COOL BY DESIGN

SOLFD: Extending SOLWEIG for Urban Design Decision Making on Outdoor Thermal Comfort

MSc Thesis in Geomatics

Jessica Monahan | 2025

TU Delft

Cool by Design

SOLFD: Extending SOLWEIG for Urban Design Decision-Making on Outdoor Thermal Comfort

by

Jessica Monahan

to obtain the degree of Master of Science in Geomatics
at the Delft University of Technology,
to be defended publicly on Thursday June 19, 2025 at 15:00 PM.

TU Delft

The work in this thesis was carried out in the:

3D geoinformation group
Delft University of Technology

Supervisors: Dr. Clara García-Sánchez
Dr. Hugo Ledoux
Co-reader: Dr. Daniela Maiullari

Abstract

As urbanization and climate change intensify, managing the urban microclimate becomes increasingly challenging, affecting outdoor thermal comfort. The practical integration of urban microclimate research into urban design remains limited, in part due to the complexity and inaccessibility of existing tools. To support early-stage, climate-sensitive urban design at the neighbourhood scale, I present SOLWEIG For Design (SOLFD): a computation framework that builds on the existing SOLWEIG tool.

SOLFD enables urban designers to visualize current microclimatic conditions and assess the impact of design interventions on outdoor thermal comfort. In particular, it focuses specifically on (in)direct solar radiation and its effect as quantified by the mean radiant temperature. Key contributions include: (1) extending SOLWEIG's 2.5D model to a layered 3D representation for improved accuracy in complex urban geometries; (2) automating the data pipeline using open Dutch geospatial datasets; (3) enabling the modification of the existing urban scene; (4) enhancing output usability through temporally grouped mean radiant temperature maps, derived physiological equivalent temperature maps, and comparison statistics; and (5) significantly reducing simulation time with GPU acceleration.

The accuracy of SOLFD was validated using sensor data, achieving an RMSE of 5.39°C. Underneath structures, the RMSE increases to 5.83 °C. The potential of SOLFD is further demonstrated with a case study across various Dutch urban typologies.

By laying the foundation for an accessible decision-support tool for outdoor thermal comfort, SOLFD takes a step toward integrating climate-responsive strategies into the urban design process.

Keywords: UMC, OTC, TMRT, SOLWEIG, UMC for Urban Design

Acknowledgments

I would like to express my heartfelt gratitude for the support and guidance I have received throughout the course of this research.

First of all, I would like to thank my supervisors, Clara García-Sánchez and Hugo Ledoux. Thank you for your confidence in me and for allowing me the freedom to shape this research around a topic that truly interests me. Your encouragement during our weekly meetings was invaluable, I really enjoyed having you as supervisors. Many thanks also to Daniela Maiullari, whose insightful comments helped me further sharpen this thesis.

Furthermore, I would like to thank my family for their love and support, during this thesis and outside of it. To my mom, thank you for helping me through the most difficult moments of my studies. To my sisters, thank you for your sisterly and doctorly advices, and to the plus ones as well. Lucas, I'm looking forward to weekends without studying. And to my dad: making me try out Scratch years ago may have led me to become a programmer after all. I wish you were here to see it.

Finally, to my friends, old and new, from elementary school to my Master's: thank you for making (thesis) life more fun.

Contents

Abstract ii

Acknowledgments iii

Nomenclature vi

1 Introduction 1
- 1.1 Background and Motivation . . . 1
- 1.2 Research Objectives and Scope . . . 2
  - 1.2.1 Objectives . . . 2
  - 1.2.2 Scope . . . 2
- 1.3 Results Overview . . . 3
- 1.4 Thesis Outline . . . 3

2 Related Work 4
- 2.1 Urban Design Process . . . 4
- 2.2 Urban Climate Tools . . . 5
- 2.3 Urban Climate Tools for Designing . . . 5
- 2.4 Urban Climate Tools used in the Design Process . . . 6

3 Theoretical Background 8
- 3.1 Mean Radiant Temperature . . . 8
- 3.2 Background on UMEP Tool and SOLWEIG . . . 9

4 Methodology 12
- 4.1 Development of SOLFD . . . 13
  - 4.1.1 Extending SOLWEIG for 3D Simulation . . . 14
  - 4.1.2 Automatic Data Collection for SOLWEIG Input Pipeline . . . 17
  - 4.1.3 Integrating User Designs and Enabling Modifications . . . 23
  - 4.1.4 Tailoring SOLWEIG Outputs for Design and Analysis . . . 30
  - 4.1.5 Computation Time Optimization . . . 32
- 4.2 Accuracy Assessment of SOLWEIG and SOLFD . . . 34
  - 4.2.1 Literature Review on Model Accuracy . . . 34
  - 4.2.2 Accuracy Assessment Using Observations . . . 34
- 4.3 Urban Test Case with SOLFD . . . 39

5 Results 41
- 5.1 SOLFD Development . . . 41
  - 5.1.1 Extending SOLWEIG for 3D Simulation . . . 41
  - 5.1.2 Automatic Data Collection for SOLWEIG Input Pipeline . . . 46
  - 5.1.3 Integrating User Designs and Enabling Modifications . . . 54
  - 5.1.4 Tailoring SOLWEIG Outputs for Design and Analysis . . . 57
  - 5.1.5 Computation Time Optimization . . . 60
- 5.2 Accuracy Assessment of SOLWEIG and SOLFD . . . 67
  - 5.2.1 Accuracy Assessment Using Observations . . . 70
- 5.3 Urban Test Case with the Developed Tool . . . 79

6 Conclusions 82
- 6.1 Conclusion . . . 82
- 6.2 Discussion . . . 85
- 6.3 Future work . . . 88

References 89

A Appendix: Supplementary Tables and Figures **93**
A.1 SOLWEIG and preprocessors interfaces . . . 93
A.2 Urban Case Study Locations . . . 98
A.3 Classification of BGT and TOP10NL types to SOLWEIG landcover types . . . 99
A.3.1 BGT . . . 99
A.3.2 TOP10NL . . . 99
A.4 Evaluated Tree Databases for Dutch Species Parameters . . . 101
A.5 Meteorology Tables and Figures . . . 102
A.6 Typologies comparison plots . . . 107

B Datasets and Reproducibility-assessment **111**
B.1 Datasets . . . 111
B.2 Reproducibility assessment . . . 111

Nomenclature

AAMS Amsterdam Atmospheric Monitoring Supersite 34
AHN Actueel Hoogtebestand Nederland 18
BGT Basisregistratie Grootschalige Topografie 22
CAD Computer-Aided Design 3
CFD Computational Fluid Dynamics 5
CHM Canopy Height Model 3
DSM Digital Surface Model vii
DTM Digital Terrain Model 3
GUI graphical user interface 93
GVF ground view factor 17
LBT Ladybug Tools 5
NDVI Normalized Difference Vegetation Index 19
OTC outdoor thermal comfort 1
PET Physiological Equivalent Temperature x
RMSE root mean square Error 3
SOLFD SOLWEIG For Design vii
SOLWEIG SOlar and LongWave Environmental Irradiance Geometry model 1
T_s surface temperature 22
SVF skyview factor 11
T_a air temperature 3
RH relative humidity 3
T_{mrt} mean radiant temperature viii
UMC urban microclimate 1
UMEP Urban Multi-scale Environmental Predictor 1
v_a air velocity 3
WCS Web Coverage Service 18

List of Figures

2.1 Simulation interaction between the design and simulation team. 7

3.1 Explanation of T_{mrt} 8
3.2 Overview of SOLWEIG 9
3.3 Overview of the required inputs for SOLWEIG. 10

4.1 Overview of the research and the research methodology. 12
4.2 Overview of SOLWEIG For Design (SOLFD) components and development. 13
4.3 Vertical sections of 2.5D and 3D situations 14
4.4 SOLWEIG shadow casting methods of DSM and CHM. 15
4.5 Illustration of the layered array concept for 3D shadow casting. 16
4.6 3D shadow casting. 16
4.7 The two test cases for optimal number of void layers testing 17
4.8 The weight for the Laplace interpolant for one neighbor 18
4.9 Conversion of raster cells to x,y coordinates. 18
4.10 Extracting the building heights of a DSM with 3DBAG 19
4.11 Points in an example scene classified as vegetation when NDVI = 0.1 20
4.12 Under-detection and misclassification of tree points 20
4.13 Goodwin et al. (2009) method to calculate wall aspect. 21
4.14 The four different kind of tree shapes included in the model; Cone, gaussian, parabolic and hemisphere, with their formulas. Resulting output tree with random noise added. . 25
4.15 Methodology for Integrating 3D into Urban Design 26
4.16 Sorted raw height dataset and 3D context 28
4.17 Void-to-Digital Surface Model (DSM) transformation 29
4.18 SOLWEIG output: (a) Tmrt in °C, (b) shadow maps indicating shaded areas, (c) upward and downward longwave and shortwave radiation maps in Wm^{-2}. 30
4.19 Computation time profile for SVF and SOLWEIG calculations across different array sizes, with and without CHM data 32
4.20 Visualization of the three different measurement cases 34
4.21 The six selected locations with overpasses in the Climate Bike dataset 35
4.22 Creation of the location datasets 36
4.23 Climate Bike orientation before and after interpolation 37
4.24 Manual correction of drifted GPS positions 38
4.25 The different types of Dutch Neighborhood typologies 39
4.26 *Wijktypologie PC6 Gebieden 2023.* Source: (Kleerekoper et al., 2023) 40

5.1 Building geometry (left) and the resulting shadow map produced by the shading algorithm (right). 41
5.2 Layered DSM for the example 42
5.3 Layered DSM input for Case 1 with 1 and 4 voids 42
5.4 Layered DSM input for Case 2 with 1 and 2 voids 43
5.5 SVF and t_{mrt} outputs for the one-void and four-void layer configurations in Case 1, along with the differences between them 44
5.6 SVF and t_{mrt} outputs for the one-void and two-void layer configurations in Case 2, along with the differences between them. 45
5.7 Input AHN DTM and DSM with missing data, filled outputs, and corresponding difference maps. 46
5.8 CHM result, validation dataset, and their differences 47
5.9 The output of both original and new building height method. 48

5.10 Wall aspect outputs using (a) the original 9x9 linear filter and (b) the new 3x3 Sobel filter. 48
5.11 Comparison of mean radiant temperature (T_{mrt}) results at 13:00 and 17:00 using the liniar filter and Sobel filter wall aspect methods. Differences are shown in the third row. . . . 49
5.12 Close-up examples showing differences caused by the new Sobel-based wall aspect method. 50
5.13 Comparison of the land cover classification output (left) with satellite imagery (right). . 51
5.14 Close-up land cover and examples of classification errors. 51
5.15 Hourly climate variable summaries for three day types: Warm (20–25°C), Hot (25–30°C), and Tropical (30°C+). 53
5.16 Example of removing existing elements from the scene based on user-specified IDs . . . 54
5.17 Example of inserting a user-defined tree into the scene. 54
5.18 Insertion of 3D urban design in SOLFD 55
5.19 Land cover update from 3D urban design in SOLFD 56
5.20 User-modified land cover classification in SOLFD 56
5.21 Overview of post-processed SOLWEIG outputs in SOLFD, including (a) output statistics, (b) classified PET based on physiological thermal stress, (c) hourly T_{mrt}, and (d) aggregated PET maps. 57
5.22 Comparison of PET values computed using the original SOLWEIG method and the newly implemented lookup table method 58
5.23 Resulting classification of the two different methods and their difference. 59
5.24 Computation time comparison between the old SOLWEIG (CPU-based) and new SOLFD (GPU-accelerated) code versions for SVF and SOLWEIG on the computer 60
5.25 Computation time comparison between the old SOLWEIG (CPU-based) and new SOLFD (GPU-accelerated) code versions for SVF and SOLWEIG on the laptop 61
5.26 Computation time comparison between SOLFD 2.5D and SOLFD 3D for SVF and SOLWEIG on the computer 63
5.27 Computation time comparison between SOLFD 2.5D and SOLFD 3D for SVF and SOLWEIG on the laptop 63
5.28 DSMs and CHMs at 0.5 m, 1.0 m, and 2.0 m resolutions 64
5.29 T_{mrt} distributions at 13:00 and 18:00 for the 0.5 m, 1.0 m, and 2.0 m resolution cases. . . . 65
5.30 Difference in T_{mrt} at 13:00 and 18:00 for the 1.0m and 2.0m resolution cases, compared to the 0.5m baseline. 65
5.31 Plot and maps of modeled and observed T_{mrt} at Location 1. 70
5.32 Plot and maps of modeled and observed T_{mrt} at Location 2. 71
5.33 Plot and maps of modeled and observed T_{mrt} at Location 3. 72
5.34 Plot and maps of modeled and observed T_{mrt} at Location 4. 73
5.35 Plot and maps of modeled and observed T_{mrt} at Location 5. 74
5.36 Plot and maps of modeled and observed T_{mrt} at Location 6. 75
5.37 Overall performance of the three evaluated models. 76
5.38 Performance of the three evaluated models underneath structures 76
5.39 Performance from observations +- 5 degrees aligned with cardinal directions 77
5.40 Incoming global radiation for 12 September 78
5.41 T_mrt based on six-directional monitoring of short- and longwave radiation for human as box and as cylinder. 78
5.42 T_{mrt} Area Percentages for the six urban typologies during different times of the day and overall, for both urban cases with trees and without trees. 81

A.1 The SOLWEIG interface from the UMEP QGIS plug-in. 94
A.2 The 'Prepare Existing Data' interface from the UMEP QGIS plug-in. 95
A.3 The 'Urban Geometry: Sky View Factor' interface from the UMEP QGIS plug-in. 96
A.4 The 'Urban Geometry: Wall Aspect and Height' interface from the UMEP QGIS plug-in. 96
A.5 The 'Spatial Data: Tree Generator' preprocessor interface from the UMEP QGIS plug-in. Used for CHM creation, points are required as input for tree positions. 97
A.6 Variations in median T_{mrt} across the six urban typologies on the 'extreme day', shown for morning, afternoon, evening, and the overall daily average. 107
A.7 Variations in mean T_{mrt} across the six urban typologies on the 'extreme day', shown for morning, afternoon, evening, and the overall daily average. 107

A.8 Variations in minimum T_{mrt} across the six urban typologies on the 'extreme day', shown for morning, afternoon, evening, and the overall daily average. 108
A.9 Variations in maximum T_{mrt} across the six urban typologies on the 'extreme day', shown for morning, afternoon, evening, and the overall daily average. 108
A.10 Variations in median T_{mrt} across the six urban typologies on the 'average day', shown for morning, afternoon, evening, and the overall daily average. 109
A.11 Variations in mean T_{mrt} across the six urban typologies on the 'average day', shown for morning, afternoon, evening, and the overall daily average. 109
A.12 Variations in minimum T_{mrt} across the six urban typologies on the 'average day', shown for morning, afternoon, evening, and the overall daily average. 110
A.13 Variations in maximum T_{mrt} across the six urban typologies on the 'average day', shown for morning, afternoon, evening, and the overall daily average. 110

List of Tables

3.1 Examples of PET values for different climate scenarios. 9

4.1 The five most common tree species in the Netherlands. 24
4.2 Thermal sensitivity and physiological stress levels based on PET thresholds 31
4.3 Included values and parameter ranges for Physiological Equivalent Temperature (PET) lookup table. 31
4.4 Body model parameters for different demographic profiles used in PET calculation. . . 32

5.1 Hardware specifications of the laptop and desktop used. 41
5.2 Computation time (in seconds) for SVF and SOLWEIG calculations by void configuration and case. 45
5.3 T_{mrt} Difference Statistics in °C at 13:00 PM and 17:00 PM 49
5.4 Comparison of original and GPU-accelerated computation times of SOLWEIG and SVF on the computer. 60
5.5 Comparison of original and GPU-accelerated computation times of SOLWEIG and SVF on the laptop. 61
5.6 Comparison of computation times of SOLFD 2.5D and SOLFD 3D of SOLWEIG and SVF on the computer. 62
5.7 Comparison of computation times of SOLFD 2.5D and SOLFD 3D of SOLWEIG and SVF on the laptop. 64
5.8 Overview of SOLWEIG model validation results by original developers. 67
5.9 Overview of SOLWEIG model validation results by independent research. 69
5.10 Average Tmrt difference for an 'average' day (with trees minus without trees) per typology, split by statistic and time of day. 79
5.11 Average Tmrt difference for an 'extreme" day (with trees minus without trees) per typology, split by statistic and time of day. 79

A.1 Case study locations grouped by urban typology, showing bounding boxes in EPSG:28992 coordinates and array sizes (in meters). 98
A.2 Mapping of terrain, road, and water types from the BGT dataset to SOLWEIG land cover classes. 100
A.3 Mapping of TOP10NL terrain, road, and water types to SOLWEIG land cover classes. . 100
A.4 Overview of tree databases evaluated for extracting species-specific parameters 101
A.5 Meteorological data for an average day with air temperatures in the range of 20-25 . . . 103
A.6 Meteorological data for an average day with air temperatures in the in the range of 25-30. 104
A.7 Meteorological data for an average day with air temperatures in the 30+ range. 105
A.8 Meteorological data for the Extreme Day in the 30+ temperature range. 106

B.1 Overview of datasets used, their sources, descriptions, and applications within the SOLFD pipeline. 112

List of Algorithms

4.1 Converting 3D model to 2D layered arrays 27

1 Introduction

1.1. Background and Motivation

As cities continue to grow rapidly due to urbanization, climate change has become an increasingly critical challenge. Their combined effects manifest in the urban microclimate (UMC); which refers to the distinct local climate within urban areas, shaped by the built environment and differing from surrounding rural conditions in temperature, wind, humidity, and solar radiation. Urbanization significantly alters these parameters, often increasing local heat stress and impacting outdoor thermal comfort (OTC) (Bherwani et al., 2020). These conditions present significant challenges for urban designers, which can be addressed by developing climate-driven urban designs aimed at reducing environmental impacts. The demand for sustainable and resilient urban design is increasing as cities must adapt to a changing climate while mitigating its effects (Cheshmehzangi et al., 2022).

Effective climate-driven urban design requires the exchange of UMC knowledge from researchers to urban designers and planners. However, as noted by Bherwani et al. (2020) and Aleksandrowicz (2022), the practical application of UMC research is limited. This is partly due to the complexity of existing tools, which are often unsuitable for practitioners. There is a need for user-friendly tools integrating UMC into city planning.

As Besserud and Hussey (2011) explain, simulation tools for urban design must provide fast results to allow for comparison between many design options, feature intuitive user interfaces accessible to non-experts, and support inputs and outputs appropriate for early-stage design.

One existing UMC tool is Urban Multi-scale Environmental Predictor (UMEP), an open-source QGIS plug-in for urban climatology and planning, designed for both researchers and service providers like architects and urban planners (Lindberg et al., 2018). However, it is not quite yet a user-friendly tool for urban design. It requires QGIS proficiency, data preparation, and does not readily accommodate the insertion of own designs or comparison between design options.

UMEP includes multiple climate models. For urban design, the effect of UMC on human OTC is particularly interesting, as it directly affects the quality of urban open spaces and plays a role in sustainable city design (Abdollahzadeh and Biloria, 2021). The major meteorological parameters affecting OTC are wind speed, air temperature, relative humidity and solar radiation (Stathopoulos et al., 2004). Of these, wind speed and solar radiation are particularly responsive to urban design interventions, as they can be significantly modified by the arrangement of building clusters and vegetation (Xie et al., 2018). UMEP includes two relevant models for urban design: *SOlar and LongWave Environmental Irradiance Geometry model (SOLWEIG)*, which focuses on solar and longwave irradiation (Lindberg et al., 2008), and *URock*, which models wind speed (Bernard et al., 2023).

Between the two models, I consider SOLWEIG as having greater potential for urban design applications for this thesis. Recent studies highlight the increasing frequency, duration and intensity of heatwaves, which worsen urban heat stress through high solar radiation and reduced wind speeds (Kong et al., 2021). To enhance urban resilience, addressing solar radiation through effective design interventions is

needed, and SOLWEIG is suited for this. In addition, SOLWEIG is a more mature and extensively evaluated model, offering a more stable basis for assessing OTC in urban environments.

This research aims to develop the framework for a fast and user-friendly design decision-support tool that enables urban designers to evaluate and incorporate UMC impacts during early design phases. Building upon the capabilities of SOLWEIG, the tool, hereafter referred to as SOLFD, will focus on assessing OTC in response to design interventions. This is important because existing UMC tools are often too complex or time-consuming for urban design practice. By translating UMC research into an accessible and practical format, SOLFD will help bridge the gap between science and design, supporting more resilient and climate-sensitive urban environments.

1.2. Research Objectives and Scope

1.2.1. Objectives

The main objective of this thesis is to develop a framework for a tool, named SOLFD, based on SOLWEIG, that enables designers to visualize current urban climate conditions at the neighborhood scale in the Netherlands and to evaluate, visualize and compare the potential impacts of proposed designs on these conditions for the human comfort. By *framework*, this thesis refers to a structured and documented set of computational functions, methods, and workflows that collectively form the foundation upon which a full tool can be built. Although the development of a complete tool interface is beyond the scope of this work, the framework provides the essential backbone for future interactive applications.

To achieve the objective, this research involves three steps. First, the devolpment of SOLFD, with a pipeline that automatically generates the necessary input for SOLWEIG and allows urban designers to modify the current situation. Second, the validation of current SOLWEIG algorithms and the modified version (SOLFD) that I will develop for this thesis. Third, the application of the designed tool on real-world use case from Dutch cities.

For the first step, the main research question will be:

> *How can a tool framework be developed for visualizing current urban climate conditions and their effect on outdoor thermal comfort at the neighborhood scale in the Netherlands using SOLWEIG, and how can design alternatives be integrated and compared within this tool?*

This step includes the following subquestions:

- To what extent can the 2.5D algorithm of SOLWEIG be transformed into a real volumetric algorithm?
- How can the input data pipeline be automated to extract and create the Dutch location data required for SOLWEIG?
- How can urban design proposals be integrated as input for SOLWEIG?
- What additional outputs or post-processing steps are required after SOLWEIG to compare and rank different urban design options based on their microclimatic impact and resulting outdoor thermal comfort?
- To what extent can SOLWEIG be optimized for computation time?

For the second step, the main research question will be:

> *What does existing research report about the accuracy of SOLWEIG, and how well do the newly implemented 3D modifications perform when validated against sensor data?*

For the third step, the main research question will be:

> *How can the tool be used to analyze the effects of urban design proposals on microclimatic conditions and outdoor thermal comfort in various Dutch urban typologies?*

1.2.2. Scope

The scope of this research is limited to developing a functional framework for a design-support tool (SOLFD) based on SOLWEIG, focusing specifically on evaluating and visualizing microclimatic conditions related to OTC. It does not include any modifications to the underlying climate physics or formulas

of SOLWEIG itself. Furthermore, the research does not include the collection of new sensor data or the development of new surface material parameterization. SOLFD will focus exclusively on solar radiation and its impact on human OTC, represented through T_{mrt}. Broader microclimatic factors affecting OTC, such as wind speed or air quality, are beyond the scope of this work, although these could be integrated into future versions to enhance decision-making. Lastly, SOLFD will not be implemented as a fully operational platform; instead, the emphasis is on developing the code foundation, recognizing that the creation of an intuitive, user-friendly interface would require additional research on interface design and usability.

1.3. Results Overview

This thesis has resulted in the development of SOLFD: a framework based on SOLWEIG, which enables integration and comparison of design alternatives for the neighborhood-scale visualization in the Netherlands. Key advancements include:

- Extension of SOLWEIG from a 2.5D to a layered 3D representation, allowing modeling of complex urban features such as building overhangs, balconies and underpasses.
- Automation of the full data pipeline by integrating open Dutch government datasets to generate all required spatial inputs, Digital Terrain Model (DTM), DSM, Canopy Height Model (CHM), and land cover, using only location bounding boxes.
- Incorporation of user input through conversion of 3D Computer-Aided Design (CAD) designs (`.OBJ`) into rasterized layers, with tools to insert or modify buildings, trees, and land cover elements.
- Enhancement of output usability by summarizing results via temporal groupings and spatial statistics for effective design comparison, and through simplified PET maps derived from modeled T_{mrt} and static inputs of air temperature (T_a), relative humidity (RH), and air velocity (v_a).
- Significant computational speed improvements through GPU acceleration and code optimization, reducing runtime from tens of minutes to a few minutes for typical neighborhood-scale simulations.

Validation against sensor data shows that SOLWEIG performs within acceptable accuracy ranges reported in the literature (root mean square Error (RMSE) ~3–7°C), with the new 3D modifications particularly improving simulation accuracy beneath overhead structures.

1.4. Thesis Outline

This thesis is composed of six chapters, including this introduction. The content is structured as follows:

Chapter 2 reviews literature and previous work relevant to this study. It explores the urban design process and the role of UMC within it, discusses general UMC tools as well as those developed specifically for urban design, and concludes with a review of studies on how urban designers interact with such tools.

In Chapter 3, the theoretical background is discussed. This includes an explanation of what T_{mrt} is, followed by an overview of how SOLWEIG operates.

Chapter 4 presents the methodology used to address the research questions. It outlines the five-step development process of SOLFD, describes how SOLWEIG is evaluated, and details the setup of the case study involving various urban typologies.

Chapter 5 provides the results from each step of the developed tool framework. It also includes the accuracy assessment of SOLWEIG through literature study, along with a validation of the 3D extension of the model using real-world sensor data. Moreover, the results of the case study are given.

Finally, Chapter 6 offers conclusions and directly answers the research questions. It reflects on the findings, highlights the contributions and limitations of the study, and ends with recommendations for future research.

2

Related Work

In this chapter, relevant literature is reviewed to understand the urban design process and the role of UMC within it. Section 2.1 introduces the urban design process, Section 2.2 provides an overview of existing UMC tools, Section 2.3 focuses on UMC tools specifically developed for design applications, and Section 2.4 discusses the practical use of UMC tools in practice.

2.1. Urban Design Process

This thesis focuses on developing the framework of a tool that integrates climate-driven decision-making into the urban design process by evaluating the impact of design options on UMC. Achieving this requires understanding the urban design process, the role of UMC in this process, and identifying when such a tool would be beneficial.

Urban design is a multifaceted process aimed at creating better places for people, with an emphasis on equity, sustainability, and the importance of place-making. It integrates creativity, exploration, and problem-solving to balance diverse objectives and constraints. This process addresses technical, functional, and aesthetic criteria holistically, ensuring designs achieve structural integrity, utility, and visual appeal while also considering economic and environmental impacts (Carmona et al., 2010).

The urban design process is iterative and cyclical, adapting to new information and changing circumstances, making it exploratory and inherently uncertain. The process typically includes goal setting, analysis, visioning, synthesis, decision-making, and evaluation. These stages are not strictly linear but overlap and repeat (Carmona et al., 2010).

Interviews by Van Esch (2015) with Dutch urban designers in 2010 show they recognize the importance of microclimates but rarely seek information on the subject, indicating a gap. Additionally, they collect data primarily during the orientation and sketch phases of the design process. Consequently, the dissemination of expert knowledge on microclimates should be targeted at these early stages.

During these early design stages, urban designers create and adjust massing models to rapidly explore various urban form options (Chen and Norford, 2017). SOLFD should enable the comparison of these generated massing models by providing performance indicators that assess their impact on UMC.

Can expert knowledge on microclimates be effectively shared with urban designers through the proposed tool? Van Esch (2015) observes that computer tools in urban design, particularly simplified expert programs, often fail to support the creative and iterative nature of the process. These tools are primarily analytical and may lead to errors or misinterpretations when designers lack the necessary theoretical knowledge. Moreover, the vast number of tools required to cover all knowledge fields can overwhelm designers and complicate decision-making. Van Esch suggests that providing basic principles, guidelines, and example projects may better support designers.

Others emphasize the value of simulation tools. For instance, Wong et al. (2011) and Ebrahimabadi (2015) argue that while simplified climate design principles aid early stages, advanced analytical

tools are crucial for later stages to evaluate and refine concepts. Additionally, Besserud and Hussey (2011) highlight how simulation tools can complement design principles, introducing greater rigor and rationality into urban design. By integrating principles with the ability to simulate and compare design options, such tools help designers access complex data intuitively, improve decision-making, and accelerate design iterations.

2.2. Urban Climate Tools

For this thesis, UMEP, described further in Section 3.2, was selected as the urban climate tool due to its open source code. However, it is not the only option available. Vurro and Carlucci (2024), have analyzed 25 different urban climate tools. These tools are predominantly commercial, with only 8 of the studied tools being open source. They can be categorized based on their primary functions as follows:

- **Computational Fluid Dynamics (CFD) Tools**: These tools simulate fluid flow and heat transfer. With 10 out of 25 tools studied, they are the most common typology. CFD tools are further divided into:
  - *General-Purpose CFD Tools* (e.g., OpenFoam, Fluent, Autodesk CFD): These tools provide extensive customization options. However, they require expertise in CFD modeling. They do not include pre-set urban item libraries (e.g., pre-made vegetation or surface models) requiring users to build simulations from scratch.
  - *Specialized Microclimate CFD Tools* (e.g., ENVI-met, Ladybug DragonFly, UrbaWind): Designed specifically for UMC modeling, these tools have limited customization options copmared to General-Purpose tools, and offer pre-set urban item libraries.
- **Surface Energy Balance Models**: Tools such as RayMan, SOLWEIG, and SUEWS focus on modeling the exchange of energy between surfaces and the atmosphere, essential for understanding urban thermal dynamics.
- **Building Energy Models (BEM)**: Tools like TAS and the BEM features of Ladybug Tools (LBT) simulate energy use and thermal behavior within buildings.
- **Air Quality Models**: Tools like the ADMS Temperature and Humidity model are designed to analyze air quality conditions in urban areas.
- **Environmental Design-Oriented Tools**: Ladybug Tools is an open-source suite designed for environmental building design, enabling climate analysis, energy modeling, airflow simulation, and urban-scale studies through integrating simulation engines in CAD environments.

As this thesis aims to develop a framework for a tool that integrates UMC and urban design, LBT is of particular interest due to its explicit focus on combining microclimate and environmental modeling within the design process. This tool will be examined in greater detail in the following section, which addresses UMC tools specifically developed for design workflows.

2.3. Urban Climate Tools for Designing

The previous sections have explored general urban climate tools and their role in the design process, particularly in relation to UMC considerations. Given the focus of this thesis on developing a framework for a tool tailored for designers, it is important to review existing literature on UMC tools specifically developed for design workflows. This includes identifying established guidelines for usability, integration within the design process, and criteria for evaluating tool effectiveness from a designer's perspective.

The interviewing and surveying of a large group of architects in Belgium by Weytjens and Verbeeck (2010) on what makes an energy evaluation simulation tool 'architect-friendly' provides insights that are assumed to be transferable to design tools more broadly, including UMC tools. The study emphasizes that such tools must prioritize simplicity, intuitive interfaces, and minimal data input, allowing designers to quickly test scenarios within time-constrained workflows. A key feature is the ability to provide feedback on the impact of design parameters, with visual and easily interpretable outputs being essential for effective communication, particularly with clients. Additionally, the tool must be flexible, transparent, and adaptable to the early stages of the design process. Notably, ease of use often outweighs analytical precision.

In line with these considerations, Erell et al. (2010) argues that UMC models should support planning decisions by addressing architects' specific concerns, using inputs relevant to their design process and allowing them to explore the impact of different design choices on the microclimate.

Van Esch (2015) highlights that UMC is influenced by urban morphology, materialization, and landscaping, through principles such as reflection, absorption, and evapotranspiration. Understanding these effects empowers designers to make informed decisions to improve the UMC. Van Esch (2015) proposes a web-based design-decision support tool for urban designers, offering categorized information on climate elements, design principles, guidelines, and example projects. The tool should focus on the early stages of the design process, providing customized information that adapts to various design approaches. Visual representations should be emphasized, while minimizing numerical data. Example projects can demonstrate how principles are applied in practice, facilitating communication among team members. The tool should be easy to navigate, allowing designers to access tailored information suited to their needs, project phase, and design style.

Taken together, these studies underscore the importance of developing designer UMC tools that are user-friendly, visually oriented, flexible, and well-integrated within early design workflows. Such tools should balance usability with sufficient analytical depth to support informed decision-making by designers.

A relevant example of an existing tool is LBT, as mentioned above in Section 2.2. LBT is a plugin for Rhino's graphical algorithm editor Grasshopper and supports climate-based analysis in early design stages. It visualizes environmental factors such as solar radiation, thermal comfort, and daylight performance using standard weather data. Integrated within parametric modeling environments, LBT provides interactive visual feedback, enabling users to test design variations efficiently. Its integration within CAD environments ensures a unified workflow, where designers can input geometry and review results (Roudsari et al., 2013). However, effective use of LBT requires an initial investment in learning both Grasshopper and LBT itself. The tool currently lacks detailed context and terrain data, and has limitations regarding scale and level of detail. Finally, although the code is openly available, Rhino itself is proprietary paid software.

2.4. Urban Climate Tools used in the Design Process

Finally, to understand the current usage of tools by practitioners, attention is given to studies that describe how these tools are applied in practice. Most studies found using UMC tools for urban design focus on comparing design options for specific cases to find the best option in regards to outdoor comfort, and to generate general guidelines for outdoor comfort. However, there is a lack of literature on the direct application of UMC tools by urban designers themselves. Some examples of the comparison of design options studies found include:

Eldarwish et al. (2020) used *ENVI-met* to compare two urban design proposals and the current situation, selecting the best option based on outdoor comfort metrics like T_{mrt} and surface temperatures. Ambrosini et al. (2014) assessed the impact of converting traditional roofs to green roofs in a historic city center using *ENVI-met*. Hsieh and Kangli (2024) applied CFD simulations to evaluate the effects of various design proposals on pedestrian wind comfort, highlighting the influence of factors such as building mass, layout, height, and open spaces on the wind environment.

Two studies focus more on the direct application of UMC tools at the beginning of the design phase, rather than using them solely for evaluation at the end.

Klemm et al. (2014) describe how climate-responsive design was integrated into a re-design proposal, incorporating UMC considerations at every stage. Initial microclimatic analysis guided the application of heat mitigation strategies like cool materials and greenery, with *RayMan* software used to simulate radiation fluxes and thermal comfort. *ENVI-met* simulations then evaluated the thermal comfort, leading to the addition of some vegetation. High-resolution *ENVI-met* models confirmed the heat reduction goals were met.

Huang et al. (2019) focus on integrating simulation tools into the design process to develop informed, evidence-based strategies for mitigating heat stress in a urban renewal project. Tools such as *+CityComfort* and *HTB2-Virvil* were employed to assess UMC conditions and pedestrian thermal stress. Simulation results were iteratively fed back to designers at early stages of the project, as seen in Figure 2.1, enabling the development of adaptive design strategies.

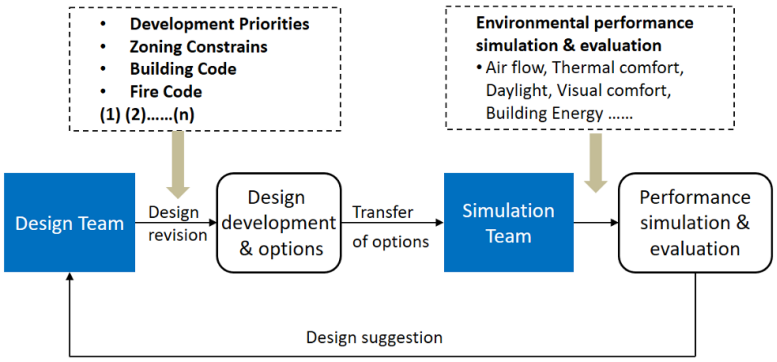

Figure 2.1: Simulation interaction between the design and simulation team. Source: (Huang et al., 2019, p. 3)

3

Theoretical Background

This chapter covers two main components of the thesis. Section 3.1 explains the concept of T_{mrt}, which is the primary variable that SOLFD aims to model. Section 3.2 introduces the tool and model on which SOLFD is built upon.

3.1. Mean Radiant Temperature

The main part of this thesis is the adaptation of SOLWEIG, which simulates the spatial variation of 3D-radiation fluxes and the T_{mrt} in urban settings (Lindberg et al., 2008). The T_{mrt} is an important parameter for evaluating OTC.

According to investigations on OTC conducted in inland cities during clear-sky weather, radiant exchange is the dominant meteorological factor influencing thermal comfort during the daytime (Lee et al., 2013). The T_{mrt} quantifies this radiant exchange and is directly influenced by built geometry, surface materials, and vegetation. This makes T_{mrt} an effective measure for evaluating the efficiency of different strategies aimed at reducing radiant heat (Thorsson et al., 2014). Consequently, it is a valuable metric for comparing various design options in urban planning.

The T_{mrt} is the total of all shortwave and longwave radiation fluxes, both direct and reflected, that the human body is exposed to. Lindberg et al. (2008) mention the definition of T_{mrt} given by the ASHRAE, which is "the uniform temperature of an imaginary enclosure in which radiant heat transfer from the human body equals the radiant heat transfer in the actual nonuniform enclosure". This principle applied to the urban environment is illustrated in Figure 3.1

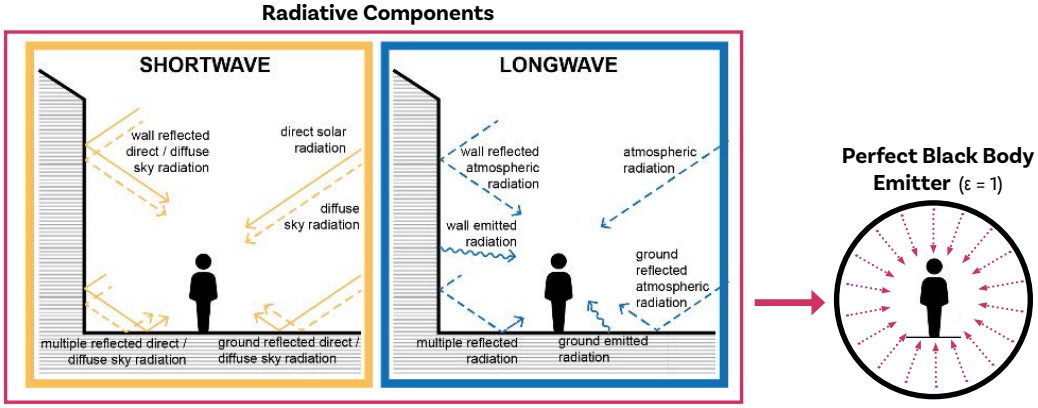

Figure 3.1: The T_{mrt} is the uniform temperature of an imaginary enclosure (perfect black body emmiter, on the right), in which the radiant heat transfer is equal to that from the nonuniform enclosure (the radiative components of the urban environment). Figure modified from (Badino et al., 2021, p. 2)

In SOLWEIG, T_{mrt} is computed by taking six radiation flux components (longwave and shortwave radiation from upward, downward, and the four cardinal directions), together with the angular and absorption factors of a person (Lindberg et al., 2008).

To make evaluation of design options possible, it helps to understand what levels of T_{mrt} are acceptable. OTC is commonly assessed using indices such as the PET, which represents how the thermal environment feels as if the person were indoors, without wind or direct sunlight. A key factor influencing PET is T_{mrt}, particularly on windless days (Höppe, 1999). However, it is not possible to define a specific minimum or maximum as ideal or problematic. This is because its effect is context-dependent, varying with other environmental factors such as air temperature, wind speed, and humidity. For instance, high values are desirable during cold weather conditions to enhance thermal comfort. This relationship is illustrated in Table 3.1.

Table 3.1: Examples of PET values for different climate scenarios.
T_a air temperature, v air velocity, VP water vapor pressure. Source: (Höppe, 1999, p. 73)

Scenario	T_a (°C)	T_{mrt} (°C)	v (m/s)	VP (hPa)	**PET** (°C)
Typical room	21	21	0.1	12	**21**
Winter, sunny	-5	40	0.5	2	**10**
Winter, shade	-5	-5	5.0	2	**-13**
Summer, sunny	30	60	1.0	21	**43**
Summer, shade	30	30	1.0	21	**29**

3.2. Background on UMEP Tool and SOLWEIG

UMEP is an open-source tool designed for researchers and service providers, for analyzing urban OTC, energy use, and climate change mitigation. It integrates 1D and 2D models, allowing data input from various sources and scales to generate outputs as maps, graphs, and datasets. Built as a QGIS plug-in, it contains three main elements: a pre-processor for data preparation, a processor for running models, and a post-processor for analyzing results (Lindberg et al., 2018). This thesis shall adapt SOLWEIG integrated in the UMEP tool. For a better understanding of SOLWEIG in its current form, figures of the UMEP interface for SOLWEIG and the associated pre-processing tools are provided and explained in Appendix A.1.

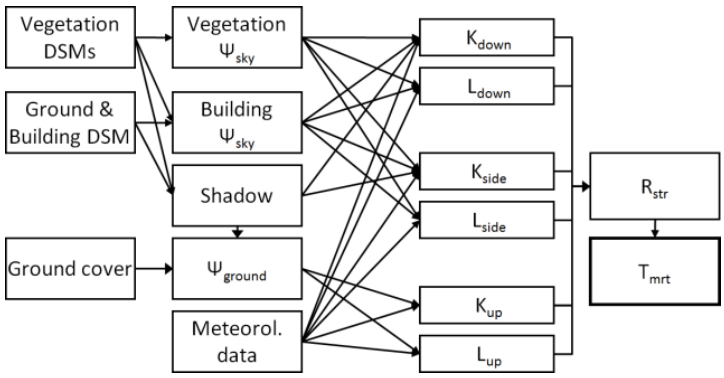

Figure 3.2: Overview of SOLWEIG with its inputs, intermediate steps and output T_{mrt}.
Source: (Lindberg and C. S. B. Grimmond, 2019)

SOLWEIG calculates T_{mrt} in urban environments. This involves determining six longwave and shortwave radiation fluxes (upward, downward, and from the four cardinal points) using inputs of global shortwave radiation, air temperature, and relative humidity (Lindberg and S. Grimmond, 2011). To compute T_{mrt} (K), when estimating the human body as a box, the mean radiant flux density (S_{str}) is calculated as follows:

$$S_{str} = \xi_k \sum_{i=1}^{6} K_i F_i + \epsilon_p \sum_{i=1}^{6} L_i F_i \tag{3.1}$$

Which is the sum of all fields of long (L_i) and shortwave (K_i) radiation components from six directions in three-dimensional space ($i = 1, \ldots, 6$), together with the angular (F), absorption (ξ_k) and emissivity (ϵ_p) factors of a person. The angular factor is 0.22 for the cardinal directions and 0.06 for the upwards and downwards directions (Lindberg et al., 2008).

The Stefan-Boltzmann law is used to calculate T_{mrt} when S_{str} is known:

$$T_{\text{mrt}} = \sqrt[4]{\frac{S_{\text{str}}}{\epsilon_p \sigma}} + 273.15 \tag{3.2}$$

To more accurately represent the shape of the human body, the option was added to consider it as a cylinder (Holmer et al., 2015). S_{str} is then calculated as:

$$S_{\text{str,cyl}} = \xi_k \cdot \left[0.28 \cdot \cos(\beta) \cdot K_{\text{dir,side}} + 0.06 \cdot (K_{\uparrow} + K_{\downarrow}) + 0.22 \cdot \sum_{i=1}^{4} K_{diff,side} \right] + \epsilon_p \cdot \sum_{i=1}^{6} w_i \cdot L_i \tag{10}$$

Where the longwave radiation is calculated as described previously, as is the shortwave radiation for the upward and downward directions. For the horizontal (cardinal) directions, shortwave radiation is divided into direct and diffuse components. $K_{dir,side}$ represents the horizontal component of the direct radiation, while $K_{diff,side}$ corresponds to the horizontal component of diffuse radiation. Direct radiation refers to sunlight reaching a surface in a straight line from the sun, whereas diffuse radiaiton is suntlight that has been scattered by the atmosphere and reaches the surface from multiple directions. In the formula, the solar elevation angle above the horizon is denoted by β (Wallenberg et al., 2023).

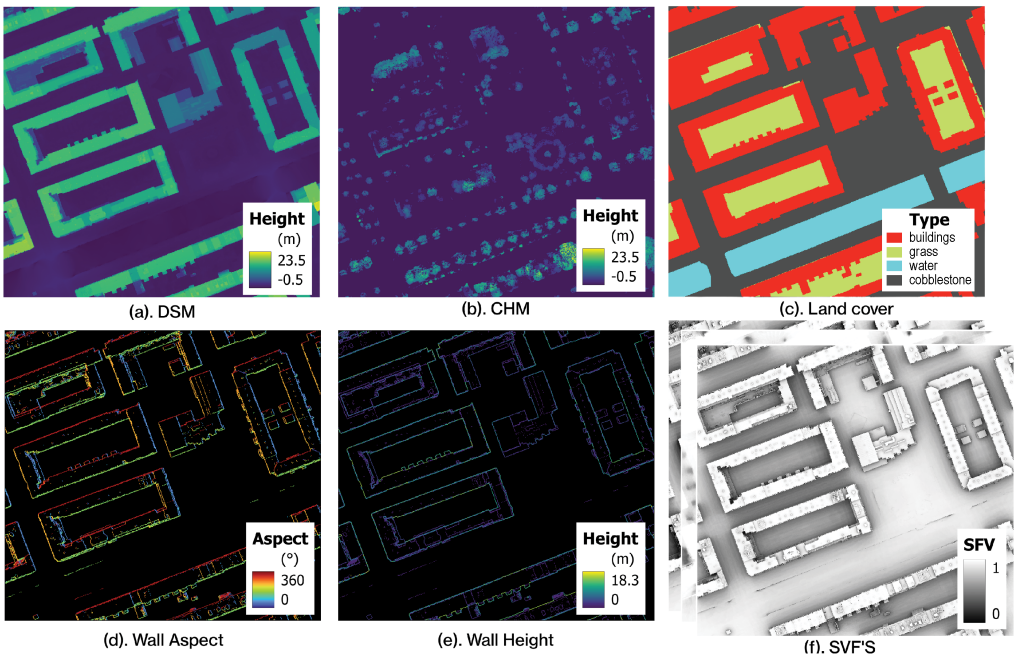

Figure 3.3: Overview of the required inputs for SOLWEIG.
(a). DSM containing building and ground heights (b). CHM containing vegetation heights (c). Land cover classification map (d). wall aspect in degrees from north (e). wall height in meters (f). SVF

SOLWEIG relies on the following inputs to calculate radiation fluxes and T_{mrt} (Figure 3.3):

- **Meteorological Data**: Includes direct, diffuse, and global shortwave radiation, air temperature, and relative humidity. If direct and diffuse radiation are unavailable, the model estimates them from global radiation data (Lindberg and S. Grimmond, 2011).
- **DSM**: Contains building and ground height values (3.3a), used to calculate spatial variations in skyview factor (SVF) (Ψ_{sky}) and shadow patterns. Shadow casting is performed by translating the DSM based on the sun's altitude and azimuth angles (Lindberg et al., 2008).
- **CHM**: Contains vegetation height values (3.3b), also used to calculate spatial variations in SVF and shadow patterns (Lindberg et al., 2016).
- **Environmental Parameters**: Surface albedo and emissivity values for walls and ground are required (Lindberg et al., 2008). Alternatively, a ***land cover raster*** (3.3c) can be used to assign cell-specific values, which also adjusts radiation fluxes and surface temperatures based on shading, sunlight exposure, and thermal properties (Lindberg et al., 2016).
- **Pre-Processor Data**: Datasets generated using UMEP preprocessor tool. These are:
  1. **Wall Aspect and Height Maps**: One raster map of the aspect of walls (3.3d), one of their height (3.3e)(Lindberg et al., 2018).
  2. **SVF Maps**: SVF quantifies the extent to which the sky is visible from a specific point. The maps (3.3f) are calculated for building and vegetation components, and the four cardinal directions. Total SVF is calculated using aforementioned shadow casting technique with 653 (or 153) uniformly distributed sun positions. The ratio that each cell is shaded determines the SVF value (Lindberg and S. Grimmond, 2010).

Over the years, improvements have been added to SOLWEIG. First, the possibility to consider the human body as a cylinder, instead of a box. The box approximation caused anomalies as a local minimum temperature was found at noon and an overestimation in later hours (Holmer et al., 2015). Second, an anisotropic diffuse shortwave radiation scheme was implemented to account for the non-uniform distribution of sky radiation, recognizing that diffuse radiation increases in the vicinity of the sun (Wallenberg et al., 2020). Finally, an anisotropic longwave scheme has been added, where emissivity increases with zenith angle (Wallenberg et al., 2023).

The output of SOLWEIG consists of six different 2D raster layers for each time step included in the meteorological data:

1. **Incoming radiation**:
   - Shortwave (W m^{-2})
   - Longwave (W m^{-2})
2. **Outgoing radiation**:
   - Shortwave (W m^{-2})
   - Longwave (W m^{-2})
3. **Shadow fraction**: Values between 0 (fully shaded) and 1 (fully sunlit), indicating shade conditions per cell.
4. T_{mrt}: The final mean radiant temperature (°C) map per cell and time step.

The first five layers represent intermediate components used in the calculation of T_{mrt}, therefore only the resulting T_{mrt} maps are directly relevant for SOLFD.

4

Methodology

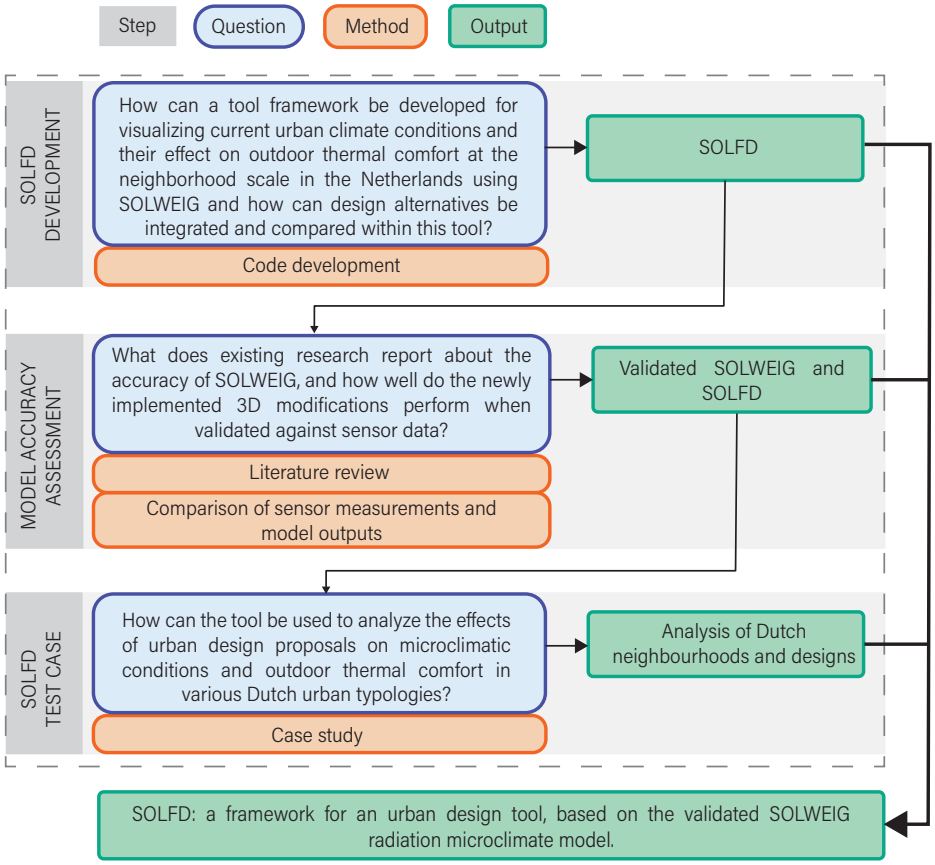

Figure 4.1: Overview of the research and the research methodology.

The methodology is divided into three main steps. A general overview is provided in Figure 4.1, which links each step to the corresponding research question, the method applied, and the resulting output. Together, these elements contribute to achieving the overall objective. The three steps are further detailed in the following sections.

Section 4.1 discusses the development of SOLFD. Section 4.2 covers the accuracy assessment of both SOLWEIG and SOLFD, based on a literature review and sensor measurements, respectively. Finally, Section 4.3 describes the setup of a case study using SOLFD.

4.1. Development of SOLFD

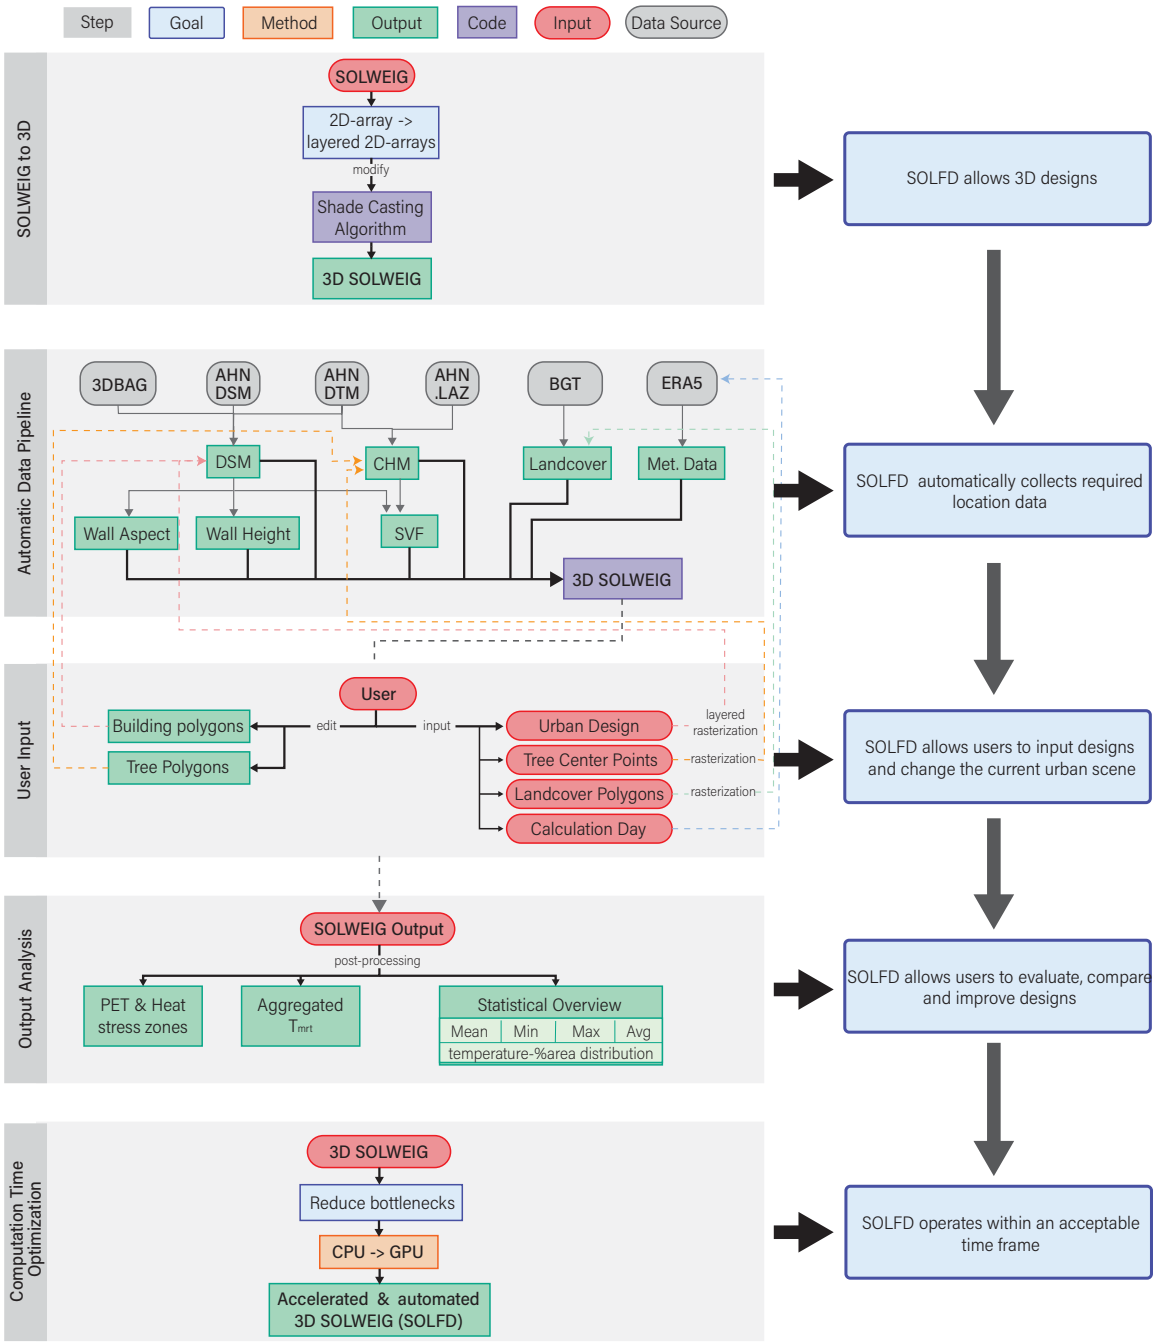

Figure 4.2: Overview of SOLFD components and development.

The main focus of this research is the development of SOLFD. Figure 4.2 presents an overview of the development process and its key components. The development is divided into five main steps, each aligned with the sub-questions outlined in Section 1.2.1. For each step, the figure highlights the methods used, the relevant data and code components, and the specific goal that the step aims to achieve. These five steps are described as follows:

1. **Extending SOLWEIG for 3D simulation** (Subsection 4.1.1): This step focuses on adapting the existing SOLWEIG model to support 3D simulations by incorporating layered 2D arrays.
2. **Automatic data collection pipeline for SOLWEIG input** (Subsection 4.1.2): A streamlined data processing pipeline is developed to automatically collect, preprocess, and rasterize all necessary input data for SOLWEIG.
3. **Integrating user designs and enabling modifications to the urban scene** (Subsection 4.1.3): Functionality is added to allow users to insert or modify design elements such as buildings, trees, and land cover directly into the SOLFD input maps. Additionally, a conversion method enables 3D urban design components to be placed in the input maps.
4. **Tailoring SOLWEIG outputs for design and analysis** (Subsection 4.1.4): The output maps from SOLWEIG are post-processed into temporally aggregated T_{mrt} and PET maps, along with summary statistics, to better support design evaluation and decision-making.
5. **Computation time optimization** (Subsection 4.1.5): Performance bottlenecks in SOLWEIG are addressed by transferring calculations from the CPU to the GPU, significantly reducing computation time.

Table B.1 provides an overview of all databases that will be used for SOLFD, and those used for its accuracy assessment.

4.1.1. Extending SOLWEIG for 3D Simulation

Currently, SOLWEIG operates on a 2D array of height data, allowing only one height value per cell, his creates a 2.5D representation of the environment (Figure 4.3a). To model designs with passages, overhangs and balconies, a full 3D representation is needed, where multiple height values are possible at the same horizontal location (Figure 4.3b).

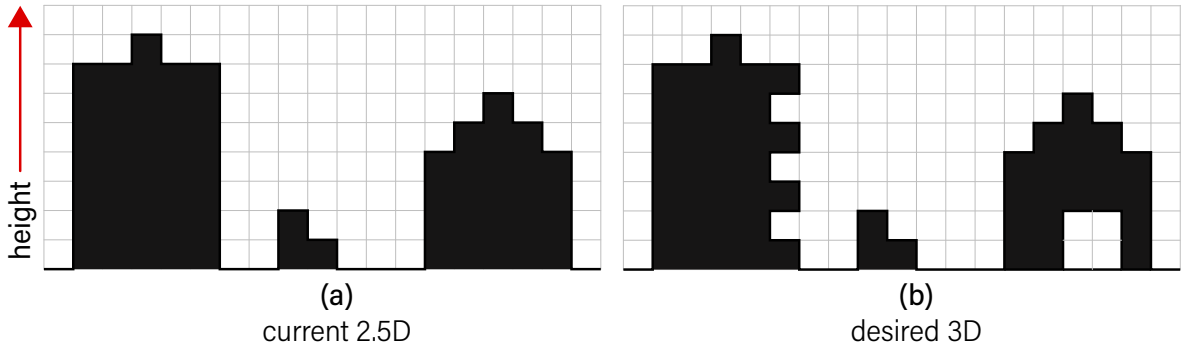

Figure 4.3: Vertical sections of height data illustrating (a). the current 2.5D situation, where a 2D array is extruded to represent height, and (b). the desired 3D situation, where multiple height values can exist at the same horizontal location.

From a climate perspective, a 3D representation is required for accurately simulating microclimatic conditions influenced by urban design. Elements such as overhangs, balconies, and passages affect how solar radiation and longwave radiation interact with surfaces and spaces. These structures can provide shading, block or reflect radiation, and create localized thermal environments that directly influence surface temperatures and T_{mrt}. A 2.5D model, which represents only the topmost surface per grid cell, cannot capture these interactions. For example, shaded areas beneath elevated structures are represented as solid masses in 2.5D, effectively erasing its possible cooling effect on the microclimate. Similarly, recessed façades or balconies that reduce solar exposure on building surfaces or adjacent ground areas are not represented. Accurately modeling these effects is important for evaluating urban design strategies aimed at mitigating urban heat stress.

To enable 3D simulations, voxelization would be a logical next step, extending the current 2D arrays into a full 3D grid. However, implementing voxel-based modelling would require a significant overhaul of SOLWEIG's existing structure. As an alternative, I propose to adapt the logic currently used for trees.

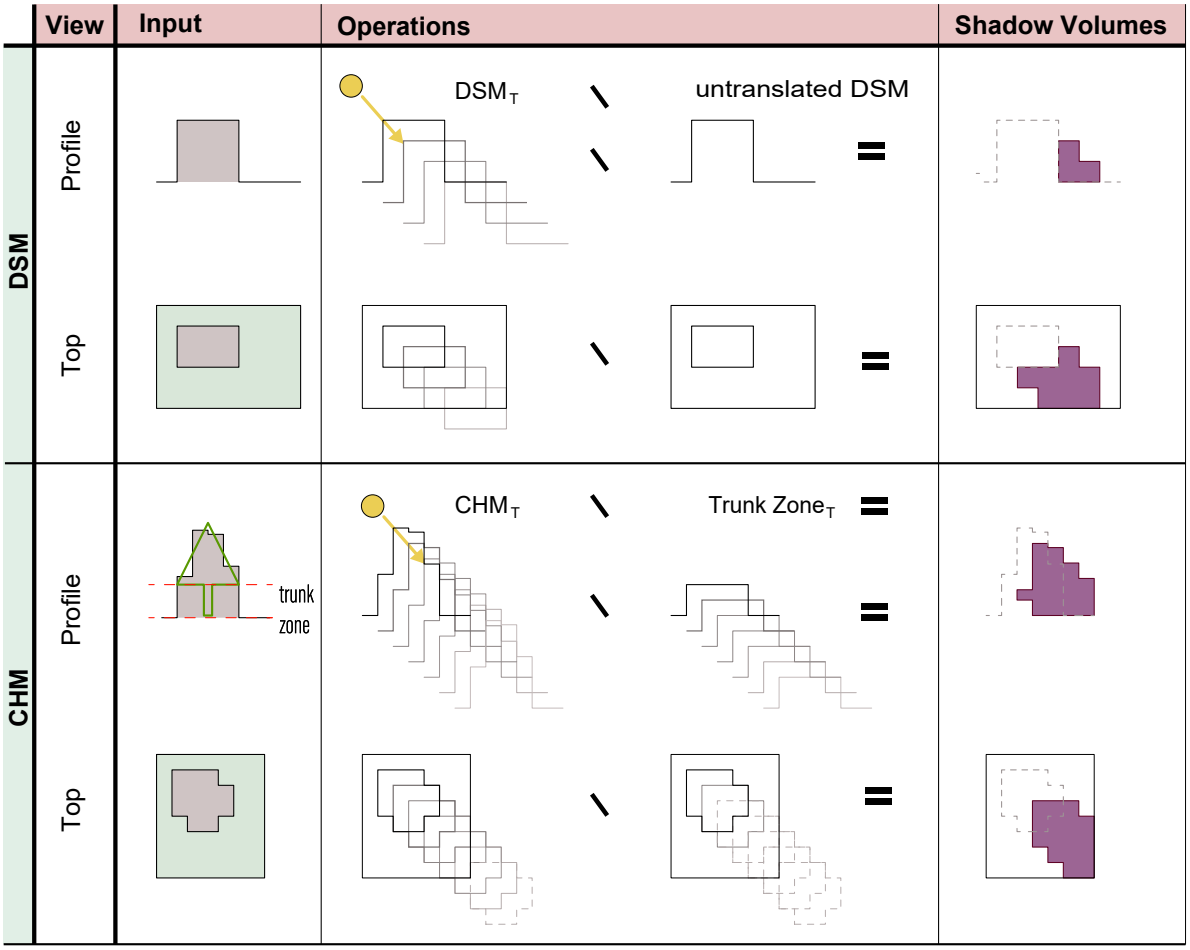

Figure 4.4: Shadow casting methods used in SOLWEIG. Top: single-step translation of the DSM to compute shadow volumes cast by buildings. Bottom: stepwise translation of the CHM and trunk zone for vegetation, where only canopy above terrain and above the trunk zone contributes to shading.

The main part of SOLWEIG affected by introducing 3D is the calculation of shadow volumes. For built structures, shadow casting is performed by translating the DSM once along the solar vector, reducing height according to the z-component of the sun's direction. The translated surface is compared to the original DSM; if the translated height is above the terrain, it casts shade. For vegetation shadow casting (Lindberg and S. Grimmond, 2011), a trunk zone is introduced to represent the vertical space beneath the canopy that does not cast shade. Both the CHM and trunk zone are translated in multiple steps along the solar vector. Shadows occur where the translated CHM lies above both the original DSM and the trunk zone, while the trunk remains below ground level. Figure 4.4 illustrates both methods, with simplified set theory notations.

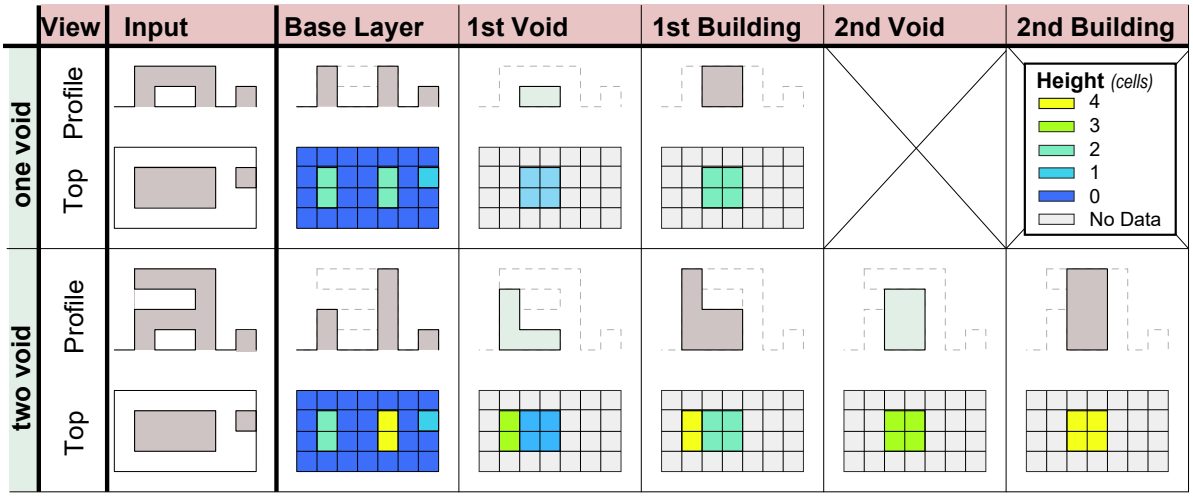

Figure 4.5: Illustration of the layered array concept for 3D shadow casting.
Two examples are shown: one with a single vertical void (resulting in three layers), and one with two voids (resulting in five layers). Profile views demonstrate how gaps and building volumes are organized across the vertical stack, top views show the corresponding heights displayed in the layered array.

For 3D shadow casting, I propose a method using a layered 2D array. Each layer is classified as either a **void layer** or a **building layer**. Void layers do not cast shadows but block shadow projection from building layers above them. Building layers represent surfaces that can cast shadows.

The structure of the layered array always follows a consistent pattern:

- **Layer 0 (base layer)**: Contains the height of the ground surface and all buildings up to the first point where where they open directly to open air.
- **Void layer(s)**: Represent the vertical extent of the open-air void using true height values.
- **Building layer(s)**: Contain building heights above voids that can cast shadows.

This structure can repeat for multiple voids. Figure 4.5 illustrates the layered array concept using two examples: one with a single void and one with two voids. Profile and top views are shown. A single void results in three layers (base, void, building), while two voids produce five layers. The resulting heights in the array layers are shown in the top view.

The shadow casting logic is similar to the CHM method, but with the possibility of additional non-shadow-casting layers. A shadow is cast in step i when the translated base layer is higher than the original base layer, or when a building layer lies above the base layer, while the corresponding void layer lies below the base layer. Figure 4.6 illustrates this principle, with simplified set theory notations.

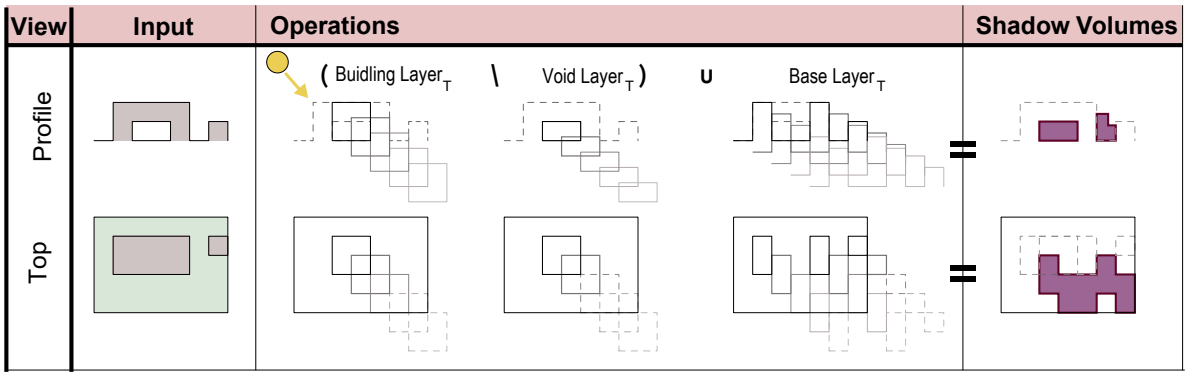

Figure 4.6: 3D Shadow casting using the layered array approach. Void layers block projected shadows from higher building layers. A shadow is cast only if the building layer lies above the base and the corresponding void layer lies below it.

For a true 3D shadow representation, the casted shadow could also be stored as a layered array. However, since SOLWEIG is currently designed to operate with single-layer inputs, implementing this would require significant changes throughout the codebase. Moreover, the main interest lies in estimating T_{mrt} at pedestrian height, and SOLWEIG itself notes that it is only validated for use at that level. Therefore, all further calculations in SOLFD will consider only the base layer. As a consequence, the shadow map will represent shade conditions only at this layer, and after shadow computation the code will continue with just the base layer.

Another necessary modification for full 3D implementation involves the 2D rasters of wall heights and wall aspects used in ground view factor (GVF) calculations (see Figure 3.3 for the wall rasters). Currently, these rasters store a single value per cell. For full 3D representation, they should be extended to the aforementioned layered 2D array format, where for each building layer the wall heights and aspects would be stored. However, as noted earlier, for further calculations SOLFD will focus exclusively on the base layer; therefore, only the wall height and aspect arrays for the base layer are generated.

Finally, my modifications do not consider surfaces above-head (ceilings) for (reflected) radiation. Incorporating these surfaces would require changes to the radiation physics, which is outside the scope.

Optimal Number of Void Layers

Theoretically, the number of void layers is only limited by the available computational resources. However, each additional modeled void layer requires tracking two extra layers: the void layer itself and its complementary building layer above it. Therefore, testing is necessary to determine the optimal number of void layers for accurate shadow representation, especially since including more voids increases computation time, a factor that should be minimized for efficient use within SOLFD.

Two different cases containing multiple voids will be tested. For each case, simulations will be run twice: once with all void layers included, and once with only the void layer at pedestrian height. Since the shadow calculations are used in generating both the SVF and T_{mrt} maps, the resulting values for SVF and T_{mrt} will be compared to assess the impact of void layer inclusion. The test cases are illustrated in Figure 4.7.

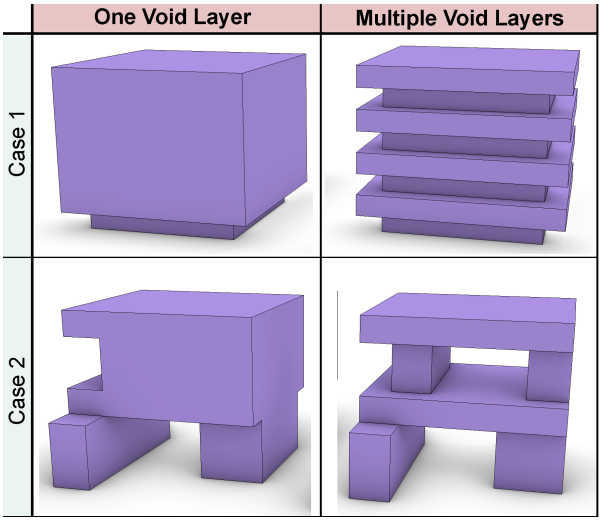

Figure 4.7: The two test cases, each shown with a single void layer (left) and multiple void layers (right).

4.1.2. Automatic Data Collection for SOLWEIG Input Pipeline

As outlined in Section 3.2, SOLWEIG requires multiple input datasets, including a DSM, CHM, land cover, SVF, wall height and aspect, and meteorological data. The SVF is generated by applying an adapted version of the existing SVF algorithm to the DSM and CHM. All other datasets are acquired and processed automatically as follows:

Digital Surface Model

The creation of the DSM builds on a previously developed methodology for generating DSM and DTM inputs for SOLWEIG using Dutch data sources (Monahan et al., 2024). The method combines high-resolution elevation data from Actueel Hoogtebestand Nederland (AHN) with the accurate building footprints from the 3DBAG dataset (`https://docs.3dbag.nl/en/`).

In the updated workflow, elevation data is retrieved on demand from the AHN Web Coverage Service (WCS) for a user-defined spatial extent. The Web Coverage Service (WCS) of AHN returns rasters with missing values (*no data*). As described in Monahan et al. (2024), these values will be filled using Laplace interpolation. This interpolation method is selected for its desirable mathematical and computational properties: it is exact (reproducing input values at known points), yields a smooth surface across most regions, is strictly local (reducing computation time), and requires no user-input (Ledoux et al., 2023).

The Laplace Interpolation method inserts a temporary point at the interpolation location into a Delaunay triangulation. The lengths of the Delaunay and Voronoi edges between this point and its natural neighbors are used to weight each neighbor's elevation value. The interpolation weight for a neighbor is proportional to the length of the Delaunay edge (orange line in Figure 4.8) and inversely proportional to the distance between the point and the neighbor (green line in Figure 4.8) (Ledoux et al., 2023).

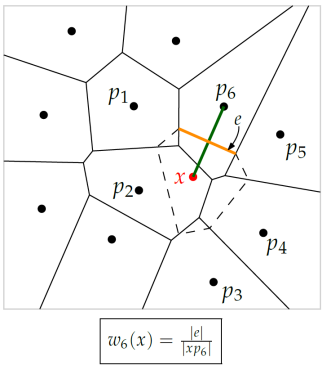

Figure 4.8: The weight for the Laplace interpolant for one neighbor. The weight is proportional to the length of the Delaunay edge (orange line) and inversely proportional to the distance between the point and the neighbor (green line).
Source: (Ledoux et al., 2023 p.55)

.

To apply the Laplace Interpolation, raster cells are converted into point coordinates, with the row and column indices becoming y and x, and the cell value as z. This is shown in Figure 4.9a–b. For full coverage, the four corner points must be included in the triangulation. If any corner points are missing, interpolation gaps may occur, as shown in Figure 4.9c. To prevent this, missing corners are filled using nearest-neighbor interpolation, relying on spatial autocorrelation to provide a reasonable approximation.

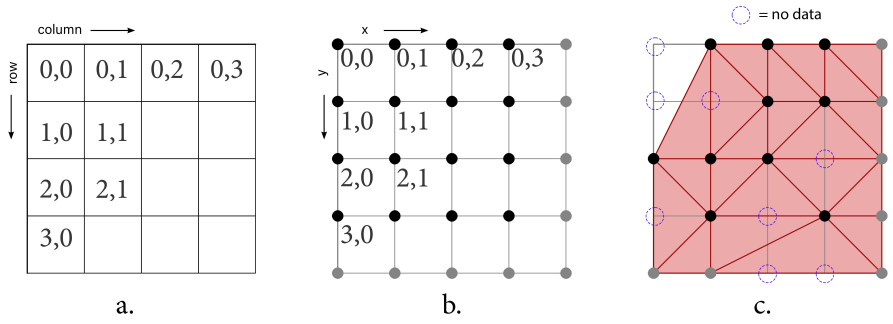

Figure 4.9: Converting the (a) raster cells to (b) x,y coordinates. (c). Interpolation issues can arise when the corner points are missing. Note: the xy-coordinates are given in yx-format.
Source: (Monahan et al. (2024), p. 22)

After filling, both DSM and DTM rasters are available for the study area. However, the DSM includes vegetation. To remove this, a cleaned DSM is created by combining the DTM with building heights extracted from the original DSM. As shown in Figure 4.10, the DSM is masked using 3DBAG building footprints: (a) shows the building polygons, while (b) presents the resulting raster in which only the elevation values within building footprints are retained.

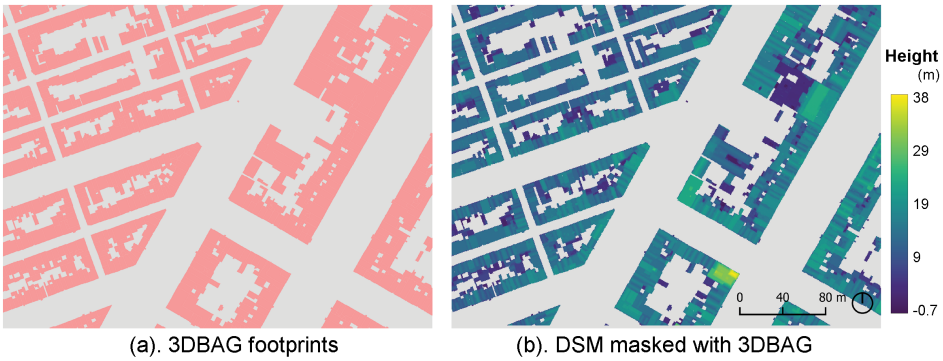

Figure 4.10: Extacting the building heights of a DSM by masking with 3DBAG building footprints. (a) The 3DBAG building polygons overlaid. (b) DSM values retained within building footprints, with all other values removed.

Canopy Height Model

The creation of the CHM also builds on a previously developed methodology described by Monahan et al. (2024), with several enhancements. The method uses AHN point cloud data in a preprocessed and structured format provided by the GeoTiles database (`https://geotiles.citg.tudelft.nl/`), which facilitates easier handling. A key improvement is the ability to retrieve point cloud tiles based on a user-defined bounding box. Tiles are selected using their spatial extents, with a preference for AHN5; if unavailable, the method automatically defaults to AHN4. The selected tiles are merged and cropped to match the bounding box and stored in .LAS format. Although this format ensures compatibility and precision, the more storage-efficient .LAZ format may be preferable in future implementations.

Vegetation points are identified in the LAS files using Normalized Difference Vegetation Index (NDVI). The NDVI is calculated from the Near-Infrared (NIR) and Red (RED) bands:

$$\text{NDVI} = \frac{\text{NIR} - \text{RED}}{\text{NIR} + \text{RED}} \tag{4.1}$$

GeoTiles provides colored AHN point clouds by combining LiDAR data with aerial imagery. This enables NDVI-based classification at high resolution. Before computing NDVI, the point cloud is filtered to exclude points classified as ground, buildings, or water. In AHN5, building façade points are included in the "unclassified" category, which reduces the reliability of this class for vegetation detection.

Figure 4.11: Points in an example scene classified as vegetation when NDVI = 0.1 (classified points in red).

To classify vegetation, a threshold must be selected for the NDVI value. Figure 4.11 shows LAS points labeled as vegetation using a threshold of 0.1. While this low threshold helps capture shaded and smaller trees, it does not eliminate all classification issues. Some vegetation is still missed, as seen in Figure 4.12(a). Conversely, lowering the threshold increases the likelihood of false positives: urban features such as street lights, walls, or signs being misclassified as vegetation, illustrated in Figure 4.12(b). Therefore, a threshold of 0.1 is retained as a balance between under-detection and misclassification.

Figure 4.12: Misclassification examples: (a) missing tree points and (b) non-tree points wrongly labeled as vegetation (classified points in red).

To generate the CHM raster, the classified vegetation points are interpolated using Laplace interpolation. A Delaunay triangulation is constructed from vegetation points, and a k-d tree is used to find the nearest vegetation point to each raster cell center. Interpolation is performed only if a vegetation point lies within 1 meter, avoiding false vegetation heights in empty areas.

Finally, the normalized CHM is created by subtracting the filled DTM from the CHM. Values below 2 meters are set to 0, as vegetation of this height does not provide effective shading and may include misclassified objects like cars. For example, a low bush would only shade the feet of a standing person. Similarly, values above 40 meters are removed to eliminate outliers and unrealistic vegetation heights for the Netherlands. Cells with values below the DSM are also set to zero. A median filter smooths the raster and fills small gaps.

This method extracts only the CHM. The trunk height layer used by SOLWEIG is not generated here, as it requires more intensive 3D analysis to cluster tree points and separate cylindrical trunks from denser canopy to identify canopy base height. This process is computationally expensive and slow. Therefore, SOLWEIG's default approximation of trunk height as 25% of total height is used, which introduces some inaccuracies.

Wall Height and Aspect

The calculation of the wall height array follows the original SOLWEIG implementation. Wall heights are derived directly from the DSM by first applying a cross-shaped maximum, which finds the highest neighboring height at each cell. The difference between this maximum and the cell value is computed, and only those differences that exceed a predefined threshold are considered as potential walls. This method results in walls being buffered by one cell from the actual building.

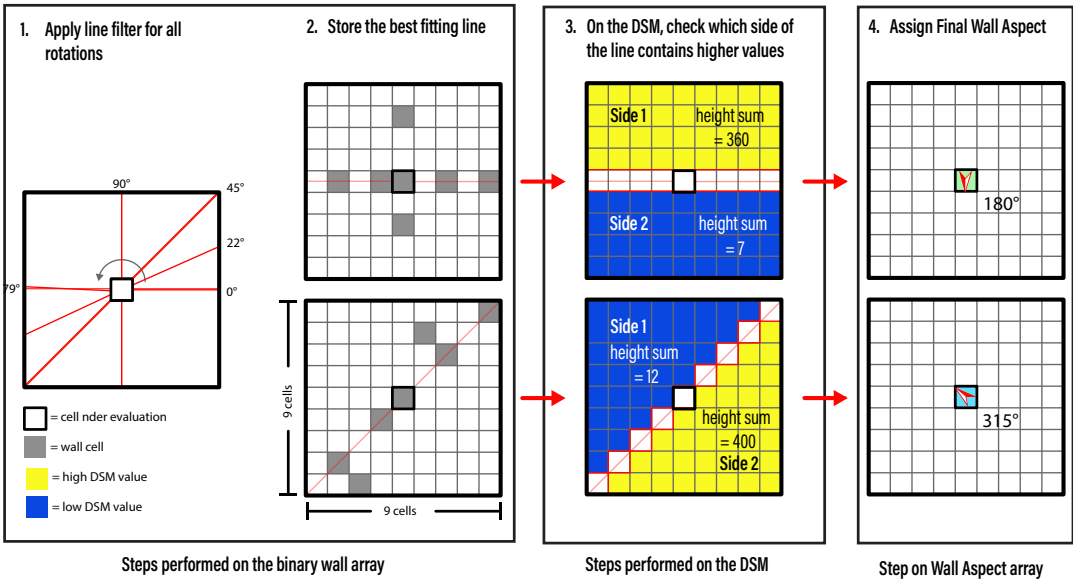

Figure 4.13: Goodwin et al. (2009) method to calculate wall aspect.

Wall aspects are then determined using an adapted linear filter method from Goodwin et al. (2009), as illustrated in Figure 4.13. The procedure consists of four steps: First (Step 1), a 9 × 9 cell moving window is rotated from 0° to 179°, applying a line filter at each angle to evaluate wall cell alignment. Second (Step 2), the best-fitting line is selected based on the highest number of aligned wall cells. Third (Step 3), the outward-facing direction is determined by comparing the sum of DSM heights on either side of the best-fit line, with the direction pointing from the higher to the lower side. Finally (Step 4), the resulting wall aspect is assigned to the central cell.

However, the original method for calculating wall aspect is computationally intensive as it evaluates 180 directions per wall cell. To improve performance, I replace it with the Sobel filter, a widely used edge-detection technique in image processing (Kanopoulos et al., 1988). The Sobel filter estimates the gradient of the DSM in both the x and y directions by convolving the raster with two 3×3 kernels:

$$G_x = \begin{bmatrix} -1 & 0 & +1 \\ -2 & 0 & +2 \\ -1 & 0 & +1 \end{bmatrix}, \quad G_y = \begin{bmatrix} +1 & +2 & +1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$$

The direction of the steepest ascent, or wall aspect angle θ, is then calculated as:

$$\theta = \arctan 2(G_y, G_x)$$

where G_x and G_y represent the horizontal and vertical gradients, respectively. The resulting angle is in radians. It is then converted to degrees and mapped to the range $[0°, 360°]$ after which it is assigned to the corresponding raster cell, producing the Wall Aspect raster.

Land Cover
Land cover is used to calculate the longwave radiation from surface temperature (T_s). T_s is estimated using a sinusoidal model. T_s varies with land cover type and is influenced by sun exposure, shadow, and the clearness index. Water temperature is derived from the diurnal air temperature. Models based on observations for bare soil, grass, asphalt, and cobblestone have been added to SOLWEIG, along with assumed values for walls and roofs (Lindberg et al., 2016).

While material differentiation is supported, many common Dutch urban materials are missing. Up to 20 materials can be included in theory, but extending the material list is out of scope (see Section 1.2.2).

For land cover raster creation, the most detailed Dutch dataset is used, which is the Basisregistratie Grootschalige Topografie (BGT) (`https://api.pdok.nl/lv/bgt/ogc/v1/`). The coarser but faster TOP10NL dataset is included as a back-up option (`https://api.pdok.nl/brt/top10nl/ogc/v1/`). Building footprints are reused from the 3DBAG dataset (Section 4.1.2). The process consists of three steps: requesting data using a bounding box, processing it, and rasterizing it. Because the BGT API returns its own 'feature types', or material classes, these are mapped to the closest matching SOLWEIG classes. Details on the conversion rules and their motivations can be found in Section A.3.

1. **Data Request and Processing:**
   (a) **Terrain polygons** are requested from the BGT and classified into SOLWEIG-compatible types based on the 'fysiek voorkomen' (physical appearance) attribute using predefined rules. This includes the collections *begroeid terreindeel*, *onbegroeid terreindeel*, *scheiding vlak*, and *ondersteunend waterdeel*.
   (b) **Road polygons** are also requested from the BGT, specifically from the *wegdeel* and *ondersteunend wegdeel* collections. These are classified by their 'fysiek voorkomen' and retain the 'relatieve hoogteligging' (relative height) attribute for use in determining rasterization order.
   (c) **Water features** are extracted from the BGT using the *waterdeel* collection and are directly assigned to the water land cover type.
   (d) **Bridging features** are obtained from the *overbruggingsdeel* collection in the BGT, and assigned the *paved* (cobblestone) land cover type. The relative height is also preserved. An 'overbruggingsdeel' is defined as "part of a movable or fixed structure connecting two points separated by water, a road, or similar, essential for the construction" (Geonovum, 2020).
2. **Rasterization Order:** Rasterization is performed in a defined sequence to ensure correct layering of overlapping land cover types. In practice, terrain polygons do not overlap with water or other features, so conflicts mainly involve roads, bridging elements, and buildings. Users can choose whether higher or lower roads appear on top at intersections. The order is:
   (a) Water polygons
   (b) Terrain polygons (would overwrite water, though no overlap occurs in the data)
   (c) Roads and bridging features (overwrite terrain and water)
      - Roads are split based on their relative height compared to bridging features.
      - With the "roads on top" option, low roads are rasterized first, then bridges, then high roads. This ensures bridges cover low roads, and high roads appear above all.
      - Otherwise, bridges are rasterized first, followed by roads from high to low. This reveals the full polygon of the lowest road.
   (d) Building footprints (overwrite all)

Each raster cell is assigned one land cover class. Small rasterization gaps between BGT features and a mismatch with the more detailed building footprints from 3DBAG can lead to unclassified cells. A default fill value can be specified to address this. The resulting land cover raster is directly usable as input for SOLWEIG.

Meteorological Data:
Currently, SOLWEIG users can download required meteorological data from ERA5 for any location and day using the ERA5 plug-in included in UMEP. ERA5 is a global climate reanalysis product from the European Centre for Medium-Range Weather Forecasts, which combines observations and numerical models to provide hourly climate variables. The data are accessible via the Copernicus Climate Data Store (Copernicus Climate Change Service (C3S), 2024).

However, to do this, users must create an ERA5 account and place an access token in specific files on their computer. Moreover, upstream issues can cause data retrieval and processing to fail SOLWEIG developers have recommend downloading meteorological data from Rokka, 2021 (`https://www.shinyweatherdata.com/`), which offers processed ERA5 data that can be converted into a format accepted by SOLWEIG.

To simplify meteorological data retrieval for designers and to provide representative weather examples for design testing, four example meteorological datasets are prepared. These are based on ERA5 hourly data downloaded from `https://www.shinyweatherdata.com/` for a central location in the Netherlands (latitude 52.25°, longitude 5.5°) spanning the period 2000 to 2024. Users requiring specific days or locations can download and preprocess data themselves using the UMEP meteorological data preprocessor.

Daily mean and maximum air temperatures are calculated, and only days with a mean temperature above 15°C are considered. Based on these temperatures, days are grouped into three categories:

- 20–25°C (mild summer days)
- 25–30°C (warm summer days)
- >30°C (hot/tropical days)

For each category, an "average day" dataset is created by averaging relevant meteorological variables (air temperature, relative humidity, pressure, global, direct, and diffuse irradiance) across all qualifying days. These averaged conditions are assigned to 1 August 2023, representing midsummer solar angles and day length.

Additionally, an 'extreme day' dataset is created using the 97th percentile of all variables from the >30°C category. This scenario is assigned to 21 June 2023 (summer solstice), when solar elevation is at its peak and shading is minimal.

4.1.3. Integrating User Designs and Enabling Modifications

The methodology for incorporating user input involves enabling users to modify the current urban landscape by adding or removing elements such as buildings, trees, and ground surface classification.

Modifying Existing Elements
Users can remove buildings and tree clusters by selection. For this, the pand id included in the BAG dataset is used for the buildings. For the trees, tree clusters are vectorized and given an unique ID. The ID is used to fetch selected polygons of objects to be removed, which are then used to mask and update the raster CHM and DSM. For the CHM, the value is set to 0. For the DSM, the values are replaced by those of the DTM.

Tree Insertion
Users can insert new trees into the CHM by selecting a cell in the array and specifying the input tree parameters. In addition to manual parameter input, preset options based on common Dutch tree species are provided, with parameter values approximated from tree databases. This section describes the tree insertion method and the process for generating these species-specific presets.

The Tree Insertion Method
Tree canopies are represented using simple mathematical models, which are rasterized and inserted into the CHM. The method requires the row and column indices of the tree center, along with the parameters: total height, crown radius, trunk height, canopy shape, and canopy base height.

The insertion method consists of the following steps:

1. **Distance calculation**: For each cell within the crown radius (plus a small buffer), the Euclidean distance from the tree center is calculated. This distance determines the height output based on the selected canopy shape in the next step.
2. **Shape application**: Four basic canopy shapes are supported: cone, gaussian, parabolic, and hemisphere. These shapes approximate common tree forms. The canopy height is calculated from the trunk height up to the tree top, with an optional `canopy_base_height` offset that allows the shape to start a bit higher or lower than the trunk itself. This gives flexibility in shaping how the canopy sits on the trunk. The hemisphere can be morphed towards more parabolic or elliptical forms by adjusting the height-to-radius ratio. The formulas and resulting canopy outputs are shown in Figure 4.14.
3. **Noise addition**: To add variability, random noise is applied to the canopy using a Gaussian distribution:
$$\text{canopy}(x, y) \mathrel{+}= \mathcal{N}(0, \text{randomness}) \tag{4.2}$$
where $\mathcal{N}(0, \text{randomness})$ is Gaussian noise with mean zero and a given standard deviation.
4. **CHM Update**: The canopy is inserted by taking the maximum of the existing CHM and the new canopy, preserving the highest value at each cell.
5. **Trunk zone update**: Using the same spatial mask as the canopy, the trunk height is inserted into the trunk height array.

Additionally, users can insert randomized trees with a chosen shape and location. In this case, height, crown radius, trunk height, and canopy base height are sampled within user-defined ranges, with a minimal canopy height enforced to ensure realistic tree shapes.

Creation of Preset Options based on common dutch trees
To simplify tree insertion and avoid manually specifying parameters, a preset tree database is included based on the most common tree species in the Netherlands. Data was collected for the five most prevalent species in the Netherlands, as listed in Table 4.1. Together these species account for 32% of 2.5 million trees in 28 Dutch municipalities.

Table 4.1: The five most common tree species in the Netherlands. (Natuur & Milieu, 2023).

Species	Dutch name	English name	Count (×1000)	Percentage
Quercus robur	Zomereik	English oak	309	12%
Fraxinus excelsior	Es	Ash	186	7%
Tilia × europaea	Gewone linde	Common lime	125	5%
Salix alba	Schietwilg	White willow	97	4%
Platanus × hispanica	Gewone plataan	London plane	92	4%

To determine suitable parameters for these species, several tree databases are examined. The goal is to obtain tree parameters not only by species but also by age category, allowing greater control during tree insertion. The datasets should provide at minimum tree height and canopy radius, and preferably trunk height as well. Table A.4 presents an overview of the consulted databases, including species coverage and data suitability.

Only two datasets contain the required parameters for the target tree species: the Urban Tree Database (McPherson et al., 2016a) includes *Fraxinus excelsior*, and the Open British Arboricultural Record Dataset (OpenBARD) (Fennell and Fay, 2025) covers all five species.

Shape	Formula	Output
Cone 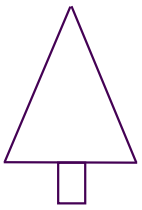	$\text{canopy}(x,y) = \max\left((\text{height} - \text{trunk\_height}) \times \left(1 - \frac{\sqrt{x^2+y^2}}{\text{crown\_radius}}\right), 0\right) + \text{trunk\_height}$	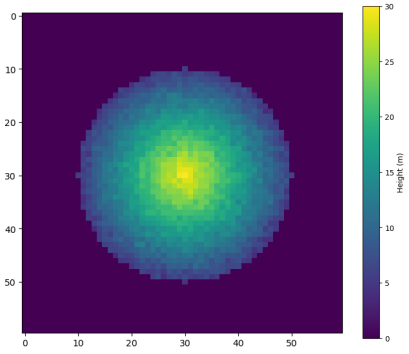
Gaussian 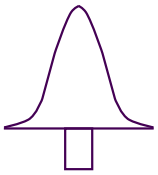	$\text{canopy}(x,y) = (\text{height} - \text{trunk\_height}) \times \exp\left(-\frac{(x^2+y^2)}{2\left(\frac{\text{crown\_radius}}{2}\right)^2}\right) + \text{trunk\_height}$	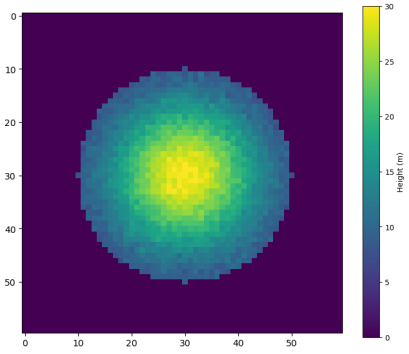
Parabolic 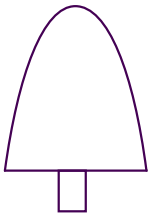	$\text{canopy}(x,y) = \max\left((\text{height} - \text{trunk\_height}) \times \left(1 - \left(\frac{\sqrt{x^2+y^2}}{\text{crown\_radius}}\right)^2\right), 0\right) + \text{trunk\_height}$	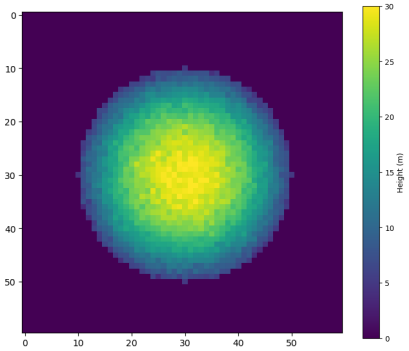
Hemisphere 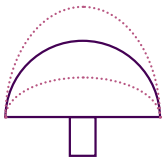	$\text{canopy}(x,y) = \frac{\sqrt{\max\left(\text{crown\_radius}^2 - (x^2+y^2), 0\right)}}{\text{crown\_radius}} \times (\text{height} - \text{trunk\_height}) + \text{trunk\_height}$	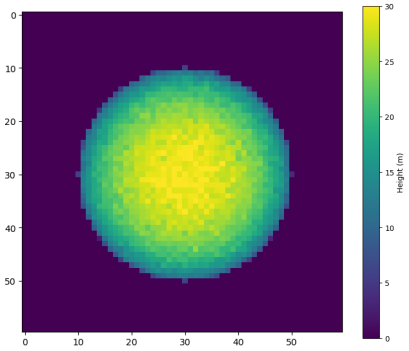

Figure 4.14: The four different kind of tree shapes included in the model; Cone, gaussian, parabolic and hemisphere, with their formulas. Resulting output tree with random noise added.

McPherson et al. (2016b) provide urban tree growth equations that calculate the required tree parameters based on age using species-specific coefficients. These equations are applied to estimate parameters for *Fraxinus excelsior* trees aged 1 to 50 years, and the results are stored in the tree database.

OpenBARD does not provide growth equations to calculate tree parameters by age, nor does it include trunk height data. Instead, it contains measurements of individual trees with recorded parameters. The age is given descriptively (e.g., young, mature, late mature) rather than as exact values. Sample sizes range from 67 to 1000 trees. For this dataset, tree height and canopy radius are averaged per age category, while trunk height is approximated as 30% of tree height, resulting in a rough estimate of tree parameters per species. Consequently, the tree database provides a coarse approximation of tree characteristics for species other than *Fraxinus excelsior*. This level of approximation is considered acceptable, as tree insertion is simplified by representing trees with a uniform, rounded crown shape.

Finally, the database functions as follows: the user specifies the tree species and age category, and the database returns the associated tree parameters for that age, which are then used in the tree insertion method. Age serves solely as a lookup key; no further growth modeling is performed, and existing trees in the CHM remain unchanged.

3D Urban Design Integration

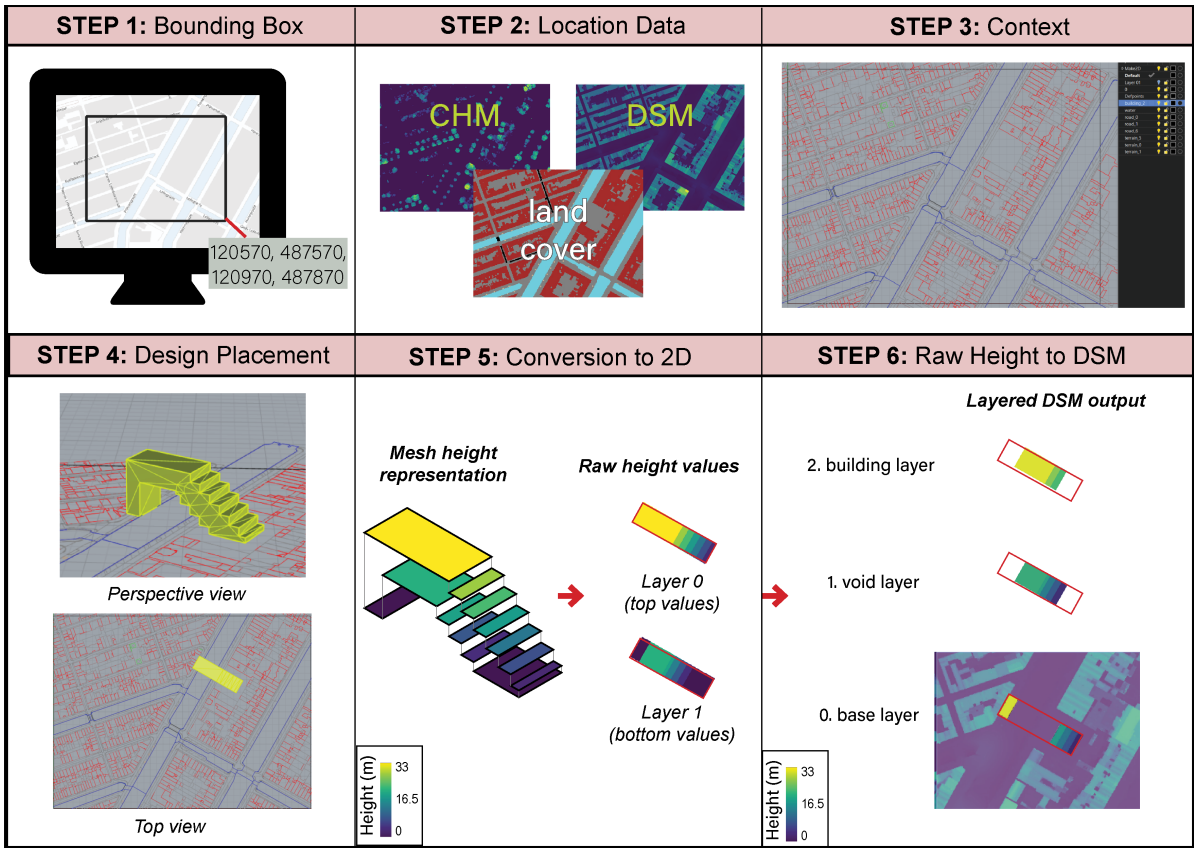

Figure 4.15: Step-by-step overview (1–6) of the 3D urban design integration process.

The methodology for enabling 3D urban design integration in SOLFD consists of the following steps, illustrated in Figure 4.15, where each step is numbered accordingly:

1. **Bounding Box Specification**: The user defines the bounding box of their design area.
2. **Location Data Generation**: SOLFD generates the necessary location data, as described in Section 4.1.2.
3. **Context Contour Line Export**: SOLFD exports contour lines of the land cover polygons derived from the relevant BGT features. Each polygon includes its land cover classification code and type (e.g., road, terrain, water). These polygons are exported into separate layers within a `.DXF` file, with each layer named and color-coded according to its type. The `.DXF` format is an open standard widely supported by CAD software, enabling users to directly import and manipulate the contextual geometry in their design environment.

4. **Design Placement and Export**: Users design or place their 3D model within the exported context and export it as a triangulated `.OBJ` file, a widely supported format in CAD applications.

   Important: Models must be clean and well-structured. As described in the following sections, the triangles in the `.OBJ` file are used to extract heights and identify void layers. Therefore, the model should only include exterior triangles of the building(s). While disconnected components are allowed, triangles must not intersect each other internally, as these will be incorrectly interpreted as part of the building envelope. Additionally, since cell centers are used for height interpolation, vertical triangles must not lie exactly on cell centers to avoid ambiguities.

5. **Conversion to Raw Height 2D Arrays**: For this I have adapted an existing Rust implementation of Pađen (2023). This code converts a triangulated 3D `.OBJ` file into a 2D raster `.TIFF` or ASCII file. To achieve this, the algorithm loops over all cell centers of the target raster grid and identifies which triangle(s) contain the point or intersect with it. Using barycentric interpolation, it calculates the height at the cell center based on the triangle's vertex heights. Initially, only the maximum intersection height per cell was retained in a single-layer height map.

   In the extended version, this approach is modified to capture all unique height values at each cell, enabling a full 3D representation. These values are sorted in ascending order from layer 1 onwards, but with layer 0 storing the topmost surface. This multi-layer raster containing raw heights is returned as a `NumPy` array, making it directly usable in Python. The procedure is summarized in Algorithm 4.1.

6. **SOLWEIG Input Preparation**: The raw height arrays are converted into the correct layered array input format for 3D SOLFD, depending on the number of void layers the user wishes to include. This process is explained in more detail below. The resulting base layer, which contains only the building footprint before insertion into the DSM, is also used to update the land cover to the building class at the footprint location.

Algorithm 4.1: Converting 3D model to 2D layered arrays

Input: OBJ file path `file`, raster size $(ncols, nrows)$, cell size $cellsize$, origin (x_0, y_0), NoData value $nodataval$
Output: 3D raster array of height layers

```
Load triangles from file into Triangles structure;
Initialize empty Raster with given size, parameters and 0 layers;
foreach triangle T in mesh do
    Get bounding box B of T;
    foreach cell (i, j) within B do
        Compute cell center (x, y);
        if (x, y) is inside or on boundary of T then
            Interpolate height z from triangle vertices;
            Add z to height set at cell (i, j);
foreach cell (i, j) with height values do
    Sort height values in descending order;
    Assign highest value to layer 0;
    Assign remaining values to layers 1, 2, ...;
    If more layers needed, resize raster;
return 3D raster array
```

Multi-layer DSM Generation from OBJ

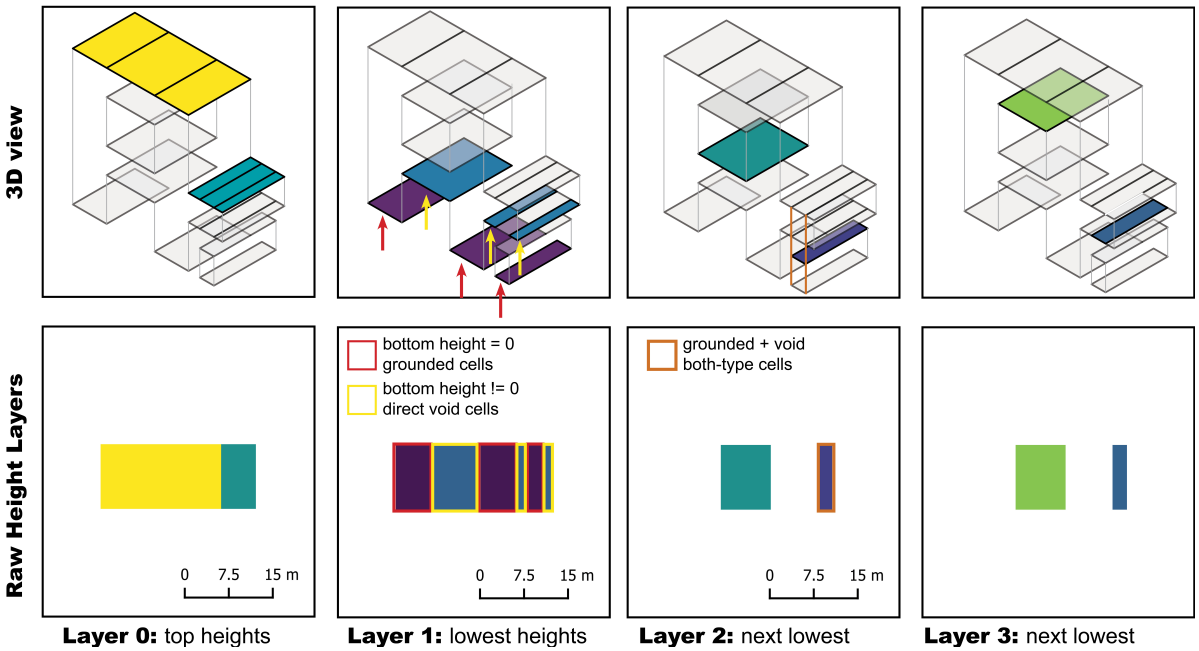

Figure 4.16: Top: 3D visualization of the inserted object showing where the raw height layers originate. Bottom: an example of the sorted raw height arrays. The diagram also illustrates the classification of grounded cells, direct void cells, and both-type cells.

I have created a method that constructs a stack of DSM layers from the sorted raw data heights at each raster cell. The logic is as following:

- **Input:** The algorithm takes as input a list of 2D arrays, each representing sorted height values at each raster cell from the triangulated OBJ file. These values are ordered from bottom to top, except for the first layer (layer 0), which stores the top surface heights. Figure 4.16 shows an example of a raw sorted height dataset along with its 3D context. The user defines the number of voids to include. The total number of layers becomes $2n + 1$, where n is the number of voids.
- **Step 1: Initialization** A 3D array is initialized to hold the DSMs, with all values set to `NaN`.
- **Step 2: Cell classification** Cells are classified using logical masks:
  - *Grounded cells* have a lowest height value of zero, indicating contact with the ground.
  - *Direct void cells* contain elevated cells not in contact with the ground.
  - *Both-type cells* are grounded but contain internal voids.

  Figure 4.16 shows examples of these three classifications.
- **Step 3: Assigning DSM layers** The first DSM layer is filled for grounded cells using the minimum of the top surface and secondary surface heights.
- For one void:
  - The second layer is filled with the internal void height for elevated and both-type cells.
  - The third layer restores the top surface for these cells.
- For multiple voids:
  - Alternating layers are filled with void and structure values from the input arrays, as long as valid data exists.
  - If a value is missing in a void layer, the top surface height is used as fallback, and this cell will be removed from the mask.
  - Both-type cells are handled similarly.

- **Step 4: Post-processing** After assigning layers, a post-processing step ensures no values are floating above empty cells:
  - Starting from the fourth layer, the algorithm checks each layer to see if it contains values above an empty cell.
  - If this occurs, the height is shifted down to fill the gap below.
  - This process repeats iteratively until no floating values remain.
  - The first three layers are exempt from this correction to preserve intentional layering near the ground.
- **Output:** The final result is a consistent DSM stack that alternates between solid surfaces and internal voids, suitable for use as input in SOLFD.

To further clarify how the method works, Figure 4.17 compares the outputs for 0-, 1-, and 2-void configurations. The left column shows the 3D overview of the raw height dataset. Each resulting DSM layer is annotated with red labels indicating which raw dataset layer its values were derived from.

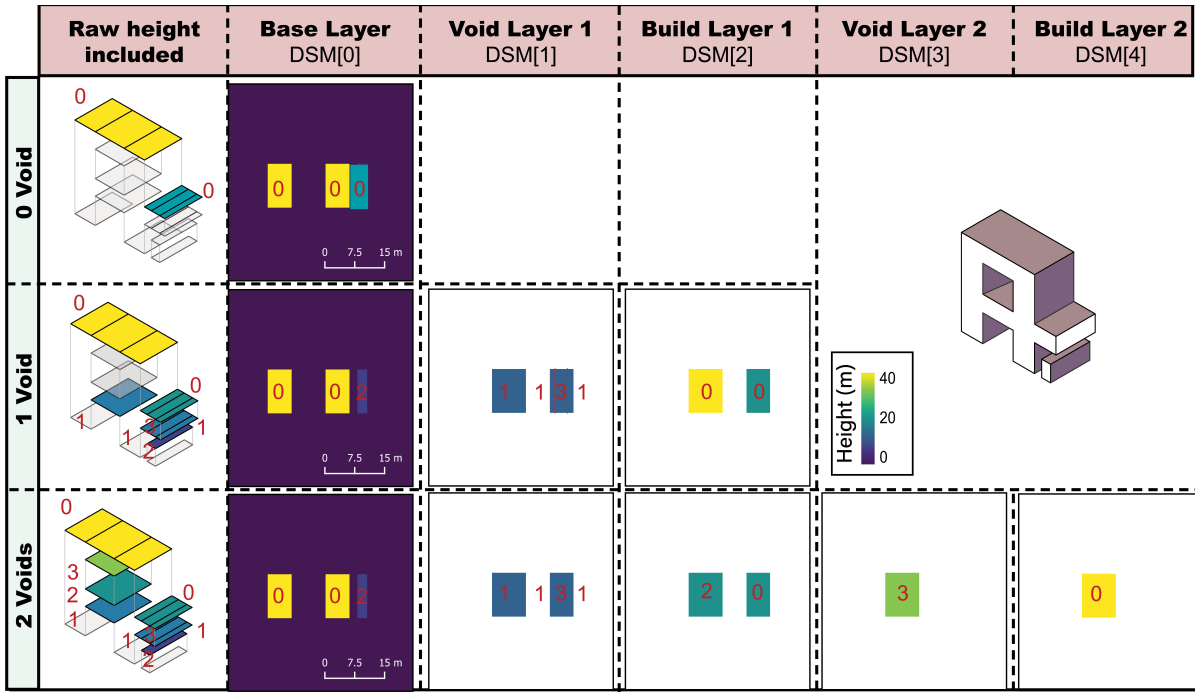

Figure 4.17: Illustration of DSM generation from raw sorted height layers for 0, 1, and 2 voids. Left: 3D representation of the raw height dataset. Right: resulting DSM layers with red labels indicating the source layer from the raw dataset.

Urban Design Integration for Materials

To change surface materials in the urban scene, the land cover layer is updated. This process corresponds to parts of steps 1–4 of the 3D urban design integration workflow (Section **??**. The new surface designs should be placed within the same context scene but represented as 2D flat surfaces instead of extruded 3D objects. These surfaces also need to be exported as a triangulated mesh in the .obj format.

Step 5 proceeds similarly to the 3D urban design integration workflow and also utilizes the algorithm that converts 3D models into 2D layered arrays. However, instead of producing a height map, this version returns an array with 0 values indicating the locations where triangles were inserted—defining where the new surface should appear in the land cover raster. The rest of the array is filled with values indicating no data, denoting areas not intersected by the mesh. This result is then converted into a binary mask.

When inserting .obj files for land cover modification, a land cover class must be specified. This class determines the type of surface material to apply. The existing land cover raster is then updated by

overwriting the masked areas with the given class value. Multiple new surfaces sharing the same land cover class can be grouped in a single .obj file. However, surfaces belonging to different classes must be inserted separately, each with its own class assignment.

4.1.4. Tailoring SOLWEIG Outputs for Design and Analysis

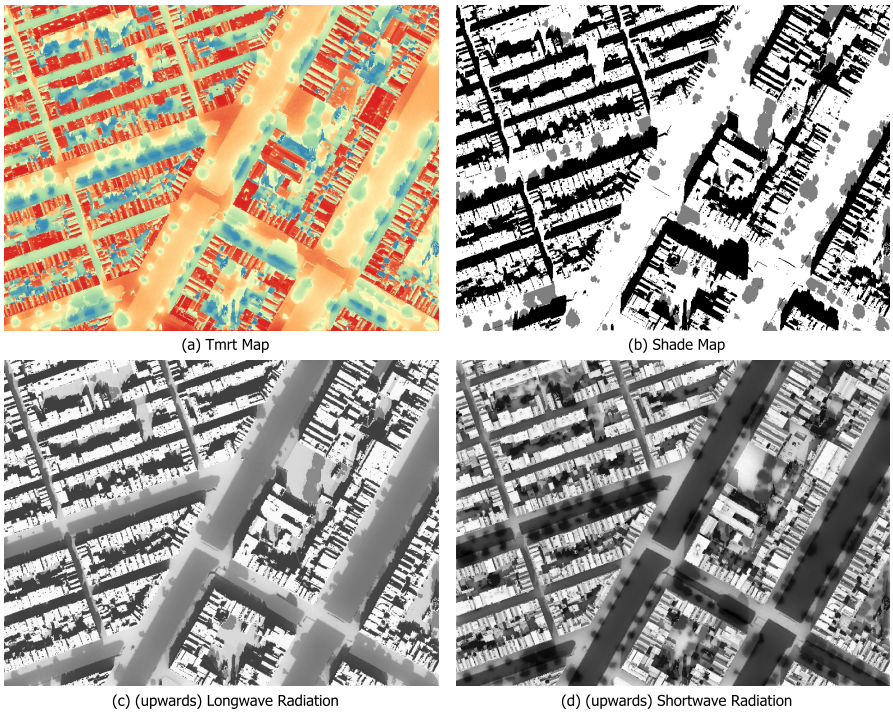

Figure 4.18: SOLWEIG output: (a) Tmrt in °C, (b) shadow maps indicating shaded areas, (c) upward and downward longwave and shortwave radiation maps in Wm^{-2}.

The output of SOLWEIG currently includes maps of T_{mrt}, shade maps, longwave, and shortwave radiation fluxes for each time step provided in the meteorological input file, as well as an overall average map (Figure 4.18). To enable comparison between design options and to identify problematic areas, these outputs require further processing.

In SOLFD, hourly T_{mrt} maps are provided. To support evaluation of thermal conditions over broader parts of the day, these maps are also aggregated into three time periods: morning (06:00–11:00), afternoon (12:00–17:00), and evening (18:00–21:00). For each of these maps, a statistical summary is generated, reporting the minimum, maximum, median, and mean T_{mrt} values. Additionally, area distributions across defined temperature bins are calculated, showing both the absolute area and the relative percentage of land surface that falls within specific T_{mrt} ranges. This enables a direct comparison between different design solutions.

Expanding the evaluation metrics to include basic statistical values, as recommended by Schneider et al., 2023, ensures spatial heterogeneity is not overlooked, which is important for urban planning.

To ensure that the statistics reflect conditions relevant to pedestrian comfort, rooftop and water surfaces are excluded by setting them to no-data. Water surfaces typically exhibit lower T_{mrt} values but are not representative of areas accessible to pedestrians, while rooftops often show high T_{mrt} values that can skew overall statistics. Moreover, SOLWEIG is explicitly designed to estimate T_{mrt} at pedestrian height, reinforcing the focus on walkable outdoor spaces in the analysis.

To fully understand OTC for pedestrians, it is important to recognize that T_{mrt} alone does not determine thermal stress. As discussed in Section 3.1, the impact of T_{mrt} on human heat (and cold) stress is influenced by additional factors such as v_a, RH, and T_a. Since T_{mrt} represents the uniform temperature of

an imaginary enclosure of a black body emitter that would result in the same net radiant heat exchange as the actual environment, it is difficult to interpret what specific T_{mrt} values "feel like" to a human. In general, reducing T_{mrt} helps mitigate heat stress, while increasing it can reduce cold stress. To help quantify and communicate how T_{mrt} translates to actual human thermal sensation and to highlight areas of potential thermal stress, maps of PET will be created alongside the T_{mrt} maps. Table 4.2 presents the relationship between PET values, thermal sensitivity levels, and grades of physiological stress.

Table 4.2: Thermal sensitivity and physiological stress levels based on PET thresholds Source (Matzarakis and Mayer, 1996, p. 8).

PET (°C)	Thermal Sensitivity	Grade of Physiological Stress
<4	very cold	extreme cold stress
4–8	cold	strong cold stress
8–13	cool	moderate cold stress
13–18	slightly cool	slight cold stress
18–23	*neutral (comfortable)*	no thermal stress
23–29	slightly warm	slight heat stress
29–35	warm	moderate heat stress
35–41	hot	strong heat stress
>41	very hot	extreme heat stress

SOLWEIG already offers the capability to calculate PET, using the inputs v_a, RH, T_a, T_{mrt}, and various body model factors (age, weight, height, activity, sex, clothing insulation). However, the current calculation method iterates cell-by-cell, which becomes computationally expensive and it is difficult to vectorize due to complex convergence logic.

To address this, first a direct translation of the existing PET calculation code to C using Cython wily be implemented, to achieve a speedup by leveraging lower-level language efficiency. For further performance improvements and usability, a precomputed lookup table will be created. Here, PET values will be pre-calculated for combinations of input parameters. During runtime, user-specified inputs (v_a and body type) and climatic data will be used to select the closest match from the lookup table by rounding to the nearest available value.

For the creation of the lookup table, selected wind speeds (v_a) were chosen to represent typical conditions: windstill, light breeze, and moderate breeze. A series of systematic tests was conducted to determine appropriate parameter ranges and step sizes, ensuring that interpolation or rounding-induced errors in PET remain below 0.5°C. The final values included in the lookup table are listed in Table 4.3.

Table 4.3: Included values and parameter ranges for PET lookup table.

Parameter	Included Values
v_a (m/s)	0.1, 2.0, 6.0
RH (%)	100, 90, 80, ..., 10, 0
T_a (°C)	40.0, 39.5, 39.0, ..., 0.5, 0.0
T_{mrt} (°C)	60.0, 59.5, 59.0, ..., 0.5, 0.0

The lookup table covers four representative body profiles: a standard adult male, a standard adult female, a young child, and an elderly individual. This selection enables the evaluation of thermal comfort for both the general population and vulnerable groups in urban design contexts. The standard male profile corresponds to the parameters included in SOLWEIG and is commonly used in other models. The other profiles are derived from this baseline, adjusted based on age-related assumptions about body size, metabolic rate, and activity level. The full set of parameter values is shown in Table 4.4.

Table 4.4: Body model parameters for different demographic profiles used in PET calculation.

Body Type	Mass (kg)	Age (years)	Height (cm)	Activity (W/m²)	Sex	Clothing (clo)
Standard Man	70	35	175	80.0	M	0.9
Standard Woman	60	35	165	80.0	F	0.9
Elderly Woman	55	75	160	60.0	F	1.0
Young Child	20	5	110	90.0	F	0.7

To ensure that the use of a lookup table does not significantly reduce the accuracy of the results, the output from the lookup approach is compared to results from the full PET calculation for validation.

Finally, to facilitate problematic area identification, PET maps will be converted into zone classification maps based on the physiological stress categories defined in Table 4.2. This classification will be generated for each hour of the day. This helps reveal consistent thermal hotspots and enables designers to evaluate how their interventions perform across different parts of the day. Important footnote here is that it's not really a true PET indication, moare a T_{mrt} result translation, as other values are not computed and especially wind speed influences largely and has great local differences.

It is important to note, however, that these classifications do not represent true PET calculations. Rather, they are approximations derived from T_{mrt}-based input, while other variables, most notably wind speed, are not dynamically computed. Given that wind conditions can vary substantially at the local scale, relying on static values introduces inaccuracies in representing actual OTC. As such, the resulting classifications should be interpreted primarily as a means to support and contextualize T_{mrt} results, rather than as precise indicators of physiological thermal stress.

4.1.5. Computation Time Optimization

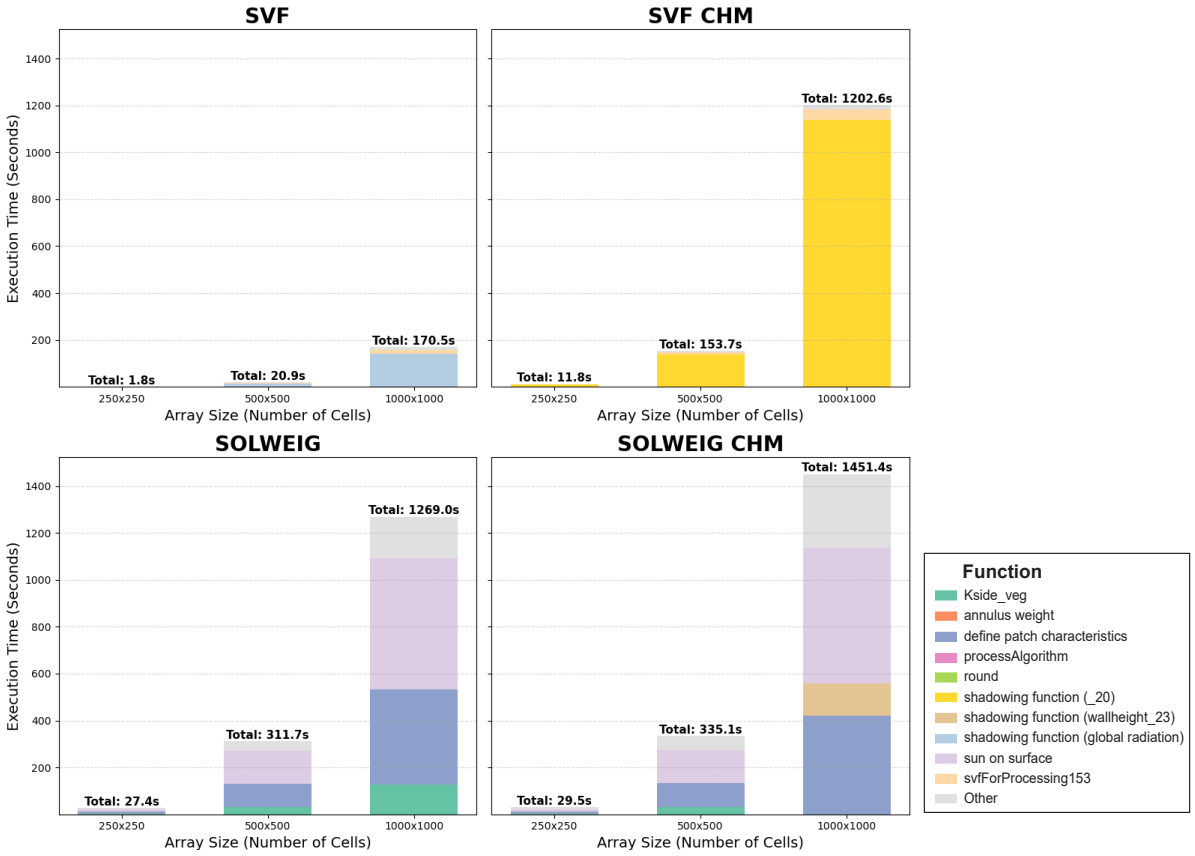

Figure 4.19: Computation time profile for SVF and SOLWEIG calculations across different array sizes, with and without CHM data. The plots highlight the execution time of the top 3 functions and the 'Other' category, along with the total computation time for each scenario.

The two important algorithms in UMEP used to calculate T_{mrt} are the SVF algorithm and SOLWEIG, with SVF serving as input for SOLWEIG. Figure 4.19 shows the computation times for SVF and SOLWEIG across three different array sizes, both modeled with and without CHM data. It reveals that using CHM data significantly slows down the SVF algorithm, and that specific functions contribute disproportionately to the overall runtime. These bottlenecks share a common feature: most of them rely on shade calculation, which is the most time-consuming operation in SOLWEIG, or other iterative functions on arrays. Specifically, "shadowing" and "sun on surface" functions rely on the shade algorithm, while "define patch characteristics" involves array iteration.

To address this, the primary optimization focuses on offloading the shade-casting algorithm to the GPU. This approach is promising because the algorithm is dominated by large array operations, tasks that GPUs are specifically optimized to handle due to their parallel architecture. For example, Li and Wang (2021) demonstrated a speed-up of over 100x in SVF calculation by isolating the relevant computation and implementing it as a custom CUDA kernel.

Most SOLWEIG and SVF functions rely heavily on NumPy for array operations. This can be substituted with CuPy, a GPU-accelerated library that offers a drop-in replacement for NumPy (Okuta et al., 2017). To fully leverage GPU acceleration, the frequency of data transfers between CPU and GPU must be minimized, as these transfers represent a major bottleneck. Therefore, the current code must be revised so that arrays are initialized on the GPU from the outset, and all subsequent operations are performed in GPU memory. Additionally, computational redundancy and memory usage should be reduced by eliminating unnecessary loops and minimizing the number of intermediate arrays wherever possible.

The goal of these optimizations is to enable multiple simulations at the same location within an acceptable timeframe, providing designers with quick feedback to iterate on design options efficiently. While real-time performance is ideal, a realistic target is to complete simulations within 5–10 minutes.

As shown in Figure 4.19, computation time is directly influenced by array size. In urban design, simulations should ideally cover a neighborhood scale. Van Esch (2015) defines a nominal radius of 300 meters for a neighborhood, corresponding to an area of approximately 600×600 meters. At the default spatial resolution of 0.5 meters (as provided by AHN data), this translates to an array size of 1200×1200 cells. However, a 0.5-meter resolution may not always be strictly necessary. In principle, designers might prefer to start with a coarser resolution to gain a rough indication of building block influence, then refining to finer resolutions as the design process progresses. Therefore, additional experiments are conducted to compare model results at 0.5 m, 1 m, and 2 m resolutions, aiming to determine the impact on accuracy for this urban spatial scale.

4.2. Accuracy Assessment of SOLWEIG and SOLFD

The second main step of the methodology involves assessing the usability of SOLFD by evaluating its accuracy as well as that of the underlying SOLWEIG model. The method for reviewing the accuracy of SOLWEIG is described in Section 4.2.1, while the evaluation of SOLFD based on observational data is described in Section 4.2.2.

4.2.1. Literature Review on Model Accuracy

Multiple studies have assessed the accuracy performance of SOLWEIG by comparing its outputs with measured data or with results from other microclimate models. In this review, an overview of papers evaluating SOLWEIG are provided, highlighting its strengths, limitations, and accuracy metrics.

The reviewed literature is categorized into two groups: (1) validation studies conducted by the model developers themselves, and (2) independent research by external groups. A primary criterion for inclusion was the use of real-world sensor data for model validation. Most studies report model accuracy using RMSE and coefficient of determination (R^2). These metrics serve as the basis for assessing SOLWEIG's predictive performance.

4.2.2. Accuracy Assessment Using Observations

To assess the accuracy of SOLFD, data collected using the climate bike will be used. The climate bike is a cargo bicycle equipped with instruments for measuring micrometeorological conditions (Heusinkveld et al., 2013). The dataset includes measurements of the six shortwave and longwave radiation fluxes, recorded along two cycling routes in Amsterdam on August 23 and September 12, 2023. These data will serve as reference measurements for validating the model's output.

Preparing SOLWEIG and SOLFD input

For comparison between SOLWEIG and SOLFD, simulations will be run for six locations featuring 3D structures such as roads and buildings with underpasses. Three distinct input scenarios are prepared and tested:

- **3D Case**: Input of the new model (SOLFD); measurements are always at pedestrian level and may have structures overhead. (Figure 4.20a)
- **Gap Case**: Input for the current model (SOLWEIG); measurements are from the pedestrian level in open spaces without overhead structures (Figure 4.20b).
- **On-top Case**: Also input for the current model; measurements are from on top of structures such as roofs or bridges (Figure 4.20c).

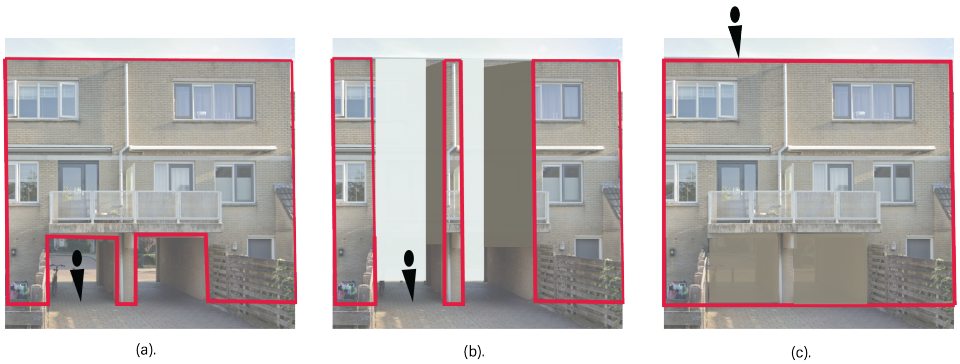

Figure 4.20: Visualization of the three different measurement cases. The location of each measurement point is depicted in the figures: (a) 3D case, (b) gap case, (c) on-top case.
Background image source: (Google, 2022).

Meteorological input data for the two measurement days are sourced from Amsterdam Atmospheric Monitoring Supersite (AAMS), a network comprising 30 weather stations, the KNMI station at Schiphol and Climate Bike measurements. The simulated radiation fluxes are compared against the observed measurements for validation. To achieve this, the steps are as follows:

Identify six locations with underpasses
First, the two climate bike routes are analyzed to identify candidate locations featuring overhead structures. Six locations are selected where the bike passes underneath buildings or bridges. These locations are visualized in Figure 4.21.

Figure 4.21: The six selected locations, shown with bounding boxes in the Amersfoort RD New coordinate system. All overpasses traversed by the climate bike are indicated.

Prepare the six locations for simulation
Second, SOLFD is used to collect the required DSM, CHM, and land cover data. The 3D input DSMs must be created manually, as the tool does not support automatic generation from existing datasets. In theory, a 3D reconstruction of roads could be incorporated by applying the method proposed by Oude Elberink and Vosselman (2009), which uses TOP10NL and AHN point clouds to reconstruct 3D road junctions.

The manual creation of the 3D data is performed using the 'overbruggingsdeel' features from the BGT dataset, which are further described in Section 4.1.2. These features are used to identify the start and end locations of bridge gaps (applied only to bridges, not buildings). Figure 4.22a shows an example of an 'overbruggingsdeel' feature highlighted in pink. It also illustrates the automated DSMs created using the method from Section 4.1.2, one with and one without using the 'overbrugginsdeel' to get heights from the DSM on the DTM. Satellite imagery was additionally used to confirm the geometry and location of these structures.

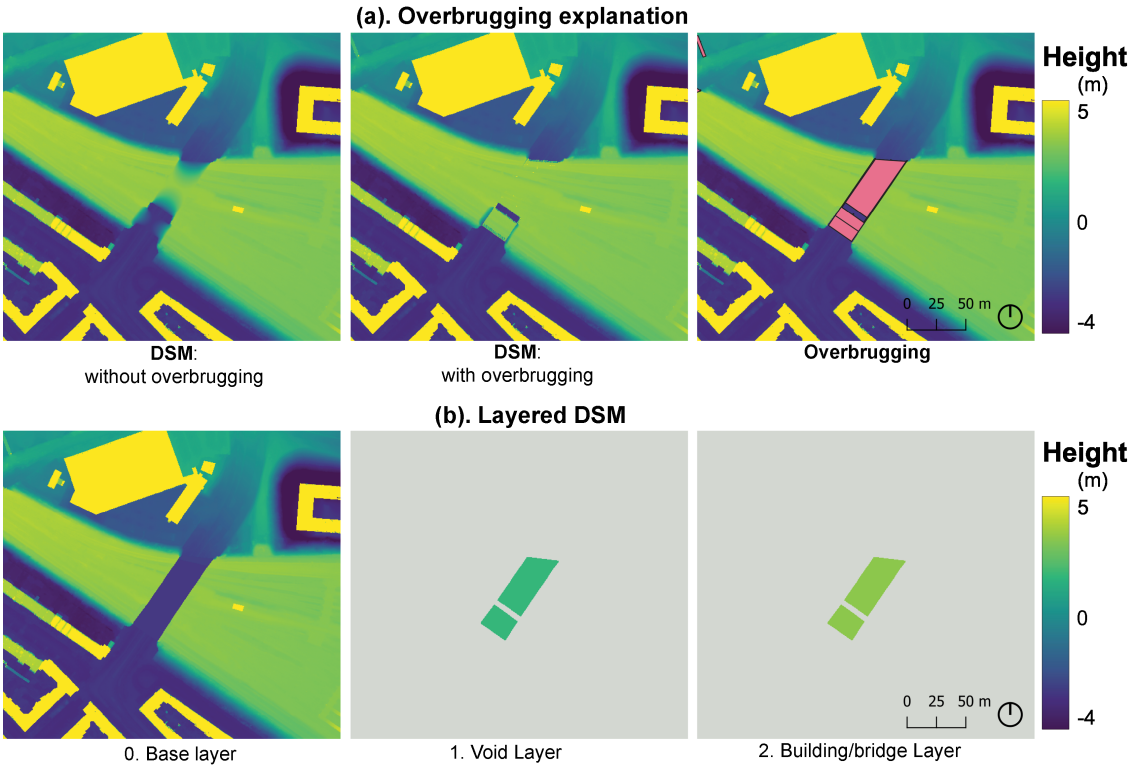

Figure 4.22: Creation of the location datasets. (a) 'Overbruggingsdeel' feature from the BGT dataset (pink) and DSM output with and without its use. (b) Manual creation of a layered DSM for a bridge location: (0) base layer, (1) void layer, and (2) top bridge or building layer.

At the bridge and building locations, a layered DSM is manually created, as shown in Figure 4.22b. The base layer (b.0) was created by replacing the bridge or building surface with the average ground height underneath. For the void layer (b.1), a vertical offset of approximately 4.8 meters was added to the ground height, based on clearance signage available at several locations. Finally, the original measured heights of the bridges or buildings were used to define the upper layer (b.2). The same approach was applied to the selected building locations.

Prepare the meteorological input data

The final step for preparing input to SOLWEIG is to assemble the meteorological input data. Specifically, the periods during which the Climate Bike traversed the six evaluation locations were isolated. Half-hourly meteorological data is taken from the start of the day up to the hour of the last recorded measurement. In addition, minute-resolution data is created for the periods where the bike is present in the selected locations.

- **Global shortwave radiation** is taken from minute data measured by AAMS.
- **Air temperature** is taken hourly from AAMS, and for specific minutes at the locations, from the Climate Bike itself.
- **Relative humidity** is taken hourly from KNMI, and for specific minutes at the locations, from the Climate Bike.
- **Atmospheric pressure** is obtained from KNMI.

Climate Bike Data Preparation

(a). Orientation before interpolation (b). Orientation after interpolation

Figure 4.23: Climate Bike direction (a). before interpolation and (b). after interpolation.

To extract useful directional information from the Climate Bike GPS data, each encoded GPS string is parsed into latitude and longitude components. The direction of movement is calculated using the arctangent of east-west (x) and north-south (y) displacements:

$$\theta = \tan^{-1}\left(\frac{x}{y}\right) \tag{4.3}$$

where x and y are derived from the GPS string, accounting for sign (E/W and N/S). When the bike is stationary (no positional displacement), the last valid direction is reused under the assumption that the orientation remains constant. The resulting orientation of climate bike shown in Figure 4.23

To handle GPS noise and inconsistencies, all direction values are smoothed using linear interpolation. This is done by fitting a line to the sine and cosine components of the direction angles from five neighboring points before and after each target point, then recomputing the interpolated angle:

$$\theta_{\text{interp}} = \tan^{-1}\left(\frac{\sum \sin(\theta_i)}{\sum \cos(\theta_i)}\right) \tag{4.4}$$

This approach maintains angular continuity while avoiding interpolation artifacts near 0°/360°. The resulting orientation of climate bike is shown in Figure 4.23b. Compared to before see orientation stays more consistent.

Subsequently, it is checked whether the resulting direction lies within ±5° of the cardinal directions (0°, 90°, 180°, 270°, 360°), in order to identify points that are aligned with the orientations used in SOLWEIG, which simulates directional radiation based on the four cardinal directions.

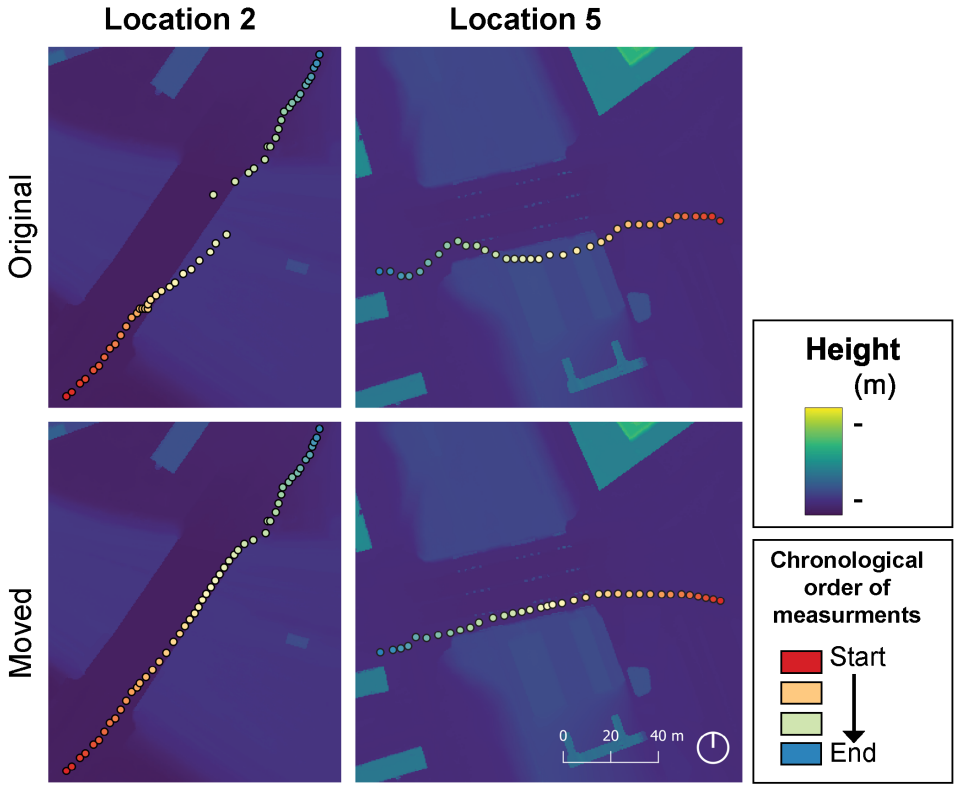

Figure 4.24: Correction of GPS position drift at two locations using manual adjustments. The points are plotted on the DSM, with color indicating chronological measurement order.

Because GPS measurements are often affected by drift and signal loss, particularly in areas with obstructed sky view or under bridges, the recorded positions can be inaccurate. To correct for this, the positions are manually adjusted using satellite imagery and maps, aligning them to the most probable road paths. This process is illustrated in Figure 4.24, which shows two examples of such corrections. The chronological order of the measurements is color-coded, and the adjusted spacing between points is based on estimated cycling speed derived from unaffected segments. Despite these corrections, the precise original locations remain uncertain, introducing potential errors when comparing the data to model output.

Finally, the observed T_{mrt} values are compared with modeled values, which requires calculating T_{mrt} from the radiation sensor readings. While direct comparison of measured and calculated longwave and shortwave radiation fluxes would be possible in principle, it is complicated by differences in orientation: SOLWEIG computes directional radiation assuming fixed cardinal orientations, whereas the Climate Bike's orientation changes continuously during movement.

For the calculation of T_{mrt} from the measured fluxes, the human body as a box approximation (Equation 3.2) will be used. It is not possible to use the cylindrical approximation (Equation 10), as it requires directional fluxes from both east and west to calculate the direct radiation.

4.3. Urban Test Case with SOLFD

SOLFD can be used to test urban areas designed from scratch. However, a significant aspect of climate adaptation for Dutch urban areas involves reorganizing existing spaces. Therefore, this test case focuses on assessing the effectiveness of heat mitigation measures proposed in the literature for existing urban areas.

Shade provision is a primary strategy for reducing urban heat stress. Among mitigation options, landscaping elements, particularly trees, are often the most effective. Trees can be integrated into various public spaces, including streets, plazas, gardens, and parks, offering scalable solutions from local interventions to broader urban applications (Aleksandrowicz, 2022).

To evaluate the cooling effect of existing trees, two scenarios are assessed for their impact on T_{mrt}:

1. *Without Trees:* The urban environment is simulated using only the DSM, omitting all vegetation.
2. *With Existing Trees:* Both the DSM and the CHM are used to include current vegetation.

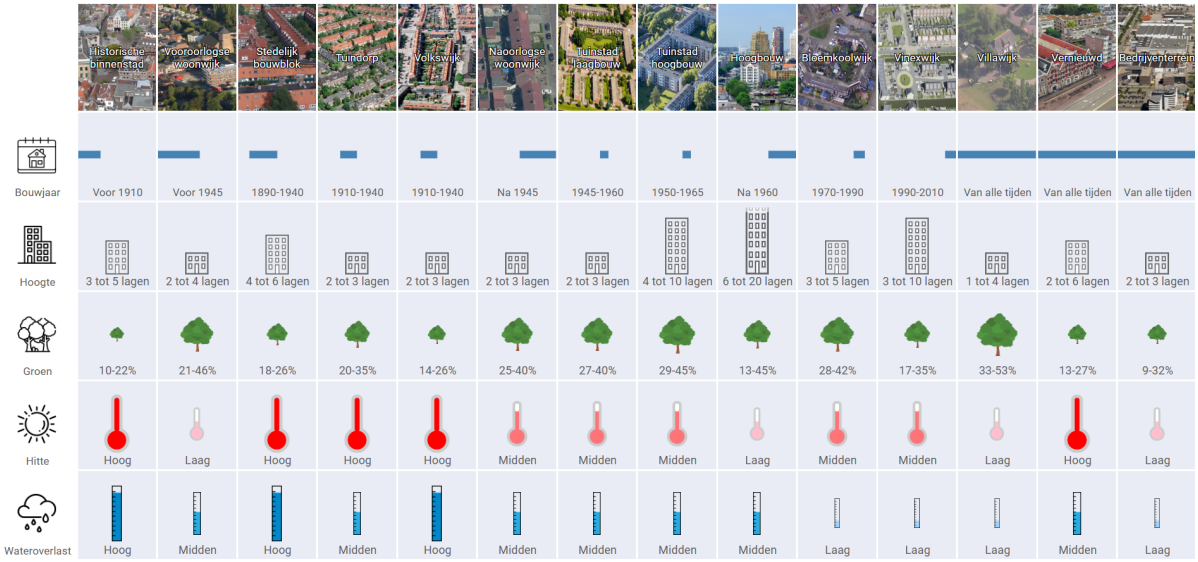

Figure 4.25: The different types of Dutch Neighborhood typologies.
Source:(Klimaateffectatlas, 2023)

Urban study areas are selected based on the neighborhood typologies defined in the Klimaateffectatlas (Figure 4.25), typologies most vulnerable to heat stress are prioritized. The *Wijktypologie PC6 Gebieden 2023* dataset (Kleerekoper et al., 2023) is used for location selection, which contains six-digit postal code areas in the Netherlands classified into the Klimaatatlas typolgies by machine learning trained on features such as building density, green space, population density, and street network morphology (Kluck et al., 2023). An example of the dataset is shown in Figure 4.26.

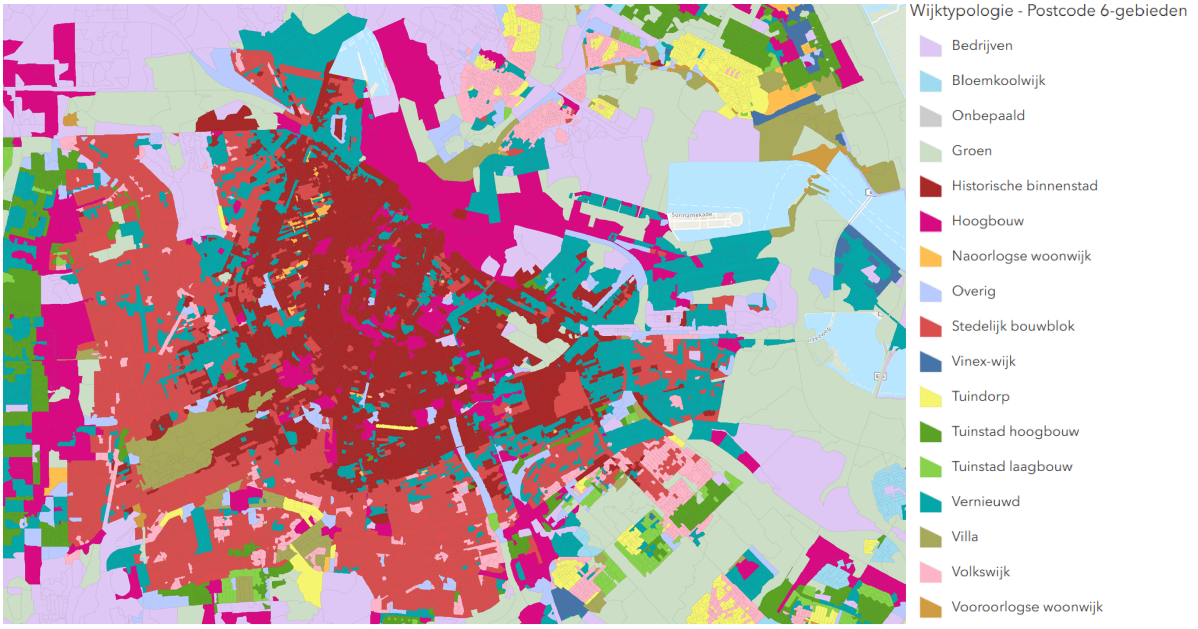

Figure 4.26: *Wijktypologie PC6 Gebieden 2023.* Source: (Kleerekoper et al., 2023)

Typologies included in the analysis are: 'historische binnenstad', 'stedelijke bouwblok', 'tuindorp', 'volkswijk', and initially 'vernieuwd'. However, due to the limited spatial extent of 'vernieuwd' areas in the dataset, 'vinex' and 'bloemkoolwijk' are included as alternatives. A total of 34 areas are selected, 6 for each typology, aiming for geographic diversity and roughly comparable area sizes. The 'stedelijk bouwblok' typology, only present in Amsterdam and Rotterdam, is represented by four areas (two per city). The complete set of study locations, along with their city, array dimensions, and bounding boxes in EPSG:28992, is presented in Table A.1.

For the evaluation of cooling effects, T_{mrt} values on tropical hot days are compared across all included typologies, considering different parts of the day (morning, afternoon, and evening) as well as the overall day. For this, SOLFD is directly applied to automatically generate the necessary input data and T_{mrt} output statistics. Because it only requires a bounding box as input, the case study setup is largely automated. As a framework, SOLFD still requires connecting the individual processing components, but once configured, it enables efficient analysis of a large number of locations.

This case study demonstrates how SOLFD can be used to assess the effects of existing trees in various neighborhood types. Due to the extensive scope of this analysis, no new urban design proposal is included, although incorporating one could further strengthen the conclusions.

As outlined in Section 4.1.4, SOLFD provides statistics of T_{mrt}, including minimum, maximum, and mean values, as well as the distribution of area across T_{mrt} bins. These outputs allow comparison across neighborhood typologies and time periods. In addition to masking out water and building areas during the computation of statistics (implemented in SOLFD), an additional mask is applied to exclude parts of the input area that fall outside neighborhood boundaries. This prevents distortion of results, since neighborhoods often do not perfectly align with the square input extents.

Simulations are performed using meteorological datasets representing both an 'average' summer day and an 'extreme' 30°C+ day. As described in Section 4.1.2, these datasets are integrated into SOLFD. This approach enables focusing on the hottest days to identify extreme temperature conditions.

5

Results

This chapter presents and analyzes the results related to the three main research questions. First, the result of SOLFD is discussed in Section 5.1, followed by the SOLWEIG and SOLFD accuracy assesment in Section 5.2, and finally the test case results in Section 5.3.

5.1. SOLFD Development

The developed code for SOLFD is available on GitHub, along with raw result data: `https://github.com/jsscmnhn/SOLWEIG_SOLFD`.

Unless stated otherwise, all tests are conducted using meteorological data from a summer day in Utrecht, the Netherlands. Specifically, July 1st, 2016, at latitude 52.25° and longitude 4.75°. Both a laptop and a computer were used for testing. Their specifications are detailed in Table 5.1.

Type	CPU	GPU
Laptop (HP ZBook Studio x360 G5)	Intel(R) Core(TM) i7-8750H, 6-Core	NVIDIA Quadro P1000
Computer (Custom Build)	AMD Ryzen 9 7950X, 16-Core	NVIDIA GeForce RTX 4090

Table 5.1: Hardware specifications of the laptop and desktop used.

5.1.1. Extending SOLWEIG for 3D Simulation

The code has been modified and enhanced to support 3D input in the form of a layered array, as described in Section 4.1.1. Generally, these layers are utilized only for shadow calculations; once shadows are computed, the simulation proceeds using just the base layer. As a result, no further updates were required elsewhere in the code.

Layered Array

Figure 5.1: Building geometry (left) and the resulting shadow map produced by the shading algorithm (right).

Figure 5.1 presents an example of a 3D building geometry alongside the corresponding shadow map produced by the shading algorithm employing the layered array method developed in this study. The gap within the structure is accurately represented, and the shadow cast on the upper portion of the building is clearly visible, demonstrating the effectiveness of the algorithm.

Figure 5.2 shows the layered DSM input used to create this example. It depicts the three layers at their actual heights in 3D, alongside the corresponding flat arrays: (0) base layer, (1) void layer, and (2) building layer, which together constitute the layered DSM.

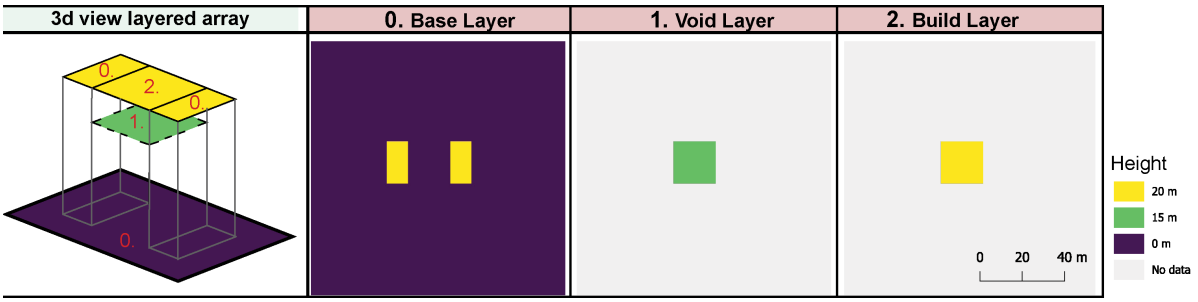

Figure 5.2: Layered input DSM for the example, showing the 3D representation of the layers at their actual heights: (0) base layer, (1) void layer, and (2) building layer.

Optimal Number of Void Layers

The input DSM for the two test cases used to determine the optimal number of void layers are shown in Figures 5.3 and 5.4, respectively. To better illustrate the layered 2D array inputs, a 3D extrusion visualization is provided that depicts the corresponding heights. The model is slightly stretched vertically to better highlight the geometries. Additionally, lines in the 3D view indicate which parts would be filled, either by the base layer (the base volume shown in orange) or by the build layers above the voids (the build volumes shown in grey).

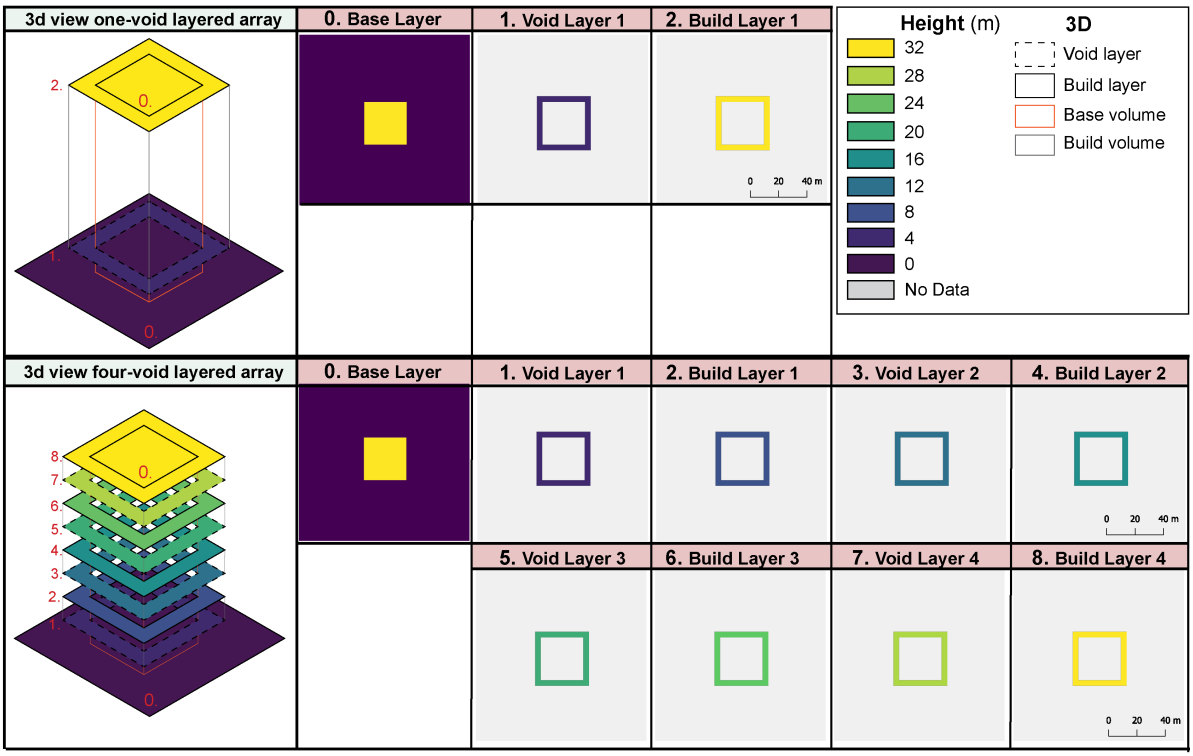

Figure 5.3: Layered input DSM for case 1. The top rows shows the height maps of the 2D layers for 1 void, resulting in 1 void and build layer. The bottom rows shows the height maps for 4 voids, resulting in 4 void and build layers. Labels in the 3D view indicate the DSM layer each height originates from.

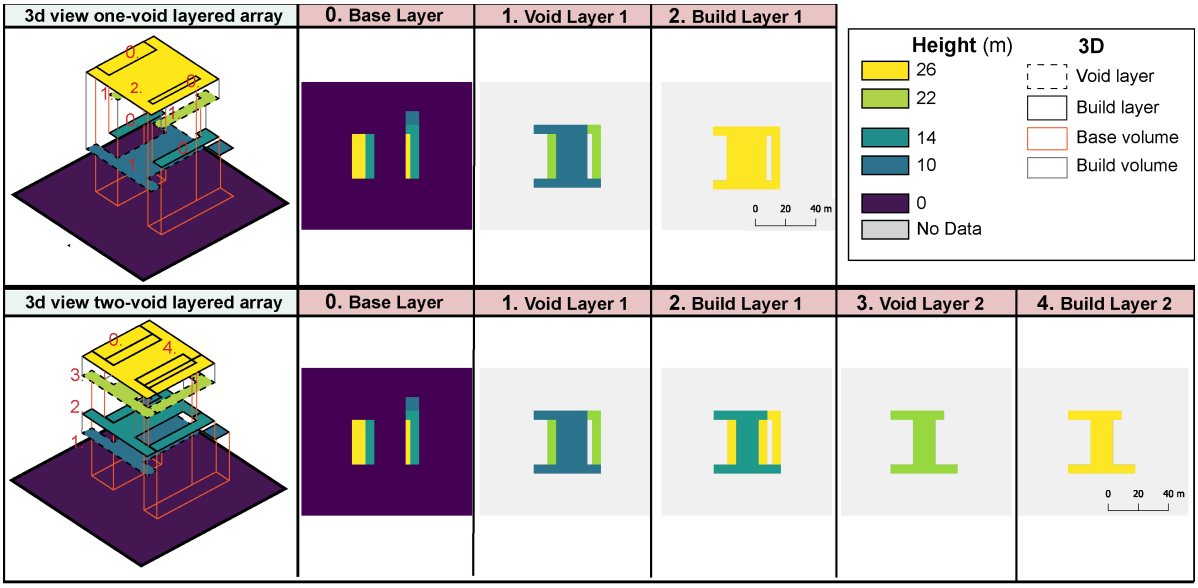

Figure 5.4: Layered input DSM for Case 2. The top rows shows the height maps of the 2D layers for 1 void, resulting in 1 void and build layer. The bottom rows shows the height maps for 2 voids, resulting in 2 void and build layers. Labels in the 3D view indicate the DSM layer each height originates from.

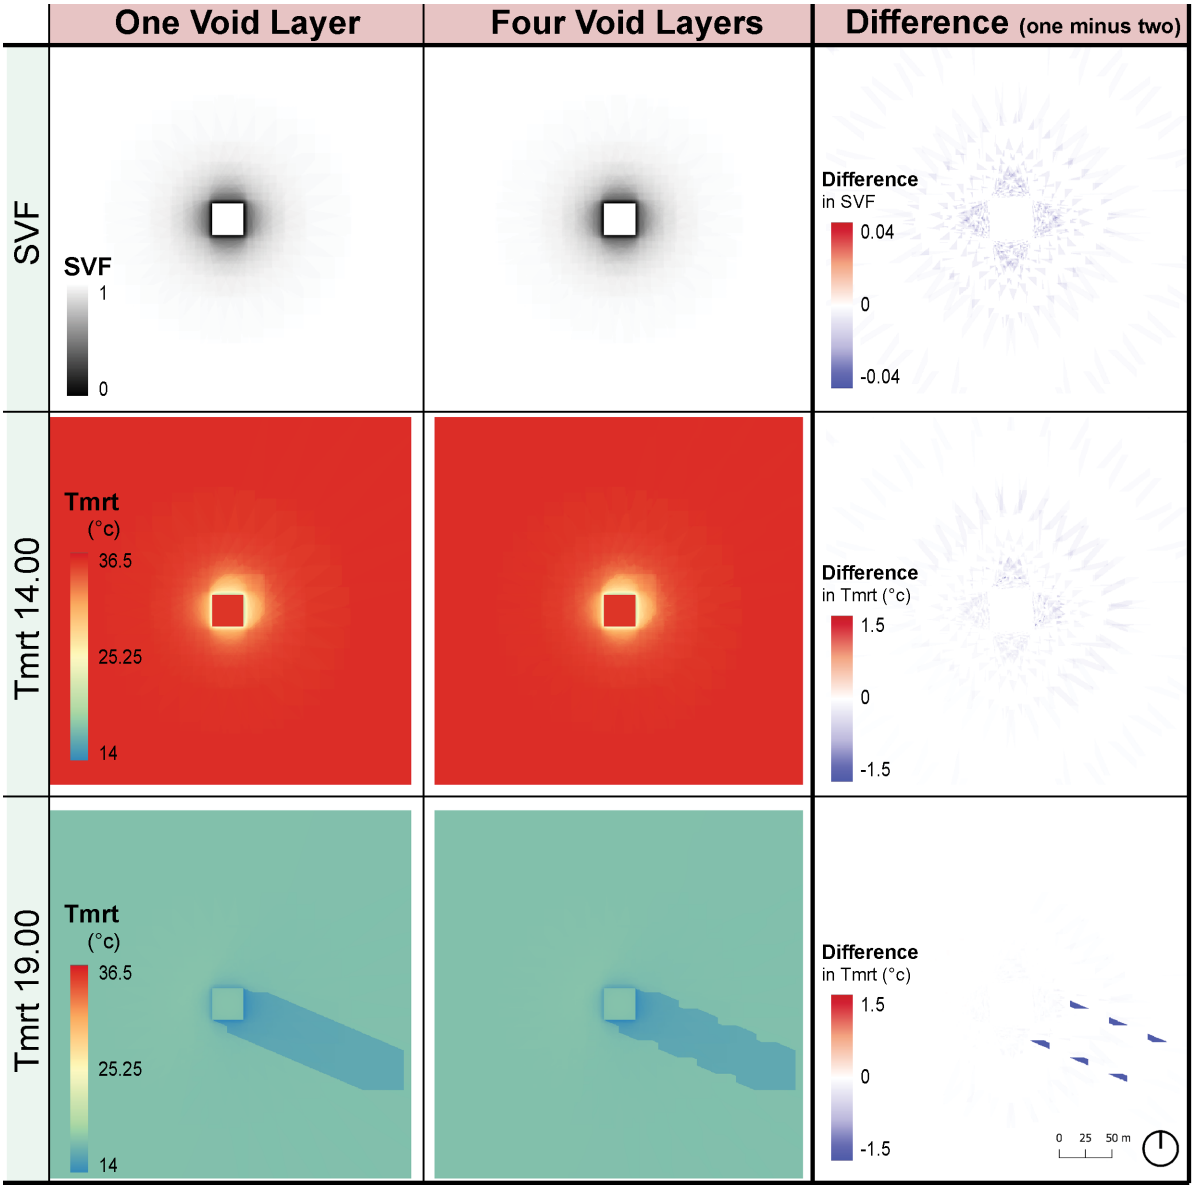

Figure 5.5: SVF and t_{mrt} outputs for the one-void and four-void layer configurations in Case 1, along with the differences between them.

Figure 5.5 presents the resulting SVF and T_{mrt} maps for the first test case. The difference in SVF between the one-void and four-void configurations is relatively small. During peak solar hours (14:00), when shadows are shortest, the differences in T_{mrt} are minimal. The largest discrepancy, up to 1.5 °C, occurs in the late afternoon (19:00), when the sun is lower and shadows are longer. However, in realistic urban settings, buildings are rarely this isolated. In such environments, additional surrounding structures are likely to cast shadows in these areas regardless, thereby diminishing the impact of the discrepancy.

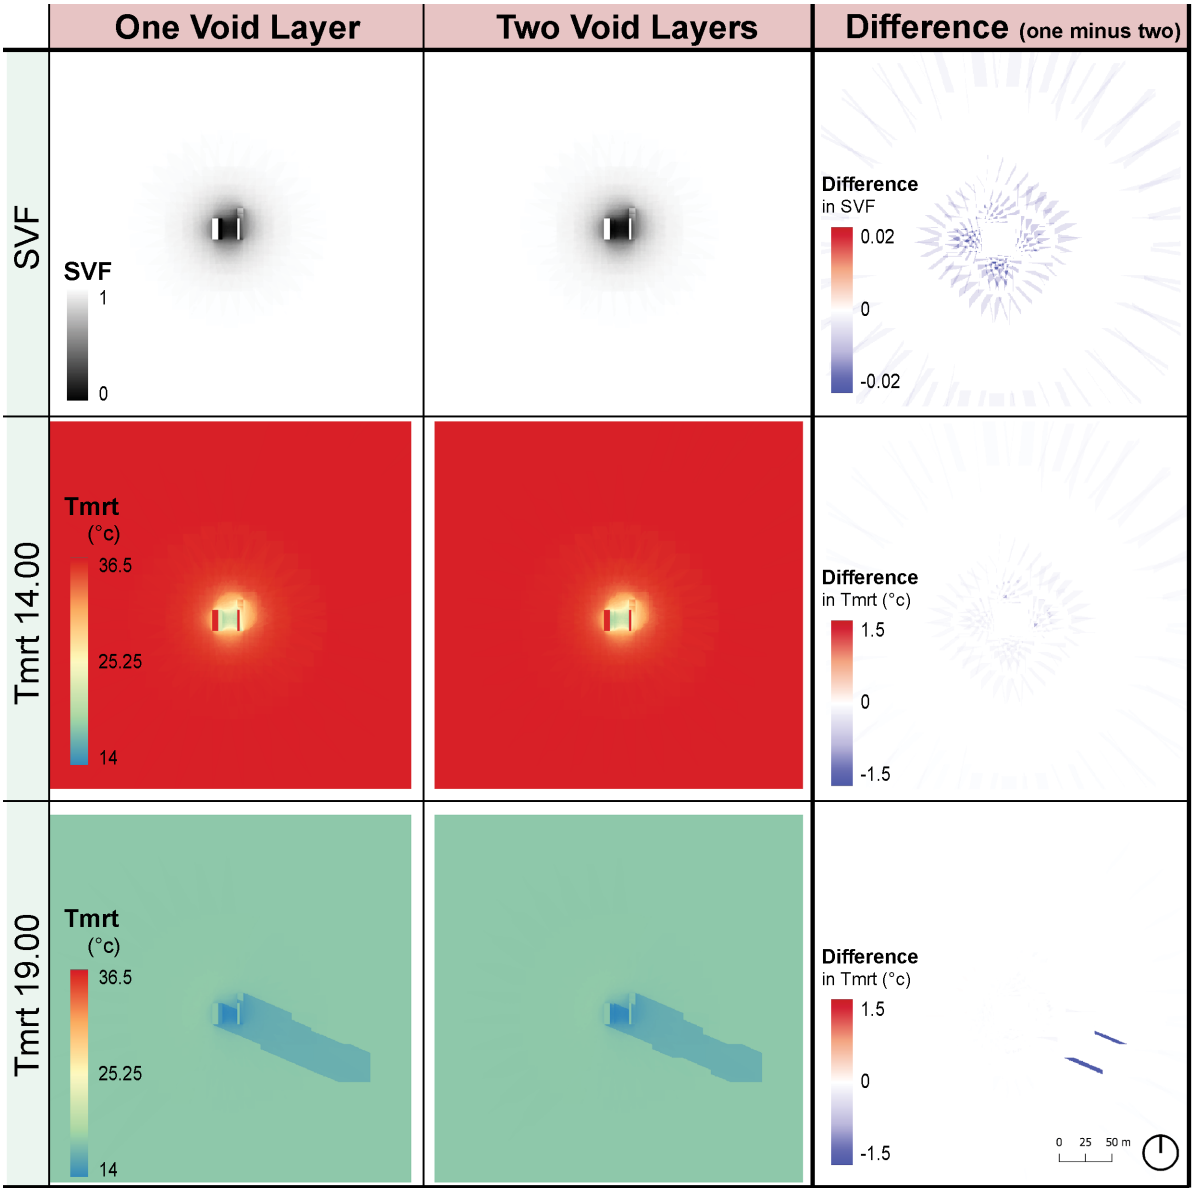

Figure 5.6: SVF and t_{mrt} outputs for the one-void and two-void layer configurations in Case 2, along with the differences between them.

Results for the second case are shown in Figure 5.6. As with the first case, differences in SVF are minimal, especially during the hottest part of the day. The largest deviations again occur in the late afternoon, when the sun is lower and casts longer shadows. Results at non-pedestrian heights have been excluded from this comparison, as these will also be omitted from the final tool output.

Table 5.2: Computation time (in seconds) for SVF and SOLWEIG calculations by void configuration and case. Comparison of SVF and SOLWEIG results by void configuration and case. Array cell size is 700x700, and for the first column Case 1 has 4 voids, Case 2 has 2 voids

Case	4 void / 2 void		1 void		no void	
	SVF	SOLWEIG	SVF	SOLWEIG	SVF	SOLWEIG
1	232.0	133.7	123.0	123.9	19.1	120.2
2	114.1	129.7	70.8	129.5	18.9	120.2

The computation times for both SVF and SOLWEIG for the different void configurations on the laptop are summarized in Table 5.2, alongside results with no void layers for comparison (2.5D case). At this array

size, the computation time for SOLWEIG is relatively unaffected by the number of void layers, as shadow casting occurs less frequently in these test cases. In contrast, SVF calculations show a more pronounced increase in computation time: from approximately 19 seconds without voids to 70–123 seconds with one void, up to 114 seconds with two voids, and nearly 4 minutes with four voids.

In conclusion, for rapid design evaluations, it is recommended to use only one void layer by default, or even none. The lowest (bottom) void layer is chosen as it is closest to pedestrian height and therefore most relevant. Differences in T_{mrt} during peak daytime hours remain small and acceptable, especially given the early design-phase focus of SOLFD. Nonetheless, the option to include multiple void layers remains available for users seeking more detailed modeling.

5.1.2. Automatic Data Collection for SOLWEIG Input Pipeline

Digital Surface Model

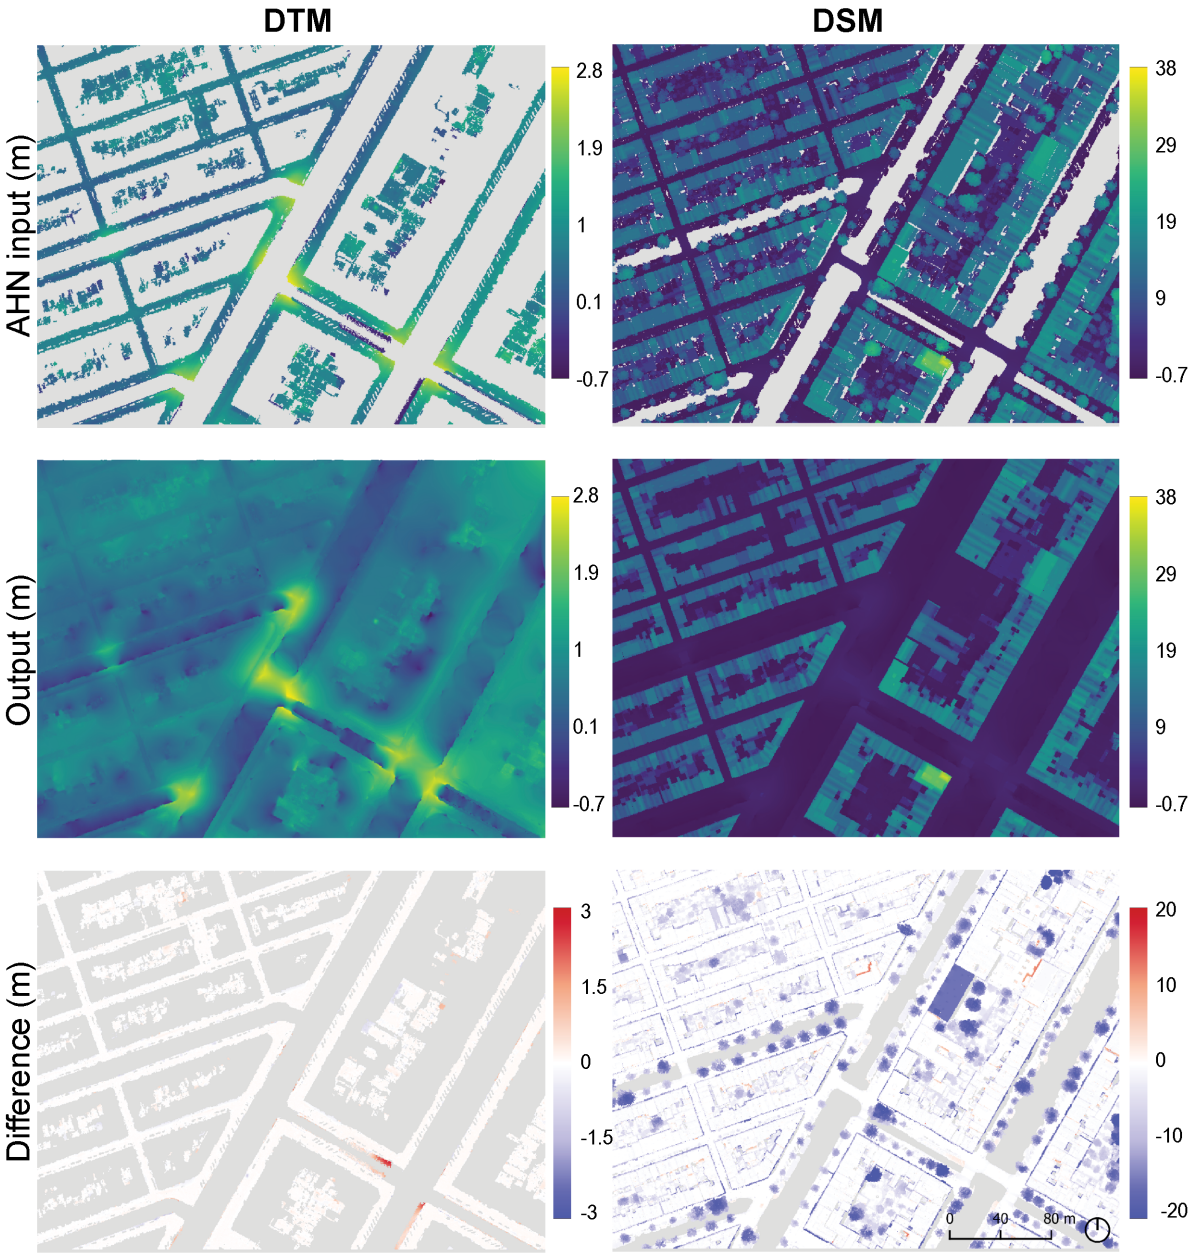

Figure 5.7: Input AHN DTM and DSM with missing data, filled outputs, and corresponding difference maps.

Figure 5.7 shows the input AHN DTM and DSM, alongside the resulting filled outputs and corresponding difference maps. For the DTM, the differences are minimal. Most areas show a slight elevation overestimation, typically between 0.01 and 0.02 meters. The largest discrepancies occur near water bodies, where data points are missing and some values are overestimated. For the DSM, trees were intentionally excluded during preprocessing and are therefore absent in the filled output. Additionally, outlines of missing height values can be seen around most buildings. This occurs because the 3DBAG building footprints are slightly smaller than the actual building extent. Furthermore, some buildings are entirely missing from the 3DBAG dataset, three such buildings are absent in this case.

Overall, the filled DTM and DSM outputs appear accurate, with the exception of a small number of missing or misrepresented buildings.

Canopy Height Model

The resulting CHM is shown in Figure 5.8. Compared to the individual tree structures visible in the input AHN DSM, most trees are successfully detected in the resulting CHM However, the heights appear lower than those in the AHN DSM. This discrepancy is due to differences in interpolation methods: the CHM is generated from raw LAS data using Laplace interpolation, where cell center heights are derived from the three corner points of the intersecting triangle. In contrast, AHN uses squared inverse distance weighting for its DSM, which considers all surrounding points inside the cell to interpolate values at the cell centers.

Additionally, a median filter is applied on the CHM, which suppresses the highest values. While using a maximum filter could restore higher values and match the original heights more closely, it would also exaggerate outliers. At the outer edges, the tree heights are often overestimated compared to AHN DSM. This is caused by the 1-meter buffer applied around the last detected vegetation point, effectively making the trees appear larger than they actually are.

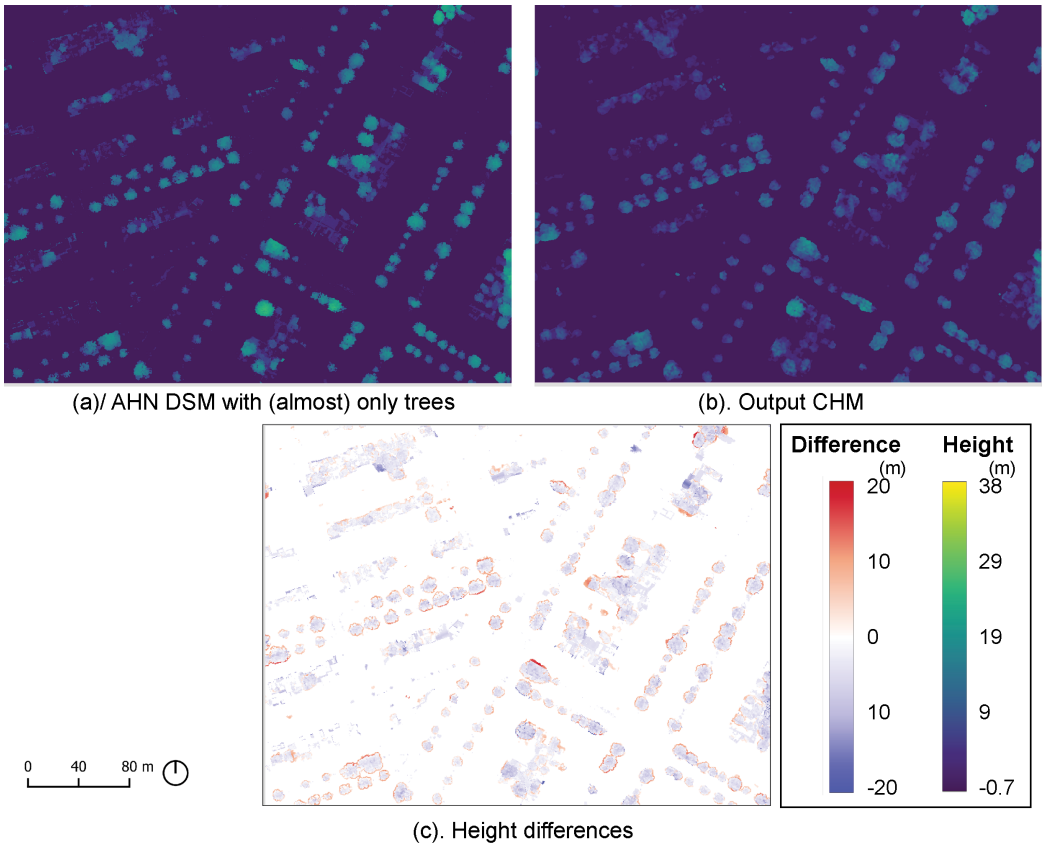

Figure 5.8: Comparison between input and processed data: (a) AHN DSM showing isolated canopy heights, (b) generated CHM based on raw LAS data using Laplace interpolation, and (c) the difference in height values between (a) and (b).

Wall Height and Aspect

Placing the wall height calculation method on the GPU reduced computation time from 6.54 seconds to 0.14 seconds for a 798x598 array. Since the core logic of the method remained unchanged, the output of both the original and new methods is the same (Figure 5.9).

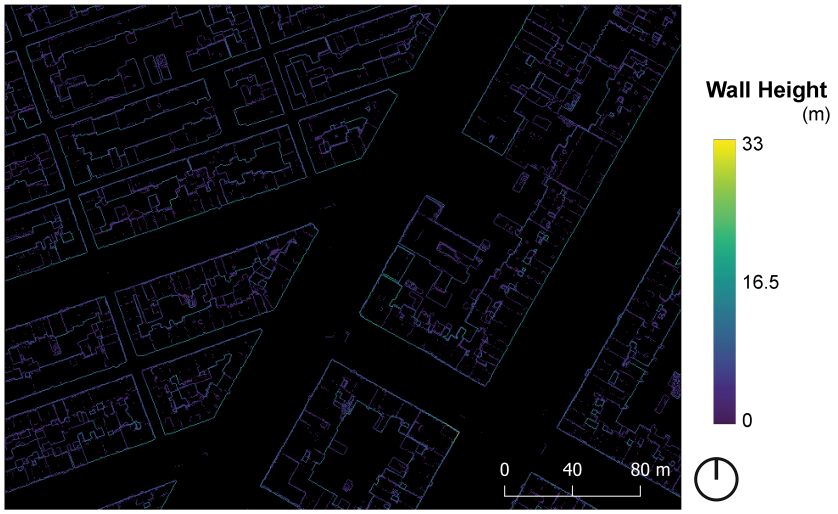

Figure 5.9: The output of both original and new building height method.

The new method for aspect calculation speeds the process up significantly, from 223.89 seconds to 0.02 seconds for a 798x598 array. Although the Sobel filter alters the aspect estimation, the overall orientation maps remain comparable, shown in Figure 5.10.

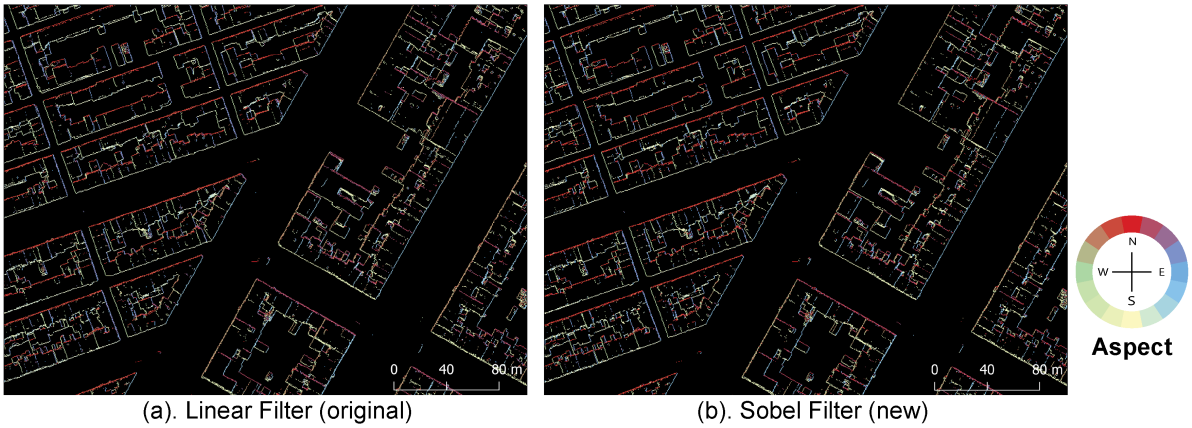

Figure 5.10: Wall aspect outputs using (a) the original 9x9 linear filter and (b) the new 3x3 Sobel filter.

To evaluate the impact of this change on thermal comfort modeling, the T_{mrt} output from SOLWEIG is compared for both methods at 13:00 and 17:00 (Figure 5.11). Differences between the outputs are mostly minimal and spatially limited to wall-adjacent cells.

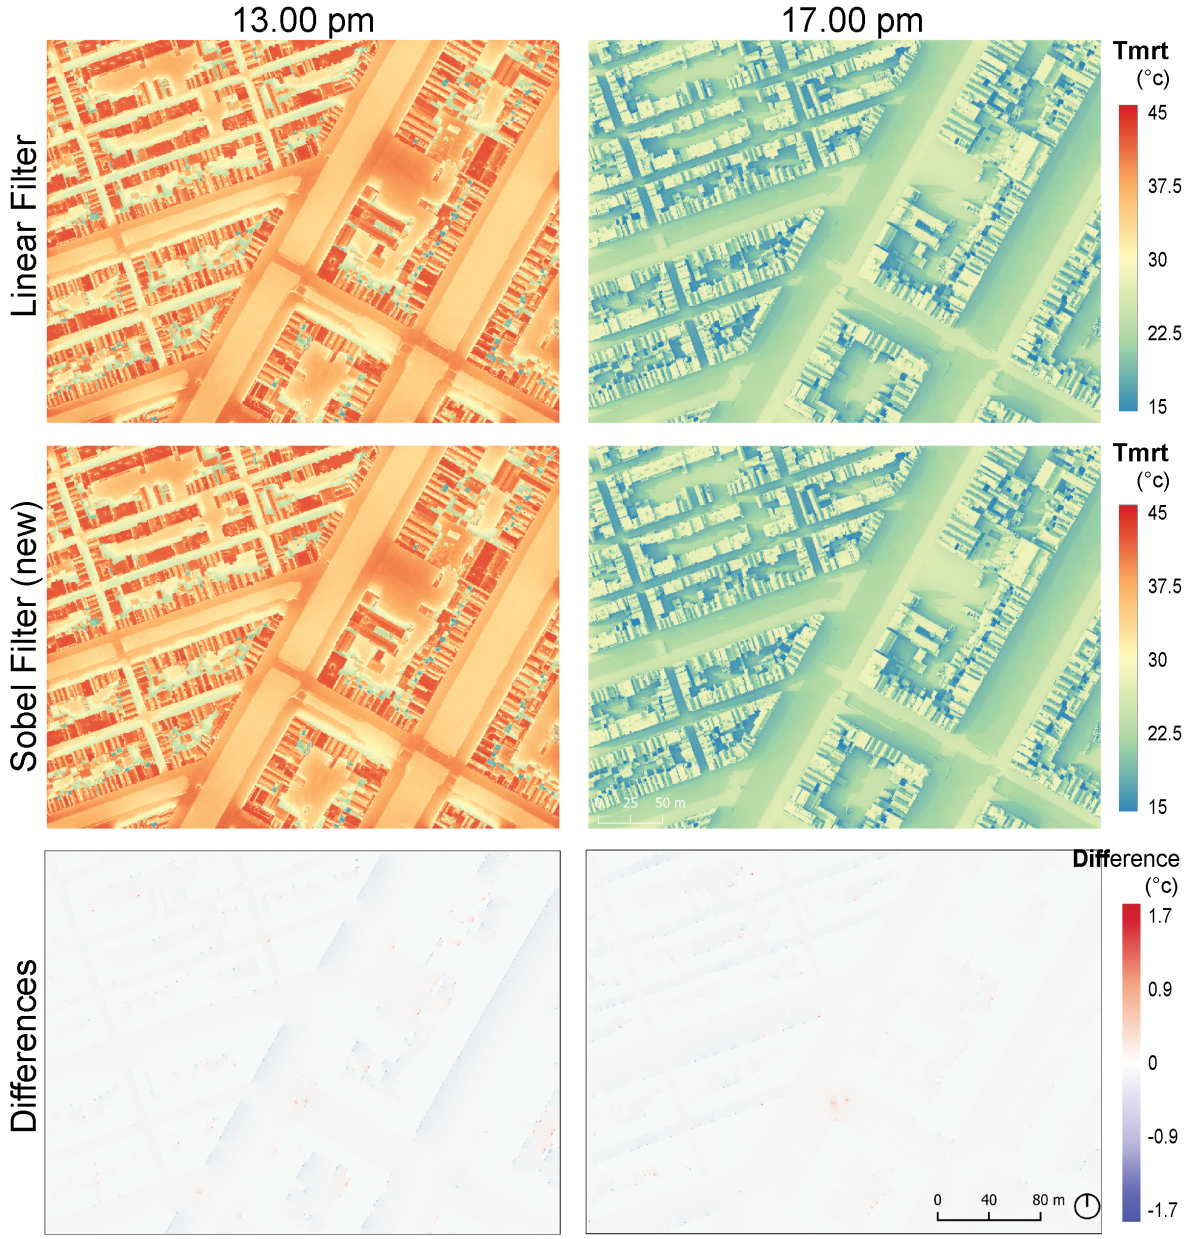

Figure 5.11: Comparison of T_{mrt} results at 13:00 and 17:00 using the liniar filter and Sobel filter wall aspect methods. Differences are shown in the third row.

Table 5.3: T_{mrt} Difference Statistics in °C at 13:00 PM and 17:00 PM

Statistic (°C)	13:00 PM	17:00 PM
Min	-1.76	-1.37
Max	0.45	1.55
Mean	-0.006	-0.003
Stddev	0.044	0.0027

Table 5.3 shows that the maximum T_{mrt} difference is -1.76 degrees. The difference is smaller later in the day, likely due to temperature stabilization as walls cool and other surfaces are shaded. The largest difference occurs directly next to the walls, and rapidly decreases with distance.

The observed differences in T_{mrt} can be explained by the method used for wall aspect calculation.

Wall aspect determines which parts of the wall are shaded or exposed to sunlight, and consequently, affects the temperature. The Sobel method uses a 3x3 filter instead of 9x9. This means the Sobel filter is more localized, which introduces inaccuracies, particularly with long wall lines. For example, a southern-eastern wall now has more southward-facing cells (Figure 5.12a), whereas in the original method, a large southeastern wall would have been taken as a single, long southwest-facing wall. This leads to errors at 13:00 PM, as incorrectly identified south-facing cells are assumed to be in sunlight rather than shaded, resulting in higher temperatures.

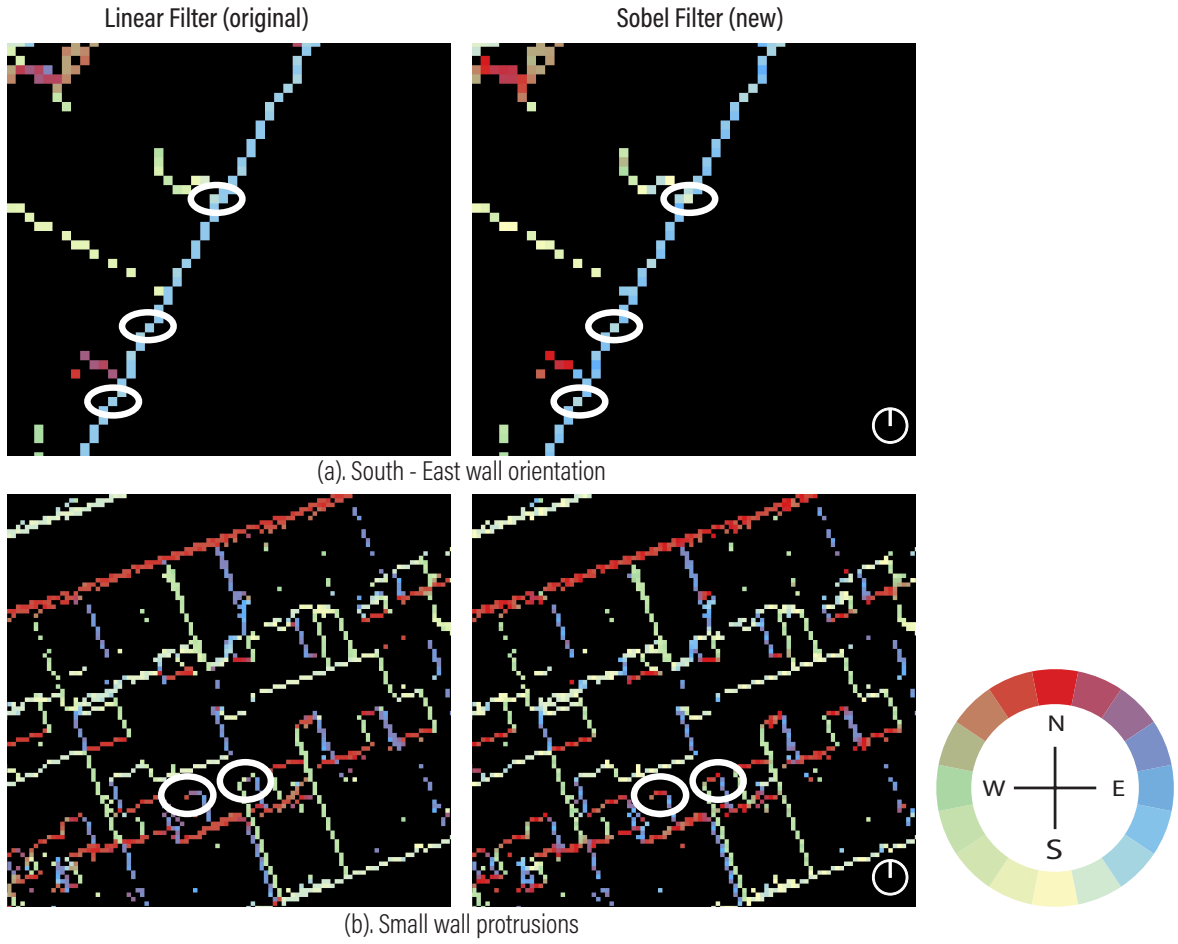

Figure 5.12: Close-up examples showing differences caused by the new Sobel-based wall aspect method. (a) An elongated wall is split into more south-facing cells. (b) A small wall protrusion is now detected as a distinct orientation

On the other hand, the Sobel filter also improves the detection of smaller protrusions, which are now identified as distinct wall orientations rather than being grouped with the larger wall (Figure 5.12b). While this results in a more accurate representation of some smaller features, it also leads to the identification of outliers as separate wall orientations.

Despite these differences, the T_{mrt} variations between the original and new methods are considered acceptable, especially given the substantial reduction in computation time.

Land Cover

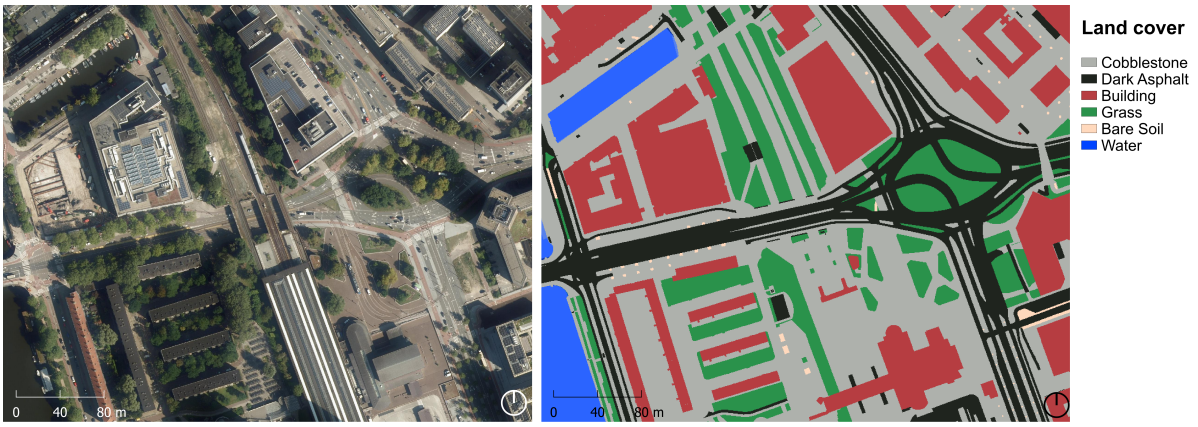

Figure 5.13: Comparison of the land cover classification output (left) with satellite imagery (right).

In Figure 5.13, the output of the land cover classification step is shown. Compared to the satellite imagery, the general shapes are preserved, and distinct patches of vegetation and road segments are clearly identifiable. However, there are some limitations in the BGT dataset that affect the classification results.

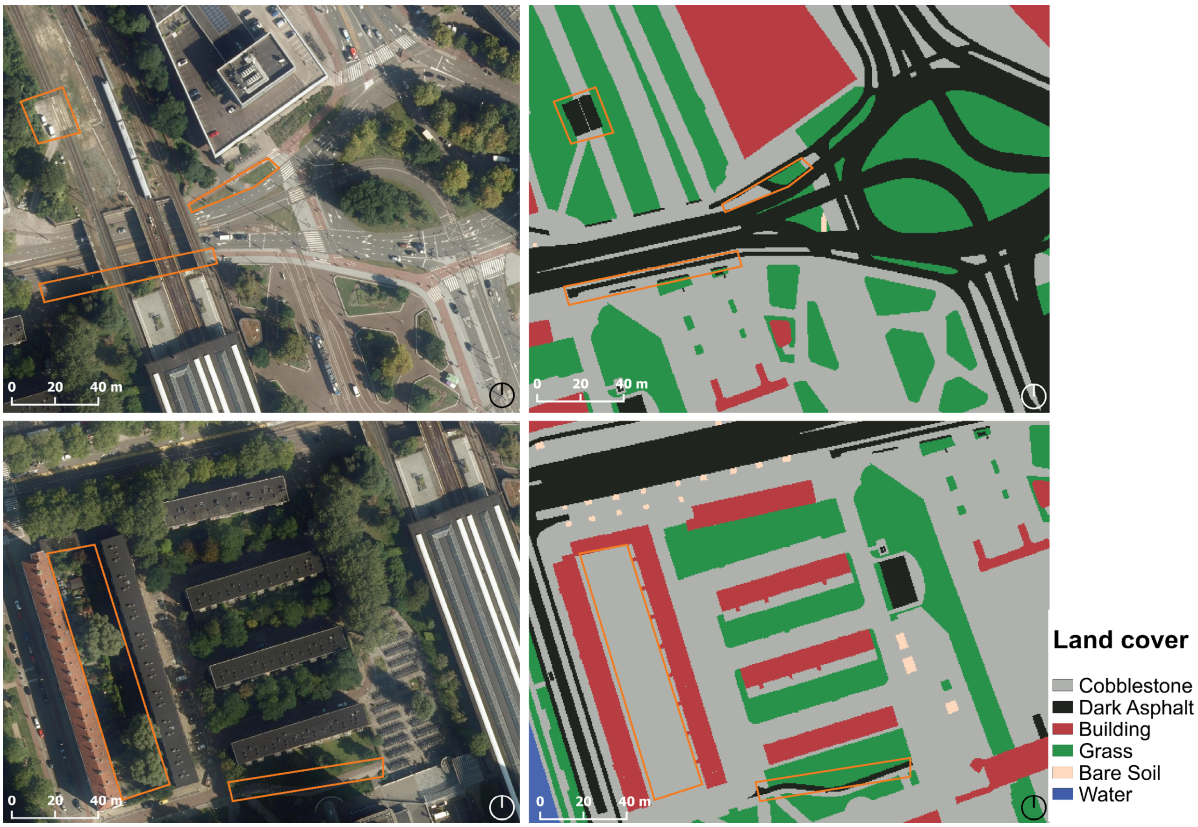

Figure 5.14: Close-up land cover and examples of classification errors. Orange outlines indicate mismatches caused by incorrect *fysiek voorkomen* labels, missing vegetation categories, or oversimplified 'erf' zones in the BGT data, leading to reduced land cover detail.

Figure 5.14 highlights two areas in more detail. The terrain classification is generally accurate, but smaller vegetated patches are sometimes missed For instance, the central roadside greenery in the top image. In other cases, such as in the bottom image on the left, the generalization of 'erf' areas, which are

private grounds surrounding buildings and for which no further information is retrieved from the BGT, leads to a loss of fine details, like individual garden patches (Geonovum, 2020).

Road shapes are mostly accurate and correctly classified, but there are occasional errors in the 'fysiek voorkomen' attribute. Examples include the bike path beneath the railway in the top image, the site above it, and the bike path in the bottom image, which were all given the 'gesloten verharding' classification, even though they are 'open verharding'.

Although the classifications mostly align with the currently supported SOLWEIG types, the limited number of surface material classes necessitates the use of approximations, which introduces generalization. Even if additional material types become available in future updates, such simplifications are likely to remain necessary due to the limited material detail in the source data. Some features do include the 'fysiek voorkomen plus' attribute, which provides more detailed surface descriptions, such as concrete blocks ('betonstenen') or baked bricks ('gebakken klinkers'). If the parametrization for such materials were added to SOLWEIG, classification accuracy could be improved. However, as discussed in the scope in Section 1.2.2, new surface parametrization is not considered within this work.

The classification of buildings and water bodies is generally reliable. However, SOLWEIG currently provides only a single generic option for both building roofs and walls, despite the fact that these elements can be constructed from various materials (e.g., glass, wood, brick). Differentiating between such materials would be possible by adding a material type attribute to the building polygons used during land cover classification, and adding their parametrization.

Finally, all no-data pixels were assigned a default land cover class. Cobblestone was chosen as the default surface type, as missing data most frequently occurred in urban areas, particularly around buildings due to mismatches between 3DBAG building footprints and those in the BGT, as well as along certain road segments. However, this assumption may not be valid in less urbanized areas, where other surface types are more likely. An alternative approach such as nearest-neighbor filling was not used, as it could lead to incorrect assignments in sensitive areas; for example, courtyards might be filled with building pixels. While the current method offers a practical solution, it introduces a level of generalization that could affect land surface modeling accuracy. A more adaptive filling strategy, informed by spatial context, could be considered in future work to improve classification.

Meteorological data

Four standardized climate datasets have been created for use in simulations, offering users the option to select from three typical summer conditions, as well as one extreme heat scenario. Figure 5.15 provides an overview of the value ranges for the relevant meteorological variables across the three typical day categories, with one std (logo thing) outlined showing that general temperature trend followed. Same for 97th percentile. Tables A.5, A.6, A.7 and A.8 show the resulting datasets.

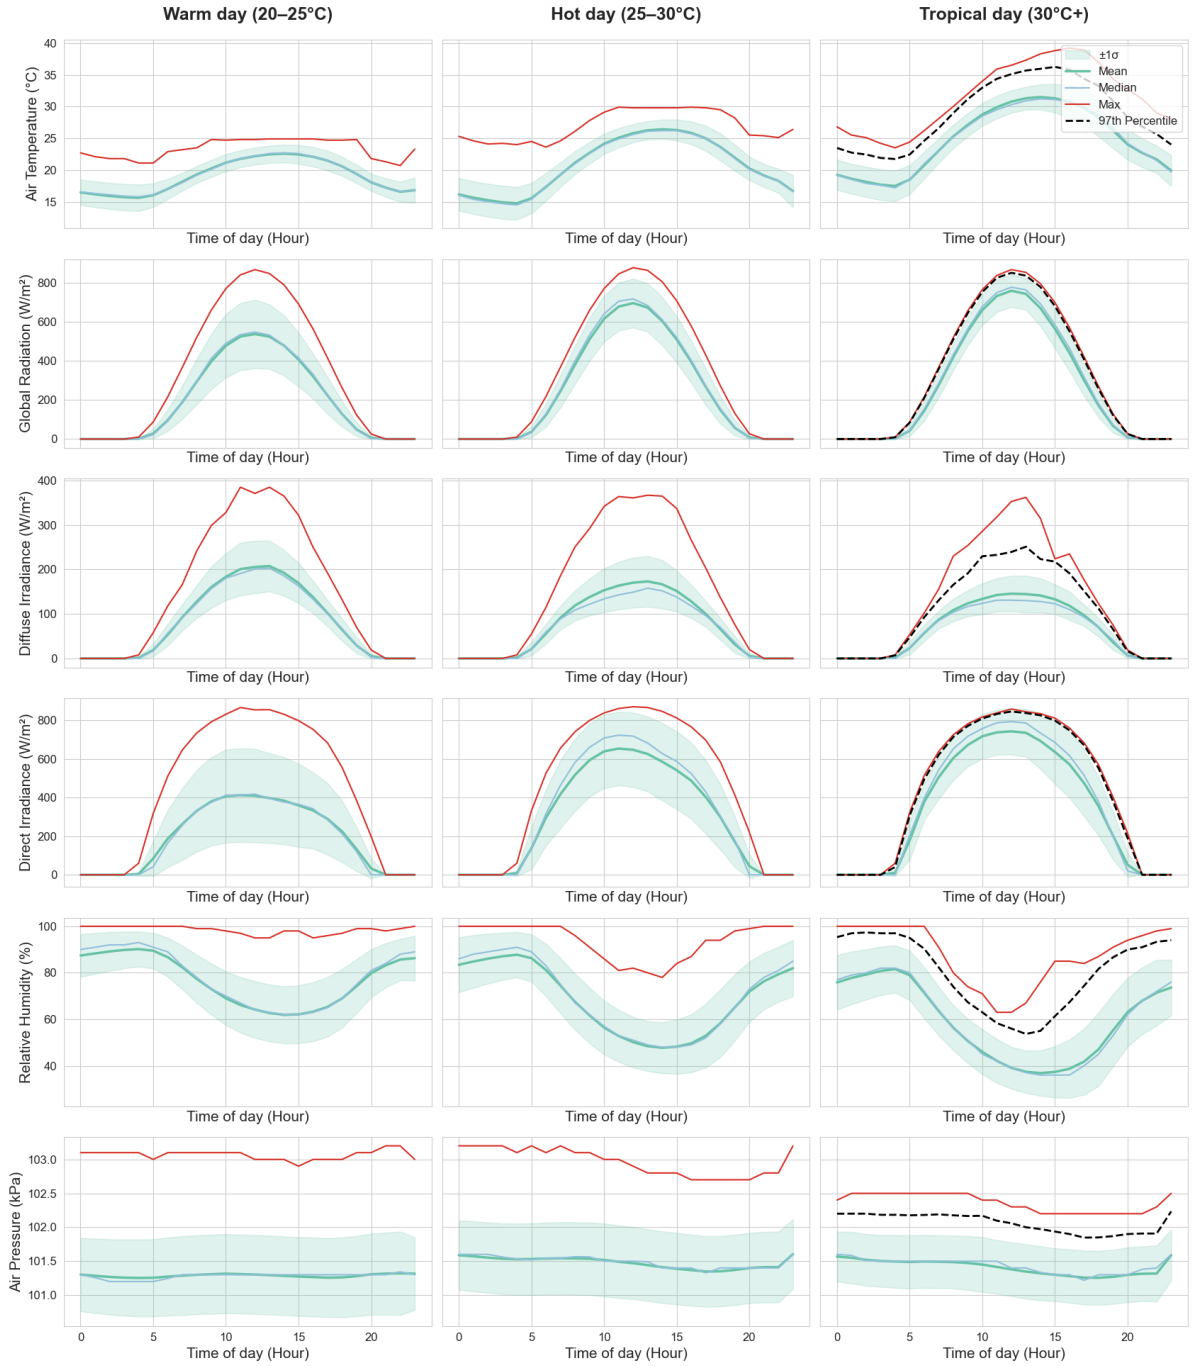

Figure 5.15: Hourly climate variable summaries for three day types: Warm (20–25°C), Hot (25–30°C), and Tropical (30°C+). Each subplot shows the mean (solid green), median (solid blue), maximum (solid red), and ±1 standard deviation (shaded region) values. The 97th percentile (dashed black line) is displayed only for the Tropical day category to highlight extreme conditions during the hottest periods. Variables include air temperature, global radiation, diffuse and direct irradiance, relative humidity, and air pressure.

Choosing mean temperature to represent the different day types is considered a valid approach, as it follows the general trends of all percentiles. The maximum air temperature shows a cutoff effect for the Warm and Hot days because any day exceeding 25°C or 30°C is categorized into the next day type. For the extreme day scenario, the maximum temperature is not used as it tends to overestimate peak heat conditions.

5.1.3. Integrating User Designs and Enabling Modifications

Modifying Existing Elements

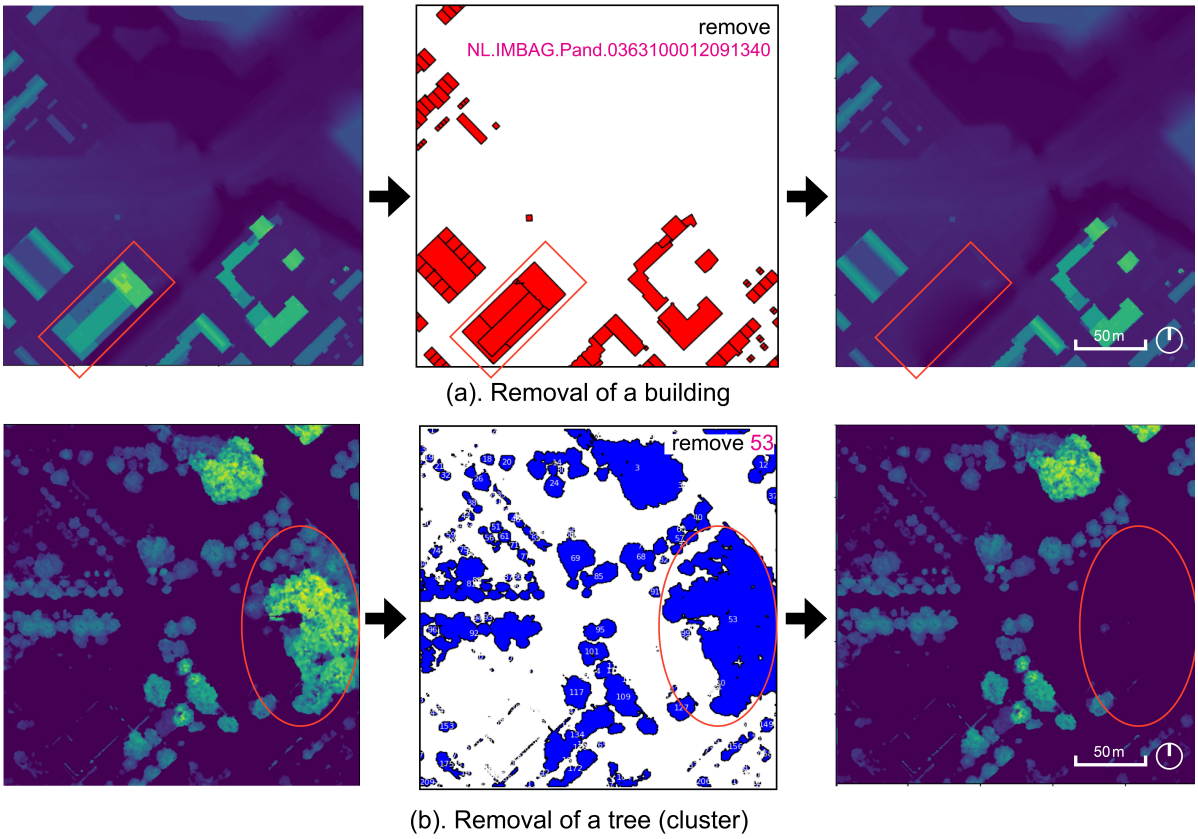

Figure 5.16: Example of removing existing elements from the scene based on user-specified IDs. (a) shows the removal of selected buildings, an (b) the removal of selected trees.

Figure 5.17 illustrates the result of deleting existing trees and buildings by specifying their IDs as input. The figure also shows the original CHM and DSM, along with the ID polygons, which serve as masks during the removal process. The original CHM and DSM datasets are stored to allow reverting changes if needed.

Tree Insertion

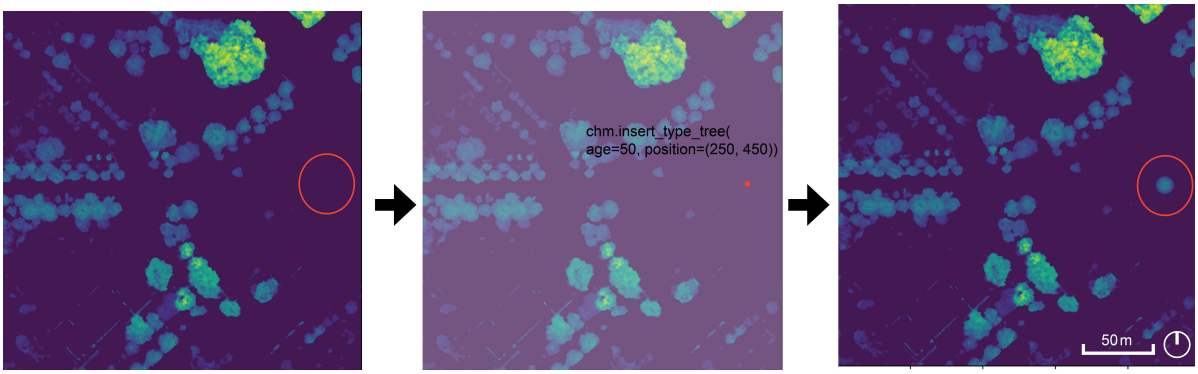

Figure 5.17: Insertion of a *Fraxinus excelsior* tree of age 50 into the CHM. Left shows the original CHM before modification, and right shows the CHM after inserting the tree at the user-defined location.

Figure 5.17 demonstrates the successful insertion of a user-defined tree into the CHM. In this example, a *Fraxinus excelsior* aged 50 years was selected from the precomputed tree database.

A current limitation of the method is that it only supports the insertion of individual trees via manually specified coordinates. Within SOLFD, this means that each tree must be defined separately, which can be time-consuming and inefficient, especially when a designer already has a full planting plan or point dataset. To improve usability, future functionality should allow batch insertion from a point layer, which would automatically convert geospatial coordinates into grid indices and apply the appropriate tree parameters from the database.

3D Urban Design Integration

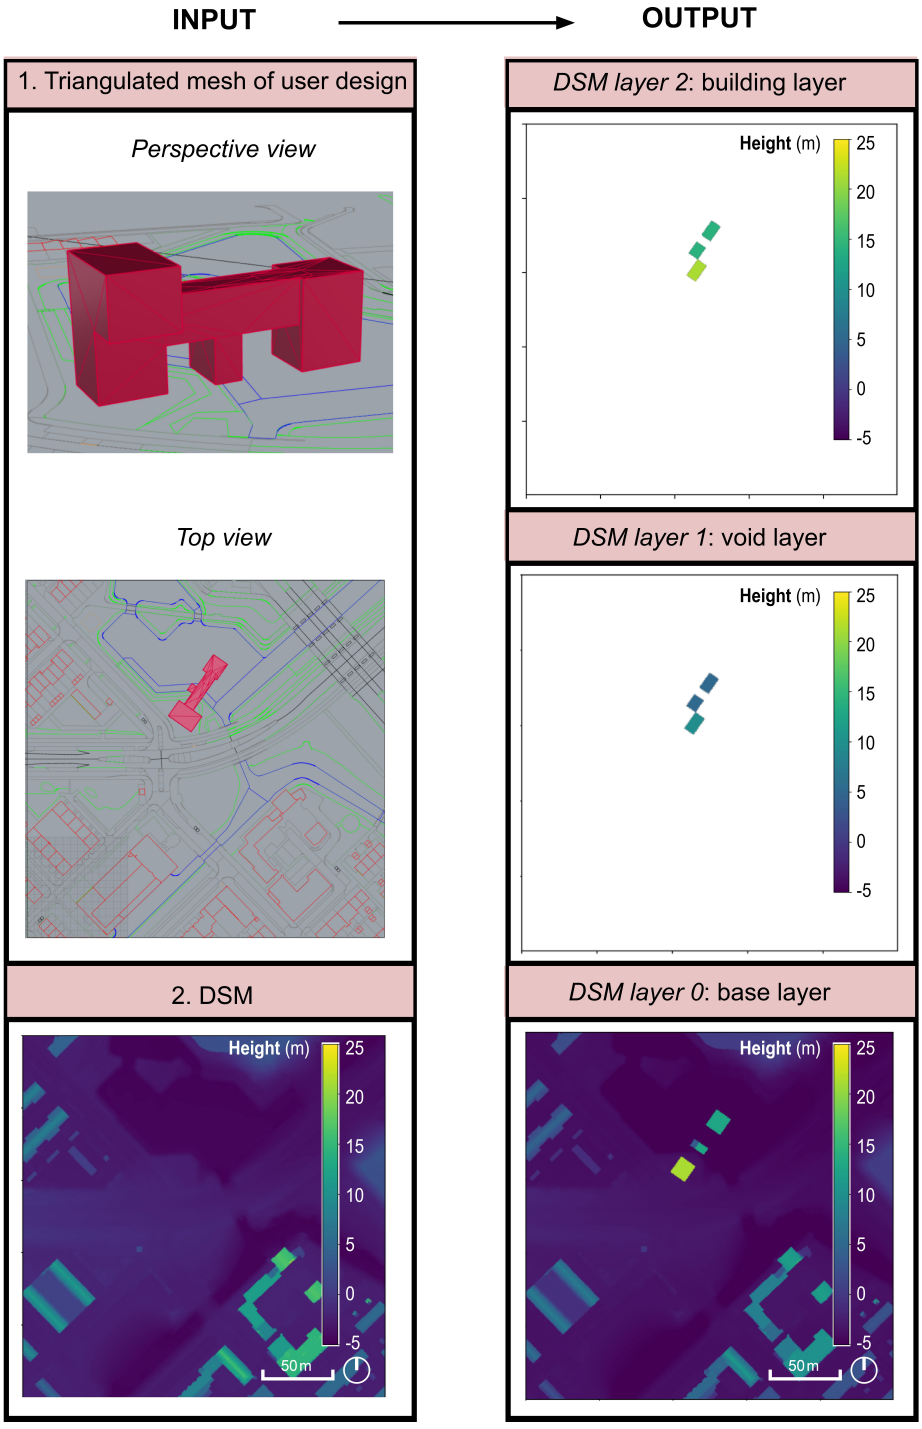

Figure 5.18: Example of integrating new 3D urban design in existing context with 1 void layer. Inputs are the user created triangulated 3D .obj mesh and the SOLFD DSM. As output, the DSM is updated to contain multiple layers; depicted here. Layer 0 contains the base layer, layer 1 the void layer and layer 2 the building layer.

The new method for 3D urban design integration in SOLFD was tested by inserting a building block from a triangulated .obj mesh file. The result, shown in Figure 5.18, demonstrates the successful integration of the 3D design into the existing DSM. As the *one-void-layer* option was selected, the output consists of three layers: the base layer, the void layer, and the building layer.

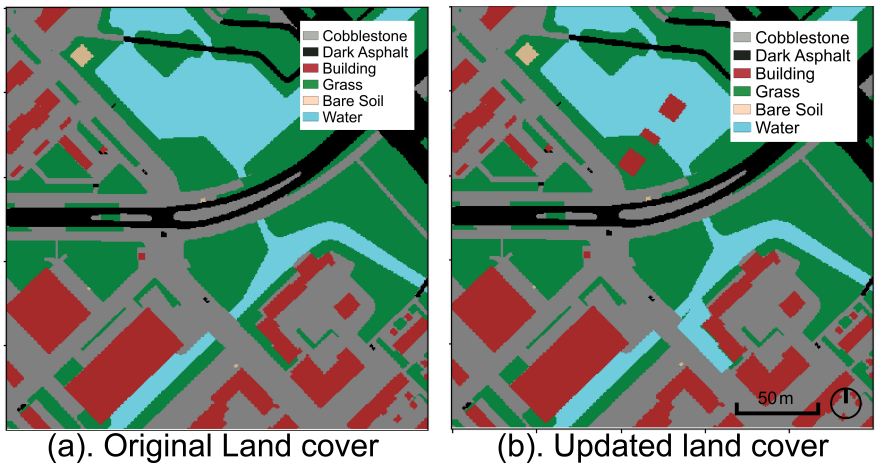

Figure 5.19: Example of updated land cover classification after integrating new 3D design in SOLFD. (a) shows the original land cover; (b) shows the updated land cover after insertion of the 3D .obj geometry.

In addition to the 3D geometry, the land cover classification is also automatically updated to reflect the newly inserted building surfaces. The original and modified land cover maps for the same location are shown in Figure 5.19. This confirms that the system correctly modifies both the elevation model and the land cover classification based on user-defined 3D input.

Urban Design Integration for Materials

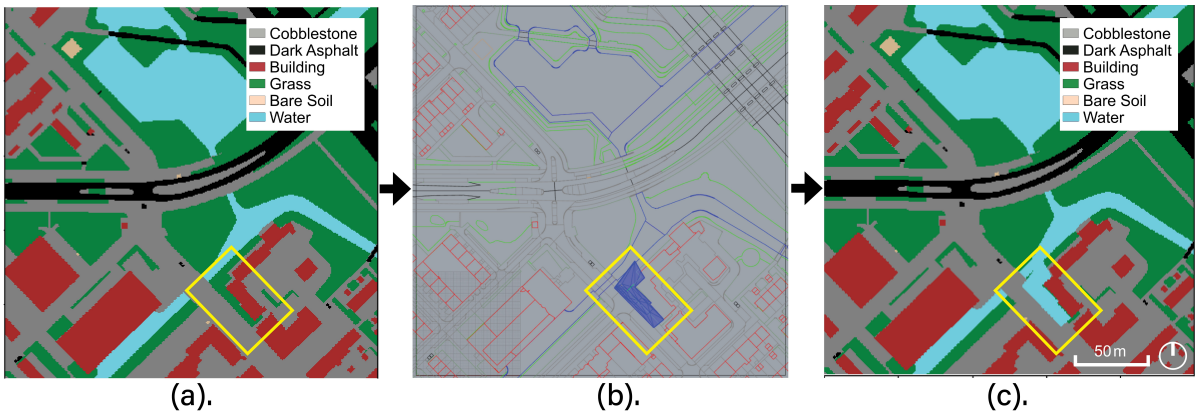

Figure 5.20: Example of land cover surface material integration in urban design, showing: (a) the original retrieved land cover; (b) the SOLFD output in a CAD program with the user-defined surface modifications overlaid as a triangulated mesh; and (c) the updated classification result.

The new SOLFD functionality to allow user-defined modifications of surface materials was tested. The input, modification process, and outcome are shown in Figure 5.20. The result demonstrates that the implemented feature correctly updates the classification.

However, material integration for vertical surfaces, such as walls, is currently missing. As noted in Section 5.1.2, SOLWEIG supports only a single wall material and does not differentiate between wall surface types. To enable complete urban design integration for material representation, support for wall material differentiation must be added. This would require the development of a new method capable of processing and assigning material changes to vertical surfaces.

Additionally, the current workflow is limited in flexibility: each surface material class and its associated 3D .obj geometry must be exported and inserted separately. This makes the design integration process less efficient. A more convenient approach would support a unified input format that allows the complete urban design. including all materials and geometry classes. to be provided in a single operation.

5.1.4. Tailoring SOLWEIG Outputs for Design and Analysis

SOLFD includes a method to post-process the outputs generated by the SOLWEIG model, enabling more effective interpretation and application of simulation results in design contexts. Figure 5.21 presents an overview of these outputs, which include summary statistics, a PET map classified according to physiological thermal stress categories, an hourly T_{mrt} map reflecting thermal conditions throughout the day, and an aggregated PET map representing overall thermal exposure.

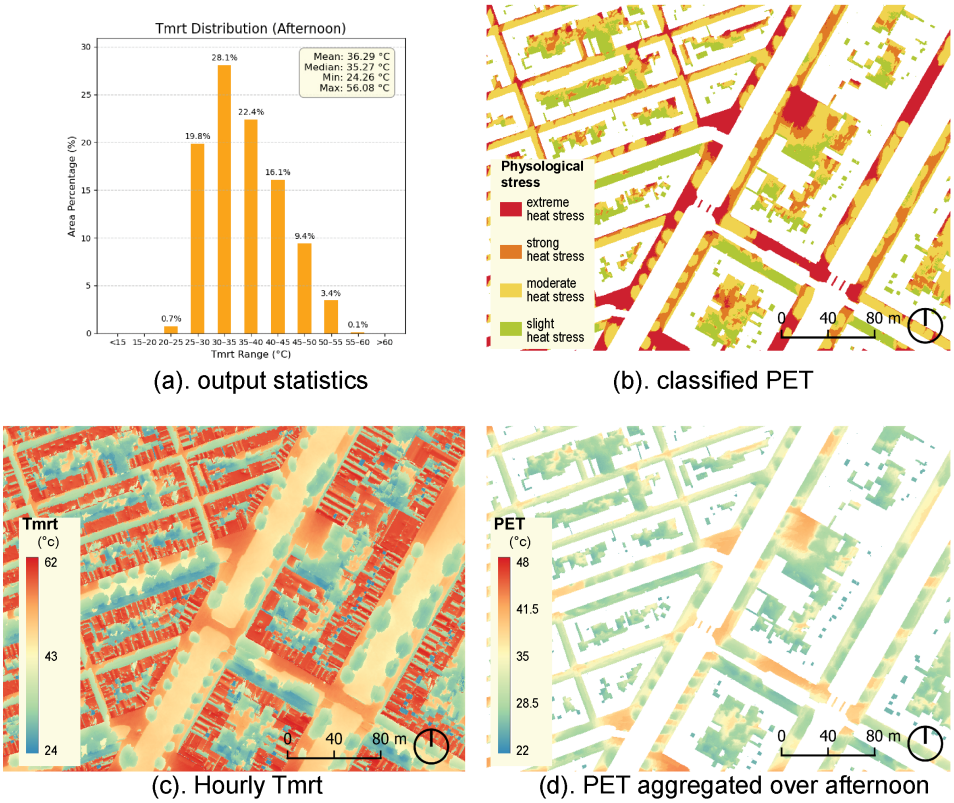

Figure 5.21: Overview of post-processed SOLWEIG outputs in SOLFD, including (a) output statistics, (b) classified PET based on physiological thermal stress, (c) hourly T_{mrt}, and (d) aggregated PET maps.

Validation of the Lookup Table.

For an input array of size 598 × 798 cells, the original (unoptimized) implementation required approximately 143 seconds per computation. Conversion to Cython significantly improved performance, reducing the runtime to 15.13 seconds. The introduction of the precomputed lookup table further accelerated the process, decreasing computation time to roughly 0.02 seconds.

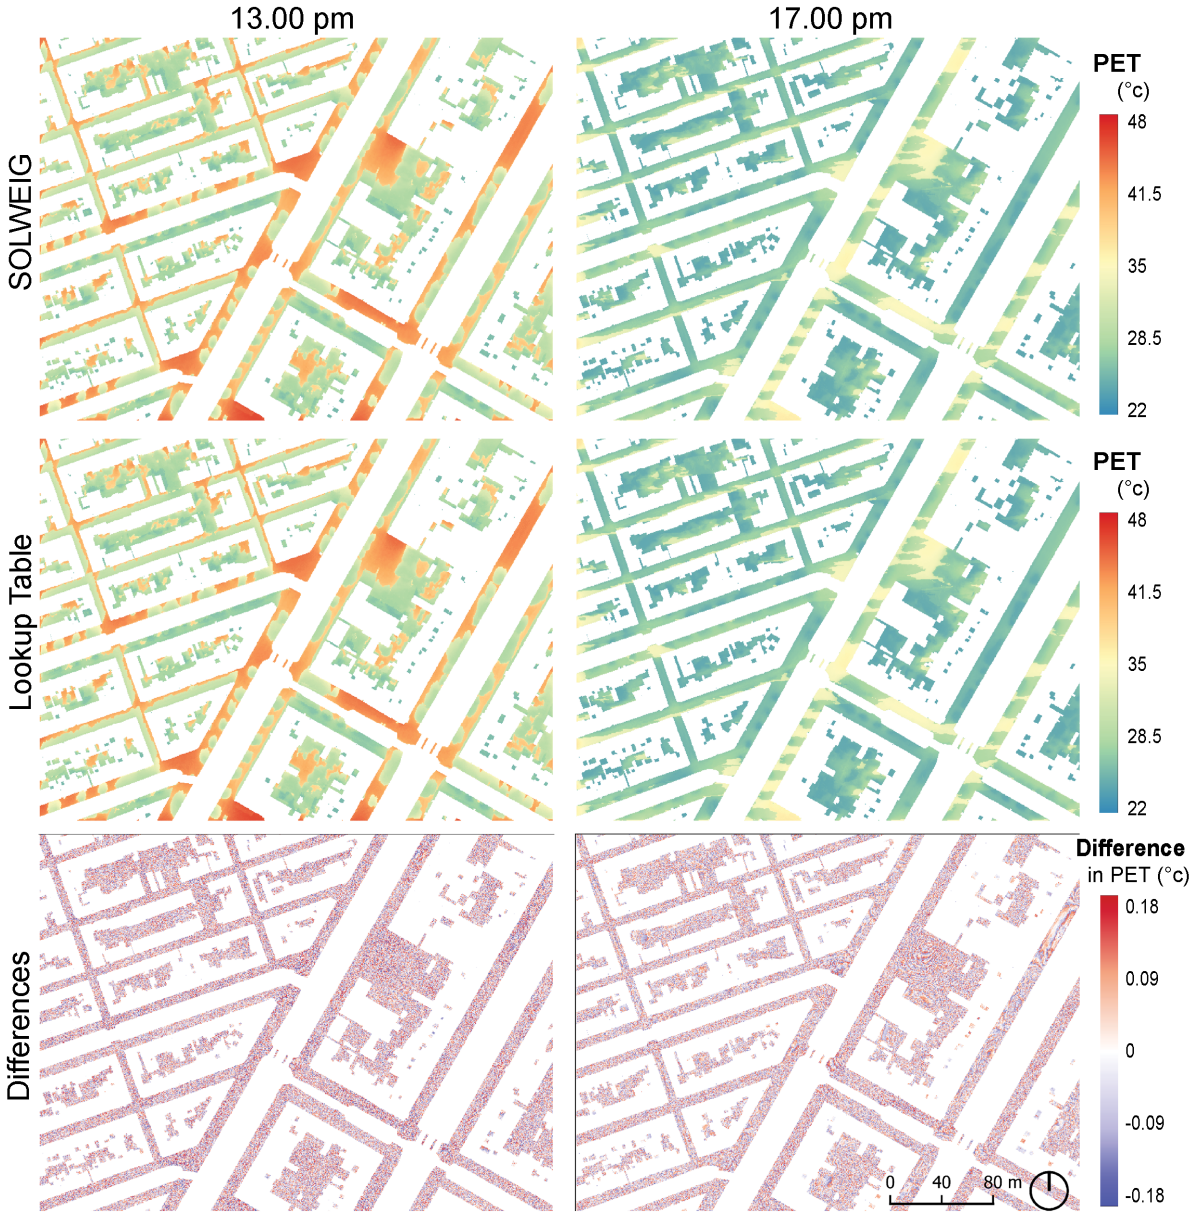

Figure 5.22: Comparison of PET values computed using the original SOLWEIG method and the newly implemented lookup table method at 13:00 and 17:00. The difference in value is given below.

For the T_{mrt} simulation, the meteorological dataset representing an average day with an air temperature range of 25 °C to 30 °C was used (see Appendix A.6). For the computation of PET, input parameters were set as follows: v_a of 0.1 m/s, T_a of 26°C, and RH of 50%.

Figure 5.22 presents the resulting PET values at 13:00 and 17:00 computed using both the original SOLWEIG method and the lookup table approach developed here, along with the difference between these two methods. The observed differences remain below 0.2°C, with values alternating slightly higher or lower spatially, producing a grainy appearance in the difference map. These differences arise because the lookup table uses discretized T_{mrt} steps rather than continuous values, as the original on-the-fly calculation does, resulting in slight spatial variations due to the limited precision of the discretized inputs.

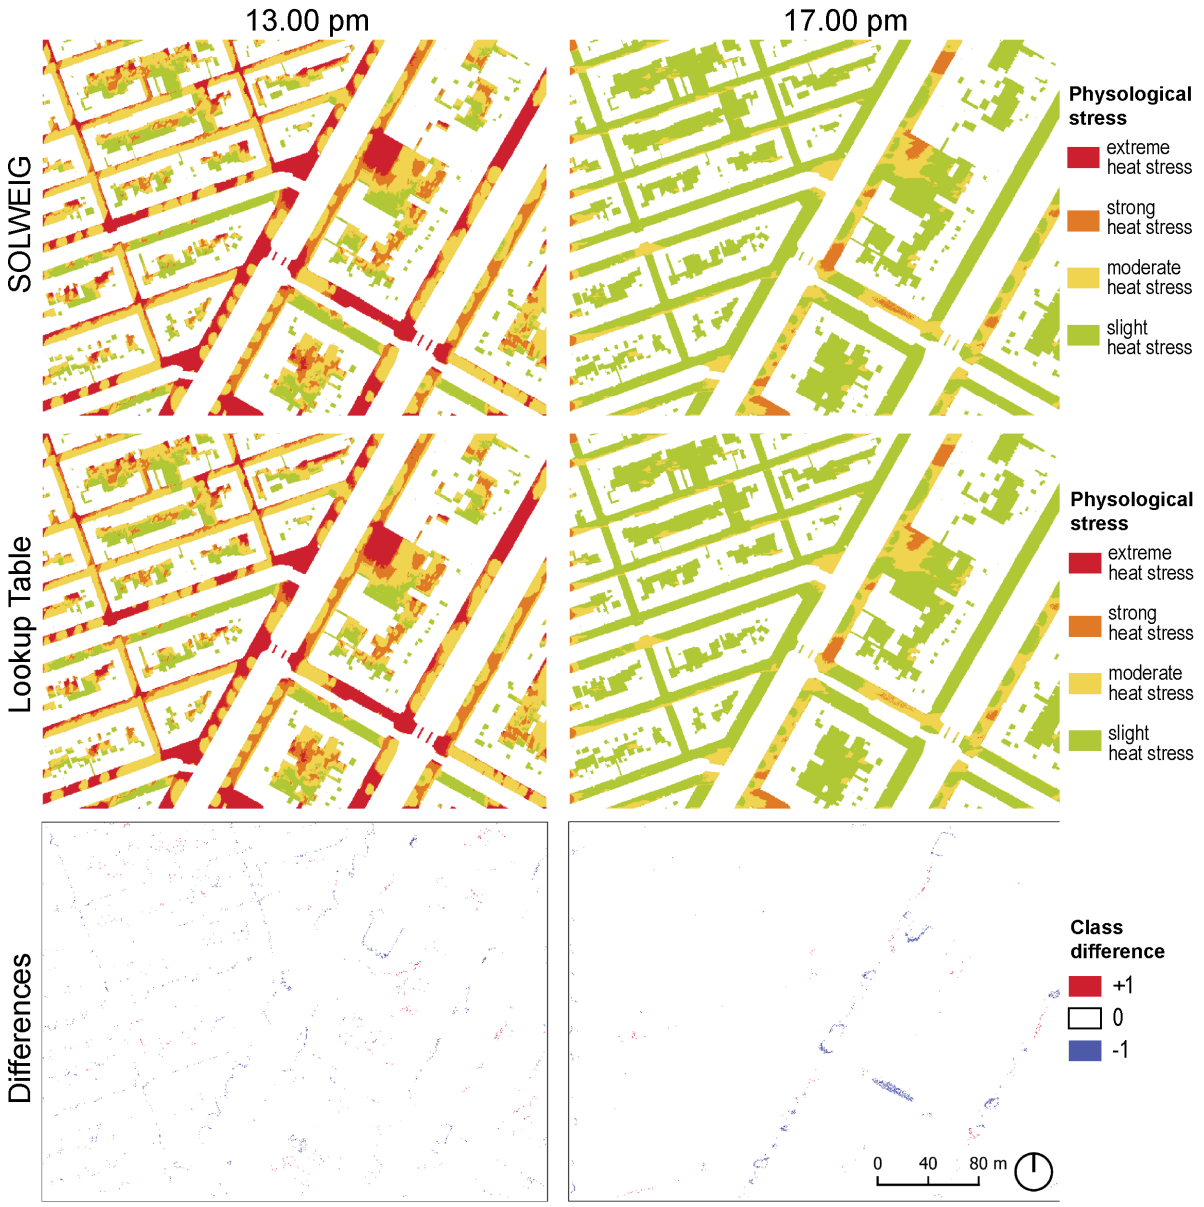

Figure 5.23: Resulting classification of the two different methods and their difference.

The physiological stress classification based on the PET values is shown in Figure 5.23, comparing the results from both the lookup table method and the original SOLWEIG calculation, as well as the difference between the two. Differences are primarily observed at the boundaries between classes, where the discretized values from the lookup table method fall into one class higher or lower compared to the continuous SOLWEIG output. Overall, the lookup table approach effectively preserves the spatial influence of building blocks and tree design on thermal patterns.

As noted in Section 4.1.4 and reiterated here, it is important to recognize that wind speed is also influenced by urban morphology, accelerating or decelerating locally. Therefore, the wind speeds included in the lookup table should be regarded as overall indications of thermal conditions under different wind scenarios. For higher accuracy, a dynamic wind simulation should be incorporated. The interaction between wind and radiation is complex: strategies for improving T_{mrt} through shading (high buildings and trees) can also adversely impact ventilation (reduced airflow under trees or increased wind speeds around tall structures).

Future extensions could include negative T_a values and higher v_a values, enabling the model to also

assess cold stress conditions.

5.1.5. Computation Time Optimization

This section presents the results of computational performance improvements achieved through GPU acceleration. It first details the speedups obtained through code optimization, followed by an evaluation of the impact of using the 3D SOLFD model with the layered array on computation time. Finally, the influence of spatial resolution on model accuracy is analyzed.

Optimization by Code Enhancement

Table 5.4: Comparison of computation times for the new (SOLFD), GPU-accelerated, and the old version of SVF and SOLWEIG on the computer.
OOM = out of memory

Array Size	SVF (Without CHM)		SVF (With CHM)		SOLWEIG (Without CHM)		SOLWEIG (With CHM)	
	New (s)	Speedup	New (s)	Speedup	New (s)	Speedup	New (s)	Speedup
250x250	3.0	0.60x	11.5	1.03x	17.7	1.6x	18.4	1.6x
500x500	5.8	3.6x	18.9	8.13x	18.7	16.7x	19.3	17.4x
1000x1000	17.9	9.5x	45.3	26.6x	26.1	48.6x	26.0	55.8x
1500x1500	34.6	12.2x	67.9	46.1x	53.1	56.0x	50.5	66.4x
2000x2000	59.0	14.8x	119.4	53.9x	102.5	54.9x	96.5	60.6x
3000x3000	195.6	10.7x	328.3	45.1x	OOM	—	OOM	—

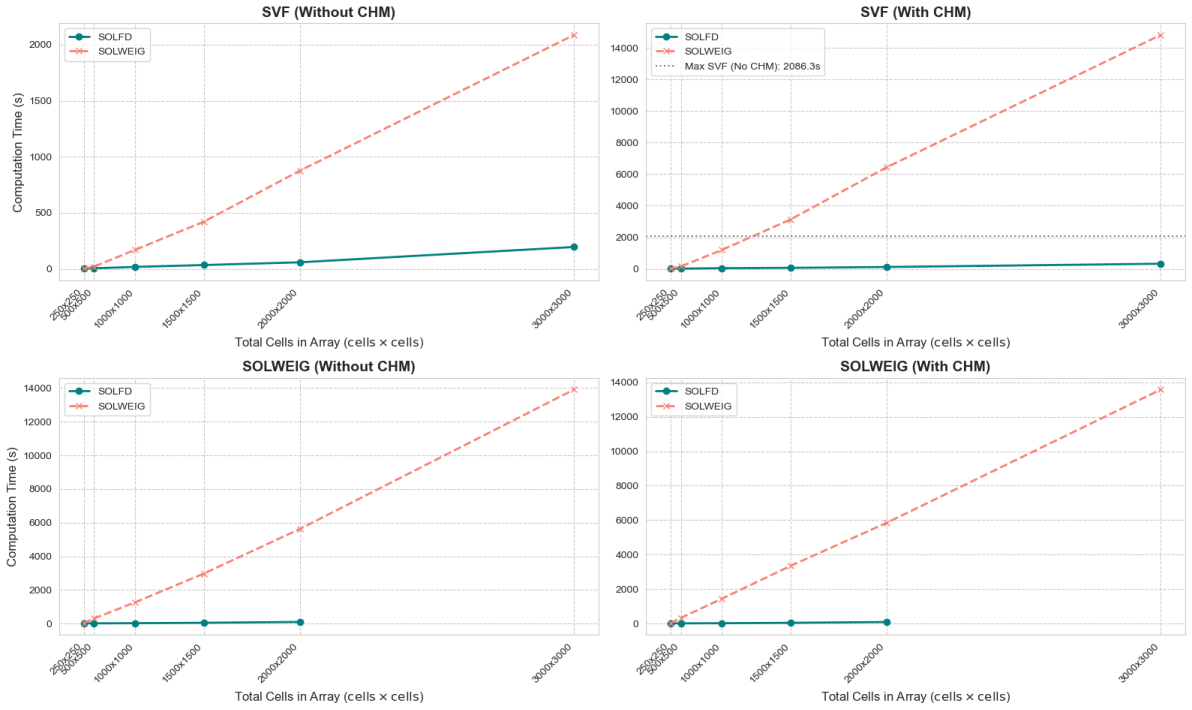

Figure 5.24: Computation time comparison between the old SOLWEIG (CPU-based) and new SOLFD (GPU-accelerated) code versions for SVF and SOLWEIG on the computer, across different array sizes, with and without CHM data.
Note: the computation time scale between the two SVF plots differ, the maximum computation time of the SVF without CHM is depicted in the SVF with CHM as a line for comparison.

Table 5.5: Comparison of computation times for the new (SOLFD), GPU-accelerated, and the old version of SVF and SOLWEIG on the laptop.
OOM = out of memory

Array Size	SVF (Without CHM)		SVF (With CHM)		SOLWEIG (Without CHM)		SOLWEIG (With CHM)	
	New (s)	Speedup	New (s)	Speedup	New (s)	Speedup	New (s)	Speedup
250x250	5.8	0.91x	32.7	0.94x	46.0	2.37x	57.5	2.0x
500x500	11.3	5.4x	50.3	8.2x	67.6	10.5x	67.3	12.3x
1000x1000	56.0	8.6x	237.3	12.8x	407.6	6.9x	351.8	9.8x
1500x1500	144.2	7.9x	649.3	11.6 x	2717.0	3.2x	OOM	—
2000x2000	OOM	—	OOM	—	OOM	—	OOM	—

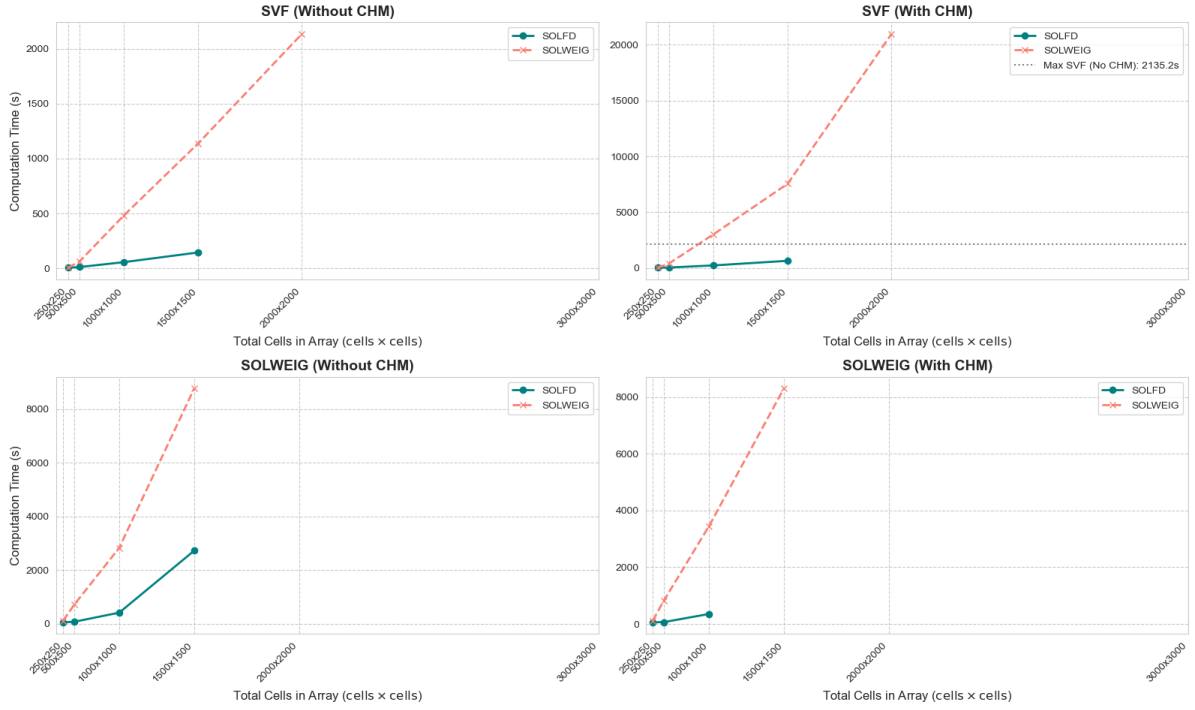

Figure 5.25: Computation time comparison between the old SOLWEIG (CPU-based) and new SOLFD (GPU-accelerated) code versions for SVF and SOLWEIG on the laptop, across different array sizes, with and without CHM data.
Note: the computation time scale between the two SVF plots differ, the maximum computation time of the SVF without CHM is depicted in the SVF with CHM as a line for comparison.

The optimization results are shown by comparing the old SOLWEIG (CPU-based) and optimized SOLFD (GPU-accelerated) code on the computer (Figure 5.24) and the laptop (Figure 5.25) for a range of array sizes. The speedup results are summarized in Table 5.4 for the computer, and Table 5.5 for the laptop.

Figures 5.24 and 5.25 show that for the CPU computation method, time generally increases linearly with array size, with more pronounced increases at larger sizes. While the GPU-accelerated code also exhibits linear increasing computation time with array size, the rate of increase is significantly lower. A sharp increase in computation time, visible as a bend towards the end of the curves, indicates near GPU memory saturation. On both systems, running out of memory significantly impacts performance, causing failures for larger array sizes, starting at 2000×2000 on the laptop and 3000×3000 on the computer when CHM is included. While the original code is significantly slower, it does complete execution. This limitation can be addressed by tiling large arrays and including a buffer zone around each tile, derived from the bounding box and the GPU's memory constraints.

The computer, with the more powerful GPU, shows significant performance gains. Notably, the speedup increases with array size for most cases, particularly for SOLWEIG (with and without CHM) and SVF with CHM. SVF with trees achieved a speedup of up to 54x, while SOLWEIG with CHM reaches up to 66x. However, SOLWEIG calculations show diminishing returns beyond 2000×2000, likely due to the

aforementioned GPU memory saturation.

In contrast, the laptop achieved more modest improvements, with speed-up generally peaking around 10× for larger arrays. Memory limitations are encountered at 2000 × 2000, where execution fails due to insufficient GPU capacity.

While raw performance varies depending on the setup and input type, the optimization consistently reduces computation times to practical levels. For array sizes up to 1500×1500 cells, computation time is reduced from 10 minutes to 1.5 hours down to just 1–5 minutes, depending on the exact array size. For example, on a computer with a powerful GPU, an urban scene with a 2000×2000 cell array including CHM data now completes in approximately 3.5 minutes, compared to over 3 hours originally. This improvement enhances SOLFD's suitability for iterative and comparative design workflows.

It is important to note that, beyond array size, computation time is also influenced by the tallest building or tree in the dataset. The shadow-casting loop continues until the shadow either exits the array bounds or reaches ground level. Consequently, scenes with high-rise buildings or tall vegetation require a higher maximum number of iterations, leading to increased computational cost.

Influence of Layered array on computation time

Enabling SOLFD 3D allows simulations with layered arrays, providing more detailed vertical representation. However, using this feature increases computation time. Tests were performed using a single void layer, which is the recommended setting for efficiency.

The results compare SOLFD without layered arrays (2.5D) and with 3D enabled, highlighting the difference in computation times. These are shown for both the computer (Figure 5.26) and the laptop (Figure 5.27) over the same range of array sizes as previously tested. The relative slowdowns when using layered arrays are summarized in Table 5.6 for the computer and Table 5.7 for the laptop.

Table 5.6: Comparison of computation times for the SOLFD 2.5D and 3D version on the computer. OOM = out of memory

Array Size	SVF (No CHM)		SVF (With CHM)		SOLWEIG (No CHM)		SOLWEIG (With CHM)	
	3D (s)	Slowdown	3D (s)	Slowdown	3D (s)	Slowdown	3D (s)	Slowdown
250x250	11.0	3.6x	19.7	1.7x	18.4	1.0x	18.8	1.0x
500x500	16.7	2.9x	31.2	1.7x	18.9	1.0x	19.8	1.0x
1000x1000	39.3	2.2x	70.7	1.6x	27.3	1.0x	28.9	1.1x
1500x1500	58.2	1.7x	97.2	1.4x	53.7	1.0x	53.9	1.1x
2000x2000	100.7	1.7x	159.1	1.3x	105.0	1.0x	103.9	1.1x
3000x3000	856.2	4.4x	OOM	—	OOM	—	OOM	—

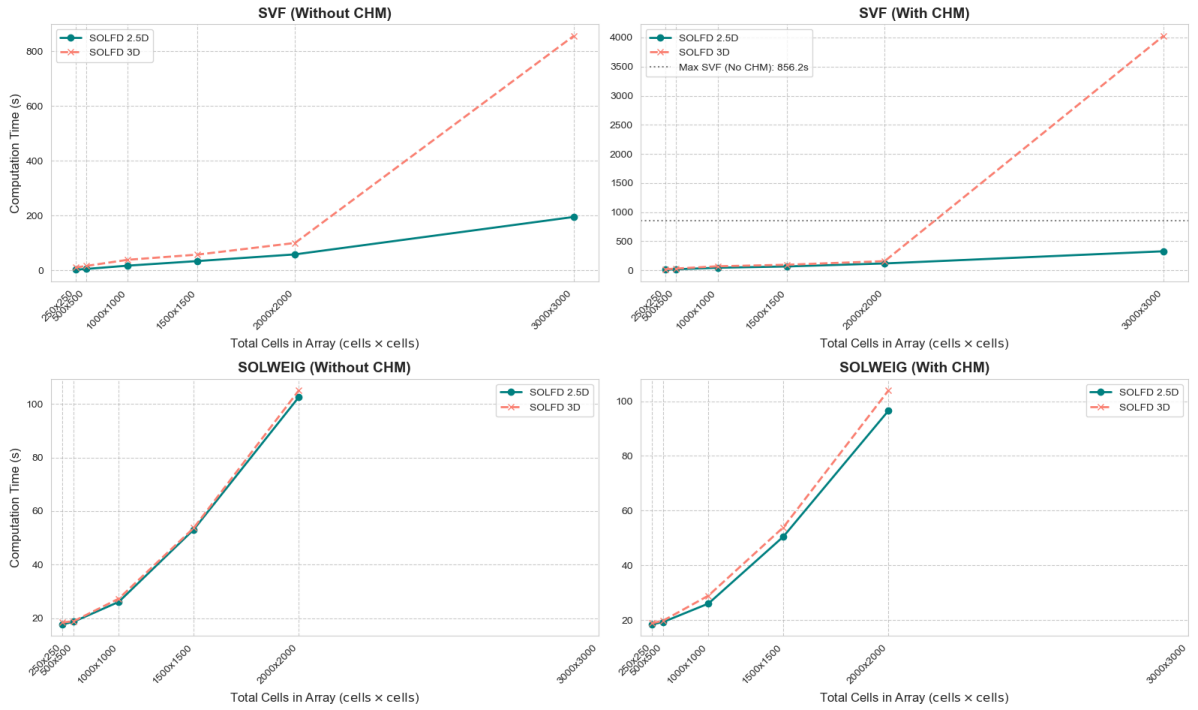

Figure 5.26: Computation time comparison between SOLFD 2.5D and SOLFD 3D for SVF and SOLWEIG on the computer, across different array sizes, with and without CHM data.
Note: the computation time scale between the two SVF plots differ, the maximum computation time of the SVF without CHM is depicted in the SVF with CHM as a line for comparison.

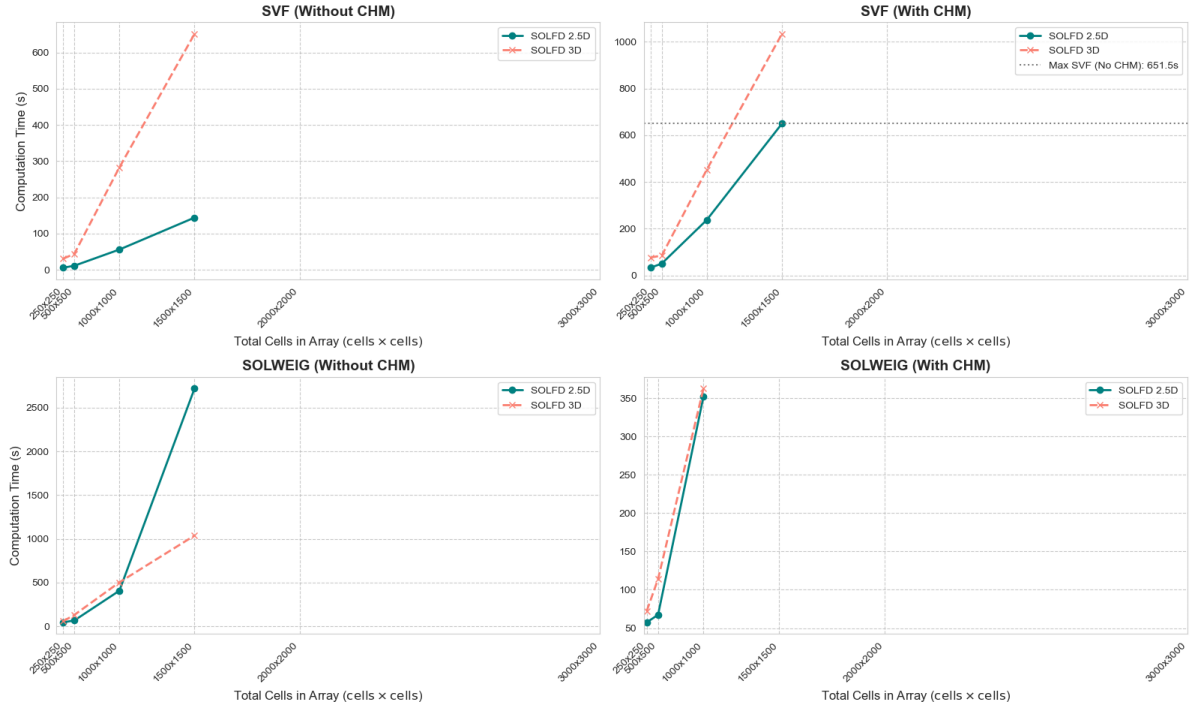

Figure 5.27: Computation time comparison between SOLFD 2.5D and SOLFD 3D for SVF and SOLWEIG on the laptop, across different array sizes, with and without CHM data.
Note: the computation time scale between the two SVF plots differ, the maximum computation time of the SVF without CHM is depicted in the SVF with CHM as a line for comparison.

Table 5.7: Comparison of computation times for the SOLFD 2.5D and 3D version on the laptop. OOM = out of memory

Array Size	SVF (No CHM)		SVF (With CHM)		SOLWEIG (No CHM)		SOLWEIG (With CHM)	
	3D (s)	Slowdown	3D (s)	Slowdown	3D (s)	Slowdown	3D (s)	Slowdown
250x250	30.5	5.3x	76.5	2.3x	58.3	1.3x	72.3	1.3x
500x500	44.2	3.9x	84.4	1.7x	130.7	1.9x	114.0	1.7x
1000x1000	283.0	5.1x	452.5	1.9x	503.2	1.2x	362.4	1.0x
1500x1500	651.5	4.5x	1033.4	1.6x	1037.7	0.4x	OOM	—

Comparing Figures 5.26 and 5.27, it is evident that for SVF on the laptop there is a sharp and steep increase in computation time: about 5× slower without CHM and 2× slower with CHM. In contrast, on the computer, the increase is more moderate: around 3× slower without CHM and 1.7× with CHM, with a sharp bend at the largest array size of 3000×3000, likely caused by GPU memory saturation. Compared to 2.5D, the GPU runs out of memory more quickly when using 3D due to the need to process more arrays, limiting computations to array sizes up to 1500×1500 for the laptop.

The SOLWEIG calculation is less affected by the introduction of 3D, especially on the computer, where speed remains nearly unchanged. The close-up view also shows that SOLWEIG computation time does not increase linearly but rather follows a more quadratic pattern.

Overall, introducing 3D slows down SOLFD by approximately 2–5× for SVF calculations and 1–2× for SOLWEIG. Nonetheless, these computations are still faster than those of the original SOLWEIG code without 3D option.

Resolution Selection for Optimization

The DSMs and CHMs for the three tested resolutions (0.5,m, 1.0,m, and 2.0,m) are shown in Figure 5.28. As expected, coarser resolutions result in a loss of detail and a blurrier appearance. Nevertheless, the primary boundaries of buildings and trees remain distinguishable.

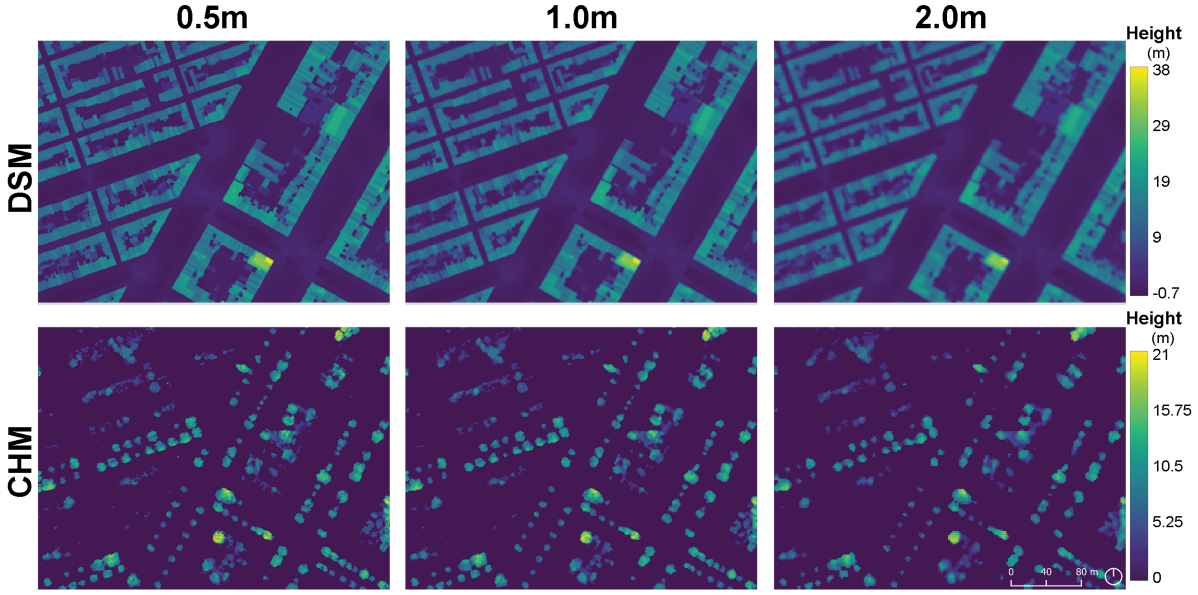

Figure 5.28: DSMs and CHMs at 0.5 m, 1.0 m, and 2.0 m resolutions

The influence of resolution on T_{mrt} calculations for a typical warm summer day (average of days exceeding 30,°C; see Table A.7) is illustrated in Figure 5.29, for 13:00 and 18:00. Overall, the T_{mrt} output remains consistent across all resolutions and time points, though finer resolutions reveal more detailed shading patterns.

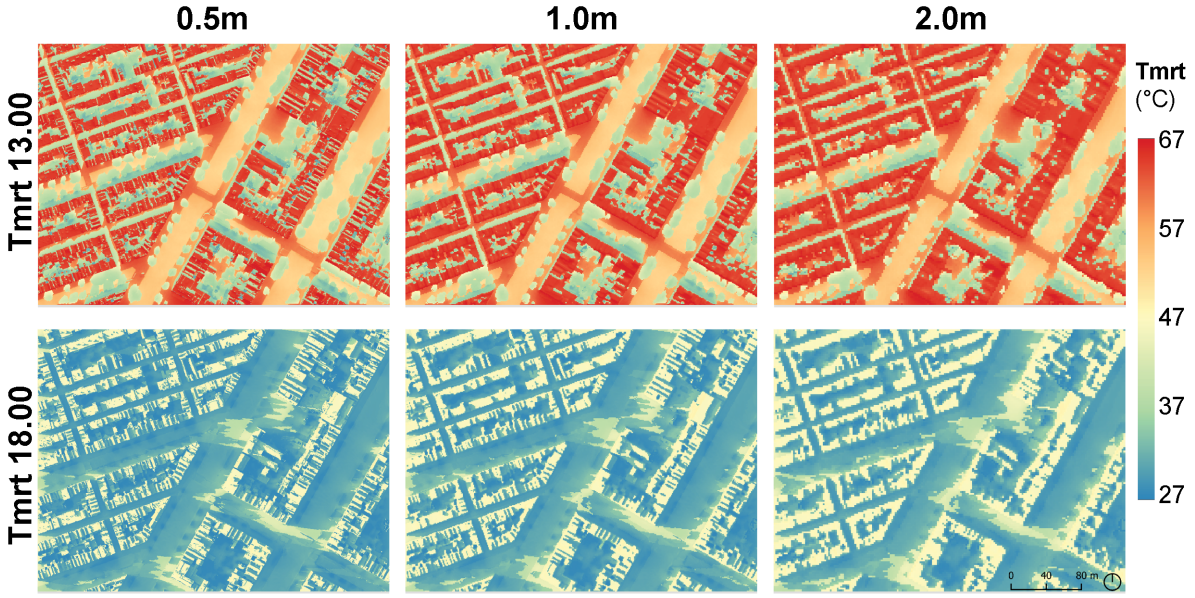

Figure 5.29: T_{mrt} distributions at 13:00 and 18:00 for the 0.5 m, 1.0 m, and 2.0 m resolution cases.

Figure 5.30 presents the difference in T_{mrt} relative to the 0.5 m resolution baseline. As resolution coarsens, fine rooftop features are increasingly lost, resulting in reduced rooftop shading and subsequently higher T_{mrt} values. However, since T_{mrt} is evaluated at pedestrian height, these rooftop discrepancies have limited practical relevance. The area of differences becomes larger at lower sun angles due to longer shadows, but the magnitude of T_{mrt} variation decreases, reflecting reduced solar radiation.

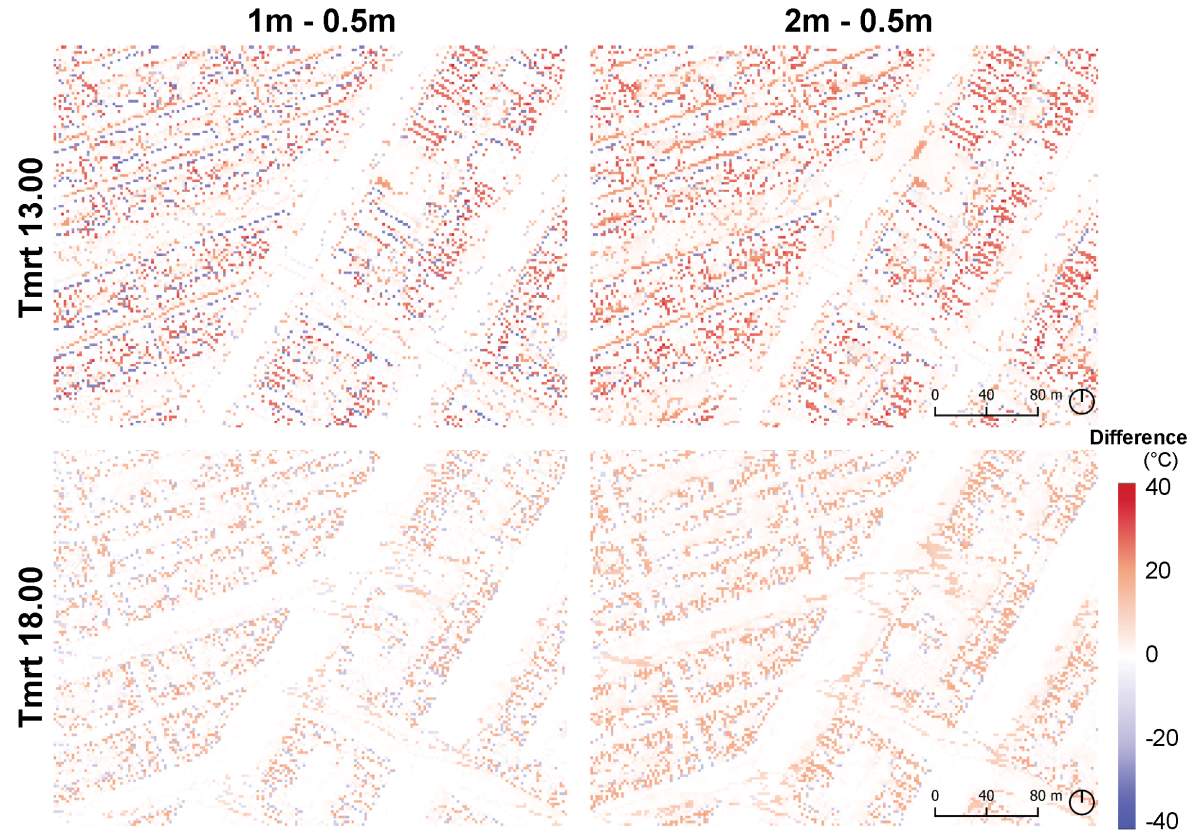

Figure 5.30: Difference in T_{mrt} at 13:00 and 18:00 for the 1.0m and 2.0m resolution cases, compared to the 0.5m baseline.

For early-stage urban design or comparative analyses, a resolution of 1.0 m offers a good trade-off between computational efficiency and spatial accuracy. Although edges of shade geometry become warmer compared to the 0.5 m case, mainly due to height loss introduced by the interpolation method, overall patterns remain sufficiently accurate for early-stage design evaluations. At 2.0 m resolution, the edges become larger, but should still be able to give a reasonable estimation for rougher calculations. One caveat is that coarser resolutions tend to fill in narrow voids in 3D modeling scenarios, which can lead to underestimation of radiation exposure. Therefore, finer resolution is generally preferred where feasible for 3D simulations.

5.2. Accuracy Assessment of SOLWEIG and SOLFD

This section provides an overview of two categories of validation studies: (1) those conducted by the developers of the model themselves, and (2) those carried out by independent researchers. A general conclusion is also presented at the end of the section.

Makers of SOLWEIG

Several studies conducted by the developers of SOLWEIG have validated and refined the model over time. These efforts primarily aim to enhance the accuracy of modeled T_{mrt} by improving the representation of vegetation, land cover, and radiative transfer processes. Observational validation is based on measurements from net radiometers, which record incoming and outgoing shortwave and longwave radiation in all cardinal directions, as well as upward and downward fluxes.

Table 5.8: Overview of SOLWEIG model validation results by original developers, showing key performance metrics (RMSE, R^2) for different model features and measurement conditions, and remarks.

Study / Version	Feature Tested	RMSE (°C)	R^2	N	Remarks
Lindberg and S. Grimmond (2011) v2.0	Vegetation scheme	*overall measurements*: 6.79 / *Measurements under trees*: 3.1	0.93 / 0.91	205 / 50	Improved 3D radiation fluxes; slight overestimation in vegetated shade; assumes vegetation temperature = air temperature
Lindberg et al. (2016), v2015a	Land cover parametrization	*building*: 4.42 / *grass*: 4.68	0.923 / 0.863	200	Land cover has small effect on Tmrt (5°C) compared to shadowing (30°C); improved 15-min temporal resolution
Wallenberg et al. (2020) v2019a	Anisotropic shortwave model	–	–	–	More realistic diffuse shortwave radiation; higher Tmrt in sunlit areas (+3°C); reduced over-/underestimation
Wallenberg et al. (2023) v2022a	Anisotropic longwave model	4.6	0.9	–	Better estimation in shaded areas; remaining errors due to uniform wall temperature assumption (e.g., material/orientation differences)

Table 5.8 summarizes key validated developments of SOLWEIG. Each enhancement; vegetation scheme, land cover parametrization, and anisotropic radiation models, improved accuracy. Version 2.0 introduced vegetation, reducing T_{mrt} overestimation and lowering RMSE from 8.63°C to 6.79°C overall and to 3.1°C under trees (R^2 = 0.93 and 0.91) (Lindberg and S. Grimmond, 2011). However, hourly input still caused overestimation in briefly sunlit areas due to coarse temporal resolution.

The land cover parametrization added surface-specific temperature profiles; though its impact on T_{mrt} (5°C) was modest compared to shading (30°C). Accuracy for this evaluation improved by using a 15-minute temporal resolution (RMSE = 4.68°C for grass, R^2 = 0.863; RMSE = 4.42°C for buildings, R^2 = 0.923) (Lindberg et al., 2016).

Later, anisotropic shortwave and longwave models replaced isotropic assumptions, capturing directional radiation effects. This raised T_{mrt} in sunlit zones (+3°C) and lowered it in shade (Wallenberg et al., 2020), while distinguishing longwave sources (e.g., sky, vegetation) further improved accuracy (RMSE = 4.6°C, R^2 = 0.9) (Wallenberg et al., 2023).

Nonetheless, limitations remain. The assumption of uniform wall temperatures, which neglects material and orientation differences (e.g., east-facing walls warm earlier than west-facing ones). Inaccuracies

also arise from shadow classification errors linked to DSM and CHM quality and temporal resolution. Overall, SOLWEIG performs best under clear sky conditions and with high temporal resolution input.

Validation by Others

Table 5.9 provides an overview of recent studies validating SOLWEIG, including the calculation of the RMSE and R^2. Overall, the results are consistent with those of SOLWEIG, with comparable RMSE (ranging from 2.22 to 7.67) and R^2 values (ranging from 0.49 to 0.95), although some R^2 values are considerably lower.

In conclusion, the accuracy of SOLWEIG depends on the time resolution used for modeling, particularly during non-clear sky days. The largest errors are caused by shade mismatching, where cells are incorrectly classified as either in the shade or not in the shade. This error stems from inaccuracies in the DSM and CHM, as well as the time resolution. Additionally, the simplification of longwave radiation fluxes from walls and the assumption that surface temperature is uniform across all wall orientations leads to an overestimation of T_{mrt} in shaded areas. The model performs most accurately during daytime modeling with clear sky conditions.

Study	Site	Time Range	n	R^2	RMSE (°C)	Remarks
Thom et al. (2016) v2014a	HD	07:30–20:00	130	0.93	4.96	Observations from black globe temperature assume equal angle factors (Fi = 0.167), leading to overestimation under high solar elevation compared to human body model angle factors used in SOLWEIG. DSM resolution and geometric shading errors may distort shade boundaries.
		10:00–16:00	65	0.91	3.50	
	RES	07:30–20:00	130	0.91	5.40	
		10:00–16:00	65	0.81	3.27	
	OG	07:30–20:00	130	0.92	3.89	
		10:00–16:00	65	0.74	3.76	
	MIX	07:30–20:00	130	0.84	7.67	
		10:00–16:00	65	0.76	5.10	
	OIP	07:30–20:00	130	0.89	5.72	
		10:00–16:00	65	0.70	5.04	
Gál and Kántor (2020) v2019a	Szeged	Aug 7–8 (26h)	98	0.92	5.02	SOLWEIG tends to overestimate T_{mrt} in shade and underestimate it in sunlit areas. Largest errors occur at low solar altitude, mainly due to shade mismatching. Nighttime T_{mrt} underestimation originates from use of domain-wide mean surface temperatures.
Buo et al. (2023) v2021a	Overall	7:00–21:00 on 9 summer days	763	0.83	5.59	Field data from 60 unique sites in Phoenix-Tempe using human-biometeorological station with sensors in the six directions. Best model performance in building canyons, worst under trees. Model overestimates T_{mrt} in shade, due to simplified wall flux estimation and use of domain-wide surface temperature. Tree shade mismatch (from crown shape errors) responsible for largest errors, also getting the correct tree transmissivity. No land cover scheme used.
	Shade		538	0.56	5.12	
	Sun-exposed		225	0.49	6.60	
	Trees		407	0.56	5.44	
	Open sites		257	0.85	6.23	
	Building canyons		99	0.88	4.38	
Guo et al. (2023) v2019a	Open areas	09:30–16:30, clear summer day	<28	0.88	1.55	Measurements in Dalian, China, using net radiometer in six directions on a clear summer day. 15-minute intervals. Evaluated in open, near-wall, and vegetation environments. Note the small sample size.
	Near-wall		<28	0.95	2.22	
	Vegetation		<28	0.82	4.32	

Table 5.9: Overview of SOLWEIG model validation results by independent researchers, showing key performance metrics (RMSE, R^2) for different model features and measurement conditions, and remarks.

5.2.1. Accuracy Assessment Using Observations

This section presents the results of the accuracy assessment using observational data. The assessment is first discussed for each of the six locations individually, followed by an overall evaluation.

Performance across Locations

For all six locations depicted in Figure 4.21, T_{mrt} was modeled using three distinct approaches: the 3D model, the gap model, and the on-top model, as detailed in Section 4.2.2. The modeled results are presented in a plot alongside observed T_{mrt} data collected by the Climate Bike for each location. To facilitate comparison between measured and modeled values, the observed T_{mrt} is emphasized using thicker black circles. Moreover, corresponding T_{mrt} maps for each model are shown with the observation points overlaid. Here, the color indicates the T_{mrt}. It should be noted that while each T_{mrt} map represents a single minute in time, the observations may cover a range of minutes.

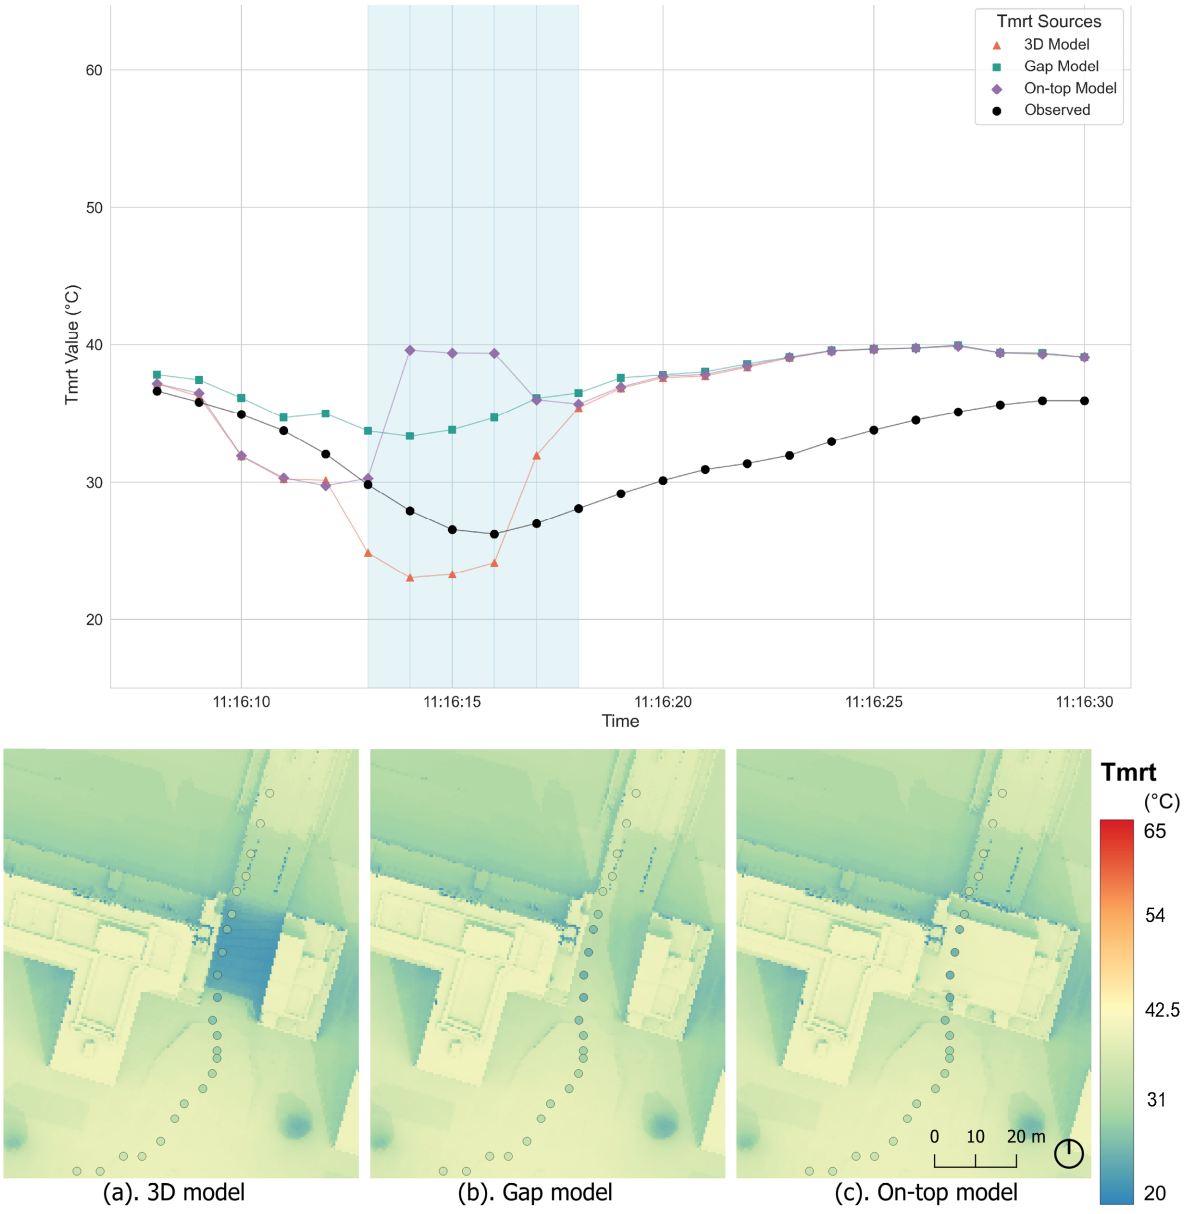

Figure 5.31: Plot of modeled and observed T_{mrt} over time, alongside maps of the three different model cases for Location 1: (a) 3D model, (b) Gap model, and (c) On-top model. The highlighted section in the T_{mrt} plot indicates measurement points located at the structure. Sample size: $n = 23$.

For Location 1, Figure 5.31 shows the plot and maps. The Climate Bike moved north to south during

measurements. Before reaching the underpass, the 3D and on-top models aligned more closely with observed T_{mrt} due to the building's effect on SVF, which is absent in the gap model.However, this shifts to underestimation once in the building's shade. Underneath the building, the 3D model underestimates T_{mrt} by 3–5°C, while the gap model overestimates it by 3–10°C, and the on-top model shows even greater overestimation. After exiting the underpass, the 3D model's temperatures rise more quickly than observed, with all models overestimating by 8°C before converging to within 3°C.

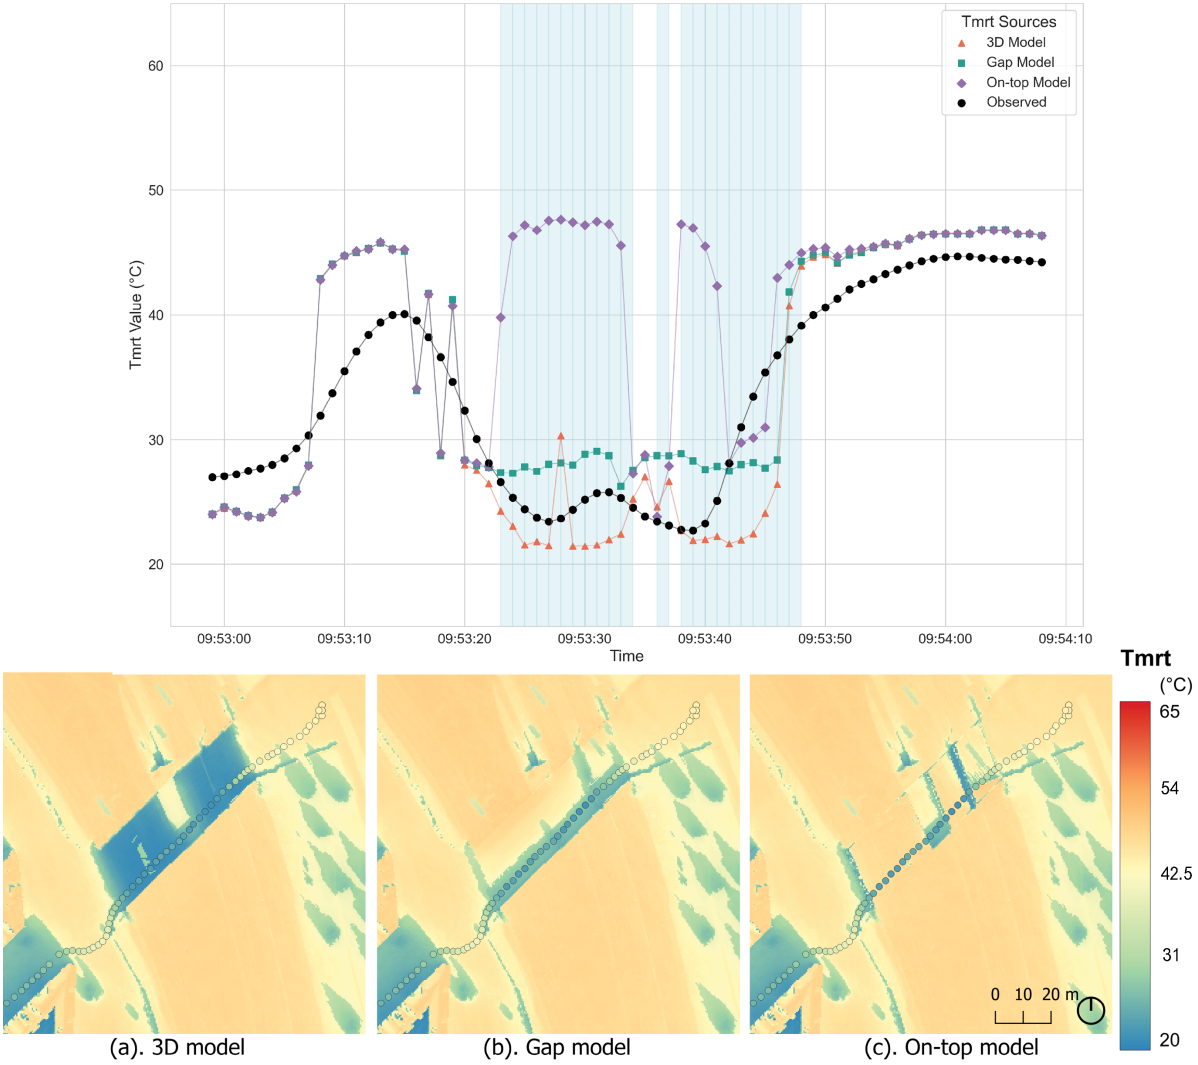

Figure 5.32: Plot of modeled and observed T_{mrt} over time, alongside maps of the three different model cases for Location 2: (a) 3D model, (b) Gap model, and (c) On-top model. The highlighted section in the T_{mrt} plot indicates measurement points located at the structure. Sample size: $n = 70$.

For Location 2, the plot and maps are presented in Figure 5.32. The Climate Bike moved from south-west to north-east. Before the underpass, all three models alternately underestimate, overestimate, and then again underestimate T_{mrt}. Inside the underpass, the Gap and 3D models show moderate deviations (over- and underestimation of 2–3°C, respectively), while the On-top model overestimates by up to 20°C. After the underpass, all models align well with observations, with a consistent overestimation of about 2°C.

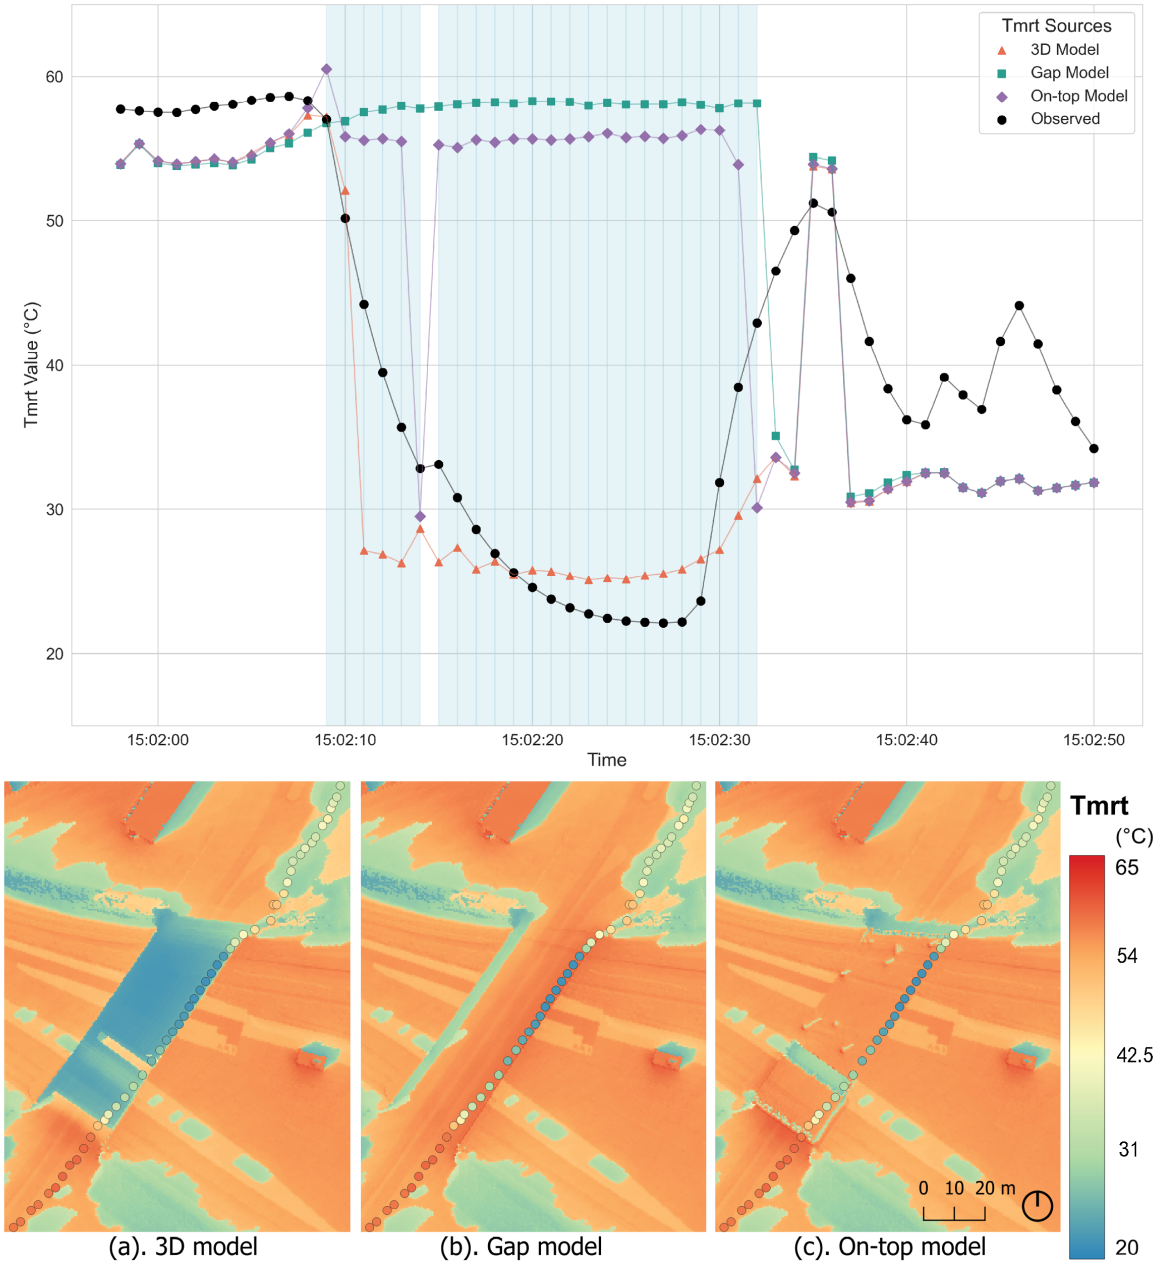

Figure 5.33: Plot of modeled and observed T_{mrt} over time, alongside maps of the three different model cases for Location 3: (a) 3D model, (b) Gap model, and (c) On-top model. The highlighted section in the T_{mrt} plot indicates measurement points located at the structure. Sample size: $n = 54$.

For Location 3, Figure 5.33 shows the plot and maps. The Climate Bike moved south to north during measurements. Before the underpass, all models underestimate T_{mrt} but stay within 3–4°C of observed values. Inside the underpass, the gap and on-top models significantly overestimate T_{mrt}, exceeding 30°C. The 3D model initially underestimates, then slightly overestimates near a turning point, likely due to radiation from the wall. A sharp transition appears in the 3D model, unlike the more gradual decrease in observations. After the underpass, all models continue to underestimate T_{mrt}, suggesting the actual shading from trees may have been less than the models assumed.

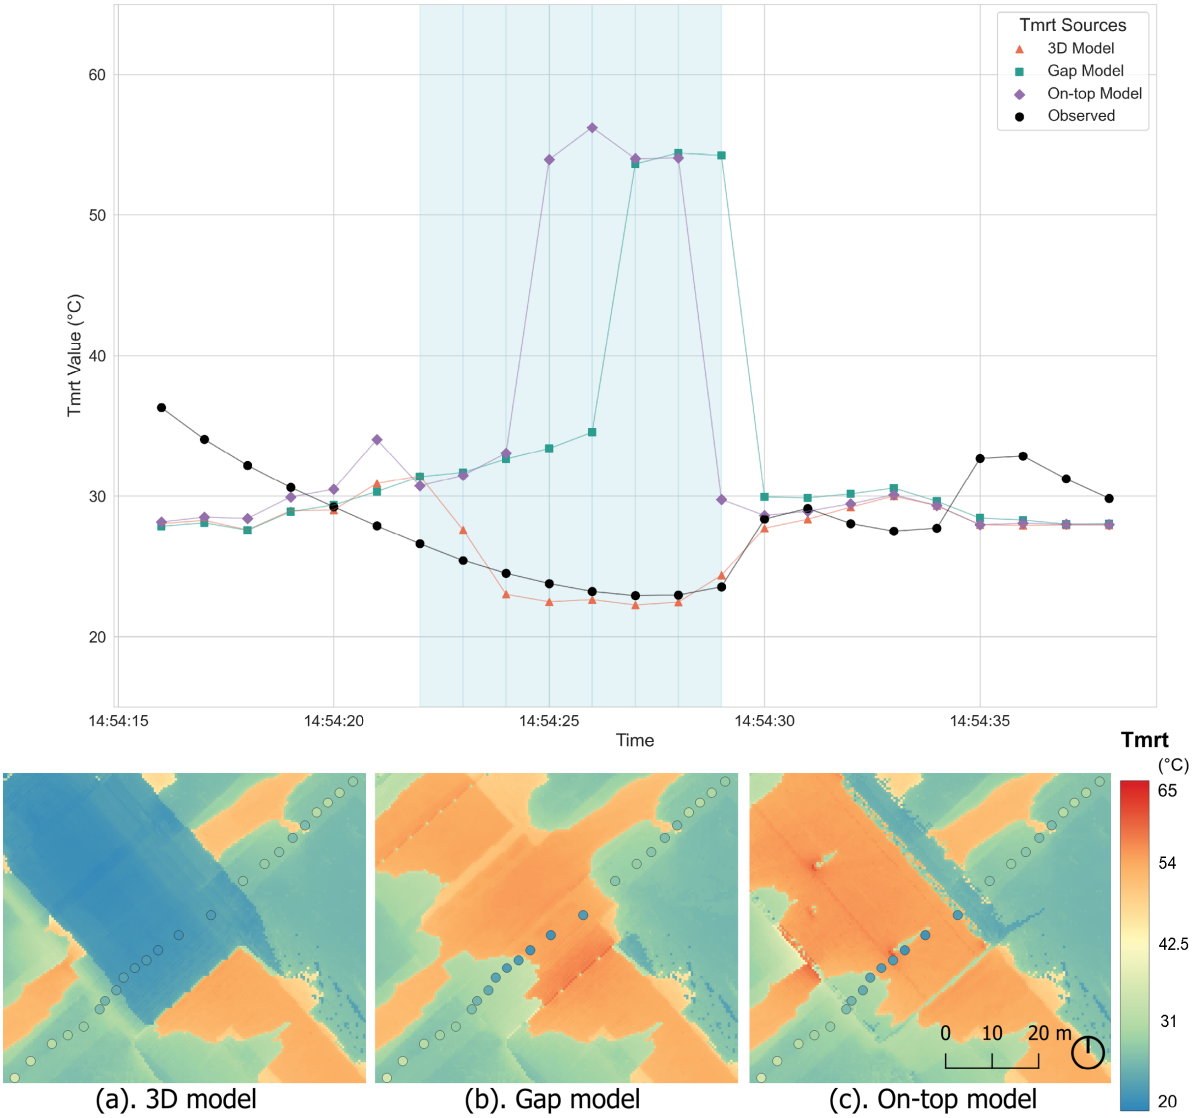

Figure 5.34: Plot of modeled and observed T_{mrt} over time, alongside maps of the three different model cases for Location 4: (a) 3D model, (b) Gap model, and (c) On-top model. The highlighted section in the T_{mrt} plot indicates measurement points located at the structure. Sample size: $n = 23$.

For Location 4, Figure 5.34 presents the corresponding plot and maps. The Climate Bike moved from south-east to north-west. Before the underpass, all models initially underestimate T_{mrt}, shifting to overestimation just before entry. Inside the underpass, the gap and on-top models overestimate T_{mrt} by more than 25°C, though this is partially mitigated near the start by shading from nearby trees. The 3D model shows slight overestimation initially but gradually aligns with observations, ending with a minor 1°C underestimation. After the underpass, all models stay relatively close to observed values, though a trend of underestimation appears toward the end of the route.

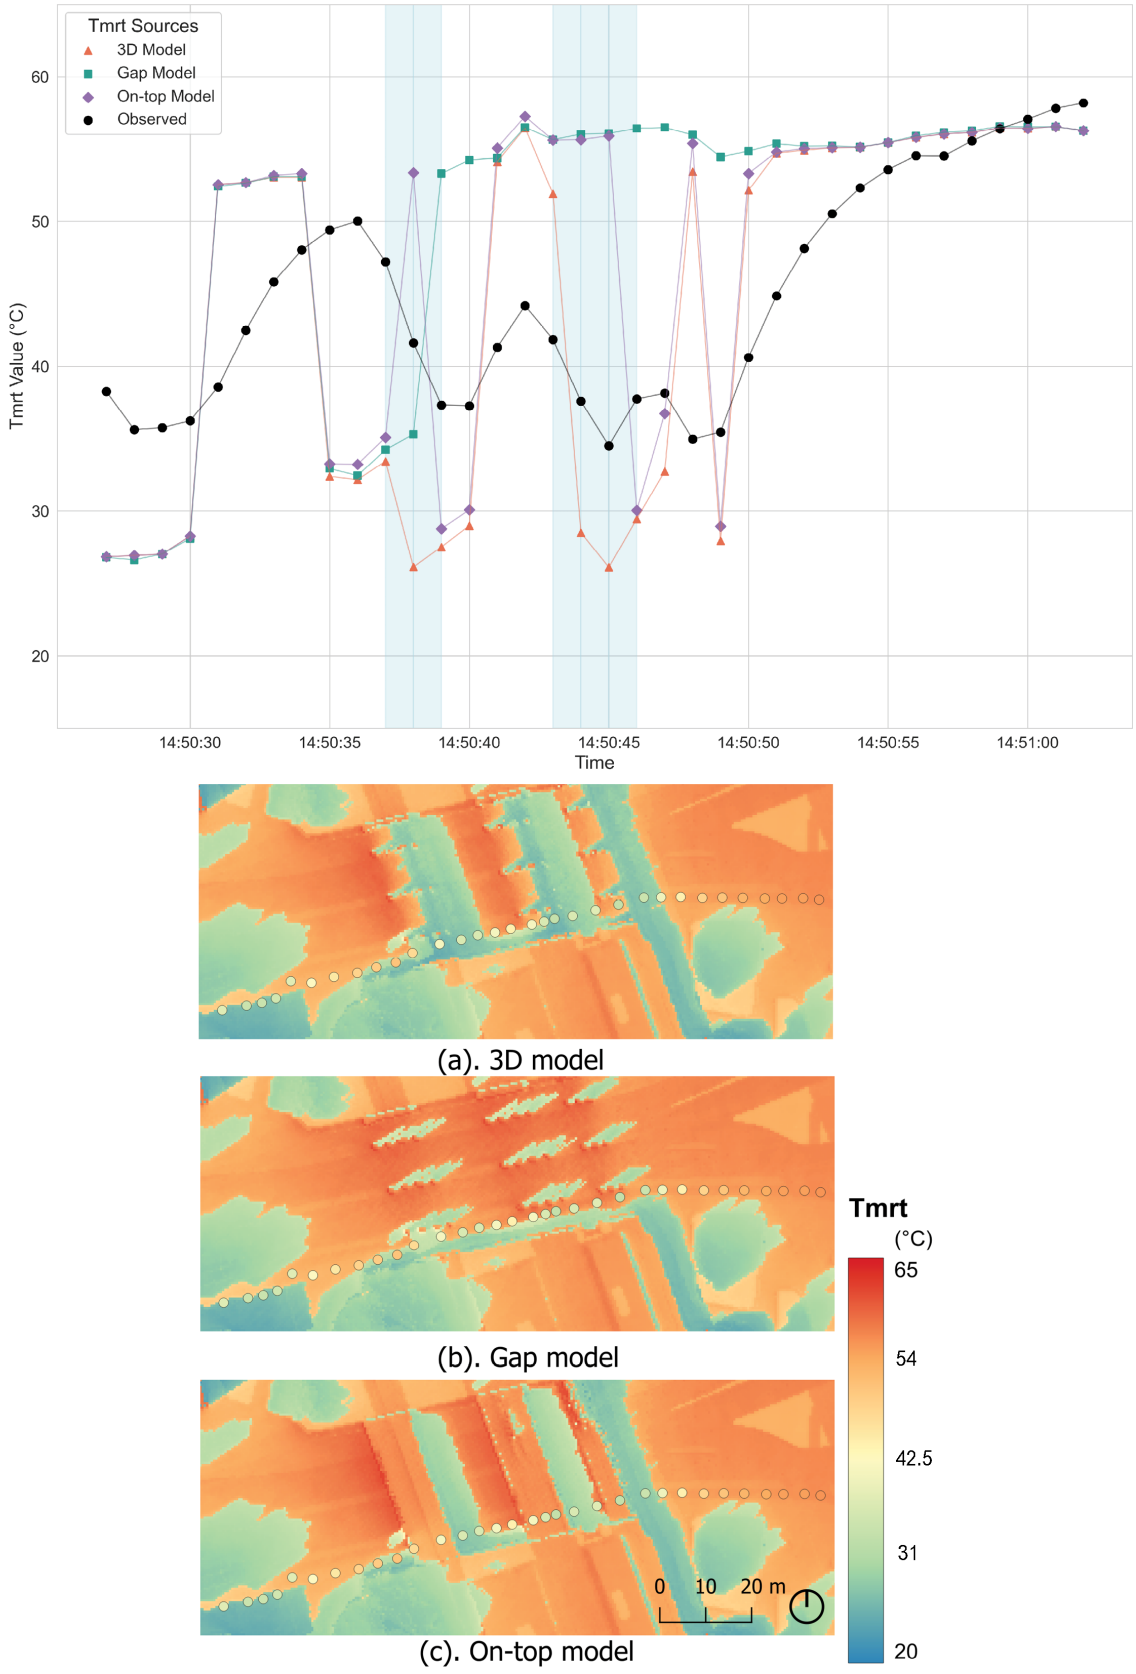

Figure 5.35: Plot of modeled and observed T_{mrt} over time, alongside maps of the three different model cases for Location 5: (a) 3D model, (b) Gap model, and (c) On-top model. The highlighted section in the T_{mrt} plot indicates measurement points located at the structure. Sample size: $n = 36$.

For Location 5, Figure 5.35 presents the plot and maps, with observations taken from west to east. The 3D model generally follows the trend of the observed T_{mrt}, but with notable over- and underestimations. The on-top model consistently overestimates and fails to capture the observed pattern. The gap model follows the trend only at certain points where shading coincidentally matches observations. Just before the underpass, a sharp drop in T_{mrt} appears in all models due to shading from a source present in the models but absent in reality.

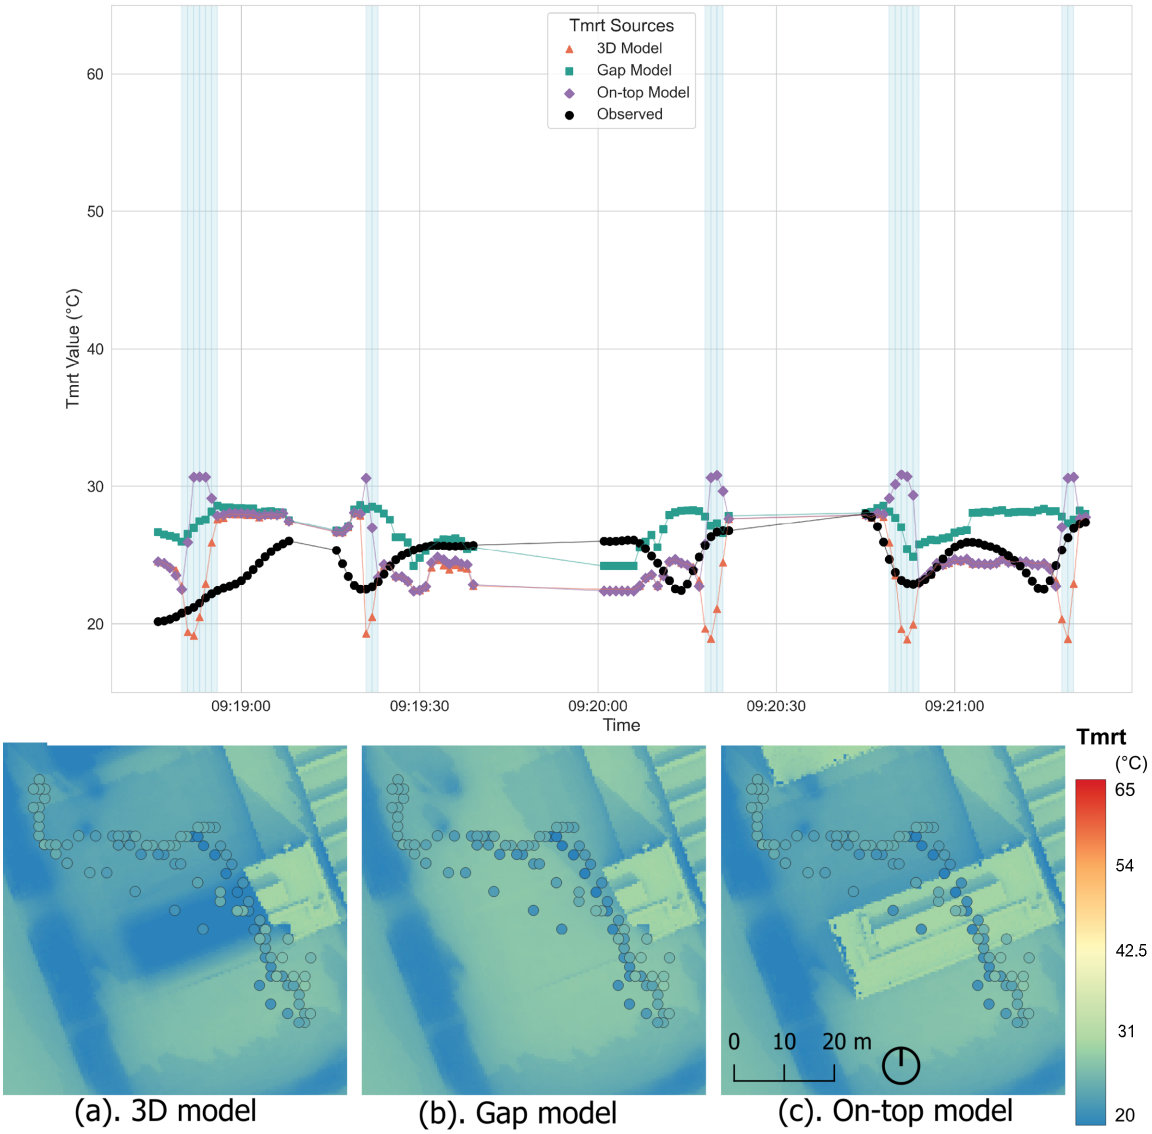

Figure 5.36: Plot of modeled and observed T_{mrt} over time, alongside maps of the three different model cases for Location 6: (a) 3D model, (b) Gap model, and (c) On-top model. The highlighted section in the T_{mrt} plot indicates measurement points located at the structure. Sample size: $n = 107$.

For Location 6, Figure 5.36 presents the plot and maps. The observations were made from south to north back to south. All models stay within 9°C of the observed T_{mrt}, mainly due to the absence of extreme temperatures. The models generally follow the trend of the observations, but only slightly.

The 3D model's performance varies notably across locations. At Location 3, it aligns closely with the observed T_{mrt} both before and beneath the gap, although with a steep slope. In contrast, Location 5 shows significant fluctuations, with the 3D model alternating between substantial over- and underestimations Overall, the 3D model generally captures T_{mrt} trends, but varies in accuracy.

When compared to the Gap and On-top models it becomes evident that the On-top model greatly

overestimates T_{mrt}, as this model places the computations above the structure rather than below. The Gap model's accuracy depends on solar angle and the orientation of surrounding structures. When properly aligned, for example at Location 2, the resulting shade covers the road surface, reducing T_{mrt} and bringing the model closer to observations. In less favorable orientations, such at Locations 1 and 4, the area is treated as open sky, and performs only slightly better, or even worse, than the On-top model.

In the 3D model, T_{mrt} is consistently lower beneath structures (highlighted in the plots), likely due to the absence of overhead surfaces that emit and reflect longwave radiation. Without this geometry, effects like trapped or reflected longwave radiation are missing, contributing to underestimation in these areas.

Overall Performance

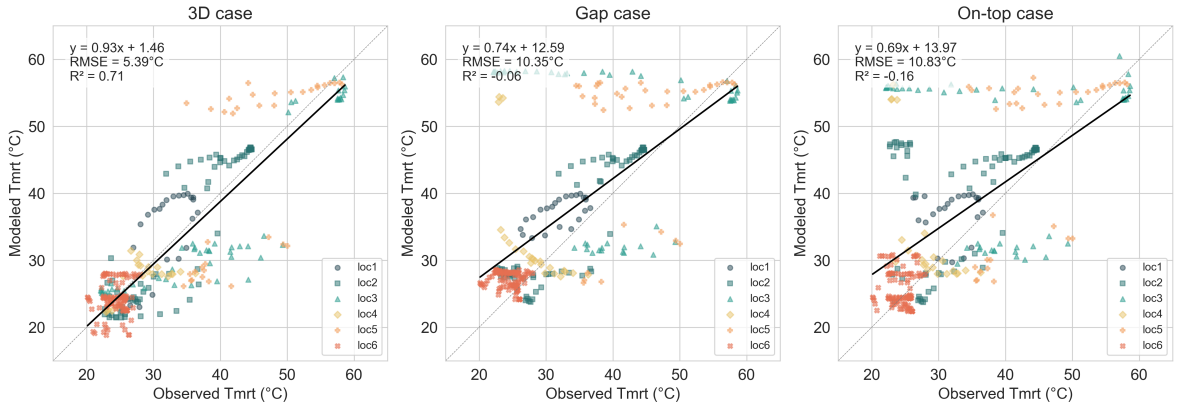

Figure 5.37: Overall performance of the three evaluated models.

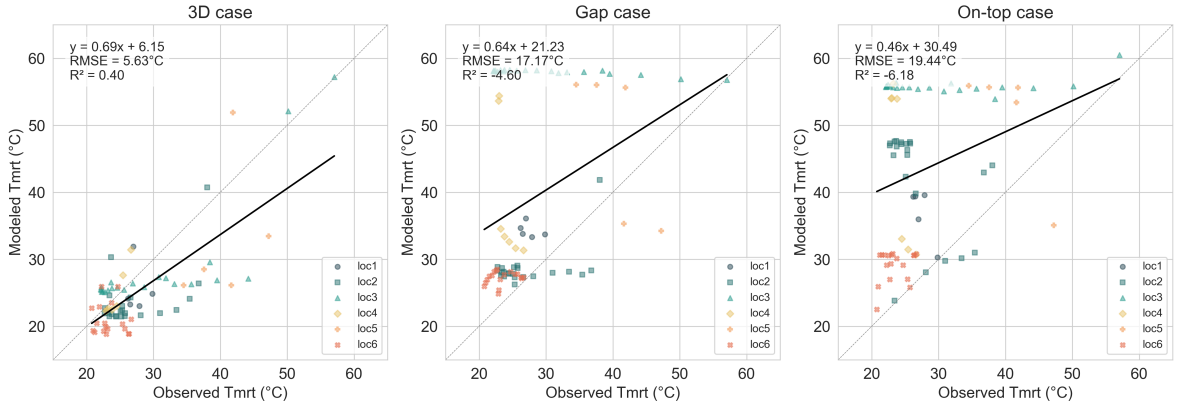

Figure 5.38: Performance of the three evaluated models underneath structures.

As shown in Figure 5.37 (n = 312), the Gap and On-top models yield high RMSE values (10.35°and 10.83°, respectively), while the 3D model performs better with an RMSE of 5.39°. The coefficient of determination R^2 also shows a notable improvement for the 3D case (0.71), compared to much lower values for the Gap (-0.06) and On-top (-0.16) models. When considering only the measurement points located directly beneath the structures (n = 79), the performance of the Gap and On-top models deteriorates significantly, with RMSE values increasing to 17.17°and 19.44°, and R^2 values dropping to -4.60 and -6.18, respectively. The 3D model performs substantially better in this subset, maintaining a lower RMSE of 5.63°and an R^2 of 0.40 (Figure 5.38).

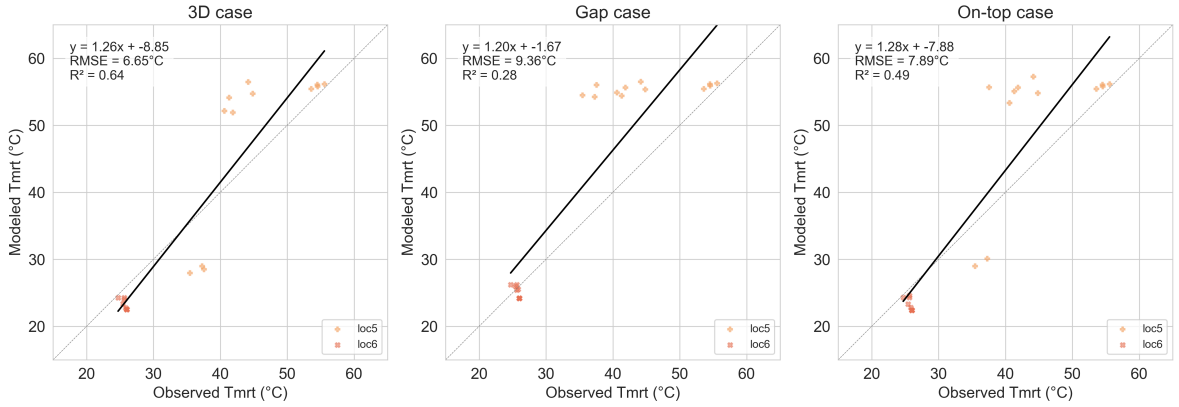

Figure 5.39: Performance from observations +- 5 degrees aligned with cardinal directions

Figure 5.39 shows the model performance for observations within 5° of cardinal alignment, where modeled radiation is expected to align most closely with measured values. Only two of the six locations contain such observations, resulting in just 21 out of 312 total observed points being included. This severely limits the strength of conclusions that can be drawn from this subset. For these points, the performance differences between the models are smaller, and again the 3D model performs best.

Given the limitations of the available validation dataset, the best conclusion that can currently be drawn is that the 3D implementation offers a closer approximation of the observed T_{mrt} compared to the other models. However, the precise locations of the measurements are unknown, and the number of points with correct cardinal orientation is insufficient. As such, it is not possible to fully quantify the performance of the 3D model. Moreover, there appears to be a systematic underestimation of T_{mrt} beneath structures, indicating that the omission of overhead surfaces has a significant impact on the results.

Identifying Sources of Error and Opportunities for Improvement

It is important to identify the reasons for discrepancies between the modeled and measured T_{mrt}, in order to improve future implementations.

- **GPS inaccuracies:** GPS data precision presents a significant limitation. The measurement location does not align perfectly with the cell of the modeled T_{mrt}. However, when examining neighboring values on the map, temperature differences are typically not substantial, unless a measurement in the shade is matched to a cell not in shade and vica versa.
- **Abrupt transitions between sun and shade:** The model shows discrete cell-level switches from sunlit to shaded areas, while in reality such transitions are often more gradual. This can be observed in plots from most locations, where the observed data changes temperature smoother.
- **Automatic input generation limitations:** Two key weaknesses are identified:
  - *Tree detection:* As discussed in Section 4.1.2, the NDVI-based method can result in missed trees or misclassification of non-tree objects as trees, leading to either underestimated or falsely added shade in the SOLWEIG and SVF calculations.
  - *Material classification:* The limited material options given with SOLWEIG do not accurately reflect Dutch street surfaces. For example, concrete tiles are common but approximated as cobblestone in the model, affecting surface temperature estimates.
- **Shading algorithm**: The shading algorithm assumes that a cell is shaded if the shadow volume intersects any part of the DSM (Section 4.1.1). However, when the vertical resolution of the shadow volumes (z-steps) is small, such as during low solar angles, relatively minor features in the DSM, like street curbs or small protrusions, may cast disproportionately large shadows. While these features have little impact on the actual T_{mrt} experienced by pedestrians, the model interprets them as casting shade at pedestrian height. As a result, the modeled T_{mrt} values can be underestimated in areas where such shading is present.

- **Weather conditions:** Measurements in Locations 1, 5, and 6 were taken on cloudy days. SOLWEIG is documented to perform best under clear-sky conditions. Figure 5.40 shows the radiance profile from the climate input data. Rather than the smooth parabolic pattern expected on clear days, the profile is jagged, reflecting frequent cloud cover and resulting in unpredictable irradiance fluctuations.

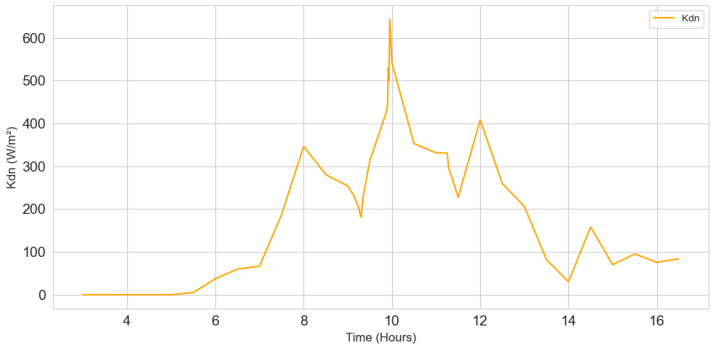

Figure 5.40: Incoming global radiation for 12 September

- **Human geometry modeling:** The model uses a cylindrical representation of the human body, as configured in SOLWEIG, but the observed data is translated to T_{mrt} using the box representation. Not computed the same way/ Figure 5.41 compares the cylindrical and box models under different solar angles (summer at high solar elevation and winter at high latitude). Although the differences are likely on the order of 0.1-2°, they are minor compared to the error found in the T_{mrt} simulation.

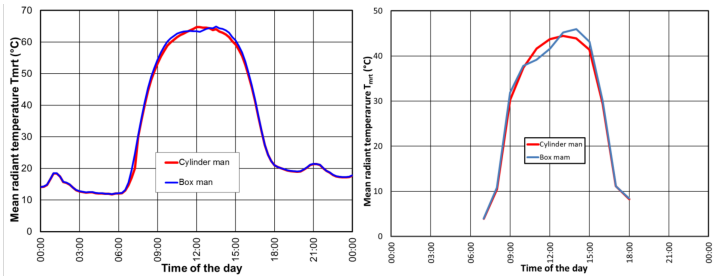

Figure 5.41: T$_{m}rt$ based on six-directional monitoring of short- and longwave radiation. Calculation of T$_{m}rt$ according to the original formula by Höppe (box man) and with the new methodology (cylinder man). Left: Ouagadougou, Burkina Faso on 10th of December 2007 and Right: Gothenburgh, Sweden on 11th of October, 2005.
Source: Holmer et al. (2015), p. 3

- **3D model inaccuracies:** The height and dimensions of the modeled gap are approximations; without exact measurements, these could introduce additional errors.
- **Uneven data distribution across locations:** Certain locations, such as Location 6, have a significantly higher number of data points compared to others, like Location 1. Since Location 6 primarily features lower temperatures that align more closely with the model (due to the absence of extreme values), this imbalance in data distribution could create a skewed perception of the model's performance.
- **Missing supporting structures:** Only for Location 5 were the columns supporting the bridge included in the model. The absence of such elements in other locations may reduce model accuracy.

5.3. Urban Test Case with the Developed Tool

Typology	Mean				Median				Min				Max			
	M	A	E	O	M	A	E	O	M	A	E	O	M	A	E	O
bloemkool	-8.16	-8.59	-3.49	-6.75	-10.87	-8.93	-5.32	-8.37	-0.04	-0.06	-0.17	-0.09	-0.61	-0.59	-0.29	-0.50
historisch	-4.21	-5.15	-1.04	-3.47	-5.79	-7.18	-0.53	-4.50	-0.05	-0.07	-0.18	-0.10	-0.17	-0.46	-0.11	-0.25
stedelijk	-6.82	-8.24	-1.50	-5.52	-9.70	-11.45	-0.64	-7.26	-0.06	-0.09	-0.27	-0.14	-0.44	-0.97	-0.30	-0.57
tuindorp	-6.23	-6.90	-2.61	-5.24	-8.05	-7.91	-3.00	-6.32	-0.02	-0.02	-0.02	-0.02	-0.61	-0.62	-0.62	-0.62
vinex	-6.24	-6.50	-2.76	-5.17	-8.30	-6.32	-4.31	-6.31	-0.01	-0.02	-0.14	-0.06	-0.51	-0.52	-0.25	-0.42
volkswijk	-5.32	-6.28	-1.56	-4.39	-7.49	-8.22	-1.12	-5.61	-0.07	-0.09	-0.18	-0.11	-0.82	-0.74	-0.98	-0.85

Table 5.10: Average Tmrt difference for an 'average' day (with trees minus without trees) per typology, split by statistic and time of day.

Typology	Mean				Median				Min				Max			
	M	A	E	O	M	A	E	O	M	A	E	O	M	A	E	O
bloemkool	-5.65	-9.93	-6.22	-7.27	-6.16	-11.28	-8.79	-8.74	-0.03	-0.07	-0.13	-0.08	-0.49	-0.59	-0.54	-0.54
historisch	-3.23	-5.91	-1.54	-3.56	-4.89	-8.20	-0.69	-4.59	-0.04	-0.09	-0.17	-0.10	-0.25	-0.56	-0.14	-0.32
stedelijk	-5.21	-9.39	-3.22	-5.94	-7.92	-14.01	-1.16	-7.69	-0.05	-0.11	-0.21	-0.13	-0.46	-0.74	-0.95	-0.72
tuindorp	-4.47	-8.02	-4.55	-5.68	-5.47	-10.63	-5.61	-7.24	-0.02	-0.03	-0.00	-0.02	-0.72	-0.50	-1.14	-0.79
vinex	-4.29	-7.52	-4.99	-5.60	-4.44	-7.91	-7.16	-6.50	-0.01	-0.02	-0.09	-0.04	-0.51	-0.55	-0.74	-0.60
volkswijk	-3.97	-7.23	-2.87	-4.69	-5.49	-10.30	-2.09	-5.96	-0.05	-0.11	-0.13	-0.10	-0.77	-0.75	-1.77	-1.10

Table 5.11: Average Tmrt difference for an 'extreme" day (with trees minus without trees) per typology, split by statistic and time of day.

Comparative plots of the statistics across typologies and time periods, for both an 'average' and an 'extreme' 30+ day, are provided in Appendix A.6. Figure 5.42 shows the percentage of area within different T_{mrt} ranges for each typology. The only difference between the 'average' and 'extreme' 30+ day scenarios lies in the size of the shadows. On the extreme day, the sun reaches its highest point in the sky, resulting in the shortest shadows of the year and therefore more areas experiencing high T_{mrt} values.

Without trees, the typologies with the lowest T_{mrt} at pedestrian height are *historisch* and *stedelijk*, with mean T_{mrt} values around 50°C in the afternoon on an average day, and about 55°C on an extreme day. This can be explained by their compact urban structure with narrow streets and some enclosed plazas. In contrast, more spacious suburban types (*bloemkool*, *vinex*, and *tuindorp*) have the highest means, reaching around 62°C on an extreme day afternoon and around 59°C on an average day afternoon. This trend is also reflected in the temperature bin distribution: historisch and stedelijk have a more balanced spread along the bins, with slightly larger in 55°C+ , while in bloemkool and vinex, nearly 90% of the area is in the 50°C+ range during the afternoon.

The inclusion of existing trees significantly reduces T_{mrt} in all typologies. Notably, *bloemkool* sees a mean afternoon reduction of 10°C and an overall reduction of about 8°C. Other typologies show an overall reduction between 4°C and 6°C. The smallest reduction is found in the *historisch* typology, its compact urban form leaves less room for trees and already limits high T_{mrt} areas. Median T_{mrt} reductions are typically 1°C to 1.5°C greater than mean reductions, indicating a strong shift in the central tendency. Large shifts are also visible in the area percentage per T_{mrt} bin.

The maximum and minimum T_{mrt} values across typologies remain roughly the same, as expected given that the same weather file was used across all simulations.

All these statistics are generated automatically by SOLFD, and its output maps allow for visual identification of hotter zones.This demonstrates that SOLFD can be used to analyze the effects of urban design interventions, such as tree plantation, across different Dutch urban typologies statistically (using mean, median, min, max, and bin distributions). Although the output maps also enable visual comparison, such comparisons become impractical with the large number of locations studied.

Interestingly, the typologies considered to have lower heat stress levels according to the Klimaateffectatlas perform worst when trees are absent. With the addition of vegetation, the *bloemkool* typology in particular approaches the performance of *historisch* and *stedelijk*, highlighting the important role of urban vegetation in reducing radiant heat exposure. However, this difference can also be explained by the fact tha only T_{mrt} is considered.

It seems that for the Klimaateffectatlas, the heat vulnerability of neighbourhood typologies was determined by the strength of their Urban Heat Island effect, which is typically measured at night using air temperature (Klimaatadaptatie Nederland, 2023). These effects are strongest at night, when dense urban structures release the heat they stored during the day. In contrast, this study focuses on daytime conditions and T_{mrt}, which is more closely linked to surface and radiant heat and varies significantly under direct solar exposure. In this context, typologies like *historisch* and *stedelijk* perform better even without trees, as their compact design naturally provides more shade during the day. However, the same dense materials and structures also trap heat, contributing to higher nighttime air temperatures. This explains their less favorable classification in the Klimaateffectatlas.

It is therefore important to note that analysis using SOLFD focuses solely on T_{mrt}; additional factors should be considered for a complete assessment of OTC.

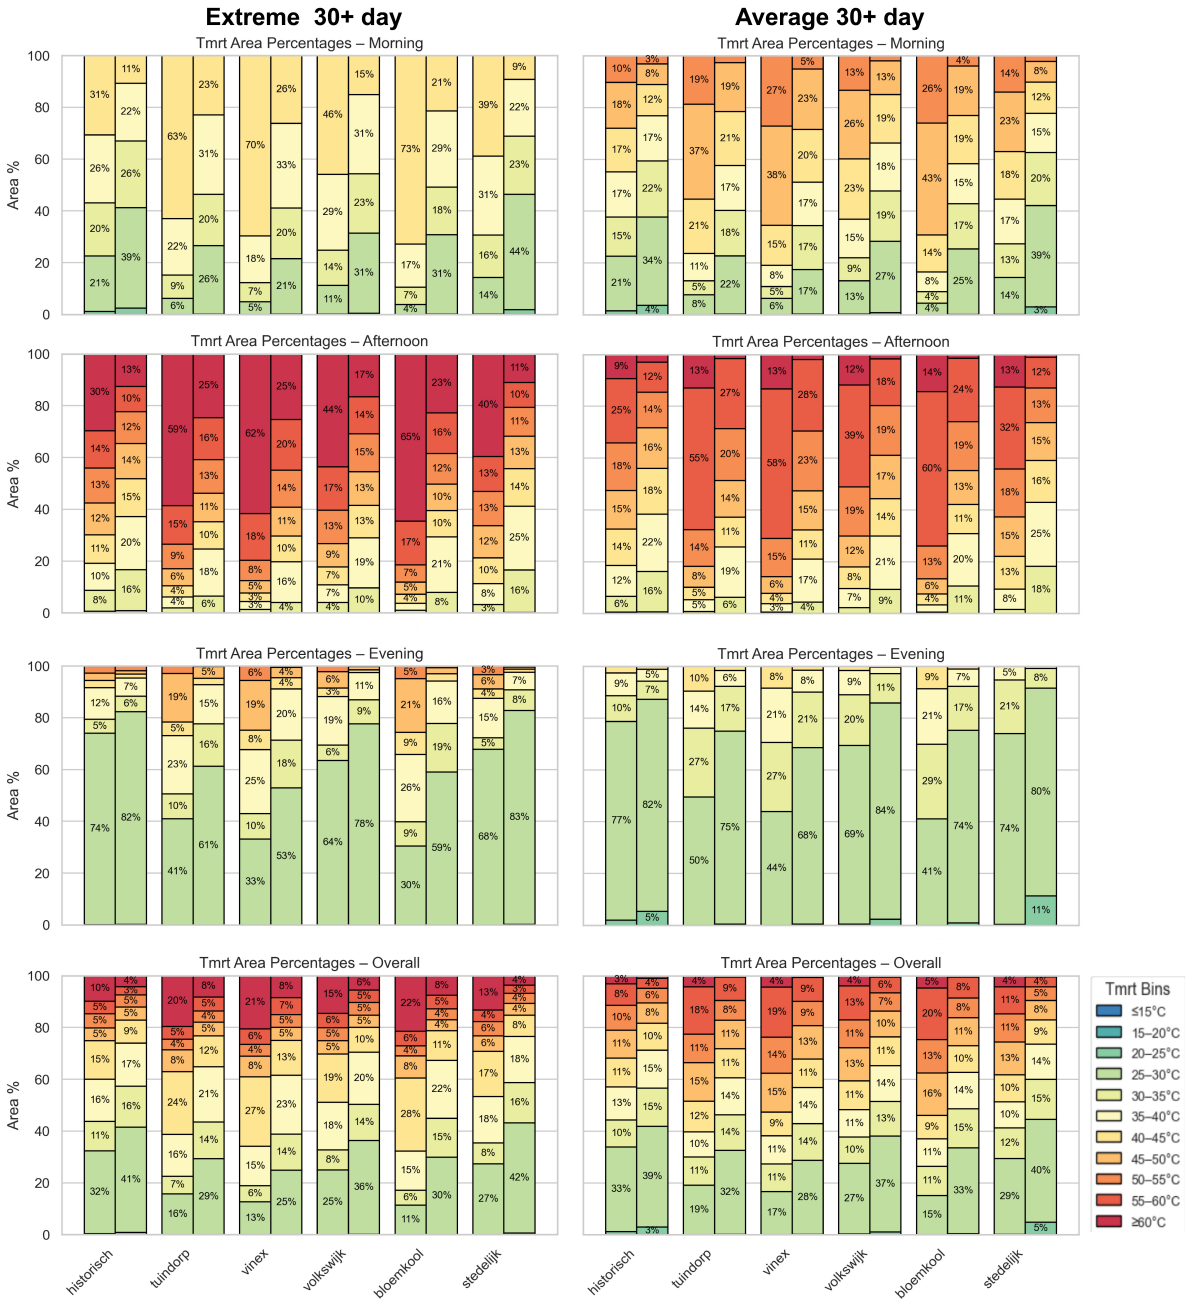

Figure 5.42: T_{mrt} area percentages for the six typologies during different times of the day (morning, afternoon, evening), and overall, for both cases with trees and without trees, on an 'average' and 'extreme' day. With Trees is always the right bar in each pair.

6

Conclusions

This chapter begins with the main conclusions drawn from the research, evaluating if the research questions have been answered. This is followed by a discussion of the results and the contributions and the limitations of the study. Finally, directions for future work are outlined.

6.1. Conclusion

This part revisits the research questions for this thesis introduced in Section 1.2. For each question, a brief response is given, supported by work from the previous chapters.

> *1. How can a tool framework be developed for visualizing current urban climate conditions and their effect on outdoor thermal comfort at the neighborhood scale in the Netherlands using SOLWEIG, and how can design alternatives be integrated and compared within this tool?*

By analyzing the current capabilities of SOLWEIG, several limitations were identified and addressed through a structured step-by-step approach. SOLWEIG requires manually generated context data, including maps of landcover, DSM, DTM, and a CHM, which are not readily available. Furthermore, SOLWEIG assumes the input maps fully represent the intended design scenario, with no built-in support for modifying or editing existing conditions. For meaningful design comparisons, the tool should allow integration of alternative configurations, such as adding, editing, or removing buildings, trees, and land cover elements. In its current form, the model also only allows 2D-arrays as input. This does not allow representation of overhangs or gaps within building blocks, features that can significantly influence microclimatic conditions through additional shading.

Additionally, the model currently outputs only hourly maps of T_{mrt}, along with shadow and radiation maps used in its calculation. Since T_{mrt} represent the net effect of radiant heat exchange in isolation, its values can be difficult to interpret in terms of actual human thermal comfort. Furthermore, complete simulations are time-consuming: generating the required SVF and running SOLWEIG for an urban scene with vegetation, can together take up to 45 minutes for an array size of 1000×1000 cells, and approximately 3.5 hours for a 2000×2000. This limits the model's practicality for iterative design analysis. These challenges are addressed through the following sub-questions:

> *1a. To what extent can the 2.5D algorithm of SOLWEIG be transformed into a real volumetric 3d algorithm?*

Without altering the core radiation physics, SOLWEIG can be extended to incorporate 3D data for shadow casting and SVF calculation at pedestrian height. This is achieved by structuring the spatial context into a layered 2D array, with the following structure:

- **Layer 0 (base layer)**: Contains the height of the ground surface and all buildings up to the first point where where they open directly to open air.
- **Void layer(s)**: Represent the vertical extent of the open-air void using true height values.
- **Building layer(s)**: Contain building heights above voids that can cast shadows.

This pattern can be repeated to accommodate multiple stacked elements. For shadow casting, a cell is considered to be shaded if, at a given iteration step, the building layer is higher than the untranslated base layer, while the void layer remains below it. This approach is adapted from tree shading in SOLWEIG, where the trunk of a tree is treated as a vertical gap.

Currently, this 3D representation is only used during calculations involving shade casting. Subsequent energy balance computations are still based solely on the base layer, thereby excluding the radiative effects of additional surfaces and vertical structures above pedestrian height.

1b. How can the input data pipeline be automated to extract and create the Dutch location data required for SOLWEIG?

The entire data pipeline has been automated to generate all required inputs for SOLWEIG using only the bounding box of a location. This is achieved through the integration of openly available Dutch government datasets.

DTMs and DSMs are sourced from the AHN WCS database. Point clouds are extracted from GeoTiles and automatically downloaded using the input bounds. Gaps in the AHN DTM and DSM are filled using Laplace interpolation. To isolate building heights, 3DBAG building footprints are used to mask and extract values from the DSM, which are then placed onto the DTM to generate a building-only DSM. Vegetation points are extracted from the 'unclassified' category in the AHN point cloud using NDVI filtering. These points are rasterized with Laplace interpolation, and a normalized CHM is produced by subtracting the DTM. Post-processing steps include height filtering and the application of a median filter to remove noise. Land cover information is retrieved from the BGT dataset, which provides terrain, road, and water polygons. These are classified according to the material categories used by SOLWEIG. For wall aspect calculation, the previously used linear kernel method is replaced with a faster Sobel filter, improving performance. Additionally, synthetic meteorological datasets have been generated for three representative summer conditions based on 20 years of historical data:

- 20–25°C (mild summer days)
- 25–30°C (warm summer days)
- >30°C (hot/tropical days)

1c. How can urban design proposals be integrated as input for SOLWEIG?

User input can be integrated into SOLWEIG by enabling modifications to the required input rasters, particularly through the insertion or alteration of buildings, vegetation, and surface classifications. This is achieved by maintaining unique IDs for individual buildings and trees, allowing users to selectively remove, modify, or reinsert elements using polygonal masks. Custom 3D designs created in CAD software and exported as .OBJ files can be rasterized and incorporated into the model. This is done through point-in-triangle intersection and barycentric interpolation, which allows for extraction of height values at raster cell centers from 3D geometries.

Additionally, new trees can be inserted into the CHM using parameterized inputs such as total height, trunk height, canopy radius, and shape, supported by a newly created species-specific database of common Dutch trees. Land cover classifications can be added by providing surface geometries as .OBJ files with the desired class, which are then converted into raster format and integrated into the land cover. By providing an export of the contextual environment at the correct scale, users can position new designs accurately within the modeled area.

1d. What additional outputs or post-processing steps are required after SOLWEIG to compare and rank different urban design options based on their microclimatic impact and resulting outdoor thermal comfort

To improve the interpretability of the results, the output T_{mrt} maps are supplemented with PET maps. These PET maps help identify actual heat stress classifications by incorporating the influence of additional factors beyond T_{mrt}, namely relative humidity, air temperature, and wind speed. The PET values are derived from the T_{mrt} maps using a precomputed lookup table applied on a per-pixel basis, while the other parameters are provided as global inputs. Air temperature and relative humidity are user-chosen or taken from the meteorological dataset, and wind speed is chosen from three discrete levels: windstill

(0.1 m/s), light breeze (2.0 m/s), and moderate breeze (6.0 m/s). However, since wind speed varies locally, this simplification limits accuracy and indicates the need for further refinement. Therefore, PET maps should be interpreted as a complementary extension to T_{mrt} maps, rather than a standalone metric.

To facilitate better comparison of urban design options, the results for T_{mrt} and PET are grouped and presented as follows:

- Values per hour
- Averages computed for each time group during the day (morning, midday, afternoon). This allows an understanding of whether a design performs best during the hottest time of day or at a specific time of day.

Additionally, statistical outputs are provided for a quick overview of the results, assisting comparison between design proposals:

- Mean, median, minimum, and maximum values for T_{mrt} and PET.
- T_{mrt} and PET temperatures are divided into bins, and the percentage of the area, as well as the total square meters, falling within each bin is calculated for each design.

1e.To what extent can SOLWEIG be optimized for computation time?

The primary computational bottleneck in SOLWEIG was found in functions related to shade calculations, which depend heavily on large-scale array operations. Since GPUs are specifically designed to handle such operations efficiently through parallel processing, significant performance gains were achieved by offloading these computations to the GPU. To maximize the benefits of GPU acceleration, the data transfers between CPU and GPU were minimized by initializing arrays directly in GPU memory and performing all subsequent operations on the GPU.

Further efficiency was obtained by reducing computational redundancy, avoiding unnecessary loops, and minimizing the creation of intermediate arrays. Depending on the array size and hardware used, the optimized implementation achieved speedups ranging from 10× to 50× compared to the original code. For array sizes up to 1500×1500 cells, this resulted in a reduction of computation times from 10 minutes to 1.5 hours down to just 1–5 minutes, depending on the exact array size. As an example, on a computer with a powerful GPU, an urban scene with a 2000×2000 cell array including CHM data now completes computation in approximately 3.5 minutes, whereas originally it would have taken over 3 hours. These improvements substantially enhance the model's suitability for iterative and comparative design workflows and meet the previously defined goal of achieving runtimes within 5–10 minutes.

2. What does existing research report about the accuracy of SOLWEIG in urban climate simulations, and how well do the newly implemented 3D modifications perform when validated against sensor data?

Both the developers of SOLWEIG and independent researchers have evaluated the model's accuracy. For the most recent version, validation results showed an RMSE of 4.6°C and $R^2 = 0.9$ for T_{mrt}. However, the model still tends to underestimate longwave radiation in shaded areas and overestimate it under direct sunlight, partly due to the assumption of uniform wall temperatures. Independent studies report comparable results, with RMSE values ranging from 2.22 to 7.67°C and R^2 values between 0.49 and 0.95. Model performance varies with resolution, input data quality, and sky conditions.

In general, SOLWEIG performs best under clear-sky, daytime conditions. The main sources of error are related to shade misclassification and simplified assumptions regarding longwave radiation fluxes.

The newly implemented 3D modifications in SOLFD appear to improve the model's performance in areas beneath overhead structures, particularly in situations involving partial shading. However, due to uncertainties in the validation dataset, the extent of this improvement cannot be precisely quantified. The 3D implementation provides a closer approximation of observed T_{mrt} than the simplified models that either remove the structure and therefore assume a full gap, or place the measurement on top of it. These simplified models yield high RMSE values (10.35°C and 10.83°C, respectively), while the 3D model performs moderately better with an RMSE of 5.9°C. The coefficient of determination R^2 also improves in the 3D case (0.71), compared to much lower values for the Gap (-0.06) and On-top (-0.16) models.

When isolating only the measurement points located directly beneath the structures, the performance of the simplified models declines sharply, with RMSE values increasing to 17.17°C and 19.44°C, and R^2 values dropping to -4.60 and -6.18. The 3D model again performs better for this subset, with an RMSE of 5.63°C, and $R^2 = 0.40$. However, due to uncertainty in the precise measurement locations and insufficient data with correct cardinal orientation, it is not possible to fully quantify the improvement. Moreover, the results suggest a systematic underestimation of T_{mrt} beneath structures, likely due to the omission of overhead surface effects.

3. How can the tool be used to analyze the effects of urban design proposals on microclimatic conditions and outdoor thermal comfort in various Dutch urban typologies?

SOLFD enables quantitative and spatial analysis of microclimatic conditions by automatically generating maps and statistical summaries of T_{mrt} and PET. This capability supports the evaluation of urban design proposals across different Dutch urban typologies, as demonstrated in the case study.

In the case study comparing six Dutch typologies, SOLFD was applied to simulate conditions both with and without existing tree vegetation. Results showed that without trees, compact typologies such as *historisch* and *stedelijk* exhibited lower T_{mrt} values due to their dense urban morphology. Conversely, more open suburban typologies like *bloemkool* and *vinex* experienced higher radiant heat exposure. When tree vegetation was included, SOLFD quantified significant reductions in T_{mrt} across all typologies, with the greatest cooling effect observed in suburban areas.

By providing statistical outputs such as mean, median, minimum, maximum, and area distribution of T_{mrt} and PET values, SOLFD enables objective assessment of the impact of design elements on OTC at pedestrian height, with a main focus on the influence of T_{mrt}. While other relevant environmental parameters affecting OTC are not currently incorporated and would be needed for a more comprehensive evaluation, SOLFD remains useful for evidence-based assessment of urban design proposals by facilitating side-by-side scenario comparisons across diverse urban typologies.

6.2. Discussion

The main question explored in this discussion is whether the objective of the research has been achieved: is SOLFD usable for urban designers? More specifically, does it meet the needs outlined in the literature? According to the criteria formulated by architects in Weytjens and Verbeeck (2010), tools intended for use during design should be simple and intuitive, require minimal data input, allow quick scenario testing, provide clear feedback on design changes, produce visual and interpretable results, be both flexible and transparent, and adaptable for early-stage design.

- **Simplicity and an intuitive interface**: Although a full working interface was not developed in this research, the framework was built with ease of use in mind. In principle, a designer would only need to select a bounding box (preferably via an interactive map), upload their design, and make edits by clicking. This interaction model is intentionally minimal and would likely be intuitive, though usability cannot be fully assessed without an actual interface.
- **Minimal data input**: SOLFD requires only building masses from the user's design and a bounding box to define the area of interest. Users select a representative day for the simulation, and all additional inputs, such as terrain height, vegetation, and weather data, are automatically generated. This reduces data input significantly.
- **Ability to quickly test scenarios**: According to Weytjens and Verbeeck (2010), acceptable evaluation times range from 5 minutes to 2 hours. Considering this study was conducted in 2009, expectations for interactivity may now be more demanding. SOLFD can process updates within 2 to 5 minutes, which supports relatively quick scenario testing, although real-time performance is not achieved.
- **Feedback on the impact of design parameters**: This remains a limitation of the current implementation. While SOLFD generates quantitative results, these are not yet directly tied back to specific design changes in a way that supports informed decision-making. Integrating feedback mechanisms to relate outcomes to design inputs would improve usefulness.
- **Visual and easily interpretable results**: SOLFD outputs T_{mrt} maps, statistics such as area distribution across temperature bins, and PET-based heat stress categories. These visualizations support straightforward interpretation of results.

- **Flexibility and transparency**: By automating data acquisition and processing, SOLFD reduces complexity for the end user, but the pipeline remains transparent. Each step, from geometry processing to rasterization and simulation, is openly available on the GitHub, allowing for debugging or customization by advanced users if needed.
- **Adaptability for early-stage design**: SOLFD can operate on simple volumetric representations of buildings and does not require detailed architectural models. This makes it well-suited for early design phases where the focus lies on massing.

Compared to the design tool LBT, described in Section 2.3, SOLFD offers fewer functionalities for OTC calculation. It focuses solely on calculating T_{mrt}, whereas LBT provide a wider range of OTC parameters and even include a form of wind simulation. Additionally, LBT can be integrated with optimization plug-ins, enabling parametric design workflows for automated design optimization. Its results are visualized in 3D, and it has a direct connection with CAD software, eliminating the need for manual import or export between tools.

On the other hand, SOLFD requires significantly less input data, offering an automated workflow that allows users to obtain results without extensive training. It also easily incorporates terrain, building, and vegetation heights automatically constructed from Dutch datasets. Notably, the inclusion of trees in SOLFD represents an upgrade over LBT, where users would need to manually create 3D models of vegetation. Furthermore, SOLFD efficiently handles large geographic areas.

A comprehensive side-by-side comparison of LBT and SOLFD, including performance and usability assessments, is necessary to draw further conclusions.

Contributions

The contributions of this thesis lie primarily in the methodological and practical extensions made to SOLWEIG, enhancing both its capabilities and applicability for design support.

- **Progress towards a truly 3D implementation of SOLWEIG:** This thesis introduces a novel layered-array logic that enables SOLWEIG to better represent three-dimensional urban geometry. By using layered 2D arrays to model height layers and developing a method to convert rasterized OBJ models, including all intersection heights, into correctly layered DSM inputs, the model can now account for architectural elements such as building overhangs and ground-level passages, at least for shadow casting and SVF calculations.
- **Substantial acceleration through GPU integration:** The computational performance of SOLWEIG has been significantly improved by transferring core array operations to the GPU. Parts of the code were adapted to support GPU execution. As a result, simulations of SVF and T_{mrt} are now performed 10 to 50 times faster, depending on the array size and maximum building height. This improvement makes it feasible to analyze larger areas at higher spatial resolutions within practical time constraints.
- **Automated pipeline for Dutch geospatial datasets:** A data pipeline was developed to automatically collect, process, and prepare spatial datasets for any location in the Netherlands using only a bounding box as input. The system retrieves DTM, DSM, CHM, and land cover data, processes them into standardized raster formats, and produces all inputs required by SOLWEIG. The workflow includes methods for filling gaps in elevation data and constructing canopy height models. Furthermore, the land cover classifications can be adjusted to suit use cases beyond the limited set supported by the original SOLWEIG implementation.
- **Support for modified and new urban scenarios:** The workflow has been extended to allow users to make direct modifications to the existing urban context. Using the given DTM, DSM, CHM, and land cover rasters as a baseline, users can insert new trees or import complete three-dimensional urban design proposals. These additions are automatically rasterized and converted into the layered input arrays required by SOLFD, which means that design proposals can be evaluated and compared without the need for extensive manual preprocessing.
- **Bringing SOLWEIG closer to practical use in urban design:** This work advances SOLWEIG toward practical application by urban designers. Although SOLFD is not yet equipped with a GUI, the included functions, for example the automated acquisition of spatial context data, already provide

valuable capabilities. With further development into a complete tool, SOLFD could meaningfully support microclimatic and thermal comfort considerations in urban design workflows.

Limitations

Despite these contributions, the thesis has several limitations that highlight directions for future improvement.

- **No validation with design practitioners:** Although SOLFD is intended to support urban design and planning processes, it has not yet been evaluated by professional practitioners. The methodology and tool development are based primarily on findings from the literature, which often lacks direct engagement with the design community. While insights from my background in Architecture and the Built Environment provided guidance, this does not replace systematic validation through user studies or expert interviews.
- **Exclusion of localized wind data:** Wind conditions are essential for comprehensive outdoor thermal comfort assessment. The current implementation does not include localized wind data, which limits its usefulness in design solutions which provoke trade-offs between solar exposure and wind speed. In such cases, it is necessary to consider both aspects simultaneously to achieve balanced design outcomes. SOLFD currently addresses only one component of the broader UMC context. More broadly, it addresses just a single subcomponent within the inherently complex and multifaceted challenge of urban design.
- **Limited support for design guidance:** While SOLFD identifies areas with unfavorable thermal conditions, it does not offer suggestions for how to improve them. As a result, users may find it difficult to translate the analysis into actionable design changes. This limitation reduces the SOLFD's practical value for design decision-making and positions it more as an evaluation tool rather than a generative design assistant.
- **Performance depends on hardware resources:** The speed improvements are strongly dependent on access to a capable GPU. Thereby, the computation will fail when the GPU-memory is full. In contrast, the original CPU-based version, although slower, would in the end finish the calculation. The computation failure caused by memory limitation can be solved by tiling.
- **Tree classification accuracy is limited:** The current NDVI method for identifying trees in point cloud data is not fully reliable. Some true tree points are missed, while other elements such as streetlights or signage are sometimes incorrectly classified as trees. Due to time constraints and the large scope of the project, a relatively simple classification approach was chosen. More advanced techniques could improve accuracy in future work, but should also remain fast.
- **Limited materials for land cover and buildings:** On one hand, the available land cover dataset lacks representation of common Dutch urban materials. On the other hand, SOLWEIG currently simplifies all buildings and walls by assigning them the same generic material properties, despite real-world variation (e.g., brick, glass, or metal). For improved accuracy, surface parametrization of representative Dutch urban materials is necessary. Additionally, the model should be extended to allow separate specification of materials for building roofs and walls.
- **3D implementation excludes radiative effects of above-head surfaces:** Although the 3D representation has been implemented for shadow and SVF calculations, it is not yet integrated into the subsequent energy balance computations. These calculations continue to rely solely on the base layer, thereby omitting the radiative contributions of elevated surfaces and vertical structures above pedestrian height. This limitation is likely a contributing factor to the underestimation of T_{mrt} observed beneath structures in the validation with sensor data.
- **Applicability is limited to the Netherlands:** The current implementation depends on Dutch national datasets for elevation, canopy height, and land cover. As a result, SOLFD cannot be applied outside of the Netherlands without adaptation of the automatic data pipeline. For a wider applicability, globally available data should be integrated, for example Open Street Map and Meta's global CHM.

6.3. Future work

Building upon the contributions and limitations of this thesis, several directions for future development can be identified to further enhance the usability and accuracy of SOLFD. First, direct engagement with urban designers is required to validate the SOLFD's relevance and usability in real-world design processes. This could be achieved through user testing, participatory workshops, or expert interviews. This should be accompanied by the development of an interface for the SOLFD.

On the technical side, opportunity lies in the integration of wind modeling into the existing framework. As wind speed significantly influences OTC, coupling SOLFD with lightweight wind simulation would allow for more comprehensive assessments. The challenge will be identifying or developing a method that is fast enough to support iterative design while maintaining sufficient accuracy. Moreover, Further 3D model enhancements, such as incorporating the radiative exchanges with vertical and overhanging surfaces, could improve accuracy in 3D enabled models. Additionally, the automatic data pipeline can be improved by refining the vegetation extraction method and by including global datasets. Furthermore, integrating SOLFD with CAD software could streamline the design–feedback loop. For greater design flexibility, functionality to assign and modify building materials should be introduced. To support this and to improve land cover representation more generally, common Dutch urban materials need to be studied and parameterized for use in the SOLWEIG algorithm.

Finally, for the accurate validation of SOLFD 3D, static sensor data is needed. The sensors should be placed underneath and adjacent to structures, and ideally collected over multiple warm days.

References

Abdollahzadeh, N. and N. Biloria (2021). "Outdoor thermal comfort: Analyzing the impact of urban configurations on the thermal performance of street canyons in the humid subtropical climate of Sydney". In: *Frontiers of Architectural Research* 10.2, pp. 394–409. ISSN: 2095-2635. DOI: https://doi.org/10.1016/j.foar.2020.11.006. URL: https://www.sciencedirect.com/science/article/pii/S2095263520300844.

Aleksandrowicz, O. (2022). "Chapter 1 - Mapping and management of urban shade assets: a novel approach for promoting climatic urban action". In: *Global Urban Heat Island Mitigation*. Ed. by A. Khan, H. Akbari, F. Fiorito, S. Mithun, and D. Niyogi. Elsevier, pp. 1–27. ISBN: 978-0-323-85539-6. DOI: https://doi.org/10.1016/B978-0-323-85539-6.00003-2. URL: https://www.sciencedirect.com/science/article/pii/B9780323855396000032.

Ambrosini, D., G. Galli, B. Mancini, I. Nardi, and S. Sfarra (2014). "Evaluating Mitigation Effects of Urban Heat Islands in a Historical Small Center with the ENVI-Met® Climate Model". In: *Sustainability* 6.10, pp. 7013–7029. ISSN: 2071-1050. DOI: 10.3390/su6107013. URL: https://www.mdpi.com/2071-1050/6/10/7013.

Badino, E., M. Ferrara, L. Shtrepi, E. Fabrizio, A. Astolfi, and V. Serra (2021). "Modelling mean radiant temperature in outdoor environments: Contrasting the approaches of different simulation tools". In: *Journal of Physics: Conference Series*. Vol. 2069. 1. IOP Publishing, p. 012186. DOI: 10.1088/1742-6596/2069/1/012186.

Bernard, J., F. Lindberg, and S. Oswald (Oct. 2023). "URock 2023a: an open-source GIS-based wind model for complex urban settings". In: *Geoscientific Model Development* 16, pp. 5703–5727. DOI: 10.5194/gmd-16-5703-2023.

Besserud, K. and T. Hussey (2011). "Urban design, urban simulation, and the need for computational tools". In: *IBM Journal of Research and Development* 55.1.2, 2:1–2:17. DOI: 10.1147/JRD.2010.2097091.

Bherwani, H., A. Singh, and R. Kumar (2020). "Assessment methods of urban microclimate and its parameters: A critical review to take the research from lab to land". In: *Urban Climate* 34, p. 100690. ISSN: 2212-0955. DOI: https://doi.org/10.1016/j.uclim.2020.100690. URL: https://www.sciencedirect.com/science/article/pii/S2212095520302364.

Buo, I., V. Sagris, J. Jaagus, and A. Middel (2023). "High-resolution thermal exposure and shade maps for cool corridor planning". In: *Sustainable Cities and Society* 93, p. 104499. ISSN: 2210-6707. DOI: https://doi.org/10.1016/j.scs.2023.104499. URL: https://www.sciencedirect.com/science/article/pii/S2210670723001105.

Carmona, M., T. Heath, T. Oc, and S. Tiesdell (2010). "Part 1 - Defining Urban Design". In: *Public Places - Urban Spaces: The Dimensions of Urban Design*. 2nd. Taylor & Francis Group, pp. 3–74.

Chen, K. W. and L. Norford (2017). "Evaluating Urban Forms for Comparison Studies in the Massing Design Stage". In: *Sustainability* 9.6. ISSN: 2071-1050. DOI: 10.3390/su9060987. URL: https://www.mdpi.com/2071-1050/9/6/987.

Cheshmehzangi, A., C. Butters, A. Cheshmehzangi, A. Dawodu, R. Jin, and E. Mangi (2022). "Urban Design in the Age of Climate Change: Paradigms and Directions". In: *Design for Vulnerable Communities*. Cham: Springer International Publishing, pp. 133–154. ISBN: 978-3-030-96866-3. DOI: 10.1007/978-3-030-96866-3_9. URL: https://doi.org/10.1007/978-3-030-96866-3_9.

Copernicus Climate Change Service (C3S) (2024). *Climate Reanalysis*. https://climate.copernicus.eu/climate-reanalysis. Accessed: 2025-06-01.

Ebrahimabadi, S. (2015). "Outdoor Comfort in Cold Climates: Integrating Microclimate Factors in Urban Design". PhD thesis. Luleå University of Technology, p. 195. ISBN: 978-91-7583-378-1.

Eldarwish, I., A. Ragheb, and A. Sherif (Aug. 2020). "The Role of Simulation in Urban Design Decisions: Microclimate and Human Comfort Considerations in Planning". In: *Architecture and Planning Journal (APJ)* 23. DOI: 10.54729/2789-8547.1048.

Erell, E., D. Pearlmutter, and T. Williamson (2010). "Modelling the Urban Microclimate". In: *Urban Microclimate: Designing the Spaces Between Buildings*. Ed. by E. Erell, D. Pearlmutter, and T. Williamson. London: Routledge, pp. 211–234. ISBN: 978-1844074679. DOI: 10.4324/9781849775397.

Fennell, J. T. and L. Fay (2025). *Open British Arboricultural Record Dataset (OpenBARD)*. Version 1.3. [Data set]. JF Scientific. DOI: 10.5281/zenodo.15593688.

Gál, C. V. and N. Kántor (2020). "Modeling mean radiant temperature in outdoor spaces, A comparative numerical simulation and validation study". In: *Urban Climate* 32, p. 100571. ISSN: 2212-0955. DOI: https://doi.org/10.1016/j.uclim.2019.100571. URL: https://www.sciencedirect.com/science/article/pii/S2212095519301804.

Geonovum (2020). *Basisregistratie Grootschalige Topografie Gegegevenscatalogus BGT 1.2*. https://docs.geostandaarden.nl/imgeo/catalogus/bgt. Accessed: 30-05-2025.

Goodwin, N. R., N. C. Coops, T. R. Tooke, A. Christen, and J. A. V. and (2009). "Characterizing urban surface cover and structure with airborne lidar technology". In: *Canadian Journal of Remote Sensing* 35.3, pp. 297–309. DOI: 10.5589/m09-015. eprint: https://doi.org/10.5589/m09-015. URL: https://doi.org/10.5589/m09-015.

Google (2022). *Google Street View in Haarlem, Netherlands*. https://www.google.com/maps/. Accessed: 06-05-2025.

Guo, F., R. Guo, H. Zhang, J. Dong, and J. Zhao (2023). "A canopy shading-based approach to heat exposure risk mitigation in small squares". In: *Urban Climate* 49, p. 101495. ISSN: 2212-0955. DOI: https://doi.org/10.1016/j.uclim.2023.101495. URL: https://www.sciencedirect.com/science/article/pii/S2212095523000895.

Heusinkveld, B. G., G.-J. Steeneveld, B. Holtslag, B. Hove, and C. Jacobs (Dec. 2013). "Spatial variability of the Rotterdam urban heat island as influenced by urban land use". In: *Journal of Geophysical Research* 119, pp. 677–692. DOI: 10.1002/2012JD019399.

Holmer, B., F. Lindberg, D. Rayner, and S. Thorsson (July 2015). "How to transform the standing man from a box to a cylinder – a modified methodology to calculate mean radiant temperature in field studies and models". In: ICUC9 - 9th International Conference on Urban Climate jointly with 12th Symposium on the Urban Environment.

Höppe, P. (Oct. 1999). "The physiological equivalent temperature - a universal index for the biometeorological assessment of the thermal environment". In: *International Journal of Biometeorology* 43.2, pp. 71–75. DOI: 10.1007/s004840050118. URL: http://dx.doi.org/10.1007/s004840050118.

Hsieh, C.-M. and W. Kangli (Dec. 2024). "Climate-Sensitive Urban Design Measures for Improving the Wind Environment for Pedestrians in a Transit-Oriented Development Area". In: *Journal of Sustainable Development* 5, pp. 46–46. DOI: 10.5539/jsd.v5n4p46.

Huang, J., T. Hao, S. Hou, and P. Jones (Oct. 2019). "Simulation-Informed Urban Design: Improving Urban Microclimate in Real-World Practice in a High Density City." In: *IOP Conference Series: Earth and Environmental Science* 329, p. 012047. DOI: 10.1088/1755-1315/329/1/012047.

Kanopoulos, N., N. Vasanthavada, and R. Baker (1988). "Design of an image edge detection filter using the Sobel operator". In: *IEEE Journal of Solid-State Circuits* 23.2, pp. 358–367. DOI: 10.1109/4.996.

Kleerekoper, L., S. Erwin, L. Corpel, P. Bons, Z. Arif, J. Kluck, K. Veenenbos, A. Koekoek, and L. Geisler (2023). *Wijktypologie PC6 Gebieden 2023*. https://research.hva.nl/en/datasets/wijktypologie-pc6-gebieden-2023/. Dataset.

Klemm, W., L. Altinisik, G. Peretti, and M. Bruse (Oct. 2014). "Integration of Microclimate-Responsive Design in the Planning of Urban Outdoor Spaces – a Case Study in Athens, Greece". In.

Klimaatadaptatie Nederland (2023). *Oorzaken van hitte*. Accessed: 2025-06-10. URL: https://klimaatadaptatienederland.nl/kennisdossiers/hitte/oorzaken/.

Klimaateffectatlas (2023). *Wijktypologie*. Accessed: 14-01-2025. URL: https://www.klimaateffectatlas.nl/nl/wijktypologie.

Kluck, J., L. Kleerekoper, S. Erwin, L. Corpel, P. Bons, Z. Arif, L. Geisler, K. Veenbos, and A. Koekoek (Sept. 2023). *Computer leert wijktypen bepalen voor heel Nederland: een studie naar het inzetten van kunstmatige intelligentie voor het bepalen van wijktypen*. Dutch. URL: https://pure.hva.nl/ws/portalfiles/portal/43807569/Rapport_wijktypen_2023_sept.pdf.

Kong, J., Y. Zhao, J. Carmeliet, and C. Lei (Sept. 2021). "Urban Heat Island and Its Interaction with Heatwaves: A Review of Studies on Mesoscale". In: *Sustainability* 13, p. 10923. DOI: 10.3390/su131910923.

Ledoux, H., K. A. Ohori, R. Peters, and M. Pronk (2023). "Spatial interpolation: deterministic methods". In: *Computational Modelling of Terrains v2023.0*. Accessed: 12-11-2023. TU Delft. Chap. 5. URL: https://github.com/tudelft3d/terrainbook/releases.

Lee, H., J. Holst, and H. Mayer (Jan. 2013). "Modification of Human-Biometeorologically Significant Radiant Flux Densities by Shading as Local Method to Mitigate Heat Stress in Summer within Urban Street Canyons". In: *Advances in Meteorology* 2013, pp. 1–13. DOI: 10.1155/2013/312572.

Li, X. and G. Wang (June 2021). "GPU parallel computing for mapping urban outdoor heat exposure". In: *Theoretical and Applied Climatology* 145, pp. 1101–1111. DOI: 10.1007/s00704-021-03692-z.

Lindberg, F. and C. S. B. Grimmond (2019). *SOLWEIG_v2019a*. Accessed: 02-12-2024. URL: https://umep-docs.readthedocs.io/en/latest/OtherManuals/SOLWEIG.html.

Lindberg, F. and S. Grimmond (July 2010). "Continuous sky view factor maps from high resolution urban digital elevation models". In: *Climate research (Open Access for articles 4 years old and older)* 42, p. 177. DOI: 10.3354/cr00882.

— (Oct. 2011). "The influence of vegetation and building morphology on shadow patterns and mean radiant temperatures in urban areas: Model development and evaluation". In: *Theoretical and Applied Climatology* 105, pp. 311–323. DOI: 10.1007/s00704-010-0382-8.

Lindberg, F., B. Holmer, and S. Thorsson (June 2008). "Solweig 1.0 – modelling spatial variations of 3D radiant fluxes and mean radiant temperature in complex urban settings". In: *International Journal of Biometeorology* 52.7, pp. 697–713. DOI: 10.1007/s00484-008-0162-7.

Lindberg, F., S. Onomura, and S. Grimmond (Feb. 2016). "Influence of ground surface characteristics on the mean radiant temperature in urban areas". In: *International Journal of Biometeorology* 60, pp. 1439–1452. DOI: 10.1007/s00484-016-1135-x.

Lindberg, F. et al. (2018). "Urban Multi-scale Environmental Predictor (UMEP): An integrated tool for city-based climate services". In: *Environmental Modelling & Software* 99, pp. 70–87. ISSN: 1364-8152. DOI: https://doi.org/10.1016/j.envsoft.2017.09.020. URL: https://www.sciencedirect.com/science/article/pii/S1364815217304140.

Matzarakis, A. and H. Mayer (Jan. 1996). "Another kind of environmental stress: Thermal stress". In: *WHO Collaborating Centre for Air Quality Management and Air Pollution Control* 18, pp. 7–10.

McPherson, E. G., N. S. van Doorn, and P. J. Peper (2016a). *Urban Tree Database*. [Data set]. Fort Collins, CO. DOI: 10.2737/RDS-2016-0005.

— (2016b). *Urban Tree Database and Allometric Equations*. General Technical Report PSW-GTR-253. Albany, CA: U.S. Department of Agriculture, Forest Service, Pacific Southwest Research Station, p. 86. DOI: 10.2737/PSW-GTR-253. URL: https://www.fs.usda.gov/psw/publications/documents/psw_gtr253/psw_gtr_253.pdf.

Monahan, J., H. Gan, Y. Gao, C. C. Andinasari, V. Tsalapati, H. Ledoux, and L. Beuster (2024). *Shady Amsterdam: Identifying the shady places and routes of Amsterdam*. Student report, Faculty of Architecture and the Built Environment, Delft University of Technology. URL: http://resolver.tudelft.nl/uuid:1878b296-b888-4043-a9a6-9050ab3e5723.

Natuur & Milieu (Dec. 2023). *Bomen voor Biodiversiteit*. https://natuurenmilieu.nl/app/uploads/Bomenonderzoek-Nederlandse-gemeenten-Natuur-Milieu.pdf. Natuur & Milieu.

Okuta, R., Y. Unno, D. Nishino, S. Hido, and C. Loomis (2017). "CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations". In: *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*. URL: http://learningsys.org/nips17/assets/papers/paper_16.pdf.

Oude Elberink, S. and G. Vosselman (2009). "3D information extraction from laser point clouds covering complex road junctions". English. In: *Photogrammetric record* 24.125, pp. 23–36. ISSN: 0031-868X. DOI: 10.1111/j.1477-9730.2008.00516.x.

Pađen, I. (2023). *rusterizer*. https://github.com/ipadjen/rusterizer. GitHub repository.

Rokka, L. (2021). *Shiny Weather Data — Preprocessed ERA5 and CAMS datasets*. Accessed: 02-03-2025. URL: https://www.shinyweatherdata.com/.

Roudsari, M. S., M. Pak, and A. Viola (Aug. 2013). "Ladybug: A Parametric Environmental Plugin For Grasshopper To Help Designers Create An Environmentally-conscious Design". In: *Proceedings of Building Simulation 2013: 13th Conference of IBPSA*. Vol. 13. Building Simulation. Chambery, France: IBPSA, pp. 3128–3135. ISBN: 978-2-7466-6294-0. DOI: https://doi.org/10.26868/25222708.2013.2499. URL: https://publications.ibpsa.org/conference/paper/?id=bs2013_2499.

Schneider, M., T. Tötzer, M. Bügelmayer-Blaschek, and R. Berg (2023). "Pitfalls and Potentials of Microclimate Simulations in Urban Planning". In: *Journal of Urban Planning and Development* 149.4, p. 04023048. DOI: 10.1061/JUPDDM.UPENG-4504. eprint: https://ascelibrary.org/doi/pdf/10.1061/JUPDDM.UPENG-4504. URL: https://ascelibrary.org/doi/abs/10.1061/JUPDDM.UPENG-4504.

Stathopoulos, T., H. Wu, and J. Zacharias (Mar. 2004). "Outdoor human comfort in an urban climate". In: *Building and Environment* 39, pp. 297–305. DOI: 10.1016/j.buildenv.2003.09.001.

Thom, J. K., A. M. Coutts, A. M. Broadbent, and N. J. Tapper (2016). "The influence of increasing tree cover on mean radiant temperature across a mixed development suburb in Adelaide, Australia". In: *Urban Forestry & Urban Greening* 20, pp. 233–242. ISSN: 1618-8667. DOI: https://doi.org/10.1016/j.ufug.2016.08.016. URL: https://www.sciencedirect.com/science/article/pii/S1618866716301297.

Thorsson, S., J. Rocklöv, J. Konarska, F. Lindberg, B. Holmer, B. Dousset, and D. Rayner (2014). "Mean radiant temperature – A predictor of heat related mortality". In: *Urban Climate* 10. ICUC8: The 8th International Conference on Urban Climate and the 10th Symposium on the Urban Environment, pp. 332–345. ISSN: 2212-0955. DOI: https://doi.org/10.1016/j.uclim.2014.01.004. URL: https://www.sciencedirect.com/science/article/pii/S2212095514000054.

Van Esch, M. (June 2015). "Designing the Urban Microclimate. A framework for a design-decision support tool for the dissemination of knowledge on the urban microclimate to the urban design process". In: *A+BE: Architecture and the Built Environment* 5. DOI: 10.7480/abe.2015.6.

Vurro, G. and S. Carlucci (2024). "Contrasting the features and functionalities of urban microclimate simulation tools". In: *Energy and Buildings* 311, p. 114042. ISSN: 0378-7788. DOI: https://doi.org/10.1016/j.enbuild.2024.114042. URL: https://www.sciencedirect.com/science/article/pii/S0378778824001580.

Wallenberg, N., B. Holmer, F. Lindberg, and D. Rayner (Feb. 2023). "An anisotropic parameterization scheme for longwave irradiance and its impact on radiant load in urban outdoor settings". In: *International journal of biometeorology* 67. DOI: 10.1007/s00484-023-02441-3.

Wallenberg, N., F. Lindberg, B. Holmer, and S. Thorsson (2020). "The influence of anisotropic diffuse shortwave radiation on mean radiant temperature in outdoor urban environments". In: *Urban Climate* 31, p. 100589. ISSN: 2212-0955. DOI: https://doi.org/10.1016/j.uclim.2020.100589. URL: https://www.sciencedirect.com/science/article/pii/S2212095519301178.

Weytjens, L. and G. Verbeeck (Apr. 2010). "Towards 'architect-friendly' energy evaluation tools". In: p. 179. DOI: 10.1145/1878537.1878724.

Wong, N. H., S. K. Jusuf, and C. L. Tan (2011). "Integrated urban microclimate assessment method as a sustainable urban development and urban design tool". In: *Landscape and Urban Planning* 100.4. Landscape and Urban Planning at 100, pp. 386–389. ISSN: 0169-2046. DOI: https://doi.org/10.1016/j.landurbplan.2011.02.012. URL: https://www.sciencedirect.com/science/article/pii/S0169204611000703.

Xie, Y., T. Huang, J. Li, J. Liu, J. Niu, C. M. Mak, and Z. Lin (2018). "Evaluation of a multi-nodal thermal regulation model for assessment of outdoor thermal comfort: Sensitivity to wind speed and solar radiation". In: *Building and Environment* 132, pp. 45–56. ISSN: 0360-1323. DOI: https://doi.org/10.1016/j.buildenv.2018.01.025. URL: https://www.sciencedirect.com/science/article/pii/S0360132318300374.

A

Appendix: Supplementary Tables and Figures

A.1. SOLWEIG and preprocessors interfaces

To provide an overview of SOLWEIG and its required preprocessors for data preparation, screenshots were taken from the QGIS UMEP graphical user interface (GUI) for each component. Figure A.1 shows the SOLWEIG interface. Users must select QGIS layers corresponding to the required rasters and specify parameters such as UTC time and the height of the trunk zone.

Figure A.2 displays the climate preprocessor tool, which allows users to convert meteorological data, such as those from shinyweatherdata.com, into the correct format for use in QGIS. Each meteorological variable must be manually assigned to the correct column. Note that settings are not saved, so the assignment must be repeated each time a new file is loaded.

The preprocessor for SVF creation, shown in Figure A.3, features a simplified interface compared to the main SOLWEIG tool, requiring only a DSM and CHM.

Figure A.4 shows the wall height and aspect preprocessor. This tool requires only a DSM, and users can define a minimum height threshold for identifying walls.

Finally, users can generate a CHM and trunk zone themselves using the tree generator preprocessor (Figure A.5). For this tool, users need to input the location of trees along with parameters such as shape, total height, trunk height, and crown diameter. It can also be used to manually insert custom trees.

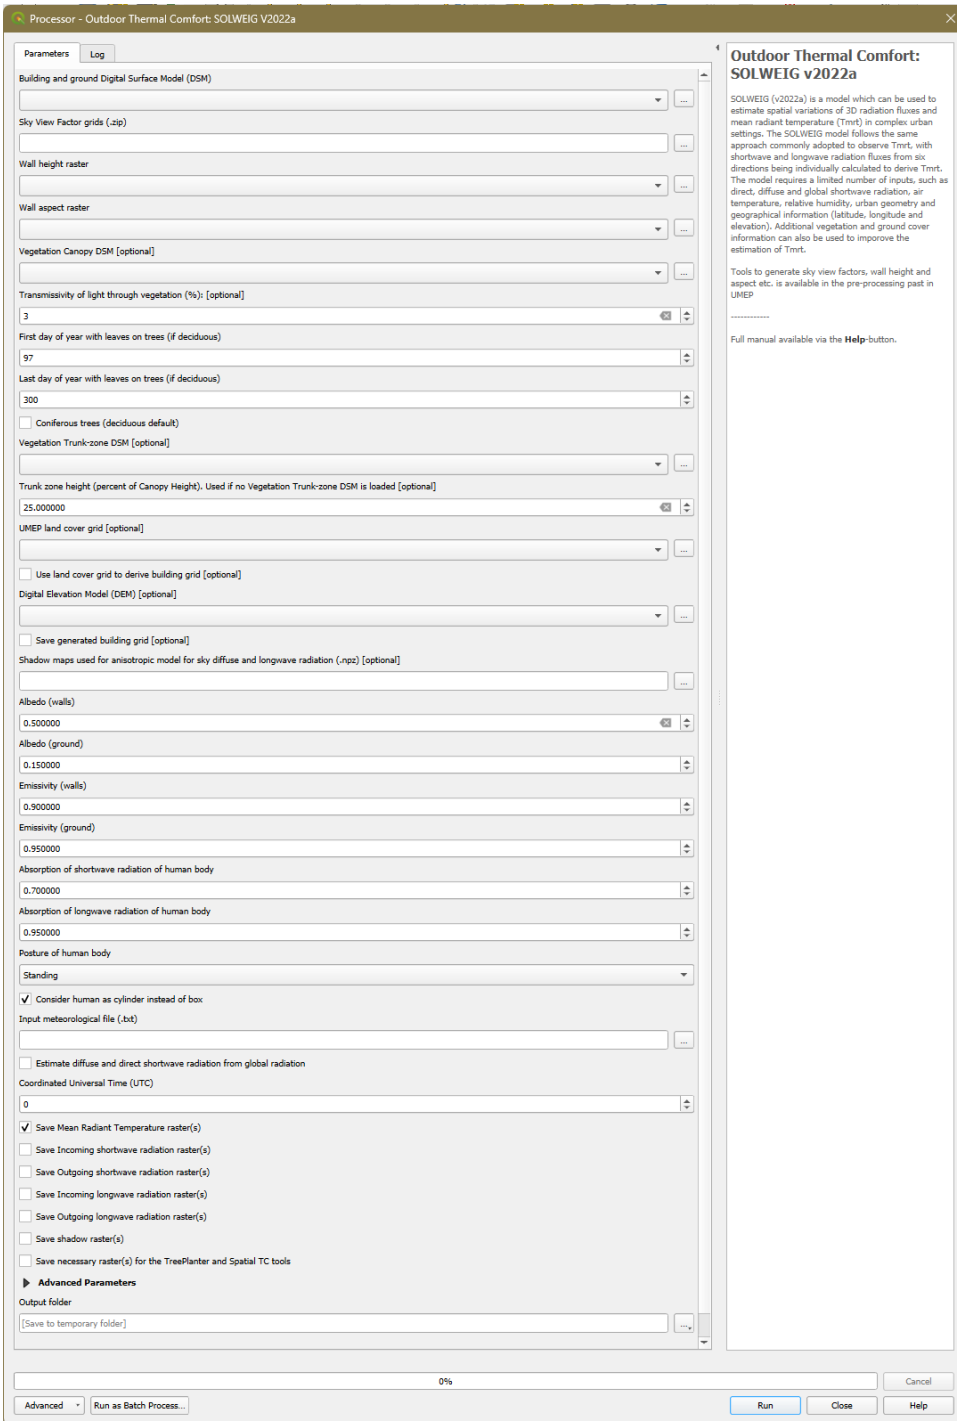

Figure A.1: The SOLWEIG interface from the UMEP QGIS plug-in.

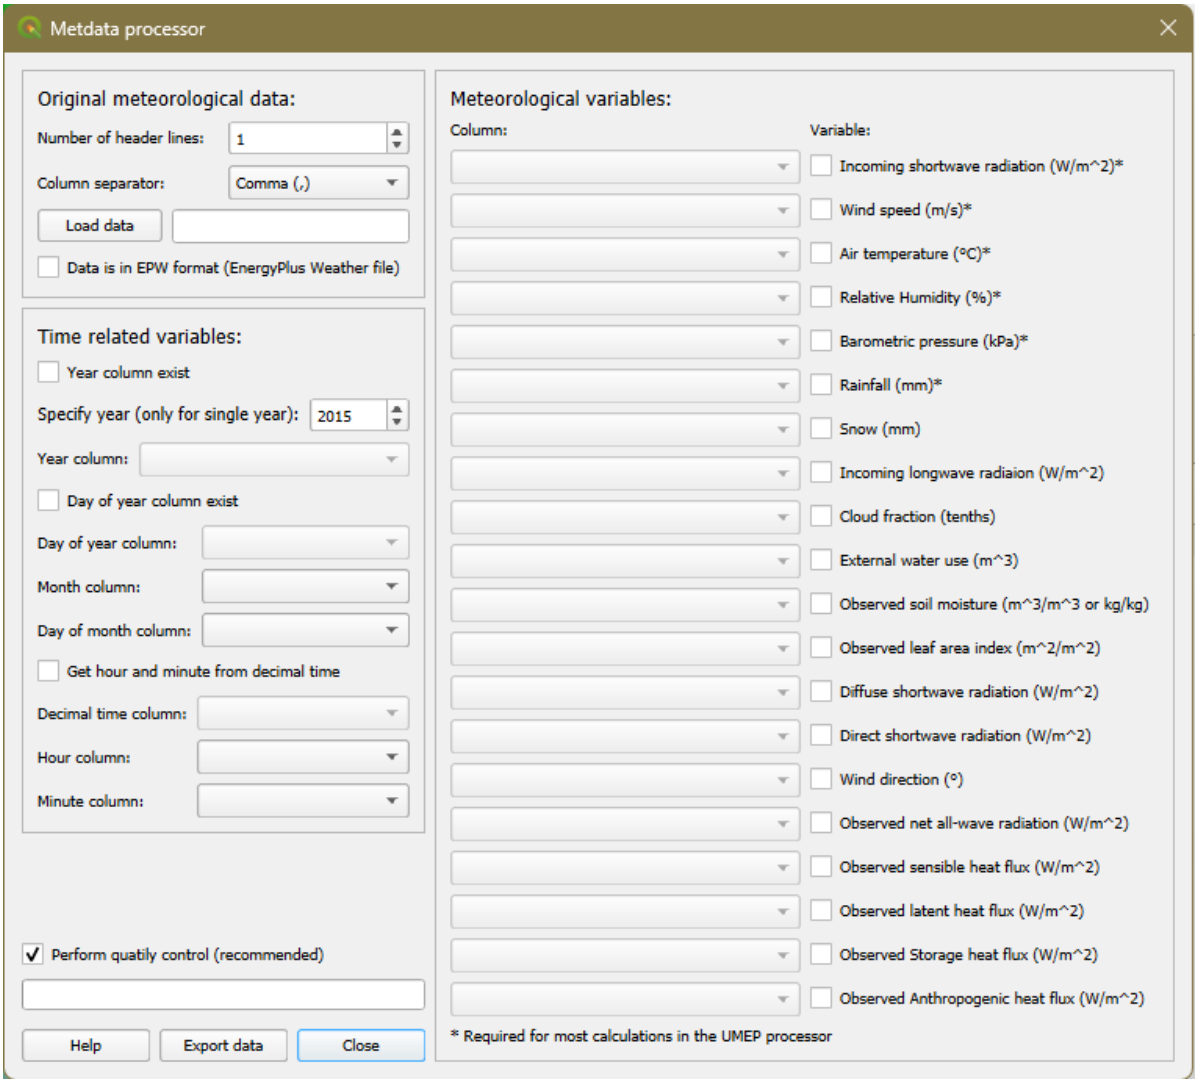

Figure A.2: The 'Prepare Existing Data' interface from the UMEP QGIS plug-in.

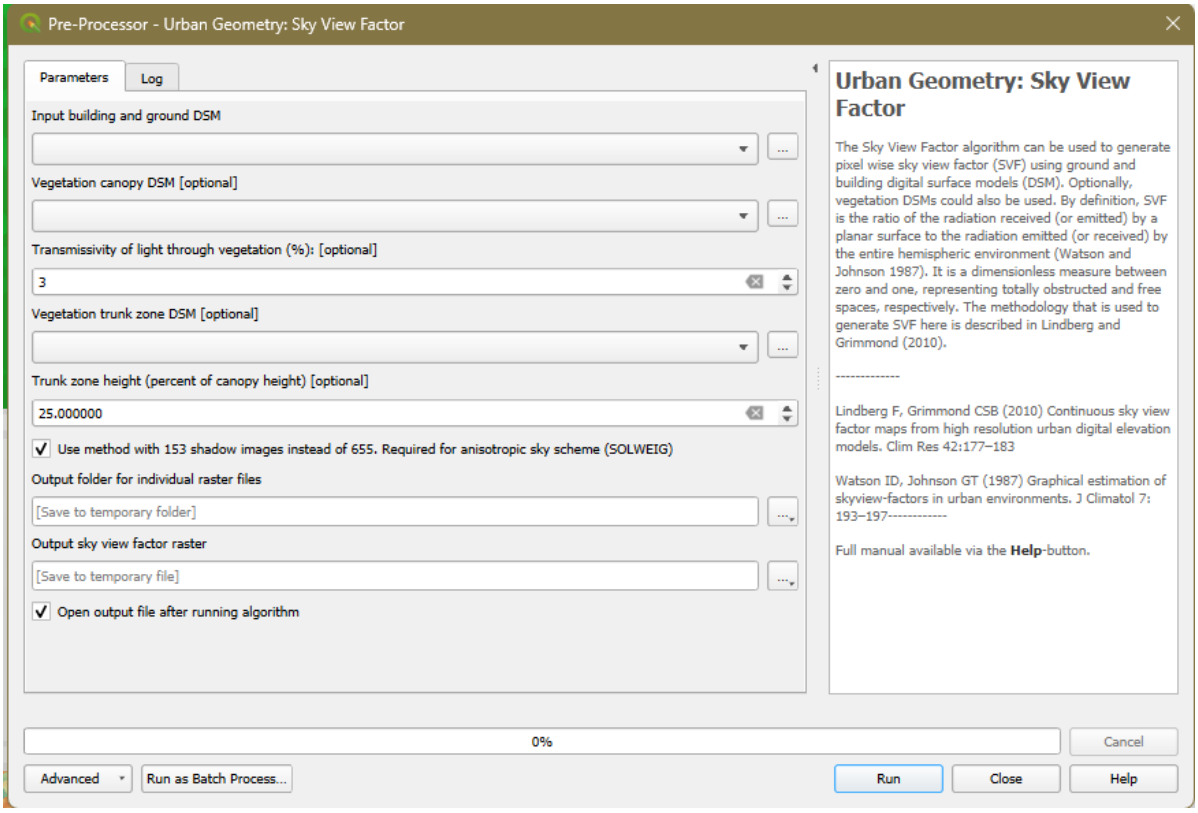

Figure A.3: The 'Urban Geometry: Sky View Factor' interface from the UMEP QGIS plug-in.

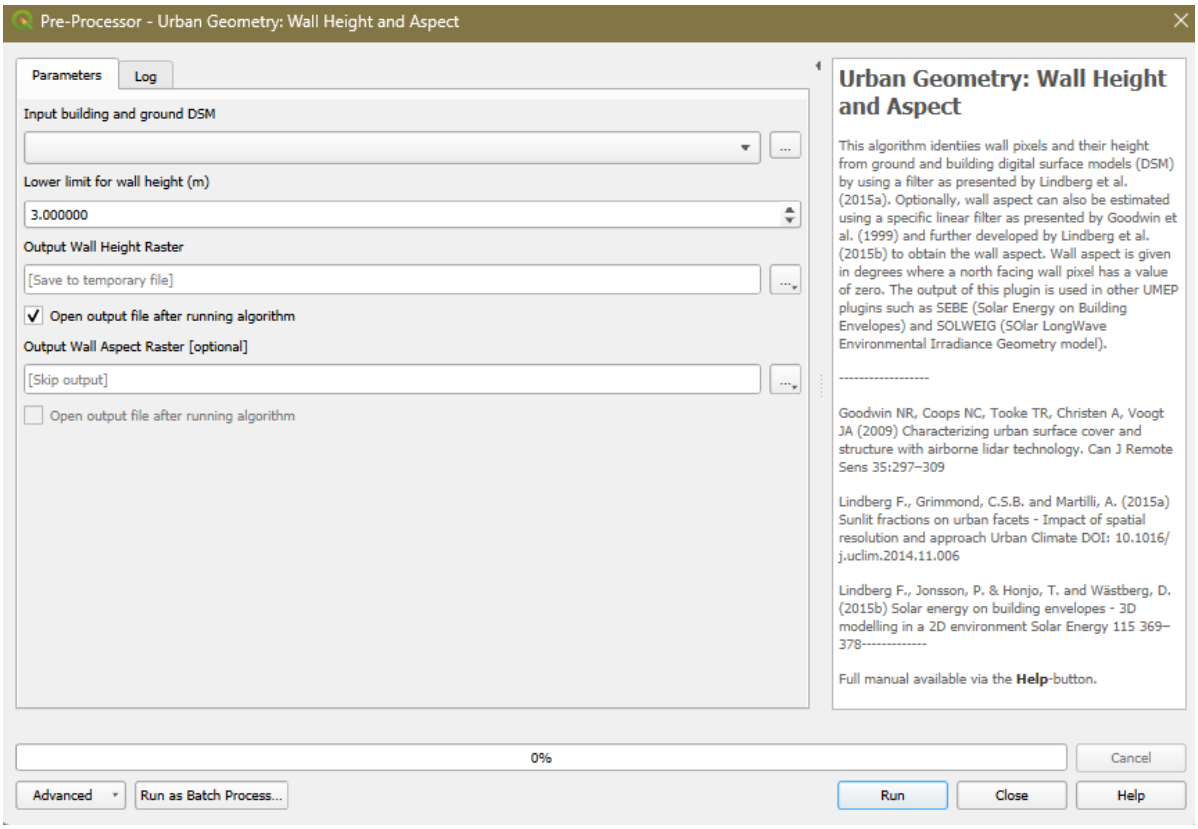

Figure A.4: The 'Urban Geometry: Wall Aspect and Height' interface from the UMEP QGIS plug-in.

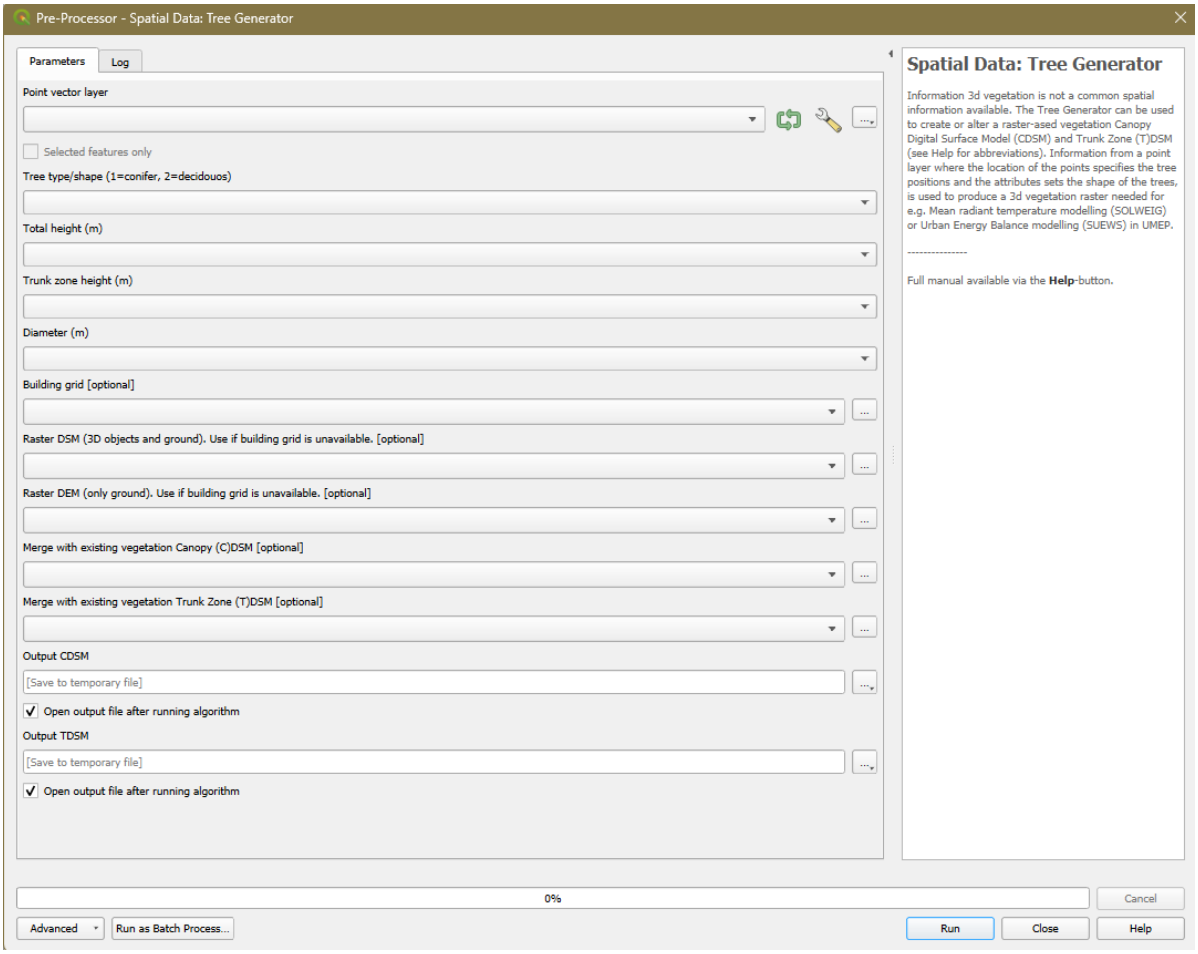

Figure A.5: The 'Spatial Data: Tree Generator' preprocessor interface from the UMEP QGIS plug-in. Used for CHM creation, points are required as input for tree positions.

A.2. Urban Case Study Locations

Table A.1: Case study locations grouped by urban typology, showing bounding boxes in EPSG:28992 coordinates and array sizes (in meters).

Entry	Historisch	Tuindorp	Vinex	Volkswijk	Bloemkool	Stedelijk
0	Maastricht, 600×600, (175905, 317210, 176505, 317810)	Scheveningen, 1400×700, (76800, 455000, 78200, 455700)	Almere, 900×900, (146100, 486500, 147000, 487400)	Haarlem, 900×900, (104200, 490550, 105100, 491450)	Spijkernisse, 1000×710, (81700, 427490, 82700, 428200)	Den Haag, 1000×700, (90300, 436900, 91300, 437600)
1	Delft, 600×600, (84050, 447180, 84650, 447780)	Amersfoort, 1300×550, (152600, 463250, 153900, 463800)	Amersfoort, 900×900, (153750, 467550, 154650, 468450)	Den Haag, 900×900, (78200, 453900, 79100, 454800)	Den Haag, 900×900, (84050, 444000, 84950, 444900)	Den Haag, 900×800, (91200, 438500, 92100, 439300)
2	Den Haag, 600×600, (80780, 454550, 81380, 455150)	Hilversum, 720×830, (139140, 469570, 139860, 470400)	Heerhugowaard, 800×850, (115300, 517400, 116100, 518250)	Delft, 550×880, (83500, 447020, 84050, 447900)	Heerhugowaard, 900×900, (116650, 518700, 117550, 519600)	Amsterdam, 900×900, (121350, 483750, 122250, 484650)
3	Groningen, 600×600, (233400, 581500, 234000, 582100)	Arnhem, 900×750, (190850, 441790, 191750, 442540)	Nieuw-Vennep, 1100×900, (102000, 475900, 103100, 476800)	Utrecht, 900×800, (136200, 456500, 137100, 457300)	Groningen, 900×900, (235050, 584950, 235950, 585850)	Amsterdam, 940×700, (118400, 486400, 119340, 487100)
4	Utrecht, 600×600, (136600, 455850, 137200, 456450)	Den Helder, 550×400, (113100, 551600, 113650, 552000)	Eindhoven, 900×900, (160750, 388450, 161650, 389350)	Leeuwarden, 1100×550, (182700, 579200, 183800, 579750)	Deventer, 900×900, (210500, 473900, 211400, 474800)	–
5	Amsterdam, 600×600, (121500, 487000, 122100, 487600)	Middelburg, 800×600, (32050, 391900, 32850, 392500)	Nootdorp, 900×900, (84350, 449800, 85250, 450700)	Groningen, 900×900, (233400, 582800, 234300, 583700)	Eindhoven, 1000×700, (154700, 381450, 155700, 382150)	–

A.3. Classification of BGT and TOP10NL types to SOLWEIG landcover types

A.3.1. BGT

Using the SOLWEIG classified materials results in a significant loss of detail. For example, materials like concrete and bricks, which are common in Dutch cities, are not represented in the dataset. Table A.2 shows the mapping choices made between BGT and SOLWEIG classes.

For the entire category of *vegetated terrain unit* (begroeid terreindeel), all features are assigned the land cover class 'grass' (5), even though this category includes a variety of vegetation-covered surfaces, such as forests and farmland. This simplification treats any vegetated surface as the only vegetated surface included. This collection also includes 'dune' (duin), which is classified as bare soil (6) due to its sandy composition.

For the category *non-vegetated terrain unit* (onbegroeid terreindeel), several physical surface types are included under the BGT attribute *physical appearance* (fysiek voorkomen). The following mapping decisions are made based on the official descriptions from Geonovum, 2020:

- **Closed paving** (gesloten verharding): "Paving made of material that cannot be removed without permanent destruction." This corresponds best to dark asphalt (class 1).
- **Open paving** (open verharding): "Paving made of interlocking small-scale elements, such as bricks or tiles." This is mapped to cobblestone (class 0).
- **Semi-paving** (half verhard): "Paving made of compacted or loose material." Since no exact match exists in SOLWEIG, cobblestone (class 0) is chosen as the closest option.
- **Unpaved** (onverhard): "Areas without paving or continuous vegetation." This corresponds best to bare soil (class 6).
- **Yard** (erf): Defined as terrain associated with buildings, consisting of a mix of vegetation, paving, and/or water. Based on its typical appearance (e.g., courtyards or sidewalks), it is mapped to cobblestone (class 0).

Roads follow the same logic as terrain, with the additional type *green space* (groenvoorziening), which is vegetated and therefore mapped to grass (class 5).

Additionally, features may include a secondary attribute, *additional physical appearance* (plus fysiek voorkomen), which provides more detailed surface material information—such as whether the material is asphalt, concrete, paving stones, or bricks. This attribute will be useful once these materials are incorporated into SOLWEIG.

A.3.2. TOP10NL

Another option for users is to use the TOP10NL dataset (`https://api.pdok.nl/brt/top10nl/ogc/v1`) to generate land cover classifications. This dataset is more generalized than BGT, consisting of larger polygons with less detail. Consequently, less data needs to be requested, making the processing somewhat faster.

Terrain, road, and water features retrieved via the PDOK API are mapped to SOLWEIG-compatible land cover classes. However, these classes offer less detail than those derived from BGT. For instance, roads are only classified as paved or unpaved, without further distinctions such as whether the paving is open or closed. To address this, users can specify the dominant road surface material, which is then uniformly applied to all road segments.

The mapping approach follows the same general principles as with the BGT dataset and is summarized in Table A.3.

Table A.2: Mapping of terrain, road, and water types from the BGT dataset to SOLWEIG land cover classes.

Feature Type	Assigned solweig Class
Terrain	
Begroeid terreindeel	5 (Grass)
Duin	6 (Bare Soil)
Zand	6 (Bare Soil)
Gesloten verharding	1 (Dark Asphalt)
Open verharding	0 (Cobble Stone)
Half verhard	0 (Cobble Stone)
Onverhard	6 (Bare Soil)
Erf	0 (Cobble Stone)
Road	
Gesloten verharding	1 (Dark Asphalt)
Open verharding	0 (Cobble Stone)
Half verhard	0 (Cobble Stone)
Onverhard	6 (Bare Soil)
Groenvoorziening	5 (Grass)
Water	
Water	7 (Water)

Table A.3: Mapping of TOP10NL terrain, road, and water types to SOLWEIG land cover classes.

Feature Type	Assigned SOLWEIG Class
Terrain	
Aanlegsteiger	0 (Cobble Stone)
Basaltblokken, Steenglooiing	0 (Cobble Stone)
Bebouwd gebied	0 (Cobble Stone)
Overig	0 (Cobble Stone)
Akkerland	6 (Bare Soil)
Duin	6 (Bare Soil)
Zand	6 (Bare Soil)
Boomgaard	5 (Grass)
Boomkwekerij	5 (Grass)
Bos: gemengd bos	5 (Grass)
Bos: griend	5 (Grass)
Bos: loofbos	5 (Grass)
Bos: naaldbos	5 (Grass)
Dodenakker	5 (Grass)
Dodenakker met bos	5 (Grass)
Fruitkwekerij	5 (Grass)
Grasland	5 (Grass)
Heide	5 (Grass)
Populieren	5 (Grass)
Braakliggend	5 (Grass)
Spoorbaanlichaam	0 (Cobble Stone)
Road	
Roads (Verhard)	1 (Dark Asphalt or user-selected material)
Roads (Half Verhard)	0 (Cobble Stone)
Roads (Onverhard)	6 (Bare Soil)
Roads (Onbekend)	0 (Cobble Stone)
Water	
Water (polygons and linestrings)	7 (Water)

A.4. Evaluated Tree Databases for Dutch Species Parameters

Table A.4 summarizes the tree databases that were examined to obtain key parameters for the five most common Dutch tree species. Essential parameters included canopy radius or diameter, tree height, and an indication of tree age. Trunk height was considered a desirable additional parameter. Only the last two databases were suitable.

Table A.4: Overview of tree databases evaluated for extracting species-specific parameters, including crown radius, tree height, and age, for the five most common Dutch tree species. The table also indicates the number of relevant species per dataset and whether the data were suitable for use.

Dataset	Source	Number of Relevant Tree Species	Data Suitability
Ebben Nurseries Tree	https://www.ebben.nl/en/treeebb/	5	No, missing parameters
Tallo - a global tree allometry and crown architecture database	https://doi.org/10.5281/zenodo.6637599	2	No, missing parameters
National register of big trees (Australia)	https://www.nationalregisterofbigtrees.com.au/pages/tree-data	4	No, missing parameters
Temperate Species - Tree database	https://www.wur.nl/en/research-results/chair-groups/environmental-sciences/forest-ecology-and-forest-management-group/education/tree-database/	3	No, missing parameters
Plant Trait Database	https://www.try-db.org/TryWeb/dp.php	4	No, missing parameters
Open British Arboricultural Record Dataset (OpenBARD)	https://doi.org/10.5281/zenodo.15593687	5	Yes, simplified
Urban tree database and allometric equations	https://doi.org/10.2737/PSW-GTR-253	1	Yes

A.5. Meteorology Tables and Figures

Year	ID	Hour	Imin	Q*	QH	QE	Qs	Qf	Wind	RH	Td	Press	Rain	Kdn	Snow	Ldow	Fcld	Wuh	Xsmc	Lai_hr	Kdiff	Kdir	Wd
2023	213	0	0							87.00	16.50	101.30		0.00							0.00	0.00	
2023	213	1	0							88.00	16.20	101.30		0.00							0.00	0.00	
2023	213	2	0							89.00	15.90	101.30		0.00							0.00	0.00	
2023	213	3	0							90.00	15.70	101.30		0.00							0.00	0.00	
2023	213	4	0							90.00	15.60	101.30		1.70							1.50	3.90	
2023	213	5	0							89.00	16.00	101.30		28.00							19.00	84.00	
2023	213	6	0							87.00	17.00	101.30		86.00							53.00	187.00	
2023	213	7	0							82.00	18.20	101.30		189.00							92.00	262.00	
2023	213	8	0							77.00	19.30	101.30		296.0							128.00	332.00	
2023	213	9	0							73.00	20.20	101.30		398.00							160.00	380.00	
2023	213	10	0							69.00	21.20	101.30		478.0							184.00	407.00	
2023	213	11	0							66.00	21.80	101.30		525.00							200.90	412.00	
2023	213	12	0							64.00	22.20	101.30		538.0							206.00	410.00	
2023	213	13	0							63.00	22.50	101.30		525.00							208.00	397.00	
2023	213	14	0							62.00	22.60	101.30		479.0							192.00	383.00	
2023	213	15	0							62.00	22.50	101.30		408.00							169.50	359.00	
2023	213	16	0							63.00	22.10	101.30		321.0							138.00	333.00	
2023	213	17	0							66.00	21.50	101.30		223.00							102.00	289.00	
2023	213	18	0							69.00	20.60	101.30		127.0							65.00	224.00	
2023	213	19	0							74.00	19.40	101.30		49.00							29.00	132.00	
2023	213	20	0							80.00	18.10	101.30		6.70							5.00	33.00	
2023	213	21	0							83.00	17.20	101.30		0.00							0.00	0.00	
2023	213	22	0							86.00	16.60	101.30		0.00							0.00	0.00	
2023	213	23	0							86.00	16.80	101.30		0.00							0.00	0.00	

Table A.5: Meteorological data for an average day with air temperatures in the range of 20-25

Year	ID	Hour	Imin	Q*	QH	QE	Qs	Qf	Wind	RH	Td	Press	Rain	Kdn	Snow	Ldow	Fcld	Wuh	Xsmc	Lai_hr	Kdiff	Kdir	Wd
2023	213	0	0							87.00	16.50	101.30		0.00							0.00	0.00	
2023	213	1	0							88.00	16.20	101.30		0.00							0.00	0.00	
2023	213	2	0							89.00	15.90	101.30		0.00							0.00	0.00	
2023	213	3	0							90.00	15.70	101.30		0.00							0.00	0.00	
2023	213	4	0							90.00	15.60	101.30		1.70							1.50	3.90	
2023	213	5	0							89.00	16.00	101.30		28.00							19.00	84.00	
2023	213	6	0							87.00	17.00	101.30		86.00							53.00	187.00	
2023	213	7	0							82.00	18.20	101.30		189.00							92.00	262.00	
2023	213	8	0							77.00	19.30	101.30		296.0							128.00	332.00	
2023	213	9	0							73.00	20.20	101.30		398.00							160.00	380.00	
2023	213	10	0							69.00	21.20	101.30		478.0							184.00	407.00	
2023	213	11	0							66.00	21.80	101.30		525.00							200.90	412.00	
2023	213	12	0							64.00	22.20	101.30		538.0							206.00	410.00	
2023	213	13	0							63.00	22.50	101.30		525.00							208.00	397.00	
2023	213	14	0							62.00	22.60	101.30		479.0							192.00	383.00	
2023	213	15	0							62.00	22.50	101.30		408.00							169.50	359.00	
2023	213	16	0							63.00	22.10	101.30		321.0							138.00	333.00	
2023	213	17	0							66.00	21.50	101.30		223.00							102.00	289.00	
2023	213	18	0							69.00	20.60	101.30		127.0							65.00	224.00	
2023	213	19	0							74.00	19.40	101.30		49.00							29.00	132.00	
2023	213	20	0							80.00	18.10	101.30		6.70							5.00	33.00	
2023	213	21	0							83.00	17.20	101.30		0.00							0.00	0.00	
2023	213	22	0							86.00	16.60	101.30		0.00							0.00	0.00	
2023	213	23	0							86.00	16.80	101.30		0.00							0.00	0.00	

Table A.6: Meteorological data for an average day with air temperatures in the in the range of 25-30.

Year	ID	Hour	Imin	Q*	QH	QE	Qs	Qf	Wind	RH	Td	Press	Rain	Kdn	Snow	Ldow	Fcld	Wuh	Xsmc	Lai_hr	Kdiff	Kdir	Wd
2023	213	0	0							76.00	19.20	101.55		0.00							0.00	0.00	
2023	213	1	0							78.00	18.60	101.52		0.00							0.00	0.00	
2023	213	2	0							79.00	18.10	101.58		0.00							0.00	0.00	
2023	213	3	0							80.00	17.70	101.50		0.00							0.00	0.00	
2023	213	4	0							82.00	17.50	101.50		2.70							2.20	9.90	
2023	213	5	0							79.00	18.50	101.50		42.50							23.50	178.20	
2023	213	6	0							71.00	20.80	101.50		144.2							57.20	377.69	
2023	213	7	0							64.00	23.10	101.49		276.90							87.20	505.60	
2023	213	8	0							56.00	25.30	101.49		419.8							108.60	604.30	
2023	213	9	0							51.00	27.10	101.48		552.90							123.90	672.10	
2023	213	10	0							46.00	28.80	101.45		661.0							133.90	716.80	
2023	213	11	0							42.00	29.90	101.42		732.70							142.70	737.50	
2023	213	12	0							39.00	30.80	101.38		759.7							145.60	742.50	
2023	213	13	0							37.00	31.30	101.35		743.60							144.60	735.00	
2023	213	14	0							37.00	31.50	101.32		670.8							141.60	692.80	
2023	213	15	0							37.00	31.30	101.30		564.30							133.00	635.80	
2023	213	16	0							39.00	30.80	101.29		441.2							118.40	571.80	
2023	213	17	0							41.80	29.80	101.26		301.90							97.20	475.30	
2023	213	18	0							47.00	28.30	101.26		171.8							70.00	355.10	
2023	213	19	0							55.00	28.70	101.27		67.30							36.40	201.70	
2023	213	20	0							63.00	26.70	101.30		9.30							6.70	51.80	
2023	213	21	0							68.00	24.20	101.32		0.00							0.00	0.00	
2023	213	22	0							71.00	22.80	101.32		0.00							0.00	0.00	
2023	213	23	0							74.00	21.50	101.60		0.00							0.00	0.00	

Table A.7: Meteorological data for an average day with air temperatures in the 30+ range.

Year	ID	Hour	Imin	Q*	QH	QE	Qs	Qf	Wind	RH	Td	Press	Rain	Kdn	Snow	Ldow	Fcld	Wuh	Xsmc	Lai_hr	Kdiff	Kdir	Wd
2023	172	0	0							95.00	23.47	102.20		0.00							0.00	0.00	
2023	172	1	0							97.00	22.74	102.20		0.00							0.00	0.00	
2023	172	2	0							97.00	22.44	102.20		0.00							0.00	0.00	
2023	172	3	0							97.00	21.90	102.18		0.00							0.00	0.00	
2023	172	4	0							97.00	21.73	102.18		9.00							7.00	40.00	
2023	172	5	0							95.00	22.43	102.18		84.00							47.34	306.34	
2023	172	6	0							90.00	24.60	102.18		210.0							92.68	494.00	
2023	172	7	0							82.00	26.20	102.19		359.34							131.00	621.36	
2023	172	8	0							74.00	28.97	102.18		510.0							165.68	713.34	
2023	172	9	0							67.00	31.17	102.17		645.36							191.40	770.68	
2023	172	10	0							63.00	32.97	102.17		753.0							229.76	809.00	
2023	172	11	0							58.00	34.37	102.10		825.68							232.88	832.68	
2023	172	12	0							56.00	35.10	102.06		851.6							239.36	845.34	
2023	172	13	0							54.00	35.67	102.00		837.02							251.20	837.38	
2023	172	14	0							55.00	35.94	101.97		779.0							223.38	825.68	
2023	172	15	0							61.00	36.24	101.94		681.00							217.68	798.34	
2023	172	16	0							67.00	35.80	101.90		552.0							191.34	747.68	
2023	172	17	0							74.00	34.37	101.85		407.68							151.70	671.02	
2023	172	18	0							81.00	33.27	101.85		256.3							112.06	552.04	
2023	172	19	0							87.00	30.84	101.87		121.34							65.34	530.60	
2023	172	20	0							90.00	28.33	101.90		26.00							16.00	380.04	
2023	172	21	0							91.00	26.84	101.91		0.00							0.00	190.06	
2023	172	22	0							93.00	25.67	101.91		0.00							0.00	0.00	
2023	172	23	0							94.00	24.00	102.23		0.00							0.00	0.00	

Table A.8: Meteorological data for the Extreme Day in the 30+ temperature range.

A.6. Typologies comparison plots

Extreme Day

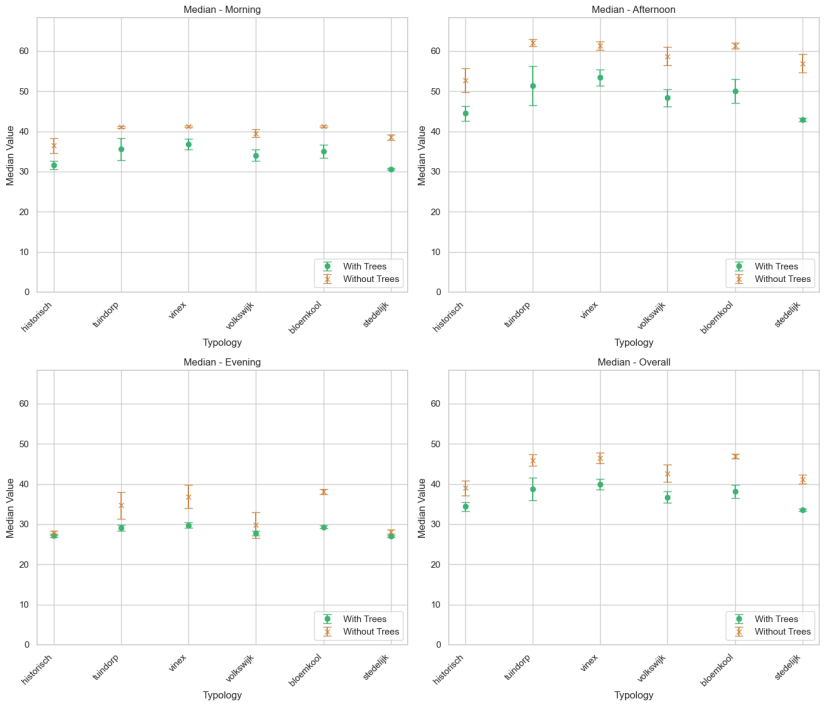

Figure A.6: Variations in median T_{mrt} across the six urban typologies on the 'extreme day', shown for morning, afternoon, evening, and the overall daily average.

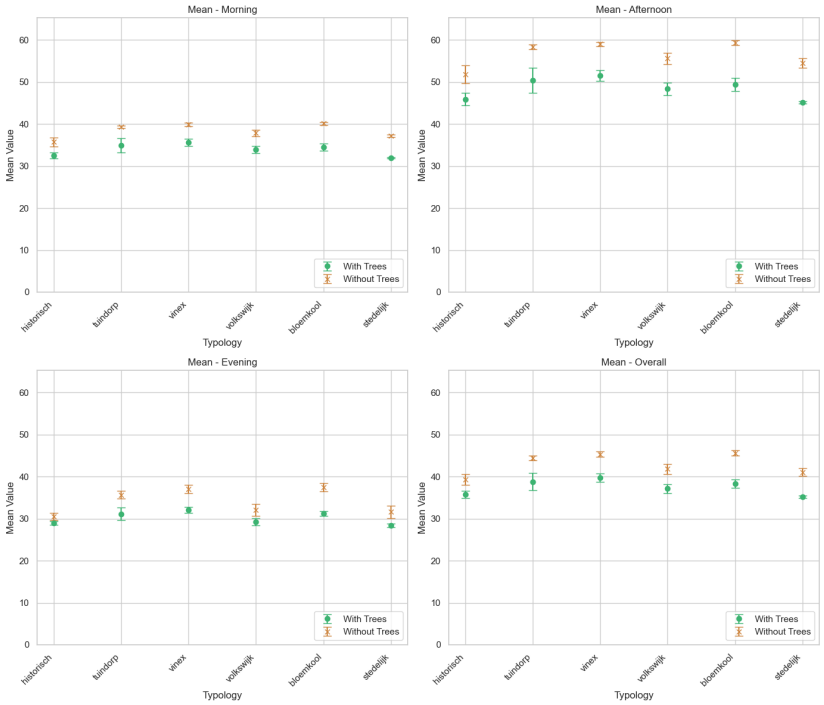

Figure A.7: Variations in mean T_{mrt} across the six urban typologies on the 'extreme day', shown for morning, afternoon, evening, and the overall daily average.

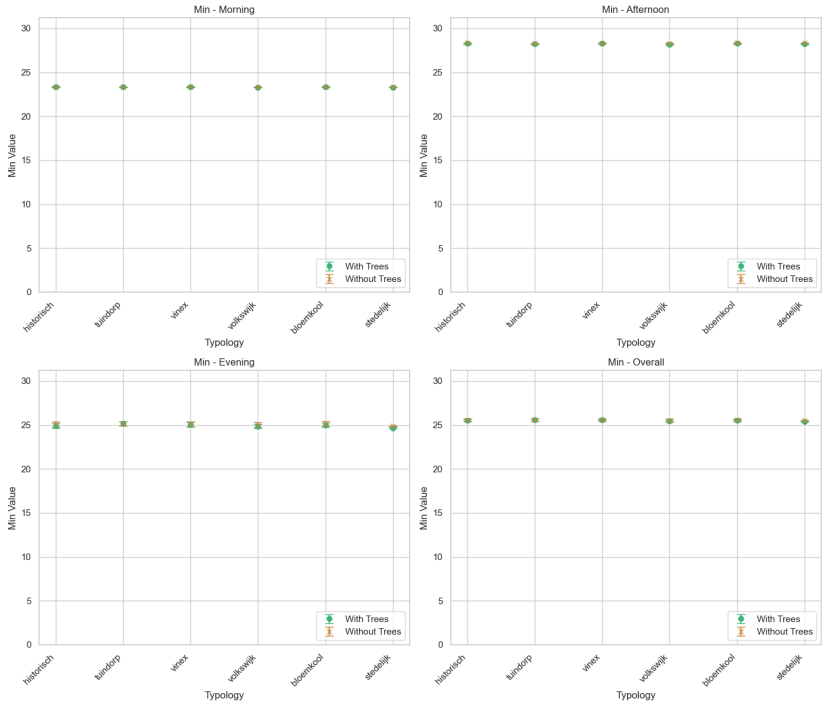

Figure A.8: Variations in minimum T_{mrt} across the six urban typologies on the 'extreme day', shown for morning, afternoon, evening, and the overall daily average.

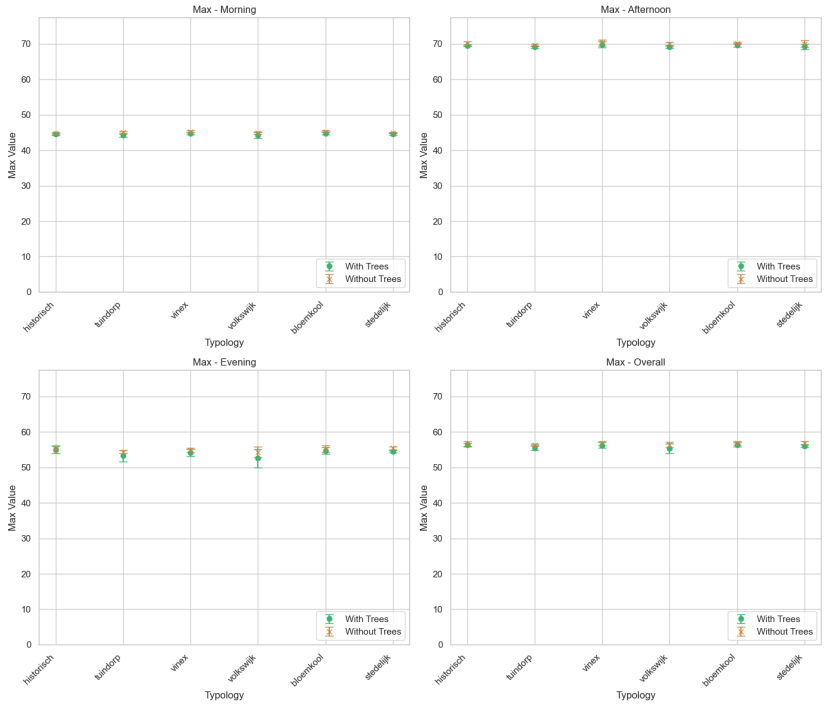

Figure A.9: Variations in maximum T_{mrt} across the six urban typologies on the 'extreme day', shown for morning, afternoon, evening, and the overall daily average.

Average Day

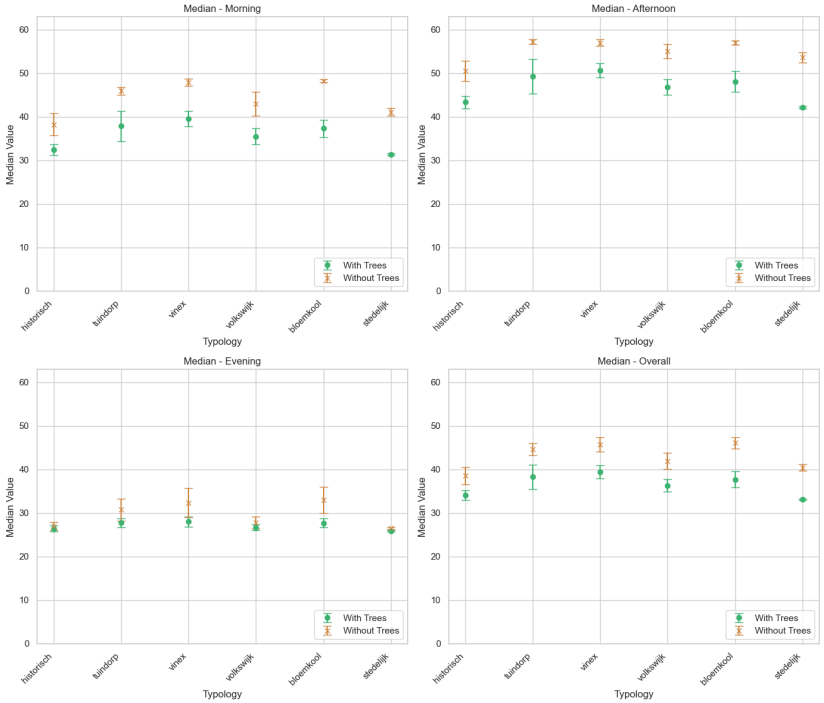

Figure A.10: Variations in median T_{mrt} across the six urban typologies on the 'average day', shown for morning, afternoon, evening, and the overall daily average.

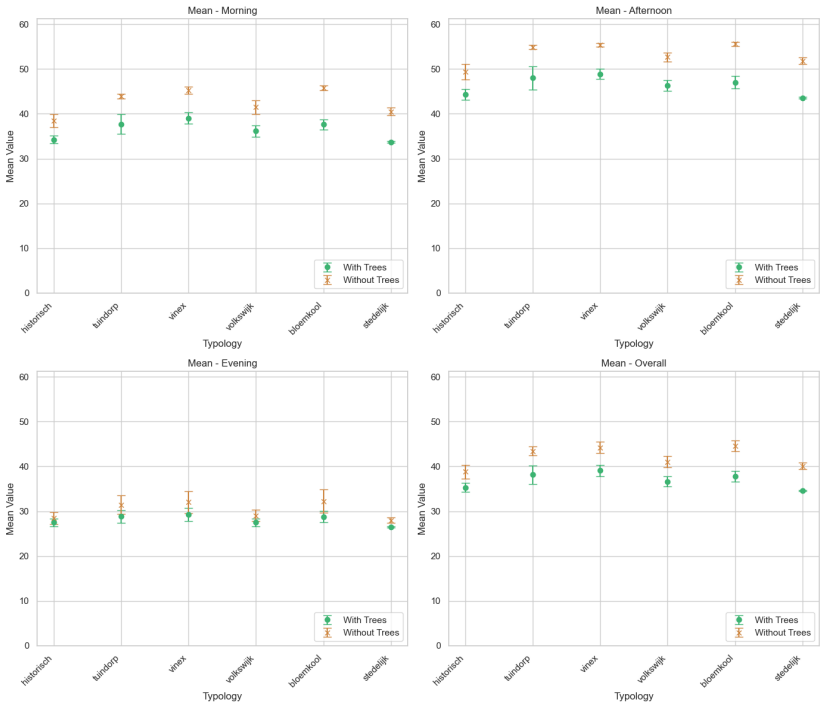

Figure A.11: Variations in mean T_{mrt} across the six urban typologies on the 'average day', shown for morning, afternoon, evening, and the overall daily average.

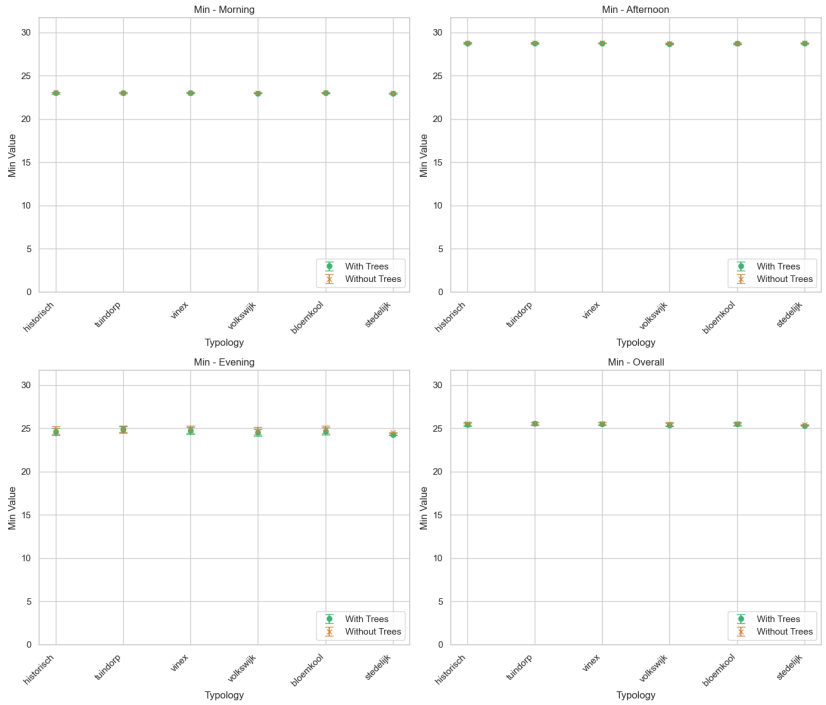

Figure A.12: Variations in minimum T_{mrt} across the six urban typologies on the 'average day', shown for morning, afternoon, evening, and the overall daily average.

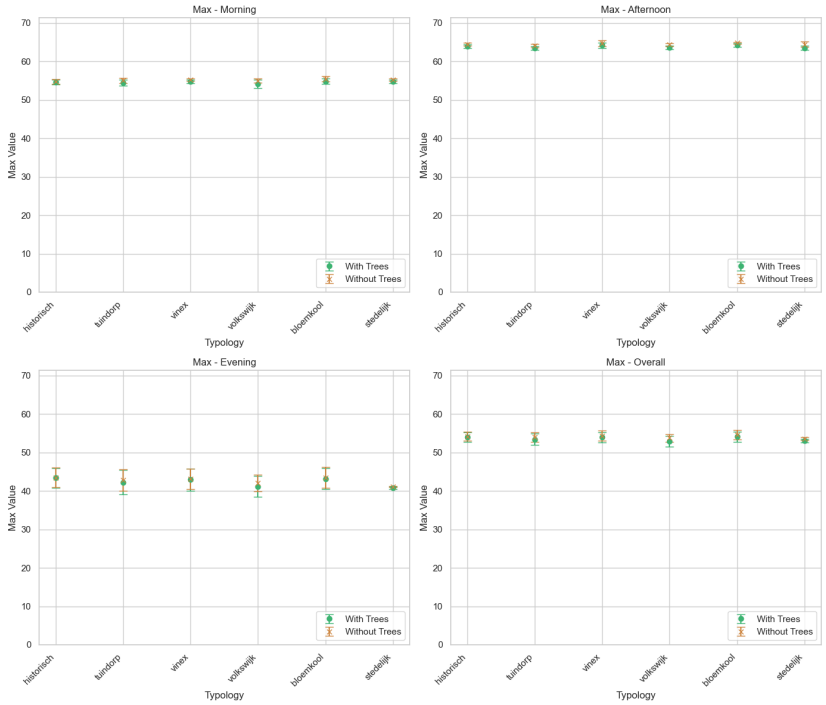

Figure A.13: Variations in maximum T_{mrt} across the six urban typologies on the 'average day', shown for morning, afternoon, evening, and the overall daily average.

B

Datasets and Reproducibility-assessment

B.1. Datasets

Table B.1 gives an overview of all the datasets that were used in this thesis, their sources and what they were used for.

B.2. Reproducibility assessment

This section reflects on the reproducibility of the research presented in this thesis by considering the input data, applied methods, and resulting outputs.

Input Data Almost all datasets used for SOLFD and in this thesis (as summarized in Table B.1) are openly accessible under public licenses, allowing for reuse and adaptation, sometimes with the requirement for attribution. Much of the data originates from Dutch government sources and open data platforms, which ensures the data remain retrievable as long as these services are maintained. The datasets used for the tree database are both available through repositories with permanent DOI links, supporting stable referencing and reuse.

On the less open side, the climate bike sensor dataset has not been used in other research and is only available by personal request from the researcher. This limits reproducibility of the validation results. The Wijktypologieën database is only accessible as a layer in ArcGIS and cannot be downloaded as a standalone dataset. Therefore, it was only used to visually verify typologies and to manually define location boundaries for this research.

Methods The original datasets and preprocessing steps performed on them will be made available online in the GitHub repository for this thesis. All implementations of the methodology into the code of SOLFD are also available there, accompanied by API documentation and examples explaining how to run the code.

Results The main result of this thesis is the code for SOLFD, which, as mentioned earlier, is openly available with documentation on GitHub. The raw outputs of the simulations that were run and used for the results will also be provided in the repository. Primarily, only SOLFD is needed to reproduce the results. However, the results presented in this thesis mainly serve as examples demonstrating how the tool functions and may not be of significant interest to fully recreate.

The only part that is not openly accessible is the validation. Due to the inaccessibility of the climate bike sensor dataset, the validation results can only be described, and the actual raw measurements cannot be shared.

Dataset	Source	Description	Usage
AHN DTM	https://service.pdok.nl/rws/ahn/wcs/v1_0	Digital Height map of the whole of Netherlands, containing only ground heights	Creation of DTM, DSM and CHM
AHN DSM	https://service.pdok.nl/rws/ahn/wcs/v1_0	Digital Height map of the whole of Netherlands, containing all measured heights	Creation of DSM
GeoTiles	https://geotiles.citg.tudelft.nl/	Tiled .LAZ pointclouds of the whole of Netherlands, overlayed with aerial photographs for RGB	Creation of CHM
3DBAG	https://data.3dbag.nl/api/BAG3D/wfs	Dataset of 3D buildings in the Netherlands. The 2D LoD 1.3 is used	Creation of DSM, Land cover
BGT	https://api.pdok.nl/lv/bgt/ogc/v1	Basisregistratie Grootschalige Topografie (BGT), the detailed large-scale digital map of the whole of the Netherlands	Creation of Land cover
TOP10NL	https://api.pdok.nl/brt/top10nl/ogc/v1	Object-oriented topographical dataset of the Netherlands, ranging from 1:5,000 to 1:25,000 scale	Creation of Land cover
Urban Tree Database	https://www.fs.usda.gov/rds/archive/Catalog/RDS-2016-0005	Urban tree growth data collected over a period of 14 years (1998–2012) in 17 cities from 13 states across the United States	Creation of Tree database
Urban Tree Allometric Equations	https://doi.org/10.5281/zenodo.15593687	Records of trees in the UK with parameters as tree height, age, crown diameter and stem width.	Creation of Tree database
Shiny Weather Data	https://www.shinyweatherdata.com/	ERA5 meteorological data processed by shinyweather data from 2000 to 2024 for the location at 52.25°N latitude and 5.5°E longitude	Creation of Meteorology datasets
Amsterdam Climate Bike	Private correspondence with G.J. Steeneveld, WUR	Radiation measurements (longwave and shortwave) and their locations collected with the Climate Bike in Amsterdam on August 23 and September 12, 2023. Air temperature and Relative humidity measurements from this dataset are also used	SOLFD accuracy evaluation
MAQ Observations	https://maq-observations.nl/data-downloads/	AAMS meteorological dataset with minute-resolution air temperature and global shortwave radiation measurements	SOLFD accuracy evaluation: meteorology dataset
Schiphol - uurgegevens weer	https://www.knmi.nl/nederland-nu/klimatologie/uurgegevens	Hourly meteorological data from KNMI the Schiphol Climate Station, including air pressure and relative humidity	SOLFD accuracy evaluation: meteorology dataset
Wijktypologie PC6 Gebieden 2023	https://tiles.arcgis.com/tiles/nSZVuSZjHpEZZbRo/arcgis/rest/services/Wijktypologie_PC6_1/MapServer	Dutch Postal code areas classified into Klimaatatlas typologies.	Urban Case Study

Table B.1: Overview of datasets used, their sources, descriptions, and applications within the SOLFD pipeline.