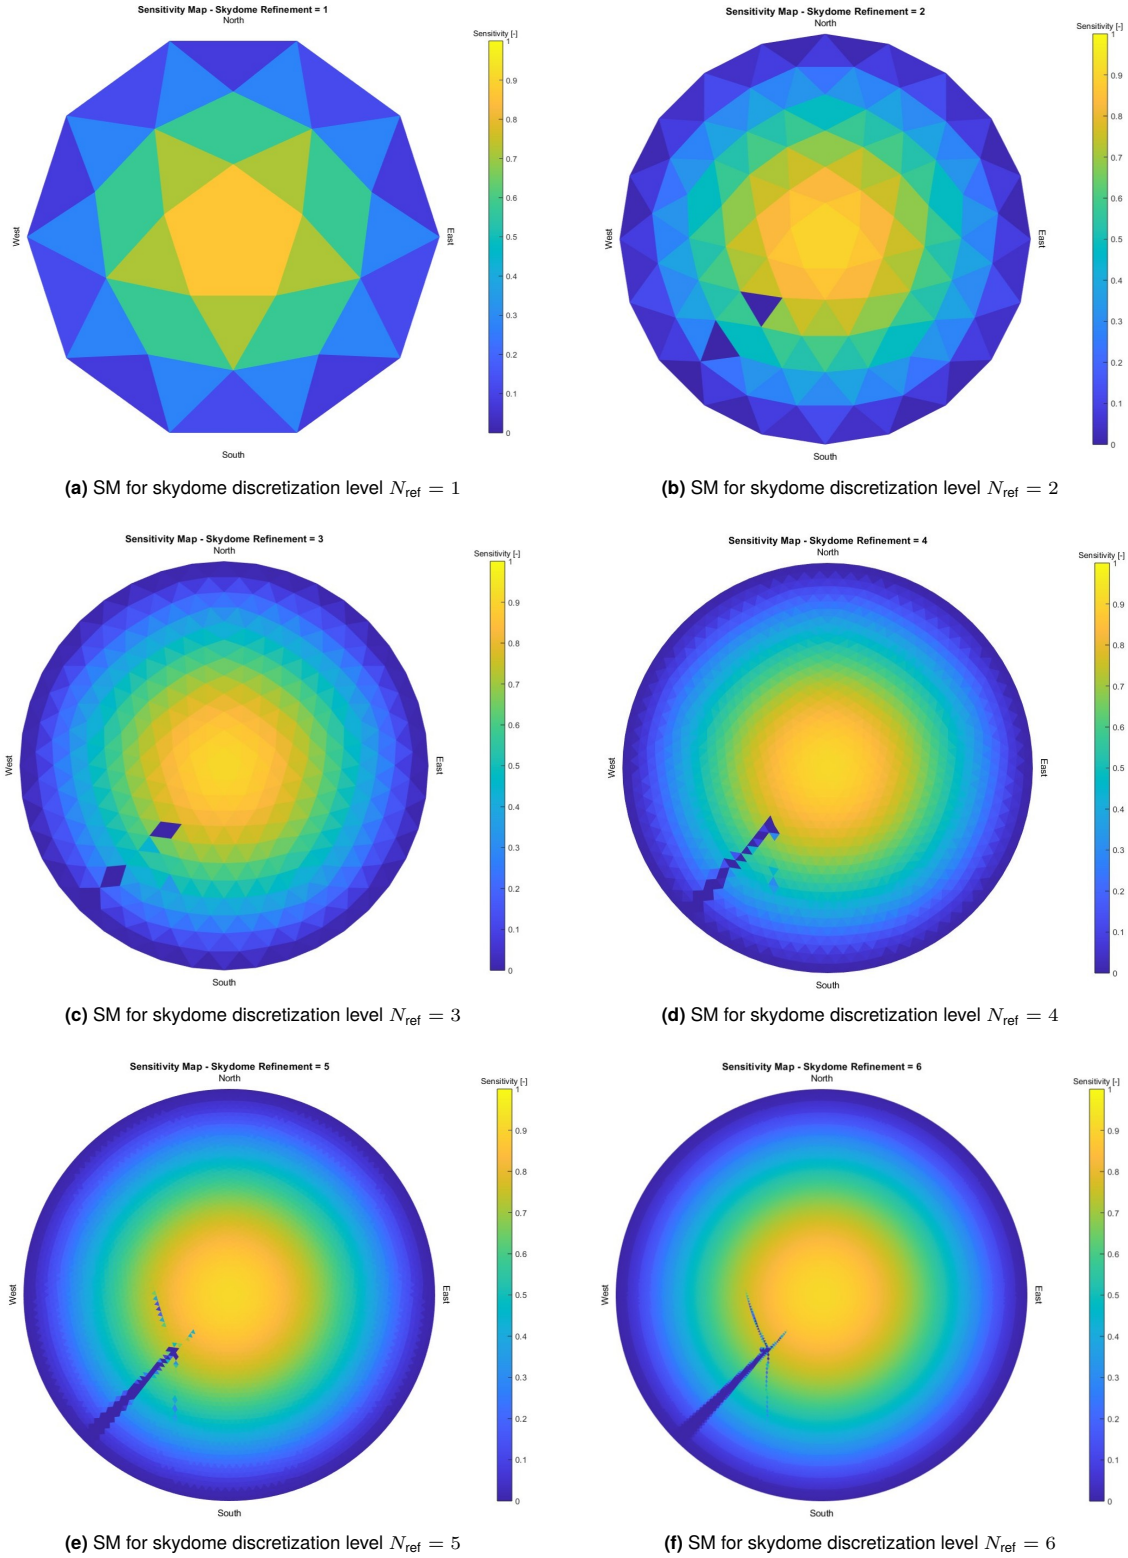
To build the skydome, the `icohemisphere` function is used. This function starts from a base icosahedron and recursively subdivides its triangular faces to form a geodesic dome. The first input for this function is the refinement level  $N_{\text{ref}}$ , which represents the number of times each face is subdivided. Starting with the icosahedron, each triangle is divided into 4 smaller ones recursively. The second input indicates whether the bottom half of the sphere should be removed to create a hemisphere. This is always selected and thus the total number of triangle faces of the hemispherical skydome is given by  $n = 10 \times 4^{N_{\text{ref}}}$ . The resulting triangles have a similar, though not identical, size and are rescaled such that their total area equals  $2\pi$ , representing the projected solid angle of the hemisphere.

The number of sky facets for each discretization level is:

- $N_{\text{ref}} = 1 \rightarrow 40$  facets
- $N_{\text{ref}} = 2 \rightarrow 160$  facets
- $N_{\text{ref}} = 3 \rightarrow 640$  facets
- $N_{\text{ref}} = 4 \rightarrow 2,560$  facets
- $N_{\text{ref}} = 5 \rightarrow 10,240$  facets
- $N_{\text{ref}} = 6 \rightarrow 40,960$  facets

Increasing the skydome discretization levels means that more rays are traced back and the shape of the objects in the environment will be depicted in greater detail. To illustrate the effect of skydome refinement, the simulation setup of the Example Case A is employed for six different refinement levels, with  $N_{\text{ref}}$  varying from 1 to 6. The resulting sensitivity maps can be seen in Figure 3.3. They demonstrate how increasing the angular resolution of the sky hemisphere enhances the detail and accuracy of the

shading representation. Naturally, this affects the final energy yield calculated. So, higher refinement would mean that the whole simulation then is more accurate.



**Figure 3.3:** Sensitivity Maps for increasing skydome refinement levels from 1 to 6.

### 3.3. Irradiance Distribution


After calculating the sensitivity array, or in other words the shading analysis of each module, the next step of the simulation is to incorporate specific weather data. This part of the simulation, **WEATHER**, is responsible for determining the amount of energy absorbed by the solar cells before electrical conversion. Consequently, in the **ELECTRIC**, the DC energy output of the module and the overall system can be estimated.

Here, the incident irradiance  $Irradiance(t)$  reaching every cell of a PV module is calculated as a function of time  $t$ . It is computed as the sum of directional irradiance contributions from the skydome weighted by the corresponding sensitivity values [44]:

$$Irradiance(t) = \sum_{i=1}^n skymap(i, t) \times sensitivity(i) \quad (3.1)$$


The sensitivity map is already explained in 3.1. The *skymap* is derived from the Perez model [46] and it represents the radiance distribution over the sky hemisphere in units of  $W m^{-2} sr^{-1}$ . This model takes as inputs the measured DHI and DNI from the weather data of the location and calculates the luminance distribution of the circumsolar and the diffuse light [44]. At each time step, every cell of the module is exposed to the entire skydome, and a radiance value is assigned to it. Furthermore, as explained before, to each cell of the module corresponds a sensitivity value for every sky triangle, which quantifies how strongly light from that direction contributes to the cell's irradiance. If a sky face is blocked by an object in the environment, this value is effectively reduced. Hence, when the luminance of this direction is multiplied by this low sensitivity value, the resulting irradiance *Irr* value is low as well.

To understand better the outputs of the **WEATHER** function and how it is related to the sensitivity analysis, the simulation setup of Example Case B is used. The 60 modules arranged in 6 rows of 10. The simulation uses hourly weather data for a single day, June 1<sup>st</sup>, and a low sky resolution setting of  $N_{ref} = 3$ . This date is chosen as it provides extended daylight hours and a high likelihood of clear sky conditions. Since the turbine is facing East and the PV modules are placed North-East of the turbine, it is then reasonable to understand that the 60-module PV system is going to experience shading at the afternoon hours, as the sun travels towards West. Figures 3.4 and 3.5 present the *Irr* over the 60 module system at 14:00 and at 15:00 pm.



**Figure 3.4:** Irradiance distribution on the 60-module PV system at 14:00, on June 1<sup>st</sup>.






**Figure 3.5:** Irradiance distribution on the 60-module PV system at 15:00, on June 1<sup>st</sup>.

At 14:00 pm, a distinct shadow from the turbine is evident that lowers the irradiance on some PV modules, as seen in the color pattern. However, in the next time step at 15:00 pm, the indication of a shadow has completely disappeared. The modules exhibit a uniform, slightly lower irradiance level due to the shift in irradiance from 14:00 to 15:00 o'clock.

Apart from the irradiance  $Irr$  equally important outputs are the absorbed irradiance  $A$  [ $W/m^2$ ] and the implied photocurrent  $J$  [ $mA/cm^2$ ]. Linking to the Sensitivity Map, the  $A$  is calculated using the first layer of the SM, representing the total module absorption  $1 - R - T$ . The estimation of the  $J$  is based on the second layer of SM, so the absorption in the active/absorption layer. To showcase how the apparent shade affects these quantities, modules 15, 22, and 46 are chosen to be examined in detail. The selected modules are shown in Figure 3.6.



**(a)** Three selected modules with different shadings at 14:00.

**(b)** Three selected modules with different shadings at 15:00.

**Figure 3.6:** Three chosen modules with different shadings at two consecutive hours on June 1<sup>st</sup>.

The effect of shading at a given time step is clearly visible when plotting the absorbed irradiance  $A$  [ $\text{W m}^{-2}$ ] and the implied photocurrent  $J$  [ $\text{mA cm}^{-2}$ ]. The absorbed irradiance and photocurrent profiles of them are depicted below, in Figure 3.7. The variations displayed in Figures 3.7a, 3.7b and 3.7c highlight the impact of shading on PV performance and the importance of accurately capturing shading effects in simulations. From the three graphs, it can be seen that abrupt drops in irradiance and photocurrent can be observed for shaded modules.

In detail, module 15, which never appears to be shaded, follows a typical irradiance pattern, as shown in Figure 3.7a. Its irradiance and corresponding photocurrent gradually increase throughout the morning, peak around midday, and then decrease gradually in the afternoon, reflecting the expected sun path and irradiance behavior. In contrast, module 22 is fully covered by shade at 14:00 pm as seen in Figure 3.4. This is leading to a significant irradiance drop in irradiance and produced photocurrent in Figure 3.7b. However, even when a module is in complete shadow at some time frame, its irradiance does not fall to zero due to the contribution of diffuse irradiance. Lastly, taking a look at module 46 that is partially shaded at the same time, we see a consistency in how the  $A$  [ $\text{W/m}^2$ ] and  $J$  [ $\text{mA/cm}^2$ ] are formed, as seen in Figure 3.7c.

### 3.4. Energy Output

One step further, the DC power output (DCP) of each module can be found. For this calculation, the sensitivity map again is very important. As mentioned earlier in 3.3, the absorptance and the photocurrent calculations are done using different layers of the SM.


The absorptance ( $A$ ) is multiplied by the active module area to yield the total absorbed power. This is relevant to the optical-to-electrical efficiency. It can be then compared against the maximum electrical output  $P_{\text{mpp}}$  and evaluate performance losses brought on by shading or mismatch.

The photocurrent ( $J$ ) is used to generate IV curves per module, under given temperature. Then  $P_{\text{mpp}}$  is computed from the simulated IV curve using the photocurrent  $J$ , cell temperature, and diode parameters. It is extracted at each timestep. The energy yield of each module is calculated by summing the DC power output ( $P_{\text{mpp}}$ ) over all time steps and multiplying by the time step duration ( $\Delta t$ ) in hours. This way the DC energy yield is obtained in watt-hours (Wh) for every kind of weather data, and here  $\Delta t = 1$  for hourly data.

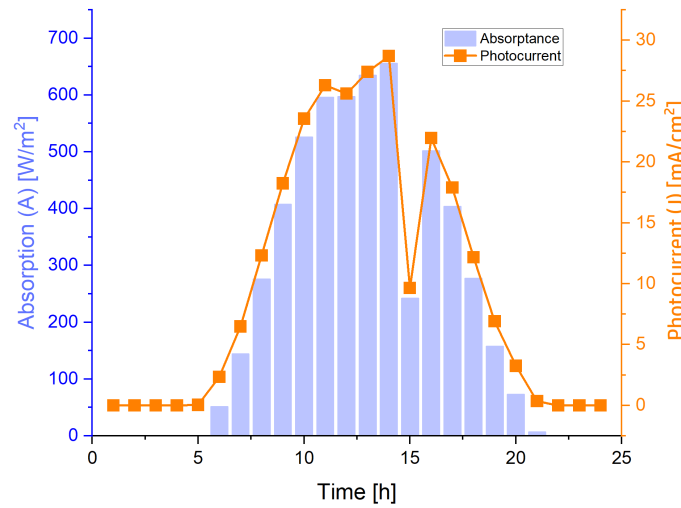
$$\text{DC Energy Output (Wh)} = \sum P_{\text{mpp}} \cdot \Delta t \quad (3.2)$$

The effect of the absorbed irradiance on the DC power produced by each module is changing accordingly, as shown in Figure 3.8.

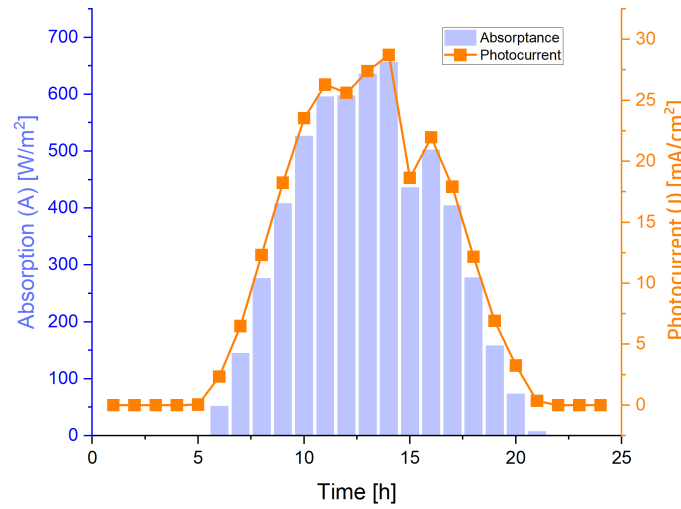
To conclude, from the results, it is seen that the shadow cast by the wind turbine for a time interval is present, but for the next and previous ones, it is nonexistent. Some modules, if not shaded at 15:00 pm, appear never to be shaded. In reality, during a day a shade from an object moves because of the sun path it goes slowly and smoothly. Thus, at some point in a day, every module would experience shading. As a result, irradiance is overestimated, as some modules that should be partially shaded at certain times appear to receive full sunlight instead. These limitations emphasize the need for a finer time resolution to improve the accuracy of the final calculated energy yield.



(a) Absorbed irradiance and photocurrent for completely unshaded module 15.

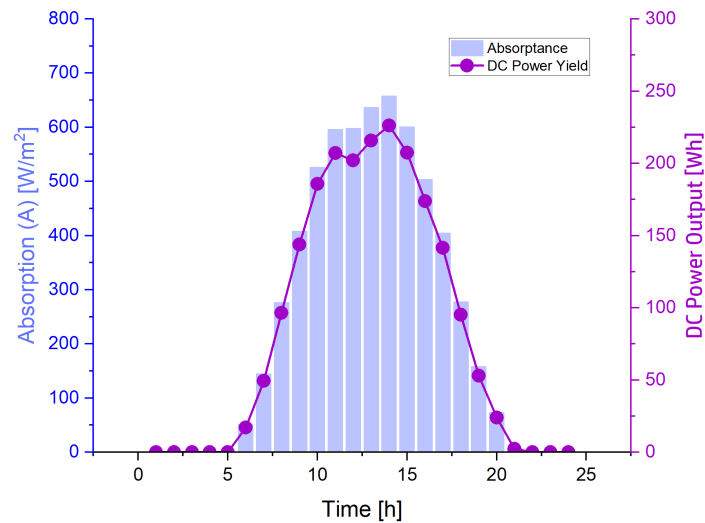


(b) Absorbed irradiance and photocurrent for completely shaded module 22.




(c) Absorbed irradiance and photocurrent for partially shaded module 46.


Figure 3.7: Comparison of irradiance and photocurrent for modules 22, 59, and 45. throughout the 1<sup>st</sup> of June



(a) Absorbed irradiance and DC power output for completely unshaded module 15.



(b) Absorbed irradiance and DC power output for completely shaded module 22.



(c) Absorbed irradiance and DC power output for partially shaded module 46.

**Figure 3.8:** Comparison of absorbed irradiance and DC power output for modules 15, 22, and 46 throughout the 1<sup>st</sup> of June.

### 3.5. Tradeoff of Shading Sensitivity Analysis

The previous Section 3.2 analyzed the concept of skydome discretization and the importance of a fine mesh. Figure 3.3 clearly shows that the skydome refinement can have a great impact on the visibility of the object and thus is critical for the shading analysis. This raises the question of why not always opting for a high skydome discretization without further consideration. There is a big tradeoff related to the increase of the sky mesh, and that is the computational burden. This section presents a computational time analysis to evaluate this trade-off for the simulation setup of Example Case A.

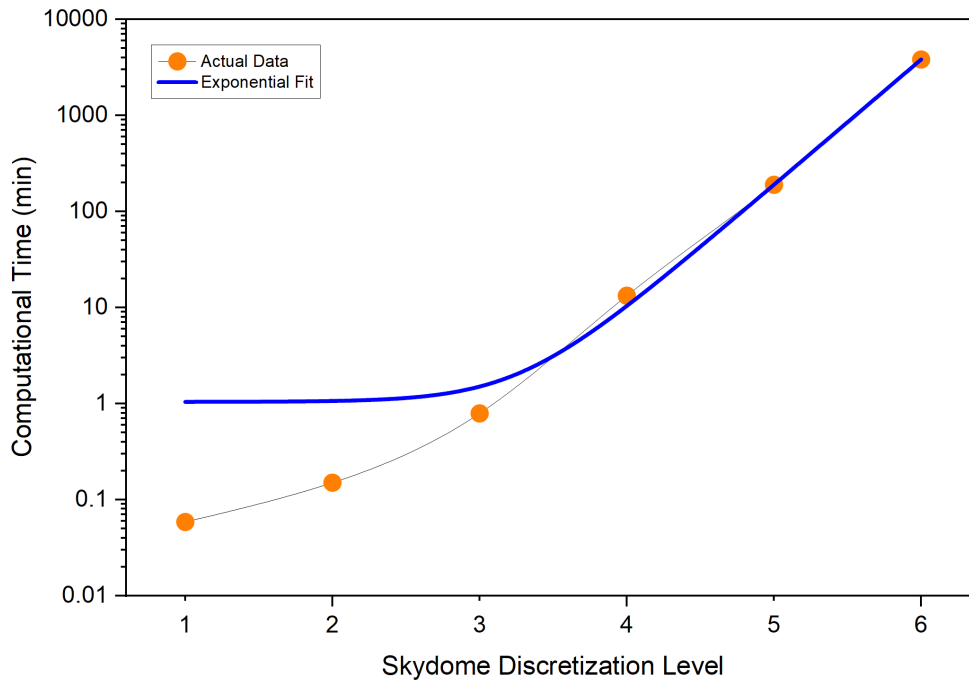
#### 3.5.1. Computational Burden Analysis

Six simulations were run for the single PV module of Example Case A, altering the skydome discretization  $N_{\text{ref}} = 1, 2, 3, 4, 5, 6$ . For each of these, the period during which the MODULE component was executed, i.e., the sensitivity analysis, was noted. The reported times, in minutes, are listed in Table 3.1 below, from which it is evident that the computational time increases immensely with increasing the skydome discretization level  $N_{\text{ref}}$ .

**Table 3.1:** Computational time in minutes for each level of skydome discretization.

Skydome Discretization Level	Computational Time (min)
1	0.0586
2	0.1501
3	0.7914
4	13.2805
5	189.9166 $\approx$ 3.2 (hours)
6	3826.0115 $\approx$ 63.8 (hours) $\approx$ 2.65 (days)

The computational time as a function of skydome refinement level is plotted in the following Figure 3.9 on a logarithmic scale, clearly illustrating the increasing trend, as well as a fitted curve.



**Figure 3.9:** Increasing computational time with higher skydome discretization levels, fitted with an exponential model.



The data of computational time as a function of skydome refinement level was fitted using a non-linear exponential model  $y = 1.04 + 5.59 \times 10^{-5} \cdot \exp(3.01x)$ , which is explained in detail in Appendix B.1.

To quantify the percentage increase in computational time between different discretization levels, we define Eq. (3.3). The computational cost  $T$  is evaluated as a function of the refinement level, from level  $i$  to level  $j = i + 1$ :

$$\Delta T_{i \rightarrow j} = \frac{T_{N_{\text{ref}}=i+1}}{T_{N_{\text{ref}}=i}} \quad (3.3)$$

To put into perspective, the increase in computational time between refinement levels is summarized below:

- $N_{\text{ref}} = 1 \rightarrow 2$ : 2.56× increase
- $N_{\text{ref}} = 2 \rightarrow 3$ : 5.27× increase
- $N_{\text{ref}} = 3 \rightarrow 4$ : 17.0× increase
- $N_{\text{ref}} = 4 \rightarrow 5$ : 14.30× increase
- $N_{\text{ref}} = 5 \rightarrow 6$ : 20.15× increase

It should be kept in mind that the computational time estimates of a simulation are not absolute. It depends not only on the code itself but also on external hardware and system-level factors, such as the computer's CPU and available RAM. Therefore, the reported values should be considered as estimates specific to this simulation setup and execution environment and may vary within a margin of error.

### 3.5.2. DC Energy Yield Analysis

It is imperative to examine how the DC energy yield is affected by the increased accuracy of the sensitivity analysis. To demonstrate that, the example simulation was carried out for two days, June 1<sup>st</sup> and January 1<sup>st</sup>, and the yield was calculated. The results are presented in the Table 3.2 below. Despite the variations in sun path and irradiance between these two dates, the general pattern is the same.

**Table 3.2:** Computational time in minutes for each level of skydome discretization.

Skydome Discretization Level	DC Energy Yield (kWh) - 1 <sup>st</sup> Jun	DC Energy Yield (kWh) - 1 <sup>st</sup> Jan
1	2.1617	0.1495
2	1.8937	0.1383
3	1.9086	0.1383
4	1.8937	0.1381
5	1.8964	0.1375
6	1.8949	0.1371

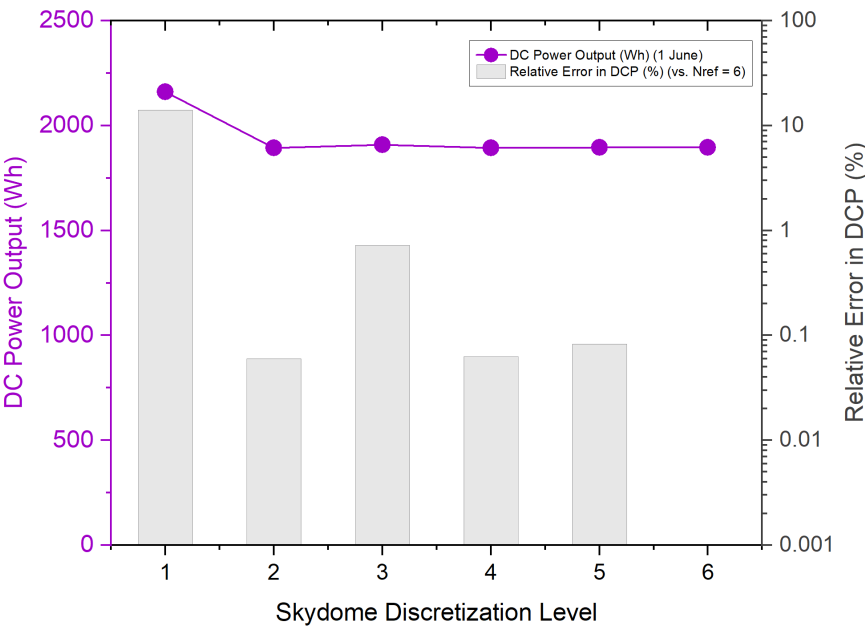
To quantify the accuracy loss due to a coarse sky discretization, the relative error between refinement levels  $N_{\text{ref}} = i$  and  $N_{\text{ref}} = j$  is defined. The following expression represents the comparison between any two levels:

$$\epsilon_{\text{rel}}(i - j) = \frac{\text{DC\_EY}(j) - \text{DC\_EY}(i)}{\text{DC\_EY}(i)} \times 100\% \quad (3.4)$$


where  $\text{DC\_EY}(k)$  is the calculated DC energy yield at refinement level  $N_{\text{ref}} = k$ . The more accurate (finely discretized) case is represented by  $i$ , and  $j$  the coarser or alternative one. A positive value of  $\epsilon_{\text{rel}}(i - j)$  indicates an overestimation of the yield due to insufficient discretization.

Here, comparisons are made between the most refined case ( $N_{\text{ref}} = 6$ ) and all coarser levels ( $N_{\text{ref}} = 1, 2, 3, 4, 5$ ), using  $i = 6$  as the reference. The results for the simulation for the two dates are plotted in Figure 3.10. The graph includes the DC energy output (Wh) for that day shown in the purple circle markers. The bars in the graph represent the relative error (%) between the DC yield at a discretization

level with respect to the most refined case ( $N_{\text{ref}} = 6$ ). This is done for the apparent reason that the most accurate result is the one with the highest refinement.



(a) June 1<sup>st</sup>



(b) January 1<sup>st</sup>

**Figure 3.10:** DC power output and relative error across skydome discretization levels. Purple data points indicate daily DC energy output (Wh) and gray bars are the errors of each value relative to  $N_{\text{ref}} = 6$ .

As the skydome discretization becomes finer, the relative error in DC energy yield tends to diminish, reflecting convergence toward the most accurate result. To be more precise, for the case of the day in June, for  $N_{\text{ref}} = 1$  the yield is overestimated by approximately 14%, and for January by 9%. For both simulations, it is seen that from  $N_{\text{ref}} = 2$  onward, the relative changes in yield remain below 1%. For June, for the cases of  $N_{\text{ref}} = 2$  and  $N_{\text{ref}} = 4$  the estimated yield is lower than the  $N_{\text{ref}} = 6$  one. For June, when increasing the sky mesh  $N_{\text{ref}} = 5$  to  $N_{\text{ref}} = 6$  the yield is only decreasing by approximately 0.08%, and for January that is 0.3%. The exact values can be seen in the Appendix C.1.

Even though the difference from  $N_{\text{ref}} = 5$  to  $N_{\text{ref}} = 6$  may seem small, it is important to discretize to a high level. When simulating entire PV farms, even minor local errors can accumulate, potentially causing significant deviations in the final results. For large-scale or long-term evaluations, it is therefore justified to use the highest available refinement level to ensure maximum accuracy.

### 3.6. Conclusions

This chapter addressed Sub-question 2.: *How can the current version of the PVMD Toolbox model and analyze shading effects of a wind turbine on a PV module, and how does skydome discretization affect accuracy and performance?*

It investigated the ability of the PVMD Toolbox to simulate the shading effects from a wind turbine on photovoltaic modules, using backward ray tracing and sensitivity analysis. The shading modeling process was analyzed step by step, from generating the sensitivity map, to computing irradiance, and finally the DC energy output. It is explained how the accuracy of the entire simulation chain is largely dependent on the angular resolution of the skydome mesh, which is determined by the discretization level  $N_{\text{ref}}$ .

A detailed computational analysis was conducted to evaluate the trade-off between skydome discretization level and simulation performance. The results in this section presented that increasing the refinement level  $N_{\text{ref}}$  enhances the angular resolution of the sensitivity map, thereby improving shading representation and energy yield estimation. However, this comes at a steep computational cost: simulation time rises from under a minute at  $N_{\text{ref}} = 1$  to over 2.5 days at  $N_{\text{ref}} = 6$  for a single module. This level of computational burden makes high-resolution simulations impractical. Especially for larger, system-level cases such as evaluating full PV systems or hybrid wind-solar parks, which is the broader goal of this thesis.

These findings demonstrate that while the PVMD Toolbox is indeed capable of accurately modeling static shading effects cast by a wind turbine, through backward ray tracing and skydome-based sensitivity analysis, this accuracy comes at an extremely high computational cost. It is therefore imperative to move away from uniform high-resolution methods in favor of more focused, effective refinement techniques. The next chapter introduces a novel selective refinement approach that aims to preserve sensitivity map precision while drastically reducing computational effort.

# 4

## Selective Skydome Refinement Method

Since the concern that high sensitivity map accuracy leads to extremely high computational time was addressed in the previous chapter, this chapter now aims to answer the third and most important Sub-question: **3. What enhancements can be introduced in the PVMD toolbox to improve the accuracy and efficiency of the shading analysis of a PV module with respect to its surroundings?** In this chapter, the concept and motivation behind the selective refinement method are first introduced in Section 4.1. A flowchart is included to facilitate the complexity of the logic. Then follows Section 4.2, which describes the full implementation of this methodology in detail.

### 4.1. Conceptual Overview of Selective Refinement

It is critical for the use of the toolbox to be able to produce high-quality sensitivity maps that accurately represent shading from the environment the modules are placed in, within a reasonable computational time. As shown earlier, computing the sensitivity arrays with a highly refined skydome mesh with the current toolbox version is very computationally heavy.

The way the sensitivity array is calculated cannot be changed within the toolbox, and from a very early stage, the skydome refinement level is selected and dictates the rest of the simulation. Thus, the proposed idea needs to modify the existing method for calculating the sensitivity map rather than inventing a completely new one.

Since both the PV modules and surrounding objects, such as wind turbines, are static, the shading direction can be geometrically determined. For example, if a module is placed south of a turbine, the shadow will appear on the north side of the sensitivity map. Similarly, placing the module east of the turbine will result in shading from the west. This allows the origin of shading to be predicted based on relative positioning (see sensitivity map 3.1). As a result, it is not necessary to finely refine the entire skydome and calculate the sensitivity of each of the 40,960 sky faces for  $N_{\text{ref}} = 6$ , since most of them are completely unshaded and will always remain so. On the contrary, only a specific area of the skymap can be selectively refined to compute the exact sensitivity values. While the rest can be left untouched as they are unaffected by the obstacle's shadow. The gain of the method is that it focuses computational effort only where shading is detected or suspected.

The flowchart of the method is given in Figure 4.1 with the following brief explanation. The process begins with an initial backward ray tracing simulation at low skydome refinement level. From this initial skydome, the triangle sky faces associated with shading or low sensitivity are identified. To ensure that no partially shaded regions are missed, the selected triangles are later expanded by adding a buffer zone of neighboring triangles. Then, the remaining faces, which are not selected for recalculation, retain the initial values. Finally, both sets are combined to produce a uniform sensitivity array with selectively refined resolution dimensions.

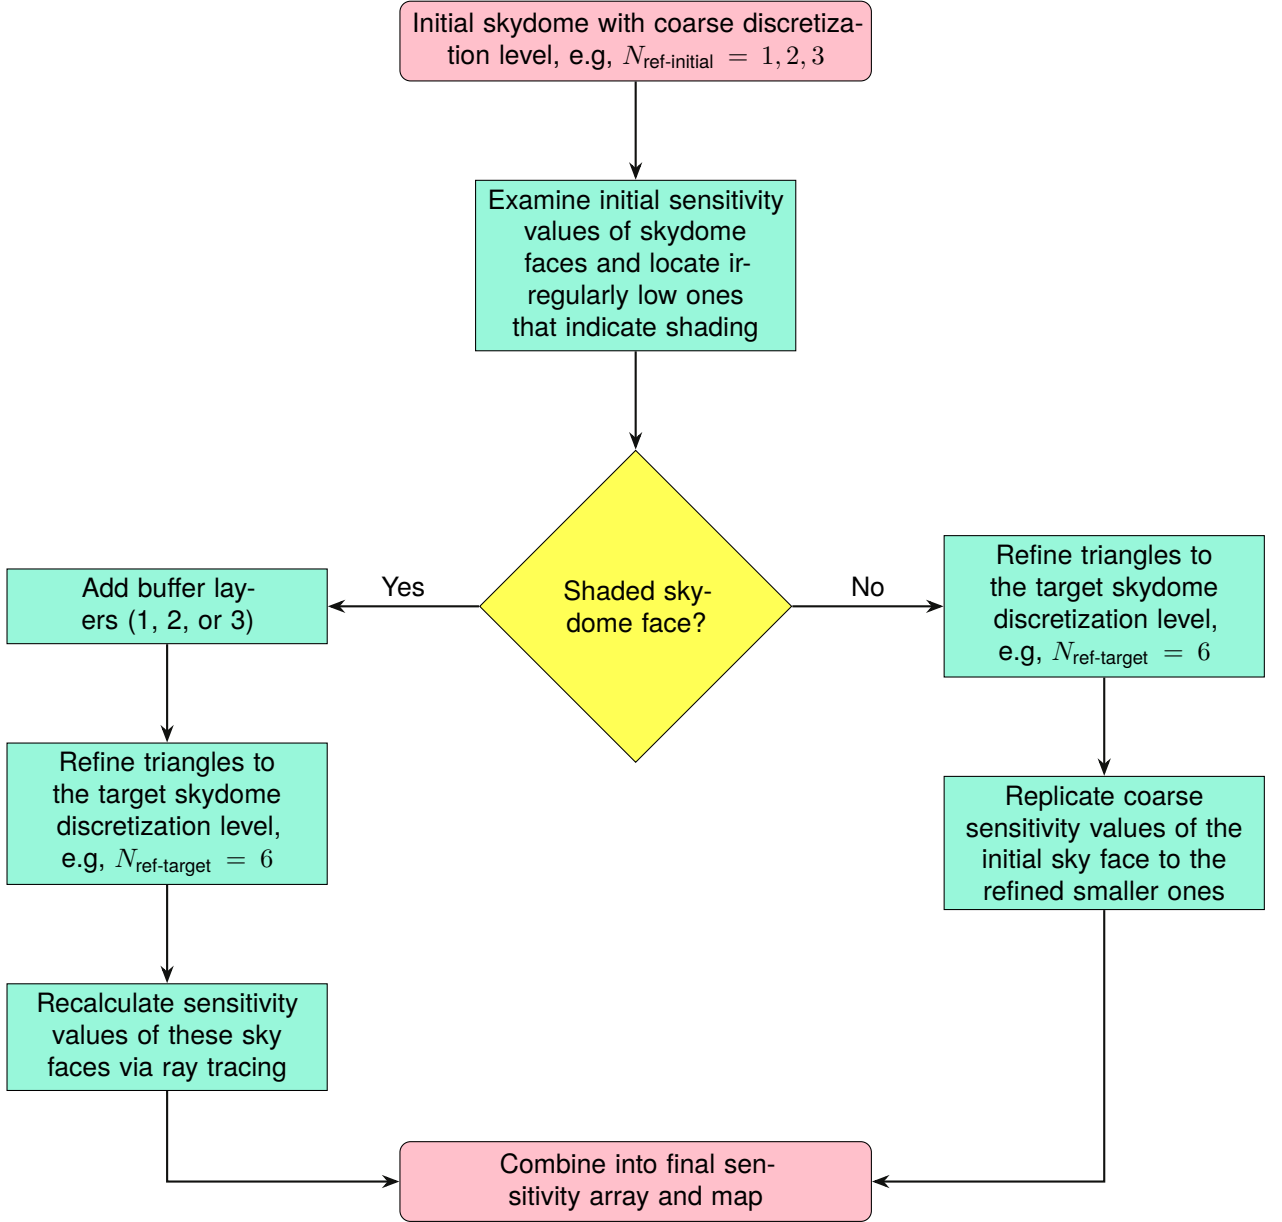


Figure 4.1: Flowchart for selective refinement and sensitivity processing

## 4.2. Selective Ray Tracing Methodology

In this section, these steps of the **Selective Ray Tracing** technique are explained in detail. The explanation begins with how the technique is integrated into the existing simulation framework, including the key variables and function modifications. Next, the process of identifying shaded regions on the skydome and constructing a buffer zone is outlined. This is followed by the refinement of selected regions and replication of sensitivity values for unshaded areas. It then explains how ray tracing is applied only to the selected regions, significantly reducing computational effort while maintaining accuracy. The final step summarizes the outputs of the method.

### 4.2.1. Integration and Key Variables

The selective ray tracing is realized within the `MODULE_nonperiodic` function, with some minor adjustments. In the original `MODULE_nonperiodic`, the `BackwardTracer` is called to perform the ray tracing based on the selected initial discretization level  $N_{\text{ref-initial}}$ . As mentioned in 3.1, the number of rays that are ray traced is equal to the number of triangular segments in the skydome facets, since each ray

represents the centroid of a facet.

After the backward ray tracing is through, the results are saved in an output variable (`MODULE_output`). This, among others, includes the skydome vertices and facets ( $V_s, F_s$ ), the skydome zenith, azimuth, and area for every triangle center ( $AZA$ ), and the sensitivity array ( $SM_f$ ) for every module.


When the selective refinement method is used, certain modifications are made within the same function. First, the outputs discussed above are marked as *initial* ( $V_{s\_initial}$ ,  $F_{s\_initial}$ ,  $AZA\_initial$ , and  $SM\_f\_initial$ ). Furthermore, the `BackwardTracer` function is modified to return two more variables:

- $N_{ref-initial}$ : the starting refinement level;
- `indexInt`: the indices of the skydome facets whose rays do not return to the sky because they intersect with an object in the environment.

These variables serve as inputs for the `SelectiveRefinement` function, where all steps (five) of the method are performed, as detailed in the following sections. To explain the steps in detail, the Example Case A (see 2.3.1) is going to be used where the initial skydome refinement level is set to  $N_{ref-initial} = 2$  and the goal level is  $N_{ref-target} = 6$ .

#### 4.2.2. Identification of Shaded Sky Triangles in the Initial Refinement

With an initial skydome discretization level of  $N_{ref-initial} = 2$  results in a skydome of  $n = 10 \times 4^{N_{ref-initial}} = 160$  triangle faces. Hence, 160 rays will be traced back for each cell, layer, and wavelength. The resulting sensitivity map is seen below in Figure 4.2.



**Figure 4.2:** Sensitivity map for skydome discretization level  $N_{ref-initial} = 2$

Then the first step is to examine which of the initial  $X = 160$  number of skydome facets could be shaded. The goal is to identify the indices of sky triangles affected by the wind turbine. That is why the `BackwardTracer` was modified to export the `indexInt`, i.e., the indices the ray was blocked, the

sensitivity values are equal to zero.


Proceeding with only these indices could lead to inaccuracies. This is because the `indexInt` checks only whether the ray from the centroid of each sky triangle intersects with the scene. In some cases, the ray may escape to the sky (indicating no direct obstruction at the centroid), yet the triangle could still receive significantly reduced irradiance due to partial shading or obstruction in neighboring directions. As a result, its sensitivity value may be low but not exactly zero. Thus, a filtering criterion was introduced.

As previously described, the plotted sensitivity map depicts the sensitivity array values for the first absorption layer averaged across all wavelengths and number of cells. In the given example, there are  $X = 160$  sensitivity values, one corresponding to each sky face. A threshold of 0.05 is applied to these values: sky indices with a mean sensitivity below this threshold are also selected.

This creates a problem, as sky triangles on the outer perimeter also have very low sensitivity values and would be selected. However, refining all of them is not necessary and they should eventually be excluded. These perimeter triangles are predefined for each initial refinement level and removed from the selection. It carries only unique entries and excludes duplicates. The final list of *important parent triangles* contains indices of:

1. fully shaded triangles with a mean sensitivity of zero, and
2. triangles with a mean sensitivity below 0.05 that do not belong to the outer perimeter.

In the present example, the final index selection would give us the sky indices `indexInt=53,138,154`, as shown in Figure 4.3.



**Figure 4.3:** Visible indexing for the sky facets of discretization level  $N_{\text{ref}} = 2$ .

The *potentially shaded triangles*, in this example indices 53, 138, and 154, are the initial triangle indices from which the selective refinement process should begin.

### 4.2.3. Buffer Zone and Final Selection


As briefly mentioned before, the method cannot proceed by only processing these initially identified potentially shaded triangles. If only these triangles were re-discretized and recalculated for their sensitivity values, the result would not be fully accurate. To ensure that all shaded regions are properly captured when increasing the skydome discretization, a buffer zone is added around them. This involves identifying one or more layers of neighboring triangles. Without this step, there is a high risk that parts of the shading geometry would be omitted, leading to a loss of accuracy.

The buffer zone of the initial *potentially shaded triangles* is constructed with the following steps:

1. For each initial *potentially shaded triangles*, identify the indices of its three vertices using the skydome facet matrix ( $Fs\_initial$ ).
2. For each of these vertices, scan the skydome facet matrix to find all other triangle facets that share the same vertex.
3. Exclude triangles already in the initial *potentially shaded triangles* set to avoid duplicates.

The buffer zone is extended iteratively for more than one layer. Finally, the initially recognized shaded triangles, along with those forming the buffer zone, are the indices that should be refined and recalculated.

In practice, using two buffer layers is typically sufficient to fully enclose the shape of the wind turbine within the refinement region for  $N_{ref} = 2 \rightarrow 6$ . The relevant graph pointing out the important area that needs recalculation of the sensitivity values is shown in Figure 4.4, with a final number of 45 triangle facets to be recalculated.



**Figure 4.4:** Initial *potentially shaded triangles* with buffer area outlined in red, including 2 layers of buffer zone.



#### 4.2.4. Increase Mesh to Target Refinement

Then it is time to increase the skydome discretization level. The whole idea behind selective backward tracing is that a very detailed skydome mesh is used, but rays are only traced for specific indices, leaving others untouched. Hence, it is necessary to know exactly which triangle face indices at the  $N_{\text{ref-target}}$  discretization level need to be ray traced again.

First the skydome needs to be recreated discretized to the desired level of  $N_{\text{ref-target}} = 6$  with the known *icohemisphere* function. This provides the vertices and the facets of a skydome with 40,960 triangles. As mentioned earlier, the *icohemisphere* function works in a way that each initial triangle, *parent triangle*, is subdivided recursively into smaller, *child triangles*. At each refinement step, every parent triangle is divided into four child triangles. As a result, a refinement jump of  $r_l$  levels produces  $4^{r_l}$  child triangles per parent.

The current example involves a transition from  $N_{\text{ref}} = 2$  to  $N_{\text{ref}} = 6$ , represented by the jump  $r_l = 4$ . Hence, every triangle in the *refined parent triangles*, should be split into  $n_{\text{children}} = 4^{r_l=4} = 256$  *child triangles*. The sensitivity value of each *parent triangle* is copied to all of its corresponding *child triangles*. To implement this replication,

1. The original 4D sensitivity array `SM_f_initial`, calculated at  $N_{\text{ref}} = 2$ , is first expanded to match the refined skydome at  $N_{\text{ref}} = 6$ . `SM_f_initial` has dimensions  $160 \times 72 \times 3 \times 46$ , which means that there are  $3 \times 46 = 138$  matrices of sensitivity values of 160 rows and 72 columns.
2. Each row of the array corresponds to one triangle facet on the skydome. Since each triangle is subdivided into 256 children, each row must be replicated  $n_{\text{children}} = 256$  times.
3. For every original triangle  $i$ , the corresponding 4D data block `SM_f(i, :, :, :)` is repeated 256 times using the `repmat` function.
4. These replicated blocks are inserted into a new array, `SM_refined_replicated`, which has dimensions  $40,960 \times 72 \times 3 \times 46$ , at the right indices, matching the total number of triangles in the fully refined skydome.

The final result is a replicated sensitivity matrix of size  $40,960 \times 72 \times 3 \times 46$ , where each of the 40,960 child skydome triangles inherits the sensitivity value of its corresponding parent triangle. Thus, there is a new pair of skydome vertices and indices ( $V_s, F_s$ ). Among these, the subset of indices labeled *retraced triangles* includes only the triangles that require re-tracing and recalculation based on the refined facets matrix ( $F_s$ ).

Having established that, the next step is to link the indices of *child triangles* that belong to a *parent triangle*, in order to know exactly which very small *child triangles* need to be ray traced again. As seen in Figure 4.4, 45 sky triangle facets are to be refined. This means that  $45 \times 256 = 11,520$  triangles should have their sensitivity values recalculated, i.e., 11,520 rays should be traced again (*retraced triangles*). In contrast, tracing the full skydome at  $N_{\text{ref}} = 6$  would require 40,960 rays per wavelength and per cell. Therefore, 29,440 ray tracing calculations are saved. This corresponds to a reduction of approximately 72% in the number of ray tracing computations. That ultimately means that more than two-thirds of the original workload is avoided through this selective refinement.

#### 4.2.5. Selective Backward Ray Tracer

The final step involves passing the *11,520 retraced triangles* to the `SelectiveBackwardTracer` function for ray tracing. This function is very similar to the original `BackwardTracer` with some alterations. There is *Primary Ray Tracing* and *Secondary Ray Tracing*.

For the *Primary Ray Tracing*, three filtering criterias are applied. Only rays that satisfy all of the following conditions are considered valid in primary the sensitivity calculation:

1. Point upward (toward the sky): The rays are first normalized, and then a basic filter is applied to exclude those pointing downwards. Only rays with a positive z-component, meaning they point upward toward the sky, are kept.
2. Hit solar cell surface from the front: The dot product between the cell normal vector and the ray direction is examined. For each solar cell in the module, only rays that strike the front side are considered. To determine this, the tilt and azimuth of each cell are used to calculate a 3D surface

normal vector that points outward from the front of the cell. The cosine of the angle of incidence (AOI) for each incoming ray is then calculated by taking the dot product of the cell normal vector and the ray direction. If this dot product is positive, the ray is approaching from the front side of the cell and is retained. If it is negative or zero, the ray comes from behind or is exactly tangent to the surface, and is therefore ignored.

3. Do not intersect anything in the 3D scene: Each valid, up to now, ray is traced from the cell towards the sky in order to detect intersections with the environment. If a ray escapes without intersection, it directly contributes to the sensitivity. If it hits an object, it spawns secondary rays.

These are the rays that provide unhindered direct irradiance from the sky. Rays that were hit the cell from behind, or did not escape upward to the sky, or were blocked by an object in the scene were filtered out, and their sensitivity contribution for the *Primary Ray Tracing* is set to zero. The rays that intersected with the environment will be addressed in the *Secondary Ray Tracing*.

Then, the valid primary rays need to be paired to the EQE profile to calculate the spectral-angular sensitivity. This is achieved by finding the closest available EQE measurement and matching it with the angle of incidence (AOI). The amount of light absorbed by the solar cell layers is calculated using this EQE. As discussed in 3.1, the first three EQE layers are used to compute how much light is absorbed in each part of the module.

The intensity of incoming light and its likelihood of being absorbed by the module are directly influenced by the angle of incidence relative to the surface normal vector [47]. Therefore, the sensitivity calculation incorporates the relevant EQE, weighted by the AOI, as shown in Equation 4.1:

$$sensitivity = \cos(\theta_{AOI}) \times EQE \quad (4.1)$$

For rays that intersect with the environment and are thus blocked, *Secondary Ray Tracing* is performed to account for reflected or scattered contributions. The type of reflection at the intersection point is determined by classifying the surface as either Lambertian or Specular.

1. Lambertian (diffuse): The surface scatters light in many upward directions [48]. When this holds true, a new set of secondary rays is cast from the point of intersection in all upward directions. These rays are checked again for intersections with the environment and only rays that escape to the sky contribute to the sensitivity map. Each contribution is scaled by the surface albedo to account for diffuse reflection. The sensitivity contribution of this ray depends on the cosine of the incoming angle and the EQE:

$$sensitivity = \cos(\theta_{AOI, \text{primary}}) \times EQE \times albedo \quad (4.2)$$

2. Specular (mirror-like): Light reflects in a single direction, as a beam [49]. This reflected ray is traced from the intersection point. If the reflected ray escapes the scene, its contribution is added to the nearest sky patch. Similar to the Lambertian case, the contribution is scaled by the surface albedo, but only a single reflected direction is considered.

$$sensitivity = \cos(\theta_{AOI, \text{primary}}) \times EQE \times albedo \quad (4.3)$$

After the function `SelectiveBackwardTracer` is called and its processes are complete, the function returns a single output. That is, the 4D sensitivity values of the designated sky faces, `F_refine_children`, that were selected to be recalculated.


### 4.3. Output

Lastly, this new sensitivity array of the recalculated faces is merged with the replicated arrays for the remaining sky faces. The resulting sensitivity array per module is a mixture of coarse and selectively refined sensitivity values. This method targets to achieve high sensitivity accuracy for shading calculations, with decreasing the computational burden to only a fraction of that required for full refinement.

The outputs of the whole selective ray tracing process are listed:


- $SM_f$ : the final selectively updated sensitivity map
- $V_s, F_s$ : the vertices and faces of the refined skydome
- $AZA$ : azimuth, zenith, and area data for each sky triangle
- $SM_f_{replicated}$ : the full refined sensitivity array before updates, for illustrative purposes

The resulting sensitivity array has dimensions  $40,960 \times 72 \times 3 \times 46$  and the sensitivity map plotted for the example case is presented below in Figure 4.5.




**Figure 4.5:** Resulting sensitivity map for selectively discretized skydome, from  $N_{ref} = 2$  to  $N_{ref} = 6$ .


For clarity, Figure 4.6 shows the evolution of the sensitivity map throughout the selective refinement process.



(a) Original sensitivity map for skydome discretization level  $N_{\text{ref}} = 2$ .



(b) Fully discretized skydome up to level  $N_{\text{ref}} = 6$ , with values replicated from the original one.



(c) Fully discretized skydome to level  $N_{\text{ref}} = 2 \rightarrow 6$ , with selective skydome refinement method

**Figure 4.6:** Evolution of the sensitivity map through the selective refinement process.

## 4.4. Conclusions

This chapter addressed Sub-question 3.: *What enhancements can be introduced in the PVMD toolbox to improve the accuracy and efficiency of the shading analysis of a PV module with respect to its surroundings?*

A *Selective Skydome Refinement* method was developed to improve the computational efficiency of sensitivity analysis in the PVMD Toolbox, and it is implemented for an example case. The method enables high resolution sensitivity analysis by *selectively* refining only the sky regions that contribute to shading. This is achieved by first running an initial low-resolution simulation to detect sky triangles that indicate that they are shaded by having zero or very low initial sensitivity. Outside of these triangles, a buffer zone is added to ensure sufficient coverage. Only then, the skydome is refined to the targeted level of discretization, and for this final subset of triangles the sensitivity values are recalculated through a second ray tracing, whereas the rest are assigned replicated initial coarse-level values.

This selective approach significantly reduces the number of rays that need to be traced, and this computational time. It is achieving accurate sensitivity mapping with high skydome resolution for the areas that matter most, at only a fraction of the computational time required by full high-resolution simulations. At the present example, the number of rays traced to a high refinement level decreased by more than 70%.

One of the major advantages of this method, apart from its efficiency, is that it fits seamlessly into the existing PVMD Toolbox architecture. The method was fully integrated into the existing toolbox through minor modifications in the `MODULE_nonperiodic` function and the creation of two dedicated functions: `SelectiveRefinement` and `SelectiveBackwardTracer`.

In the following chapter, concrete validation and performance analysis of the method will be demonstrated.

# 5

## Performance Evaluation and Application of the Developed Skydome Refinement Method

This chapter addresses Sub-question: **4. What are the performance benefits of the developed method, and how can it be further optimized? Furthermore, how does wind turbine shading influence spatial energy yield variability across a large-scale PV layout?** First, in Section 5.1, in-depth simulation results are presented for an example case, aiming to validate the accuracy of the selective skydome refinement method introduced in the previous chapter. The performance gains are assessed in terms of reduced computational burden and preserved energy yield accuracy. Then, in Section 5.2 further optimization is carried out with respect to what is the optimal starting skydome resolution for selective refinement. In the final Section 5.3 of this chapter, the selective refinement method is applied to a larger photovoltaic (PV) system layout surrounding a wind turbine.

### 5.1. Validation of the Selective Refinement Method

In order to present the validation of the introduced sensitivity analysis method, Example Case A will be used again. For this specific module, the sensitivity analysis of the original unrefined case of skydome discretization  $N_{\text{ref}} = 6$  will be compared to a selective refinement discretization, which starts with a base of  $N_{\text{ref}} = 2$  and selectively refines specific parts of the skymap to  $N_{\text{ref}} = 6$ . First, to assess how much the computational time for the sensitivity analysis was reduced and then to evaluate the effect on the energy yield calculation accuracy.

#### 5.1.1. Computational Cost Analysis

For the Example Case A, the computational time for the calculation of the sensitivity array in the original way as mentioned in 3.5 is 229560.69 (s)  $\approx$  3826.0115 (min)  $\approx$  63.8 (h)  $\approx$  2.65 (d). After following the exact methodology that was described in 4.2 and refining only specific selected sky triangles, with a base of  $N_{\text{ref}} = 2$ , the computational time, for that specific case, reduced to 8430.53 (s)  $\approx$  140.51 (min)  $\approx$  2.34 (h). For clarity, it is more meaningful to compare the computational time in hours. The total time was remarkably reduced from 63.8 (h) in the fully refined case to 2.34 (h) using the selective refinement approach. To make the comparison more transparent, Table 5.1 summarizes this significant improvement. The representation of the sensitivity maps before and after refinement are shown in Figure C.1, in Appendix C.2.

**Table 5.1:** Comparison of total computational time between the original and selective refinement methods.

Method	Seconds (s)	Minutes (min)	Hours (h)	Days (d)
Original (fully refined) $N_{\text{ref}} = 6$	229,560.69	3,826.01	63.80	2.65
Selective refinement $N_{\text{ref}} = 2 \rightarrow 6$	8,430.53	140.51	2.34	–

The percentage reduction in computational time is calculated using the following formula:

$$\Delta T = \left( \frac{T_{\text{original}} - T_{\text{refined}}}{T_{\text{original}}} \right) \times 100\% \quad (5.1)$$

For this case :  $T_{\text{original}} = 63.8$  h and  $T_{\text{refined}} = 2.34$  h. The computational time was reduced by approximately 96.3%:

$$\Delta T = \left( \frac{63.8 - 2.34}{63.8} \right) \times 100\% = \left( \frac{61.46}{63.8} \right) \times 100\% \approx 96.33\%$$

As previously mentioned, it needs to be clear that these numbers correspond to this specific simulation and it would be wrong to generalize. As previously mentioned the simulation time depends also on the capacity of the computer that it is running on, or other applications running in the background, etc. The computational time for the sensitivity analysis for a simulation with selective refinement from  $N_{\text{ref}} = 2$  to  $N_{\text{ref}} = 6$  on different devices can vary within a range of 2 (h) to 5 (h). Nevertheless, that would mean that the computational time reduction  $\Delta T$  from Eq. 5.1 lies within the range  $\Delta T \in [92.16\%, 96.86\%]$ , depending on the exact runtime, which is in any case a substantial decrease.

### 5.1.2. Energy Yield Analysis

However, there is naturally a tradeoff of the selective refinement method and that is the reduction in the energy yield accuracy calculation. In other words, this refers to the divergence between the energy yield computed using a fully refined skydome and the yield obtained after applying the selective refinement method, for a given simulation period.

To demonstrate that, the example simulation A was carried out for two days, June 1<sup>st</sup> and January 1<sup>st</sup>, and for a whole year. For the three time frames, the the DC energy yield was calculated using the initial version of the shading analysis and the refined one. The results are presented in the Table 5.2 below:

**Table 5.2:** Comparison of DC energy output (in kWh) for unrefined and selectively refined methods.

Method	DC Energy Output (kWh)		
	1 <sup>st</sup> January	1 <sup>st</sup> June	Full year
Original (fully refined) $N_{\text{ref}} = 6$	0.1371	1.8949	288.8022
Selectively refined $N_{\text{ref}} = 2 \rightarrow 6$	0.1369	1.9109	289.0628

The accuracy of the method is assessed using the relative error formula in Equation 3.4:

$$\epsilon_{\text{rel}}(i - j) = \frac{\text{DC\_EY}(j) - \text{DC\_EY}(i)}{\text{DC\_EY}(i)} \times 100\% \quad (5.2)$$

where  $i$  refers to the unrefined case with full skydome resolution ( $N_{\text{ref}} = 6$ ) and  $j$  to the selectively refined case ( $N_{\text{ref}} = 2 \rightarrow 6$ ). This metric quantifies the relative difference in DC energy yield between the original and selectively refined sensitivity analyses. The value of  $\epsilon_{\text{rel}}(i - j)$  is calculated for the three selected simulation time frames respectively and the results are displayed in Table 5.3.

$$\epsilon_{\text{rel}}(N_{\text{ref}} = 6 \rightarrow N_{\text{ref}} = 2 \rightarrow 6) = \frac{\text{DC\_EY}(N_{\text{ref}} = 2 \rightarrow 6) - \text{DC\_EY}(N_{\text{ref}} = 6)}{\text{DC\_EY}(N_{\text{ref}} = 6)} \times 100\% \quad (5.3)$$

**Table 5.3:** Relative error (%) between unrefined and selectively refined methods for selected periods.

Period	1 <sup>st</sup> January	1 <sup>st</sup> June	Full Year
$\epsilon_{\text{rel}}(N_{\text{ref}} = 6 - N_{\text{ref}} = 2 \rightarrow 6)$	-0.116%	0.847%	0.090%

As seen, the values of  $\epsilon_{\text{rel}}$  are below 1%, whether it leads to a slight overestimation or underestimation. This confirms that the selective refinement method maintains a high level of accuracy in the DC energy yield estimation.

However, it is important to note that these results are indicative. They are specific to the example case presented here and should not be taken as absolute. Additionally, the initial backward ray tracing at  $N_{\text{ref}} = 2$  plays a critical role, as it influences all subsequent steps of the simulation. Since this initial sensitivity mapping step depends on geometric alignment and is inherently outside the user's control, variations may occur between simulations. Nevertheless, the observed results are indicative of the method's strong overall performance in terms of accuracy and computational efficiency.

## 5.2. Optimal Starting Resolution for Selective Refinement

Since the method-model is validated it is important to also optimize. Assessing whether  $N_{\text{ref}} = 2$  is the most optimal starting level for skydome discretization is a key optimization step that will help to find the best tradeoff between computational efficiency and energy yield accuracy. It makes sense to compare starting the selective refinement from  $N_{\text{ref}} = 2$  or 3, but not higher, since the aim is to reduce computational time.

The number of triangles, *refined parent triangles*, in the initial skydome discretization that need recalculation are the *important parent triangles* combined with the *buffer area triangles*. Starting from the lower resolution  $N_{\text{ref}} = 2$ , the number of *important parent triangles* identified is smaller compared to starting from  $N_{\text{ref}} = 3$ . However, the coarser mesh of  $N_{\text{ref}} = 2$  consists of larger triangles that each cover a bigger angular area on the skydome, than of  $N_{\text{ref}} = 3$ . This can also be seen in Figure 5.1b, 5.1a.

Then, as explained in the previous Chapter 4, the triangles that are actually re-sent to the backward ray tracer and recalculated are the *child triangles* of the *refined parent triangles*. The *child triangles* for each case result in the following way:

- Starting from  $N_{\text{ref}} = 2$ : Two buffer layers are required.  
For  $N_{\text{ref}} = 2 \rightarrow 6$  a jump of 4 levels takes place, thus each *refined parent triangles* produces  $n_{\text{children}} = 4^4 = 256$  *child triangles*.
- Starting from  $N_{\text{ref}} = 3$ : Three buffer layers are required.  
For  $N_{\text{ref}} = 3 \rightarrow 6$ , a 3 level jump occurs which means that each *refined parent triangles* produces  $n_{\text{children}} = 4^3 = 64$  *child triangles*.

The final number of triangles that need to be re-traced backward for the selective sensitivity analysis is found by multiplying the number of *refined parent triangles* with the  $n_{\text{children}}$ . The more these triangles are, the more time is needed for tracing back the rays, hence the bigger the computational cost of the simulation. Thus, it must be examined which one of the two gives, on average, smaller total amount of triangles that need recalculation. And furthermore, take into account how this difference affects the accuracy of energy yield and computation time.

For this purpose, a sample of 32 modules are tested, derived from Example Case D (see 2.3.4). For each of these modules the sensitivity maps of  $N_{\text{ref}} = 2$  and  $N_{\text{ref}} = 3$  were examined. It was then documented for each the number of parent triangles and child triangles to refine. The results are documented and displayed in Appendix C.3.

For  $N_{\text{ref}} = 2$ , an average of 7,336 skydome facets were identified for recalculation per module. However, in 7 out of the 32 cases, no *important parent triangles* were detected at all, thus zero skydome triangles would be sent for ray tracing. This outcome occurs due to the coarse resolution of the skydome at  $N_{\text{ref}} = 2$ . The skydome at  $N_{\text{ref}} = 2$  consists of only 160 relatively large triangles and each covers a larger angular area. As a result, it becomes less likely that a ray will intersect exactly through the center



of a triangle, making it more prone to missing shaded regions. Consequently, these cases where no *important parent triangles* were found biased the average of 7,336 refined skydome facets for  $N_{\text{ref}} = 2$  downward. Apart from that, this also highlights a potential drawback of starting from  $N_{\text{ref}} = 2$ : the coarse resolution may lead to missed shading detection and thus unreliable refinement results.

Then, examining the same 32 modules for  $N_{\text{ref}} = 3$  showed an average of 7,720 skydome facets marked for recalculation (child triangles), with no cases where zero facets were detected, unlike the  $N_{\text{ref}} = 2$  case. Although the average number of facets marked for recalculation is slightly higher for  $N_{\text{ref}} = 3$ , this starting level, it provides a more reliable and accurate basis for refinement, as it eliminates the risk of overlooking initially shaded triangles. Which is a factor that is ultimately more critical than minimizing the number of retraced rays.

In order to compare the computational time for the sensitivity analysis of the two skydome discretization points, the module of Example Case A is used. As previously mentioned in Section 5.1.1, for a simulation with selective refinement from  $N_{\text{ref}} = 2$  to  $N_{\text{ref}} = 6$  on different devices can vary within a range of 0.5 (h) to 4 (h). Using Eq. 5.1 the computational time reduction  $\Delta T$  lies within the range  $\Delta T \in [93.73\%, 99.22\%]$ . For the case that the selective refinement from  $N_{\text{ref}} = 3$  to  $N_{\text{ref}} = 6$ , the computational time for this module was found to be 0.1 (h) to 3 (h). Consequently, compared to the original unrefined  $N_{\text{ref}} = 6$  analysis the reduction in time from Eq. 5.1 can be found to be within the range  $\Delta T \in [95.3\%, 99.84\%]$ . The conclusion that starting from initial skydome discretization  $N_{\text{ref}} = 3$  has in most cases a faster response can be drawn.

As before, apart from the computational cost reduction, the energy yield needs to be looked at. Six simulations for the Example Case A were carried out to estimate and compare the DC energy output values. Three for  $N_{\text{ref}} = 2 \rightarrow 6$  and three for  $N_{\text{ref}} = 3 \rightarrow 6$  for the simulation time frames of two days, June 1<sup>st</sup> 2021 and January 1<sup>st</sup> 2021, as well as the entire year of 2021. For all, the yield was calculated, and the results are presented in the Table 5.4.

**Table 5.4:** Comparison of DC energy output (in kWh) for unrefined and the two selectively refined methods.

Method	DC Energy Output (kWh)		
	1 <sup>st</sup> January	1 <sup>st</sup> June	Full year
Unrefined	0.1371	1.8949	288.8022
Selectively refined $N_{\text{ref}} = 2 \rightarrow 6$	0.1369	1.9109	289.0628
Selectively refined $N_{\text{ref}} = 3 \rightarrow 6$	0.1366	1.9112	290.1053


The relative error  $\epsilon_{\text{rel}}(i - j)$  in DC energy yield is calculated using Equation 3.4, comparing the fully refined skydome at  $N_{\text{ref}} = 6$  with the selectively refined cases starting from  $N_{\text{ref}} = 2$  and  $N_{\text{ref}} = 3$ . The resulting values of  $\epsilon_{\text{rel}}(i - j)$  are evaluated for the three simulation time frames are seen in Table 5.5.

**Table 5.5:** Relative error (%) between unrefined and the two selectively refined methods for selected periods.

Period	1 <sup>st</sup> January	1 <sup>st</sup> June	Full Year
$\epsilon_{\text{rel}}(N_{\text{ref}} = 6 - N_{\text{ref}} = 2 \rightarrow 6)$	-0.116%	0.847%	0.090%
$\epsilon_{\text{rel}}(N_{\text{ref}} = 6 - N_{\text{ref}} = 3 \rightarrow 6)$	-0.359%	0.861%	0.451%

From this comparison it is seen that while the relative error  $\epsilon_{\text{rel}}$  of the calculated DC energy yield for  $N_{\text{ref}} = 3$  is slightly higher than that of the  $N_{\text{ref}} = 2$  starting point, the difference remains well below 1% across all simulation periods and can therefore be considered acceptable. It should be again reminded that these numerical results are specific to the simulated case and should not be interpreted as universally applicable. It does not mean that starting from  $N_{\text{ref}} = 3$  will always under or over estimate the yield this much.

An additional comparison between the two initial refinement levels is illustrated in Figure 5.1, which shows the sensitivity maps for the Example Case C simulation of a module located north of the wind turbine (see 2.3.3).



**Figure 5.1:** Sensitivity maps of the module of Example Case C, for four different skydome discretization levels: (a)  $N_{\text{ref}} = 2$ , (b)  $N_{\text{ref}} = 3$ , (c)  $N_{\text{ref}} = 2 \rightarrow 6$ , and (d)  $N_{\text{ref}} = 3 \rightarrow 6$ .

To conclude, starting the selective refinement process from  $N_{\text{ref}} = 3$  appears to be the more robust and efficient choice. Refinement starting from  $N_{\text{ref}} = 3$  consistently avoids the risk of missing shaded regions, an issue observed in several  $N_{\text{ref}} = 2$  cases, where no *important parent triangles* were detected due to the coarse angular resolution. Additionally, the computational time for  $N_{\text{ref}} = 3 \rightarrow 6$  is lower on average than for  $N_{\text{ref}} = 2 \rightarrow 6$ . And although for this example the calculated DC power yield is slightly less accurate compared to the not refined one, it can be concluded that beginning the refinement from  $N_{\text{ref}} = 3$  offers a better overall tradeoff between accuracy, consistency, and computational efficiency.

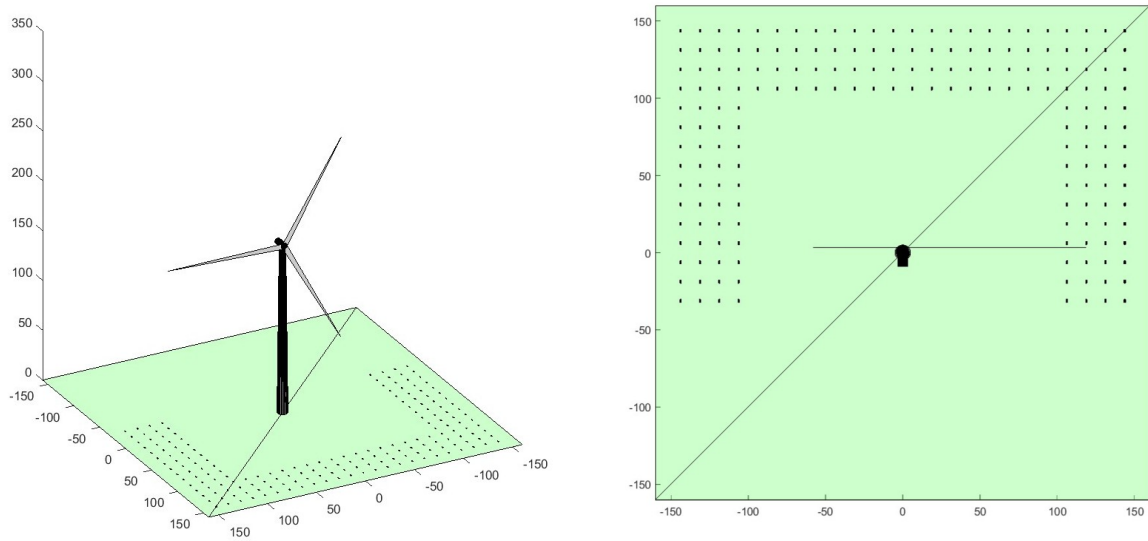
Hence, for the next and final section of this chapter the simulations will be done by using the selective refinement technique for  $N_{\text{ref}} = 3 \rightarrow 6$ .

### 5.3. Shading Analysis for a larger PV system

This section applies the developed selective refinement method to a more realistic PV system layout in order to investigate spatial variability in energy yield caused by shading from a wind turbine. Two main objectives are addressed:

- To quantify how much the simulated energy yield can vary between adjacent module positions, where all other conditions are identical, and the only influencing factor is dynamic shading from the turbine.
- To analyze overall shading patterns across the layout during the simulation period and explore how these patterns could inform preliminary PV placement strategies.

The simulation that is employed corresponds to the Example Case D (see 2.3.4), with 184 positions where modules are placed, as shown in Figure 5.2. The placement of these modules is not arbitrary but follows the following rationale, which is explained and visualized in detail in Appendix D.1.



**Figure 5.2:** The 184-module PV system and 1 wind turbine used in the simulation.

Each module is individually simulated with the selective refinement method  $N_{\text{ref}} = 3 \rightarrow 6$  and the DC energy output is calculated for the 1<sup>st</sup> June.

#### 5.3.1. Energy Yield Variability Between Adjacent Modules

Modules placed close together are expected to experience similar irradiance and shading conditions, but shading from the turbine introduces small-scale spatial variability. To quantify this variability, a statistical uncertainty analysis is performed. Two types of uncertainty are evaluated:

- **Intra-cell uncertainty:** Each PV module is placed in the center of a  $12.5 \text{ m} \times 12.5 \text{ m}$  grid cell. In reality in an area like that, multiple PV modules could fit. As a result, simulating just one (the central) position may not fully capture the local variability in shading and yield within that cell. That is why a margin of error is calculated to quantify how representative the simulated module is for the entire cell, in terms of shading-induced yield variability.
- **Intra-cluster uncertainty:** The total amount of 184 modules can be split into several clusters. In this case, it can be examined how much the simulated energy yield differs between these positions. In this case, an estimation of a margin of error captures the spread in simulated yield values across modules separated by distances up to approximately 50 m. This gives a direct indication of how yield differences can emerge within short spatial scales due to turbine shadowing alone.

The margin of error (MoE) throughout the analysis below is calculated using the following equation:

$$\text{MoE} = t^* \cdot \frac{s}{\sqrt{n}} \quad (5.4)$$

where  $t^*$  is the critical value from the Student's  $t$ -distribution [50] corresponding to the desired confidence level and  $\text{df} = n - 1$  degrees of freedom,  $s$  is the sample standard deviation, and  $n$  is the number of samples (modules) considered. This formulation is appropriate for small sample sizes ( $n < 30$ ) where the population standard deviation is unknown and must be estimated from the sample [51]. In the present work, a two-tailed 95% confidence interval is used to estimate the uncertainty in the average DC yield. This interval captures the central 95% of the distribution of possible sample means, leaving 2.5% in each tail. The critical value  $t^*$  corresponds to the 97.5th percentile of the distribution, denoted as  $t^* = t_{1-\alpha/2, \text{df}} = t_{0.975, n-1}$ , where  $\alpha = 0.05$ . This ensures that the resulting confidence interval:

$$\bar{x} \pm t^* \cdot \frac{s}{\sqrt{n}}$$

has a 95% probability of containing the true population mean.

#### Uncertainty within a single grid cell

The simplification that only one PV module is simulated at the center of each grid square raises the question of how accurately this single point can represent the energy yield across the entire  $12.5 \text{ m} \times 12.5 \text{ m}$  area of the square. To evaluate this assumption, a dedicated simulation is performed for one grid square, where 25 PV modules are placed within the same square area, arranged in a  $5 \times 5$  grid. For more detailed visual explanation, see D.3.

The goal of this analysis is to quantify how much the DC energy yield fluctuates within a single grid square and to measure the relative deviation from the center module. And furthermore, to determine the statistical confidence in using only the center module's result to represent the entire square.

To quantify spatial variation, the relative error of each module is calculated with respect to the center module using the Equation 3.4. In this context,  $\epsilon_{\text{rel}}$  is the relative error (in percent),  $\text{DC\_EY}(i)$  is the DC energy yield of the reference module (here, the center module), and  $\text{DC\_EY}(j)$  is the yield of any other of the 25 modules  $j$  within the square. The values of the yield for all positions and the corresponding relative errors are shown in Table D.1.

The results, Table 5.6, show that the mean DC energy yield across all 26 modules is 2.0268 kWh, while the center module yields 2.0331 kWh. The maximum absolute relative error among the modules is 0.82 %, and the mean absolute relative error is only 0.35 %. The standard deviation  $s$  of the yield is 0.005 48 kWh, resulting in a margin of error of  $\pm 0.002 \text{ 21 kWh}$ , with 95% confidence. According to these results, the relative error is below 1%, which indicates that placing and simulating one module per grid square is a sufficiently accurate approach.

**Table 5.6:** Statistical overview of center-based sampling accuracy within a  $12.5 \text{ m} \times 12.5 \text{ m}$  grid square.

Metric	Value
Number of simulated modules	25
Grid square size	$12.5 \text{ m} \times 12.5 \text{ m}$
Central module DCP	2.0331 kWh
Mean of all 26 DCPs	2.0268 kWh
Standard deviation $s$	0.005 48 kWh
Critical value: $t_{0.975, 25}^*$	2.068
Margin of error (95% CI)	$\pm 0.002 \text{ 21 kWh}$
Max absolute relative error	0.82 %
Mean absolute relative error	0.35 %

#### Uncertainty within a cluster of adjacent modules

Next, a cluster of 16 modules is selected and simulated, an area that spans across  $50\text{ m} \times 50\text{ m}$  (see Appendix D.3), and their energy yield is calculated for the specified day. It is examined how much each of these modules placed in the 16 grid squares deviates from the center value. Again, the statistical uncertainty (MoE) associated with this simplification is quantified.

The margin of error is, as mentioned, calculated according to Equation 5.4, with  $s$  the standard deviation of the yield values across the 16 grid squares, and  $n = 17$  is the number of samples. Furthermore, the relative difference between the 16 positions and the central one is computed using the known Equation 3.4, where  $\text{DC\_EY}(i)$  denotes the energy yield at the batch center and  $\text{DC\_EY}(j)$  the yield at grid square  $j$ .

The results of these calculations for the relative error and the margin of error can be found in Table D.2. A summary of the most relevant statistical values is displayed in Table 5.7.

**Table 5.7:** Statistical overview of center-point sampling accuracy within a  $50\text{ m} \times 50\text{ m}$  module cluster.

Metric	Value
Number of simulated modules	16
Total batch area	$50\text{ m} \times 50\text{ m}$
Central (batch) module DCP	2.0350 kWh
Mean of all 17 DCPs	2.0335 kWh
Standard deviation $s$	0.0023 kWh
Critical value: $t_{0.975, 17}^*$	2.119
Margin of error (95% CI)	$\pm 0.0012\text{ kWh}$
Max absolute relative error	0.31 %
Mean absolute relative error	0.12 %

The average energy yield across the cluster is 2.0335 kWh, closely matching the cluster center yield of 2.0350 kWh. The maximum relative deviation observed among modules is 0.31 %, while the average deviation is only 0.12 %. With a standard deviation of 0.0023 kWh, the margin of error at a 95% confidence level is  $\pm 0.0012\text{ kWh}$ .

#### 5.3.2. Analysing shading patterns

The second part of this section analyzes the shading patterns affecting a slightly more realistic PV system. A clear way to visualize the impact of turbine-induced shading across the layout is by calculating the percentage loss in DC energy yield over a defined period. The percentage loss is calculated by comparing the energy yield of each module in the presence of the turbine to its yield in the same case but without the turbine, as defined in Equation 5.5. In this way, the change in the estimated DC power yield (DCP) value is only affected by the presence of the turbine. For each PV module, the loss is computed as:

$$L_{\text{shading}} = \left( \frac{|\text{DC\_EY}_{\text{no turbine}} - \text{DC\_EY}_{\text{with turbine}}|}{\text{DC\_EY}_{\text{no turbine}}} \right) \times 100\% \quad (5.5)$$

A way of visually illustrating the losses due to shading can be with a contour plot. In order to do that, Example Case D is used, where 184 PV modules are positioned around a wind turbine and a simulation is performed for one day, the 1<sup>st</sup> June, using the selective refinement method for  $N_{\text{ref}} = 3 \rightarrow 6$ .

### Sampling and Interpolation Strategy

In principle, to do that, all the 184 modules should be simulated fully to find their DC electrical output. However, since modules located close together are likely to experience similar irradiance conditions and shading patterns, as shown in the previous section, it is reasonable to simulate only a selection of representative positions. Based on this assumption, the rest can be estimated using an interpolation, with a margin of error included to reflect the expected variation.

To proceed efficiently, the 184 module positions are divided into 10 clusters, as shown in Appendix D.4. This structure helps to divide and identify which modules are simulated and which are estimated. For all clusters of 16 or 20 modules covering an area of  $50\text{ m} \times 50\text{ m}$  or  $50\text{ m} \times 62.5\text{ m}$ , the four corner modules and the central one are simulated. Their DC energy yield is found, and for the rest in-between modules the DC yield is estimated using an interpolation method. In addition to that, based on the variation among the five simulated modules, a margin of error is calculated and applied uniformly to all interpolated values within the cluster, given in Eq. 5.4.

The method of *two-dimensional quadratic surface interpolation* is employed to estimate the DCP at non-simulated positions. This method assumes that the shading effect, and therefore energy yield, varies smoothly but possibly with some curvature across a small spatial area. It based on classical 2D polynomial fitting techniques, as described in [52–54]. The interpolation relies on the known DCP values of five modules in each batch: the four corner modules and the central module. These are treated as known function values  $f_{00}$ ,  $f_{10}$ ,  $f_{01}$ ,  $f_{11}$ ,  $f_c$ , corresponding to the bottom-left, bottom-right, top-left, top-right corners, and the center of the cluster area, respectively.

To apply this interpolation method, the cluster's physical coordinates (such as a  $50\text{ m} \times 50\text{ m}$  or  $50\text{ m} \times 62.5\text{ m}$  cluster area) are first transformed into a local, normalized coordinate system. The bottom-left corner is set as the origin  $(0, 0)$ , and vectors to the bottom-right and top-left corners define the local basis. Each module's position is then mapped to normalized coordinates  $(x, y)$  within the unit square.

These known values  $f_{00}$ ,  $f_{10}$ ,  $f_{01}$ ,  $f_{11}$ ,  $f_c$  constitute the constraints used to determine the six polynomial coefficients  $a$ ,  $b$ ,  $c$ ,  $d$ ,  $e$ ,  $f$  in Equation 5.6. The DCP values are then estimated using a quadratic polynomial surface of the form:

$$f(x, y) = a + bx + cy + dx^2 + ey^2 + fxy \quad (5.6)$$


The coefficients  $a$ ,  $b$ ,  $c$ ,  $d$ ,  $e$ ,  $f$  are computed by solving a linear system that fits the surface to the five known data points. These results for each cluster can be found in Appendix D.5.

### Visualization


After simulating the indicated positions out of the 184 and applying the interpolation method for the rest, all the DCP values for the 184 modules in place are known. Plotting these, a contour figure of the daily variation of the percentage of DCP reduction due to shading can be constructed.

Here, two of these plots are presented. Figure 5.3a, which uses an hourly irradiance weather data set, and Figure 5.3b which is based on a 10-minute irradiance dataset. The 10-minute irradiance data derived through linear interpolation of the original hourly dataset, which was constructed by Sathya Shanka Vasuki.

Glancing at the figures, it is seen that higher percentage losses, represented by red and yellow regions in the contour plot, indicate significant shading effects, leading to a greater reduction in energy yield. Conversely, lower percentage losses, shown in blue regions, correspond to locations where PV modules experience minimal shading impact, making them more suitable for installation. The excluded grid positions are set to zero, as no PV modules can be installed there.



(a) Contour plot with hourly weather data.



(b) Contour plot with 10 minute weather data.

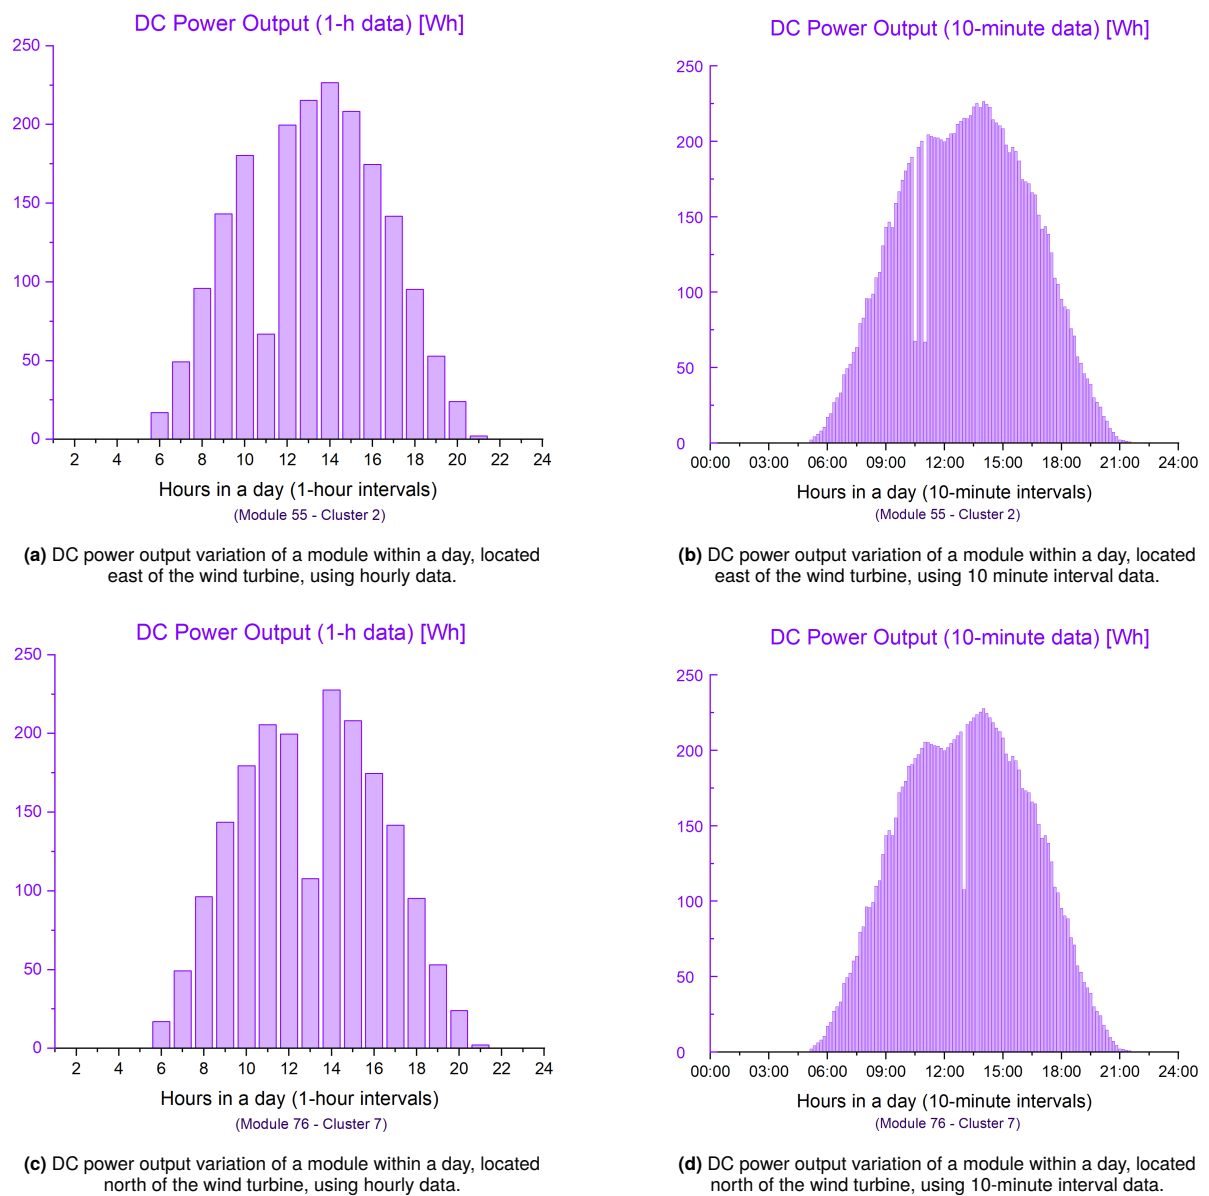
**Figure 5.3:** Comparison of contour plots for illustrating the DC power losses due to a wind turbine shadow over 184 modules, using 1h and 10 minute resolution data.

### Analysis & Interpretation

A first observation is that for the hourly data case, the average loss is 1.74 % and a median of 1.04 %, whereas for the 10-minute data, the average increases to 2.04 % and a median of 1.95 %. It is worth mentioning that the maximum loss observed with the hourly data is 7.5 %, but for the 10-minute data set the maximum drops to 4.20 %.

In the case of using hourly irradiance data, the difference between the average percentage losses and the maximum is very big, higher than 5 %. On the contrary, for the 10-minute data set this difference is only about 2 %. This suggests that using coarser weather data may be less accurate, as it can both under and over estimate the duration of apparent irradiance on a module.

In more detail, shading conditions significantly change within an hour. So when a module appears to be shaded at one time step in hourly data, it is assumed to be shaded for the entire hour, leading to overestimation of the DC power drop. This can be seen in Figure 5.4, where the daily variation of the DC power output of two modules is illustrated. One module is placed east and the other north of the wind turbine.




**Figure 5.4:** Contour plots for DC power output variation within a day for two modules, with 1h and 10min resolution, showing how shading conditions differ with the weather data resolution.



From these figures, it is observed that in both cases the shading moment is captured, but in the case of hourly data, it is assumed that the module is shaded for the entire hour, whereas in the more realistic 10-minute data case, the shading only lasts for a single time interval, i.e., 10 minutes.

For the same reason, the energy yield can also be underestimated because short-duration shading events may not be captured at all. A module may appear unshaded throughout the day in hourly datasets if a shadow passes over it briefly, for a few minutes, and this is not recorded in the hourly sample. On the other hand, with 10-minute irradiance data there is a more uniform pattern and the difference between average and maximum loss due to shading is not that extreme. This is seen in the case of a module placed west of the wind turbine, depicted in Figure 5.5. Thus, it can be said that it captures more shading moments, resulting in more consistent losses across modules and a better representation of the actual energy impact.



**Figure 5.5:** Contour plots for DC power output variation within a day for a module, where with 1h the yield is overestimated compared to 10min resolution.


A detailed analysis computing the solar elevation angle and the resulting wind turbine tower shadow length throughout 1 June in Delft was conducted and can be found in Appendix D.6. First, the solar elevation angles for the day of June 1 were found with 1-minute intervals, and they were used to compute the shadow length, which was calculated using Eq. D.1. Then, using Eq. D.2, the shadow-tip velocity was approximated using finite-difference derivatives.

The shadow-tip speed analysis revealed that the turbine shadow moves with a median velocity of approximately  $0.019 \text{ m s}^{-1}$ , which corresponds to a displacement of around 68 m over the course of one hour. Based on the median shadow-tip speed of  $0.019 \text{ m s}^{-1}$ , a shadow typically covers a single PV module in about 52 s (Eq. D.4). To move across the 12.5 m spacing between adjacent modules, it takes approximately 11 min (Eq. D.6). Hence, if irradiance data are sampled at a coarse temporal resolution (e.g., hourly), such brief shading events may be entirely missed, leading to an underestimation of the resulting energy yield loss.

It should be noted that this approximation of the shadow-tip speed is not completely precise. Normally, the path of the shadow is not linear, but it traces an arc across the ground due to the Sun's changing position. However, this is an indication of how quickly the shadow boundary shifts. Moreover, this analysis in Appendix D.6 considers only the turbine tower shadow, whereas the simulation results also account for shading from the stationary blades. However, the tower-based estimate still provides a clear indication of how quickly the shadow boundary shifts.

Moreover, it can be justified why, in both cases, the shadow losses are less in front of the turbine and bigger on the sides. Given that the turbine is oriented facing north and the site is located in the Northern Hemisphere, the sun appears to be in the southern sky [55]. This geometry results in the turbine casting its shadow primarily toward the north, and only during early morning or late afternoon is the shadow extended south of the turbine. That is why also in the Example Case D, there are no modules placed south of the turbine.

When the solar elevation angle reaches its peak at 13:45 local time, the shadow cast by the 165 m wind turbine tower is shortest, measuring approximately 94.90 m. The nearest PV rows sit 100 m or more from the base, so the tower shadow never reaches them. The modules placed north of the turbine show the lowest DC power loss. An example of the DC daily power variation of a module placed in the furthest back row of the layout is depicted in Figure 5.6. However, there are also the blades that exceed the height of the tower (in different directions also) that have not been taken into account in the shadow length analysis, so there are still losses.



**Figure 5.6:** Contour plots for DC power output variation within a day for a module, with 1h and 10min resolution. North of the turbine, the shadow influences less the energy yield.

Excluding early afternoon, when there is the peak, so the shortest shadow, the shadow of the turbine tower is much longer during most of the day. From the moment sunlight appears at 05:37 until sunset, the shadow remains long, except around midday when it's shortest, so it can reach the PV modules, as seen in Figures 5.5, 5.4. That might explain why we see higher losses on the modules located near the edges, especially in the late morning and early afternoon. Also, in the very early morning or very late evening, the shadow is extremely long, but the irradiance is also very low, so the impact of shading on energy yield isn't that significant.

## 5.4. Conclusions

This chapter addressed Sub-question 4.: *What are the performance benefits of the developed method, and how can it be further optimized? Furthermore, how does wind turbine shading influence spatial energy yield variability across a large-scale PV layout?*

The accuracy of the proposed method is validated by extensively analyzing an example case. The findings show that the selective skydome refinement approach provides notable performance advantages by cutting down on computation time while preserving high energy yield estimation accuracy. In the validation case, it was shown that the original full refinement sensitivity analysis of  $N_{\text{ref}} = 6$  required approximately 63.8 hours of computation, while the selective refinement method starting from  $N_{\text{ref}} = 2$

to  $N_{\text{ref}} = 6$  reduced this to 2.34 hours, a reduction of 92.2% to 96.9%, depending on the runtime environment. Three cases, two days and a year, were examined to investigate the corresponding DC energy yield and the deviation in energy yield remained below 1% across daily and annual simulations, with relative errors ranging from -0.12% to 0.85%. These results confirm that the method provides an effective tradeoff between speed and accuracy, allowing high-resolution shading representation only where needed. Thus, the selective refinement technique greatly enhances the efficiency of the PVMD Toolbox without compromising the reliability of energy yield predictions.

Furthermore, it was examined which starting resolutions for the selective skydome refinement process,  $N_{\text{ref}} = 2$  and  $N_{\text{ref}} = 3$ , in order to determine the optimal tradeoff between computational efficiency and energy yield accuracy. Starting from  $N_{\text{ref}} = 3$  ensures consistent shading detection, unlike  $N_{\text{ref}} = 2$ , which can miss shaded regions due to its coarse resolution. In terms of computational time, the refinement from  $N_{\text{ref}} = 2$  to 6 led to time savings ranging from 92.2% to 96.9%, while  $N_{\text{ref}} = 3 \rightarrow 6$  achieved slightly better reductions, between 96.1% and 98.75%. Although the relative error in annual DC energy yield was slightly lower for the  $N_{\text{ref}} = 2$  case (0.09%) than for  $N_{\text{ref}} = 3$  (0.45%), both were well below 1%, confirming acceptable accuracy. These results highlight that  $N_{\text{ref}} = 3$  provides a more robust and consistent basis for refinement, avoiding the pitfalls of under-detection while achieving greater time savings, making it the preferred starting level for future simulations.

This section's findings show that the spatial variation in energy yield between adjacent modules stays relatively low in a system of modules that are shaded at different times of the day by the turbine. With error margins below 0.0022 kWh, the intra-cell and intra-cluster analyses showed maximum relative deviations below 1%. As a result, one centrally located module per grid cell is adequate to accurately capture the shading impact throughout the layout, allowing for effective large-scale simulations through the use of the selective refinement method. Finally, the method's application to a larger PV-wind turbine layout demonstrated distinct spatial patterns in shading-induced energy loss, which were successfully captured through the use of interpolation and selective simulations. The contour plots that were produced demonstrated how using more precise weather data results in yield predictions that are more reliable and consistent throughout the layout. By highlighting areas with little exposure to shade, these insights can help with early-stage PV placement decisions in hybrid parks.

# Conclusions and Recommendations

The objective of this MSc thesis, as introduced in Chapter 1 and addressed throughout the subsequent chapters, was:

**“Enhancing the PVMD Toolbox for Efficient Shading Sensitivity Analysis through a Method of Selective Skydome Refinement, Applicable to Static Shading Simulations.”**

Section 6.1 summarizes how the research addressed the four sub-questions that support the main objective. This is followed by Section 6.2, which presents recommendations for future work.

## 6.1. Conclusions

This section revisits the four research sub-questions introduced in Chapter 1, summarizing how each has been addressed throughout the thesis.

### 1. RQ 1. How is a hybrid wind–solar simulation environment, including a detailed wind turbine, constructed and configured in the PVMD Toolbox?

In order to simulate the shading effects of a wind turbine in a hybrid PV–wind system, a comprehensive 3D environment was constructed in the PVMD Toolbox. A script-based method was used to develop the wind turbine’s tower, nacelle, and blades geometry. The environment setup, PV system design, weather information, and electrical behavior are all specifically documented, as are all other pertinent simulation parameters and assumptions. Three reference cases (A, B, and C) are established that are consistently used in the thesis for the rest of the analysis.

### 2. RQ 2. How can the current version of the PVMD Toolbox model and analyze shading effects of a wind turbine on a PV module, and how does skydome discretization affect accuracy and performance?

The PVMD Toolbox performs shading analysis using a backward ray tracing approach, where rays are traced in reverse from each PV cell toward the sky. Each ray corresponds to the centroid of a triangular patch within the hemispherical skydome. Sensitivity values are then calculated to quantify the irradiance reaching the module from each sky direction, considering whether rays are blocked by objects in the environment. The angular resolution of these sensitivity maps, and thus the accuracy of the shading simulation, is determined by the skydome refinement level.

The trade-off between simulation performance and skydome discretization level was assessed through a thorough computational analysis. Raising the refinement level  $N_{\text{ref}}$  increases the sensitivity map’s angular resolution, improving energy yield accuracy and shading representation. But the computational cost is high: for a single module, simulation time goes from less than a minute at  $N_{\text{ref}} = 1$  to more than 2.5 days at  $N_{\text{ref}} = 6$ , rendering high-resolution simulations unfeasible. These results verify that uniform high-resolution techniques are computationally

prohibitive, even though the PVMD Toolbox can accurately model static shading effects using skydome-based sensitivity analysis and backward ray tracing. This emphasizes the necessity of more focused refinement techniques. A new selective refinement technique that maintains accuracy while drastically cutting down on simulation time is presented in the following chapter.

**3. RQ 3. What enhancements can be introduced in the PVMD toolbox to improve the accuracy and efficiency of the shading analysis of a PV module with respect to its surroundings?**

To increase the computational efficiency of sensitivity analysis in the PVMD Toolbox, a *Selective Skydome Refinement* technique was created. The technique *selectively* refines only the sky regions that contribute to shading, allowing for high resolution sensitivity analysis. This is accomplished by first executing a low-resolution simulation to identify sky triangles that have zero or extremely low initial sensitivity, indicating that they are shaded. To guarantee adequate coverage, a buffer zone is added outside of these triangles. The skydome is only refined to the desired level of discretization after which the sensitivity values for this last subset of triangles are recalculated using a second ray tracing, while the initial coarse-level values are replicated for the remaining triangles.

The gain in computational time comes from the fact that by focusing only on the most relevant regions of the skydome, the selective refinement method drastically lowers the number of rays traced and, consequently, the overall computation time. In the example presented, the number of rays traced at high refinement was reduced by over 70%, and the simulation time by more than 90%. A key advantage of this method is that it integrates easily into the existing PVMD Toolbox framework, requiring only minimal changes to the current simulation structure.

**4. RQ 4. What are the performance benefits of the developed method, and how can it be further optimized? Furthermore, how does wind turbine shading influence spatial energy yield variability across a large-scale PV layout?**

A thorough example case confirmed the accuracy of the suggested approach. The findings show that selective skydome refinement offers notable performance advantages by reducing computation time while maintaining high energy yield accuracy. In the validation case, full-resolution analysis at  $N_{\text{ref}} = 6$  took 63.8 hours, while the selective method ( $N_{\text{ref}} = 2 \rightarrow 6$ ) reduced this to 2.34 hours, a reduction of 92.2% to 96.9%, depending on the system. Energy yield comparisons over two days and a year showed relative errors between -0.12% and 0.85%, all below 1%. These results confirm that the method strikes an effective balance between accuracy and efficiency, enabling high-resolution shading analysis only when needed and greatly enhancing PVMD Toolbox performance.

Starting resolutions  $N_{\text{ref}} = 2$  and  $N_{\text{ref}} = 3$  were compared to find the optimal tradeoff. While  $N_{\text{ref}} = 2$  may miss shaded regions due to coarser resolution, starting from  $N_{\text{ref}} = 3$  ensures more reliable detection. Computational time was reduced by 92.2% to 96.9% for  $N_{\text{ref}} = 2 \rightarrow 6$  and slightly more, 96.1% to 98.75%, for  $N_{\text{ref}} = 3 \rightarrow 6$ . Both provided acceptable accuracy, though annual yield error was slightly lower for  $N_{\text{ref}} = 2$  (0.09%) compared to  $N_{\text{ref}} = 3$  (0.45%). Overall,  $N_{\text{ref}} = 3$  offers a more robust and trustworthy refinement starting point by combining better detection with greater time savings.

## 6.2. Future Work Recommendations

The current model has significantly improved the feasibility of conducting high-resolution sensitivity analyses within the MODULE component of the toolbox by balancing computational efficiency with accurate energy yield estimations. Nevertheless, there is still room for improvement. Enhancing various elements of the implementation and simulation pipeline could enhance flexibility and performance.

The interpretation of the obstacle in the environment is important. Specifically for hybrid wind-solar analysis, the wind turbine shape is quite complex and with the current method described in 2.1.2 only approximates in simpler shapes. To accurately capture the shape of the wind turbine, one possible improvement would be to use the actual geometry from a real turbine design. For example, coordinates could be extracted from a manufacturer's CAD model and directly inserted into the toolbox.

By now, it is known that in the skydome is uniformly refined, but sensitivity values outside the recalculated region are replicated from parent triangles, giving the visual impression of large, blocky areas despite the fine mesh. Although the replicated regions do not contribute to shading, as they are always clean, it is undeniable that the repetition in values impacts the overall accuracy of the estimated DC power yield, as discussed in the previous chapter.

An idea to tackle this can be to inject some degradation inside the replicated children to make the transitions smoother and more realistic. This could be done by first calculating the average sensitivity of each parent triangle and its neighboring ones. Their comparison reveals which sides of the patch are adjacent to brighter or darker regions. Then, the children triangles inside the parent patch can be grouped into two or three zones. Each of these zones is assigned a value, through interpolation, based on the range of the values of the neighboring parent triangles. This way, a softer fade can be achieved to visually blend the transitions between neighboring patches. And thus resembles the skydome's natural intensity gradient, without additional ray tracing.

Another potential limitation lies in how the sun is modeled in the current framework of the toolbox. The Perez sky model is used to construct the luminance distribution (both circumsolar and diffuse) of the sky [44, 46]. It takes the DNI, DHI, and the sun's position (azimuth and altitude) from the weather data and it compares it to sky facets. It locates and matches it to the sky triangle that points closest to the sun's position. All the direct sunlight (DNI) is added to just this triangle, the one closest to the sun's direction. As a result, the sun behaves like a point source, with only one triangle receiving the full beam irradiance, while the rest of the sky has only diffuse brightness.

This simplification arises due to the angular resolution of the skydome, which is determined by the refinement level  $N_{\text{ref}}$ . For example,  $N_{\text{ref}} = 2$  generates 160 triangular sky facets, each covering roughly  $9^\circ$ , while  $N_{\text{ref}} = 6$  produces 40,960 triangles, with each spanning about  $0.6^\circ$ . Even at this resolution, the angular size of a triangle remains slightly larger than the sun's true angular diameter of approximately  $0.53^\circ$  [56, 57], meaning the sun's disk typically fits within a single triangle, concentrating all DNI into one direction. While computationally efficient, this introduces inaccuracies in irradiance and shading modeling.

A more realistic alternative would be to spread the DNI across several sky triangles instead of assigning it all to just one. This can be done by defining a small cone (angular region) around the sun's position in the sky. Then, the angular distance  $\gamma$  between the center of each triangle and the sun can be calculated, and triangles whose direction lies within a chosen cone angle are considered to be part of the sun's visible disk. The DNI is then divided among these triangles instead of just one. To make this even more realistic, the distribution can be weighted so that triangles closer to the center of the sun's disk receive more irradiance. A common choice for this is a Gaussian function, which smoothly spreads the light while making sure that the total amount of energy remains the same.

Last but not least, since this thesis is focused on the static shading analysis, a promising next step would be to extend this work toward dynamic shading. This would take into consideration the turbine parts, like yawing nacelles or rotating blades, that move over time. The environment has already been adapted to allow separate rotation of these components, providing a foundation for such development. The orientation of the nacelle can be adjusted to simulate different wind directions, and the blades can be rotated independently to represent their real motion. However, as it is, the toolbox cannot handle environments that change over time. Enabling simulations with time-varying geometry would be a significant step toward more realistic modeling in hybrid renewable systems. A practical fix would be to run several static simulations at various time steps, like taking "snapshots" of the environment, and evaluating the shading effects in each case.

# Bibliography


- [1] I. R. E. A. (IRENA). "Global renewables outlook: Energy transformation 2050." Accessed: 2024-12-12. (2020), [Online]. Available: <https://www.irena.org/Publications/2020/Apr/Global-Renewables-Outlook-2020>.
- [2] Council of the European Union. "Fit for 55." Accessed: 2025-05-22, Council of the European Union. (2023), [Online]. Available: <https://www.consilium.europa.eu/en/policies/fit-for-55/>.
- [3] E. Commission and D.-G. for Climate Action, *Going climate-neutral by 2050 – A strategic long-term vision for a prosperous, modern, competitive and climate-neutral EU economy*. Publications Office, 2019. DOI: doi/10.2834/02074.
- [4] Government of the Netherlands. "Step by step, the netherlands is transitioning to sustainable energy." Accessed: 2025-05-22, Government of the Netherlands. (2024), [Online]. Available: <https://www.government.nl/topics/renewable-energy/step-by-step-the-netherlands-is-transitioning-to-sustainable-energy>.
- [5] Netherlands Enterprise Agency (RVO), TKI Urban Energy, Holland Solar, and Energy Storage NL. "Solar energy and storage guide 2024." Accessed: 2025-05-22, Netherlands Enterprise Agency (RVO). (2024), [Online]. Available: [https://english.rvo.nl/sites/default/files/2024-07/Solar%20Energy%20Storage%20Guide%202024\\_0.pdf](https://english.rvo.nl/sites/default/files/2024-07/Solar%20Energy%20Storage%20Guide%202024_0.pdf).
- [6] Netbeheer Nederland. "Integrale infrastructuurverkenning 2030 - 2050." Accessed: 2025-02-28. (2023), [Online]. Available: <https://www.energiekompas2050.nl/wp-content/uploads/2024/01/II3050-Main-Report-20240118-1.pdf>.
- [7] Netbeheer Nederland. "Het energiesysteem van de toekomst: De ii3050-scenario's." Accessed: 2025-02-28. (2023), [Online]. Available: [https://www.netbeheernederland.nl/sites/default/files/Rapport\\_II3050\\_Scenario%2527s\\_280.pdf](https://www.netbeheernederland.nl/sites/default/files/Rapport_II3050_Scenario%2527s_280.pdf).
- [8] Netherlands Enterprise Agency (RVO). "Sde++ 2022: Stimulation of sustainable energy production and climate transition." Accessed: 2025-05-22, Netherlands Enterprise Agency. (2022), [Online]. Available: <https://english.rvo.nl/sites/default/files/2023-08/English%20brochure%20SDE%2B%2B%202022%20-%2020juli%202022.pdf>.
- [9] L. Essak and A. Ghosh, "Floating photovoltaics: A review," *Clean Technologies*, vol. 4, no. 3, pp. 752–769, 2022, ISSN: 2571-8797. DOI: 10.3390/cleantechnol4030046. [Online]. Available: <https://www.mdpi.com/2571-8797/4/3/46>.
- [10] C. Netherlands. "Land use in the netherlands (2015)." Accessed: 2025-03-01. (2015), [Online]. Available: <https://www.clo.nl/en/indicators/en006111-land-use-in-the-netherlands-2015>.
- [11] Rijksoverheid. "Zonne-energie - rijksoverheid.nl." Accessed: 2025-03-01. (2025), [Online]. Available: <https://www.rijksoverheid.nl/onderwerpen/duurzame-energie/zonne-energie>.
- [12] Euractiv, *The netherlands, unquestionable solar energy leader of 2022: Study*, Accessed: 2025-02-28, 2022. [Online]. Available: <https://www.euractiv.com/section/politics/news/the-netherlands-unquestionable-solar-energy-leader-of-2022-study/>.
- [13] Energy Storage News, *Wind-solar-storage hybrid project with 12mwh bess online in netherlands*, <https://www.energy-storage.news/wind-solar-storage-hybrid-project-with-12mwh-bess-online-in-netherlands/>, Accessed: 2025-06-19, 2024.
- [14] E. Bellini. "Model agreement to make solar and wind share same connection point in the netherlands." Accessed: 2025-06-26, pv magazine. (2021), [Online]. Available: <https://www.pv-magazine.com/2021/04/19/model-agreement-to-make-solar-and-wind-share-same-connection-point-in-the-netherlands/>.

- [15] Vattenfall, *Vattenfall's largest hybrid energy park is taking shape in the netherlands*, <https://group.vattenfall.com/press-and-media/newsroom/2020/vattenfalls-largest-hybrid-energy-park-is-taking-shape-in-the-netherlands>, Accessed: 2025-06-19, 2020.
- [16] E. Bellini, "Shell to build 100 mw hybrid wind-solar project," *pv magazine*, May 2022, Accessed: 2025-06-19. [Online]. Available: <https://www.pv-magazine.com/2022/05/24/shell-to-build-100-mw-hybrid-wind-solar-project/>.
- [17] Oceans of Energy, *Oceans of energy expands floating solar deployment in north sea*, <https://oceansofenergy.blue/2025/05/20/11074/>, Accessed: 2025-06-19, 2025.
- [18] Q. Hassan, S. Algburi, A. Z. Sameen, H. M. Salman, and M. Jaszczur, "A review of hybrid renewable energy systems: Solar and wind-powered solutions: Challenges, opportunities, and policy implications," *Results in Engineering*, vol. 20, p. 101621, 2023, issn: 2590-1230. DOI: <https://doi.org/10.1016/j.rineng.2023.101621>.
- [19] J. Huang and G. Iglesias, "Hybrid offshore wind-solar energy farms: A novel approach through retrofitting," *Energy Conversion and Management*, vol. 319, p. 118903, 2024, issn: 0196-8904. DOI: <https://doi.org/10.1016/j.enconman.2024.118903>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0196890424008446>.
- [20] H. Khurshid, B. S. Mohammed, A. M. Al-Yacoubi, M. Liew, and N. A. W. A. Zawawi, "Analysis of hybrid offshore renewable energy sources for power generation: A literature review of hybrid solar, wind, and waves energy systems," *Developments in the Built Environment*, vol. 19, p. 100497, 2024, issn: 2666-1659. DOI: <https://doi.org/10.1016/j.dibe.2024.100497>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2666165924001789>.
- [21] X. Costoya, M. deCastro, D. Carvalho, B. Arguilé-Pérez, and M. Gómez-Gesteira, "Combining offshore wind and solar photovoltaic energy to stabilize energy supply under climate change scenarios: A case study on the western iberian peninsula," *Renewable and Sustainable Energy Reviews*, vol. 157, p. 112037, 2022, issn: 1364-0321. DOI: <https://doi.org/10.1016/j.rser.2021.112037>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1364032121012995>.
- [22] O. Delbeke, J. D. Moschner, and J. Driesen, "The complementarity of offshore wind and floating photovoltaics in the belgian north sea, an analysis up to 2100," *Renewable Energy*, vol. 218, p. 119253, 2023, issn: 0960-1481. DOI: <https://doi.org/10.1016/j.renene.2023.119253>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0960148123011680>.
- [23] S. Golroodbari *et al.*, "Pooling the cable: A techno-economic feasibility study of integrating offshore floating photovoltaic solar technology within an offshore wind park," *Solar Energy*, vol. 219, pp. 65–74, 2021, Special Issue on Floating Solar: beyond the state of the art technology, issn: 0038-092X. DOI: <https://doi.org/10.1016/j.solener.2020.12.062>.
- [24] G. Ren, W. Wang, J. Wan, F. Hong, and K. Yang, "A novel metric for assessing wind and solar power complementarity based on three different fluctuation states and corresponding fluctuation amplitudes," *Energy Conversion and Management*, vol. 278, p. 116721, 2023, issn: 0196-8904. DOI: <https://doi.org/10.1016/j.enconman.2023.116721>.
- [25] N. J. Dekker *et al.*, "Wind turbine dynamic shading: The effects on combined solar and wind farms," *Journal of Renewable and Sustainable Energy*, vol. 15, no. 6, p. 063703, Dec. 2023, issn: 1941-7012. DOI: [10.1063/5.0176121](https://doi.org/10.1063/5.0176121). [Online]. Available: <https://doi.org/10.1063/5.0176121>.
- [26] M. López, N. Rodríguez, and G. Iglesias, "Combined floating offshore wind and solar pv," *Journal of Marine Science and Engineering*, vol. 8, no. 8, 2020, issn: 2077-1312. DOI: [10.3390/jmse8080576](https://doi.org/10.3390/jmse8080576).
- [27] C. Tripp, D. Guittet, J. King, and A. Barker, "A simplified, efficient approach to hybrid wind and solar plant site optimization," *Wind Energy Science*, vol. 7, no. 2, pp. 697–713, 2022. DOI: [10.5194/wes-7-697-2022](https://doi.org/10.5194/wes-7-697-2022).
- [28] J. Robledo, J. Leloux, B. Sarr, C. Gueymard, and P. Darez, "Dynamic simulation of the shading cast by a wind farm on an adjacent photovoltaic plant," *EU PVSEC Proceedings*, Sep. 2021. DOI: [10.4229/EUPVSEC20212021-5CV.3.8](https://doi.org/10.4229/EUPVSEC20212021-5CV.3.8).



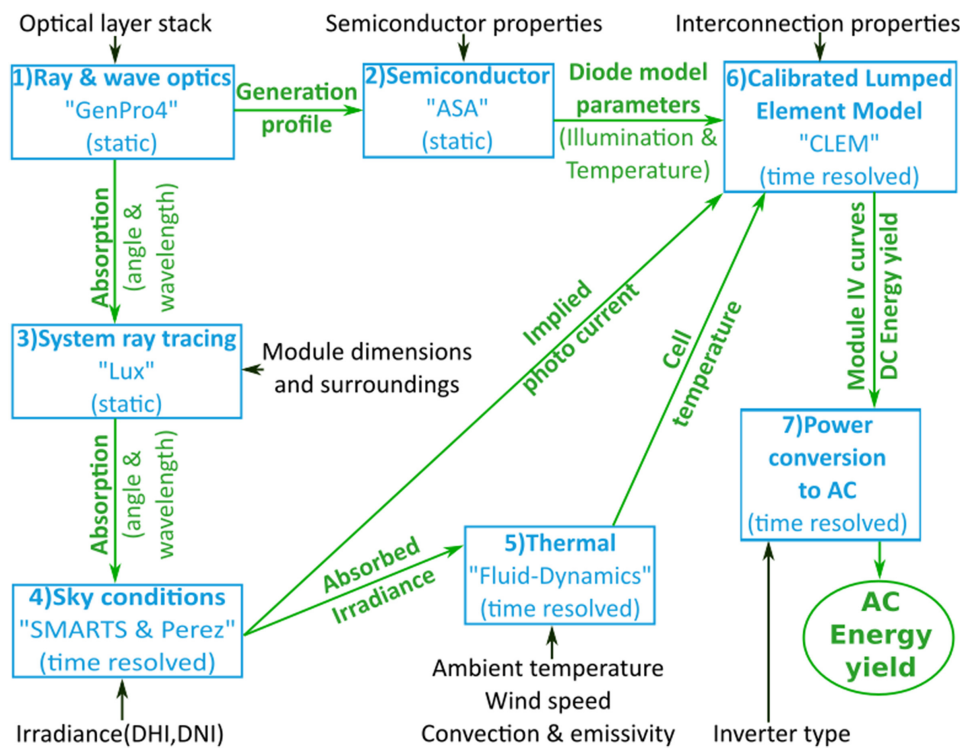
- [29] C. Nazir, "Techno- economics of a hybrid photovoltaic wind turbine," *SCIREA Journal of Energy*, Dec. 2021. doi: 10.54647/energy48149.
- [30] TNO. "Shading effects on solar panels caused by wind turbines." Accessed: 2024-12-16. (2023), [Online]. Available: <https://www.tno.nl/en/newsroom/2023/12/shading-solarpanels-windturbines/>.
- [31] TNO. "Research on the effects of wind turbine shade." Accessed: 2024-12-16. (2021), [Online]. Available: <https://www.tno.nl/en/newsroom/2021/02/research-effects-wind-turbine-shade/>.
- [32] S. R. Mohanrajan and V. Hamsadhwani, "Optimal photovoltaic capacity planning for windfarm expansion," *International Journal of Engineering and Advanced Technology*, vol. 8, no. 5, pp. 2126–2131, 2019.
- [33] A. Routray, K. Mistry, S. R. Arya, and B. Chittibabu, "Power output evaluation of a wind–solar farm considering the influence parameters," *IET Renewable Power Generation*, vol. 15, pp. 1613–1623, May 2021. doi: 10.1049/rpg2.12018.
- [34] S. Gat, R. Ben-Arie, and A. Liberman-Shalev, "Decision support tools for improving land footprint and power generation of a photovoltaic field by the deployment of wind turbines in the same designated area: Eilat district as a case study," *Environment, Development and Sustainability*, vol. 25, no. 11, pp. 12 495–12 526, 2023. doi: 10.1007/s10668-022-02576-0.
- [35] S. Shanghavi, W. Mack Grady, and B. Schwarz, "Evaluating the impact of wind turbine shadows on an integrated wind and solar farm," in *2012 3rd IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*, 2012, pp. 1–6. doi: 10.1109/ISGTEurope.2012.6465711.
- [36] D. Ludwig, S. Asfaw, and R. Seguin, "Evaluation of an onsite integrated hybrid pv-wind power plant," *AIMS Energy*, vol. 8, pp. 988–1006, Oct. 2020. doi: 10.3934/energy.2020.5.988.
- [37] A. Benson, H. Darling, and R. Whitten, "Pre-construction modeling of shade loss due to wind turbine shadow flicker on solar pv," *IEEE Photovoltaic Specialists Conference (PVSC)*, pp. 1231–1233, Jun. 2024. doi: 10.1109/PVSC57443.2024.10749295.
- [38] I. Mamia and J. Appelbaum, "Shadow analysis of wind turbines for dual use of land for combined wind and solar photovoltaic power generation," *Renewable and Sustainable Energy Reviews*, vol. 55, pp. 713–718, 2016, issn: 1364-0321. doi: <https://doi.org/10.1016/j.rser.2015.11.009>.
- [39] M. Vogt *et al.*, "Introducing a comprehensive physics-based modelling framework for tandem and other pv systems," *Solar Energy Materials and Solar Cells*, vol. 247, p. 111 944, 2022, issn: 0927-0248. doi: <https://doi.org/10.1016/j.solmat.2022.111944>.
- [40] E. Gaertner *et al.*, "Definition of the IEA 15-megawatt offshore reference wind turbine," National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-5000-75698, Mar. 2020. [Online]. Available: <https://www.nrel.gov/docs/fy20osti/75698.pdf>.
- [41] V. W. S. A/S. "V236-15.0 mw offshore wind turbine." Accessed: 2025-02-17. (2024), [Online]. Available: [https://ludwiglorenz.ch/images/V236-150\\_MW\\_brochure.pdf](https://ludwiglorenz.ch/images/V236-150_MW_brochure.pdf).
- [42] T. Takatani, Y. Mukaigawa, Y. Matsushita, and Y. Yagi, "Decomposition of reflection and scattering by multiple-weighted measurements," *IPSI Transactions on Computer Vision and Applications*, vol. 10, Dec. 2018. doi: 10.1186/s41074-018-0049-4.
- [43] R. Energy. "Rsm72-6-360-380m datasheet." Accessed: 2025-02-17. (2019), [Online]. Available: <https://betsolar.es/wp-content/uploads/2020/03/RSM72-6-360-380M-IECUL1500V-40mm-2019H1-4-EN.pdf>.
- [44] R. Santbergen, V. Muthukumar, R. Valckenborg, W. van de Wall, A. Smets, and M. Zeman, "Calculation of irradiance distribution on pv modules by combining sky and sensitivity maps," *Solar Energy*, vol. 150, pp. 49–54, 2017, issn: 0038-092X. doi: <https://doi.org/10.1016/j.solener.2017.04.036>.
- [45] D. V. Schroeder. "Seasonal beam spreading diagram." Licensed under CC BY-NC 3.0. Accessed: 2025-07-07. (n.d.), [Online]. Available: <https://www.saltergrove.org/weather/>.

- [46] R. Perez, R. Seals, and J. Michalsky, "All-weather model for sky luminance distribution—preliminary configuration and validation," *Solar Energy*, vol. 50, no. 3, pp. 235–245, 1993, ISSN: 0038-092X. DOI: [https://doi.org/10.1016/0038-092X\(93\)90017-I](https://doi.org/10.1016/0038-092X(93)90017-I).
- [47] A. Smets, K. Jäger, O. Isabella, R. van Swaaij, and M. Zeman, *Solar Energy: The physics and engineering of photovoltaic conversion, technologies and systems*, English. UIT Cambridge Limited, 2016, ISBN: 978-1-906860-32-5.
- [48] Firebird Optics. "The basic guide to lambertian reflectance." Accessed: 2025-05-27. (2024), [Online]. Available: <https://www.firebirdoptics.com/blog/the-basic-guide-to-lambertian-reflectance>.
- [49] R. T. Marcus, "Chapter 2 - the measurement of color," in *Color for Science, Art and Technology*, ser. AZimuth, K. Nassau, Ed., vol. 1, North-Holland, 1998, pp. 31–96. DOI: [https://doi.org/10.1016/S1387-6783\(98\)80005-6](https://doi.org/10.1016/S1387-6783(98)80005-6).
- [50] Wikipedia contributors, *Student's t-distribution — wikipedia, the free encyclopedia*, Accessed: 2025-06-20, 2024. [Online]. Available: [https://en.wikipedia.org/wiki/Student%27s\\_t-distribution](https://en.wikipedia.org/wiki/Student%27s_t-distribution).
- [51] Study.com, *How to determine the margin of error for a one-sample t-interval*, Accessed: 2025-06-20, 2024. [Online]. Available: <https://study.com/skill/learn/how-to-determine-the-margin-of-error-for-a-one-sample-t-interval-explanation.html>.
- [52] Wikipedia contributors, *Polynomial regression — Wikipedia, the free encyclopedia*, Accessed: 2025-06-23, 2024. [Online]. Available: [https://en.wikipedia.org/wiki/Polynomial\\_regression](https://en.wikipedia.org/wiki/Polynomial_regression).
- [53] H. McIntosh, *Two dimensional quadratic approximation*, Accessed: 2025-06-17, 2000. [Online]. Available: <https://delta.cs.cinvestav.mx/~mcintosh/comun/contours/node8.html>.
- [54] MIT OpenCourseWare, *Polynomial interpolation – multivariable calculus*, Accessed: 2025-06-17, 2005. [Online]. Available: <https://ocw.mit.edu/ans7870/18/18.013a/textbook/HTML/chapter11/section01.html>.
- [55] Wikipedia contributors, *Sun path — wikipedia, the free encyclopedia*, [https://en.wikipedia.org/wiki/Sun\\_path](https://en.wikipedia.org/wiki/Sun_path), Accessed: 2025-06-24, 2025.
- [56] N. Janotte, S. Wilbert, F. Sallaberry, M. Schroedter-Homscheidt, and L. Ramirez, "2 - principles of csp performance assessment," in *The Performance of Concentrated Solar Power (CSP) Systems*, P. Heller, Ed., Woodhead Publishing, 2017, pp. 31–64, ISBN: 978-0-08-100447-0. DOI: <https://doi.org/10.1016/B978-0-08-100447-0.00002-X>.
- [57] WebAssign. "Angular size and distance." Accessed: 2025-05-21. (n.d.), [Online]. Available: <https://www.webassign.net/seedfoundations/ebook/CH03-3.html>.
- [58] D. Koblick, *Vectorized solar azimuth and elevation estimation*, <https://www.mathworks.com/matlabcentral/fileexchange/23051-vectorized-solar-azimuth-and-elevation-estimation>, MATLAB Central File Exchange. Retrieved June 23, 2025, 2025.
- [59] F. Sandnes, "Determining the geographical location of image scenes based on object shadow lengths," *Signal Processing Systems*, vol. 65, pp. 35–47, Oct. 2011. DOI: 10.1007/s11265-010-0538-x.
- [60] H. S. University, *Introduction to air photo interpretation*, [https://gsp.humboldt.edu/olm/Courses/GSP\\_216/lessons/air-photo.html](https://gsp.humboldt.edu/olm/Courses/GSP_216/lessons/air-photo.html), Accessed: 2025-06-23, 2025.
- [61] Wikipedia contributors, *Finite difference*, [https://en.wikipedia.org/wiki/Finite\\_difference](https://en.wikipedia.org/wiki/Finite_difference), Accessed: 2025-06-23, 2025.



## A.1. PVMD Toolbox Flowchart

The simulation model flowchart for energy yield calculation in the PVMD Toolbox is presented in Fig. A.1 into seven key pillars. It highlights the main models (blue), the information flow and intermediate results (green), and the input data (black).



**Figure A.1:** Flowchart of PVMD Toolbox describing the simulation methodology. Reproduced from [39].

# B

## B.1. Non-linear Exponential Model Used for Computational Time Fitting

The exponential growth in runtime is modeled by:

$$y = y_0 + A \cdot \exp(R_0 \cdot x) \quad (\text{B.1})$$

where:

- $y$  is the computational time in minutes,
- $x$  is the skydome refinement level,
- $y_0$ ,  $A$ , and  $R_0$  are model parameters determined through regression.

The estimated values of the model parameters above and their standard errors are summarized in Table B.1. These values were obtained through a nonlinear least squares fitting procedure. This minimizes the residual sum of squares between the measured data and the model.

**Table B.1:** Estimated parameters for the exponential model  $y = y_0 + A \cdot \exp(R_0 \cdot x)$ .

Parameter	Estimated Value	Standard Error
$y_0$	1.04177	0.97041
$A$	$5.59452 \times 10^{-5}$	$3.80207 \times 10^{-6}$
$R_0$	3.00674	0.01131
$y = 1.04 + 5.59 \times 10^{-5} \cdot \exp(3.01x)$		

The exponential model yielded a reasonably good fit to the six computed points. The model achieved an adjusted coefficient of determination  $R^2 = 1.00$ , indicating an almost perfect agreement between the predicted and observed values. Nonetheless, the perfect fit should be interpreted with caution due to the small sample size and the number of model parameters.

### C.1. Relative Errors in Energy Yield Between Levels of Full Refinements


The following table presents the relative errors in DC energy yield for each skydome discretization level, compared to the most refined case ( $N_{\text{ref}} = 6$ ), for two representative days in June and January.

**Table C.1:** Relative error (%) in DC energy yield at each level of skydome discretization, compared to  $N_{\text{ref}} = 6$ , for two representative dates.


Skydome Discretization Level	Relative Error (%) - 1 <sup>st</sup> Jun	Relative Error (%) - 1 <sup>st</sup> Jan
1	14.08%	9.07%
2	−0.06%	0.87%
3	0.72%	0.84%
4	−0.06%	0.77%
5	0.08%	0.32%
6	0.00%	0.00%

## C.2. Comparison of Sensitivity Maps for Example Case A

Figure C.1 below shows the visual validation of the selective sensitivity analysis comparing the original sensitivity map with  $N_{\text{ref}} = 6$  (C.1a) and the selectively refined one from  $N_{\text{ref}} = 2$  to 6 (C.1b).



(a) SM for skydome discretization level  $N_{\text{ref}} = 6$



(b) SM with selective refinement starting from  $N_{\text{ref}} = 2$  to  $N_{\text{ref}} = 6$

**Figure C.1:** Comparison of sensitivity maps using full versus selective skydome refinement at level  $N_{\text{ref}} = 6$ .

### C.3. Number of Refined Triangles per Module for $N_{\text{ref}} = 2, 3 \rightarrow 6$


The table below summarizes the number of parent and child sky triangles that require recalculation per module, comparing two refinement levels:  $N_{\text{ref}} = 2 \rightarrow 6$  with 2 buffer layers and  $N_{\text{ref}} = 3 \rightarrow 6$  with 3 buffer layers.

**Table C.2:** Number of refined parent triangles and total child triangles to recalculate per module for 32 modules.

Module	$N_{\text{ref}} = 2 \rightarrow 6$ (2 buffer layers)		$N_{\text{ref}} = 3 \rightarrow 6$ (3 buffer layers)	
	Refined Parent Triangles	Child Triangles	Refined Parent Triangles	Child Triangles
1	45	11520	119	7616
2	45	11520	120	7680
3	49	12544	154	9856
4	36	9216	130	8320
5	0	0	132	8448
6	48	12288	124	7936
7	55	14080	145	9280
8	20	5120	137	8768
9	26	6656	127	8128
10	26	6656	84	5376
11	0	0	138	8832
12	33	8448	102	6528
13	33	8448	128	8192
14	0	0	85	5440
15	26	6656	127	8128
16	0	0	104	6656
17	43	11008	126	8064
18	0	0	92	5888
19	26	6656	105	6720
20	0	0	116	7424
21	46	11776	140	8960
22	26	6656	133	8512
23	36	9216	117	7488
24	38	9728	104	6656
25	36	9216	125	8000
26	43	11008	126	8064
27	45	11520	118	7552
28	20	5120	114	7296
29	50	12800	135	8640
30	30	7680	89	5696
31	36	9216	129	8256
32	0	0	135	8640
<b>Average</b>		<b>7,336</b>		<b>7,720</b>

### D.1. Example Case C: Grid Analysis Setup

The Example Case C presented in 2.3.3 is created by the following logic. The  $300\text{ m} \times 300\text{ m}$  area is divided into a  $24 \times 24$  uniform grid, resulting in 576 square cells, each measuring  $12.5\text{ m} \times 12.5\text{ m}$ . These represent the 576 possible PV module positions if a single PV module is conceptually placed at the center of each grid cell.



**Figure D.1:** Grid Layout with 576 squares for PV placements



In Figure D.1, red numbers indicate valid positions, while blue numbers represent excluded positions due to placement regulations. An exclusion zone of 100 m around the wind turbine is introduced, based on clearance distance criteria discussed in Section 2.1.4. This way, the simulated scenario aligns with literature-based recommendations for scaled-down simulations like this one. Apart from that, to avoid unnecessary calculations and computational burden, grid positions located more than 40 m south of the turbine are also excluded. Since the simulation assumes a northern hemisphere location with the wind turbine facing north, the majority of sunlight comes from the southern sky. Therefore, those modules south of the turbine would rarely experience significant shading. This yields the remaining 184 modules to simulate of the Example Case C.

The 184 PV modules that are actually simulated and constitute the Example Case C(2.3.3) seen in Figure D.2.




Figure D.2: The 184 grid squares with simulated PV modules.

D.2. Single Grid Square Uncertainty Analysis

Spatial layout of the 25 PV modules simulated inside Position 406 (Grid Square 110), arranged in a 5×5 configuration within a 12.5 m × 12.5 m square. Each blue rectangle represents a module; the black dots indicate center points used for simulation.




Figure D.3: Spatial layout of the 25 PV modules within a 12.5 m × 12.5 m grid square.


For the above layout of positions for PV modules the DC power yield is calculated and presented for the 1<sup>st</sup> June. Also, the relative error of each value compared to the central position of the 12.5 m × 12.5 m grid square.

**Table D.1:** DC yield and relative error (vs. central position) for all 25 modules in Position 406 D.1 or 110 D.2.

Module	Row	Column	DCP [kWh]	Relative Error
1	1	1	2.0176	0.762%
2	1	2	2.0328	0.014%
3	1	3	2.0331	0.002%
4	1	4	2.0234	0.473%
5	1	5	2.0355	0.120%
6	2	1	2.0253	0.381%
7	2	2	2.0211	0.587%
8	2	3	2.0341	0.050%
9	2	4	2.0218	0.553%
10	2	5	2.0248	0.408%
11	3	1	2.0254	0.378%
12	3	2	2.0243	0.430%
13	3	3	2.0234	0.477%
14	3	4	2.0336	0.025%
15	3	5	2.0237	0.460%
16	4	1	2.0247	0.410%
17	4	2	2.0251	0.391%
18	4	3	2.0253	0.381%
19	4	4	2.0251	0.389%
20	4	5	2.0346	0.078%
21	5	1	2.0245	0.420%
22	5	2	2.0247	0.409%
23	5	3	2.0164	0.819%
24	5	4	2.0356	0.123%
25	5	5	2.0263	0.332%
<b>Central Position (Module 406)</b>			<b>2.0331</b>	—

### D.3. Analysis of a cluster of 16 Modules

An area 50 m × 50 m is simulated, with a spatial layout of 16 PV modules, each placed in one of the hypothetical 12.5 m × 12.5 m squares. Each blue rectangle represents a module.



**Figure D.4:** Spatial layout of the 16 PV modules within a 50 m × 50 m.

For the above layout of positions for PV modules the DC power yield is calculated and presented for the 1<sup>st</sup> June. Also, the relative error of each value compared to the central position of the 50 m × 50 m area.

**Table D.2:** DC yield and relative error (vs. central position) for all 16 modules in Batch 10.

Module	DCP [kWh]	Relative Error
109	2.0290	0.095%
110	2.0330	0.312%
111	2.0300	0.231%
112	2.0290	0.283%
113	2.0340	0.064%
114	2.0340	0.055%
115	2.0340	0.042%
116	2.0360	0.040%
117	2.0340	0.065%
118	2.0310	0.179%
119	2.0360	0.050%
120	2.0360	0.049%
121	2.0340	0.049%
122	2.0350	0.019%
123	2.0330	0.109%
124	2.0360	0.057%
<b>Central Position (Batch Center)</b>		<b>2.0350</b>

### D.4. Clustering the 184 modules

The 184 module positions are organized into clusters. Two types of clusters are defined: (a) six batches consisting of 20 modules each, primarily located at the sides of the turbine, and (b) four batches of 16 modules each, selected from other representative regions in the grid.




Figure D.5: Clustering of PV modules for simulations.

## D.5. Uncertainty analysis of clusters

For each cluster, the corner points are simulated, and the rest are estimated using 2D quadratic surface interpolation. For each, metrics for the uncertainty inside the cluster with 5 simulated points are found from the Equation 5.4 for Margin of Error, with 95% two-tailed critical value for 5 sample points equal to 2.7764.

**Table D.3:** DCP simulation metrics for clusters 1–3 and statistical summary.

	Cluster 1	Cluster 2	Cluster 3
Simulated module count	5	5	5
Cluster area [m × m]	50 × 62.5	50 × 62.5	50 × 62.5
DCP bottom-left [kWh]	2.024	2.020	2.027
DCP top-left [kWh]	1.929	2.027	2.036
DCP bottom-right [kWh]	1.992	2.023	2.024
DCP top-right [kWh]	1.971	1.894	2.035
DCP center [kWh]	1.9659	2.0170	2.0355
Mean $\mu$ [kWh]	1.9764	1.9962	2.0315
Standard deviation $s$ [kWh]	0.0349	0.0572	0.0056
Margin of error (95% CI) [kWh]	±0.0434	±0.0711	±0.0069

**Table D.4:** DCP simulation metrics for clusters 4–6 and statistical summary.

	Cluster 4	Cluster 5	Cluster 6
Simulated module count	5	5	5
Cluster area [m × m]	50 × 62.5	50 × 62.5	50 × 62.5
DCP bottom-left [kWh]	2.037	2.045	1.914
DCP top-left [kWh]	1.949	2.035	2.034
DCP bottom-right [kWh]	2.016	2.046	1.918
DCP top-right [kWh]	2.046	2.043	2.038
DCP center [kWh]	2.0466	2.0415	1.9770
Mean $\mu$ [kWh]	2.0189	2.0421	1.9762
Standard deviation $s$ [kWh]	0.0409	0.0043	0.0600
Margin of error (95% CI) [kWh]	±0.0509	±0.0053	±0.0745

**Table D.5:** DCP simulation metrics for clusters 7–10 and statistical summary.

	Cluster 7	Cluster 8	Cluster 9	Cluster 10
Simulated module count	4	4	4	16
Cluster area [m × m]	50 × 50	50 × 50	50 × 50	50 × 50
DCP bottom-left [kWh]	2.033	2.029	2.022	2.029
DCP top-left [kWh]	2.035	2.033	2.029	2.029
DCP bottom-right [kWh]	2.031	2.025	2.022	2.034
DCP top-right [kWh]	1.927	2.030	2.024	2.036
–	–	–	–	2.033
–	–	–	–	2.030
–	–	–	–	2.034
–	–	–	–	2.034
–	–	–	–	2.034
–	–	–	–	2.036
–	–	–	–	2.034
–	–	–	–	2.031
–	–	–	–	2.036
–	–	–	–	2.036
–	–	–	–	2.035
–	–	–	–	2.033
DCP center [kWh]	2.030	2.028	1.9966	2.035
Mean $\mu$ [kWh]	2.0112	2.0290	2.02425	2.0335
Standard deviation $\sigma$ [kWh]	0.0471	0.0029	0.0619	0.0023
Margin of error (95% CI) [kWh]	±0.0584	±0.0036	±0.0769	±0.0012



## D.6. Contour plot analysis using different weather data resolution

To evaluate the shading effect of a wind turbine on the surrounding area, a MATLAB script was developed to compute the solar elevation angle and the corresponding shadow length for the full day in Delft, on 1 June 2005. This analysis follows the methodology presented in [58] for finding the solar position and [59, 60] for the shadow length computation.

### D.6.1. Solar elevation - Shadow length

The solar azimuth and elevation angles were computed at 1-minute resolution using the geographic coordinates of Delft (latitude  $52.0116^\circ$  N, longitude  $4.3571^\circ$  E, and sea level altitude), as described by Koblick [58]. To simulate this, a time vector was initially constructed in Coordinated Universal Time (UTC), covering the full day of 1 June 2005 (00:00–24:00 UTC). However, since Delft observes Central European Summer Time (CEST) during June (UTC+2), this range corresponds to 02:00 on 1 June to 02:00 on 2 June in local time. To properly represent 00:00–24:00 local time, the input UTC vector was shifted to start at 22:00 on 31 May 2005 UTC. This adjustment ensures that the plotted solar elevation aligns correctly with the local diurnal cycle.


After computing the solar elevation angle, it was used to calculate the shadow length cast by a vertical wind turbine tower using the following trigonometric relationship:

$$L = \frac{H}{\tan(\theta_{el})} \quad (D.1)$$

where:

- $L$  is the shadow length in m,
- $H$  is the height of the turbine tower (assumed to be 165 m),
- $\theta_{el}$  is the solar elevation angle in  $^\circ$ .

Over the course of the simulated day, the solar elevation angle increases from just above zero shortly after sunrise, peaks around midday, and then decreases again toward sunset.



**Figure D.6:** Solar position and resulting shadow length throughout 1<sup>st</sup> June 2005 in Delft. As the elevation angle increases, the turbine shadow shortens, reaching a minimum at solar noon.

In Figure D.6, the left subplot presents the solar elevation angle as a function of time throughout 1<sup>st</sup> June 2005 in Delft. The curve is plotted in red, with solid lines indicating periods when the elevation

angle exceeds  $0^\circ$ , and dashed lines representing early morning and late evening periods when the Sun remains below this threshold. The right subplot displays the corresponding shadow length cast by the 165 m wind turbine tower. Here, the shadow length is plotted in blue, following the same convention of solid lines for valid solar elevation angles above  $0^\circ$ , and dashed lines for low-angle periods near sunrise and sunset. As expected, the turbine shadow decreases in length as the elevation angle increases, reaching a minimum around solar noon.

The Sun first reaches a meaningful elevation angle above  $0^\circ$  at minute 337 (05:37 local time), with an elevation of  $0.20^\circ$  and an associated shadow length of approximately  $1.13 \times 10^5$  m. As shown in Figure D.6, this moment marks the transition from the dashed to the solid line in both subplots. The elevation angle continues to rise throughout the morning, reaching its maximum of  $60.09^\circ$  at minute 821, i.e., 13:41 local time. Delft lies at approximately  $4.36^\circ$  E, while Central European Summer Time (CEST) is based on the  $15^\circ$  E meridian. This east–west difference causes solar noon in Delft to occur about 42–45 minutes later than 12:00 GMT. As a result, solar noon falls at approximately 12:45 GMT, which corresponds to 13:45 local time (CEST, GMT+2). At this point, the shadow cast by the 165 m wind turbine's tower is shortest, measuring approximately 94.90 m. As the Sun descends in the late afternoon and evening, the elevation angle steadily decreases. At minute 1307 (21:47 local time), the angle drops to  $0.13^\circ$ , resulting in a longer shadow of about  $6.90 \times 10^4$  m, just before sunset when the Sun dips below the horizon and the shadow becomes undefined.

**Table D.6:** Summary statistics of solar elevation angle and turbine shadow length on 1<sup>st</sup> June 2005 in Delft (threshold =  $0^\circ$ ).

	Min	Max	Mean	Median
Solar elevation angle [ $^\circ$ ]	0.02	60.09	33.79	35.19
Shadow length [m]	94.90	$6.15 \times 10^5$	1668.83	233.97

### D.6.2. Velocity of the Turbine Tower Shadow Tip

To estimate how fast the wind turbine tower shadow moves throughout the day, the shadow tip speed is approximated using the forward finite difference method [61]. Based on the calculus definition of the derivative, the instantaneous rate of change in shadow length  $L(t)$  can be estimated as:

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} \Rightarrow \text{Shadow tip speed: } v(t_i) \approx \frac{L(t_{i+1}) - L(t_i)}{\Delta t}, \quad \Delta t = 60 \text{ s} \quad (\text{D.2})$$

Here,  $f(x)$  corresponds to the shadow length  $L(t)$ , and  $h$  is the 1-minute sampling interval. This provides a numerical approximation of the shadow tip velocity as it traverses the ground. The shadow-tip speed is negative in the morning as the shadow shortens and positive in the afternoon as it lengthens. This change reflects the sun's path across the sky. The velocity of the shadow tip is highest during the early morning and late evening hours. This behavior is physically intuitive: near sunrise and sunset, the solar elevation angle is very low, and the Sun appears to move nearly parallel to the horizon. As a result, even a small increase in elevation angle causes a significant reduction in shadow length, leading to a rapid displacement of the shadow tip across the ground.

Over the course of the day, the shadow-tip speed ranged from a minimum of  $0.000 \text{ m s}^{-1}$  to a peak of  $0.839 \text{ m s}^{-1}$ , with an average value of  $0.067 \text{ m s}^{-1}$ . To characterize the typical motion of the shadow tip throughout the day, the *median* shadow-tip speed is used. Unlike the mean, the median is less sensitive to extreme outliers, such as the artificially high speeds observed near sunrise and sunset. The calculated median shadow-tip speed during the day, considering only times when the solar elevation angle exceeds  $5.0^\circ$ , is  $0.019 \text{ m s}^{-1}$ . This corresponds to a shadow displacement of:

- $0.019 \text{ m s}^{-1} \times 3600 \text{ s} = 68.4 \text{ m per hour}$
- $0.019 \text{ m s}^{-1} \times 600 \text{ s} = 11.4 \text{ m per 10 minutes}$

These values provide a practical sense of how gradually the wind turbine's tower shadow moves across the ground under normal daylight conditions.

An estimate of how long a single PV module remains fully shaded by the turbine tower shadow can be derived from a simple kinematic relationship. This follows from the definition of velocity:

$$v = \frac{\Delta x}{\Delta t} \Rightarrow \Delta t \approx \frac{w}{v} \quad (\text{D.3})$$

where:

- $\Delta x = w = 0.992 \text{ m}$  is the width of the PV module in the direction of shadow movement,
- $v$  is the shadow-tip speed in  $\text{m s}^{-1}$ .

Using the median shadow-tip speed of  $v_{\text{med}} = 0.019 \text{ m s}^{-1}$ , the estimated full-shading duration is:

$$\Delta t_{\text{med}} \approx \frac{0.992 \text{ m}}{0.019 \text{ m s}^{-1}} \approx 52 \text{ s} (\approx 0.9 \text{ min}). \quad (\text{D.4})$$

Shadow movement becomes significantly slower near sunrise and sunset. Taking a representative lower value of  $v_{\text{sh, min}} = 0.005 \text{ m s}^{-1}$ , the maximum full-shading duration becomes:

$$\Delta t_{\text{max}} \approx \frac{0.992 \text{ m}}{0.005 \text{ m s}^{-1}} \approx 200 \text{ s} (\approx 3.3 \text{ min}). \quad (\text{D.5})$$

Thus, even under the slowest conditions, complete shading of a module typically lasts only a few minutes, and for most of the day it remains well under one minute. Given the  $12.5 \text{ m}$  spacing between adjacent module centers in the simulated layout, the shadow front travels from one module to the next in approximately:

$$\Delta t_{\text{spacing}} \approx \frac{12.5 \text{ m}}{0.019 \text{ m s}^{-1}} \approx 658 \text{ s} (\approx 11 \text{ min}). \quad (\text{D.6})$$

As a result, irradiance data sampled at coarse temporal resolutions (e.g., hourly) may fail to capture these brief but potentially impactful shading events, leading to an underestimation of the associated energy yield losses.