MSc thesis in Geomatics

A Comparison of Wind Velocity Predictions from Wind-Only and Thermally Coupled OpenFOAM Solvers in Urban Models with Complex Terrain

Carmem Aires

2026

MSc thesis in Geomatics

A Comparison of Wind Velocity Predictions from Wind-Only and Thermally Coupled OpenFOAM Solvers in Urban Models with Complex Terrain

Carmem Aires

June 2026

A thesis submitted to the Delft University of Technology in partial fulfilment of the requirements for the degree of Master of Science in Geomatics

Carmem Aires: *A Comparison of Wind Velocity Predictions from Wind-Only and Thermally Coupled OpenFOAM Solvers in Urban Models with Complex Terrain* (2026)

The work in this thesis was carried out in the:

3D geoinformation group
Delft University of Technology

Supervisors: Prof.dr. Clara Garcia Sanchez
Dr. Themistoklis Vargiemezis
Frank van Gool (Haskoning)
Co-reader: Dr. Akshay Patil

Abstract

Climate change is intensifying urban heat stress, and microclimate simulation can support climate-responsive planning. Computational fluid dynamics (CFD) is the dominant approach at the microscale, but urban CFD studies are predominantly isothermal, neglecting the buoyancy effects. Yet these effects become non-negligible under the low-wind, thermally-active conditions that climate change is making more frequent. Coupled solvers such as urbanMicroclimateFoam (uMF), built on OpenFOAM, resolve the airflow, heat and moisture transport, radiation, and vegetation interactions that wind-only formulations omit. However, it remains unclear under which conditions a wind-only simulation is sufficient and when the additional complexity of a coupled solver is justified. Addressing this question is hindered by three further gaps: published applications of uMF to realistic urban sites with complex terrain are scarce, its computational cost are not well characterised, and rigorous validation against field measurements is limited.

To address these gaps, this thesis asks three questions: whether uMF can be feasibly applied to realistic urban models with non-flat terrain; how its mesh requirements, solver parameters, and computational demands compare to those of the wind-only solver simpleFoam (SF) on the same case; and how the two solvers' wind-speed predictions compare against street-level field measurements during a representative heatwave hour.

Applied to the Carnegie Mellon University campus in Pittsburgh, USA, during one hour of a heatwave on 27 August 2024, with a 3D model reconstructed in City4CFD from LiDAR and building footprints and validated against four MORICHI street-level stations, three findings emerge. First, applying urbanMicroclimateFoam (uMF) to realistic terrain is feasible but requires time-consuming, case-specific geometric adaptations compared to geometry-robust SimpleFOAM (SF). Second, on identical geometry and 20 cores, uMF took 3.6× longer than SF (6 h 28 min versus 1 h 49 min), which does not represent a barrier to the usage of the solver. Third, uMF improved wind-speed agreement across every aggregate indicator: FB moved from $+0.65$ to -0.11, MG from 1.89 to 0.83, NMSE was reduced by 65%, and FAC2 rose from 0.50 to 0.75. These results indicate that uMF merits adoption when buoyancy effects are non-negligible and the analyst can absorb the case-setup overhead. The findings rest on a single one-hour window under one inflow direction, with the observed wind speeds at three of the four stations falling below the anemometer accuracy threshold, and vegetation modelled as a porous medium. Addressing these limitations is identified as a priority for future work.

Acknowledgements

I would like to begin by thanking Clara and Frank for their invaluable guidance. Your supervision was always very helpful, pragmatic and insightful, and I have learned a great deal from working with you both. Thank you, Miguel, for the co-supervision during the first months of this project and for helping me get started, and Themis, for your follow-up and constructive feedback.

I am grateful to my colleagues at Haskoning and to my fellow students in the MSc Geomatics programme for the discussions, the shared struggles, and the sense of community that made these past years much lighter. My thanks also go to the Geomatics staff, from whom I have learned so much, and especially to Gina, for her support with the Foote server.

I am also grateful for my friends in The Netherlands, Brazil and Germany, who have supported me and reminded me to relax and enjoy life outside of work and studies. I am deeply thankful to my family for their unconditional love and encouragement, which has been a constant source of strength and confidence no matter the distance.

Finally, I dedicate this work to you, Guillaume, my wonderful and beloved partner. Thank you for your love and unconditional patience throughout the entire process of this master's journey. Thank you for your constant support of every endeavour I take on. Everything makes more sense alongside you, and none of this would have been possible without you.

Contents

List of Figures

List of Tables

Acronyms

CFD	Computational Fluid Dynamics	xi
RANS	Reynolds-Averaged Navier–Stokes	xi
LES	Large Eddy Simulation	25
URANS	Unsteady Reynolds-Averaged Navier–Stokes	23
uMF	urbanMicroclimateFoam	v
SF	SimpleFOAM	v
HAM	Heat, Air, Moisture	xi
RAD	Radiation	15
VEG	Vegetation	xi
LAD	Leaf Area Density	15
RMSE	Root Mean Square Error	25
LoD	Level of Detail	xii
SIMPLE	Semi-Implicit Method for Pressure-Linked Equations	15
UTCI	Universal Thermal Climate Index	6
PET	Physiological Equivalent Temperature	6
UHI	Urban Heat Island	1
ABL	Atmospheric Boundary Layer	5
GCI	Grid Convergence Index	34

1. Introduction

Cities are getting hotter. Global mean surface temperature in 2024 reached 1.60 °C above the pre-industrial baseline, making it the first calendar year to exceed the 1.5 °C threshold of the Paris Agreement (Copernicus Climate Change Service (C3S), 2025), and dense built environments amplify this warming through the Urban Heat Island (UHI) effect (Oke, 2017; Roy et al., 2025). At the same time, an ever larger share of the world's population lives in cities and is therefore directly exposed to heat-related health risks and reduced thermal comfort (World Health Organization, 2024; UN-Habitat, 2024). Closing the urban climate data gaps to support informed planning, design, and adaptation has become a central concern of climate action (UN-Habitat, 2024).

Microclimate simulation tools are one valuable instrument for closing these gaps: they make it possible to evaluate how a given urban configuration will respond thermally before it is built, and to compare alternative design or adaptation strategies under realistic atmospheric conditions (Toparlar et al., 2017; Blocken, 2015). Among the available methods, CFD has become the dominant approach for studies at the microscale (under 2 km). The vast majority of urban CFD studies, however, treat the atmosphere as isothermal: they solve only for mass and momentum and neglect temperature altogether. This wind-only assumption is acceptable when wind forcing dominates, but it loses validity when buoyancy effects (vertical air motion driven by temperature-induced density differences) become comparable to wind shear, which is precisely the regime that extreme heat events and calm summer days bring about. As climate change increases the frequency of such conditions, the limitations of the wind-only approach become more consequential.

In parallel, two complementary fields have been evolving rapidly. On the geomatics side, reconstruction pipelines that turn real-world data such as LiDAR point clouds and 2D building footprints into simulation-ready 3D geometries is evolving towards full automation (Pađen et al., 2022, 2024), making it increasingly practical to model real urban sites rather than idealised geometries. On the CFD side, urban microclimate research has expanded tenfold since 2010 (Roy et al., 2025), and open-source solvers have begun to offer the physical complexity that closed wind-only formulations cannot. uMF (Kubilay et al., 2017), developed within the OpenFOAM framework, is one such solver: it couples airflow, heat and moisture transport in porous building materials, short- and long-wave radiation, and vegetation, providing a multiphysics approach to the urban microclimate.

Despite this promise, published applications of uMF to realistic urban sites including complex terrain are scarce, the additional preprocessing and computational cost relative to a wind-only baseline are not well characterised, and rigorous validation against high-resolution field measurements is limited. Closing these gaps would allow us to determine when the added complexity of a coupled solver is justified in practice.

This thesis addresses three questions:

1. How feasible is the application of uMF to realistic urban models that include non-flat terrain, given current geometric reconstruction and meshing tools?

2. How do mesh requirements, solver parameters, and computational demands differ between uMF and the wind-only solver SF when applied to the same case study?

3. How do wind-speed predictions from SF compare to those of uMF when both are evaluated against field measurements from four stations over a one-hour window under a single inflow condition during a representative heatwave?

The Carnegie Mellon University campus in Pittsburgh, USA, is used as a case study, with reference observations drawn from the MORICHI dataset (Martin et al., 2026), acquired during a heatwave in August 2024. The validation is carried out over a single one-hour window, against four street-level stations under one inflow direction, with vegetation represented as a porous medium.

From a disciplinary standpoint, this work is grounded in geomatics, whose core competencies are precisely those required by CFD, and specifically by uMF pipelines. The study draws on the full chain of geomatics operations: acquiring heterogeneous geographic data, transforming and integrating it through automated 3D reconstruction pipelines, debugging the resulting models, running physical simulations on them, and analysing, visualising, and statistically validating the predictions against field measurements. The contribution of the field to CFD applications is especially relevant because CFD is unusually sensitive to geometric malformations compared with most other uses of 3D data. Details that are harmless for, for instance, volume calculations can produce can produce invalid mesh cells, unphysical values, and numerical instabilities, any of which can break the solver, and uMF is especially demanding in this regard. Catching such defects upstream, at the geometry stage, yields a disproportionate saving in time, since a single simulation run can take hours or days to reach the point at which the defect manifests. Both fields, geomatics and CFD, therefore stand to gain from this exchange.

Three findings emerged from this study. First, applying uMF to realistic terrain is feasible but requires time-consuming, case-specific geometric adaptations compared to geometry-robust SF. Second, on identical geometry and 20 cores, uMF took 3.6× longer than SF (6 h 28 min versus 1 h 49 min), which does not represent a barrier to the usage of the solver. Third, uMF improved wind-speed agreement across every aggregate indicator: FB moved from $+0.65$ to -0.11, MG from 1.89 to 0.83, NMSE was reduced by 65%, and FAC2 rose from 0.50 to 0.75. These results indicate that uMF merits adoption when buoyancy effects are non-negligible and the analyst can absorb the case-setup overhead. The findings rest on a single one-hour window under one inflow direction, with the observed wind speeds at three of the four stations falling below the anemometer accuracy threshold, and vegetation modelled as a porous medium. Addressing these limitations is identified as a priority for future work.

This work contributes (i) a documented catalogue of the geometric adaptations required to apply uMF to a realistic urban model with complex terrain; (ii) a quantitative comparison of SF and uMF on identical geometry, covering geometry requirements, solver setup, and computational cost; (iii) evidence that uMF improves aggregate wind-speed agreement with field measurements in a low-wind, buoyancy-influenced regime; (iv) a methodology for validating and evaluating isothermal and non-isothermal solvers against high-resolution field measurements during a heatwave; and (v) a set of practical recommendations for advancing uMF-based urban microclimate workflows, from geometric reconstruction through preprocessing, solver behaviour, and validation.

Structure of the thesis

This thesis is composed of three main parts: Preliminary materials, the Article, and the closing Discussions. A reader familiar with urban microclimate studies, CFD and uMF may proceed directly to the Article; readers from a broader background are encouraged to read the Preliminary materials first. For those interested in advancing this topic further or in understanding its link with geomatics, please refer to the Discussions.

Preliminary materials This part provides the technical background needed to follow the scientific article. It introduces the physical processes governing urban microclimate, the principles of CFD modelling, the specific coupling of heat, air, moisture, radiation, and vegetation that distinguishes uMF from a wind-only solver such as SF, and the practical workflow of mesh generation, pre-processing, and solver steps in uMF. It also presents the datasets and tools used in the study, and reviews the related work that situates the present research within the broader literature on urban microclimate CFD.

Article This part is the scientific contribution of the thesis, written as a stand-alone article. It introduces the research problem and gaps, describes the methodology applied to the CMU case study, presents the results on feasibility, computational cost, and wind-speed validation, and discusses their implications for the practical use of coupled urban microclimate solvers.

Discussions This closing part reflects on how the work connects to the MSc Geomatics for the Built Environment, expands on specific suggestions for future work that fall outside the scope of the article, particularly regarding the radiation pre-processing chain, the geometric reconstruction pipeline, and the solver itself, and discusses the general challenges encountered during the project.

2. Preliminary materials

This chapter presents the background information needed to understand the core of this work in Chapter 3. It begins with the broader context of urban climate and microclimate (Section 2.1), including the typical applications of microclimate simulation and the different approaches available to study it. It then narrows down to the methodology used in this work (Section 2.2), covering the fundamentals of urban CFD, the main modelling choices (RANS turbulence modelling, wind-only versus thermally coupled formulations, HAM transport in porous materials, and radiation), the practical setup of a CFD case, and validation methods. The specific tools and datasets employed are presented next (Sections 2.3 and 2.4), including a broader overview of uMF. The chapter closes with a review of related work in CFD-based urban microclimate simulation (Section 2.5), covering the current state of the field, best practices, practical examples, and existing applications of uMF.

2.1. Urban climate and Microclimate Studies

Strong urbanization over the past century has been capable of changing local, regional, and even global climates. Urban climate describes the physics of these human-modified environments, or as Oke (2017) put it, ”urban effects on climate”.

Urban climate develops within the Atmospheric Boundary Layer (ABL), the lowest layer of the atmosphere directly influenced by the Earth’s surface, whose thickness varies from about 100 m to 3 km depending on time and location (Stull, 1988; Oke, 2017). The ABL is subdivided into two parts. The outer part, called the mixed layer, is where the atmosphere is well mixed and properties such as temperature, humidity, and wind are nearly uniform. The inner part, which makes up roughly the lowest 10%, is called the surface layer; this is where surface friction, heat exchange, and human activity have the strongest direct effect on the flow. Over a city, the ABL is termed the urban boundary layer, which may remain confined above the city as an urban dome under calm conditions, or be advected as an urban plume when the wind is stronger (Oke, 2017).

Inside the urban boundary layer, surface friction from buildings, heat exchange with materials, anthropogenic emissions (from traffic, combustion, and air conditioning), and evapotranspiration from vegetation directly affect the atmospheric properties. The resulting urban climate is shaped by two types of controls. The extrinsic controls are inherent to the geographic location (latitude, altitude, relief, proximity to water bodies) and cannot be changed by urban design. The intrinsic controls are those introduced by the city itself – the land cover, the urban fabric, the three-dimensional structure of the built environment, and the emission of pollutants and heat (metabolism) (Oke, 2017).

The urban surface influences the atmosphere through several distinct physical effects: radiative exchange, thermal exchange, moisture exchange, and aerodynamic interaction. Radiative exchange involves the absorption, reflection, transmission, and emission of shortwave and longwave radiation, and depends on surface geometry, albedo, transmissivity, and emissivity. Thermal exchange involves conduction into the substrate, heat storage, and convective release to the air, and is governed by specific heat, heat capacity, thermal conductivity, diffusivity, and admittance. Moisture exchange involves interception, storage, evaporation, and transpiration, and depends on interception and storage capacity, permeability, stomatal characteristics, and chemical nature. Aerodynamic interaction involves drag, turbulence generation, and flow channelling, and is characterised by roughness, zero-plane displacement, and porosity. Different urban materials (asphalt, concrete, brick, vegetation) exhibit very different combinations of these properties (Oke, 2017), so the choice and arrangement of materials directly shapes the magnitude of each effect. In particular, the urban structure defines the turbulence pattern by elements in the surface, which is characterised by the aerodynamic roughness. See Section 2.2.

The phenomena described above span the micro, local, and meso scales, which Oke (2017) define as ranging from 0.01 m to about 100 km. Blocken (2015) follow the American Meteorological Society classification (American Meteorological Society, 2026), which distinguishes three regimes: the macroscale for atmospheric motions spanning many hundreds of kilometres (e.g. cyclones); the mesoscale for motions of a few hundred kilometres (e.g. thunderstorms); and the microscale for motions of 2 km or less. The microscale is formally defined as the regime in which the Lagrangian Rossby number exceeds 200, i.e. where inertial forces dominate over those induced by the Earth’s rotation (Coriolis forces), so that the latter can be neglected. In this work, microclimate refers to phenomena within this 2 km microscale range.

2.1.1. Applications of microclimate studies

Microclimate simulation supports a wide range of applications. They can play a central role in the early-design and planning stages, providing quantitative evidence to support decisions for future planning and adaption strategies (Toparlar et al., 2017; Mei and Yuan, 2022; Blocken, 2015; Tominaga et al., 2008).

- UHI assessment quantifies how much warmer the city is than the surrounding rural area, which materials and street configurations drive this warming, and how it interacts with regional heatwaves;
- Pedestrian wind comfort and safety assesses whether the wind around planned or existing buildings is acceptable for walking, sitting, or cycling;
- Outdoor thermal comfort predicts how pedestrians experience temperature, humidity, wind, and radiation together at street level, often using indices such as Universal Thermal Climate Index (UTCI) or Physiological Equivalent Temperature (PET);
- Air quality and pollutant dispersion predicts how traffic emissions, industrial plumes, or accidental releases spread through streets and where they accumulate;
- Building energy demand assesses how the microclimate around a building (wind, shading, surface temperatures) directly affects its heating, cooling, and ventilation loads, and therefore its energy use and carbon emissions;
- Climate adaptation and urban design compares different design alternatives (greenery, reflective materials, building geometry, urban morphology) to identify which ones effectively cool the city, improve ventilation, or reduce pollutant exposure.

- Building envelope and durability applications address phenomena such as wind-driven rain on façades, snow drift around buildings, and the transport of sand or debris, all of which affect material weathering, moisture damage, and structural loading.

The work presented in this thesis is most directly relevant to UHI assessment, outdoor thermal comfort, and climate adaptation and urban design, since it evaluates a coupled microclimate solver under heatwave conditions and compares it against a wind-only baseline. The methodology, however, is transferable to the other applications listed above, as all of them rely on the same underlying CFD machinery and benefit from the same investigation of when buoyancy and surface-energy effects must be resolved explicitly.

2.1.2. Types of microclimate studies

The urban microclimate can be studied through several complementary techniques, here the three main ones are listed: observational techniques, wind-tunnel experiments, and numerical simulations.

Observational techniques include field measurements and remote sensing. Field measurements observe the microclimate directly using fixed or mobile sensors, which are highly accurate but limited to point locations or trajectories. Remote sensing uses airborne or satellite sensors, providing larger spatial coverage at lower resolution. Both approaches are essential for characterising existing conditions and for validating other methods, but they cannot evaluate urban design alternatives nor alternative conditions (e.g. strong storms, heatwaves), and the instruments and field campaigns involved are often expensive and time-consuming (Oke, 2017; Roy et al., 2025; Toparlar et al., 2017; Mirzaei and Haghighat, 2010; Blocken, 2015).

Wind-tunnel experiments use a scaled physical model of the urban area placed in a controlled laboratory, where different fields are measured with techniques such as particle image velocimetry or tracer gases (Tominaga et al., 2008; Zhao et al., 2021; Roy et al., 2025). They can compare design alternatives and provide the reference data commonly used to validate numerical simulations (Oke, 2017; Blocken, 2015). Their main limitations are the construction cost of the physical model, the difficulty of satisfying all relevant similarity criteria at reduced scale, especially for non-isothermal flows, and the difficulty of accurately reproducing thermal stratification (Mirzaei and Haghighat, 2010).

Numerical simulations solve the governing equations on a computer, using a digital 3D model, and can simulate both existing and hypothetical scenarios. In the microscale, CFD is the most used tool. Meteorological microclimate phenomena are explicitly solved (e.g. turbulence, wakes, thermals). Typical domain sizes range from 100m to 5km. Resolution of cells range from a few centimetres to a few meters (about 0.1 to 100m) and explicit buildings, trees, terrain are modelled (Toparlar et al., 2017; Mirzaei and Haghighat, 2010; Blocken, 2015; Roy et al., 2025).

Data-driven approaches, especially machine learning, have recently emerged as a complementary technique (Roy et al., 2025). Rather than solving the governing equations, these methods learn statistical relationships between urban morphology, meteorological inputs, and microclimate variables from datasets generated by the other three techniques. Once trained, they produce predictions at a fraction of the cost of CFD, but their accuracy depends on the coverage of the training data and they extrapolate poorly beyond it.

This work focuses on microscale CFD for resolving the interaction between individual buildings, trees and terrain in a range within 2 km. The validation uses observational techniques.

2.2. Fundamentals of Urban Computational Fluid Dynamics

2.2.1. Governing equations and incompressibility

CFD is a computational technique that uses numerical analysis to predict how fluids flow and interact with surfaces (e.g. a building wall, a terrain). It is based on three fundamental physical principles: the conservation of mass, of momentum, and of energy. Conservation means that mass, energy and momentum do not simply appear or disappear. For an incompressible flow (where density does not vary), the air that flows into a region must flow out. Air is treated as incompressible, since the wind speeds in urban simulations are low (below 100 m/s) and pressure differences are small, resulting in very small density differences. Conservation of momentum is based on Newton's second law applied to a fluid. The change in a fluid's momentum equals the sum of forces acting on it. The forces acting on a fluid are pressure (high to low pressure), viscous forces (internal friction) and body forces (gravity). These forces are defined in the Navier-Stokes equations. And finally, the conservation of energy: energy cannot be created, it can only be converted between forms. Energy in a fluid changes only through heat transfer (hot to cold), pressure work (compression or expansion), viscous dissipation (friction converting motion into heat), or external sources (e.g. radiation) (Anderson, 1995; ANSYS, Inc., n.d.b).

2.2.2. The ABL and turbulence

Key characteristics of the ABL are constantly turbulent flows and the near logarithmic wind speed profile in the surface layer. Turbulence is random, chaotic, unsteady (flow quantities change with time) and three-dimensional. Turbulence is characterised by eddies, swirling structures of fluid motion of many different sizes; and they cause strong mixing of mass, momentum, and energy. Turbulence happens when inertial forces dominate over viscous forces, measured by the Reynolds number (Stull, 1988; Versteeg, 2007).

2.2.3. RANS and turbulence modelling

There is no closed-form mathematical solution for turbulence. For most engineering applications, however, time-averaged flow quantities are sufficient and it is unnecessary to resolve the details of the instantaneous structure of turbulence (Versteeg, 2007). The most common practical compromise is the RANS approach (Toparlar et al., 2017; Roy et al., 2025), in which all flow variables are decomposed into a mean and a fluctuating part, and the equations are then time-averaged. An intuitive analogue is shown in Figure 2.1: a long-exposure photograph of a waterfall blurs the chaotic, instantaneous motion of the water into a smooth, time-averaged field, which is conceptually what a RANS simulation produces for a turbulent flow. This averaging procedure produces additional unknown terms, the Reynolds stresses in the momentum equations, and analogous turbulent fluxes in the equations for scalars

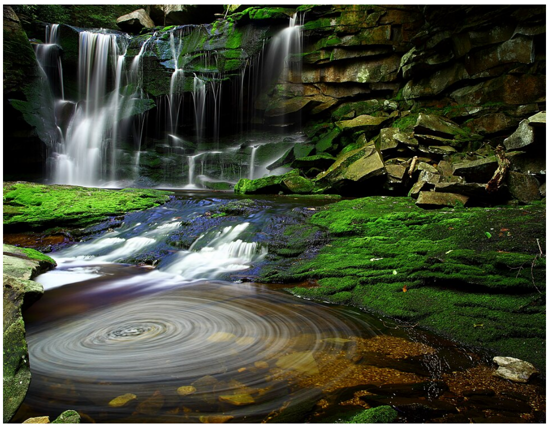

Figure 2.1.: Long-exposure photograph of a waterfall: the smooth, blurred water represents the time-averaged motion rather than the instantaneous turbulent structures. This provides a visual analogue to what a RANS simulation outputs. Photograph by Forest Wander

such as temperature. A turbulence model is then introduced to close the system by expressing these unknown terms as functions of the mean-flow quantities. The turbulence model represents the effect of turbulence on the mean flow (Anderson, 1995; Versteeg, 2007). $k-\epsilon$ is a widely used model in urban simulations, which solves two additional transport equations for turbulent kinetic energy (k) and its dissipation rate (ϵ) to estimate the eddy viscosity (Toparlar et al., 2017; Roy et al., 2025).

The realisable $k-\epsilon$ model was selected for this study, consistent with the tutorials of the tool used for the microclimate simulations, uMF. This option presents a balance between robustness and the relative ease of obtaining converged solutions on complex urban geometries.

2.2.4. Wind-only vs. thermally coupled

Regardless of the turbulence model chosen, an urban CFD simulation must also decide which governing equations are actually solved. In climatic analysis, it makes it possible to simulate and predict different properties of fluid flow, such as velocity, pressure, temperature, turbulence properties and pollutant concentration. In urban climate studies, two main types of simulations are commonly used: wind-only (isothermal) and thermally coupled (non-isothermal) CFD analyses.

In a wind-only analysis, the air is assumed to be at a uniform temperature throughout the domain. The energy equation is switched off, and only the conservation of mass and momentum are solved. This corresponds to a neutral atmospheric boundary layer (i.e. no significant thermal stratification, the layering of warm and cool air caused by surface heating), where wind-driven turbulence dominates and buoyancy effects (vertical air motion driven by temperature-induced density differences) are negligible. Wind-only simulations are computationally cheap. This approach is most commonly used, under non-calm wind conditions, to study pedestrian wind comfort and safety, wind loads on buildings for structural analysis, general urban ventilation under windy conditions, pollutant dispersion (Franke et al., 2007).

However, for UHI studies, the atmosphere is often not neutrally stratified, meaning that the temperature differences between the surface and the air in a higher position generate buoyancy forces that are no longer negligible, and these can drive local flow circulations when the wind is calm. The relative importance of mechanical (shear-driven) and buoyant (thermally-driven) forcing in the urban boundary layer can be quantified by the bulk Richardson number (Mei and Yuan, 2022; Zhao et al., 2021),

$$Ri_b = \frac{g\,H\,(\theta_a - \theta_g)}{\theta_a\,u_a^2}, \tag{2.1}$$

where g is gravitational acceleration, H a characteristic height (taken here as the average building height), θ_a and θ_g are the potential temperatures of the reference air and the ground surface respectively, and u_a is the reference wind speed at height H (Mei and Yuan, 2022).

Wind tunnel and numerical studies of urban street canyons (i.e. streets with buildings on either side) have shown that thermal stratification can substantially modify the flow. The swirling wind pattern that normally forms between two buildings (the canyon vortex) gradually weakens as the surfaces become hotter, and is eventually replaced by upward-rising plumes of warm air when the bulk Richardson number $|Ri_b|$ becomes of the same order as one, meaning that buoyancy is comparable to or stronger than wind forcing (Zhao et al., 2021). In this regime, the wind at pedestrian level, the way the street canyon ventilates itself, and the air exchange across the rooftops are all fundamentally altered relative to the neutral (unheated) case (Zhao et al., 2021; Mei and Yuan, 2022).

In a thermally coupled analysis, the energy equation is solved together with mass and momentum, and buoyancy is included as a body force in the momentum equation (Versteeg, 2007; Ferziger et al., 2020). Air density is allowed to vary with temperature, e.g. via the Boussinesq approximation, in which density variations are neglected everywhere except in the gravitational term of the momentum equation, where they drive the buoyancy force. This way, warm air rises and cool air sinks. Temperature becomes an active variable, coupled with the velocity, turbulence and pressure fields (Ferziger et al., 2020; Mei and Yuan, 2022).

Thermally coupled simulations require additional inputs, such as surface temperatures (or a surface energy balance model), solar radiation, material properties (albedo, emissivity, thermal conductivity, density, and specific heat), and thermal boundary conditions at walls and ground. They are also inherently transient: while wind-only simulations can in many cases be run as steady-state, the surface energy balance evolves with the diurnal cycle of solar radiation and the corresponding thermal storage and release by walls, roofs, and ground, so thermally coupled models must be advanced in time to model this behaviour. This combination of additional inputs, larger set of governing equations, and transient time-stepping makes thermally coupled simulations significantly more complex to set up, more computationally expensive, and more sensitive to input data quality than wind-only analyses.

2.2.5. Heat, Air, Moisture in solids

Urban surfaces themselves store and release heat and moisture over much longer timescales. Building walls, pavements, and soil are porous materials: they absorb solar radiation, conduct heat into their interior, store rainwater, and slowly release both back to the air through conduction and evaporation. This behaviour is governed by coupled HAM transport.

The two basic principles are conservation of energy and conservation of moisture mass within the porous material. Heat is transported through the solid structure, and by latent effects associated with phase change (evaporation and condensation of water inside the pores). Moisture is transported in two phases: liquid water moves through the pore network driven by capillary pressure, while water vapour diffuses driven by vapour pressure gradients. The two transport equations are coupled because temperature affects how much water vapour the air in the pores can hold, and phase change releases or absorbs latent heat. Each material is characterised by a set of properties: density, specific heat, thermal conductivity, porosity, and moisture-retention and -permeability curves that describe how much water the material holds at a given capillary pressure and how easily it moves (Janssen et al., 2007).

Typical urban materials (brick, concrete, asphalt, soil) have very different properties, which is why surface temperature and moisture response vary so strongly across a city. In urban microclimate simulations, HAM models are essential to model the daily cycle of surface heating and cooling, evaporative cooling from wet surfaces, and the thermal inertia that drives the nocturnal urban heat island.

2.2.6. Radiation

Radiation is the dominant energy input to the urban surface during the day and the dominant loss at night. Two electromagnetic wavelength ranges are relevant. Short-wave radiation (0.1–3 μm) or solar radiation comes from the Sun and reaches a surface through three components: the direct component, arriving in a straight line from the solar disc; the diffuse component, scattered by the atmosphere and arriving from the sky hemisphere; and the reflected component, bouncing between urban surfaces. The fraction of incoming short-wave radiation (irradiance) that a surface reflects is its albedo. Long-wave radiation (3–100 μm) is the thermal emission from the earth-atmosphere system, i.e. any surface at non-zero temperature, including buildings, ground, vegetation, and the atmosphere. Its intensity follows the Stefan–Boltzmann law and depends on the surface emissivity (Oke, 2017).

In an urban environment, radiative exchange between surfaces is strongly geometric. A street canyon traps a portion of the radiation reflected and emitted by its walls and ground, which is one of the main physical mechanisms behind the UHI (Oke, 2017). The key geometric quantity is the view factor F_{ij}, defined as the fraction of radiation leaving surface i that reaches surface j. The set of view factors between all surface pairs forms a matrix that, combined with surface temperatures and emissivities, allows the net radiative flux at every surface to be computed (the radiosity method). For short-wave radiation, the direct component additionally depends on the solar position and on which surfaces are sunlit or shaded, typically computed by ray tracing (Kubilay et al., 2017).

Two simplifying assumptions are common in urban radiation models. Surfaces are usually treated as gray and diffuse, meaning that emissivity and albedo are independent of wavelength and direction. The air between surfaces is treated as radiatively transparent, neglecting absorption, scattering, and emission by the atmosphere within the domain (Kubilay et al., 2017).

2.2.7. Setting up a Computational Fluid Dynamics case

3D Model

In practice, to construct a CFD case, a fluid flow domain is first defined, i.e. a volume representing the air in which the calculations are performed. The domain is then discretised into a mesh of small cells, also referred to as finite volume. At each cell, a discretised version of the governing equations is used to solve for the different fields (velocity, pressure, temperature, etc.). The spatial derivatives in the governing partial differential equations describe how each field varies between neighbouring cells, and the fineness of the mesh directly affects both the accuracy of the solution and the computational cost (Franke et al., 2007; Moukalled et al., 2016).

The computational domain must be large enough to avoid the boundaries artificially influencing the flow around the buildings of interest. The recommended dimensions for urban CFD follow the so-called *5H–15H* rule defined by the Franke et al. (2007) best-practice guidelines. The inlet, lateral, and top boundaries are placed at least $5H$ away from the buildings of interest, and the outlet is placed at least $15H$ downstream, where H is the height of the tallest building in the explicitly modelled area. Tominaga et al. (2008) recommends a slightly shorter downstream distance of $10H$. Both guidelines also recommend a blockage ratio (the proportion of the inlet area occupied by the buildings) below 3%, above which the boundaries begin to artificially accelerate the flow around the buildings, see Figure 3.4 in Chapter 3.

An area of the domain has its buildings and vegetation explicitly modelled. For the rest of the area up to the boundaries, rather than meshing these elements explicitly, their aggregate effect is represented by the aerodynamic roughness length z_0. Typical values range from ~ 0.002 m for smooth surfaces such as grass, up to ~ 2 m for dense urban centres, depending on the morphology of the environment; tabulated values are available in the literature (e.g. Wieringa (1992)).

Boundary conditions

Once the domain is defined, the next step is to describe what happens at its boundaries. For each side (inlet, outlet, top, ground, and lateral boundaries), the behaviour of each flow variable (velocity, pressure, turbulence quantities, etc.) must be specified. These specifications, called boundary conditions, tell the solver where the flow enters the domain, where it leaves, and how it interacts with the ground and the buildings (Blocken, 2015).

For microclimate simulations, thermal, moisture, radiation, and modified-pressure boundary conditions must be prescribed in addition to the wind boundary conditions. Air temperature and humidity at the inlet are taken from local meteorological data, typically available from the same sources as wind speed. Surface temperature and moisture initial values are set at the terrain and on the building surfaces. The general settings are described in Section 2.3.3, and those specific to this work in Section 3.3.6.

2.2.8. Validation methods

A microclimate model can be validated in several ways. Roy et al. (2025) identify nine complementary approaches: comparison against field measurements, satellite observations, statistical metrics, cross-model comparison, temporal and spatial validation, sensitivity analysis, calibration with historical data, and validation against independent datasets.

In this study, the simulation results are compared to field measurements. Model performance is then quantified using five dimensionless statistical indicators for all the stations together recommended by Sabatino et al. (2011) and first proposed by Chang and Hanna (2004): the fractional bias (FB), normalised mean square error (NMSE), geometric mean bias (MG), geometric variance (VG) and fraction of predictions within a factor of two of the observations (FAC2). In addition, the relative error is calculated for each station. Their definitions are:

$$\mathrm{FB} = \frac{\overline{C_o} - \overline{C_p}}{0.5(\overline{C_o} + \overline{C_p})}, \tag{2.2}$$

$$\mathrm{NMSE} = \frac{\overline{(C_o - C_p)^2}}{\overline{C_o}\,\overline{C_p}}, \tag{2.3}$$

$$\mathrm{MG} = \exp(\,\overline{\ln C_o} - \overline{\ln C_p}\,), \tag{2.4}$$

$$\mathrm{VG} = \exp(\,\overline{(\ln C_o - \ln C_p)^2}\,), \tag{2.5}$$

$$\mathrm{FAC2} = \text{fraction of data that satisfy } 0.5 \leq \frac{C_p}{C_o} \leq 2.0, \tag{2.6}$$

$$\varepsilon_{\mathrm{rel}} = \frac{|C_p - C_o|}{C_o}. \tag{2.7}$$

where:

C_p: model predictions;

C_o: observations;

$\overline{C}$: average over the dataset;

$\varepsilon_{\mathrm{rel}}$: relative error at each station reported as a percentage.

Positive FB and MG > 1 indicate model under-prediction, whereas negative FB and MG < 1 indicate over-prediction.

The acceptance bands cited by (Sabatino et al., 2011) are $|\mathrm{FB}| < 0.3$, NMSE < 4, $0.7 <$ MG $<$ 1.3, VG < 1.6, FAC2 ≥ 0.5. It is important to note that these bands were developed primarily for atmospheric dispersion model evaluation against wind tunnel data. The difference is that wind tunnel experiments produce observations with very small variability, while field measurements include the natural noisy fluctuation of the wind.

2.3. Tools

This section describes the main tools used in this research.

2.3.1. Model reconstruction

City4CFD is a specialised tool developed by the 3D Geoinformation Research Group at Delft University of Technology (Pađen et al., 2022, 2024) that automates the generation of geometries for CFD simulations in urban environments. Unlike general-purpose 3D models, City4CFD reconstructs terrain and building surfaces with CFD requirements in mind, significantly reducing the manual, time-consuming work traditionally required for geometry preparation. The tool uses point cloud data to extract terrain elevations, calculate building heights and roof form, and 2D building polygons to define building roofprints. Additionally, vegetation and water polygons can be input and the terrain will be subdivided into these surfaces. City4CFD's underlying reconstruction algorithm, Roofer, requires input data in a metric coordinate system, a good match between point-cloud and building footprints, ideally a classified point cloud with at least Building and Ground classes. With these requirements fulfilled, the algorithm produces a surface mesh ready for CFD simulations. The City4CFD tool is indispensable for this study, as it facilitates precise reconstruction of the site geometry.

Surface fixes that could not be performed automatically, or that were specific to the uMF input requirements, were done manually in Rhinoceros 3D (Robert McNeel & Associates, n.d.), a commercial CAD package with extensive mesh-editing capabilities.

2.3.2. Simulation tools

Two OpenFOAM solvers are used in this work: SF (OpenFOAM Wiki, n.d.) for wind-only simulations, and uMF (Kubilay et al., 2017) for coupled urban microclimate simulations. **OpenFOAM** (Open Field Operation and Manipulation) is an open-source CFD framework widely used in engineering and scientific research (OpenFOAM Foundation, n.d.). It provides a library of finite-volume solvers and utilities for fluid flow, heat transfer, and related phenomena. Its modular, object-oriented architecture allows users to develop custom solvers and extend existing capabilities, which makes it well suited to research applications.

SimpleFoam — SF

SF is a steady-state solver for incompressible, turbulent flow (OpenFOAM Wiki, n.d.). It uses the Semi-Implicit Method for Pressure-Linked Equations (SIMPLE) algorithm to solve the RANS equations, iteratively coupling pressure and velocity so that the resulting velocity field satisfies mass conservation. In this work, SF is used exclusively for wind-only CFD simulations, with no thermal or buoyancy effects.

UrbanMicroclimateFoam — uMF

uMF models urban microclimate based on coupled physical processes. It is a multi-region solver, consisting of an air region, which represents the subdomain for turbulent transport, and multiple solid regions, which represent urban surfaces such as buildings and terrain. It is based on the standard OpenFOAM solver chtMultiRegionFoam and extends it with additional coupled physical models. uMF integrates four interacting sub-models: CFD, HAM, Radiation (RAD) and VEG.

The following description is based on the solver documentation in Kubilay et al. (2017).

CFD - air region. A steady RANS model solves continuity, momentum, and energy for the turbulent air flow, plus a transport equation for humidity.

HAM - porous solid regions. HAM is a coupled model for heat and moisture transport in porous building materials, accounting for absorption, transport, and storage. The pore-scale structure is not resolved; instead, the material is modelled as a continuum in which temperature and moisture content vary smoothly throughout the solid. Surfaces can alternatively be modelled as impermeable, or moisture transport can be neglected entirely.

RAD - long- and short-wave radiation. The RAD model accounts for short-wave radiation from the sun and long-wave radiation emitted by all surfaces. Both are computed with a view-factor radiosity method that includes inter-surface reflections; the short-wave direct component additionally uses ray casting based on the sun's position. All surfaces are assumed to be gray, diffuse, and opaque (so windows on façades are not represented), and the air between surfaces is treated as radiatively transparent.

VEG - vegetation. Trees are represented as porous fluid zones with a prescribed Leaf Area Density (LAD), in which source and sink terms for momentum, heat, and moisture are added. The leaf heat balance is formulated in terms of sensible, radiative, and latent fluxes. Transpiration depends on stomatal resistance, which reflects the physiological response of trees to species, incoming radiation, vapour pressure deficit, and available soil moisture. Leaf temperatures are obtained through an iterative energy balance, from which the cooling effect is evaluated.

uMF case steps

Figure 2.2 summarises the steps required to build and run a uMF case. The utilities involved are grouped into four stages: air mesh, solid mesh, radiation preparation, and the solver itself.

Air region mesh.

`blockMesh` Generates the background hexahedral mesh defining the domain extents and base cell size. All subsequent mesh operations refine this grid.

`surfaceFeatures` Extracts feature edges from the STL/OBJ geometry files, producing `.eMesh` files used by `snappyHexMesh` to preserve building and terrain edges.

`snappyHexMesh` Cuts and snaps the background mesh to the building and terrain surfaces in three stages: castellated (removes cells inside/outside surfaces), snap (moves vertices onto surfaces), and layers (optional boundary-layer cells). The output is the final air-region mesh.

`changeDictionary` (air side) Converts building and terrain wall patches to type `mappedWall`, enabling each region to sample field values from the coupled solid region.

Solid region mesh.

`extrudeMesh` Builds each solid mesh by extruding the surface geometry inward by a specified thickness, producing a shell beneath the air–solid interface. Other approaches, such as `blockMesh`, can also be used.

`createPatch` Renames the extruded patches to interface-specific names (e.g. `buildings_to_air`, `terrain_to_air`) so that boundary conditions can target them explicitly.

`changeDictionary` (solid side) Converts these patches to `mappedWall` on the solid side, mirroring the air-side step and completing the two-way coupling.

`setSet` / `topoSet` Defines cell zones (e.g. Soil) that mark which cells belong to the active material region for the HAM transport equations.

Radiation preparation.

`faceAgglomerate` Groups fine mesh faces into coarser clusters, reducing the view-factor matrix from millions of faces to thousands of clusters. Per-patch settings control cluster size and feature-angle sensitivity.

`viewFactorsGen` Computes long-wave view factors between agglomerated clusters. Each factor F_{ij} represents the fraction of radiation leaving cluster i that reaches cluster j. The resulting matrix is reused at every timestep.

`solarRayTracingGen` Performs short-wave ray tracing to determine which surfaces are sunlit and which are in shadow at each solar angle, accounting for shadowing between buildings and on the terrain.

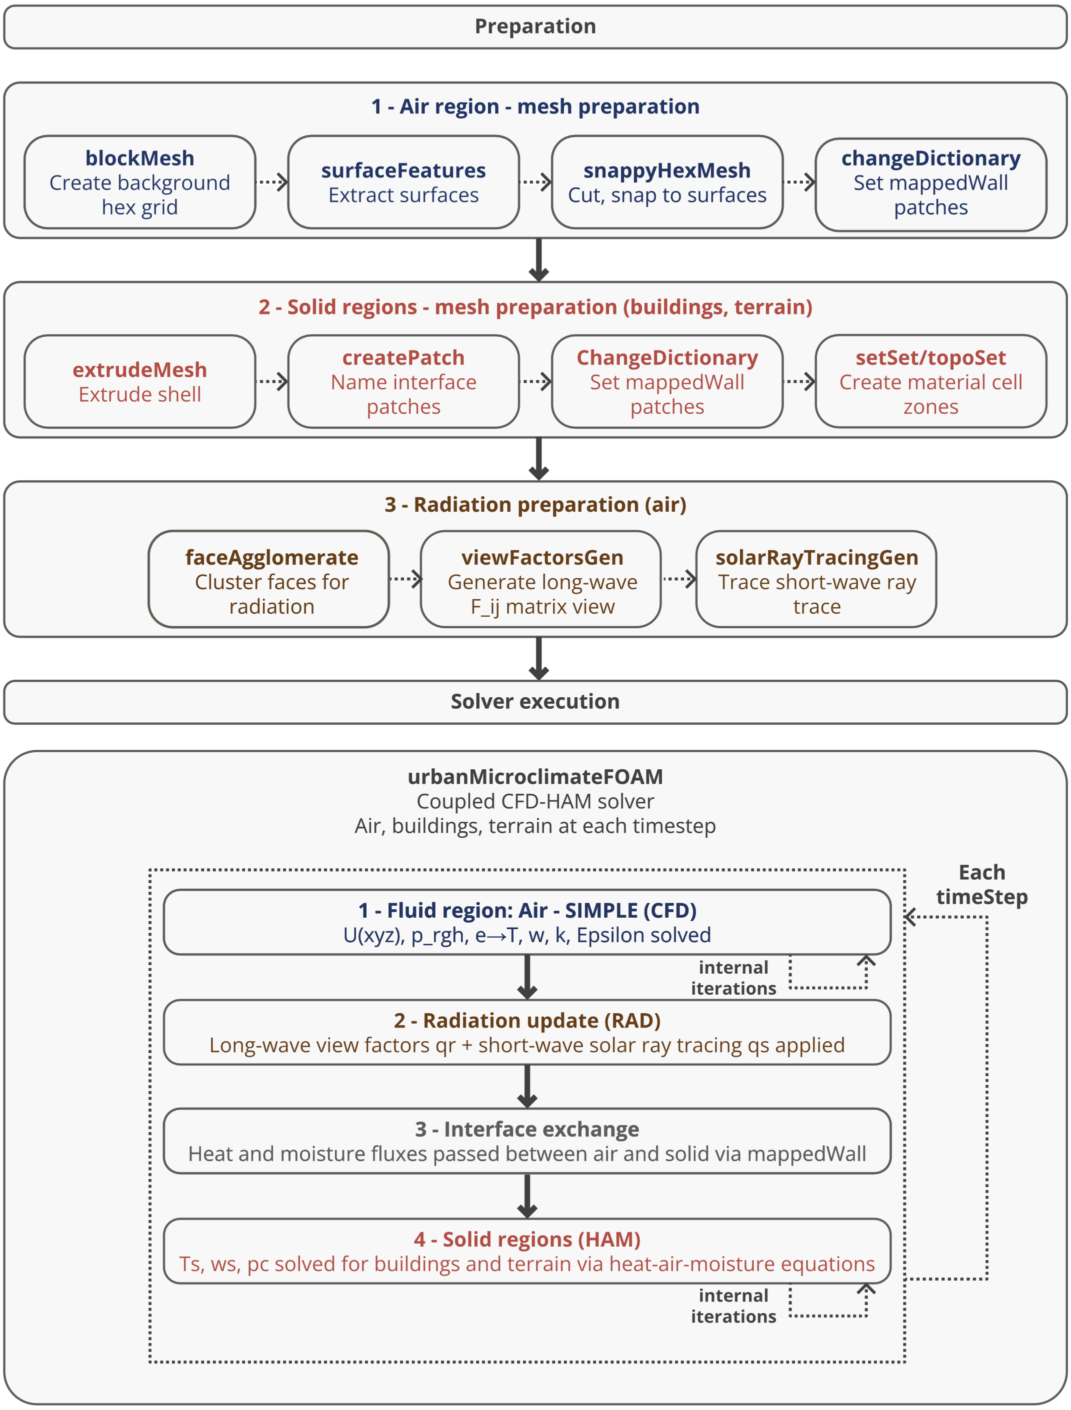

Figure 2.2.: Solver steps for a uMF case of CFD-HAM coupling (note: VEG not included).

Spatial and temporal coupling

Thermal and moisture coupling between the air and solid domains occurs at `mappedWall` boundary patches, which sample field values from a corresponding `samplePatch` in the partner region; these mappings are declared in the boundary files of both meshes. Data is exchanged at discrete exchange timesteps, set by the `deltaT` value in `controlDict`. The coupling is sequential: at each exchange step, the CFD model applies Dirichlet boundary conditions using temperature and humidity values from the HAM model, while the HAM model applies Neumann conditions using heat and moisture fluxes derived from the CFD convective transfer coefficients.

The CFD model is steady-state, whereas the HAM model is transient. This split reflects the difference in transport timescales: air has a low thermal and moisture storage capacity and reaches a new equilibrium within seconds to minutes, while solid building materials such as concrete or brick respond over hours to days. Running both models transiently at the same timestep would therefore be both expensive and unnecessary. Instead, the CFD model provides effectively instantaneous boundary conditions to HAM at each exchange step, under the assumption that the air has already adjusted to the current surface conditions.

Solver execution

At each outer timestep, which is a period of the day, `urbanMicroclimateFoam` (1) solves the air region for velocity, pressure, temperature, moisture, and turbulence using SIMPLE with a defined number of internal iterations; (2) updates long- and short-wave radiation; and (3) solves each solid region for temperature and moisture transport, exchanging heat and moisture fluxes with the air at the coupled interfaces. The simulation marches forward in time until `endTime`.

2.3.3. Comparison of simpleFoam and urbanMicroclimateFoam solvers

SF treats the simulation as a single fluid region, while uMF uses a multi-region formulation with separate air, buildings, terrain, and vegetation regions. Each region requires its own mesh, fields, and numerical settings. Table 2.1 summarises the differences in case setup for the `0/`, `constant/`, and `system/` folders.

The SF case solves only for velocity (`U`), pressure (`p`), and the turbulence quantities `k`, `epsilon`, and `nut`. The uMF air region adds several fields to model heat, moisture, and radiation transport. Thermal and moisture transport require air temperature `T`, humidity ratio `w`, turbulent thermal diffusivity `alphat`, and modified pressure `p_rgh` for buoyancy-driven flow. Radiation and rain are represented by the net short- and long-wave fluxes `qs` and `qr`, the non-dimensional rainwater mass flux `gcr`, and the hourly `cloudCover`, while `Tambient` and `wambient` provide ambient reference values. Each solid region (`buildings`, `terrain`, `vegetation`) additionally carries the three primary fields of the HAM solver: solid temperature `Ts`, moisture content `ws`, and capillary pressure `pc`. Coupled boundary conditions at every air–solid interface exchange temperature, moisture, and radiation fluxes, an exchange that has no equivalent in SF.

In `constant/`, SF requires only `transportProperties` (kinematic viscosity) and `turbulenceProperties` (RANS model selection), together with a `triSurface/` directory holding the

Table 2.1.: Differences in configuration elements required by SF and uMF

Element	**SF**	**uMF**
folder *0/*		
Fields air	U, p, k, ε, ν_t	+ T, w, α_t, p_rgh, qs, qr, gcr, cloudCover
Fields solid	–	Ts, ws, pc (per solid region)
folder *constant/*		
Turbulence model	turbulenceProperties	turbulenceProperties (air)
Thermophysics	–	thermophysicalProperties (air), g
Radiation	–	radiationProperties, solarLoadProperties, viewFactorsDict
Material models	transportProperties	transportProperties per region, g
Solar inputs	–	sunPosVector, IDN, Idif
folder *system/*		
fvSchemes / fvSolution	1 pair	1 pair per region
fvOptions	1	1 (air)
changeDictionaryDict	–	1 per region
createPatchDict	–	1 per region
topoSetDict/setSet.batch	–	1 per solid region

surface meshes. uMF requires considerably more. The air region needs `thermophysicalProperties` for the equation of state and thermophysical coefficients, `radiationProperties` for long-wave model selection (`viewFactorSky`), `solarLoadProperties` for direct and diffuse short-wave radiation, `viewFactorsDict` for the face-agglomeration parameters of view-factor computation, and a gravitational acceleration vector g. Each solid region adds its own `transportProperties` file specifying the HAM material model: density, specific heat, thermal conductivity, and moisture-retention curves (e.g. `Hamstad5Brick`, `AsphaltConcrete`, `Soil`). Three top-level files, `sunPosVector`, `IDN`, and `Idif`, supply the time-varying solar position and the direct and diffuse irradiance values consumed by the radiation model.

Because fluid and solid regions use different discretisation schemes, uMF requires separate `fvSchemes` and `fvSolution` files per region. Several additional dictionaries handle the multi-region setup: `changeDictionaryDict` maps solid surfaces onto the corresponding air-side patches for the coupled exchange step, `createPatchDict` renames patches after solid construction with tools such as `extrudeMesh`, and `topoSetDict` (or `setSet.batch`) defines the cell sets that identify material regions.

2.4. Datasets

This section describes the main datasets used in this study. The MORICHI dataset provides field measurements from the CMU campus during a 2024 heatwave; LiDAR point cloud data are measurements that allow geometric reconstruction of the study area.

MORICHI Dataset

MORICHI (Meteorological Observations and Radiometric Images for Cold Climate Heat Wave Investigations) is an open-access dataset acquired in the Carnegie Mellon University (CMU) campus in Pittsburgh, Pennsylvania, USA (Martin et al., 2026). The data enable studies of urban overheating during extreme heat in a hot-summer humid continental climate (Köppen Dfa), similar to regions in Northeast USA, Eastern Europe, Central Asia, and East Asia. Measurements were acquired from July 26 to September 20, 2024. The data registered the extreme heat event from August 26 to September 1, 2024, when temperatures reached 36°C during the day and 24°C at night.

MORICHI contains measurements from four street-level "Davis Vantage Vue" weather stations (see Figure 2.3). The stations measured air temperature, relative humidity, wind speed/direction, rainfall at 2m height every 5 minutes. Moreover, it includes micro-scale thermal images captured by the infrared thermal camera "FLIR A50" mounted at 30m high, taking images at 1 Hz in 5-minute sequences (640×480 pixels, 1 image/second). Ground truth measurements from contact sensors were used for calibration and achieved $< 1.7\,°C$ RMSE accuracy. The dataset exceeds 1 terabyte and is stored in the 4TU.ResearchData repository under CC BY 4.0 license.

The dataset provides the boundary conditions required to initialize the CFD simulations and enables validation of both SF and uMF predictions against observed field data during the heatwave period.

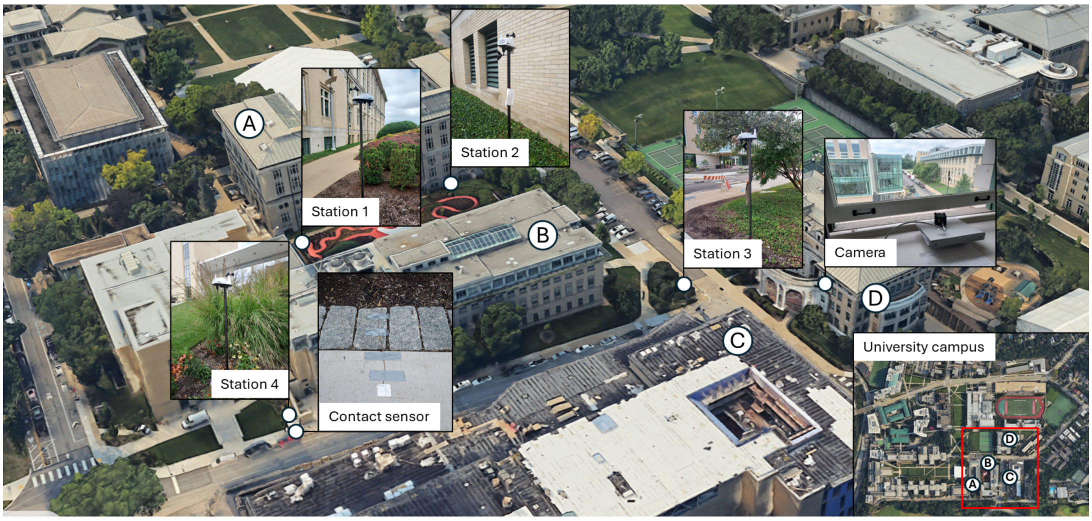

Figure 2.3.: The four MORICHI weather stations. Source: Martin et al. (2026)

Airport Meteorological Data

To define the atmospheric boundary conditions for the numerical simulations, meteorological observations were obtained from the Iowa State University Environmental Mesonet (ISU Mesonet). The data were retrieved, based on the *MADIS HFMETAR / 5-Minute ASOS* dataset.

This dataset aggregates quality-controlled surface observations from the Meteorological Assimilation Data Ingest System (MADIS), incorporating routine and special METAR reports from Automated Surface Observing System (ASOS) stations.

Wind speed, wind direction, and air temperature data were downloaded on January 7, 2026, for the whole year of 2024. Observations from two nearby airport stations were downloaded: Pittsburgh International Airport (PIT) and Allegheny County Airport (AGC). These stations provide representative meteorological conditions for the Pittsburgh metropolitan area.

The dataset provides wind observations at 5-minute temporal resolution at 10 m height, with wind speed reported in miles per hour and wind direction given in degrees measured clockwise from true north, see an example of the data in Figure 2.4. Air temperature is reported in degrees Celsius. Each record includes the station identifier, observation timestamp, station elevation, air temperature (`tmpc`), wind direction (`drct`), and wind speed (`sped`). Routine hourly reports are also included.

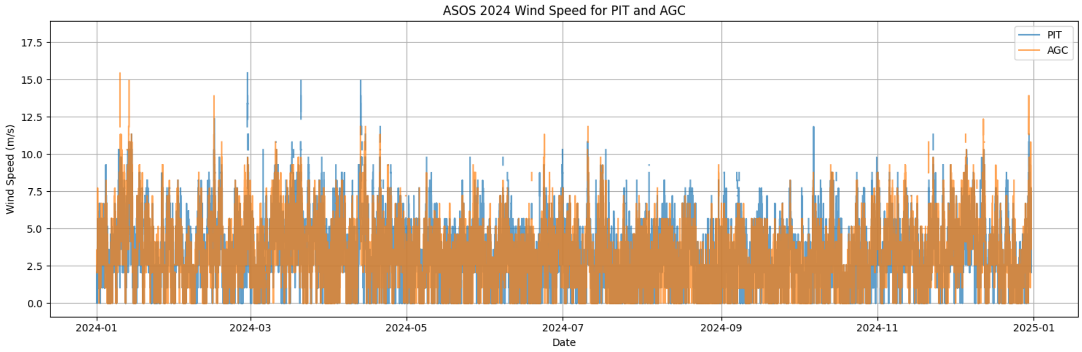

Figure 2.4.: ASOS Wind-speed during 2024 for PIT and AGC Airports

Point-Cloud Dataset

In order to accurately reconstruct the geometry of the study location using City4CFD, it is necessary to employ a point cloud that has been classified with at least Building and Ground classes. The LiDAR survey of Allegheny County, Pennsylvania, acquired in 2015 is openly available. The dataset is georeferenced in the Pennsylvania State Plane coordinate system (south zone, US survey feet). There are two releases of the same data. The 2015 dataset contains approximately 1.3 million classified points distributed across four categories: unclassified points (296,925), ground surface returns (361,004), noise artifacts (1,016), and overlap points from adjacent flight lines (662,958). An additional version from 2017 is available, which builds upon the 2015 dataset by incorporating the building class with 26,003 points. The building classification is, however, not of a high quality, as evidenced by the significant number of unclassified building points, see Figure 2.5.

The ground classification in both datasets provides a good basis for the terrain levelling, but the building-classified points required enhancement to enable proper reconstruction of buildings. Algorithms, such as CANUPO within CloudCompare (CloudCompare, n.d.), were used for this purpose. Furthermore, the dataset was reprojected to a metric coordinate system, using PDAL (PDAL Contributors, 2024), as this is a prerequisite of City4CFD's underlying algorithm, Roofer.

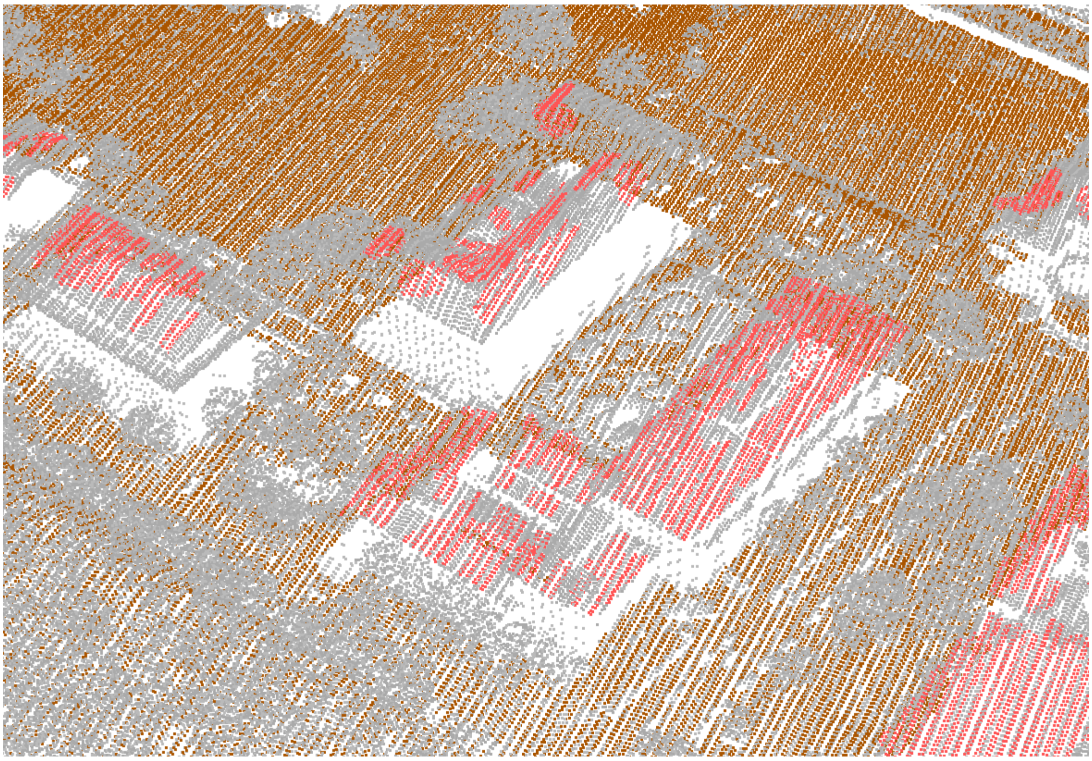

Figure 2.5.: LiDAR survey of Allegheny County (2017 release), coloured by point classification: red for buildings, brown for terrain, and grey for unclassified. A significant number of roof points is unclassified, showing that the building class is incomplete.

Building footprints

Building footprint data for the study area were obtained from the Pennsylvania Spatial Data Access (PASDA) portal, which provides an openly available vector dataset of all buildings and outbuildings in Allegheny County (PASDA, 2026). The footprints were originally digitised from a 2004 aerial flyover and have since been periodically updated; the version used in this study was released on May 18, 2026. Buildings are represented as polygons traced along the roof edge, with near-orthogonal corners squared and minor roof features such as dormers omitted. Minor attached structures (carports, decks, patios, stairs) are excluded, and only buildings larger than 400 square feet are captured, with the exception of garages between 200 and 400 square feet.

For this study, the vector footprints complement the LiDAR point cloud described in Section 2.4 to reconstruct the buildings with City4CFD (Section 2.3.1).

2.5. Related work

This section reviews CFD-based urban microclimate research along the three threads most relevant to this thesis: the state of the field, including its validation gap and representative examples; the modelling choices that determine when buoyancy and geometric fidelity must be resolved; and existing applications of uMF.

2.5.1. State of the field and the validation gap

Urban microclimate CFD has grown rapidly over the past two decades. Toparlar et al. (2017), reviewing 183 studies published between 1998 and 2015, report that 96% of these studies relied on RANS simulations and that 58% lacked any comparison with measurements. They recommend that validation in future work move from simple percentage differences toward dimensionless statistical indicators of the kind proposed by Chang and Hanna (2004) and adopted with acceptance bands by Sabatino et al. (2011).

A decade later, Roy et al. (2025) reports that the field has expanded roughly tenfold but that the validation gap, while narrower, has not closed: field-data validation remains inadequate, and most tools are still too complex for routine use by planners. Practical guidance for credible CFD setups is consolidated in the best-practice guidelines of Tominaga et al. (2008) and Blocken (2015), whose recommendations on domain sizing, grid quality, roughness specification, and inlet conditions form the methodological basis adopted in this work. Schneider et al. (2023) complement this by showing that even with well-validated tools, inappropriate boundary conditions or evaluation parameters can produce misleading conclusions, emphasising that simulation objectives must be clearly defined before setup.

2.5.2. Examples of urban microclimate simulations

Two published studies illustrate how these principles translate to real urban sites. Toparlar et al. (2015) simulated the Bergpolder Zuid district in Rotterdam using 3D unsteady Unsteady Reynolds-Averaged Navier–Stokes (URANS) simulations with the realisable k–ε model on a 6.6-million-cell grid to reproduce a five-day heatwave in July 2006. The study addressed six UHI causes, and was validated against spatially-averaged surface temperatures from NOAA-AVHRR thermal satellite imagery, reporting an average deviation of 7.9%. The results illustrated two key UHI mechanisms: heat storage in construction materials, with surface temperatures of 21.8 °C persisting at 23:00 when air temperatures had dropped to roughly 17.0 °C, and reduced turbulent heat transport in streets oriented perpendicular to the prevailing wind. The acknowledged limitations are the omission of trees, uniform material properties across surfaces, and the use of neutral-stratification inlet profiles despite low wind speeds during the heatwave.

Brozovsky et al. (2021) present a complementary case at the NTNU campus in Trondheim, Norway, addressing the relative scarcity of microclimate studies in high-latitude cold-climate cities. They applied 3D URANS with the realisable k–ε model in ANSYS Fluent on a 9.1-million-cell polyhedral grid, including geometrically explicit buildings and trees, short- and long-wave radiation, heat storage, and evapotranspirational cooling from trees and grass via the Penman–Monteith equation. Validation against six weather stations during two 48-hour periods in autumn 2019 yielded mean R^2 of 0.60 for wind speed, 0.63 for wind direction, and 0.96 for air temperature at the reference station. Wind speed at the mobile stations was more difficult to assess, since recorded values were below the sensors' starting threshold of 1 m/s for more than 95% of the time in both validation periods. This issue is closely paralleled by the present study, where three of the four MORICHI stations recorded hourly means within the manufacturer's stated anemometer accuracy. Brozovsky et al. (2021) also report that the Penman–Monteith formulation slightly overestimated evapotranspirational cooling on grass-covered surfaces, producing a systematic temperature underestimation at the lower stations. Together with the Rotterdam case, this study illustrates that even rigorously validated CFD models retain discrepancies linked to vegetation modelling, inlet-profile assumptions, and the limits of low-wind-speed instrumentation.

2.5.3. Buoyancy and the wind-only versus coupled question

Other studies have examined when buoyancy effects can no longer be neglected in urban flows. Reviewing wind-tunnel and CFD studies of heated urban geometries, Mei and Yuan (2022) identify regime thresholds based on the bulk Richardson number: reattachment lengths behind isolated buildings begin to shorten under buoyancy once $|Ri_b| > 0.2$, and the canyon vortex reorganises into a two-vortex structure when $|Ri_b| > 2$. Zhao et al. (2021) provide complementary wind-tunnel evidence that thermal stratification fundamentally alters pedestrian-level wind and canyon ventilation when buoyancy is comparable to wind forcing. Together, these studies establish that the wind-only approximation loses validity once $|Ri_b|$ exceeds roughly 0.2.

2.5.4. The role of geometric fidelity

A parallel line of research has asked how strongly the geometric representation of buildings and terrain affects CFD wind predictions. Hågbo et al. (2021) compared four building-geometry sources and report significant differences in pedestrian-level wind fields when a simple footprint-extrusion model (equivalent to LoD 1.2) is used instead of LiDAR-based geometry. Patil and García-Sánchez (2025) systematically compared LoD 1.2 and LoD 2.2 across 1440 RANS simulations on two sites, reporting wind-speed differences of 10-20% and a systematic under-prediction of joint velocity–turbulence risk indicators at LoD 1.2. García-Sánchez et al. (2021) report similar effects on the TU Delft campus and additionally show that semantic surfaces (water, vegetation) affect pedestrian-level wind speeds. On the terrain side, Talwar and Yuan (2024) compared Large Eddy Simulation (LES) with natural and flat terrain under both isothermal and non-isothermal conditions, finding a statistically significant impact under buoyancy-driven flow, with local deviations reaching 40-45% and 55% in low-density and windward open-space areas respectively. Together, these studies establish that geometric fidelity, in both buildings and terrain, significantly affects CFD wind predictions.

2.5.5. Applications of urbanMicroclimateFoam

uMF (Kubilay et al., 2017) couples airflow, heat and moisture transport in porous solids, short- and long-wave radiation, and vegetation within the OpenFOAM framework. The most recent published application, Rahimi et al. (2025), validates the solver against a network of weather stations on the Georgia Tech campus and reports Root Mean Square Error (RMSE) as low as 1.03 °C for air temperature, with vegetation reducing local temperatures by up to 2.1 °C. A 24-hour simulation completed in 2.5 hours on a high-performance system, suggesting that coupled CFD is computationally viable for design-oriented studies.

The study acknowledges limitations in vegetation modelling and radiation assumptions, and is restricted to a single 24-hour cycle. A further limitation, not discussed by the authors, is that the case is set on flat terrain. To the authors' knowledge, no published application of uMF has evaluated its behaviour on a realistic urban site with substantial topographic variation.

2.5.6. Research gaps and motivation

Synthesizing these studies reveals three critical gaps in current urban microclimate modelling research that motivates this research. First, although geometric fidelity has been shown to significantly affect wind-only CFD predictions, the feasibility of applying thermally-coupled uMF solver to realistic urban environments with complex terrain is under-explored. Second, the relative benefits and limitations of wind-only versus fully coupled simulations seem insufficiently quantified under heat-stress-relevant conditions. Third, validation practices often lack statistically robust comparisons against field measurements. Addressing these gaps is essential to determine when advanced coupled solvers such as uMF are both necessary and practically viable.

3. Article

3.1. Abstract

Climate change is intensifying urban heat stress, and microclimate simulation can support climate-responsive planning. CFD is the dominant approach at the microscale, but urban CFD studies are predominantly isothermal, neglecting buoyancy effects. Yet these effects become non-negligible under the low-wind, thermally-active conditions that climate change is making more frequent. Coupled solvers such as uMF, built on OpenFOAM, resolve the physics interactions (HAM, RAD, VEG) that wind-only formulations omit. Despite this promise, published applications of uMF to realistic urban sites including complex terrain are scarce, the additional preprocessing and computational cost relative to a wind-only baseline are not well characterised, and rigorous validation against high-resolution field measurements is limited. This research addresses these gaps to help determine when the added complexity of a coupled solver is justified in practice.

This study evaluates uMF on the Carnegie Mellon University campus (Pittsburgh, USA) during one hour of a heatwave in August 2024, comparing it against a wind-only SF baseline on identical geometry. A 3D model integrating LiDAR point clouds and official building footprints was reconstructed using City4CFD, and both simulations were validated against four street-level weather stations from the MORICHI dataset under low-wind, thermally-active conditions where buoyancy is non-negligible relative to wind shear.

Three findings emerge. First, applying uMF to realistic terrain is feasible but requires time-consuming, case-specific geometric adaptations compared to geometry-robust SF. Second, on identical geometry and 20 cores, uMF took 3.6× longer than SF (6 h 28 min versus 1 h 49 min), which does not represent a barrier to the usage of the solver. Third, uMF improved wind-speed agreement across every aggregate indicator: FB moved from $+0.65$ to -0.11, MG from 1.89 to 0.83, NMSE was reduced by 65%, and FAC2 rose from 0.50 to 0.75. These results indicate that uMF merits adoption when buoyancy effects are non-negligible and the analyst can absorb the case-setup overhead. The findings rest on a single one-hour window under one inflow direction, with the observed wind speeds at three of the four stations falling below the anemometer accuracy threshold, and vegetation modelled as a porous medium. Addressing these limitations is identified as a priority for future work.

3.2. Introduction

Climate change is increasing the frequency, duration and intensity of heatwaves worldwide. Urban areas are especially susceptible to this phenomenon, as temperature extremes are intensified in cities (IPCC, 2023). Global mean surface temperature in 2024 reached 1.60 °C above the pre-industrial (1850–1900) baseline, making it the first calendar year to exceed the 1.5 °C threshold of the Paris Agreement (Copernicus Climate Change Service (C3S), 2025).

Compounding this, dense built environment tends to amplify temperatures through the UHI effect (Roy et al., 2025; Oke, 2017).

At the same time, an increasing proportion of the global population is living in cities, thereby more people are exposed to heat-related health risks (World Health Organization, 2024) and reduced thermal comfort: more than two billion people currently living in cities could be exposed to an additional temperature increase of at least 0.5 °C by 2040 relative to present levels, and up to 36% of the global urban population may experience mean annual temperatures of 29 °C or higher. This growing heat exposure is further exacerbated by the decline of urban green spaces, as observed globally from an average of 19.5% in 1990 to 13.9% in 2020, reducing cities' capacity to mitigate heat stress through shading and evapotranspiration (UN-Habitat, 2024).

Given this escalating urban heat challenge, UN-Habitat (2024) argues that cities must be placed at the centre of global climate action. Urban and land-use planning and design are defined by critical instruments for climate action, as climate-responsive planning can support progress towards net-zero targets while strengthening resilience to extreme heat and other climate-related shocks. However, effective planning requires robust data: UN-Habitat (2024) identifies that closing urban climate data gaps is crucial to enable evidence-based decision-making, attract investment, and ensure that adaptation strategies address the needs of the most vulnerable populations.

To bridge these data gaps and support evidence-based planning, microclimate simulation tools offer one valuable approach for assessing urban thermal environments and evaluating adaptation strategies (Toparlar et al., 2017; Blocken, 2015; Roy et al., 2025). By modelling the interactions between wind flow, solar radiation, surface materials, vegetation, and urban morphology, such simulations help researchers and planners to assess the thermal behaviour of existing environments and evaluate potential adaptation strategies throughout the planning process, helping cities improve outdoor thermal comfort and resilience under financial and spatial constraints. They inform decisions relating to material choices, shading design, vegetation placement, and broader urban form interventions or planning aimed at mitigating heat stress (Toparlar et al., 2015; Schneider et al., 2023; Kubilay et al., 2020; Yang et al., 2023).

Among the available simulation approaches, CFD models are predominant in urban microclimate research (Toparlar et al., 2017; Roy et al., 2025). Isothermal scenarios that neglect thermal buoyancy effects dominate the literature, given their focus on wind comfort and urban ventilation. Even thermal simulations often fail to include critical factors such as long-wave radiation between building surfaces, detailed thermal storage properties of buildings and soil, and comprehensive vegetation effects (Yang et al., 2023), while there is evidence that these factors can significantly alter microclimatic conditions (Rahimi et al., 2025).

Several tools have sought to address this gap. Among the most widely used are ANSYS Fluent (ANSYS, Inc., n.d.a), ENVI-met (ENVI-met GmbH, n.d.), and OpenFOAM (OpenFOAM Foundation, n.d.). Commercial software such as ANSYS Fluent dominated early urban microclimate studies. Among dedicated microclimate software, ENVI-met has long been the reference for the urban planning community thanks to its user-friendly interface and built-in modules for vegetation, radiation, and thermal comfort indices. OpenFOAM, as an open-source alternative, has supported a range of research-oriented workflows. Within this landscape, uMF (Kubilay et al., 2017) stands out as an open-source OpenFOAM solver that integrates the four physical subsystems most relevant to urban heat-island studies under one common framework: airflow (CFD), heat and moisture transport in porous materials

(HAM), short-and long-wave radiation (RAD), and vegetation (VEG). uMF' most recent published application (Rahimi et al., 2025) validated the solver against on-site weather stations on the Georgia Tech campus, reporting RMSE as low as 1.03 °C for temperature and demonstrating that vegetation modelling and humidity coupling affect predicted microclimate dynamics.

The relative importance of mechanical (shear-driven) and buoyant (thermally-driven) forcing in the urban boundary layer can be quantified by the bulk Richardson number (Mei and Yuan, 2022; Zhao et al., 2021). For the conditions of the present study ($u_a \approx 1.6$ m/s, $H \approx 22$ m, $\theta_a \approx 299$ K, $\theta_g \approx 302$ K), the bulk Richardson number is $Ri_b \approx -0.85$. Reviewing buoyancy effects on urban airflow, Mei and Yuan (2022) report two relevant regime thresholds from wind-tunnel and CFD studies of heated urban geometries: reattachment lengths behind isolated buildings begin to shorten under buoyancy once $|Ri| > 0.2$, while the single canyon vortex reorganises into a two-vortex structure under heated windward walls when $|Ri| > 2$. The present value of $Ri_b \approx -0.85$ lies well beyond the first threshold and below the second, placing the flow in a regime where buoyancy is no longer negligible relative to wind shear but is not yet strong enough to fundamentally restructure the canyon-scale flow. The wind-only assumption is therefore expected to lose accuracy in this regime. The comparison between the wind-only and non-isothermal simulations presented in this work aims at quantifying the difference between the two approaches in this regime, using field measurements as the reference.

Despite the relevance of uMF, three important gaps remain in its application to realistic urban case studies.

First, realistic models involving non-flat topography are under-explored in published work using uMF, even though high-detail, realistic reproduction of the 3D environment is known to ensure accurate wind simulations.

Patil and García-Sánchez (2025) and García-Sánchez et al. (2021) examine the impact of building LoD, comparing LoD 1 and LoD 2 representations (and their subcategories). Both studies reports that these differences affect pedestrian level wind speeds. Patil and García-Sánchez (2025) finds a systematic over-prediction of LoD 1.2 when compared to LoD 2.2, among other findings, and recommends using higher geometric detail where possible. Hågbo et al. (2021) also reports differences between highly detailed models and extruded models (which can be defined as LoD 1), especially in the downstream region of the buildings.

Regarding terrain, Talwar and Yuan (2024) examined the influence of natural terrain on pedestrian level wind predictions under both isothermal and non-isothermal conditions and reported statistically significant impact in buoyancy-driven flows. In wind-driven flows, over high- and medium-density neighbourhoods, the flow is more affected by buildings than by terrain. The effect of terrain is most significant around open spaces and in low- and medium-density areas with shallow and steep slopes. Deviations in average wind speed reach 45% under buoyancy-driven conditions and 40% under wind-driven conditions. The authors further observe that flat-terrain cannot capture the disturbances introduced by immediate upstream topography: at a windward open-space site, the flat-terrain simulation over-predicted pedestrian wind speed by 55% relative to the natural-terrain case. Both conditions are present in the case study examined here: substantial terrain variation upstream of the buildings of interest and low to middle built-up density.

Evolving solutions from the geomatics field make it possible to reconstruct realistic, simulation-ready geometries by integrating heterogeneous spatial data such as point clouds and 2D topographic data, and modern pipelines such as City4CFD (Pađen et al., 2022, 2024) now largely automate the reconstruction of detailed 3D city models for CFD purposes. However, it remains uncertain to what extent uMF can ingest such highly realistic inputs while maintaining computational feasibility and numerical stability.

Second, urban microclimate simulations are more computationally expensive than wind-only simulations, and it is unclear under which conditions wind-only simulations would be sufficient, especially for urban heat stress assessment, where the additional cost of fully-coupled uMF simulations must be justified.

Third, rigorous validation against field measurements remains inadequate: Toparlar et al. (2017) found that 58% of urban microclimate CFD studies between 1998 and 2015 lacked any comparison with measurements, and Roy et al. (2025) notes that this gap has narrowed but not closed. Where validation is performed, it often relies on satellite imagery (Toparlar et al., 2015; Moediartianto et al., 2026) or a small number of ground-based sensors providing coarse temporal resolution. The MORICHI dataset (Martin et al., 2026), used in this study, combines four street-level weather stations recording at five-minute intervals with infrared thermal imagery acquired during an extreme heat event. While limited in spatial coverage, this combination of high temporal resolution with co-located thermal imaging supports the evaluation against dimensionless indicators (Sabatino et al., 2011; Chang and Hanna, 2004) of the kind Toparlar et al. (2017) recommended as the standard for the next generation of urban CFD studies.

Closing these gaps is necessary to determine the circumstances in which wind-only or non-isothermal simulations can be reliably and practically applied in real-world urban climate analyses. Accordingly, the aim of this study is to evaluate the feasibility and performance of uMF when applied to a realistic urban environment with complex terrain during an extreme heat event, compare it with a wind-only simulation, and validate both against high-resolution field measurements.

Three questions guide the investigation.

1. How feasible is the application of uMF to realistic urban models that include non-flat terrain, given current geometric reconstruction and meshing tools?
2. How do mesh requirements, solver parameters, and computational demands differ between uMF and the wind-only solver SF when applied to the same case study?
3. How do wind-speed predictions from SF compare to those of uMF when both are evaluated against field measurements from four stations over a one-hour window under a single inflow condition during a representative heatwave?

Together, these questions address the three gaps identified above: practicability on realistic terrain, the conditions under which the additional cost of coupled simulations is justified, and rigorous validation against high-resolution measurements.

This article is structured as follows. Section 3.3 describes the methodology, data, and tools used in this study. Section 3.4 reports the results, covering the feasibility of applying uMF to realistic urban geometries, the computational cost relative to SF and the comparison of wind-speed predictions against the MORICHI field measurements. Section 3.5 shows the limitations of the analysis. Section 3.6 discusses the findings in light of the three research questions, reflects on the conditions under which fully coupled simulations are warranted. Finally, Section 3.7 concludes with recommendations for future development of uMF-based urban microclimate workflows.

3.3. Methodology

This study adopted a comparative, case-study methodology comprising three main stages: geometric reconstruction and mesh generation, setup and execution of SF and uMF simulations, and evaluation and comparison of the results.

3.3.1. Case study

The Carnegie Mellon University (CMU) campus in Pittsburgh, Pennsylvania, was selected as a case study (Figures 3.1-3.2). This location was selected for three reasons. First, high-temporal-resolution field measurements are available through the MORICHI dataset (Martin et al., 2026), see sensor locations in Figure 3.3, enabling validation against meteorological data. Second, detailed geospatial datasets covering the campus and its surroundings are publicly available, allowing accurate 3D reconstruction of the simulation domain. Third, the site has substantial terrain variation (83m within the domain area), making it a suitable test case for evaluating the performance of uMF under these topographic conditions.

The CMU campus is located in the Squirrel Hill North district of Pittsburgh, approximately 4 km east of the downtown. Pittsburgh has a hot-summer humid continental climate (Köppen Dfa), characteristic of the Northeast United States and comparable to regions of Eastern Europe, Central Asia, and East Asia (Martin et al., 2026). The site lies in hilly topography that produces strong terrain variation both within the campus itself and across its surroundings.

The buildings within the campus are predominantly two to five storeys in height. To the southeast, the campus borders Schenley Park, a roughly 1 km open green space, with patches of dense vegetation. Beyond the park, between 1 km and 5 km, the terrain is occupied by a hilly, well-vegetated single-family residential neighbourhood of predominantly two-storey houses, before reaching the Monongahela River. Further south, from approximately 5 km to 10 km, the landscape continues as forested hilly terrain. These characteristics are reflected in the aerodynamic roughness length values assigned at the simulation inflow boundary (Wieringa, 1992).

3.3.2. Geometry reconstruction

The terrain and buildings 3D geometry was reconstructed using City4CFD (Pađen et al., 2022, 2024) and adjustments to address specific requirements of uMF were performed with Rhino (Robert McNeel & Associates, n.d.). The geometric alterations required for uMF stability are documented in Section 3.4.1. Vegetation was reconstructed with a C++ pipeline developed by TU Delft using CGAL alpha wrapping tool (Alliez et al., 2026) and adjustments were also performed in Rhino.

City4CFD allows combining LiDAR point cloud data (PASDA, 2015) and official 2D building footprint dataset (PASDA, 2026) for the terrain and building reconstruction. The original point cloud in the EPSG:2272+5702 coordinate reference system (imperial units) was first converted to the metric EPSG:32129+5703 system using the PDAL library (PDAL Contributors, 2024), compatible with City4CFD. The LiDAR point cloud underwent supervised classification using CANUPO, a plugin of CloudCompare (CloudCompare, n.d.), which separated most building points from the unclassified point data and improved classification

Figure 3.1.: Site plan showing the location of the CMU campus within the city and the extent of the computational domain used in this study.

(a) Campus of the Carnegie Mellon University. Image source: Bing Maps 3D.

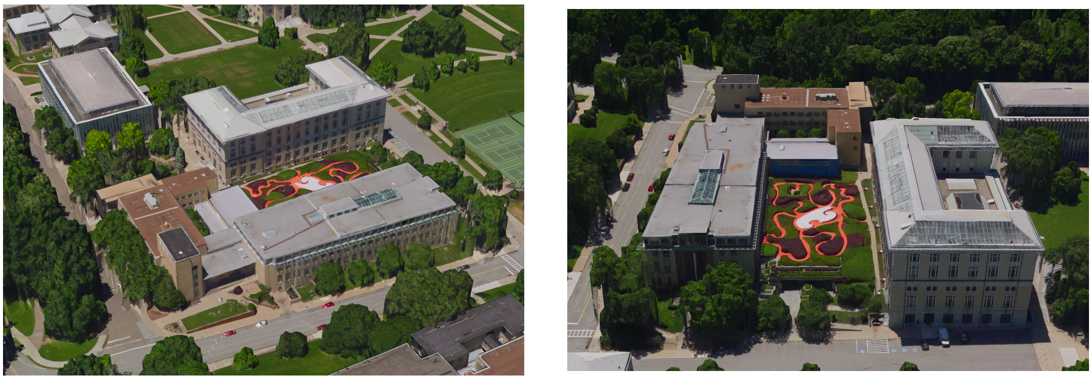

(b) Building block where MORICHI measurements were acquired. Image source: Bing Maps 3D.

(c) Building block where MORICHI measurements were acquired. Image source: Bing Maps 3D.

Figure 3.2.: Case study location: the Carnegie Mellon University campus, Pittsburgh, Pennsylvania. (a) Aerial view of the campus. (b–c) Building block at which MORICHI street-level measurements were acquired. Image source: Bing Maps 3D.

quality for building reconstruction. Tree points were isolated using the HOMED vegetation filter algorithm (Alting, 2025). Buildings within a 300 m radius of the weather station locations were explicitly reconstructed in LoD 1.2.

3.3.3. Domain and mesh

Domain sizing

The simulation domain is centred on the CMU campus at coordinates [414 009.007, 125 342.985] in EPSG: 32129. It is sized following Franke et al. (2007); Blocken (2015) guidelines for atmospheric boundary layer CFD simulations, with a lateral extent of approximately 5H on each side of the area of interest, an upstream fetch of 5H, a downstream fetch of 15H, and a vertical extent higher than 6H, where H denotes the height of the tallest building within the area of interest. The resulting domain is a rectangular block extending 1,670 m in the streamwise direction, 1,230 m in the lateral direction, and 555 m in the vertical direction, discretised into 16,214,790 cells, see Figure 3.4. Further justification of the domain shape and size is provided in Section 3.4. Local refinement and three levels of boundary layers around buildings and terrain are applied with the `snappyHexMesh` utility, see Figure 3.5. Boundary layers are approximately 20 cm thick, and the smallest cells in the most refined regions range from 60 to 90 cm. This ensures that the weather stations are located at least in the third cell from the ground, avoiding wall effects and high gradients. Moreover, for uMF HAM, the solid surfaces, buildings and terrains were constructed with the `extrudeMesh` utility at thickness of 0.1 m with 10 layers. Terrain output 482,680 cells and buildings, 261,770, see Figure 3.5
.

This geometry served as a common basis for both wind-only and thermally-coupled simulations, ensuring that differences in results can be fully attributed to model physics.

Grid Convergence Index (GCI)

To verify that the simulation results are sufficiently independent of the mesh resolution, a grid convergence test was carried out following the methodology of Celik et al. (2008). The test was performed with the SF solver only, since the two solvers share the same air-region mesh. Three meshes of increasing refinement were generated for the same domain and the same surface geometry: a coarse, a nominal (the one used for the final simulations), and a fine mesh. The representative grid size was computed as

$$\Delta h_r = \left(\frac{V_{\text{domain}}}{N} \right)^{1/3}, \tag{3.1}$$

where V_{domain} is the total volume of the simulation domain and N is the total number of cells. See Table 3.1

Table 3.1.: Mesh characteristics for the grid convergence study. The representative grid size is computed using Eq. 3.1 with $V_{\text{domain}} = 1.140 \times 10^9$ m^3.

Mesh	Cell count N	Δh_r (m)	Refinement ratio
Fine	35 524 066	3.18	–
Nominal	16 214 605	4.13	1.30
Coarse	7 859 572	5.26	1.30

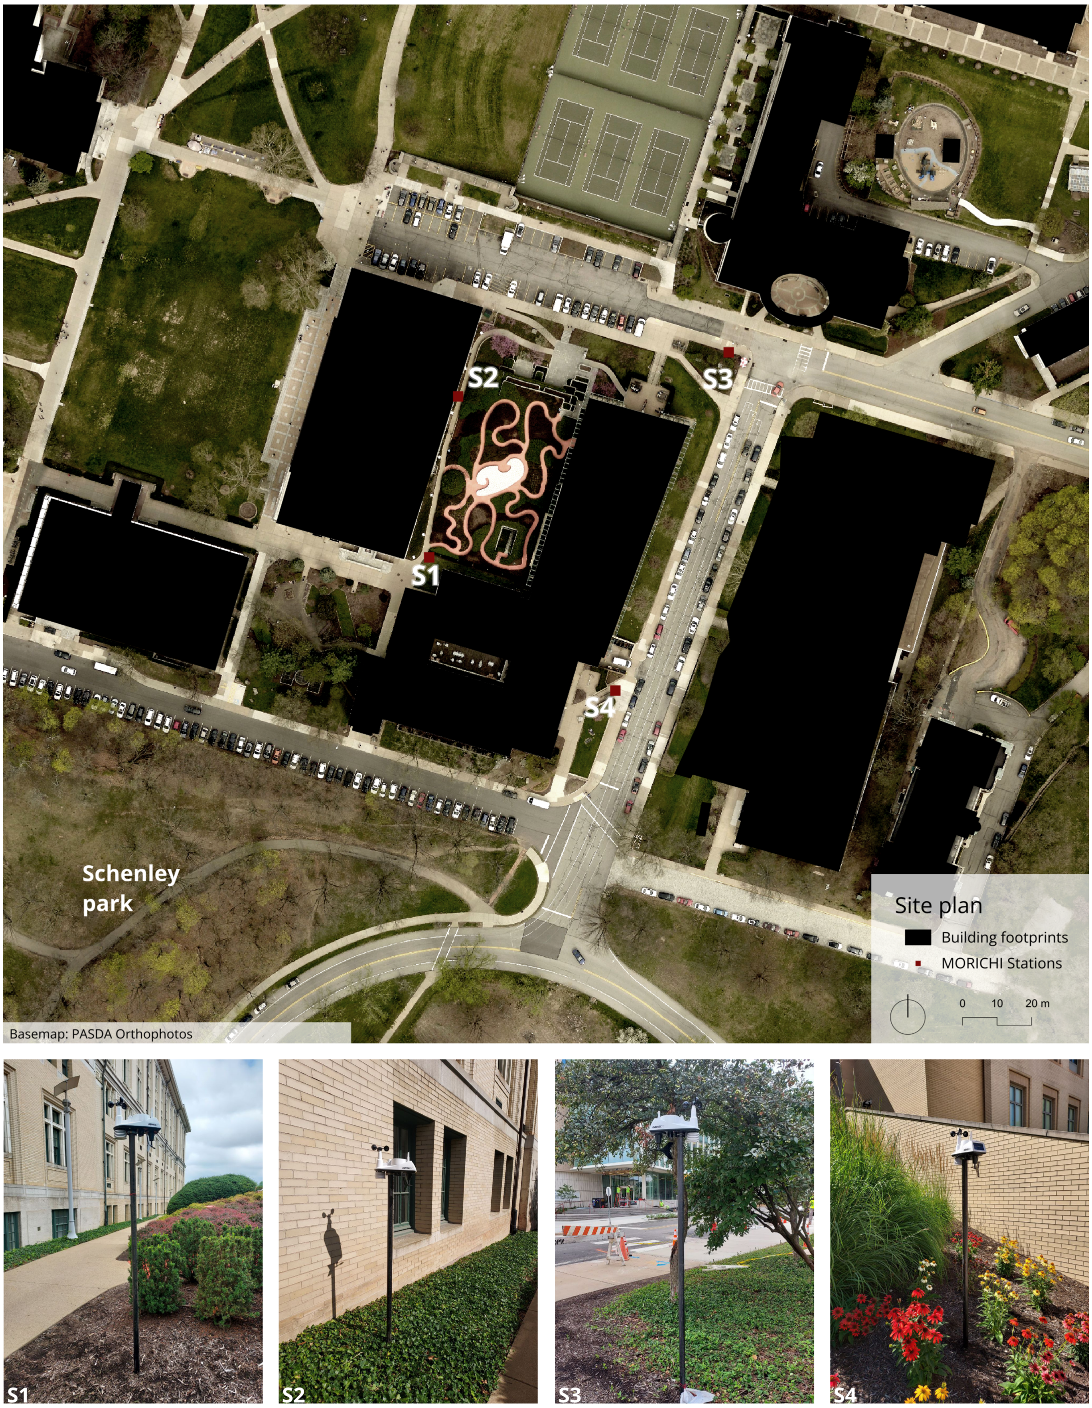

Figure 3.3.: MORICHI stations. Source: station photos provided by Miguel Martin Fehlmann; orthomosaic imagery by PASDA (2017).

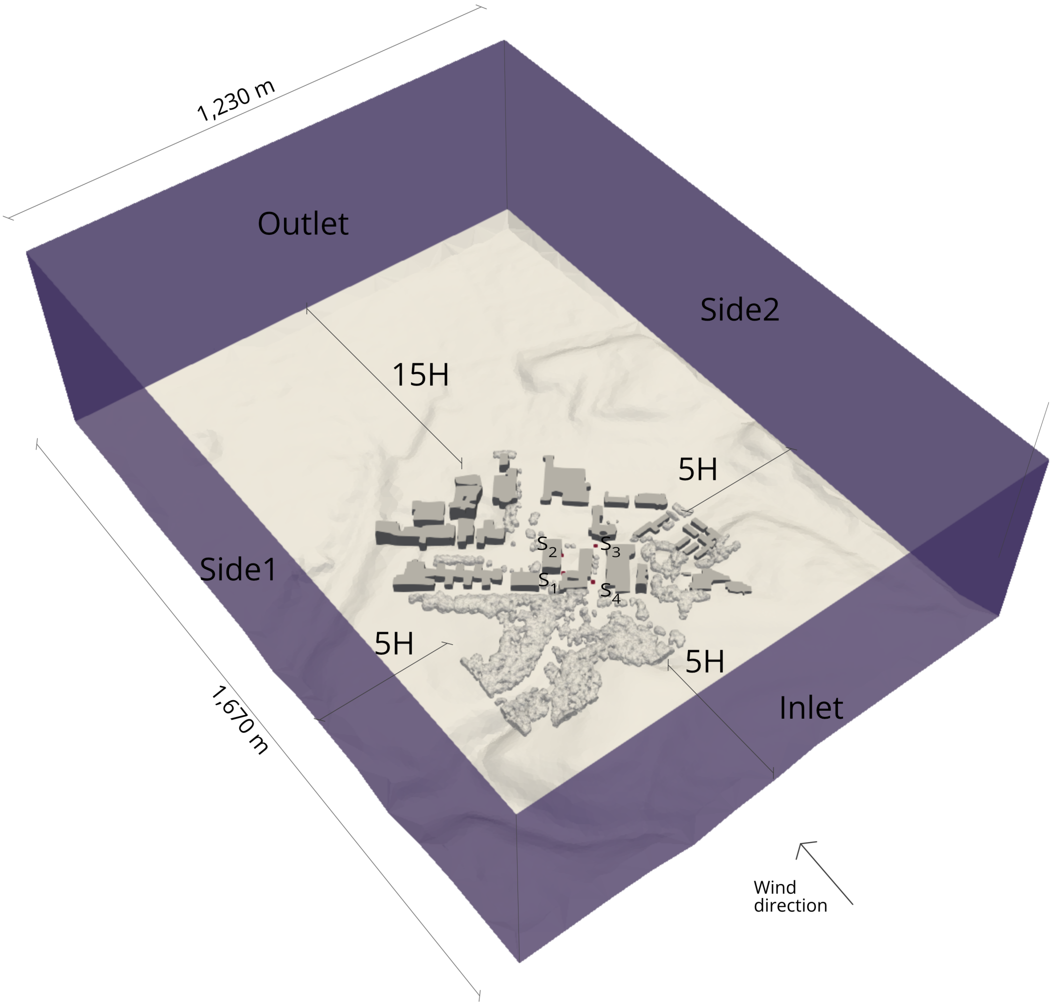

Figure 3.4.: 3D Geometry reconstruction including vegetation and buildings within 300 meters of MORICHI stations (red dots) and terrain.

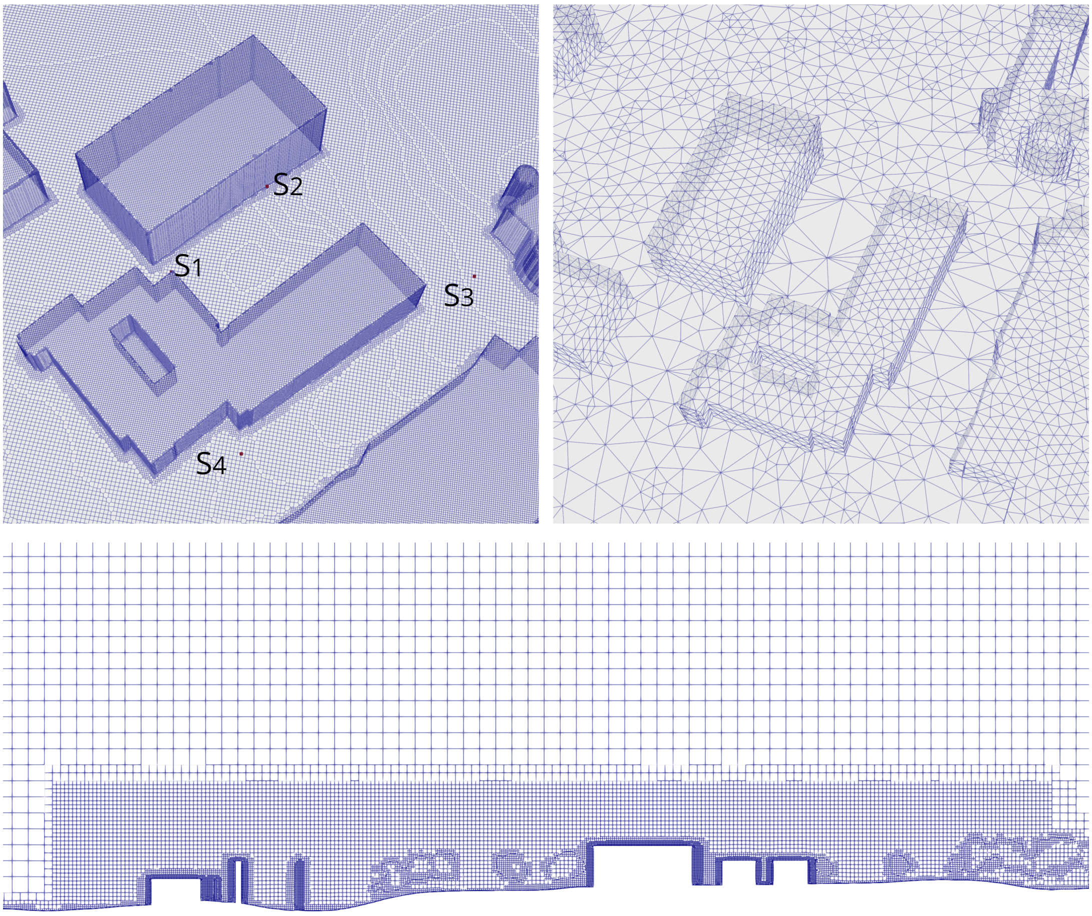

Figure 3.5.: Air mesh at surfaces (top left) and station locations, Building and terrain mesh (top right) and Air mesh resulting from `snappyHexMesh` (bottom)

The GCI was computed at 60 probes distributed throughout the domain. For each variable, the apparent order of convergence, the relative error between fine and nominal solutions, the extrapolated error, and the GCI were computed at each probe (Celik et al., 2008). To avoid contamination of the grid-sensitivity estimate by local effects near walls and near-zero velocity regions, the median across probes is reported as the primary indicator (Patil and García-Sánchez, 2025).

Table 3.2.: Grid convergence indicators for the four primary flow variables, computed over 60 probes.

Convergence parameter	U	p	k	ε
Apparent order of convergence (median)	3.6	3.9	4.2	4.4
Relative error (%, median)	0.9	1.7	1.7	3.5
Extrapolation error (%, median)	0.7	0.9	0.9	1.7
GCI (%, median)	0.9	1.2	1.2	2.1

The median GCI values are all below 3% for the four primary flow variables, in line with the indices reported by Patil and García-Sánchez (2025). The nominal mesh is therefore considered sufficiently grid-independent for the analyses presented in this work.

3.3.4. Simulation tools and solvers

Given that one of the central aims of this study is to assess the performance of wind-only simulations relative to thermally coupled simulations, OpenFOAM offers a well-suited environment: uMF provides the coupled HAM and RAD solvers, while SF can be applied to the identical geometry, mesh, and boundary configuration to produce a directly comparable wind-only baseline.

3.3.5. Effects of vegetation

For simplification of the case, the VEG module of uMF was not used. Instead, the trees were modelled as porosity zones. A porous resistance region was defined using the OpenFOAM `explicitPorositySource` model with a Darcy–Forchheimer formulation. The Forchheimer coefficient relates to canopy properties through

$$f = C_d \cdot \mathrm{LAD}$$

where C_d is a sectional drag coefficient and LAD is the leaf area density. Following values reported in the literature ($C_d \approx 0.2$ for urban trees (Katul and Albertson, 1998; Massman, 1987)) and a representative LAD of 3 $\mathrm{m}^2/\mathrm{m}^3$ for the deciduous canopies present on the CMU campus, an isotropic Forchheimer coefficient of

$$f = (0.6,\ 0.6,\ 0.6)\ \mathrm{m}^{-1}.$$

was prescribed for porous zones. Isotropy (tree offers the same flow resistance horizontally through or vertically) was assumed in the absence of species-specific anisotropy data for the site.

3.3.6. Time window and boundary conditions

A heatwave period measured by the MORICHI dataset (Martin et al., 2026) was selected for analysis. The hottest day (peak of air temperature) recorded across all four street-level weather stations, 27 August 2024, was identified as the representative day. Although peak temperatures occurred between 15:00 and 16:00, the interval from 13:00 to 14:00 was selected for analysis because it offered minimal wind-direction variability while also presenting high air temperature.

Inflow boundary conditions are collected from the airport meteorological stations (Iowa Environmental Mesonet, 2026) both from PIT and AGC airports and extracted for this time interval. These measurements are taken at 10 m height (z_{ref}), which is ideal to represent the upstream atmospheric conditions. During the selected time window, the prevailing wind direction was 157° azimuth, which defines the orientation of the simulation domain inlet and the average wind speed was 1.6 m/s, used as U_{ref}.

For SF the atmospheric boundary layer inlet profile is assumed to have a neutral stratification, neglecting temperature. For both SF and uMF, the streamwise mean wind speed follows the logarithmic law (Franke et al., 2007),

$$U(z) = \frac{u^*}{\kappa} \ln\left(\frac{z + z_0}{z_0}\right) \tag{3.2}$$

where $U(z)$ is the mean horizontal wind speed at height z above the ground, u^* is the friction velocity, $\kappa \approx 0.42$ is the von Kármán constant, and z_0 is the aerodynamic roughness length set at 0.5 m corresponding to Wieringa (1992) reference. The friction velocity u^* is obtained by fitting Equation 3.2 to the reference wind speed U_{ref} at the reference height z_{ref}, and results in $u^* \approx 0.22$ m/s.

Realisable $k - \epsilon$ turbulence model was used, with inlet profiles for the turbulent kinetic energy k and its dissipation rate ϵ defined as follows to maintain a horizontally homogeneous atmospheric boundary layer through the domain:

$$k(z) = \frac{u^{*2}}{\sqrt{C_\mu}} \tag{3.3}$$

$$\epsilon(z) = \frac{u^{*3}}{\kappa\,(z + z_0)} \tag{3.4}$$

where $C_\mu = 0.09$ is a model constant of the k–ϵ model. The first expression assumes that k is constant with height in the surface layer (a common simplification).

For microclimate simulations, in addition to the boundary conditions already described, thermal, moisture, radiation, and modified-pressure boundary conditions must be prescribed at every boundary of the domain. Air temperature and humidity at the inlet are taken from same sources as the wind speed (Iowa Environmental Mesonet, 2026). Air temperature 26 °C, 299 K and air humidity at w = 0.014 kg/kg for the same hour window. Both quantities are imposed uniformly with height at the inlet, no vertical profile is prescribed, which represents a simplification of the actual atmospheric stratification.

The radiation fluxes are also prescribed at every boundary. The short-wave flux `qs` (solar radiation) is computed via the `solarLoadRadiationViewFactor` boundary condition, which combines ray-cast direct solar radiation with diffuse and reflected radiation weighted by the surface albedo. Surfaces have different albedos depending on their type: 0.4 for terrain and buildings, and 0.1 for the open boundaries of the domain. The long-wave flux qr is computed via the `greyDiffusiveRadiationViewFactor` boundary condition, which assumes each surface emits and absorbs long-wave radiation as a gray, diffuse body with a prescribed emissivity (0.9 for urban surfaces and 1.0 for the open boundaries).

On the fluid side, CFDHAMfluidTemperatureCoupledMixed was used for the temperature of terrain and building surfaces, which uses temperature from solid side and CFDHAMfluidMoistureCoupledImpermeable for the humidity ratio, a zero gradient for an impermeable wall. On the solid side, the corresponding CFDHAMsolidTemperatureCoupledMixed and CFDHAMsolidMoistureCoupledImpermeable were used. MORICHI provides surface temperature measurements for different building materials. Surface temperatures were prescribed using these measurements: Brick wall and Road, which were respectively 32 °C, 305 K and 29 °C, 302 K. The air temperature for this hour registered by the MORICHI stations was around 31 °C, 304 K. Solid thickness was maintained as 0.1 m, following tutorial properties.

Cloud coverage was extracted from Weather Spark (2026), but it was open sky for the time simulated. Materials transport properties for the heat-and-moisture transport model of the solid regions were defined following the uMF tutorial examples, see Table 3.3.

Table 3.3.: Thermal properties used in the urbanMicroclimate OpenFOAM setup, prescribed in the case files. Moisture transport properties are inherited from the respective `buildingMaterialModel` classes.

Material	Density (kg/m^3)	Specific heat (J/kgK)	Thermal conductivity (W/mK)
AsphaltConcrete	1800	1200	0.8
Soil	1150	650	1.5
Hamstad5Brick	1600	1000	0.682

3.3.7. Convergence criteria

For both solvers second order schemes were used. The simulations were run for 800 iterations in SF and 2500 for uMF, reaching residuals in the order of 10^{-4}–10^{-6} for most fields and 10^{-2} for p_{rgh} in uMF. More details about these differences are presented in 3.4.2.

3.3.8. Evaluation procedure

Assessment of feasibility on realistic terrain

The feasibility of applying uMF to a realistic urban configuration with complex terrain is assessed qualitatively through an incremental development procedure. Because published applications of uMF to non-flat terrain are, to the authors' knowledge, absent and the documentation of its geometric and meshing requirements is limited, the configuration is built up in three successive stages of increasing complexity: (i) buildings on flat terrain, (ii) buildings on a reduced domain with complex terrain, and (iii) the full domain described in Section 3.3.3 with complex terrain at the recommended size.

At each stage, the geometric and meshing modifications required for the case to pass the pre-processing pipeline (notably `snappyHexMesh`, `faceAgglomerate`, and `viewFactorsGen`) and to reach a stable solution are recorded.

The same geometry that produces a stable SF run is used as the starting point, and every adaptation required by uMF is documented relative to that baseline.

Comparison of mesh, solver parameters, and computational cost

The two solvers are compared on the same geometry, mesh, and inflow conditions, so that any differences in mesh requirements, numerical configuration, and computational cost can be attributed to the solver physics rather than to case setup. The numerical configuration of each solver, including linear solvers, relaxation factors, convergence tolerances, and the set of transported fields, is documented and contrasted. Computational cost is reported as wall-clock time on a fixed number of processors (20 cores) for a one-hour simulated `timeStep`, together with the number of outer iterations required to reach the residual thresholds specified. The relative cost is expressed as the ratio of uMF to SF wall-clock time on identical hardware.

Wind speed validation against field measurements

Simulation outputs were evaluated by extracting wind-speed probes at the four MORICHI street-level sensor locations and comparing them against the distribution (mean, median, minimum and maximum) of the 5-minute observations recorded for the selected time window. For each station, the observed hourly mean U_o, standard deviation σ_o, and the predicted U_p from each solver are reported, together with the relative error $\varepsilon_\mathrm{rel} = |U_\mathrm{p} - U_\mathrm{o}| / U_\mathrm{o}$. The anemometer's specified accuracy of ±0.89 m/s (Martin et al., 2026) exceeds three of the four observed hourly means. The 5-minute standard deviations and bounding values provide a tighter estimate of within-hour variability and are used as the primary uncertainty scale in this evaluation.

Model performance is also quantified using five dimensionless statistical indicators: fractional bias (FB), normalised mean square error (NMSE), geometric mean bias (MG), geometric variance (VG), and the fraction of predictions within a factor of two of the observations (FAC2). The acceptance bands recommended by Sabatino et al. (2011), are positive FB and MG > 1 indicate model under-prediction, whereas negative FB and MG < 1 indicate over-prediction. NMSE and VG measure the magnitude of scatter regardless of direction; they equal zero (NMSE) or one (VG) for perfect agreement and grow as the discrepancy increases.

Together, the errors per-station and the aggregate indicators allow the wind-only and coupled solvers to be compared locally and globally.

Summary

A summary of the assumptions adopted in this study is presented in Table 3.4.

Table 3.4.: Summary of the main modelling assumptions adopted in this study.

Category	**Assumption**	**Justification / Source**
Geometry and domain		
Building detail	Buildings reconstructed as LoD 1.2 within 300 m radius of the weather stations	Limitation imposed by uMF
Vegetation	Trees modelled as porous zones (Darcy–Forchheimer), VEG module not used	Simplification
Porous coefficient	Isotropic Forchheimer coefficient $f = (0.6, 0.6, 0.6)\ \mathrm{m}^{-1}$	$C_d \approx 0.2$ (Katul and Albertson, 1998; Massman, 1987), LAD $= 3\ \mathrm{m}^2/\mathrm{m}^3$
Domain size	Lateral 5H, upstream 5H, downstream 15H, vertical >6H	Standard guidelines Franke et al. (2007); Blocken (2015)
Solid thickness	0.1 m for terrain and buildings with 10 layers	uMF tutorial default; simplification
Flow and turbulence		
Turbulence model	Realisable k–ε	uMF tutorial default; Standard for ABL urban CFD
Inlet wind profile	Logarithmic law (Eq. 3.2)	Standard for ABL urban CFD Franke et al. (2007)
Roughness length	$z_0 = 0.5$ m	Wieringa (1992) reference for the surrounding fetch
Inlet TKE	$k(z)$ constant with height (Eq. 3.3)	Standard for ABL urban CFD
Numerical schemes	Second-order schemes for both solvers	More accurate output
Thermal and moisture (uMF only)		
ABL stratification (SF)	Neutral; temperature neglected	wind-only baseline
Inlet T and humidity	Uniform with height (no vertical profile)	uMF tutorial default
Surface temperatures	Brick wall 305 K; Road 302 K (constant)	MORICHI measurements (Martin et al., 2026)
Wall moisture BC	Impermeable (zero gradient) HAM still solved inside solids	uMF tutorial default
Material properties	AsphaltConcrete, Soil, Hamstad5Brick (Table 3.3)	uMF tutorial examples
Radiation		
Sky condition	Open sky (no clouds)	Weather Spark (2026)
Albedo	0.4 terrain/buildings; 0.1 open boundaries	Prescribed per surface type
Emissivity	0.9 urban surfaces; 1.0 open boundaries	Grey-body assumption
Short-wave model	`solarLoadRadiationViewFactor` (direct + diffuse + reflected)	OpenFOAM RAD utility
Long-wave model	`greyDiffusiveRadiationViewFactor` (grey, diffuse)	OpenFOAM RAD utility

Time window and boundary data		
Representative period	13:00–14:00 on 27 August 2024	Minimal wind-direction variability; high T
Wind direction	157° azimuth (uniform inlet orientation)	Average of PIT/AGC airport data (Iowa Environmental Mesonet, 2026)
Reference wind speed	$U_{\text{ref}} = 1.6$ m/s at $z_{\text{ref}} = 10$ m	Average of PIT/AGC airport data (Iowa Environmental Mesonet, 2026)
Friction velocity	$u^* \approx 0.22$ m/s (derived)	Fit of log law to U_{ref}
Inlet temperature/humidity	$T = 299$ K, $w = 0.014$ kg/kg	(Iowa Environmental Mesonet, 2026)

3.4. Results

3.4.1. Application of urbanMicroclimateFoam to realistic urban models

This subsection evaluates the feasibility of applying uMF to realistic urban configurations, including detailed building and the terrain mesh. The investigation demonstrates that simulations over realistic terrain are achievable, however, it is subject to several constraints discussed below.

The model that performed well in SF could not be used directly with uMF; the transition required progressive adaptation of the geometry to achieve numerical stability. The following modifications were required:

Removal of non-radiating faces The ground faces of buildings, the shared walls between adjacent buildings, and the building footprints in the terrain mesh had to be removed (Figures 3.6 and 3.7). These faces do not exchange radiation with the sky or with other exposed surfaces, and including them causes the `viewFactorsGen` step to fail or to produce non-physical view factor coefficients. Specially party-walls limited the use of higher LoD levels, which include these faces, because the manual removal would be very time-consuming.

Separating building surfaces Building surfaces had to be separated from each other prior to meshing to avoid face pyramid errors introduced by the `extrudeMesh` utility. A face pyramid error occurs when a cell centre lies on the wrong side of one of its own faces, producing a negative pyramid volume and rendering the cell geometrically invalid. Surface separation can be achieved by duplicating shared vertices between adjacent surfaces so that each surface owns its own independent set of vertices. Although the duplicated vertices occupy identical positions in space, the surfaces are no longer topologically connected, which prevents the face pyramid errors. This was achieved using the `Unweld` tool in Rhino, see Figure 3.7.

Surface mesh refinement Refinement of the solid surface meshes is not strictly required for a successful uMF run; however, the geometric discretisation affects the local results because field values are defined at the centroid of each cell (Moukalled et al., 2016). As shown in Figure 3.8, local surface temperatures differed by up to 1 °C between the coarse and refined cases. Such refinement requires re-meshing.

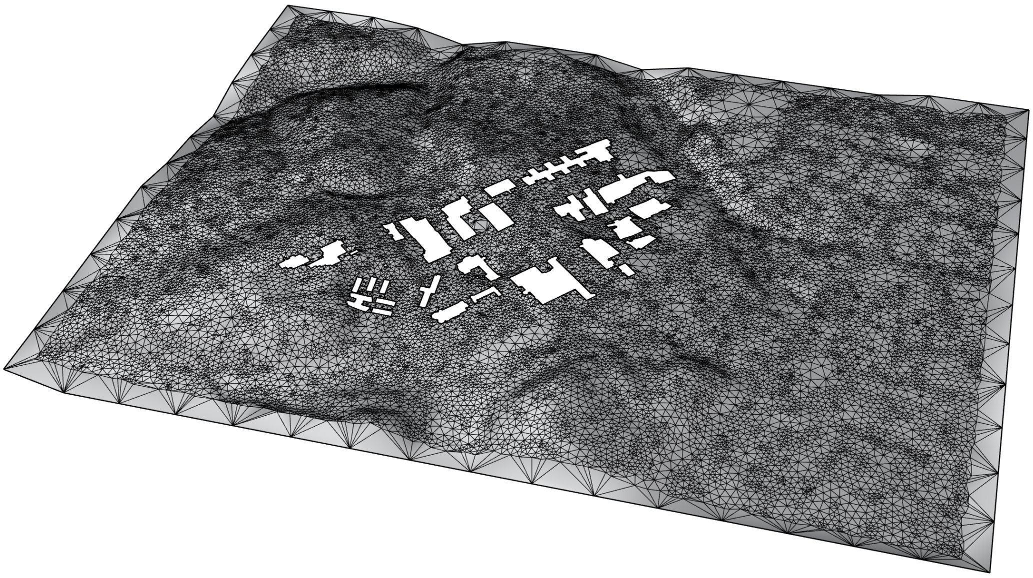

Figure 3.6.: Terrain mesh with building footprints removed, as required for uMF.

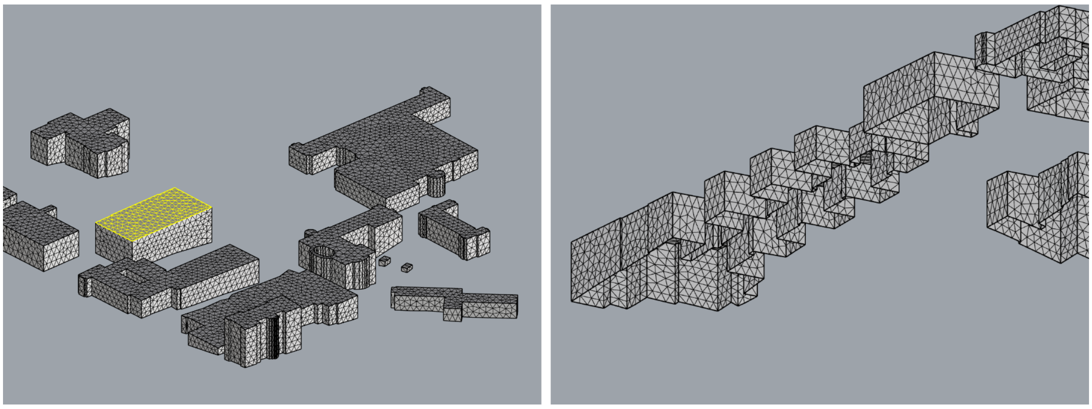

Figure 3.7.: Building surfaces are separated from each other (the yellow surface shares no vertices with adjacent surfaces), and bottom faces and party walls are removed, as required by the uMF solver.

Isotropic re-meshing is available in City4CFD (for buildings), Rhino and PyMeshLab (for arbitrary input meshes). However, none of these tools reliably produced a usable mesh: re-meshing modifies the discrete representation of solid boundaries and can introduce topological issues, such as gaps between adjacent surface sources (e.g. building and terrain). Figure 3.11 shows an example of this, two mesh files were re-meshed independently, and the resulting edges no longer aligned along their shared boundary. Controlling the mesh edge size can help, however it is not guaranteed to work. This proved especially problematic in steep terrain areas. These gaps cause `snappyHexMesh` to create air region below the terrain. This required manual geometric correction of the surfaces as a workaround.

Mesh quality at small geometric features Numerical instabilities occurred in regions of high mesh non-orthogonality near small geometric features of the solid surfaces, typically building and terrain corners or details smaller than the final cell size after castellation (Figure 3.9). For example, a building detail of approximately 60 cm caused the simulation to fail with a 1 m background mesh, but ran successfully when the mesh was refined to 60 cm. A similar issue arose at building–terrain interfaces, where small or sliver-shaped terrain face (Figure 3.10) produced low-quality cells after castellation and unphysical surface temperatures. However, controlling the sizes and form of terrain mesh faces after automatic reconstruction pipelines to avoid the aforementioned problems is still not possible with the available tools. Therefore this also required manual geometric correction. Within this study, such instabilities occurred more frequently at terrain surfaces than at building surfaces. The underlying cause could not always be identified, suggesting that this remains an open practical issue when applying uMF to realistic, complex terrain.

Radiation preprocessing failures A geometry that runs successfully through all stages of SF recurrently failed during the `faceAgglomerate` or `viewFactorsGen` steps of the radiation pre-processing. Existing non-orthogonality in the mesh breaks the `face Agglomerate` algorithm capacity of finding neighbouring faces. Moreover, `faceAgglomerate` is prone to produce non-manifoldness, which breaks the geometry topological structure, see Figure 3.12. Although `faceAgglomerate` itself completes without error, the resulting non-manifold geometry is not accepted by `viewFactorsGen`. Overcoming these issues required a trial-and-error tuning of the `faceAgglomerate` parameters (`nFacesInCoarsestLevel` and `featureAngle`), which proved time-consuming.

Domain shape Although a circular domain was initially adopted to enable multi-directional inflow analysis, this configuration proved incompatible with uMF due to numerical instabilities. The clear reason for these instabilities was not found within the scope of this research. Circular domains additionally entail a higher cell count and therefore a greater computational cost. Therefore, the rectangular domain was adopted in its place.

Table 3.5 summarises all necessary geometric and meshing adaptations for uMF to run. Together, they represent a substantial increase in preprocessing effort compared with the workflow that is accepted by SF. The available algorithms and tools do not currently support a fully automated path from a realistic urban geometry with complex terrain to a stable uMF case. The corrections required are case-specific and largely manual, and represent a practical barrier to broader adoption of uMF for non-flat urban sites.

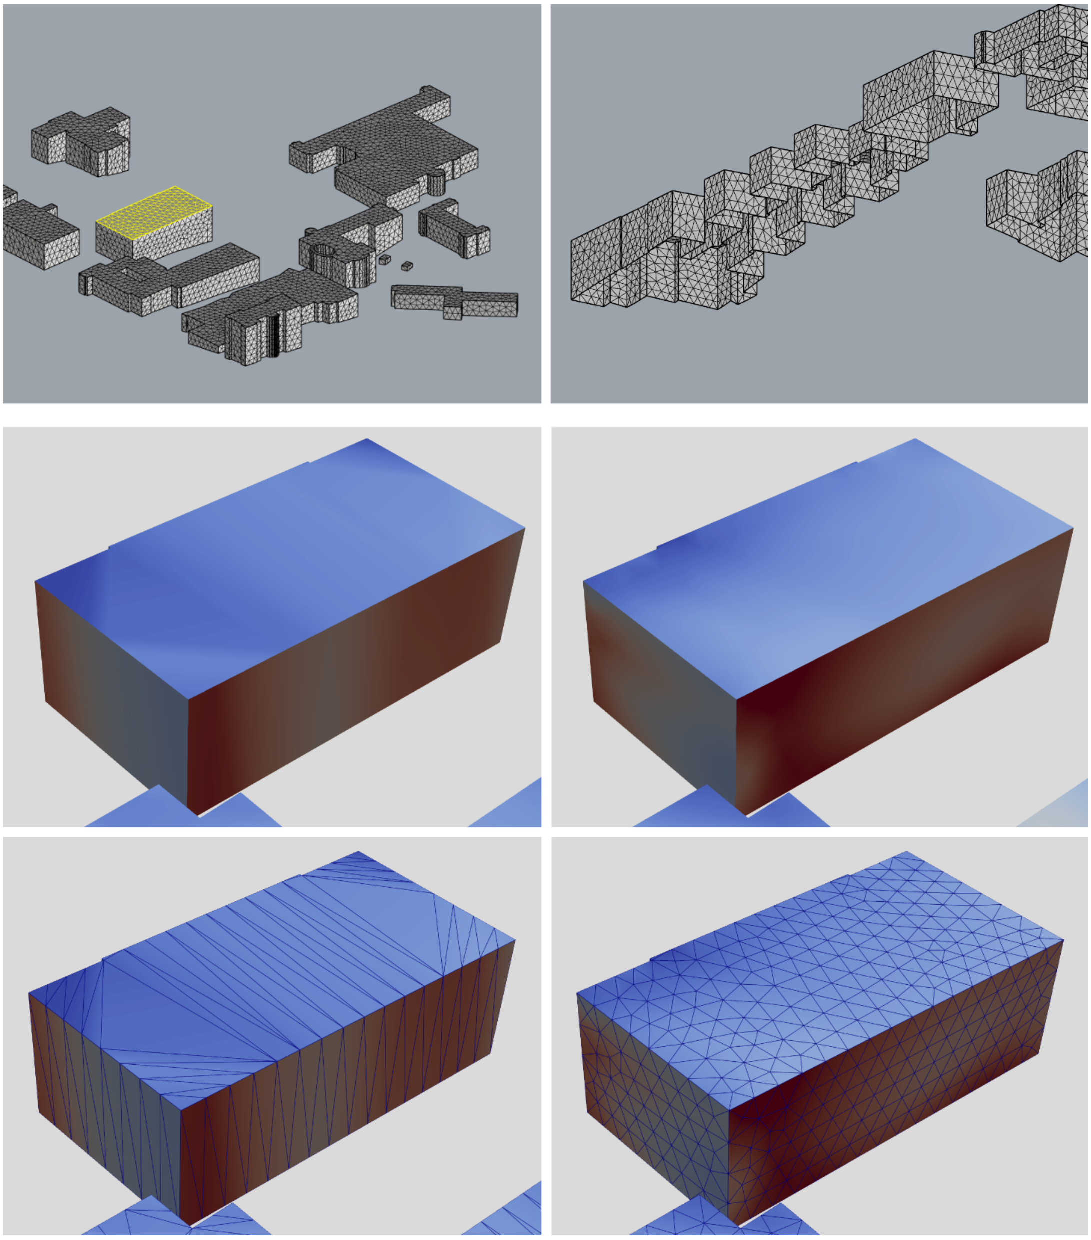

Figure 3.8.: Left column: buildings with large unrefined mesh, Right column: refined isotropic mesh. Colours represent surface temperature in Kelvin

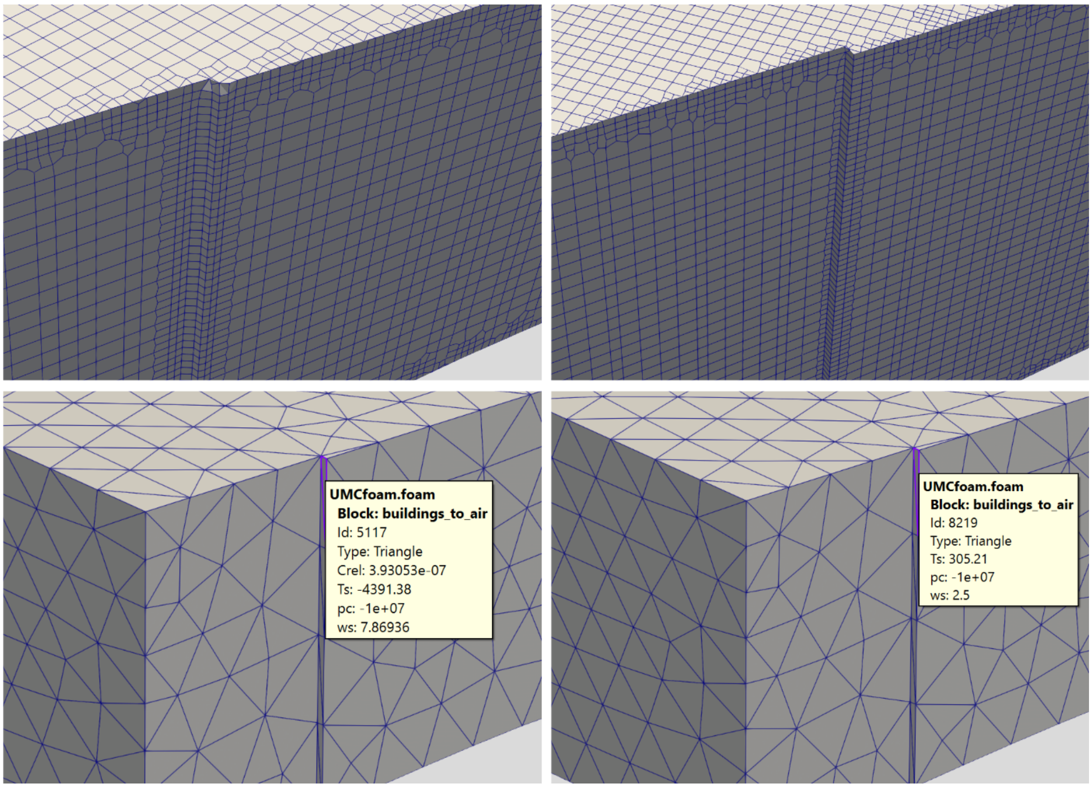

UMCfoam.foam
Block: buildings_to_air
Id: 5117
Type: Triangle
Crel: 3.93053e-07
Ts: -4391.38
pc: -1e+07
ws: 7.86936

UMCfoam.foam
Block: buildings_to_air
Id: 8219
Type: Triangle
Ts: 305.21
pc: -1e+07
ws: 2.5

Figure 3.9.: Effect of mesh resolution on HAM stability at a building corner. Left column: coarser mesh cells (1 m) fail to conform to the building geometry (top), producing unphysical surface temperatures (bottom) and causing HAM solver failure. Right column: the same corner reconstructed with finer cells (60 cm) closely follows the building detail (top), yielding physically consistent surface temperatures (bottom).

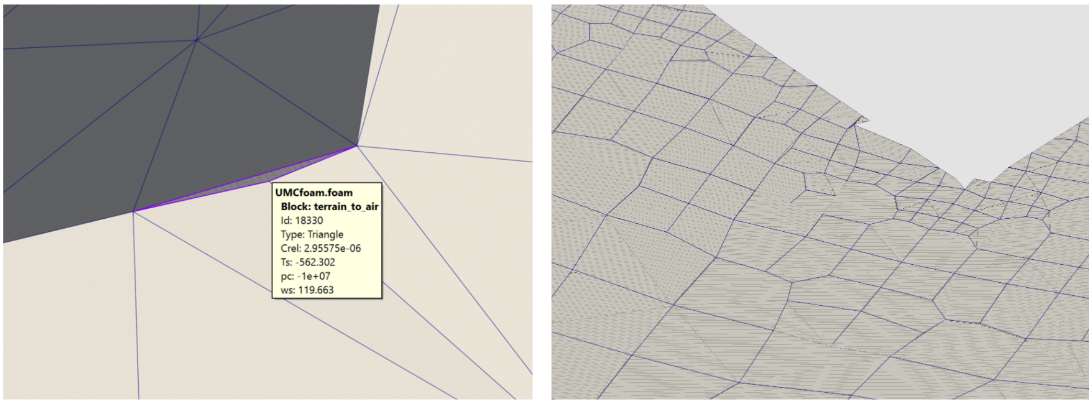

UMCfoam.foam
Block: terrain_to_air
Id: 18330
Type: Triangle
Crel: 2.95575e-06
Ts: -562.302
pc: -1e+07
ws: 119.663

Figure 3.10.: Low-quality `snappyHexMesh` cells (right), which were generated around a small terrain face at a building interface (left). This combination produces unphysical surface temperatures analogous to those shown in Figure 3.9.

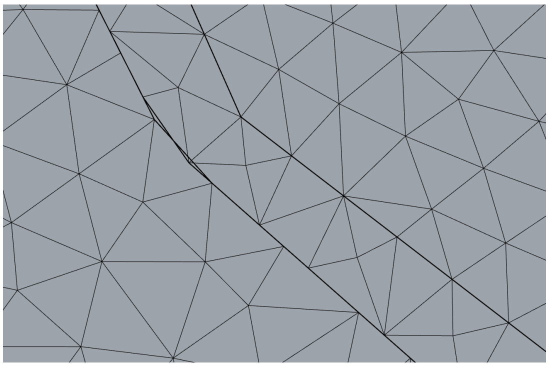

Figure 3.11.: Topological inconsistency introduced by isotropic re-meshing in PyMeshLab. The surfaces were re-meshed independently, altering their edges and producing gaps or overlapping each other.

Figure 3.12.: `faceAgglomerate` output for one of its agglomerations. This non-manifold face is going to fail in the `viewFactorsGen` step.

Table 3.5.: Summary of geometric and meshing adaptations required for uMF stability on a realistic urban model with complex terrain, relative to the SF baseline.

Adaptation	**Issue addressed**	**Workaround / Cost**
Removal of non-radiating faces	Building bottoms, party walls, and building footprints in terrain cause `viewFactorsGen` to fail or produce non-physical view factors	Remove bottom option (City4CFD); Limits use of higher LoD levels
Separation of building surfaces	Shared vertices between adjacent building surfaces cause face pyramid errors after `extrudeMesh`, rendering cells geometrically invalid	Duplicate shared vertices so each surface owns independent vertices at the same geometric position (`Unweld` in Rhino); small additional preprocessing step
Surface mesh refinement	Coarse surface discretisation shifts cell centroids and changes local surface temperatures (up to 1 °C)	Isotropic re-meshing (City4CFD, Rhino, PyMeshLab); not strictly required for stability, affects accuracy
Topological consistency after re-meshing	Independent re-meshing of adjacent surfaces (e.g. building–terrain) creates gaps or overlaps, causing `snappyHexMesh` to generate air cells below terrain	Manual geometric correction; no reliable automated tool found
Mesh quality at small features	Building/terrain details smaller than the background cell size and sliver terrain faces produce high non-orthogonality, unphysical surface temperatures, and solver failure	Local refinement to match feature size; manual cleanup of sliver faces, more frequent on terrain
Radiation preprocessing tuning	`faceAgglomerate` fails on non-orthogonal meshes and produces non-manifold faces that `viewFactorsGen` rejects	Trial-and-error tuning of `nFacesInCoarsestLevel` and `featureAngle`; time-consuming
Domain shape	Circular domain (initially adopted for multi-directional inflow) caused unresolved numerical instabilities and higher cell count	Switched to rectangular domain; multi-directional analysis sacrificed

3.4.2. Computational cost

The fact that the model cells had to be decreased in size to overcome failures like the one presented in Figure 3.9 already indicates a higher computational cost relative to the regular SF case. Nevertheless, wall-clock time on a fixed processor count (20 cores) was tracked and compared for both solvers using the same geometry.

Both solvers were set to use the SIMPLEC (SIMPLE-Consistent) algorithm. However, they differed in the numerical configuration, see Table 3.6. uMF solves additional, coupled equations (energy, humidity, and air–solid exchange) requires smaller update steps, and tighter convergence tolerances. Relaxation factors had to be kept lower to avoid instability. This increases the cost of a single iteration and increases the number of iterations necessary for convergence, contributing to the higher overall computational cost of uMF.

For a one-hour `timeStep` and with 20 processors, uMF took 6 hours, 28 minutes, and 20 seconds to complete 2500 inner iterations for the air, exchange, and solid steps, whereas SF took 1 hour, 48 minutes, and 39 seconds for 830 fluid iterations. This represents an increase of approximately 260% in wall-clock time, i.e. uMF was roughly 3.6× slower than SF on the same geometry and processor count.

Table 3.6.: Numerical configuration differences between SF and uMF.

Aspect	**SF**	**uMF**	**Cost effect**
Linear solver (U, k, ε)	`smoothSolver` (GaussSeidel)	`PBiCGStab` + `DILU`	More work per inner iteration in uMF
Linear solver (pressure)	`GAMG` on p `relTol` $= 0.1$	`GAMG` on p_{rgh} `relTol` $= 0.05$	Tighter pressure solve in uMF
`relTol`	0.1	0.1	Comparable inner tolerance
Relaxation (fields)	$p = 0.9$	$p_{rgh} = 0.5$ $\rho = 1.0$	Stronger under-relaxation in uMF
Relaxation (equations)	$U = 0.9$ $k = 0.9$ $\varepsilon = 0.9$	$U = 0.7$ $k = 0.5$ $\varepsilon = 0.5$ $h = 0.9$ $w = 0.9$ $T_l = 0.3$	Stronger under-relaxation in uMF
Number of fields	U, p, k, ε	$+ w$, p_{rgh}, T_l, + solid fields $, q_s, q_r, \ldots)$	More equations to solve each step

3.4.3. Wind speed predictions comparison

Figure 3.13 shows the distribution of the MORICHI 5-minute measurements (maximum, minimum, mean and median) within the target hour (13:00–14:00) compared against the simulation predictions at the sensor location for each of the four stations. Table 3.7 reports these values together with the standard deviations σ_o, relative error ε_{rel}, and difference of simulations at each station.

uMF achieves a FAC2 of 0.75, meaning that three of the four predictions fall within a factor of two of the observed mean, compared with two of four for SF (FAC2 = 0.50). Both solvers achieve the 0.5 acceptance threshold recommended by Sabatino et al. (2011), but uMF does so by a slightly wider margin (one more station). The other bias-related indicators show the same tendency: FB drops from $+0.65$ in SF (failing the $|\text{FB}| < 0.3$ acceptance band) to -0.11 in uMF (passing it), and MG moves from 1.89 to 0.83, crossing the acceptance band in the same direction.

The two solvers show different error tendencies: as evidenced in Figure 3.13 and Table 3.8, SF under-predicts wind speed at every station, while uMF over-predicts on average. SF produces a calmer wind field overall, as visible in the cross-sections of Figures 3.14–3.17, which show U_x, U_z, U for both solvers and air temperature T for uMF. The vertical component U_z is generally stronger in the uMF predictions, consistent with the inclusion of buoyancy: warm air rises, driving vertical motion that is much weaker in the wind-only solver. See also plan views 3.18 and 3.19.

At the per-station level, uMF reduces the relative error in predicted wind speed by 29, 15, and 35 percentage points at stations 1, 3, and 4 respectively. Station 2 is the exception: SF matches the observation closely ($\varepsilon_{\text{rel}} = 4\%$) while uMF over-predicts substantially ($\varepsilon_{\text{rel}} = 139\%$).

The relative errors should be interpreted with caution, since the observed wind speeds at all four stations are low (0.44–1.29 m/s), so small absolute differences translate into large percentage errors. Moreover, the anemometer's manufacturer-specified accuracy of ±0.89 m/s (Martin et al., 2026) exceeds the observed hourly mean at three of the four stations: only Station 1 (observed mean 1.29 m/s) clears this threshold, and is therefore the only location at which a per-station comparison can be drawn on independently of instrument uncertainty. At Stations 2, 3, and 4, per-station error figures are dominated by instrument accuracy rather than by solver behaviour, and the more reliable basis for solver comparison is the aggregate indicators reported above.

At Station 1, uMF predicts a value within one standard deviation of the observed mean, though at the upper edge of that range (Figure 3.13). Station 1 is where SF performs the worst, strongly under-predicting the wind speed. Among the four stations, this is where the vertical component U_z predicted by uMF is most pronounced (Figure 3.14), and where local temperatures are highest. Strong vertical motion and elevated local temperatures are consistent with buoyancy contributing non-negligibly to the flow at this location, which explains why the coupled solver outperforms the wind-only baseline here. The observed standard deviation at this station is also the highest of the four ($\sigma_o = 0.52$ m/s), indicating greater flow variability during the analysis window.

At Station 2, both solvers accelerate the flow near the building wall adjacent to the station, as shown in Figures 3.13 and 3.15. This acceleration is in line with the documented downwash effect, in which wind striking the windward face of a building is deflected downwards, amplifying near-surface wind speeds. The strength of this effect scales with the height of the building Mittal et al. (2018), since a taller windward facade intercepts faster upper-level flow. The building adjacent to station 2 was reconstructed at LoD 1.2, which in this case converts the original gabled roof into a flat roof at a height above the original eave, as visible in Figure 3.20. Together with the downward-deflected flow visible along the windward facade in Figure 3.15 ($U_z < 0$), this suggests that the geometric simplification amplifies the wall-driven acceleration at Station 2 in both solvers. This was also documented in the wind-tunnel measurements of Zhao et al. (2021) for tall windward buildings in steep street canyons. In SF, this acceleration coincidentally compensates for the solver's general tendency to under-predict, producing a value close to the observation. In uMF, the same acceleration is augmented by buoyancy-driven motion at the warm wall, producing the over-prediction. The current dataset cannot, however, fully prove this hypothesis and a sensitivity analysis with the original gabled-roof geometry is necessary.

At Station 3, the two solvers are apart in opposite directions and by similar amounts: SF under-predicts and uMF over-predicts, with both predictions falling close to one standard deviation from the observed mean ($\sigma_o = 0.25$ m/s). The observed variability at this station is among the lowest of the four. The results can be considered comparable.

At Station 4, uMF is very accurate (within 1%), while SF falls within one standard deviation below the observed mean. This station's measurements also have a low standard deviation of $\sigma_o = 0.24$ m/s. At this station, U_x and U_z show little difference between the two solvers (Figure 3.17); the small remaining discrepancy is attributable to U_y.

These results support the hypothesis that in the studied low-wind, thermally-active regime, the wind-only solver SF loses accuracy due to the absence of buoyancy-driven mixing, while the non-isothermal uMF captures the missing physics. The evidence is most robust at the aggregate level, with measurable improvement across all five validation indicators. Among the per-station results, Station 1, the only location whose observed mean clears the anemometer accuracy threshold, provides the strongest per-station support, with a 29 percentage point reduction in relative error. The remaining stations are consistent with this picture but cannot be drawn on independently of instrument uncertainty.

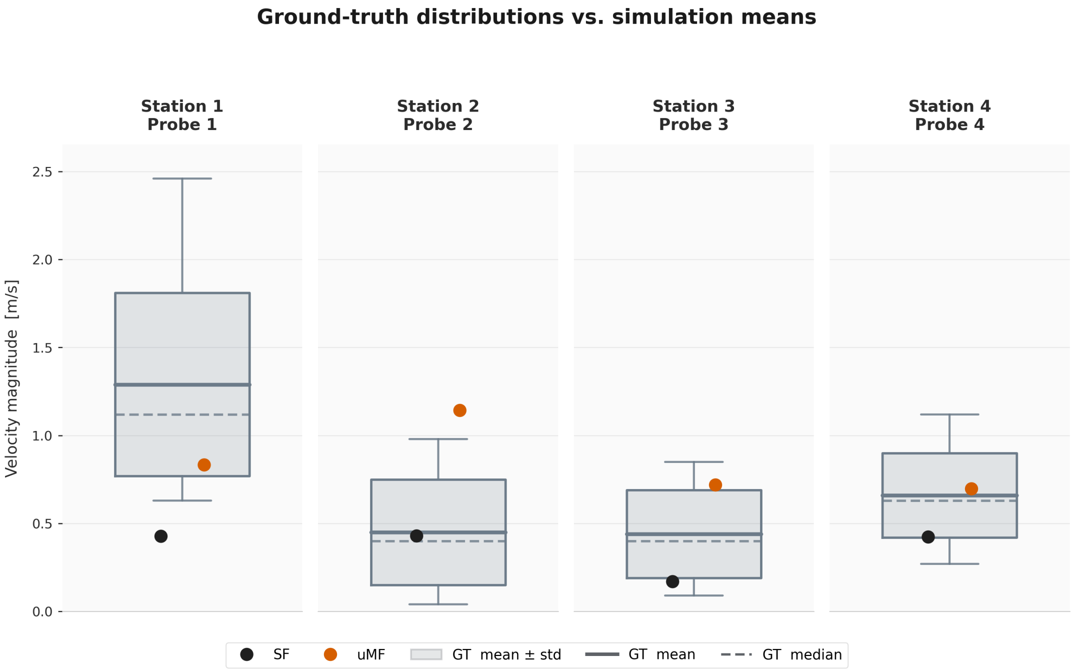

Figure 3.13.: Distributions of the MORICHI 5-minute wind-speed measurements at the four street-level stations during the 13:00–14:00 window on 27 August 2024 (minimum, maximum, mean, and median shown as box-and-whisker plots). Simulation predictions are overlaid as markers: steady-state mean velocity magnitudes from SF and uMF, computed using inflow conditions representative of the same window.

Table 3.7.: Observed and predicted mean wind speed at the four MORICHI street-level stations during the 13:00–14:00 window on 27 August 2024. U_o and σ_o are the hourly mean and standard deviation of the 5-minute observations; U_SF and U_uMF are the simulation predictions. The relative error is defined as $\varepsilon_\mathrm{rel} = |U_p - U_o|/U_o$. The last column reports the change in relative error from SF to uMF in percentage points (pp); negative values indicate that uMF is more accurate, positive values that SF is more accurate.

Station	U_o [m/s]	σ_o [m/s]	U_SF [m/s]	$\varepsilon_\mathrm{rel,SF}$ [%]	U_uMF [m/s]	$\varepsilon_\mathrm{rel,uMF}$ [%]	Difference (SF→uMF)
Station 1	1.29	0.52	0.43	67	0.80	38	−29 pp
Station 2	0.45	0.30	0.43	4	1.07	139	+135 pp
Station 3	0.44	0.25	0.17	61	0.64	46	−15 pp
Station 4	0.66	0.24	0.42	36	0.66	1	−35 pp

Table 3.8.: Validation metrics for SF and uMF computed against the four MORICHI street-level stations. Acceptance bands follow Sabatino et al. (2011). The "Reduction" column reports the change in each metric from SF to uMF, expressed as the relative reduction in distance from the ideal value (ideal: 0 for FB and NMSE, 1 for MG, VG, and FAC2). Negative values indicate improvement; for FAC2, where higher is better, the column reports the direct relative increase, and a positive value indicates improvement.

Metric	**Acceptance**	**SF**		**uMF**		**Reduction**
		Value	Pass	Value	Pass	(SF→uMF)
FB	$\lvert\mathrm{FB}\rvert < 0.3$	+0.65	No	−0.11	Yes	−83%
NMSE	$\mathrm{NMSE} < 4$	0.85	Yes	0.30	Yes	−65%
MG	$\mathrm{MG} < 1$	1.89	No	0.83	Yes	−80%
VG	$\mathrm{VG} < 1.5$	1.79	No	1.32	Yes	−59%
FAC2	$\mathrm{FAC2} \geq 0.5$	0.50	Yes	0.75	Yes	+50%

3.5. Limitations

Several simplifications constrain the generality of these findings.

The anemometer accuracy specified by the manufacturer (±0.89 m/s, Martin et al. (2026)) exceeds the observed hourly mean wind speed at three of the four stations (0.44–0.66 m/s). Strictly interpreted, this means that at those three stations any individual prediction within ±0.89 m/s of an observation is statistically indistinguishable from it on the basis of the instrument alone. Station 1 (observed mean 1.29 m/s) is the only location whose observed mean clears the accuracy threshold, and is therefore the most reliable per-station comparison point. For the remaining stations, the more robust evidence for solver performance comes from the aggregate indicators across all four stations, which describe systematic trends that are less sensitive to single-sensor instrument noise.

Beyond the instrument limitation, the comparison rests on a single one-hour window under one inflow direction, and the threshold above which the wind-only approximation becomes acceptable cannot be established from this dataset alone. Moreover, broader claims about diurnal or seasonal performance would require longer simulation horizons.

Vegetation was represented as Darcy-Forchheimer porosity zones rather than through the full uMF VEG module, omitting evapotranspiration and the associated latent-heat sink. This limitation is particularly relevant for the CMU campus in Pittsburgh, where extensive green areas surround the domain of interest. Therefore, the omission of these effects is likely non-trivial, and may contribute to the average over-prediction observed with uMF.

Buildings within the area of interest were reconstructed at LoD 1.2 to satisfy uMF requirements (e.g. the absence of party walls), which alters roof geometry and may affect local-scale results, such as the flow acceleration observed at Station 2.

Material properties were drawn from the uMF tutorial library rather than calibrated to the specific properties of CMU materials, and identical albedo and emissivity values were assigned to vegetated and non-vegetated surfaces. Moreover, the thicknesses of solid regions were likewise retained from the tutorial cases rather than matched to the actual building envelopes and terrain. Since solid thickness governs thermal mass and the time lag of heat conduction, this simplification may affect the simulated surface temperature evolution. However, prescribing surface temperatures from MORICHI measurements is assumed to provide a reasonable approximation to the real thermal boundary conditions.

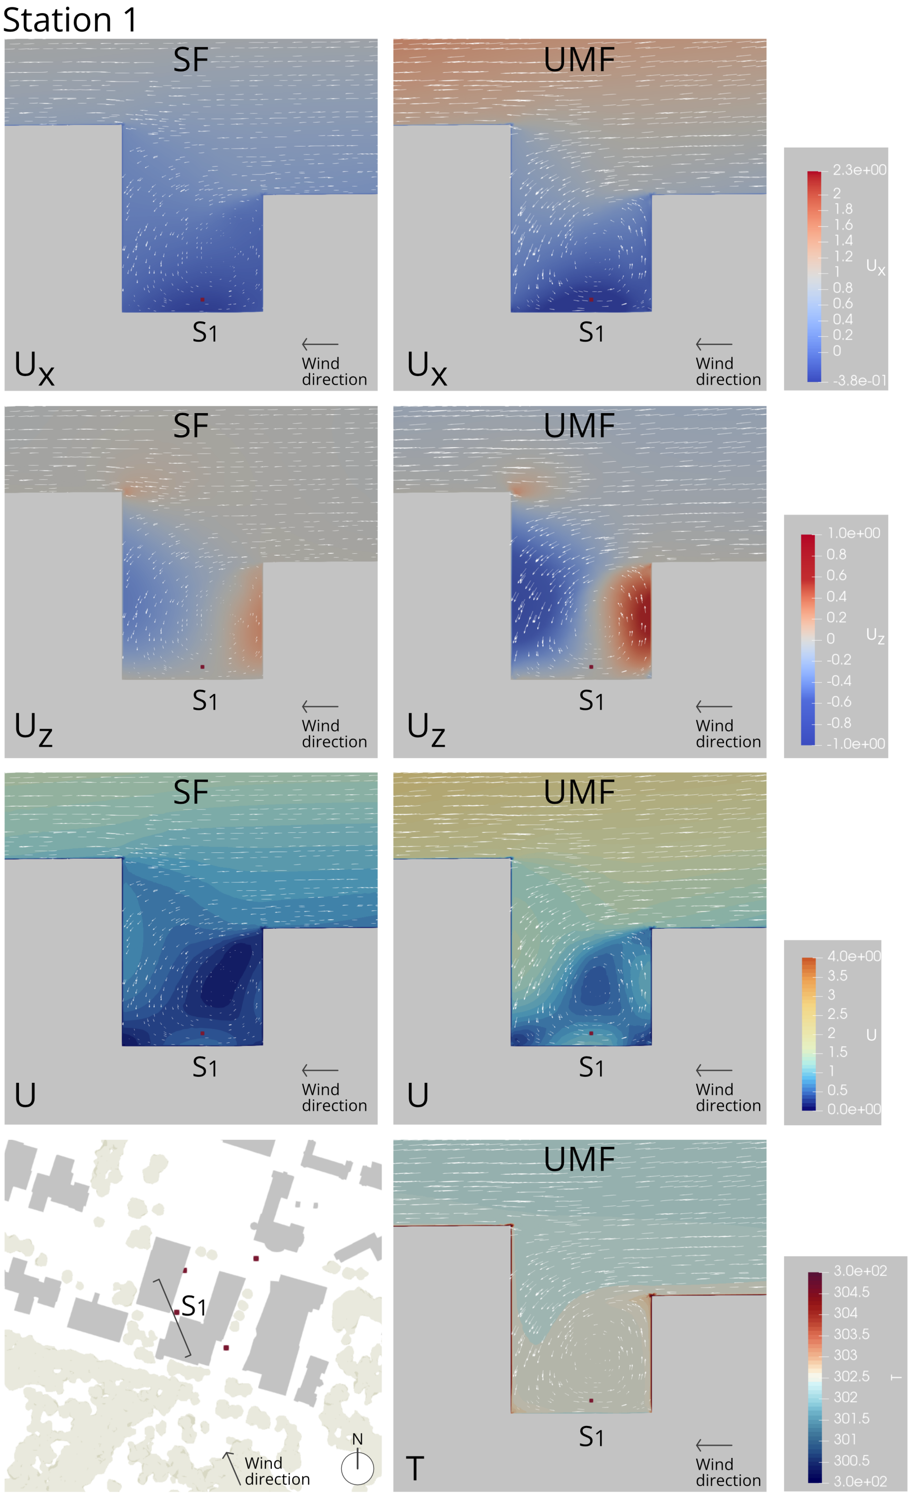

Figure 3.14.: Station 1: cross-sections of the horizontal velocity U_x, vertical velocity U_z, and velocity magnitude U for both solvers, together with air temperature T for uMF only, computed for the 13:00–14:00 time window on 27 August 2024.

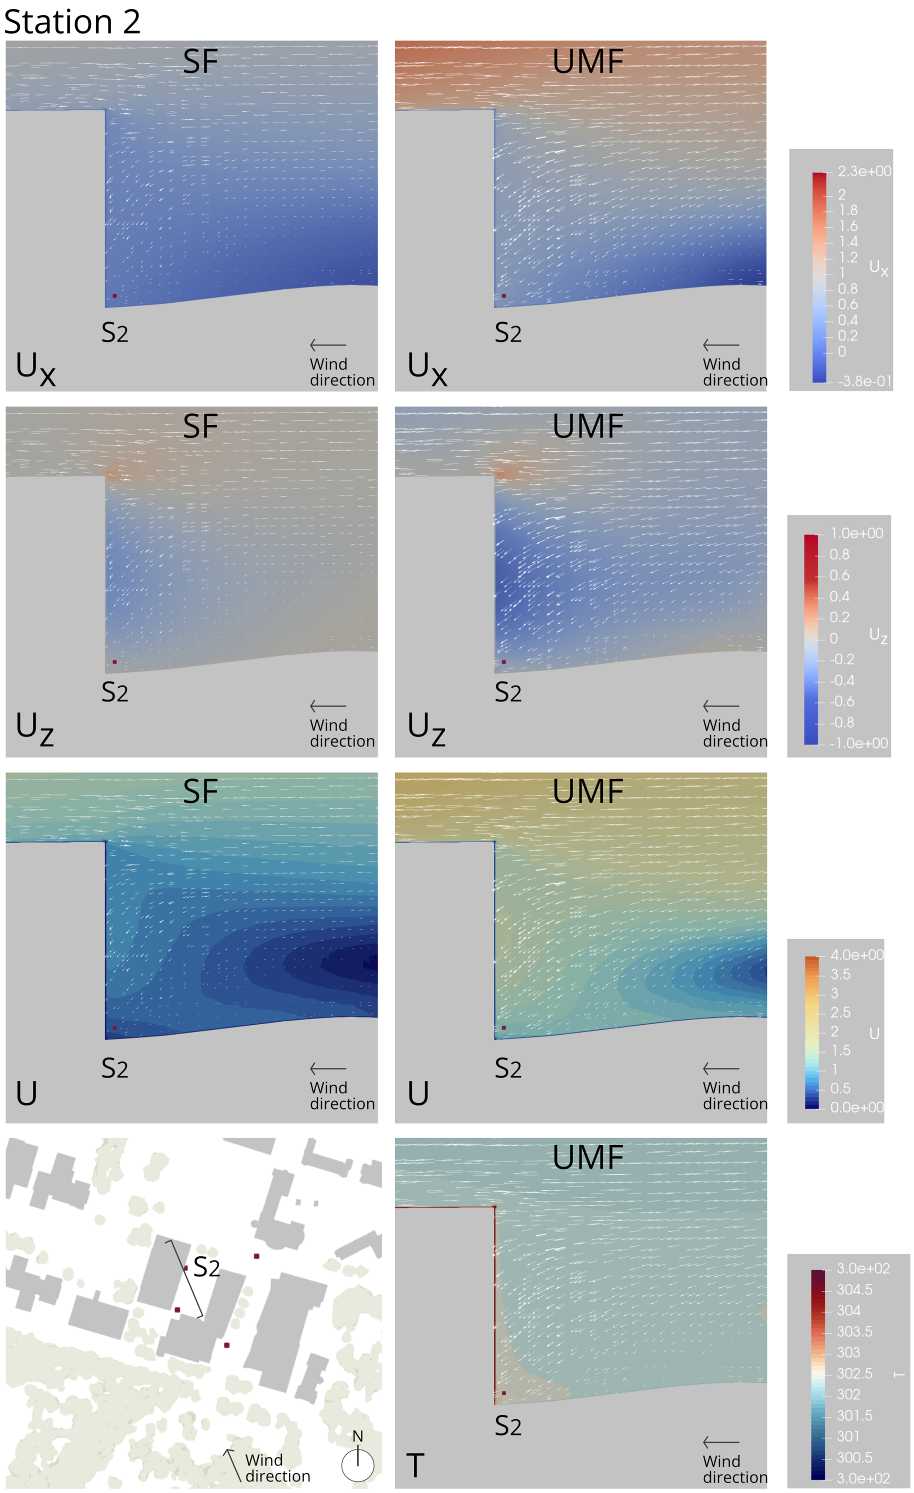

Figure 3.15.: Station 2: cross-sections of the horizontal velocity U_x, vertical velocity U_z, and velocity magnitude U for both solvers, together with air temperature T for uMF only, computed for the 13:00–14:00 time window on 27 August 2024.

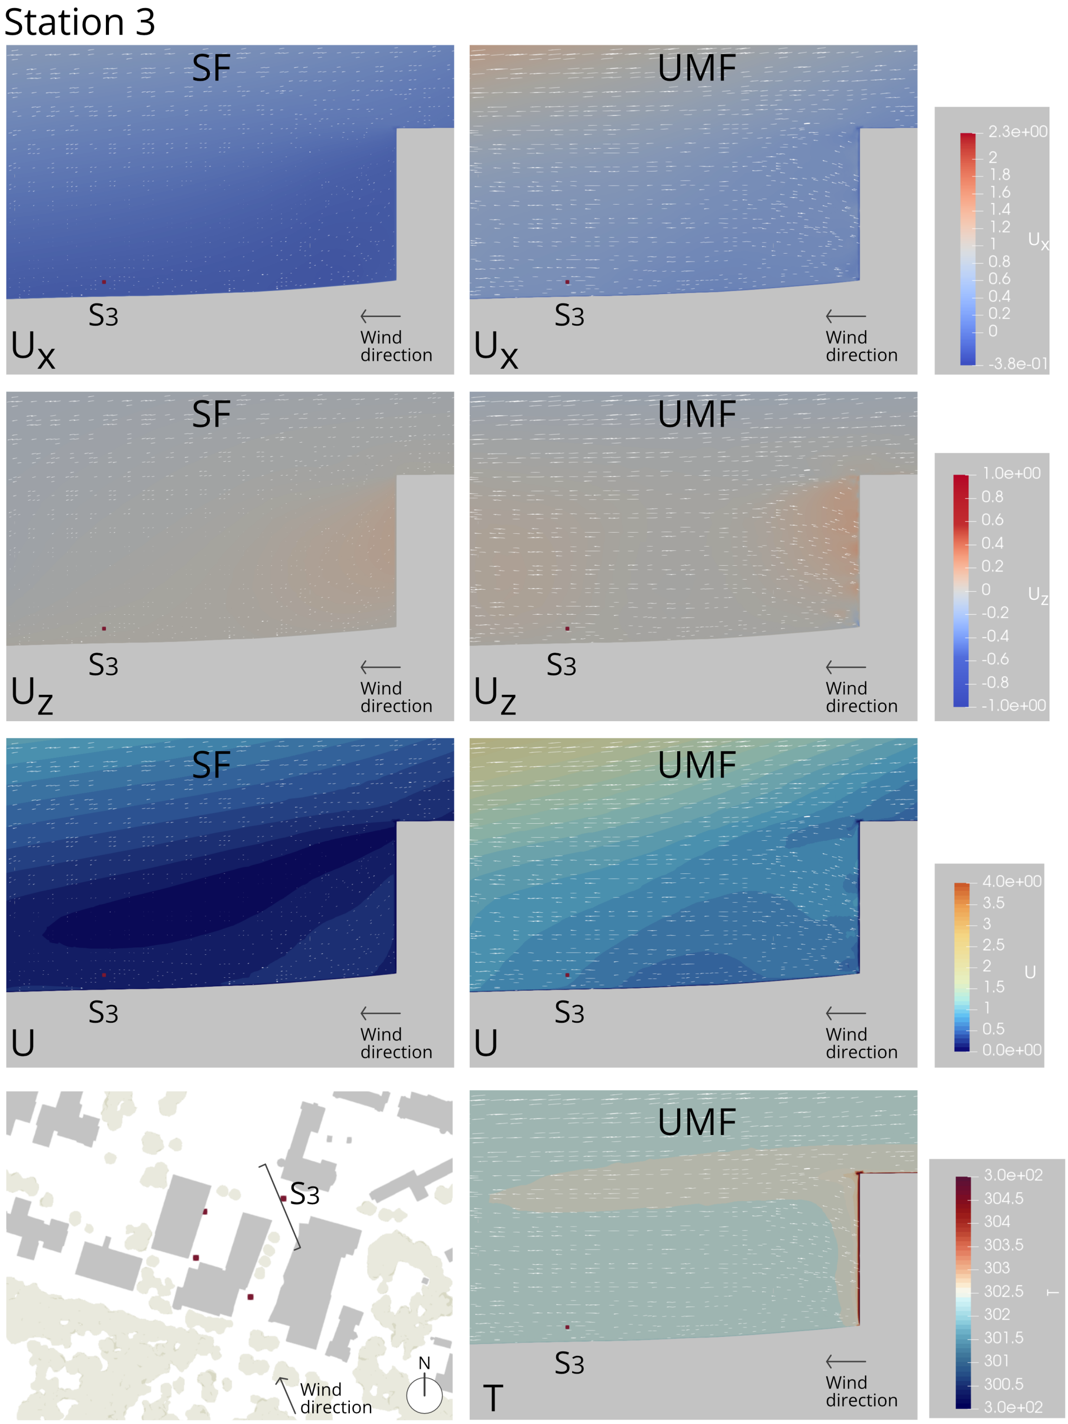

Figure 3.16.: Station 3: cross-sections of the horizontal velocity U_x, vertical velocity U_z, and velocity magnitude U for both solvers, together with air temperature T for uMF only, computed for the 13:00–14:00 time window on 27 August 2024.

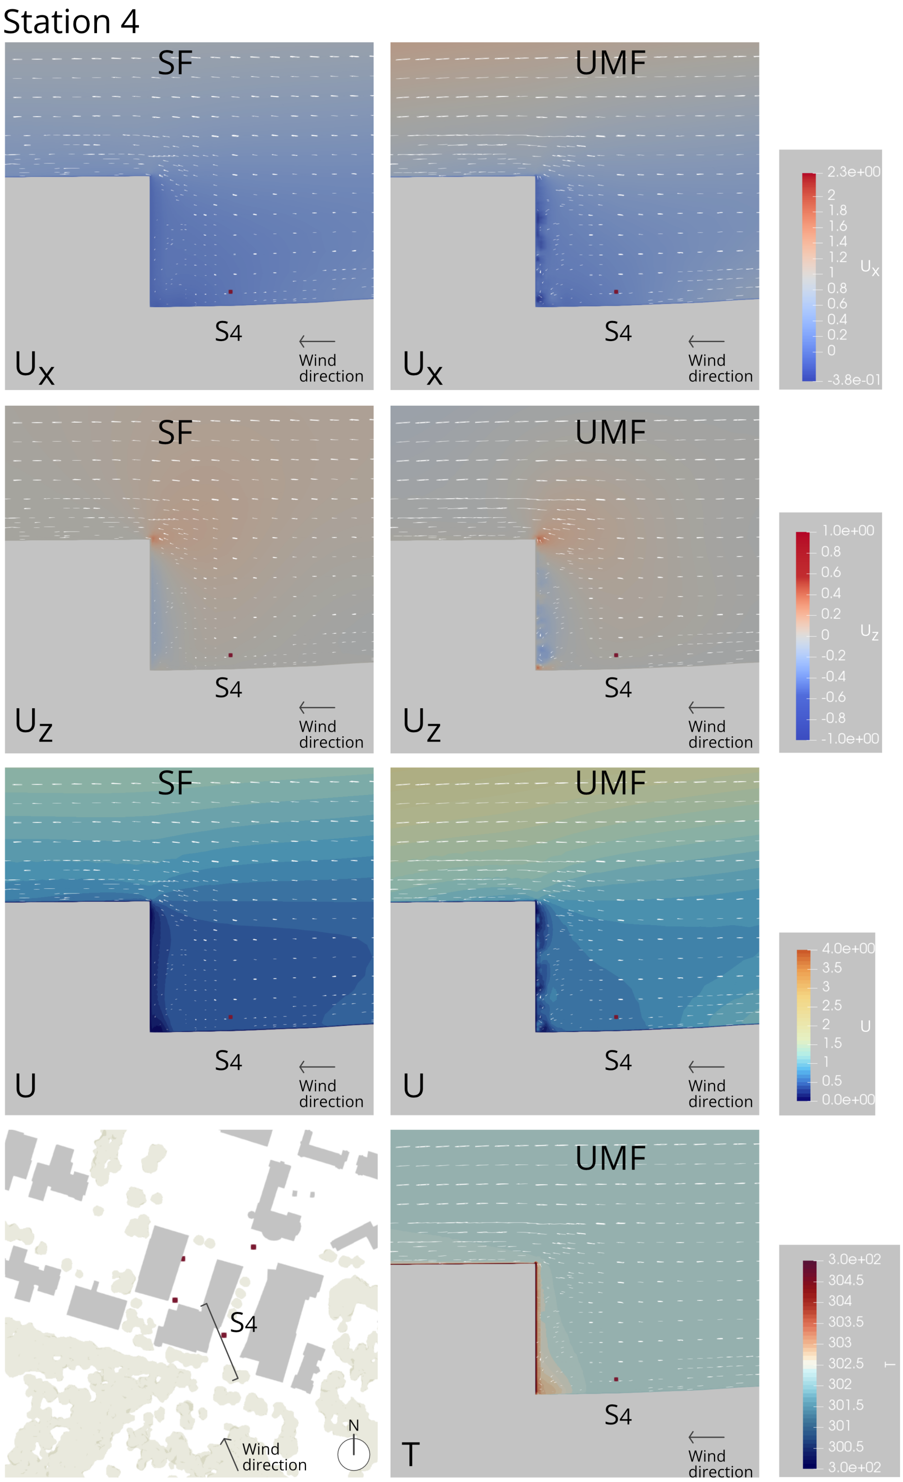

Figure 3.17.: Station 4: cross-sections of the horizontal velocity U_x, vertical velocity U_z, and velocity magnitude U for both solvers, together with air temperature T for uMF only, computed for the 13:00–14:00 time window on 27 August 2024.

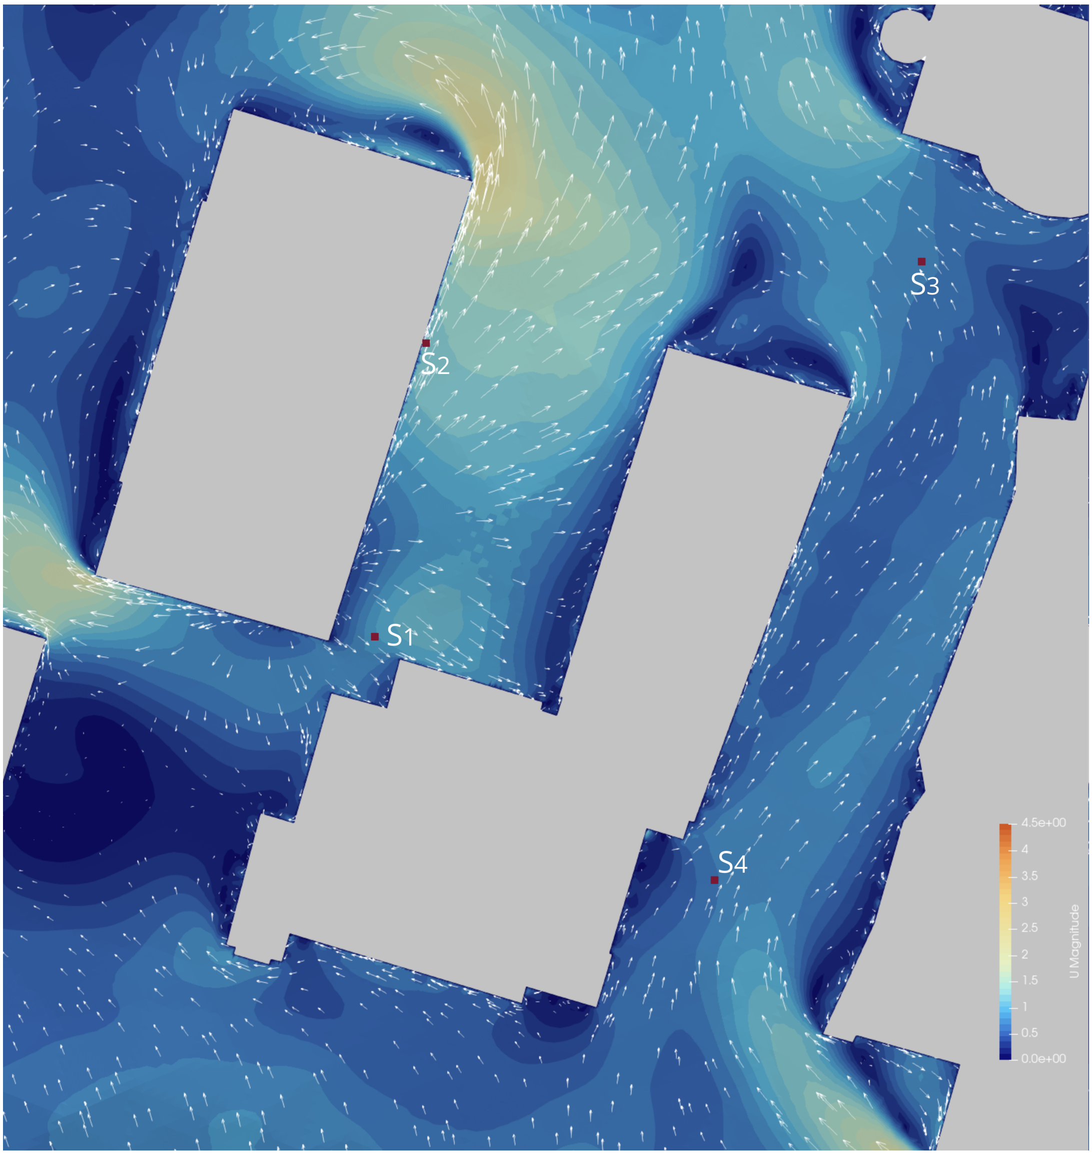

Figure 3.18.: Plan view of the horizontal wind speed U for uMF solver at 1.8 m above ground level

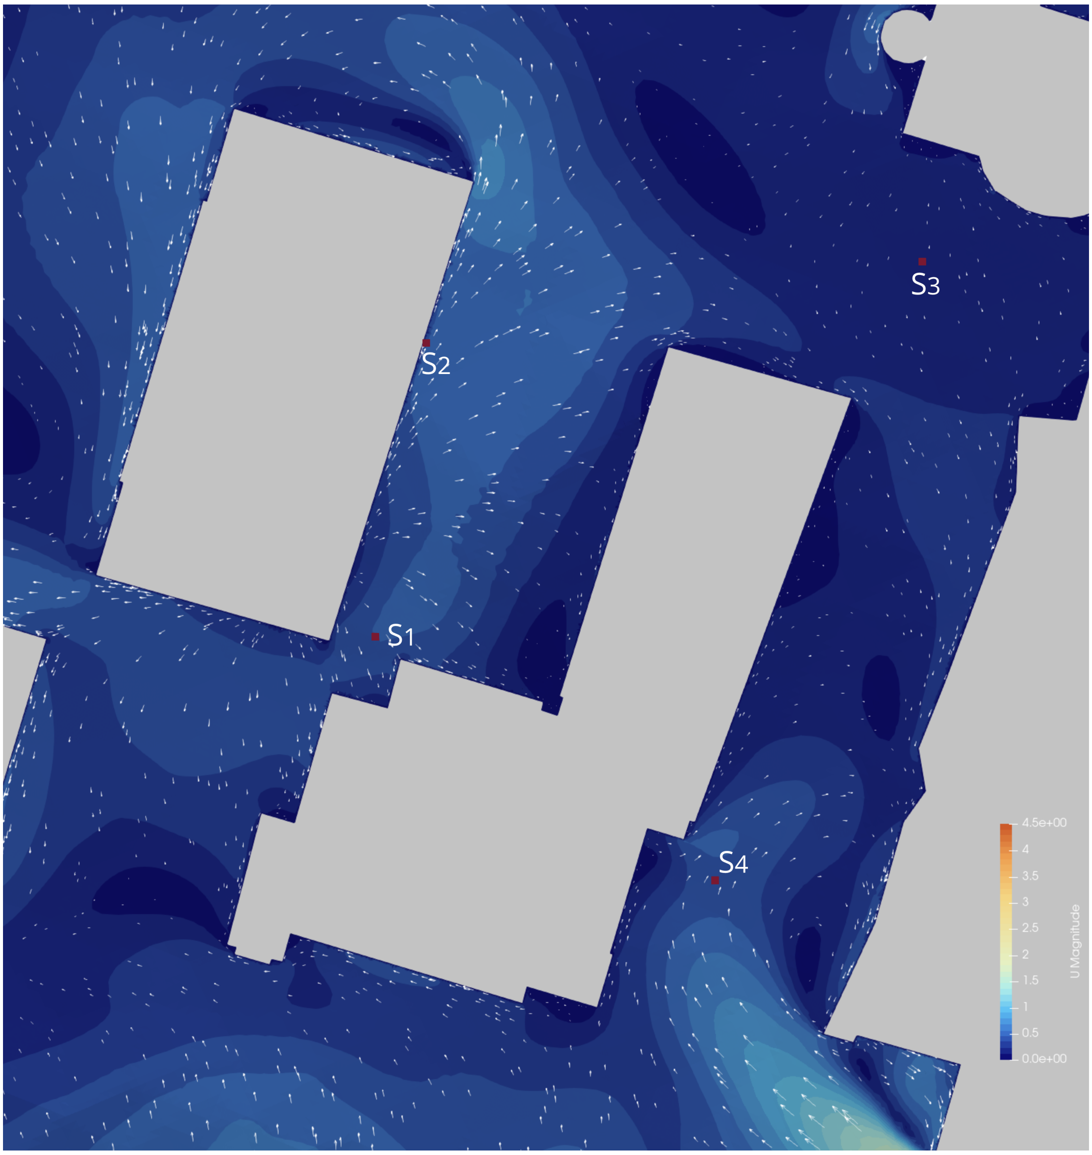

Figure 3.19.: Plan view of the horizontal wind speed U for SF solver at 1.8 m above ground level

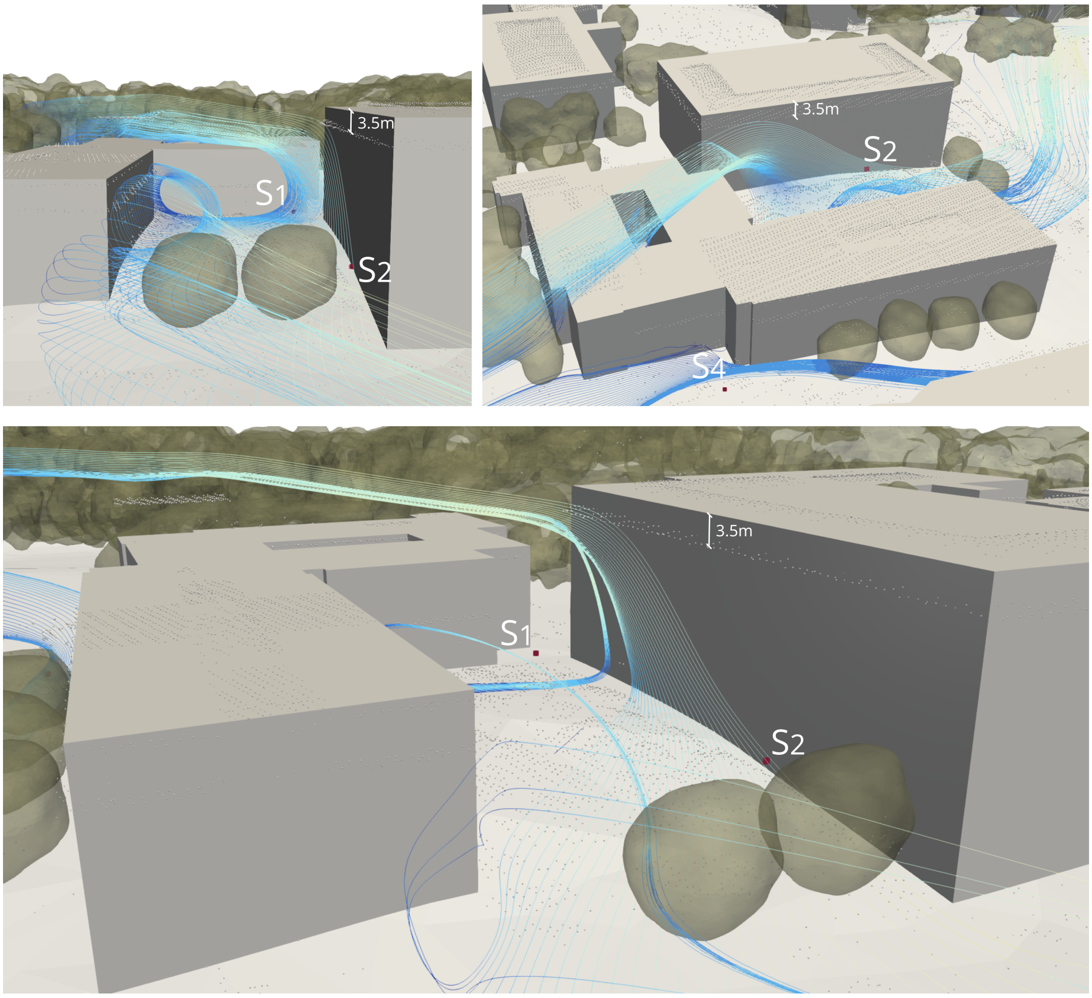

Figure 3.20.: Streamlines through Station 2 shown alongside the point cloud overlaid with the reconstructed building model. The LoD 1.2 reconstruction represents the original gabled roof as a flat roof located 3.5 m above the true eave height. This geometric simplification raises the effective windward facade and can locally alter the near-surface wind speed through an downwash effect.

3.6. Discussions and Conclusions

This study evaluated the feasibility and performance of the coupled HAM–RAD–AIR solver uMF on a realistic urban environment with complex terrain, applied to the Carnegie Mellon University campus during a heatwave hour on 27 August 2024, and compared it against a wind-only SF baseline validated against the MORICHI weather-station network. The discussion is organised around the three research questions posed in Section 3.2, followed by a synthesis of what the combined picture implies for practice and an explicit statement of which claims generalise beyond this case.

Feasibility on realistic terrain (RQ1) Applying uMF to realistic urban models with complex terrain is achievable, but the path from a reconstructed 3D model to a stable case is neither short nor automated. The same City4CFD output that ran directly in SF required case-specific edits before uMF would converge. Namely, removal of non-radiating faces, surface separation, local refinement at small features, manual correction of sliver terrain faces, and trial-and-error tuning of `faceAgglomerate` parameters. To the authors' knowledge, no published application of uMF reports a comparable workflow on non-flat terrain: existing validations and tutorial examples (Rahimi et al., 2025; Kubilay et al., 2017) use flat-ground configurations, which sidesteps the building–terrain interface issues that dominated the preprocessing effort here. This gap matters because Talwar and Yuan (2024) showed that natural terrain produces statistically significant deviations from flat-terrain predictions under buoyancy-driven flows, with average wind-speed differences up to 45% in certain conditions. If coupled solvers cannot reliably ingest realistic terrain, the conditions under which they offer the largest advantage over wind-only baselines could remain out of practical reach.

Computational cost (RQ2) On identical geometry and 20 cores, uMF took approximately 3.6× longer than SF (6 h 28 min versus 1 h 49 min) for a one-hour time step, with roughly three times as many inner iterations and tighter convergence tolerances on additional transported fields. This overhead has two sources that are worth distinguishing. The first is the additional physics itself: solving for humidity, modified pressure, and solid-region temperature and moisture, with the associated under-relaxation required for stability, raises the cost of each iteration and the number of iterations needed. The second is geometry-induced: both complex terrain and detailed buildings require finer cells near small features to remain stable, which inflates the total cell count beyond what the physics alone would dictate. In the present case, the 83 m of terrain variation within the domain forced local refinement that a flat-terrain case would avoid, contributing to the overhead reported here.

Despite this overhead, the cost remains within the same order of magnitude as SF. For a research workflow, 3.6× runtime is not, by itself, a barrier. The barriers identified in RQ1, manual geometric correction, trial-and-error pre-processing, consume far more time over the lifecycle of a study than the solver itself.

Accuracy under low-wind, thermally-active conditions (RQ3) In the regime studied here ($Ri_b \approx -0.85$, $U_{\text{ref}} \approx 1.6$ m/s), uMF improved wind-speed predictions over SF on every aggregate validation indicator: FB moved from $+0.65$ to -0.11, MG from 1.89 to 0.83, NMSE was reduced by 65%, VG by 59%, and FAC2 rose from 0.50 to 0.75. SF failed the Sabatino et al. (2011) acceptance bands on FB, MG, and VG; uMF passed all five.

This is consistent with the regime threshold reported by Mei and Yuan (2022), who identified $|Ri| > 0.2$ as the point beyond which buoyancy begins to shorten reattachment lengths behind isolated buildings, and below the $|Ri| > 2$ threshold at which canyon-scale vortex structures reorganise. The present case sits in the intermediate band where buoyancy is non-negligible but does not restructure the canyon flow. The systematic under-prediction by SF at all four stations is consistent with the absence of buoyancy-driven vertical mixing: the wind-only solver cannot generate the upward motion that warm surfaces drive in reality, and so under-estimates the near-surface wind speed. The vertical velocity component U_z is visibly stronger in the uMF cross-sections (Figures 3.14–3.17), most pronounced at Station 1, where SF performed worst and uMF reduced the error by 29 percentage points. This is the empirical indication that buoyancy is the missing physics in the SF predictions, and supports extending the regime reasoning of Mei and Yuan (2022) from wind-tunnel geometries to a realistic urban site using CFD.

The Station 2 anomaly, where SF matched the observation closely and uMF over-predicted by 139%, is revealing. The building immediately upwind of Station 2 was reconstructed at LoD 1.2, which converts its original gabled roof into a flat roof located approximately 3.5 m above the true eave (Figure 3.20). This raises the effective windward facade, which intercepts faster upper-level flow and amplifies downwash along the wall (Mittal et al., 2018; Zhao et al., 2021). In SF, this geometry-induced acceleration coincidentally compensates for the solver's general under-prediction tendency, producing apparent agreement with the observation. In uMF, the same downwash is augmented by buoyancy-driven motion at the warm wall, and the two effects compound into the observed over-prediction. The implication is that Station 2 is probably not evidence that uMF is less accurate than SF in general, it is evidence that LoD sensitivity propagates differently through the two solvers, and that apparent agreement under a wind-only model can mask compensating errors. However, whether this also extends the LoD-sensitivity findings of Patil and García-Sánchez (2025); García-Sánchez et al. (2021) and Hågbo et al. (2021) to the coupled-solver setting cannot be established from a single station. A sensitivity analysis comparing the original gabled-roof geometry against the LoD 1.2 reconstruction under both solvers would be needed to test this hypothesis.

Why uMF over-predicts on average Although uMF passes every acceptance band, the mean prediction lies above the observations at three of the four stations. The most plausible single contributor is the representation of vegetation as Darcy–Forchheimer porous zones, which captures momentum extraction but omits evapotranspiration and the associated latent-heat sink. The CMU campus borders Schenley Park to the southeast and sits in a well-vegetated residential surrounding; the omitted cooling and humidity flux is therefore unlikely to be a small correction. A canopy that cools and humidifies the near-surface air would reduce the buoyancy forcing reaching the campus from upwind, which would in turn reduce the vertical mixing currently driving the uMF over-prediction. Quantifying this requires running the uMF VEG module on the same case.

When is uMF worth the cost? The aggregate indicators show that uMF delivers a meaningful accuracy improvement in the studied regime, and the additional runtime is modest relative to the preprocessing investment required to set up the case. The practical decision rule that emerges is therefore not framed in terms of runtime. The decision depends instead on (i) whether the regime is one in which buoyancy matters, i. e. heatwave conditions with light wind, $|Ri_b| \gtrsim 0.2$ being a reasonable first criterion; and (ii) whether the analyst can absorb the case-setup overhead. Where both conditions hold, the evidence from this study supports adopting the coupled solver. Where the regime is dominated by mechanical forcing, or where the analysis must be repeated across many directions or many time windows in a planning workflow, the case for SF is, at the moment, stronger, given the challenges presented here. However, a more robust conclusion would require simulating additional hours of the day, preferably those with wind speeds above the anemometer accuracy threshold, and including vegetation through the uMF VEG module.

In conclusion and returning to the three questions posed at the outset: uMF can be applied to realistic urban geometries with complex terrain, but the preprocessing burden is substantial and largely manual; on identical geometry it costs roughly 3.6× the runtime of a wind-only baseline, which is not a barrier in itself; and in the low-wind, thermally-active regime examined here, it improves wind-speed agreement against street-level measurements across every aggregate indicator considered, with the strongest per-station support coming from Station 1, the only location whose observed mean exceeds the anemometer accuracy specification.

3.7. Recommendations

Several developments would improve the practical accessibility of uMF for urban applications.

The most pressing need is for the preparation utilities of the solver to become more robust to realistic geometries: algorithms such as `faceAgglomerate` currently struggle with the meshes that `snappyHexMesh` produces from complex building and terrain inputs. Making them robust to these inputs would remove the largest source of manual intervention in the current workflow. Complementing this, a better understanding of which specific mesh features induce solver instability, besides the ones identified in this study, would allow upstream pipelines to avoid them.

On the geometric-reconstruction side, automated pipelines with utilities tailored to uMF would minimize the current reliance on time-consuming trial-and-error. This could include, for example, building-footprint simplification to remove unnecessary details, terrain-mesh refinement, more robust refinement of building meshes, and diagnostic tools that flag problematic regions before a run.

On the solver side, circular domains would make simulations for varying wind directions more practical, though the boundary-condition formulations may need adaptation to support this.

Finally, for the broader evaluation of tools applied to UHI studies, extending the validation to longer time windows would bring more robust evidence when the wind-only approximation is acceptable. The methodology presented in this study can be applied systematically to that end, ideally including windows with stronger wind forcing above the anemometer's accuracy threshold and incorporating vegetation through the uMF VEG module.

4. Discussions

This closing chapter reflects on the connection between the work and the MSc Geomatics for the Built Environment, expands on the future-work directions outlined in the article, and discusses the broader challenges encountered during the project.

4.1. Connection to MSc Geomatics for the Built Environment

This thesis sits at the intersection of geomatics and CFD. The geometry that drives the simulations, namely terrain elevation, building footprints, and vegetation extents, originates from the sources that the Geomatics for the Built Environment programme takes as its primary objects of study: cadastral data, point clouds, and observational techniques. The simulation tools themselves, by contrast, originate in a community that, to a certain extent, treats geometry as a given input.

CFD for urban and environmental applications increasingly depends on combining heterogeneous datasets, each with its own coordinate system, level of detail, and definitions of accuracy and precision. Reprojection, generalisation, and dataset integration, to name a few, are core competencies developed in the Geomatics for the Built Environment programme. Much of the work reported here consisted of bridging that gap: translating real-world urban geometry into a form that an existing CFD chain can ingest, identifying the points at which the chain breaks down when the geometry is more topologically and geometrically complex than the toolchain's developers anticipated, and then mitigating those issues.

The motivation for catching such problems early is practical. Diagnosing a geometry-induced failure in a CFD chain requires constructing a full case, running it to the point of failure, reconstructing the partitioned output, and loading it into a visualisation tool, a cycle that takes hours. Each defect caught at the geometry stage rather than at the solver stage therefore represents a disproportionate saving in time.

CFD imposes specific requirements on a 3D reconstruction, and its tolerance for geometric defects, e.g. invalid geometry, is much lower than in other uses of 3D data. Consider built-up volumetry calculations: real-world data such as building footprints may contain façade details, see Figure 3.9, that are perfectly acceptable for computing a volume, but the air mesh that results from the same geometry after `snappyHexMesh` can, depending on the refinement level, produce non-uniform volumetric cells exhibiting non-orthogonality or skewness of faces, which in turn affect numerical results. In the worst case, the output contains topologically broken and invalid cells that block some algorithms entirely, as seen in Section 3.4. In the example of Figure 3.9, the issue could be mitigated by removing the offending detail, since for an urban-scale simulation, and given its location far from the area of interest, the omission would not affect the results. This is precisely the nature of most of the issues encountered during this thesis.

Because real-world datasets are rarely optimal for direct use and require pre-processing, achieving a realistic 3D model reconstruction, producing complex geometry, and building a pipeline that fits automatically into a simulation workflow is a substantial challenge. This is especially true for uMF, which relies on several different utilities compared to SF and is therefore built on specific, sometimes restrictive, requirements.

A further consideration is that CFD involves an inherent trade-off between the level of detail, which determines the number of cells in the domain, and runtime. Generalisation techniques, specifically simplification, selection, and smoothing, are standard preprocessing steps for reducing excessive, non-necessary detail in real-world data. When applied, they lower the level of domain cell refinement required, which in turn helps optimise the computational cost.

Finally, taking simulation results and analysing them against observational data is a skill built throughout the programme: understanding the accuracy and precision of measurements, and assessing predictions against measured values using statistically robust techniques.

Both fields, geomatics and CFD, therefore stand to gain from this exchange. The conceptual vocabulary required to recognise and remedy mesh defects is part of the standard training of geomatics, and framing it as upstream of CFD benefits both sides. Several of the recommendations that follow can be read in that light: they are not only software improvements, but specific places where a geomatics perspective on robustness to real-world geometric input would yield concrete benefits.

4.2. Suggestions for future work

4.2.1. Errors on the `faceAgglomerate` step

Among the recommendations for future work, the most pressing concerns the radiation pre-processing chain in OpenFOAM, and `faceAgglomerate` in particular. The utility takes the air-side boundary mesh produced by `snappyHexMesh` and greedily merges neighbouring fine faces into larger coarse faces (agglomerations), so that the subsequent view-factor computation in `viewFactorsGen` operates on the reduced count of coarse faces rather than on the full set of fine faces. The grouping is produced by face-normal alignment and edge connectivity: two faces become candidates for merging if they share a mesh edge, and the merge is favoured when their normals are close to parallel. In practice, however, the utility is not robust even on geometries that complete every previous step of the SF chain without warnings or numerical problems, and two distinct failure modes were encountered in this work.

The first failure is an internal hash-table lookup error within `faceAgglomerate`. The utility maintains internal bookkeeping of fine-face pairs that must remain consistent as the agglomeration proceeds. Under certain configurations of local face topology, this bookkeeping becomes inconsistent and the run aborts with a fatal error: `(i j) not found in table`. The problem appears particularly sensitive to the local topology produced near processor boundaries: the same case that failed in parallel decomposition succeeded when run in serial. Parallel decomposition is, however, effectively a requirement for uMF cases like the one studied here, given the computational cost associated with domains of this size. Because the focus of this thesis lay elsewhere, the failure was worked around rather than investigated in depth, but it warrants attention: a more defensive treatment of the affected lookup would likely eliminate the abort.

The second failure mode appears when the agglomeration map produced by `faceAgglomerate` is subsequently consumed by `viewFactorsGen`, which requires every coarse face to be a non-manifold surface. The agglomeration step can nevertheless produce coarse faces that fail this requirement. Such groupings are rejected only at the `viewFactorsGen` stage, where the run aborts. Figure 3.12 shows an example of a non-manifold grouping. The topological validity that `viewFactorsGen` requires of each coarse face should instead be enforced at the `faceAgglomerate` stage, with any rejected groupings retried there. The utility would also benefit from a diagnostic mode that, on failure, writes the offending coarse face and its constituent fine faces to disk for visual inspection. At present, identifying the failure location requires manual inspection of per-processor outputs, which is a significant barrier to debugging.

4.2.2. Improvements on the geometric reconstruction

On the geometric-reconstruction side, automated pipelines with utilities tailored to uMF would minimize the reliance on trial-and-error. This could include, for example, building-footprint generalisation tools such as simplification to remove unnecessary detail. Control over the terrain mesh is at present difficult, and both isotropic refinement and diagnostic tools that flag problematic faces (such as sliver faces) would help. A more robust refinement of building meshes should also avoid the creation of gaps such as those seen in Section 3.4. Related to this, a better understanding of which specific mesh features induce solver instability, beyond those identified in this study, would allow upstream pipelines to avoid them altogether. For solid-mesh construction, alternatives to `extrudeMesh` are worth exploring, for example construction from patches or via `snappyHexMesh`.

4.2.3. Improvements on the solver

On the solver side, circular domains would make simulations for varying wind directions more practical, though the boundary-condition formulations may need adaptation to support this. Particularly p_{rgh} is incompatible with having one patch for inlet, outlet and sides. During this work, it was attempted to define regions for these subdividing the cylinder sides into several patches (inlet, outlet and sides). However, this was also not numerically stable. A closer look would be beneficial to identify the cause of the issues and mitigate them.

4.2.4. Improvements on validation with more time windows

The final recommendation, extending the validation to longer time windows, depends partially on the previous challenges being resolved, or at least stands to benefit substantially from progress on them.

The first requirement is a domain that can adapt to inlet and outlet directions (for example a circular domain) and that supports an extent of 15H in every direction, so that multiple wind directions can be simulated, as is realistic for a run spanning several hours.

To this end, the use of a varying U per timestep is also necessary. This appears achievable through vector fields, specifically `Utarget`, which sets boundary conditions from a vector field that assigns U for every face on a patch. A preprocessing tool for constructing the `Utarget` files would be practical. The use of `Utarget` vector fields also imposes that the patches be named exactly `north`, `south`, `east`, and `west`, and an algorithm that adapts to user-defined patch names would be a useful addition. An alternative worth exploring is a pipeline that updates the inlet U value after each run and before the next `timeStep`, instead of using `Utarget`.

A larger domain (15H in every direction) also implies a greater chance of encountering problematic terrain and building faces of the kinds presented in this study. Overcoming the `faceAgglomerate` challenge, whether manually or through an improved version of the algorithm, therefore becomes essential, since this step is a hard requirement of the chain.

From a strategic perspective, given the anemometer accuracy of ±0.89 m/s reported in Martin et al. (2026), future validation windows should be selected to ensure that observed wind speeds clearly exceed this threshold, so that the instrument's absolute accuracy no longer dominates the station uncertainty.

4.3. General challenges

Beyond the specific issues named above, it is worth reflecting on the lack of documentation surrounding the uMF algorithm and the scarcity of examples in the literature that make use of it. This was a barrier both to the learning path required to use the tool and to avoid and mitigate bugs. As an intrinsically novel approach, and one applied here to complex geometry, this research required (many) careful iterations between constructing the geometry and testing it within the uMF pipeline.

Many tests were required, and each had to be tightly controlled, varying only one or two settings at a time, such as a single boundary condition, so that the cause of any change in behaviour could be isolated. This trial-and-error process was sufficiently time-consuming that it constrained the scope of the validation: simulating additional time windows, which would have strengthened the generality of the conclusions, was not feasible within the available timeframe.

Overcoming the issues presented here would broaden the accessibility of the solver, particularly for specialists from fields other than CFD and geomatics. At present, a successful uMF run on a realistic urban site requires familiarity with mesh diagnostics, the internal logic of `snappyHexMesh` and `faceAgglomerate`, and the geometric assumptions baked into each preprocessing utility. This combination of skills is uncommon, which restricts adoption to research groups that can afford time dedication. Removing the manual interventions documented in this work would lower the entry threshold to a level closer to that of established tools such as ENVI-met, opening this Open Source solver to urban planners, architects, and building-physics researchers whose questions the coupled HAM–RAD–AIR formulation is designed to support.

It is hoped that this thesis will contribute to broader adoption and refinement of the uMF tool.

A. Declaration of AI/LLM usage

Claude was used to assist with debugging OpenFOAM, although its help was very limited. Additionally, it was also used to break down difficult concepts. For grammar checking and clarity improvements in the text, DeepL and Claude were used. I confirm that I adhere to academic integrity and that AI/LLM was not used to generate falsified, fictional, or plagiarized content.

B. Reproducibility self-assessment

All data and code related to this thesis are openly accessible at `https://github.com/CarmemAires/Thesis_uMF-SF_wind-comparison.git` and at 4TUResearchData under the DOI: 10.4121/fbdbda4e-e4e5-46eb-9521-505b6a15889f. Complete instructions and workflow are provided in the project README. I rate the reproducibility of this thesis as **High** according to the provided scale.

Table B.1.: Reproducibility criteria and their grades.

Criteria	**Grade**	**Justification**
Input data	3	Case folders, including geometry and pipeline, available at the 4TU repository and github
	3	MORICHI dataset publicly available (Martin et al., 2026)
	2	Building footprints and LiDAR point cloud publicly available (PASDA, 2015, 2026)
Preprocessing	3	2D geometry, point cloud, and City4CFD configuration file available at the 4TU repository; preprocessing steps documented in this report
Methods	3	Methods documented in detail in this report and published in TU Delft's repository
Computational environment	3	City4CFD (open source)
	3	OpenFOAM (open source)
	3	Linux (open source)
	3	PyMeshLab (open source)
	0	Rhino (proprietary, requires a license)
Results	2	Results are not available, but they can be reproduced from the case folders and geometry provided

Bibliography

Alliez P, Cohen-Steiner D, Hemmer M, Portaneri C, and Rouxel-Labbé M (2026). 3D alpha wrapping. In *CGAL User and Reference Manual*. CGAL Editorial Board, 6.1.1 edition.

Alting N (2025). From point clouds to porous crowns: A scalable approach for cfd-ready urban tree reconstruction. `https://github.com/NoahAlting/CFTree`. Accessed: February 2026.

American Meteorological Society (2026). Glossary of Meteorology. `https://glossary.ametsoc.org/wiki/microscale/`. Accessed: 2026-05-15.

Anderson JD (1995). *Computational fluid dynamics*. McGraw-Hill series in aeronautical and aerospace engineering. McGraw-Hill, New York [u.a.], internat. ed., [repr.] edition. ISBN 0070016852.

ANSYS, Inc (n.d.a). Ansys fluent. `https://www.ansys.com/`. Accessed: May 2026.

ANSYS, Inc (n.d.b). What is Computational Fluid Dynamics (CFD)? `https://www.ansys.com/simulation-topics/what-is-computational-fluid-dynamics`. Accessed: May 2026.

Blocken B (2015). Computational fluid dynamics for urban physics: Importance, scales, possibilities, limitations and ten tips and tricks towards accurate and reliable simulations. *Building and Environment*, 91:219–245. ISSN 0360-1323. doi:https://doi.org/10.1016/j.buildenv.2015.02.015. Fifty Year Anniversary for Building and Environment.

Brozovsky J, Simonsen A, and Gaitani N (2021). Validation of a cfd model for the evaluation of urban microclimate at high latitudes: A case study in trondheim, norway. *Building and Environment*, 205:108175. ISSN 0360-1323. doi:https://doi.org/10.1016/j.buildenv.2021.108175.

Celik IB, Ghia U, Roache PJ, Freitas CJ, Coleman H, and Raad PE (2008). Procedure for estimation and reporting of uncertainty due to discretization in cfd applications. *Journal of Fluids Engineering*, 130(7):078001. ISSN 0098-2202. doi:10.1115/1.2960953.

Chang JC and Hanna SR (2004). Air quality model performance evaluation. *Meteorology and Atmospheric Physics*, 87(1-3). ISSN 1436-5065. doi:10.1007/s00703-003-0070-7.

CloudCompare (n.d.). CANUPO plugin documentation. `https://www.cloudcompare.org/doc/wiki/index.php/CANUPO_(plugin)`. Accessed: January 2026.

Copernicus Climate Change Service (C3S) (2025). Global climate highlights 2024. `https://climate.copernicus.eu/global-climate-highlights-2024`. European Centre for Medium-Range Weather Forecasts (ECMWF). Accessed: May 2026.

ENVI-met GmbH (n.d.). ENVI-met. `https://envi-met.com/`. Accessed: May 2026.

Ferziger JH, Perić M, and Street RL (2020). *Computational Methods for Fluid Dynamics*. Springer, Cham, Switzerland, 4th edition. ISBN 9783319996912.

Franke J, Hellsten A, Schlünzen H, and Carissimo B (2007). Best practice guideline for the CFD simulation of flows in the urban environment. COST Action 732, COST Office, Brussels.

García-Sánchez C, Vitalis S, Paden I, and Stoter J (2021). The impact of level of detail in 3d city models for cfd-based wind flow simulations. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVI-4/W4-2021:67–72. ISSN 2194-9034. doi:10.5194/isprs-archives-xlvi-4-w4-2021-67-2021.

Hågbo TO, Giljarhus KET, and Hjertager BH (2021). Influence of geometry acquisition method on pedestrian wind simulations. *Journal of Wind Engineering and Industrial Aerodynamics*, 215:104665. ISSN 0167-6105. doi:https://doi.org/10.1016/j.jweia.2021.104665.

Iowa Environmental Mesonet (2026). ASOS-AWOS-METAR data download. `https://mesonet.agron.iastate.edu/request/download.phtml?network=RAOB`. Accessed: February 2026.

IPCC I (2023). *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. doi:10.59327/ipcc/ar6-9789291691647.

Janssen H, Blocken B, and Carmeliet J (2007). Conservative modelling of the moisture and heat transfer in building components under atmospheric excitation. *International Journal of Heat and Mass Transfer*, 50(5):1128–1140. ISSN 0017-9310. doi:https://doi.org/10.1016/j.ijheatmasstransfer.2006.06.048.

Katul G and Albertson JD (1998). An investigation of higher-order closure models for a forested canopy. *Boundary-Layer Meteorology*, 89(1):47–74. ISSN 1573-1472. doi:10.1023/a:1001509106381.

Kubilay A, Allegrini J, Strebel D, Zhao Y, Derome D, and Carmeliet J (2020). Advancement in urban climate modelling at local scale urban heat island mitigation and building cooling demand. *Atmosphere*, 11(12):1313. ISSN 2073-4433. doi:10.3390/atmos11121313.

Kubilay A, Ferrari A, and Manickathan L (2017). Urban microclimate foam. Accessed: December 2025.

Martin M, Garcia-Sanchez C, Stoter J, and Berges M (2026). Morichi: a dataset to study urban overheating during extreme heat in a hot-summer humid continental climate. *Scientific Data*, 13(1). ISSN 2052-4463. doi:10.1038/s41597-026-06763-w.

Massman W (1987). A comparative study of some mathematical models of the mean wind structure and aerodynamic drag of plant canopies. *Boundary-Layer Meteorology*, 40(1-2):179–197. ISSN 1573-1472. doi:10.1007/bf00140075.

Mei SJ and Yuan C (2022). Urban buoyancy-driven air flow and modelling method: A critical review. *Building and Environment*, 210:108708. ISSN 0360-1323. doi:https://doi.org/10.1016/j.buildenv.2021.108708.

Mirzaei PA and Haghighat F (2010). Approaches to study urban heat island – abilities and limitations. *Building and Environment*, 45(10):2192–2201. ISSN 0360-1323. doi:https://doi.org/10.1016/j.buildenv.2010.04.001.

Mittal H, Sharma A, and Gairola A (2018). A review on the study of urban wind at the pedestrian level around buildings. *Journal of Building Engineering*, 18:154–163. ISSN 2352-7102. doi:10.1016/j.jobe.2018.03.006.

Moediartianto A, Montazeri H, and Blocken B (2026). On the use of uav-thermal imaging for cfd validation of urban thermal microclimate. *Sustainable Cities and Society*, 136:106968. ISSN 2210-6707. doi:https://doi.org/10.1016/j.scs.2025.106968.

Moukalled F, Mangani L, and Darwish M (2016). *The Finite Volume Method in Computational Fluid Dynamics: An Advanced Introduction with OpenFOAM® and Matlab*. Springer International Publishing. ISBN 9783319168746. doi:10.1007/978-3-319-16874-6.

Oke TR (2017). *Urban climates*. Cambridge University Press, Cambridge. ISBN 9781139016476.

OpenFOAM Foundation (n.d.). OpenFOAM v7. `https://cpp.openfoam.org/v7/`. Accessed: December 2025.

OpenFOAM Wiki (n.d.). SimpleFoam. `https://openfoamwiki.net/index.php/SimpleFoam`. Accessed: December 2025.

Pađen I, García-Sánchez C, and Ledoux H (2022). Towards automatic reconstruction of 3D city models tailored for urban flow simulations. *Frontiers in Built Environment*, 8. ISSN 22973362. doi:10.3389/fbuil.2022.899332.

Pađen I, Peters R, García-Sánchez C, and Ledoux H (2024). Automatic high-detailed building reconstruction workflow for urban microscale simulations. *Building and Environment*, 265:111978. ISSN 0360-1323. doi:https://doi.org/10.1016/j.buildenv.2024.111978.

PASDA (2015). Allegheny county lidar classified LAS point cloud. `https://www.pasda.psu.edu/download/alleghenycountyimagery2015/`. Accessed: February 2026.

PASDA (2017). Allegheny county imagery 2017 - 3 inch pixels. `https://www.pasda.psu.edu/download/alleghenycountyimagery2017/OrthoMosaicTiles/`. Accessed: February 2026.

PASDA (2026). Allegheny county building footprints. `https://www.pasda.psu.edu/uci/DataSummary.aspx?dataset=1195`. Accessed: February 2026.

Patil A and García-Sánchez C (2025). Quantifying the impact of urban geometric detail for urban air mobility risk forecasting. *Sustainable Cities and Society*, 132:106750. ISSN 2210-6707. doi:https://doi.org/10.1016/j.scs.2025.106750.

PDAL Contributors (2024). PDAL: Point Data Abstraction Library. `https://pdal.org/en/2.9.3/`. Version 2.9.3. Accessed: January 2026.

Rahimi S, Alvarez M, Berardi U, Stone B, and Kastner P (2025). Coupled urban physics in microclimate modeling: Validating and enhancing simulation tools. *Building and Environment*, 285:113637. ISSN 0360-1323. doi:10.1016/j.buildenv.2025.113637.

Robert McNeel & Associates (n.d.). Rhinoceros 3D. `https://www.rhino3d.com`. Accessed: May 2026.

Roy TB, Middey A, and Krupadam R (2025). Unveiling the microclimate: A comprehensive review of tools, techniques, and future directions for sustainable cities. *Building and Environment*, 274:112726. ISSN 0360-1323. doi:https://doi.org/10.1016/j.buildenv.2025.112726.

Sabatino SD, Buccolieri R, Olesen HR, Ketzel M, Berkowicz R, Franke J, Schatzmann M, Schlunzen KH, Leitl B, Britter R, Borrego C, Costa AM, Castelli ST, Reisin TG, Hellsten A, Saloranta J, Moussiopoulos N, Barmpas F, Brzozowski K, Goricsan I, Balczo M, Bartzis JG, Efthimiou G, Santiago JL, Martilli A, Piringer M, Stanzer KB, Hirtl M, Baklanov AA, Nuterman RB, and Starchenko AV (2011). Cost 732 in practice: the must model evaluation exercise. *International Journal of Environment and Pollution*, 44(1/2/3/4):403. ISSN 1741-5101. doi:10.1504/ijep.2011.038442.

Schneider M, Tötzer T, Bügelmayer-Blaschek M, and Berg R (2023). Pitfalls and potentials of microclimate simulations in urban planning. *Journal of Urban Planning and Development*, 149(4):04023048. doi:10.1061/JUPDDM.UPENG-4504.

Stull RB (1988). *An Introduction to Boundary Layer Meteorology*. Kluwer Academic Publishers, Dordrecht.

Talwar T and Yuan C (2024). Impact of natural urban terrain on the pedestrian wind environment in neighborhoods: A cfd study with both wind and buoyancy-driven scenarios. *Building and Environment*, 261:111746. ISSN 0360-1323. doi:https://doi.org/10.1016/j.buildenv.2024.111746.

Tominaga Y, Mochida A, Yoshie R, Kataoka H, Nozu T, Yoshikawa M, and Shirasawa T (2008). AIJ guidelines for practical applications of CFD to pedestrian wind environment around buildings. *Journal of Wind Engineering and Industrial Aerodynamics*, 96(10–11):1749–1761. doi:10.1016/j.jweia.2008.02.058.

Toparlar Y, Blocken B, Maiheu B, and van Heijst G (2017). A review on the cfd analysis of urban microclimate. *Renewable and Sustainable Energy Reviews*, 80:1613–1640. ISSN 1364-0321. doi:10.1016/j.rser.2017.05.248.

Toparlar Y, Blocken B, Vos P, van Heijst G, Janssen W, van Hooff T, Montazeri H, and Timmermans H (2015). Cfd simulation and validation of urban microclimate: A case study for bergpolder zuid, rotterdam. *Building and Environment*, 83:79–90. ISSN 0360-1323. doi:https://doi.org/10.1016/j.buildenv.2014.08.004. Special Issue: Climate adaptation in cities.

UN-Habitat UNHSP (2024). *World Cities Report*. World cities report. UN-Habitat, Nairobi, Kenya. ISBN 978-92-1-132955-1.

Versteeg HK (2007). *An introduction to computational fluid dynamics*. Pearson Education Limited, 2 edition. ISBN 9780131274983.

Weather Spark (2026). The weather year round anywhere on earth. `https://weatherspark.com/`. Accessed: 2026-03-01.

Wieringa J (1992). Updating the Davenport roughness classification. *Journal of Wind Engineering and Industrial Aerodynamics*, 41(1–3):357–368. doi:10.1016/0167-6105(92)90434-C.

World Health Organization (2024). Climate change: Heat and health. `https://www.who.int/news-room/fact-sheets/detail/climate-change-heat-and-health`. Accessed: 2024-12-13.

Yang S, Wang LL, Stathopoulos T, and Marey AM (2023). Urban microclimate and its impact on built environment – a review. *Building and Environment*, 238:110334. ISSN 0360-1323. doi:https://doi.org/10.1016/j.buildenv.2023.110334.

Zhao Y, Li H, Kubilay A, and Carmeliet J (2021). Buoyancy effects on the flows around flat and steep street canyons in simplified urban settings subject to a neutral approaching boundary layer: Wind tunnel piv measurements. *Science of The Total Environment*, 797:149067. ISSN 0048-9697. doi:https://doi.org/10.1016/j.scitotenv.2021.149067.

Colophon

This document was typeset using LaTeX, using a modified the KOMA-Script class `scrbook` available at `https://github.com/tudelft3d/msc_geomatics_thesis_template`. The main font is Palatino.