UNIVERSITÀ DEGLI STUDI
DI MODENA E REGGIO EMILIA

Dipartimento di Ingegneria "Enzo Ferrari"

Corso di Laurea Magistrale in Ingegneria dl Veicolo

Ciclo XXXVIII

Development and Validation of Advanced Physical Models for Urban Wind Microclimate Assessment

Master's Thesis

Candidate: Samuele **Patelli**

Supervisor: Prof. **Andrea Cimarelli**

Academic Year 2025/2026

ABSTRACT

The quality of the pedestrian outdoor environment is increasingly recognized as a public health and planning priority. Wind-induced discomfort, outdoor thermal stress, and urban heat island effects are the subject of a growing body of regulatory guidance that requires developers and urban planners to demonstrate compliance with quantitative comfort criteria at pedestrian level. Meeting these requirements within the timescales of architectural practice demands simulation tools that are physically rigorous, geometrically faithful, and computationally efficient.

This thesis presents a series of validated contributions to ArchiWind, a CFD-based platform for pedestrian wind comfort assessment built on the open-source solver OpenFOAM and developed at NablaFlow. The work addresses three main areas.

First, the platform is validated against the Architectural Institute of Japan (AIJ) Benchmark Case E — a canonical dense urban environment in the Niigata city centre — demonstrating good agreement with wind tunnel measurements across 80 pedestrian-level probe points and 16 inflow directions (RMSE = 0.174, Pearson r = 0.769 for the after-construction scenario using the Realizable k-ε turbulence model on the Detailed mesh). The platform correctly resolves the spatial pattern of the construction effect $\Delta U/U_{ref}$, confirming its suitability for before/after planning impact assessments. Sensitivity studies show that the Draft mesh quality setting yields essentially the same global accuracy as the Detailed setting (ΔRMSE = −0.002), and that the k-ω SST turbulence model performs comparably to Realizable k-ε.

Second, two methodological enhancements to the CFD workflow are introduced. A physics-based porous media model for tree canopy representation replaces the previous static default height assignment with a Voxel-Based Canopy Profiling (VCP) algorithm that extracts canopy height directly from the input STL geometry and introduces a geometry-scaled numerical stability limiter ($K_{stability}$), eliminating the floating-point exceptions and divergence episodes previously associated with high-density vegetation zones. A two-stage physics-driven convergence criterion is developed and validated across eight real-world test cases, replacing traditional global residual monitoring with a sequential algorithm that monitors the stagnation of pedestrian-level physical fields at 1.5

m AGL; runtime reductions of 27–62 % are achieved with comfort map differences below 1.76 % and probe-level errors below 0.1 %.

Third, the physical scope of the platform is extended. Thermal buoyancy is implemented via the Boussinesq approximation through a transition from OpenFOAM's `simpleFoam` to `buoyantBoussinesqSimpleFoam`; solver equivalence with the isothermal baseline is confirmed with RMSE = 0.00024 m s^{-1} and Pearson r = 1.000, and active buoyancy produces physically consistent tilted plumes, surface-driven updraughts, and horseshoe vortex patterns. A complete outdoor thermal comfort assessment module is developed, computing Mean Radiant Temperature (MRT) and the Universal Thermal Climate Index (UTCI) from first principles using the CFD wind atlas, hourly Meteoblue climate records, and three-dimensional ray-casting for solar shading and sky view factor computation. Applied to a representative London urban site, the model produces seasonal compliance maps in accordance with the five-tier GLA London Plan classification, with mean summer MRT values of 18–25 °C and winter values of 2–5.5 °C.

Taken together, the contributions advance ArchiWind from a wind-only assessment tool to a fully integrated urban microclimate platform capable of meeting current and emerging regulatory requirements for pedestrian comfort analysis in architectural and urban planning practice.

SOMMARIO

LIST OF FIGURES

LIST OF TABLES

1 INTRODUCTION

1.1 BACKGROUND AND MOTIVATION

The quality of the outdoor environment at pedestrian level is a critical determinant of the livability and usability of urban public spaces. Wind speed and direction, solar radiation, air temperature, and humidity collectively shape the thermal and mechanical comfort experienced by pedestrians in streets, squares, and open spaces. None of these variables acts in isolation: strong wind amplified at a building corner can be tolerable on a warm summer afternoon but dangerously uncomfortable in winter, while elevated solar radiation that provides welcome warmth in spring can produce thermal stress conditions in summer that suppress outdoor activity entirely. The interaction between these variables, mediated by the geometry of the built environment, defines the concept of outdoor thermal comfort — and its failure represents a direct cost to urban livability, public health, and economic vitality of city spaces.

As European cities densify, the aerodynamic and thermodynamic impacts of new construction on the surrounding pedestrian environment become progressively harder to predict without quantitative simulation. Tall buildings deflect and accelerate the ambient wind at street level, producing corner jets, downwash zones, and recirculation regions that can exceed thresholds for pedestrian discomfort or safety. At the same time, the urban heat island effect — driven by the replacement of vegetated surfaces with impervious materials, the thermal mass of dense building stock, and reduced sky view factors — elevates surface temperatures and compounds thermal stress during summer heat events. The European heat waves of 2003 and 2019 resulted in tens of thousands of excess deaths across the continent, giving concrete human weight to what had previously been treated as an environmental amenity question.

The regulatory response to these challenges is accelerating. In the United Kingdom, the Greater London Authority (GLA) Wind Microclimate Guidelines [1] require applicants for major planning consents to demonstrate compliance with a five-tier annual classification of outdoor comfort based on the Universal Thermal Climate Index (UTCI), covering not only mechanical wind speed but the full thermodynamic environment experienced by pedestrians. Equivalent requirements are codified in Germany under VDI 3787 [2] and are emerging across the European Union through the Urban Adaptation

Strategy of the EU Green Deal. The practical consequence is that any development team working on a major urban project in northern Europe now requires simulation tools capable of delivering spatially resolved, standards-compliant pedestrian comfort assessments as a routine part of the design workflow.

Computational Fluid Dynamics (CFD) provides the most direct route to accurate pedestrian-level wind speed predictions in complex urban geometries. Steady-state Reynolds-Averaged Navier–Stokes (RANS) methods offer an attractive balance between physical fidelity and computational cost: they resolve the time-mean flow field on realistic three-dimensional building geometries at a fraction of the cost of wind tunnel physical modelling and are well-suited to the multi-direction sweeps required by comfort standards [3, 4]. The present thesis is situated within this context: it develops, validates, and extends ArchiWind — a CFD-based pedestrian comfort platform developed at NablaFlow — to address the full scope of current regulatory requirements.

1.2 THE ARCHIWIND PLATFORM

ArchiWind is a CFD-based simulation platform developed at NablaFlow, a company specializing in computational fluid dynamics for the built environment. The platform is designed to make rigorous pedestrian wind comfort assessments accessible to architectural practices without requiring in-house CFD expertise and was already deployed in production use for UK planning applications at the outset of this work.

The technical architecture rests on three layers. The computational backend uses OpenFOAM configured for steady-state RANS with the `SIMPLE` pressure-velocity coupling algorithm. The geometry and meshing layer use SnappyHexMesh to generate body-fitted hexahedral meshes on a cylindrical domain topology — following the approach of Hågbo and Giljarhus [5] — which allows the same mesh to be reused for all wind directions by rotating the inlet boundary conditions, rather than rebuilding the geometry for each direction. The orchestration layer is a Python pipeline that manages geometry ingestion (STL and OBJ formats), domain generation, boundary condition setup, cloud-based solver execution, and post-processing. Boundary conditions follow standard analytical atmospheric boundary layer profiles, with roughness length z_0 derived automatically from upstream terrain data for each wind direction. Post-processing extracts the pedestrian-level wind speed field at 1.5 m AGL and exports it as normalized greyscale PNG maps — one per wind direction — forming the wind atlas used for comfort classification.

The work described in this thesis extended ArchiWind along four directions. First, while validation studies had been conducted in earlier stages of the platform's development, these required updating to reflect the current solver configuration and to establish a rigorous quantitative baseline against a canonical international benchmark. Second, the vegetation model — based on a porous media approach — used a static default canopy height regardless of the input geometry; replacing this with a dynamic, geometry-aware formulation and a numerical stability limiter was identified as a priority for physical accuracy and solver robustness. Third, solver termination relied on a fixed maximum iteration count with no mechanism to detect convergence of the pedestrian-level physical fields of interest, leaving computational resources routinely over- or under-consumed. Fourth, the platform addressed only mechanical wind comfort and produced no thermal comfort output; extending it to compute MRT and UTCI — and separately, to account for thermal buoyancy effects in the CFD solver — was necessary to meet the full scope of current regulatory requirements such as the GLA London Plan guidelines [1].

1.3 RESEARCH OBJECTIVES

The work presented in this thesis pursues three main objectives, each directly addressing one of the identified gaps in the ArchiWind capability.

Objective 1: Benchmark Validation. To validate the ArchiWind RANS CFD workflow against the AIJ Benchmark Case E, a canonical dense urban environment in Niigata with 80 pedestrian-level measurement points, 16 wind directions, and paired before/after construction scenarios. Specific sub-objectives are: (a) to quantify global accuracy metrics (RMSE, bias, Pearson r) for both construction scenarios; (b) to evaluate the platform's ability to detect the construction effect $\Delta U/U_{ref}$; and (c) to assess sensitivity to mesh quality (Detailed vs. Draft) and turbulence model (Realizable k-ε vs. k-ω SST).

Objective 2: CFD Methodology Enhancements. To improve the CFD workflow through two developments: (a) a physics-based, geometry-aware porous media model for tree canopy representation incorporating a numerically stable leaf area density formulation; and (b) a two-stage physics-driven convergence criterion that replaces global residual monitoring with physical field stagnation detection at pedestrian level, validated across eight real-world test cases.

Objective 3: Advanced Physics and Thermal Comfort. To extend the platform through: (a) implementation of thermal buoyancy via the Boussinesq approximation and numerical validation against the isothermal baseline; and (b) development of a complete outdoor thermal comfort module computing MRT and UTCI from first principles and producing GLA London Plan compliant seasonal compliance maps.

1.4 THESIS STRUCTURE

Chapter 2 presents the validation of ArchiWind against the AIJ Case E benchmark, covering the case description, CFD setup, global and per-direction accuracy metrics, construction effect detection, and sensitivity studies on mesh quality and turbulence model. Chapter 3 describes the two CFD enhancements: the dynamic porous media vegetation model with the Voxel-Based Canopy Profiling algorithm and $K_{\mathrm{stability}}$ limiter (Section 3.1), and the two-stage physics-driven convergence criterion with validation across eight test cases (Section 3.2). Chapter 4 presents the Boussinesq buoyancy implementation, its three-step validation protocol, and an analysis of the physical flow modifications induced by active thermal buoyancy. Chapter 5 describes the thermal comfort assessment module — climate data processing, sky view factor computation, radiation model, UTCI pipeline, and results for a representative London site. Chapter 6 draws conclusions, acknowledges limitations, and identifies directions for future work.

2 VALIDATION AGAINST THE AIJ CASE E BENCHMARK

2.1 INTRODUCTION

Validation against well-established benchmark datasets is a prerequisite for the responsible deployment of any CFD platform in engineering practice. Without quantitative evidence that a solver correctly reproduces known physical behaviour, the numerical results it produces cannot be trusted as a basis for planning decisions. For urban pedestrian wind assessments, the benchmark programme established by the Architectural Institute of Japan (AIJ) provides the most comprehensive publicly available reference dataset: a suite of wind-tunnel test cases covering isolated buildings, building arrays, and realistic urban environments, accompanied by detailed measurement protocols and recommended CFD best-practice guidelines [4, 6].

The AIJ programme has become the de facto international standard for validating CFD tools in the pedestrian wind comfort domain. Over the past two decades, dozens of research groups worldwide have published their results against the AIJ cases, creating a well-populated accuracy envelope against which new tools can be positioned. Blocken [7] reviews this literature and reports that properly configured steady-state RANS solvers consistently achieve RMSE values of 0.15–0.25 on normalized wind speed U/U_{ref} for complex urban geometries, with systematic negative bias arising from the known tendency of RANS to underpredict wake recovery length. This range serves as the accuracy target against which ArchiWind is assessed in this chapter.

Case E of the AIJ benchmark suite was selected as the reference test case. It represents a 500 × 500 m area of the Niigata city centre with a dense, irregular urban morphology and features 80 pedestrian-level measurement points, 16 inflow directions, and a paired before/after construction dataset — making it the most demanding and most widely cited case for dense urban environments. It directly tests the capabilities most relevant to ArchiWind's production use case: prediction of pedestrian-level wind speeds in complex irregular building arrangements, and assessment of the wind impact of new construction.

2.2 CASE DESCRIPTION AND CFD SETUP

2.2.1 NIIGATA URBAN ENVIROMENT

Case E of the AIJ benchmark represents a 500 × 500 m area of the Niigata city centre, modelled at a geometric scale of 1:250 in the wind tunnel. The urban fabric consists of a dense, irregular arrangement of low- to mid-rise buildings with plan shapes typical of Japanese city centres: compact rectangular blocks of varying heights, narrow street canyons, and occasional taller landmark structures. Measurements were conducted at 80 pedestrian-level probe points located at a full-scale equivalent height of z = 1.5 m above ground level, for 16 inflow directions at 22.5° intervals (N, NNE, NE, ENE, E, ESE, SE, SSE, S, SSW, SW, WSW, W, WNW, NW, NNW). The reference velocity is U_{ref} = 3.928 m s^{-1}, defined at z = 15.9 m (height of the Niigata meteorological station).

The dataset is provided in two versions: the Before-Construction configuration represents the existing urban fabric, and the After-Construction configuration introduces a new building complex at the centre of the domain. Comparing the two configurations provides a direct assessment of ArchiWind's ability to resolve the wind impact of new construction — the primary use case in pedestrian wind comfort practice.

FIGURE 2.1 *Case E geometry with the 80 pedestrian-level probe point locations (full scale). Points are located at z = 1.5 m AGL*

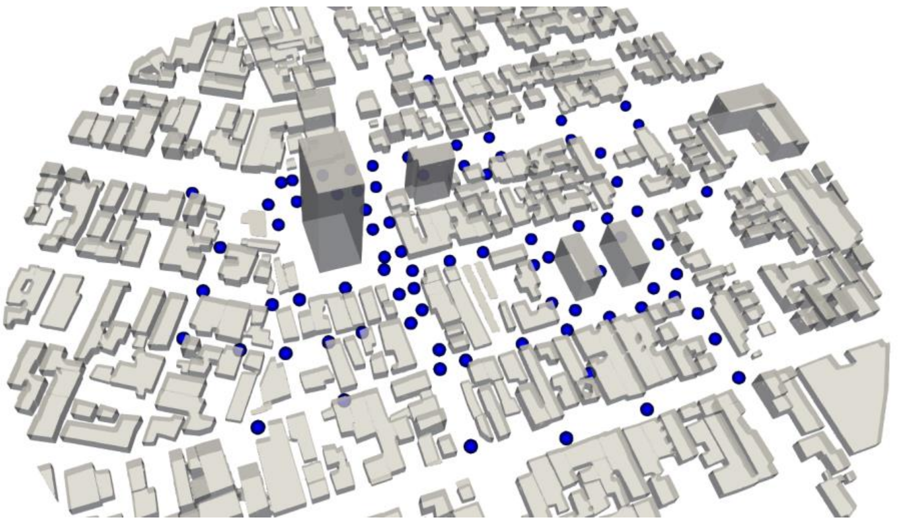

FIGURE 2.2 *3D view of the Case E urban geometry with the 80 pedestrian-level probe points (blue dots) shown in their spatial arrangement*

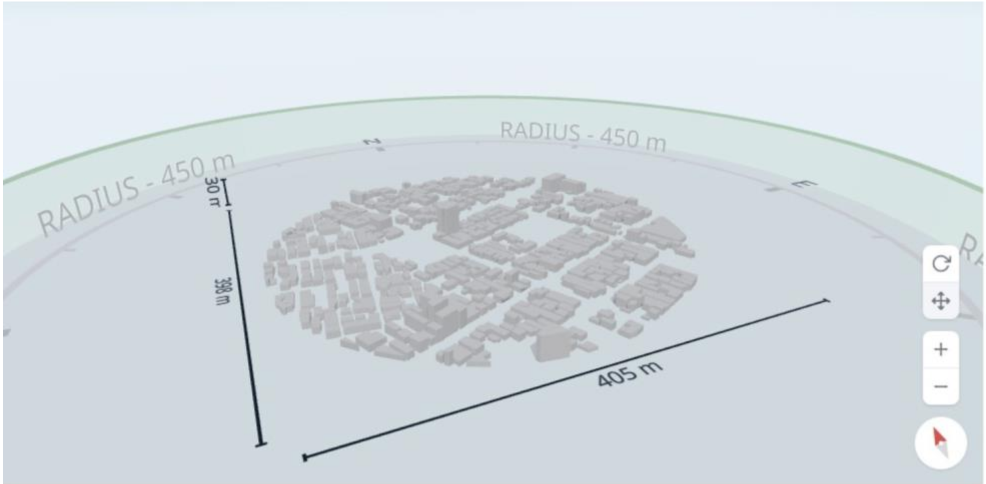

FIGURE 2.3 *ArchiWind computational domain — Before-Construction configuration.*

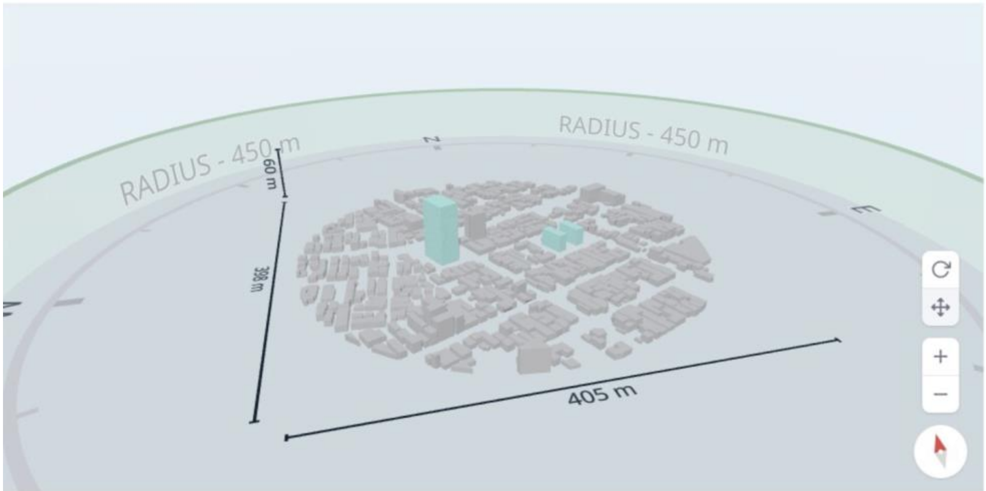

FIGURE 2.4 *ArchiWind computational domain — After-Construction configuration (new building visible at centre).*

2.2.2 COMPUTATIONAL DOMAIN

Mesh generation and boundary condition assignment are handled automatically by ArchiWind's preprocessing pipeline. The platform employs a cylindrical domain geometry following the approach of Hågbo and Giljarhus [5], which allows the same mesh to be reused for all 16 simulated wind directions by rotating the inlet boundary conditions rather than rebuilding the geometry.

While AIJ best-practice guidelines [1] specify distinct inlet and outlet distances for traditional box-shaped domains, those recommendations assume a fixed inflow direction and are not directly applicable to multi-direction studies. To resolve this, the radial distance from the central building cluster to the domain boundary is set to 15H in all directions (where H is the height of the tallest building), which is consistent with the AIJ 15H downstream outlet requirement [3]. This ensures that the inlet boundary remains placed well away from the region of influence of the buildings regardless of the wind angle. Within this domain, results are evaluated inside an inner zone of radius 5H from the building cluster, adhering strictly to the local resolution criteria specified by the benchmark.

2.2.3 TURBULENCE MODELS

Two RANS turbulence models were evaluated:

Realizable *k-ε* (production model): a well-established model for bluff-body flows, featuring improved predictions of spreading rates and recirculation zones compared to the standard k-ε formulation [8].

***k-ω* SST** (comparison model): a hybrid model combining k-ω near walls and k-ε in the free stream, often recommended for flows with adverse pressure gradients [6].

2.2.4 INLET BOUNDARY CONDITION AND REFERENCE VELOCITY

In standard applications, ArchiWind automatically derives atmospheric boundary layer profiles from upstream terrain roughness for each wind direction. For this benchmark, prescribed roughness parameters were used to reproduce the inflow conditions reported in the AIJ wind tunnel dataset [3].

The inlet velocity and turbulence profiles were prescribed from the AIJ-published wind tunnel inflow measurements. The reference velocity $U_{\mathrm{ref}} = 3.928\ \mathrm{m\ s^{-1}}$ is defined at z = 15.9 m (height of the Niigata meteorological station), consistent with the AIJ normalization convention.

To verify that the inlet profile is preserved through the fetch region upstream of the measurement area, an ABL profile check was performed at distances of 220 m to 900 m upstream of the building cluster. The figure below shows the velocity profile at each fetch distance compared to the AIJ target.

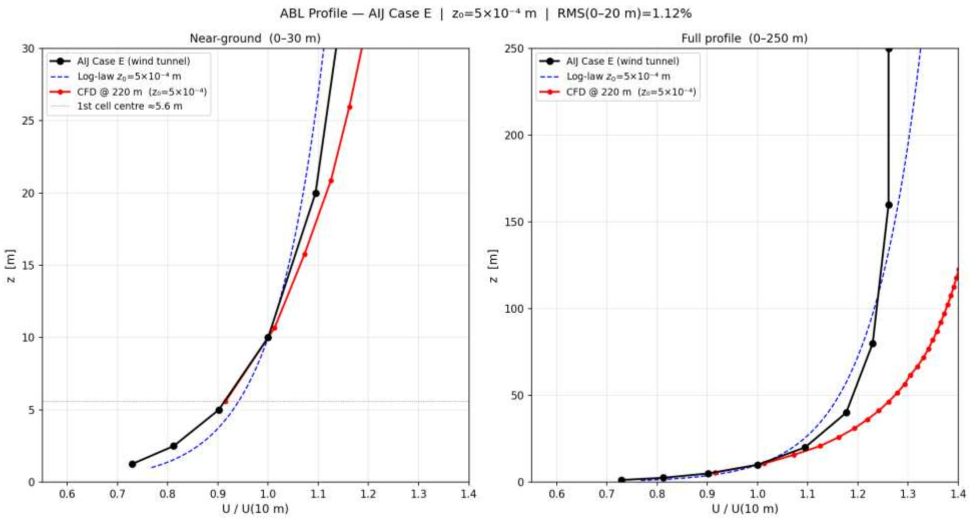

FIGURE 2.5 *Fetch profile analysis: simulated ABL velocity profile at multiple upstream distances (220–900 m) compared to the AIJ target inflow profile.*

2.2.5 VALIDATION METRICS

The primary quantity of comparison is the normalized scalar wind speed

$$\frac{V_{\mathrm{s}}}{U_{\mathrm{ref}}} = \frac{\sqrt{U_{\mathrm{x}}^2 + U_{\mathrm{y}}^2 + U_{\mathrm{z}}^2}}{U_{\mathrm{ref}}}$$

which is directly comparable to the normalised values provided by the AIJ datasets. Performance is quantified by three metrics computed over all N available measurement points, where CFD_{i} and Exp_{i} denote the normalized wind speed predicted and measured at probe i:

RMSE: root mean square error of the normalized wind speed

$$\mathrm{RMSE} = \sqrt{\frac{\sum(CFD_i - Exp_i)^2}{N}}$$

Bias: systematic over- or underprediction averaged across all probes. Positive means CFD overestimates, negative means CFD underestimates.

$$\mathrm{Bias} = \frac{1}{N}\sum(CFD_i - Exp_i)$$

R²: spatial correlation between CFD and experiment: $\boldsymbol{R}^2 = 1$ indicates perfect agreement, $\boldsymbol{R}^2 = 0$ means the CFD has no predictive power for the spatial wind pattern, $\boldsymbol{R}^2 < 0$ means it performs worse than simply predicting the spatial mean.

$$R^2 = 1 - \frac{\sum(CFD_i - Exp_i)^2}{\sum(Exp_i - \bar{E})^2}$$

where $\bar{E}$ is the mean of the experimental values.

2.3 AFTER-CONSTRUCTION VS. EXPERIMENT

This section compares ArchiWind results for the After-Construction configuration (Detailed mesh, Realizable *k-ε* model) against wind tunnel measurements.

2.3.1 GLOBAL ACCURACY

Figure 2.6 shows the global scatter plot pooling all 80 probe points and 16 wind directions (N = 1280 data pairs). The dashed line represents perfect agreement.

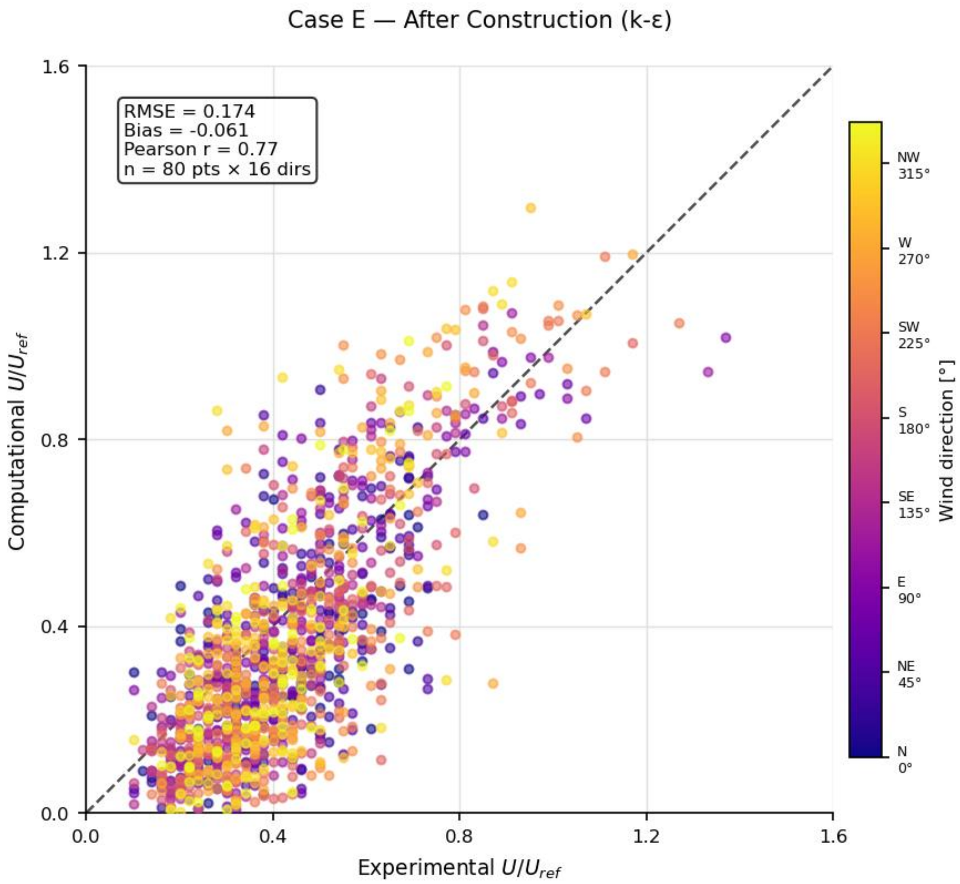

FIGURE 2.6 *Global scatter — After-Construction: Experimental (x) vs. Computational (y) normalised wind speed* U/U_{ref} *All 80 probe points × 16 directions. Dashed line = perfect agreement.*

TABLE 2.1 *Global validation metrics — Case E, After-Construction, Detailed mesh, Realizable k-ε*

Configuration	N pts	N dirs	RMSE	Bias	Pearson r
After — Detailed (k-ε)	80	16	0.174	−0.061	0.769

The RMSE of 0.174 and Pearson *r* of 0.769 are consistent with the accuracy range of 0.15–0.25 reported in the literature for steady-state RANS applied to complex urban environments [4, 7]. The systematic negative bias of −0.061 is physically expected: RANS closures underpredict wake recovery in the lee of buildings, leading to wider sheltered regions than observed in the wind tunnel. This bias is conservative for comfort assessments, as underestimation of wind speed never generates false alarms of dangerous conditions.

2.3.2 RESULTS BY CARDINAL DIRECTION

Figure 2.7 and ***Figure 2.8*** show line plots and scatter plots for the four cardinal wind directions (N, E, S, W). Each line plot displays ArchiWind (blue solid) against experiment (grey dashed) across all 80 probe points.

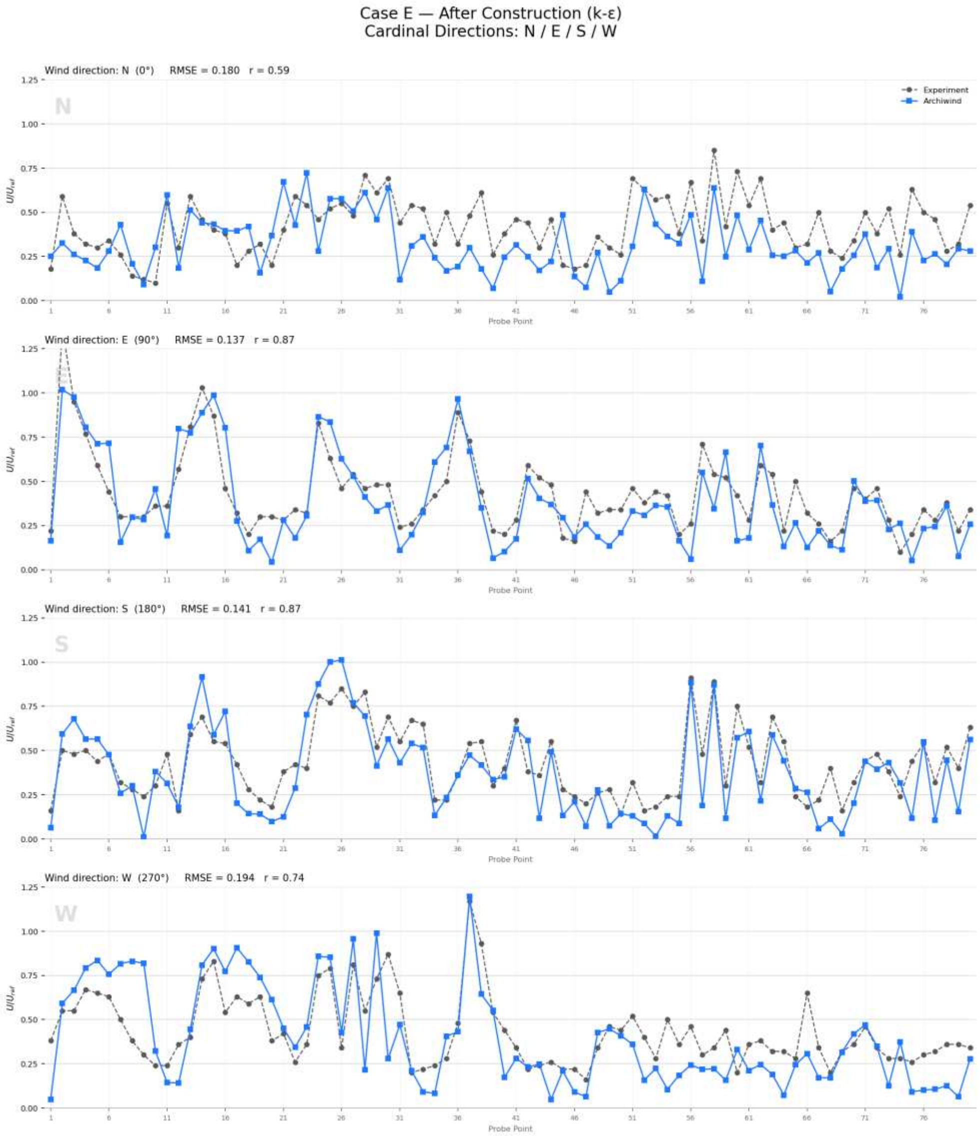

FIGURE 2.7 *Line plots — After-Construction, cardinal directions (N, E, S, W). Grey dashed: Experiment. Blue solid: ArchiWind (Realizable k-ε). RMSE and Pearson r are shown for each direction.*

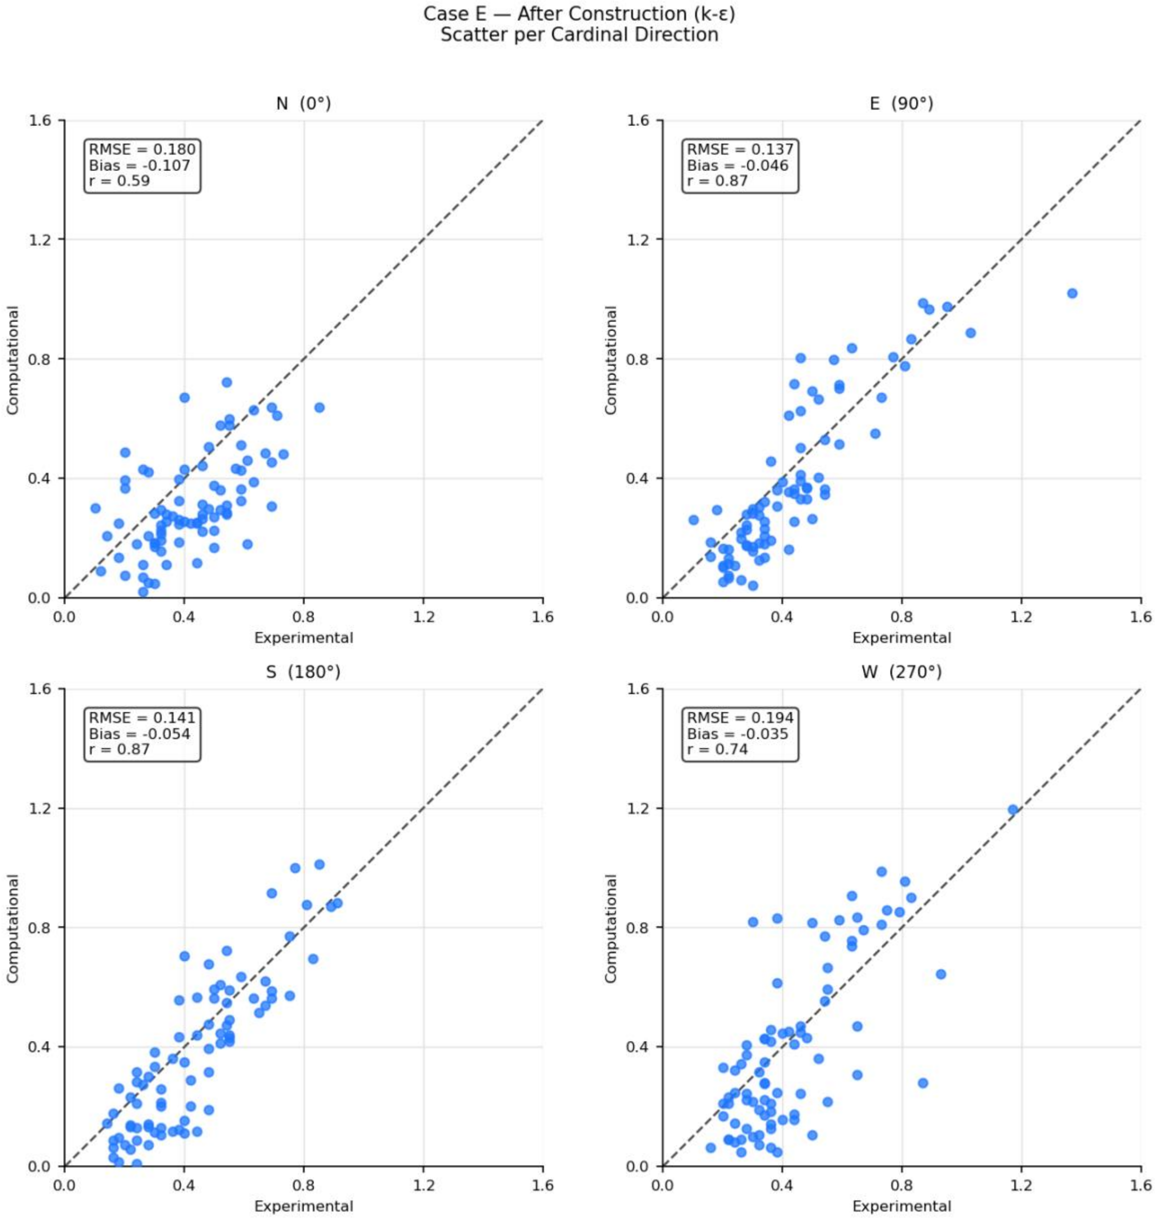

FIGURE 2.8 *Scatter plots per cardinal direction — After-Construction. x-axis: Experimental. y-axis: Computational.*

2.3.3 ERROR ANALYSIS

Figure 2.9 shows RMSE and bias for each of the 16 wind directions. ***Figure 2.10*** shows the normalized velocity field at pedestrian level for the North wind direction, providing a qualitative view of the CFD output.

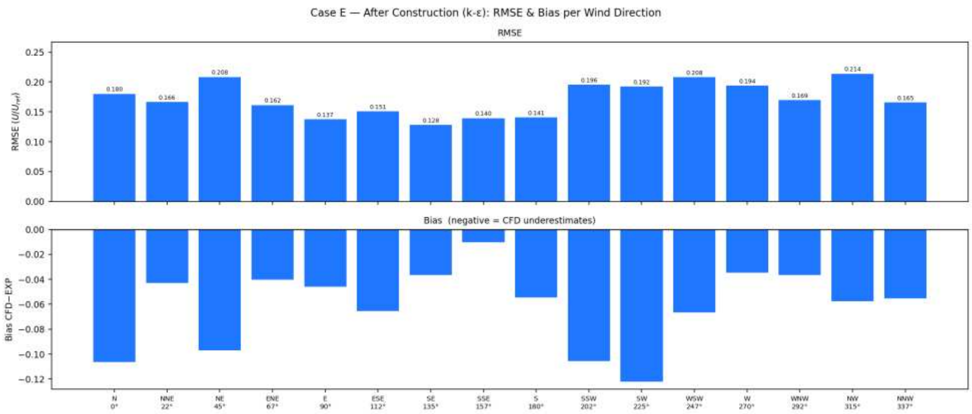

FIGURE 2.9 *Per-direction RMSE and bias — After-Construction, all 16 wind directions.*

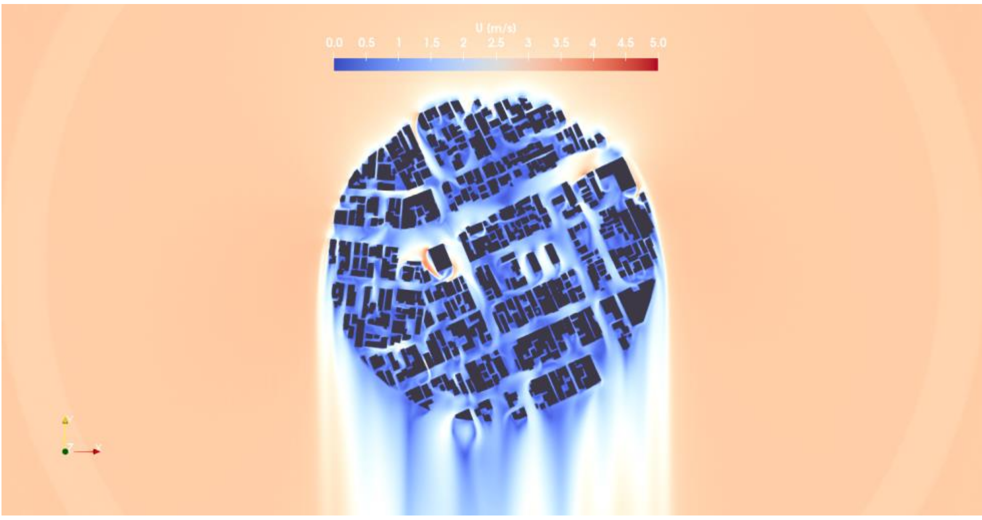

FIGURE 2.10 *Normalized velocity field U/U_{ref} at pedestrian level — After-Construction, North wind direction (0°). Blue: sheltered zones; warm colours: accelerated flow.*

2.4 BEFORE-CONSTRUCTION VS. EXPERIMENT

This section compares ArchiWind results for the Before-Construction configuration against wind tunnel measurements of the original urban fabric.

2.4.1 GLOBAL ACCURACY

Figure 2.11 shows the global scatter for the Before-Construction configuration.

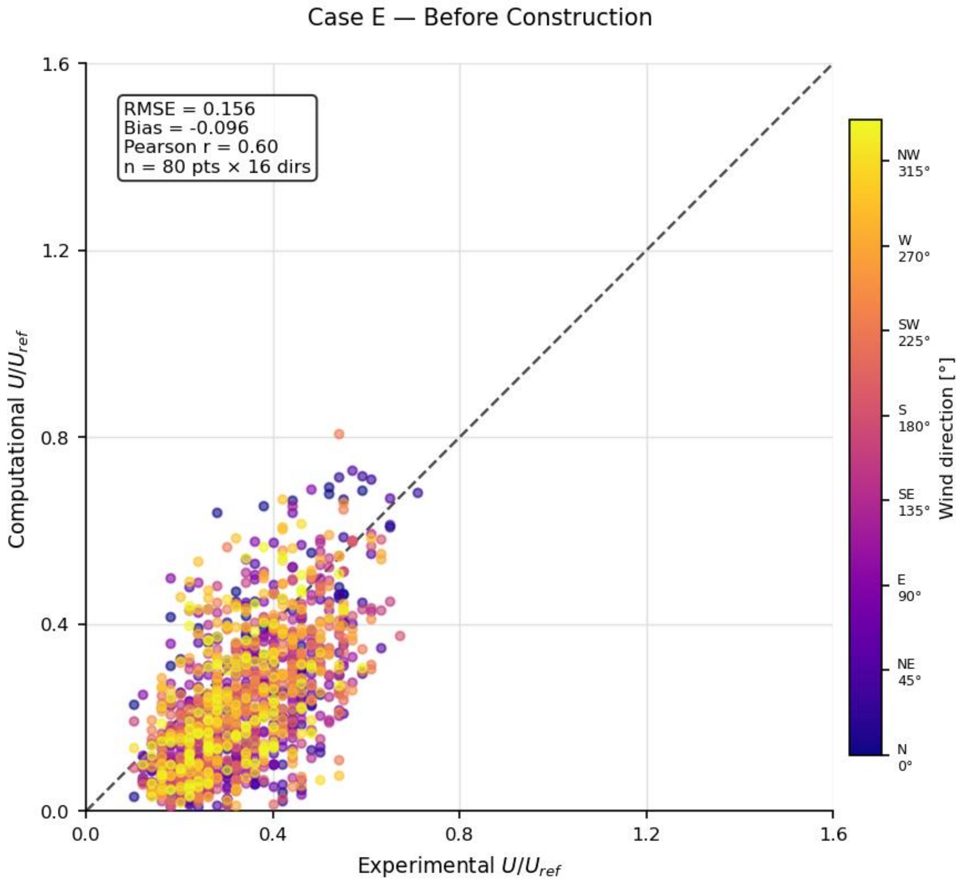

FIGURE 2.11 *Global scatter — Before-Construction: Experimental (x) vs. Computational (y) normalised wind speed* U/U_{ref}*. All 80 probe points × 16 directions.*

TABLE 2.2 *Global validation metrics — Case E, Before-Construction, Detailed mesh, Realizable k-ε.*

Configuration	N pts	N dirs	RMSE	Bias	Pearson r
Before—Detailed (k-ε)	80	16	0.156	−0.096	0.603

The Before-Construction scenario yields a lower RMSE (0.156) but also a lower Pearson *r* (0.603) compared to the After-Construction case. Without the new building, the dominant spatial features that create high-amplitude probe variability — corner acceleration zones and deep wakes — are absent, making the remaining subtle gradients harder to rank correctly in relative terms. The larger negative bias (−0.096) is consistent with more extensive open-space fetch regions where steady-state RANS underestimates wake recovery most severely.

2.4.2 RESULTS BY CARDINAL DIRECTION

Figure 2.12 and ***Figure 2.13*** show line plots and scatter plots for the four cardinal directions.

FIGURE 2.12 *Line plots — Before-Construction, cardinal directions (N, E, S, W). Grey dashed: Experiment. Blue solid: ArchiWind.*

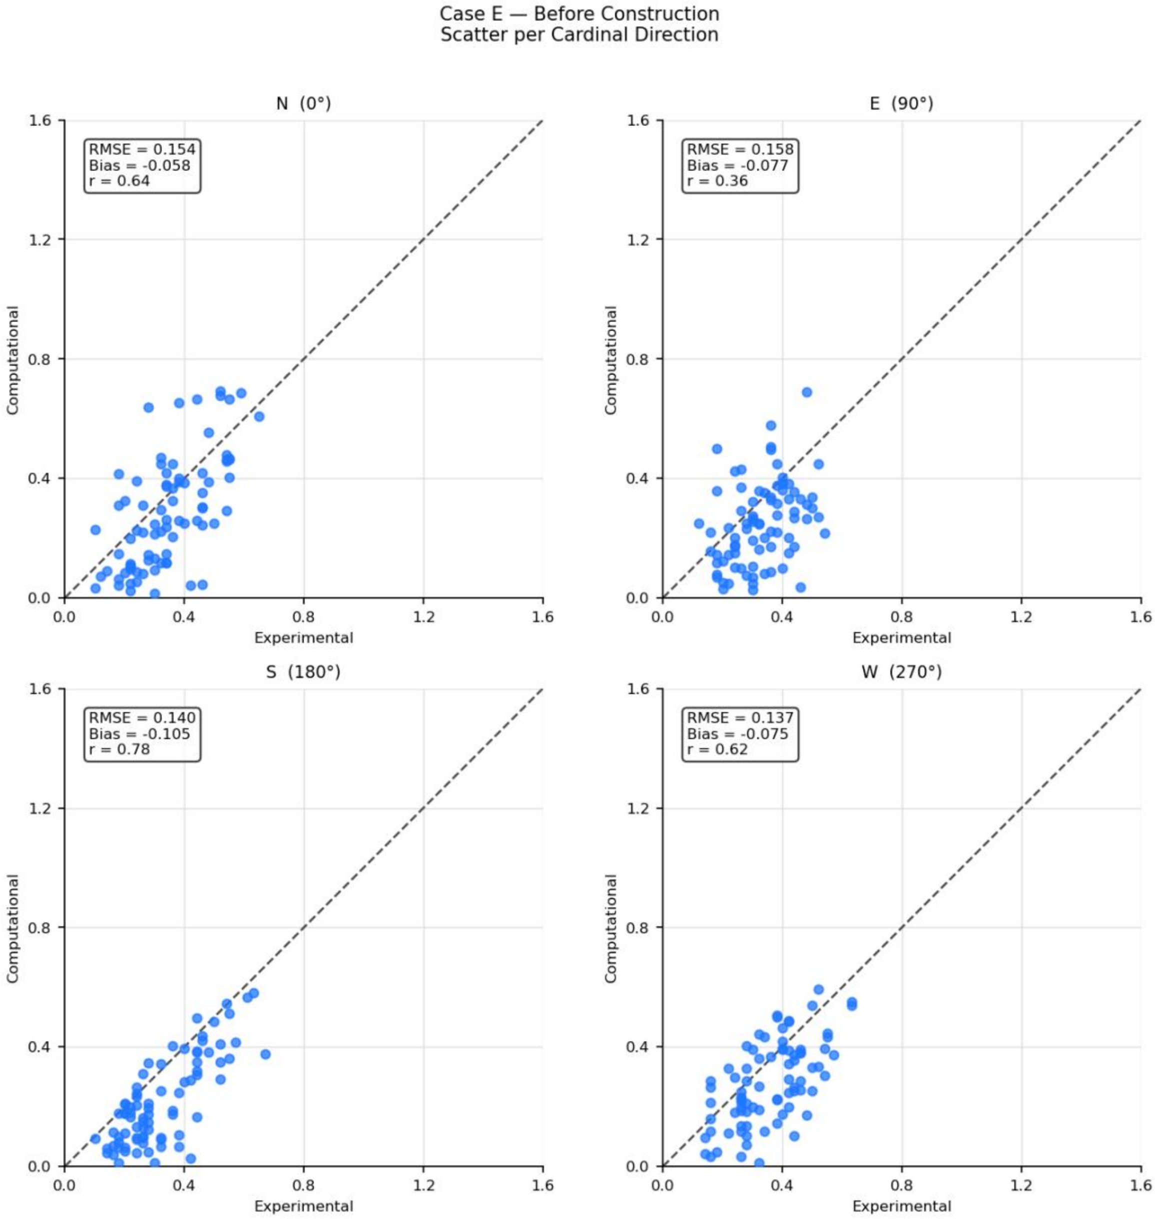

FIGURE 2.13 *Scatter plots per cardinal direction — Before-Construction. x-axis: Experimental. y-axis: Computational.*

2.4.3 ERROR ANALYSIS

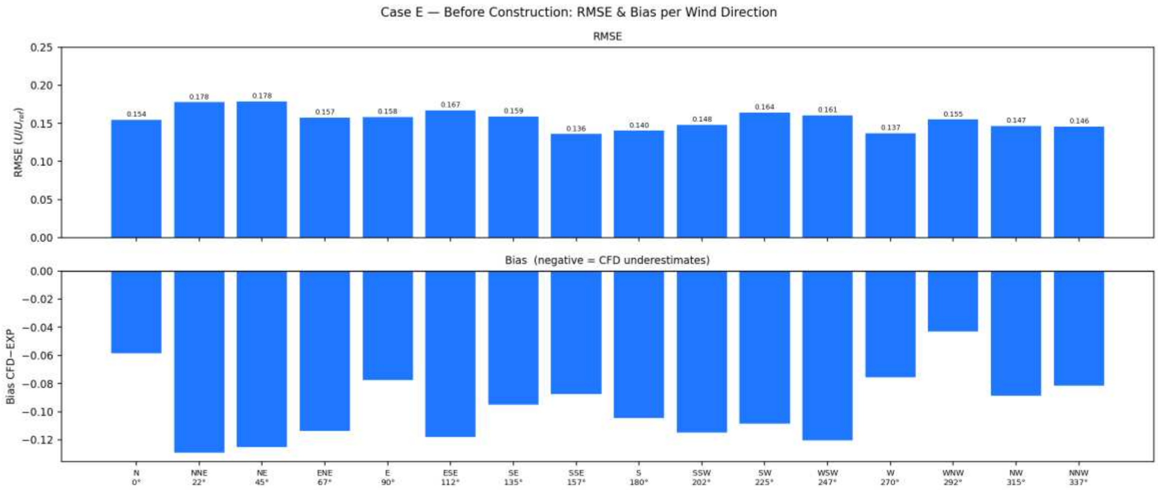

FIGURE 2.14 *Per-direction RMSE and bias — Before-Construction, all 16 wind directions.*

2.5 CONSTRUCTION EFFECT

A key use case for ArchiWind is assessing the wind impact of a new building on the surrounding pedestrian environment. This section evaluates whether the CFD-predicted change in wind speed $\Delta U = U_{\text{after}} - U_{\text{before}}$ matches the change measured in the wind tunnel.

2.5.1 BEFORE VS. AFTER

Figure 2.15 places the two global scatter plots side by side to illustrate how overall accuracy compares between the two construction scenarios.

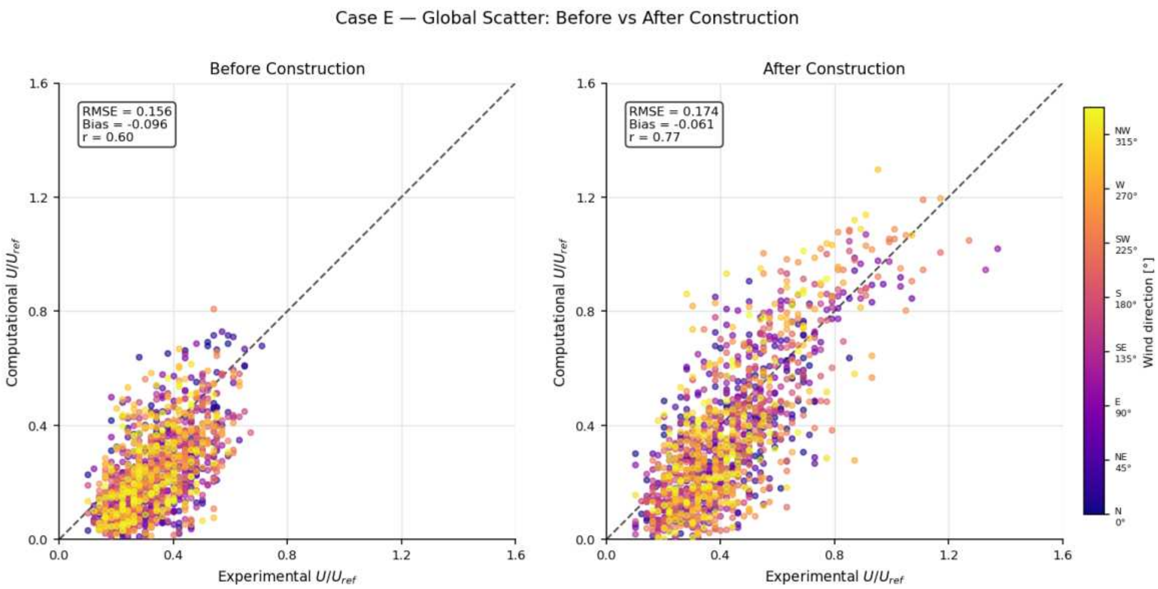

FIGURE 2.15 *Global scatter — Before (left) and After (right) construction. Axes: Experimental (x) vs. Computational (y) normalised wind speed U/U_{ref}..*

2.5.2 ΔU SCATTER AND SPATIAL MAPS

Figure 2.16 shows the scatter of CFD-predicted $\Delta U/U_{\mathrm{ref}}$ versus experimentally measured $\Delta U/U_{\mathrm{ref}}$ for all 80 probe points and 16 wind directions, together with spatial maps of mean ΔU per probe point averaged over all 16 directions. Positive ΔU indicates wind speed increase (channelling or corner acceleration); negative values indicate sheltering. The scatter confirms that ArchiWind correctly identifies the sign and approximate magnitude of the construction effect at most probe locations.

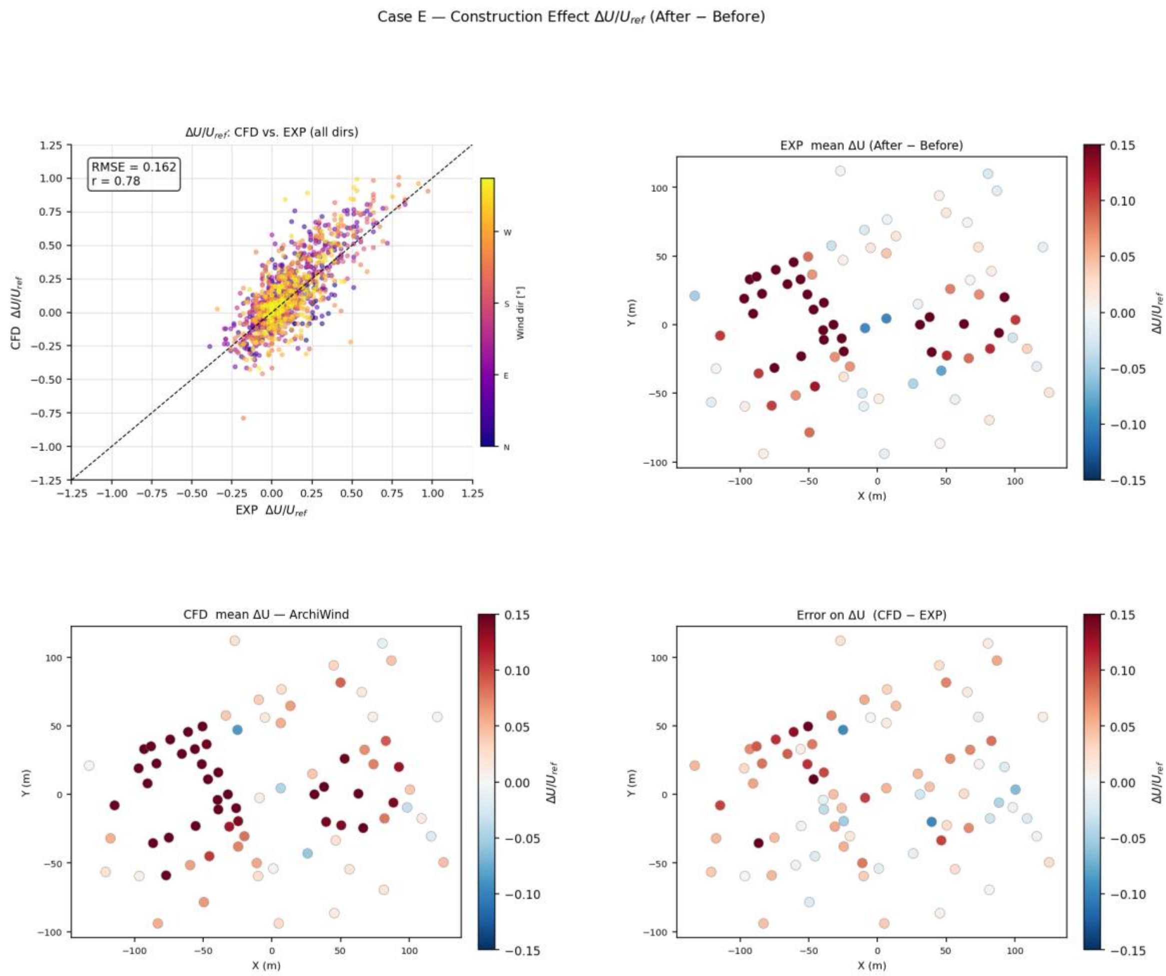

FIGURE 2.16 Construction effect $\Delta U/U_{ref}$. (After − Before). Top-left: scatter CFD vs. experiment coloured by wind direction. Top-right: experimental mean ΔU per probe. Bottom-left: ArchiWind mean ΔU. Bottom-right: error on ΔU (CFD − EXP). Spatial maps averaged over 16 wind directions.

2.5.3 MEAN VELOCITY AND RESIDUAL MAPS

Figure 2.17 and ***Figure 2.18*** provide spatial context by showing the mean velocity ratio averaged over all 16 directions and the mean signed residual (CFD − EXP) per probe point for both construction scenarios.

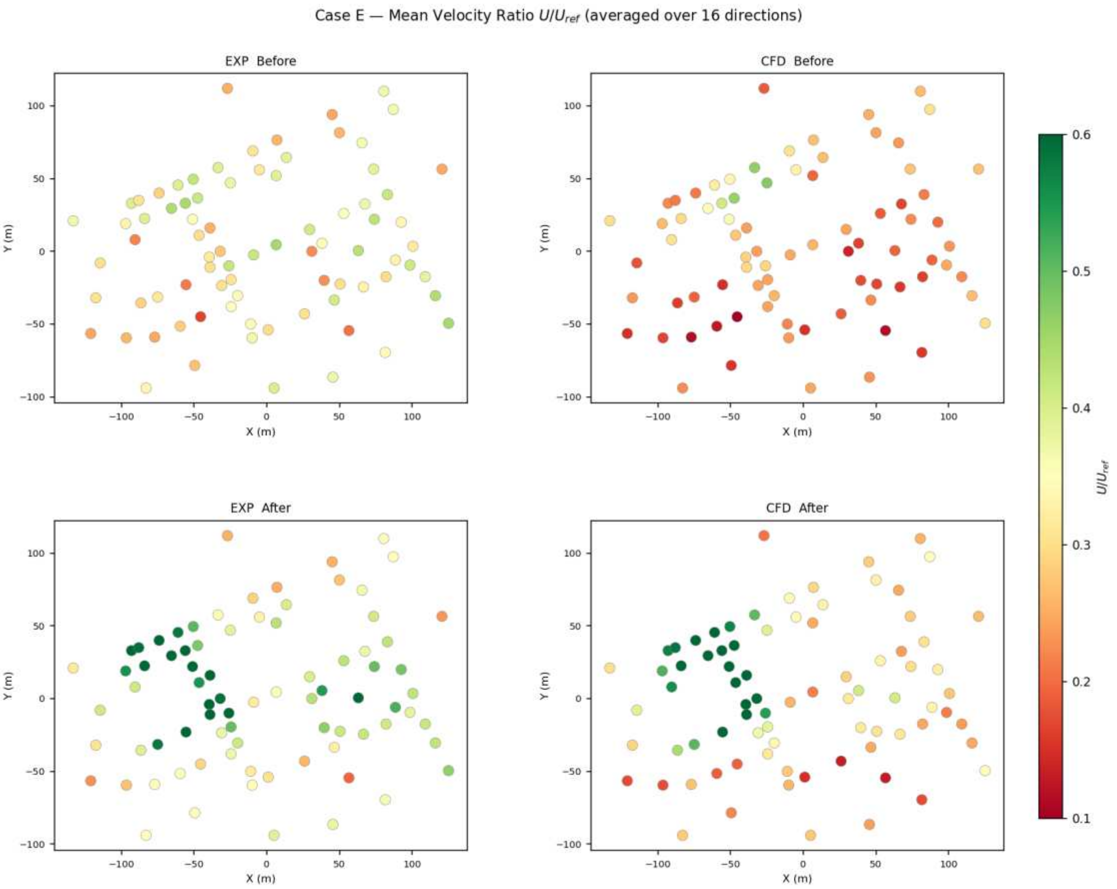

FIGURE 2.17 *Spatial map of mean normalised wind speed U/U_{ref} averaged over 16 directions. Rows: Before (top) / After (bottom). Columns: Experiment (left) / ArchiWind (right).*

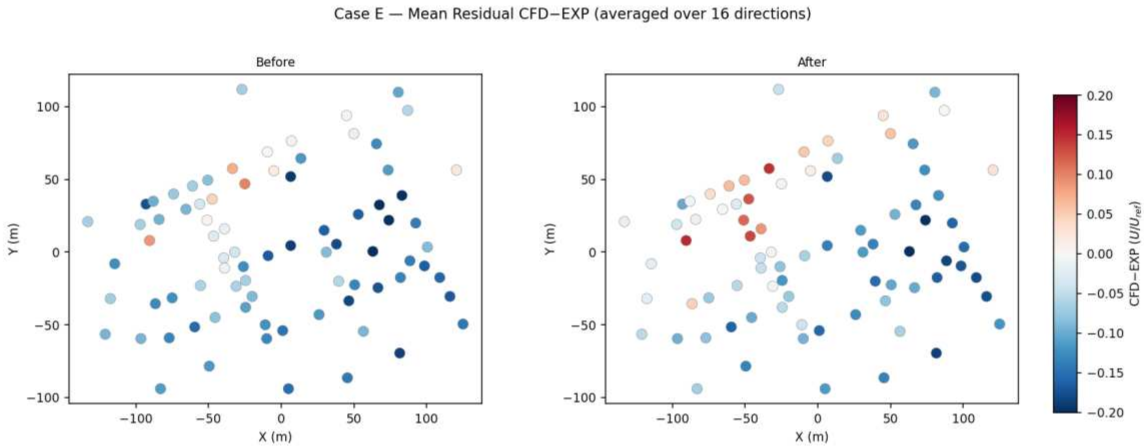

FIGURE 2.18 *Mean residual (CFD − EXP) per probe point averaged over 16 directions. Left: Before-Construction. Right: After-Construction. Red = CFD overestimates; blue = CFD underestimates.*

2.6 MESH QUALITY: DETAILED VS. DRAFT

ArchiWind offers two mesh quality settings: Detailed and Draft. The Detailed setting generates a finer hexahedral mesh with tighter near-wall refinement around building surfaces and layered inflation cells on all solid surfaces. The Draft setting uses a coarser base mesh for faster computation. ***Figure 2.19*** and ***Figure 2.20*** show mesh cross-sections for each setting. This section quantifies the impact of mesh quality on prediction accuracy for the After-Construction scenario.

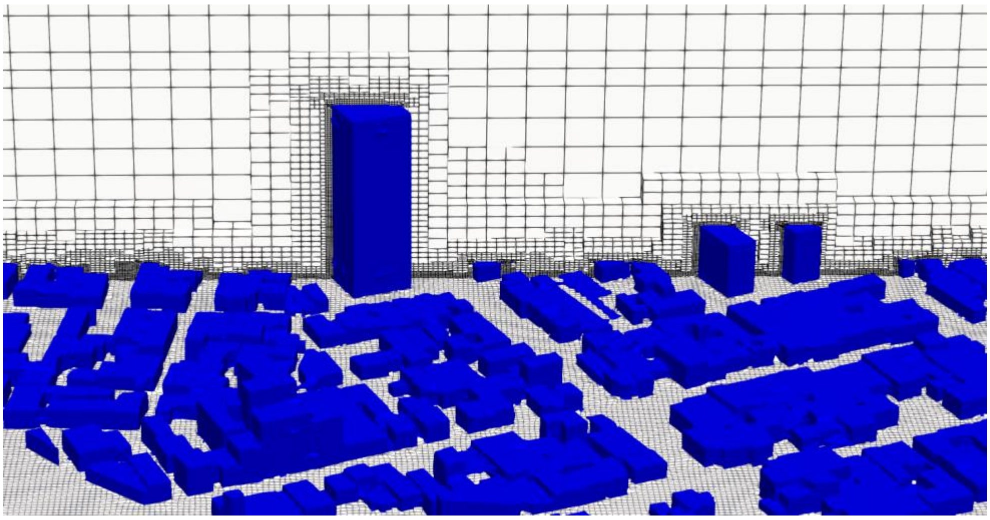

FIGURE 2.19 *Mesh cross-section — Detailed quality setting. Fine hexahedral cells near building surfaces with layered refinement.*

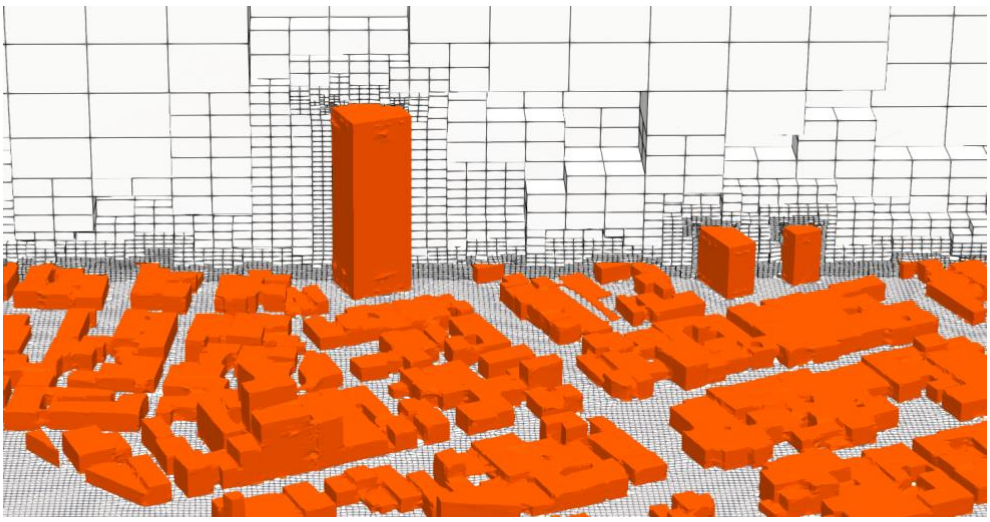

FIGURE 2.20 *Mesh cross-section — Draft quality setting. Coarser cell distribution with reduced near-wall refinement.*

2.6.1 GLOBAL ACCURACY

Figure 2.21 shows the global scatter for both mesh quality levels side by side.

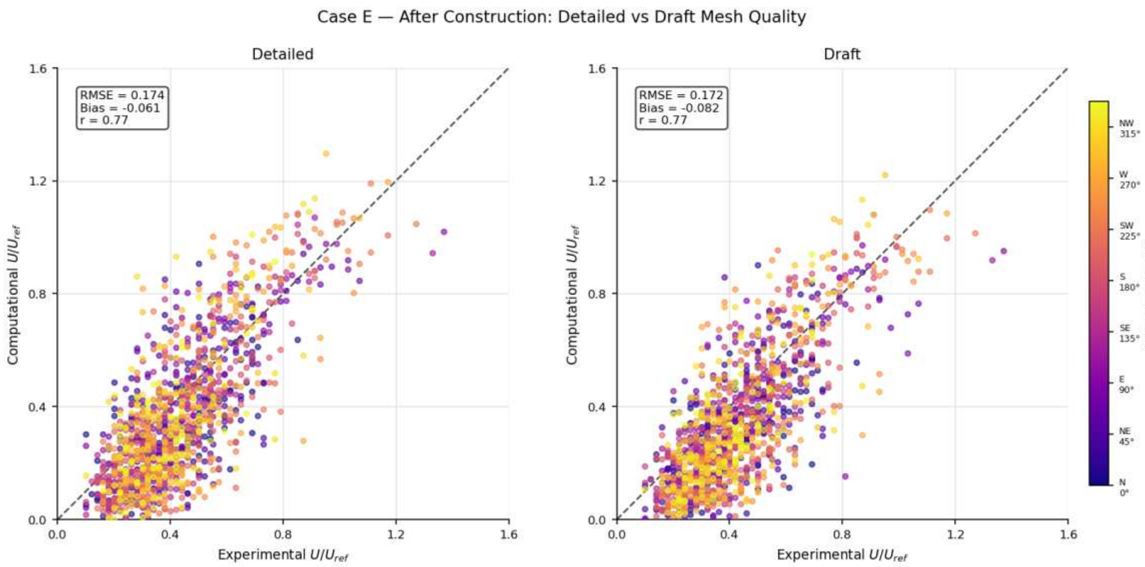

FIGURE 2.21 *Global scatter — Detailed (left) vs. Draft (right) mesh quality, After-Construction. Axes: Experimental (x) vs. Computational (y).*

TABLE 2.3 *Metric comparison — Detailed vs. Draft mesh quality, After-Construction.*

Mesh Quality	RMSE	Bias	Pearson r
Detailed	0.174	−0.061	0.769
Draft	0.172	−0.082	0.769

The Draft mesh yields global metrics virtually identical to the Detailed setting: ΔRMSE = −0.002 and identical Pearson r = 0.769. The slightly more negative bias (−0.082 vs. −0.061) reflects reduced near-wall resolution, with negligible operational impact on comfort classification.

2.6.2 RESULTS BY CARDINAL DIRECTION

Figure 2.22 overlays Detailed and Draft results on the same line plots for the four cardinal directions. Figure 2.23 shows per-direction RMSE and bias for both quality settings.

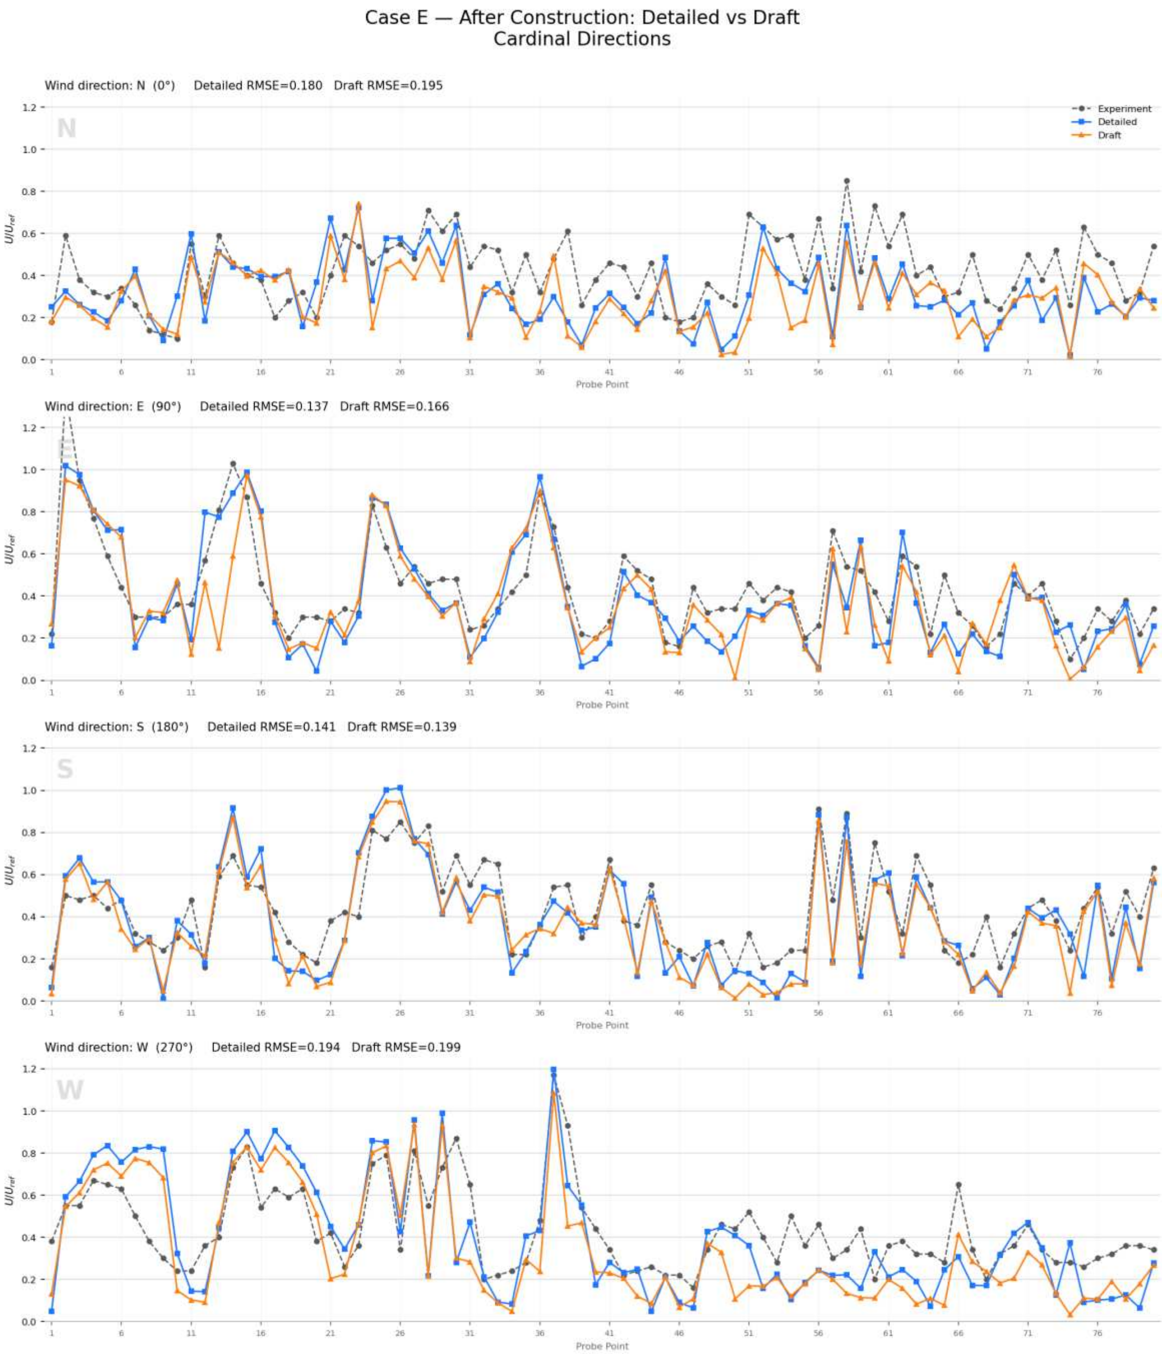

FIGURE 2.22 *Line plots — After-Construction, cardinal directions: Experiment (grey dashed), Detailed (blue), Draft (orange). RMSE shown per direction for each quality setting.*

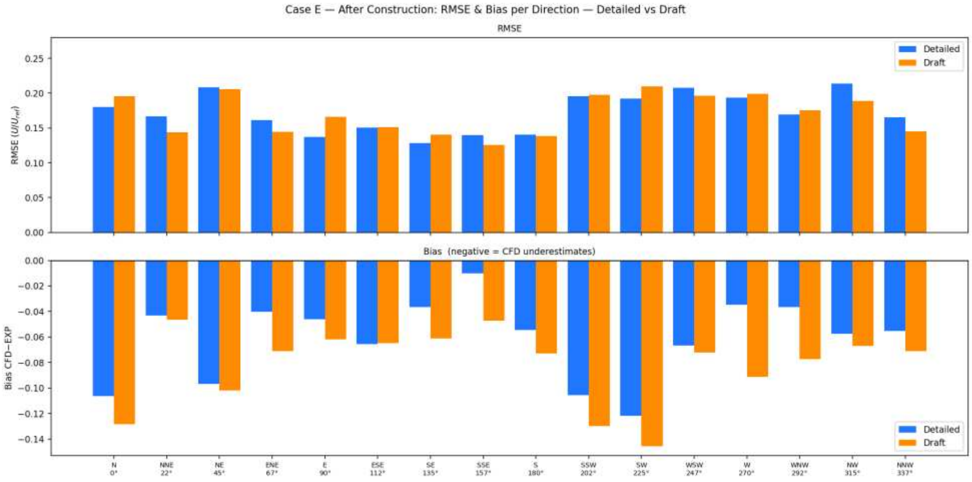

FIGURE 2.23 *Per-direction RMSE and bias: Detailed (blue) vs. Draft (orange), After-Construction.*

2.6.3 ERROR DISTRIBUTION

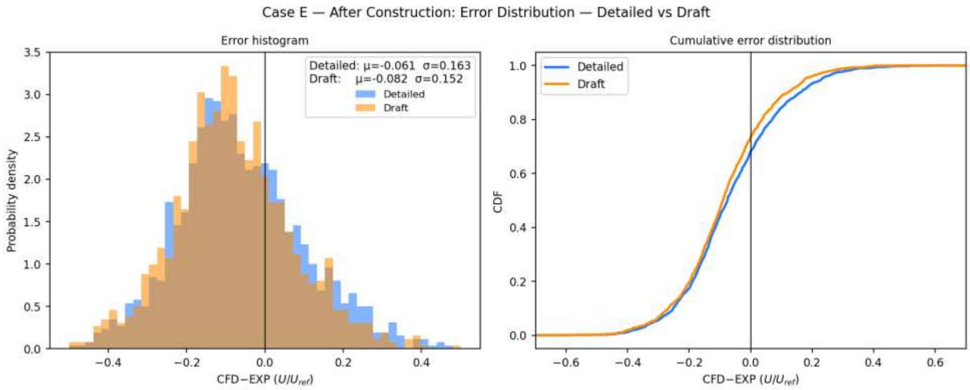

FIGURE 2.24 *Error distribution (CFD − EXP) — Detailed (blue) vs. Draft (orange). Left: histogram. Right: cumulative distribution.*

2.7 TURBULENCE MODEL: K- ε VS. K-Ω SST

This section compares the Realizable *k-ε* model against the *k-ω* SST model for the After-Construction configuration. Both simulations use the same Detailed mesh and boundary conditions; only the turbulence closure differs.

2.7.1 GLOBAL ACCURACY

Figure 2.25 shows the global scatter for both turbulence models side by side.

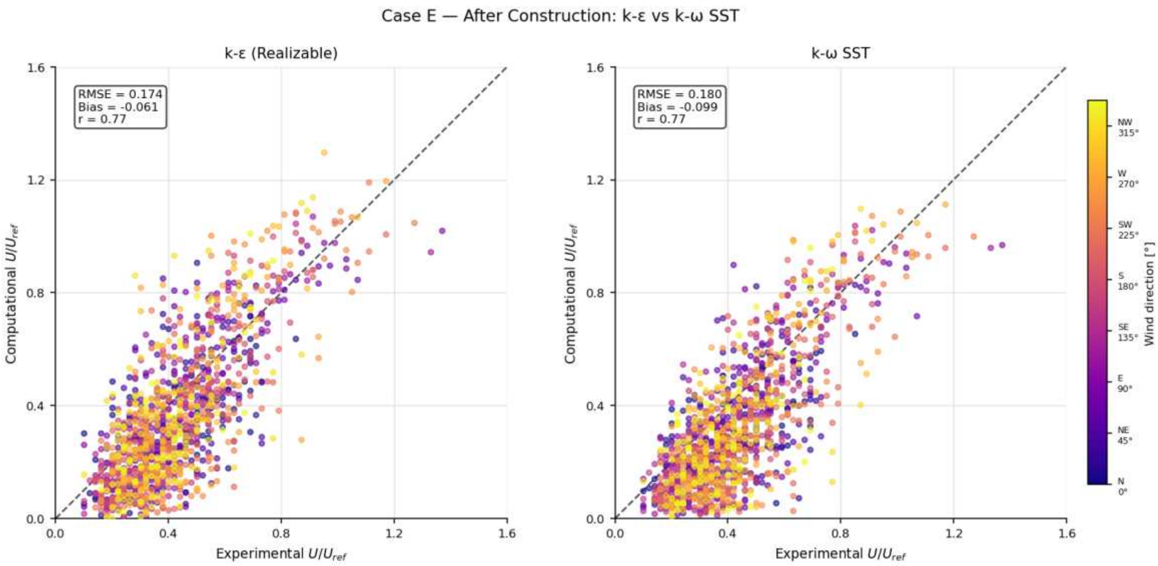

FIGURE 2.25 *Global scatter — k-ε (left) vs. k-ω SST (right), After-Construction. Axes: Experimental (x) vs. Computational (y). Points coloured by wind direction.*

TABLE 2.4 *Global validation metrics — Realizable k-ε vs. k-ω SST, After-Construction.*

Turbulence Model	RMSE	Bias	Pearson r
Realizable k-ε	0.174	−0.061	0.769
k-ω SST	0.180	−0.099	0.771

The *k-ω* SST model yields comparable global accuracy (RMSE = 0.180, r = 0.771) but a substantially larger negative bias (−0.099 vs. −0.061). The lower bias of the Realizable *k-ε* model is the deciding factor in its selection as the default ArchiWind turbulence model.

2.7.2 RESULTS BY CARDINAL DIRECTION

Figure 2.26 overlays *k-ε* and *k-ω* SST results on the same line plots for the four cardinal directions. ***Figure 2.27*** shows per-direction RMSE and bias for both models.

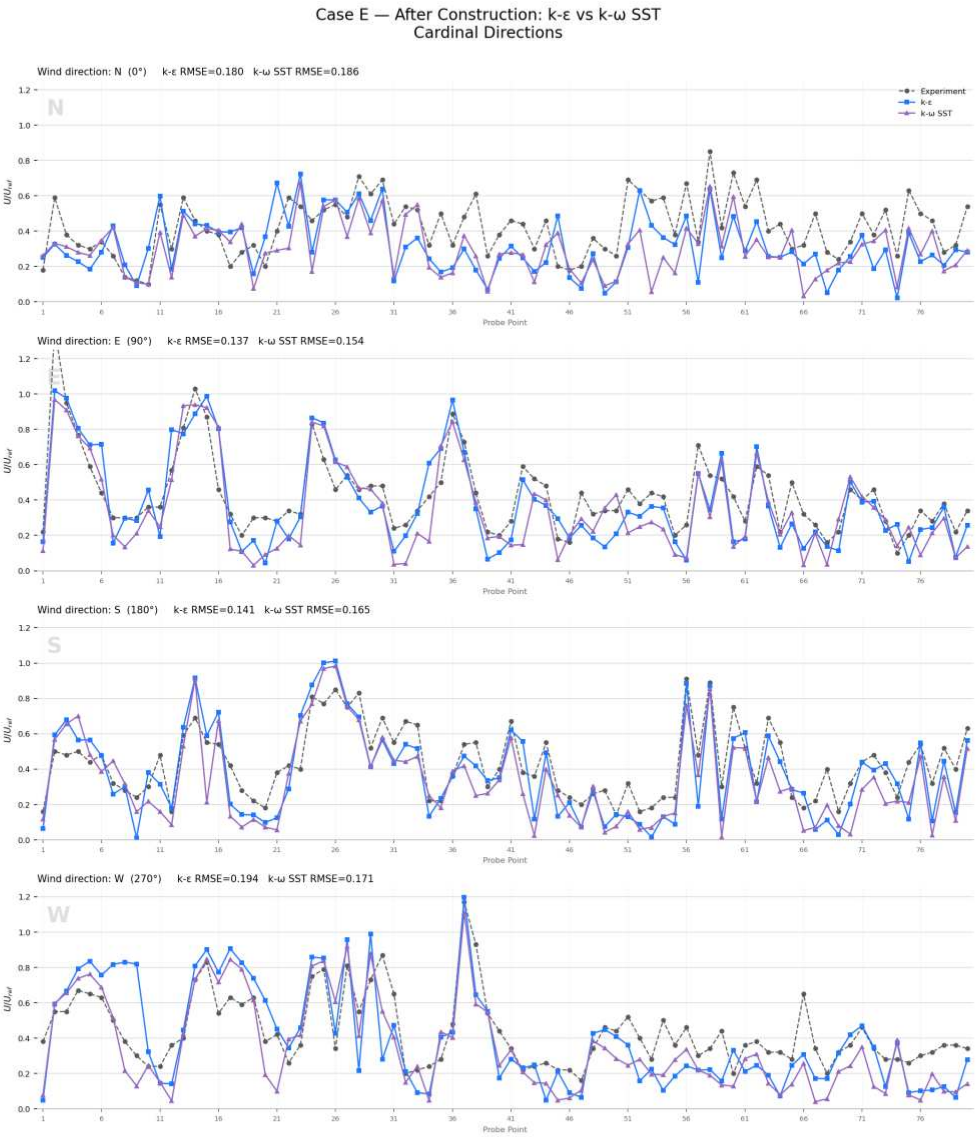

FIGURE 2.26 *Line plots — After-Construction, cardinal directions: Experiment (grey dashed), k-ε (blue), k-ω SST (violet). RMSE shown per direction for each model.*

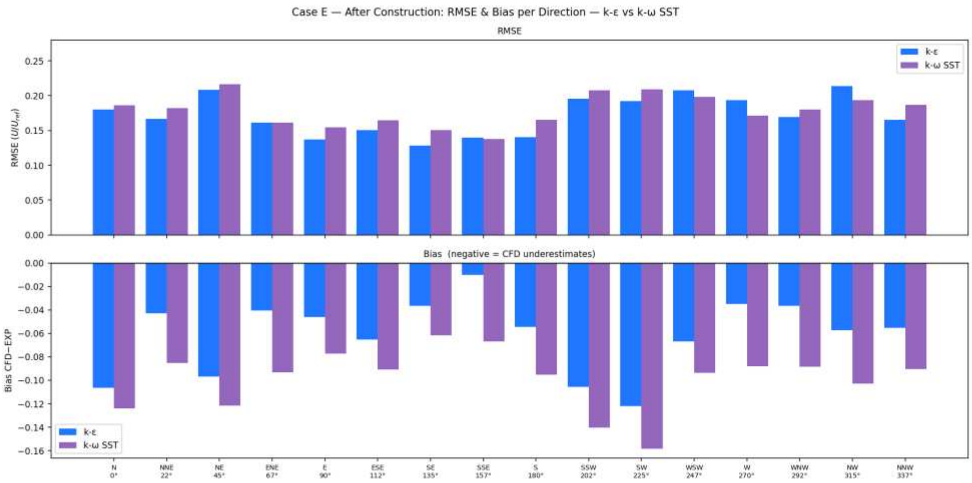

FIGURE 2.27 *Per-direction RMSE and bias: k-ε (blue) vs. k-ω SST (violet), After-Construction.*

2.7.3 ERROR DISTRIBUTION

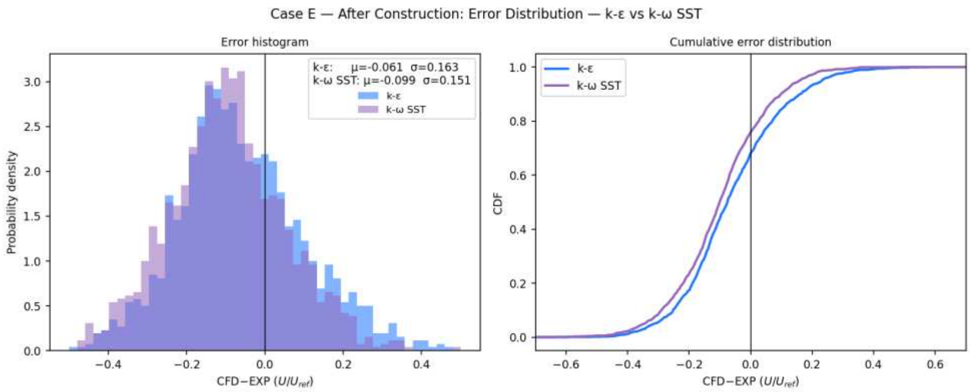

FIGURE 2.28 *Error distribution (CFD − EXP) — k-ε (blue) vs. k-ω SST (violet). Left: histogram. Right: cumulative distribution.*

2.8 CONCLUSIONS

The quantitative evaluation of the platform yields four primary conclusions regarding its operational deployment.

Mesh sensitivity and screening efficiency. The Draft mesh quality setting provides near-identical global accuracy to the Detailed configuration ($\Delta\text{RMSE} = -0.002$, $\Delta r = 0.000$). This result strongly supports the integration of a multi-fidelity workflow in architectural practice, where the faster Draft quality is utilized for rapid, iterative screening assessments, while the Detailed mesh is reserved exclusively for final regulatory compliance submissions.

Turbulence model selection. The Realizable k- ε closure model outperforms the k-ω SST formulation primarily through a significantly lower systematic bias (−0.061 versus −0.099). This statistical advantage confirms the selection of Realizable k-ε as the default turbulence model within the ArchiWind simulation engine for dense urban contexts.

Impact assessment capabilities. The platform consistently detects both the correct sign and the approximate magnitude of the construction effect ΔU across most probe-direction combinations. This structural fidelity confirms the suitability of the automated CFD pipeline for conducting before-and-after construction impact assessments in dense city centres.

Operational conservatism. The documented systematic negative bias remains physically consistent with well-known RANS streamline curvature limitations. Operationally, this behaviour introduces a built-in conservative margin for pedestrian safety; underestimating local wind speeds ensures that potential critical discomfort zones are flagged during the design phase rather than generating false negatives in comfort classifications.

3 CFD METHODOLOGY ENHANCEMENTS

This chapter presents two independent enhancements to the ArchiWind CFD workflow introduced during this work. Section 3.1 describes the transition from a static vegetation model to a physics-based, geometry-aware porous media formulation with a numerical stability limiter. Section 3.2 describes the development and validation of a two-stage physics-driven convergence criterion that replaces the conventional monitoring of residuals and max iterations with a sequential logic checking for residual thresholds or stagnation, followed by the autonomous detection of physical field stagnation at pedestrian level.

3.1 POROUS MEDIA VEGETATION MODEL

3.1.1 INTRODUCTION

The assessment of urban wind comfort and microclimate requires an accurate representation of the aerodynamic interaction between the atmospheric boundary layer and vegetated elements. Canopy structures, such as trees and hedges, play a crucial role in mitigating local wind speeds and enhancing the pedestrian-level microclimate.

However, explicitly resolving the complex, multi-scale geometry of foliage within Computational Fluid Dynamics (CFD) frameworks is computationally prohibitive. This section details the transition from a legacy static height allocation to a dynamic, geometry-based numerical methodology, specifically developed to enhance both the physical fidelity and the numerical stability of vegetation modeling within the ArchiWind platform.

3.1.2 THEORETICAL AND MATHEMATICAL FOUNDATIONS

Aerodynamic mechanisms of urban canopies. Natural vegetation alters the local flow field primarily through momentum dissipation. This macro-level windbreak effect scales directly with canopy density, spatial morphology, and aerodynamic porosity, resulting in localized velocity reductions and downstream wake regions.

In urban microclimate planning, these effects are governed by two primary geometric and biological features: the canopy aspect ratio (height-to-width ratio), which dictates the spatial extent of the downwind sheltered wake, and the seasonal variation in foliage density. Even during winter periods, the bare branch networks of deciduous species maintain a residual aerodynamic drag capable of moderating local microclimates.

Porous media formulation. To integrate these mechanisms into steady-state RANS simulations without resolving explicit leaf geometries, vegetation is modeled as a continuous, permeable porous medium. The hydrodynamic resistance exerted by the foliage is represented mathematically by introducing a volumetric momentum sink term S_{i} directly into the Navier–Stokes equations via the OpenFOAM class `atmPlantCanopyUSource`:

$$S_{\mathrm{i}} = -\alpha\rho C_{\mathrm{d}}\mathrm{LAD}|U_0|U_{\mathrm{i}}$$

where α is the fluid phase fraction (set to 1 for single-phase atmospheric flows), ρ is the fluid density, C_{d} is the dimensionless canopy drag coefficient, $|U_0|$ is the magnitude of the velocity field from the previous iteration, U_{i} is the current velocity component in the i–th direction, and LAD is the leaf area density (a user specified value representing the plant canopy density). This source term acts as a continuous drag force opposing the flow velocity within the plant canopy region, modeling the pressure drop and kinetic energy extraction characteristic of fluid flow through dense vegetative structures.

Turbulence transport source terms. The physical presence of a porous canopy also significantly alters the local turbulence field, generating small-scale vortices in the wakes of leaves and branches while accelerating the dissipation of larger shear-layer eddies. To capture this behavior within the *k*- ε framework, the numerical setup interfaces with the standard OpenFOAM implementation (`atmPlantCanopyTurbSource`), which acts on the turbulence transport equations as follows:

Turbulent Kinetic Energy equation (*k*): No explicit source term is added directly to the *k* equation (S_{k} = 0). The canopy does not directly inject or extract turbulent kinetic energy via an explicit mathematical term; instead, the TKE field is modified indirectly. The reduction in mean velocity caused by the aerodynamic drag S_{i} naturally lowers the overall production of TKE through the standard shear production term $G_{\mathrm{k}} \propto \partial U_{\mathrm{i}}/\partial x_{\mathrm{j}}$

Dissipation Rate equation (ε): To account for wake-induced dissipation acceleration an explicit source term S_{ε} is introduced:

$$S_{\varepsilon} = \alpha\rho(\mathrm{C}_1 - \mathrm{C}_2)\mathrm{C}_{\mathrm{canopy}}$$

with:

$$C_{\text{canopy}} = \frac{1}{2.0} \mathrm{C}_{\mu} \mathrm{C}_{\mathrm{d}} \mathrm{LAD} |U_0|$$

where C_1 and C_2 are closure constants governing canopy-induced turbulence production and destruction, and Cμ is the standard empirical model constant. By substituting their standard numerical values ($C_1 = 1.44$, $C_2 = 1.92$, and $C_\mu = 0.09$), the difference inside the parentheses results in a negative value $(C_1 - C_2) = (1.44 - 1.92) = -0.48$. Since this term is plugged directly into the k- ε transport solver framework, the combination of signs resolves into a net positive source term. Mathematically, this yields an effective multiplier of +1.728 ahead of the remining physical parameters (C_d, LAD, $|U_0|$, and ε). Physically, this positive contribution ensures that the presence of the vegetation acts as a net source of turbulence dissipation. These terms systematically modify the effective eddy viscosity inside the vegetated zone, preventing the overestimation of wind penetration through dense shelters by modelling how the small-scale wakes around leaves and branches accelerate the breakdown of larger turbulent structures (wake production).

Leaf Area Density (LAD) and Index (LAI). The structural porosity of the canopy is quantified via the Leaf Area Density (LAD), which is the leaf area per unit volume of foliage. It determines the porosity of the tree and influences how much air resistance the foliage provides.

$$\mathrm{LAD} = \mathrm{LAI} \cdot \frac{h_\mathrm{t} - z_\mathrm{m}}{h_\mathrm{t} - z} n \exp\left(n \left(1 - \frac{h_\mathrm{t} - z_\mathrm{m}}{h_\mathrm{t} - z} \right) \right)$$

While it is common to make the hypothesis of horizontal homogeneity, LAD will depend on the vertical location LAD(z). The empirical relation allows to link LAD(z) to a global quantity LAI.

Representative experimental values for common European tree species are documented in the literature to calibrate these numerical domains:

TABLE 3.1 *Leaf Area Index values for common European tree species.*

Tree Species	LAI
Sycamore	2.97
Silver Birch	3.24
Quercus	5.16
Chestnut	5.19
Platanus	5.28

Rather than implementing individual botanical species, ArchiWind categorizes vegetation into four operational aerodynamic bands: VERY_DENSE (LAI = 5.28, calibrated on *Platanus*), DENSE (LAI = 3.24, *Silver Birch*), MEDIUM_DENSE (LAI = 2.97, *Sycamore*), and LOW_DENSE (LAI = 1.07, representing sparse herbaceous cover).

Although the vertical profile LAD(z) provides a highly detailed representation of the canopy structure, its point-by-point implementation within numerical solvers can be computationally intensive and requires specific empirical parameters (z_m, n) that are often difficult to calibrate for general urban environments.

To address this, numerical wind engineering tools such as OpenFOAM via the standard `atmPlantCanopy` source terms, frequently adopt a simplified, volume-averaged approached. Under the assumption that the foliage density is homogeneously distributed along the vertical direction, the local Leaf Area density is generalized as a constant average value over the entire height of the tree, defined as:

$$\mathrm{LAD} = \frac{\mathrm{LAI}}{h}$$

where h is the height of the vegetation canopy. This formulation allows for a more efficient allocation of the porous media properties within the numerical grid cells while maintaining a global aerodynamic resistance consistent with the leaf area index of the specific tree species.

3.1.3 NUMERICAL IMPLEMENTATION

Dynamic geometric analysis: Voxel-Based Canopy Profiling. The core of the new methodology lies in the transition from a legacy static height h assignment, previously fixed at a default value regardless of geometry, to a Voxel-Based Canopy Profiling (VCP) approach. This algorithm ensures that the applied LAD is calculated based on the effective biological thickness of the vegetation, regardless of its placement or scale in the digital domain.

Slope-invariant height extraction. The VCP algorithm was implemented to resolve two critical geographical challenges: complex topography (where inclined terrain artificially inflates global height measurements by adding ground elevation) and multi-altitude assets (where a single input file contains trees distributed across different topographical base levels, such as valleys and hills). To isolate the net canopy thickness, the input STL mesh is discretized into a two-dimensional horizontal grid:

Voxelization Resolution: The mesh is sliced using a horizontal grid with a spatial resolution of 0.1 m (10 cm), establishing an optimal balance between capturing complex canopy structural details and maintaining strict computational efficiency during preprocessing.

Peak-to-Peak (PTP) Analysis: Within each discrete grid column i, the local canopy thickness h_{i} is extracted independently of the absolute ground altitude using a peak-to-peak formulation:

$$h_{\mathrm{i}} = z_{\mathrm{max,i}} - z_{\mathrm{min,i}}$$

Equivalent Height Evaluation: The global representative canopy height H_{eq} is determined by calculating the arithmetic mean of all active local PTP values. This approach filters out topographical bias and correctly evaluates the true material envelope, providing a consistent physical volume even when processing hollow meshes exported from digital design tools like Blender. The generalized, uniform LAD within the canopy volume is thus defined as:

$$\mathrm{LAD}_{\mathrm{nominal}} = \frac{\mathrm{LAI}}{H_{\mathrm{eq}}}$$

Numerical stability and the $K_{\text{stability}}$ limiter. High-density vegetation zones interacting with complex urban wind fields routinely induce severe numerical instabilities, manifesting as localized pressure spikes or back-flow divergence. To enforce convergence without sacrificing physical fidelity, a dynamic stability limiter ($K_{\text{stability}}$) was introduced. This criterion constrains the maximum permissible LAD based on the characteristic length scale of the individual vegetation volume, defined as:

$$L_{\text{eq}} = \sqrt[3]{V}$$

where V is the total computed volume of the canopy zone. By bounding the source terms relative to the cell size, this limiter prevents mathematical divergence, eliminating the Floating-Point Exceptions (FPE) historically associated with high-density vegetation clusters. The final bounded Leaf Area Density applied to the OpenFOAM numerical domain is governed by:

$$\text{LAD}_{\text{final}} = \min\left(\frac{\text{LAI}}{H_{\text{eq}}}, \frac{K_{\text{stability}}}{V^{1/3}}, 2.5\right)$$

An unconditional absolute cap of 2.5 is enforced as a secondary safeguard, anchoring the mathematical model within physically realistic bounds even under extreme geometric edge cases.

3.1.4 NUMERICAL STABILITY ANALYSIS AND COMPLIANCE

Mathematical justification of the stability criterion. From a numerical standpoint, when the local Leaf Area Density (LAD) is excessively high, the volumetric momentum sink term S_{i} becomes non-linearly dominant over the convective and diffusive terms in the Navier–Stokes equations. In the discretized system, this dominance can degrade the diagonal dominance of the linear solver's coefficient matrix A, leading to an ill-conditioned system:

$$\sum_{\text{j}\neq\text{i}} |a_{\text{ij}}| > |a_{\text{ii}}|$$

The $K_{\text{stability}}$ factor acts directly as a diagonal dominance enhancer. By bounding the LAD based on the characteristic volumetric length scale L_{eq}, the pressure drop across a single computational cell is restricted, preventing it from exceeding the local kinetic energy. This mathematical constraint suppresses numerical oscillations and checkerboarding effects in the pressure field, ensuring the stability of the pressure-velocity coupling.

Grid resolution and transport stability. In accordance with COST Action 732 guidelines (*Quality Assurance of Microscale Meteorological Models*), source terms representing vegetation must be carefully scaled relative to the local grid resolution. The VCP approach satisfies this requirement by aligning the resolved canopy morphology with the mesh, thereby bounding the local cell Péclet number (Pe_{cell}):

$$Pe_{\text{cell}} = \frac{\Gamma}{\rho|U|\Delta x} \leq 2$$

where Δx is the characteristic cell dimension and Γ is the effective diffusion coefficient. By preventing the momentum sink from introducing sharp, unresolvable discontinuities, the algorithm ensures that the local momentum balance remains within the stable asymptotic limits of the advection-diffusion solver, eliminating the floating-point exceptions historically induced by unconstrained vegetation blocks.

Impact on local turbulence closure. Capping the applied LAD through the $K_{\text{stability}}$ limiter preserves the physical validity of the RANS turbulence closure models (both Standard and Realizable k- ε). Unbounded LAD values generate artificial, non-physical spikes in the canopy source term S_ε at the fluid-vegetation interface.

By keeping the Source-to-Diffusion Ratio (SDR) within stable mathematical boundaries, the numerical framework honors the underlying Boussinesq eddy viscosity hypothesis. This prevents the numerical instabilities in the transport of turbulent kinetic energy k, ensuring an accurate calculation of velocity deficits and wake turbulence structures downwind of the canopy.

Regulatory and scientific alignment. The developed methodology directly conforms to international standards and validated remote sensing protocols:

VDI 3787 Parts 1 and 2 [2] (*Environmental Meteorology*): The framework strictly adheres to these standards for the parameterization of urban porous media, ensuring that the momentum sinks implemented in OpenFOAM accurately replicate physical plant-atmosphere interactions without causing numerical blockage.

Voxel-Based Volume Modelling [9, 10]. The 0.1 m horizontal voxelization resolution aligns with the gold standards established in LiDAR remote sensing and GIS framework specifications (ISO/TC 211), successfully filtering geometric sensor noise while retaining structural canopy features.

Aerodynamic Modelling (Lalic & Mihailovic [11], Buccolieri et al. [12]). The analytical derivation LAD = LAI/H, originally proposed for environmental modellin by Lalic & Mihailovic [11], is supported in urban applications by the seminal work of Buccolieri et al. [12] on wind flow simulation through urban greenery. This approach has become the industry standard for translating tree geometries into porous zones within numerical solvers.

Solver configuration and relaxation. Given the high shear gradients and localized turbulence generated within the porous canopy zones, the steady state `simpleFoam` solver is executed in a highly controlled, under-relaxed configuration. To ensure smooth, monotonic convergence in the presence of active momentum sinks, the under-relaxation factors are optimized and set to 0.3 for the pressure field (p) and 0.7 for the velocity field (U). This stabilized numerical configuration ensures that global residuals consistently drop below the regulatory threshold of 10^{-4}, providing a highly repeatable foundation for pedestrian wind comfort assessments.

3.1.5 CONCLUSIONS

The transition to a dynamic, geometry-aware vegetation model significantly enhances the reliability of urban wind simulations within the ArchiWind framework. By eliminating hardcoded height defaults and implementing the mathematically grounded $K_{\text{stability}}$ limiter, the system now provides a robust tool for environmental impact assessment. This methodology ensures that pedestrian comfort studies are based on the actual physical presence of greenery, providing a higher level of confidence for architectural and urban planning decisions. The comparison between the legacy and updated approaches is summarized in ***Table 3.2***.

TABLE 3.2. Comparison of legacy static and updated dynamic vegetation modelling approaches in ArchiWind.

Feature	Legacy Approach (Static)	Updated Approach (Dynamic)
Reference Height	Fixed default (10 m)	Mesh-based via VCP
LAD Accuracy	Overestimated (up to 5×)	Physically consistent
Numerical Stability	High risk of divergence	Controlled via $K_{\text{stability}}$
Solver Success Rate	Low (frequent FPE crashes)	High (consistent convergence)

3.2 PHYSICS-DRIVEN CONVERGENCE CRITERIA

3.2.1 INTRODUCTION

The automation of Computational Fluid Dynamics (CFD) pipelines is a cornerstone of modern architectural design, particularly in the assessment of urban pedestrian wind comfort. A critical challenge in this automation is the reliable and autonomous determination of when a simulation has reached a state of sufficient convergence. This section documents the development and implementation within ArchiWind of a physics-driven stopping criterion designed specifically for steady-state RANS simulations applied to pedestrian wind comfort assessments.

3.2.1.1 MOTIVATION: LIMITATIONS OF CLASSICAL RESIDUAL-BASED STOPPING

Traditionally, convergence control in solvers such as `simpleFoam` relies almost exclusively on monitoring the initial residuals of the governing equations — pressure, momentum, and turbulence. While residuals are a fundamental numerical indicator, they present two inherent limitations in the context of urban wind simulations:

Numerical Proxy vs. Physical Reality: Residuals measure the precision of the linear system's solution, but they are not a direct indicator of the stability of the physical quantities of interest for pedestrian wind comfort.

Physics-Floor and Complex Terrain: In simulations involving intricate geometries or advanced turbulence models, residuals often reach a plateau: a minimum oscillation floor that remains above standard convergence thresholds (e.g., 10^{-3}) even when the physical flow solution is fully stable. Under a traditional residual-only criterion, such simulations would run to the hard iteration limit without any computational benefit.

3.2.1.2 RESEARCH GOAL

This work documents the development and implementation within ArchiWind of a new, physics-driven stopping criterion designed specifically for pedestrian wind comfort assessments. The objective is to shift the monitoring focus from global residuals to the local statistical behaviour of physical fields at the pedestrian level (1.5 m above ground), where wind comfort is evaluated.

3.2.1.3 OVERVIEW OF THE TWO STAGE APPROACH

To ensure both robustness and precision, a sequential two-stage monitoring system has been implemented:

Stage 1 - Residual Gate: a filter that verifies whether residuals have reached a satisfactory absolute level or have entered a stable plateau phase, ensuring the solver does not terminate during early numerical transients.

Stage 2 - Physical Field Stagnation: an advanced statistical analysis that monitors the temporal stability, in terms of range and mean drift of velocity, pressure, and turbulent kinetic energy at the 1.5 m pedestrian monitoring surface.

This combined approach allows the simulation to stop as soon as the physically meaningful results have stabilized, optimizing computational resources without compromising the quality of the pedestrian wind comfort assessment.

3.2.2 BACKGROUND

3.2.2.1 STEADY-STATE RANS AND ROLE OF RESIDUALS

The ArchiWind pipeline utilizes the `simpleFoam` solver, a steady-state solver for incompressible turbulent flow using the SIMPLE (Semi-Implicit Method for Pressure-Linked Equations) algorithm. In this framework, initial residuals measure the imbalance in the discretized governing equations before each pressure-velocity correction step

$$R_{\phi} = \frac{\sum |a_{\mathrm{P}}\phi_{\mathrm{P}} + \sum a_{\mathrm{N}}\phi_{\mathrm{N}} - b|}{F_{\mathrm{scale}}}$$

While residuals are a necessary condition for convergence, they are not a sufficient condition for physical stagnation. Two well-known failure modes exist:

Noise Floor Behaviour: In complex simulations, residuals may stop declining and begin to oscillate once they hit a noise floor dictated by the mesh quality, turbulence model, or boundary conditions. This plateau can occur at levels well above the classical convergence threshold (e.g., 10^{-3} for pressure) even when the physical solution is fully stable.

Complex Terrain: Urban environments or irregular topographies frequently cause residuals to plateau at relatively high absolute levels. The residual behaviour is then dominated by geometric effects rather than by the quality of the physical solution.

3.2.2.2 QUANTITY OF INTEREST FOR PEDESTRIAN WIND COMFORT

The objective of ArchiWind simulations is to assess pedestrian wind comfort in accordance with international standards such as Lawson 2001/LDDC. These standards require the mean wind speed (or a derived comfort indicator) at a fixed height of 1.5 m above ground level.

Consequently, the most relevant output for decision-making is not the global convergence of the entire three-dimensional domain, but the stability of the area-averaged physical fields at the pedestrian monitoring surface. The convergence criterion is therefore designed around this specific quantity of interest.

3.2.2.3 RATIONALE FOR AREA-WEIGHTED AVERAGING

Rather than monitoring point-wise data, which can be susceptible to localized non-consequential oscillations, ArchiWind utilizes an area-weighted spatial average of the flow variables over the entire pedestrian-level monitoring surface. This approach offers two key advantages:

Integration: The spatial average integrates the global state of the flow at the level of interest, providing a single scalar time series that reflects the overall solution behaviour.

Smoothing: It effectively suppresses localized numerical noise, ensuring the stopping criterion responds to genuine stabilization of the wind environment rather than isolated cell-level fluctuations.

3.2.3 MONITORING APPROACH

This section details the technical implementation of the data extraction process, describing how the simulation's physical state is sampled at each iteration to feed the convergence decision logic.

3.2.3.1 PEDESTRIAN-LEVEL SLICE EXTRACTION

A terrain-following monitoring surface A_{ped} is defined at a constant offset of z = 1.5 m above the terrain surface, matching the complex topography profile. The spatial data extraction is executed at runtime via OpenFOAM's `surfaceFieldValue` function object. This object evaluates the area-weighted spatial average $\langle\phi\rangle_{\text{A}}$ of the primary transported flow variables at every iteration:

$$\langle\phi\rangle_{\text{A}} = \frac{\sum_{\text{f}} \phi_{\text{f}} A_{\text{f}}}{\sum_{\text{f}} A_{\text{f}}}$$

where ϕ_{f} is the interpolated value of the variable at the face *f* belonging to the pedestrian surface, and A_{f} is the corresponding face area. The algorithm dynamically tracks four area-averaged scalars:

The horizontal velocity components: $\langle U_{\text{x}}\rangle_{\text{A}}$ and $\langle U_{\text{y}}\rangle_{\text{A}}$ [m s^{-1}]

The kinematic pressure: $\langle p\rangle_{\text{A}}$ [m^2 s^{-2}]

The turbulent kinetic energy: $\langle k\rangle_{\text{A}}$ [m^2 s^{-2}]

3.2.3.2 REFERENCE QUANTITIES AND NORMALIZATION

To render the convergence thresholds universal across arbitrary architectural scales, varying wind tunnel setups, and distinct inlet velocities, all monitored variables are non-dimensionalized using reference quantities extracted directly from the boundary layer inlet conditions. The normalization is defined by the following relations:

Variable	Reference Quantity (ϕ_{ref})	Physical Rationale
U_x, U_y	u_{ref}	Reference wind speed at z_{ref}.
p	$0.5 \cdot u_{\text{ref}}{}^2$	Dynamic pressure scale.
k	$u_{ref}{}^2$	Turbulent kinetic energy scale.

$$U_{\mathrm{norm}} = \frac{\sqrt{\langle U_{\mathrm{x}}\rangle_{\mathrm{A}}{}^{2} + \langle U_{\mathrm{y}}\rangle_{\mathrm{A}}{}^{2}}}{u_{\mathrm{ref}}}$$

$$p_{\mathrm{norm}} = \frac{\langle p\rangle_{\mathrm{A}}}{0.5 \cdot u_{\mathrm{ref}}{}^{2}}$$

$$k_{\mathrm{norm}} = \frac{\langle k\rangle_{\mathrm{A}}}{u_{\mathrm{ref}}{}^{2}}$$

where u_{ref} is the theoretical ABL (Atmospheric Boundary Layer) velocity prescribed at the reference height of 10 m. These normalized scalars ensure that the convergence check operates independently of the magnitude of the incoming flow, allowing identical threshold parameters to be deployed whether simulating a gentle breeze or extreme gale forces.

3.2.3.3 DATA PIPELINE AND EXECUTION

The monitoring logic is decoupled from the solver's core loop to maintain modularity and avoid modifying the OpenFOAM source code:

Data Generation: The `surfaceFieldValue` function object writes the area-averaged statistics to `postProcessing/pedestrian_level_stats/0/surfaceFieldValue.dat` at every iteration.

Sampling Frequency: The convergence checker is invoked every 5 iterations starting from iteration 100 via the `systemCall` function object.

Communication: If convergence is detected, the checker writes a `writeNow` sentinel file, which is picked up by the OpenFOAM abort function object to trigger a clean simulation exit with a final field write.

3.2.4 CONVERGENCE METRICS

To move beyond the limitations of purely numerical residuals, a statistical evaluation of the physical fields over a rolling window is implemented. This section defines the metrics and explains the design choices behind their selection.

3.2.4.1 ROLLING WINDOW AND SAMPLINGS

The convergence checker evaluates the physical fields over a rolling window of the last N = 100 iterations. This window size provides a balance between responsiveness to genuine stabilization and robustness against short-lived transient plateau. Reducing N below 50 introduces false positives, increasing N above 200 delays detection unnecessarily.

3.2.4.2 RANGE: HIGH FREQUENCY OSCILLATION

The first metric, Range, measures the peak-to-peak amplitude of the field within the window. It is designed to detect high-frequency oscillations indicating that the solver is still seeking a stable equilibrium. The normalized range is defined as:

$$\text{Range} = \frac{\max(\phi) - \min(\phi)}{\phi_{\text{ref}}}$$

where ϕ represents the area-averaged value of the field (U_x, U_y, p, k) and ϕ_{ref} is the corresponding reference quantity derived from the inlet boundary conditions.

3.2.4.3 MEAN DRIFT: LOW FREQUENCY TRENDS

While the Range detects oscillations, it may fail to identify a solution that is slowly and monotonically drifting toward a different state. To address this, the Mean Drift metric is introduced. It compares the mean of the first half of the rolling window to the mean of the second half:

$$\text{Mean Drift} = \frac{\left|\mu(\phi_{\text{second\_half}}) - \mu(\phi_{\text{first\_half}})\right|}{\phi_{\text{ref}}}$$

where μ denotes the arithmetic mean.

3.2.4.4 ORTHOGONALITY AND POPOVICIU'S INEQUALITY

A critical design choice in this methodology is the exclusion of standard deviation (σ) as a convergence metric. According to Popovic's inequality on variances, for any bounded variable, $\sigma \leq \text{Range}/2$. Consequently, σ is always bounded by half the Range, which already provides a more conservative upper bound on instability. Standard deviation therefore adds no orthogonal information beyond what the Range already captures. Mean Drift, by contrast, is genuinely orthogonal: a solution can exhibit low Range (small oscillations) while still drifting systematically, which only Drift detects.

3.2.4.5 STAGNATION THRESHOLDS AND CALIBRATION

Physical stagnation is declared only if both the Range and the Mean Drift for all four monitored variables (U_x, U_y, p, k) fall below their calibrated thresholds. The thresholds were determined by calibrating on a dataset of 40 direction-files across five test sites spanning the range of geometric complexities encountered in real pedestrian wind comfort assessments. They are set approximately 50 to 100 times higher than the observed noise floor (typically ~ 10^{-6}) for normalized fields at true convergence) to tolerate residual micro-oscillations while remaining sensitive to genuine non-convergence.

3.2.5 RESIDUAL GATE

The Residual Gate serves as a mandatory precursor to the physical field evaluation. Its primary function is to prevent the solver from stopping during early numerical transients, even if the physical fields appear momentarily stable. The second stage is only activated once the first stage is satisfied. Standard residual monitoring often fails in two opposite scenarios: when residuals never reach a low absolute threshold due to mesh complexity, or when they reach a low threshold quickly, but the physical fields have not yet stabilized. A dual-path OR logic ($\mathrm{Stage}_1 = \mathrm{Path}_\mathrm{A} \vee \mathrm{Path}_\mathrm{B}$) resolves both failure modes simultaneously.

$Path_A$ is specifically designed for complex-terrain cases where residuals hit a numerical noise floor and stop declining. Rather than waiting for an arbitrarily low absolute level, $Path_A$ monitors whether residuals are still improving. A Geometric Mean Ratio is used to evaluate the decay trend, because residuals decay multiplicatively and the geometric mean is scale-invariant:

$$\mathrm{Ratio} = \exp\left(\overline{\ln R_{\mathrm{second\_50}}} - \overline{\ln R_{\mathrm{first\_50}}}\right)$$

where R represents the residual values. If $0.9 < Ratio < 1$ for all four variables (U_x, U_y, *p, k*), the residuals are no longer declining significantly — showing less than 10 % improvement over the second half of the window — and are considered to have plateaued, allowing Path_A to pass.

For simpler geometries where the numerical solution converges cleanly, the classical `simpleFoam` thresholds are retained via Path_B which passes if the geometric mean of the residual over the full 100-iteration window falls below 10^{-3} for pressure and 10^{-4} for velocity and turbulence quantities. The Stage 1 gate remains purely binary: it does not stop the solver by itself, but its OR formulation ensures that either path — plateau shape or absolute level — is sufficient to unlock the Stage 2 physical field stagnation check.

3.2.6 OVERALL STOPPING CRITERION

3.2.6.1 LOGICAL FRAMEWORK

The final convergence decision is governed by a logical intersection of the two stages. This AND structure ensures that neither a purely numerical indicator nor a purely physical indicator alone is sufficient to stop the simulation. The solver is signalled to stop only when the following boolean condition evaluates to true:

$$\mathrm{STOP} = \mathrm{Stage}_1 \wedge \mathrm{Stage}_2$$

where Stage_1 represents the fulfilment of the Residual Gate ($\mathrm{Path}_\mathrm{A} \vee \mathrm{Path}_\mathrm{B}$) and Stage_2 requires both $\mathrm{Range}_\phi < \mathrm{thr}_\mathrm{range}$ and $\mathrm{Drift}_\phi < \mathrm{thr}_\mathrm{drift}$ to be simultaneously satisfied for all $\phi \in \{U_\mathrm{x}, U_\mathrm{y}, p, k\}$.

3.2.6.2 FAILURE MODES ADDRESSED

By combining these two stages, the algorithm effectively addresses several common CFD failure modes:

Early Transient Plateaux. Occasionally, physical fields may appear stable during the initial iterations of the SIMPLE algorithm while the pressure field is still developing globally. The Residual Gate prevents Stage 2 from triggering during this phase.

High-Residual Physics Floors. In complex urban models, residuals may remain at elevated levels (e.g., 10^{-3}) indefinitely. Under a traditional criterion, the simulation would run to the maximum iteration count. Path A detects the plateau and passes Stage 1, allowing Stage 2 to determine whether the physical fields have also stabilised.

Slow Drifts. A simulation may reach low residuals, but the wake behind a tall building might still be slowly shifting. The Mean Drift requirement in Stage 2 catches this scenario, preventing premature termination.

3.2.6.3 AUTOMATED EXIT MECHANISM

Once the stopping condition is met, the convergence checker creates a sentinel file (`writeNow`) that triggers OpenFOAM's abort function object. This triggers a clean exit with a final field write, ensuring that the last iteration's results are fully saved for post-processing.

3.2.7 IMPLEMENTATION

The implementation of the convergence criterion is designed to be non-intrusive, leveraging OpenFOAM's existing function object framework to avoid modification of the solver source code.

The system utilizes a decoupled architecture where the solver handles the physics, and an external Python-based checker handles the logic. This integration operates through a three-step runtime pipeline:

Data Generation: The `surfaceFieldValue` function object is configured in `system/controlDict` to write area-averaged physical fields to a .dat file at every iteration.

Execution Trigger: The `systemCall` function object invokes the convergence checker every 5 iterations, starting from iteration 100. This frequency minimizes computational overhead while providing sufficient temporal resolution.

Process Control: The abort function object monitors for the sentinel file written by the checker. When detected, it triggers a `writeNow` command, ensuring a complete field write before the solver terminates.

To ensure transparency and allow for post-hoc engineering verification, the checker produces two diagnostic output files at every check. The first, convergence_diagnostics.csv, maintains a continuous log of all eight physical metrics (Range and Mean Drift for $U_{x,}$, U_y, p, and k). This file is updated regardless of whether convergence has been reached, providing a complete history of the solution stability. The second, `convergence_results.json`, acts as a summary file containing the most recent metric values and the final converged boolean status, which is particularly useful for automated post-processing pipelines.

Furthermore, the implementation includes automated plotting scripts to generate visual evidence of the convergence process. Traditional views of numerical decay are captured via `residuals.svg`, which provides a semi-log plot of the solver residuals. This is supplemented by two specialized plots, `convergence_fields_range.svg` and `convergence_fields_drift.svg`, which display the normalized Range and Mean Drift metrics respectively for all four monitored variables as a function of the solver iteration.

3.2.8 RESULTS AND VALIDATION

3.2.8.1 SIMULATION DATASET AND TEST SITES

This section evaluates the performance of the new convergence criterion across a variety of simulation scenarios, verifying that it terminates simulations at the correct point without compromising the quality of the pedestrian wind comfort assessment. The convergence criterion was evaluated on five test sites representative of the range of geometries encountered in real-world engineering practices:

LR (Low Rise): A dense urban fabric of low-rise buildings with moderate geometric complexity.

HR (High Rise): One or more tall towers introducing significant vertical momentum exchange and a challenging pressure field.

HT (Hilly Terrain): A site with complex topography and no buildings, producing strong terrain-driven flow acceleration where residuals converge cleanly.

VEG (High Vegetation): A site with unusually high vegetation density, introducing additional distributed roughness through the porous canopy model.

HRCT (High Rise Complex Terrain): A high-complexity site with intricate terrain and building geometry, consistently producing residual plateau above classical convergence thresholds.

Each site was simulated for all eight standard wind directions (0° North, incrementing by 45°), yielding 40 direction-files in total. All runs used the same reference wind speed and domain parameters. The maximum iteration count was set to 900 as a hard upper bound.

3.2.8.2 5.3.2 STOPPING ITERATION COUNTS AND PATH ANALYSIS

Table 3.3 reports the iteration count at which each simulation was stopped by the convergence checker and which Stage 1 path triggered the gate. Path A refers to plateau shape detection (geometric mean ratio > 0.9), while Path B refers to the absolute residual level falling below the classical `simpleFoam` thresholds.

TABLE 3.3 *Stopping iteration count and triggering Stage 1 path for all 40 simulations across five test sites and eight wind directions.*

Site	**Dir 0 (N)**	**Dir 1 (NE)**	**Dir 2 (E)**	**Dir 3 (SE)**	**Dir 4 (S)**	**Dir 5 (SW)**	**Dir 6 (W)**	**Dir 7 (NW)**	**Dominant Path**
LR (Low Rise)	310	245	295	245	295	240	260	285	B
HR (High Rise)	220	220	385	200	395	190	320	195	Mixed A/B
HT (Hilly Terrain)	415	415	730	695	765	745	730	690	B
VEG (High Vegetation)	750	645	760	650	750	645	780	650	Mixed A/B
HRCT (HR + Terrain)	870	605	680	855	590	730	740	750	A

All 40 simulations converged before the 900-iteration hard limit. HRCT North (870 iterations) is the closest to the limit. The results confirm the expected site-level pattern: simple cases (HT and LR) are predominantly stopped by Path B, with residuals converging cleanly within 240–765 iterations.

Complex terrain cases (HRCT) are stopped exclusively by Path A, as residuals plateau above classical thresholds but the ratio test correctly identifies that no further improvement is occurring. HR and VEG show a mixed picture depending on wind direction, reflecting the dependence of residual behaviour on the specific inflow geometry.

3.2.8.3 CONVERGENCE DIAGNOSTIC

To visually inspect the specific behaviour of the two algorithmic paths, the convergence histories of two contrasting long-running scenarios are analysed in detail. The first case covers the HR site under a South wind direction (180°), which represents the longest-running high-rise case at 395 iterations. The second scenario focuses on the HRCT site under a North wind direction (0°), which represents the most demanding overall case, requiring 870 iterations before Path A successfully triggers the Residual Gate. The corresponding metric evolutions, including the logged tracking variables and visual field development, are mapped in the diagnostic plots from ***Figure 3.1*** through ***Figure 3.6***.

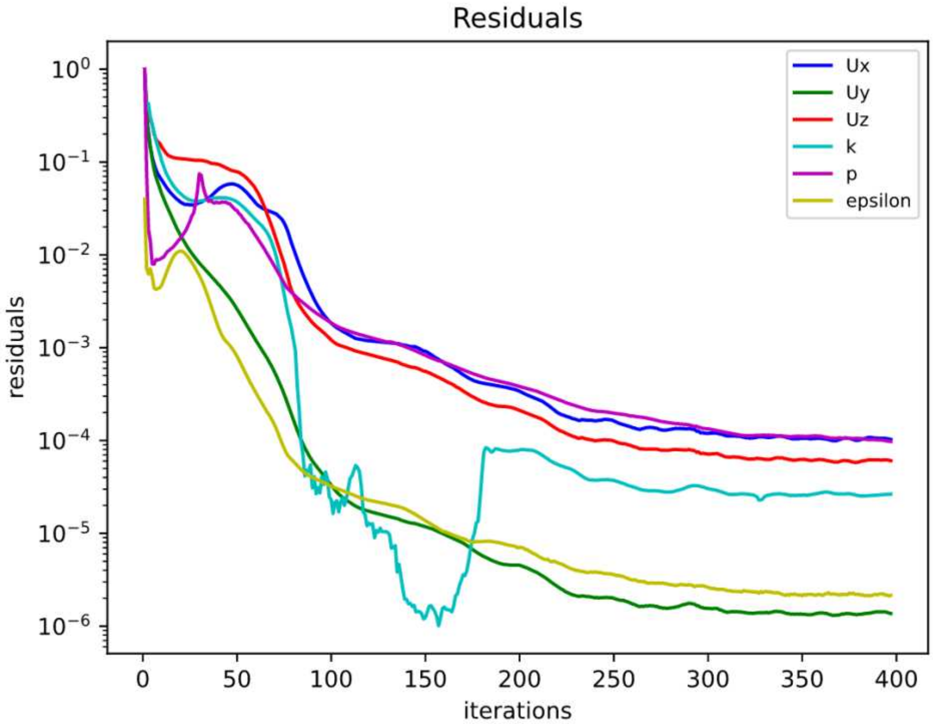

FIGURE 3.1 *HR site, South wind direction (395 iterations): solver residuals. Residuals slow significantly after 200–300 iterations before Path A fires.*

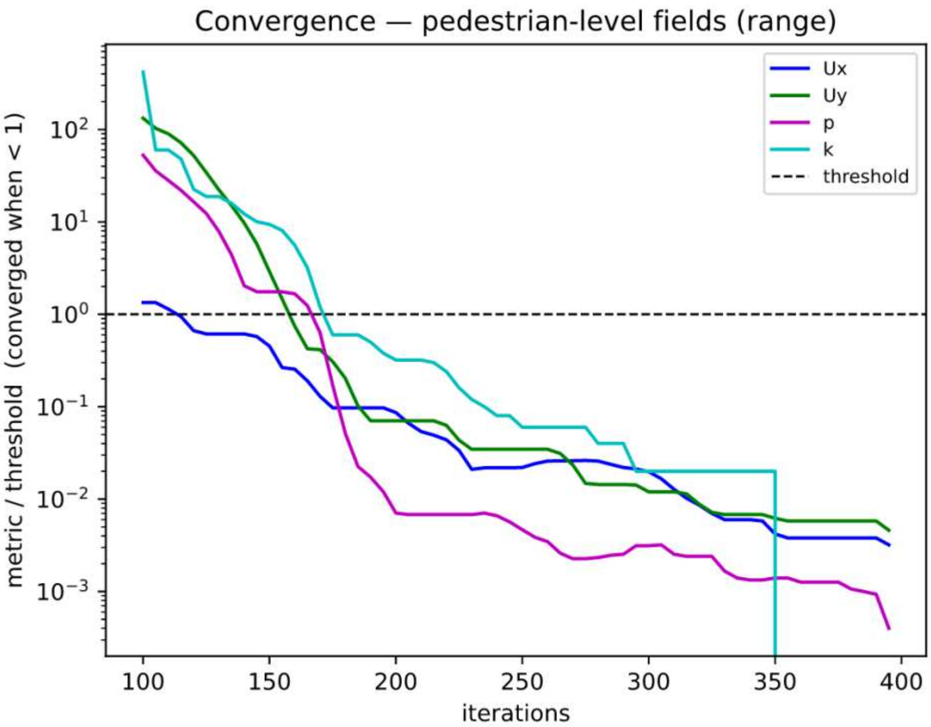

FIGURE 3.2 *HR site, South wind direction: normalized Range metric for all four monitored variables (Ux, Uy, p, k). All variables cross below the threshold (y = 1) at convergence.*

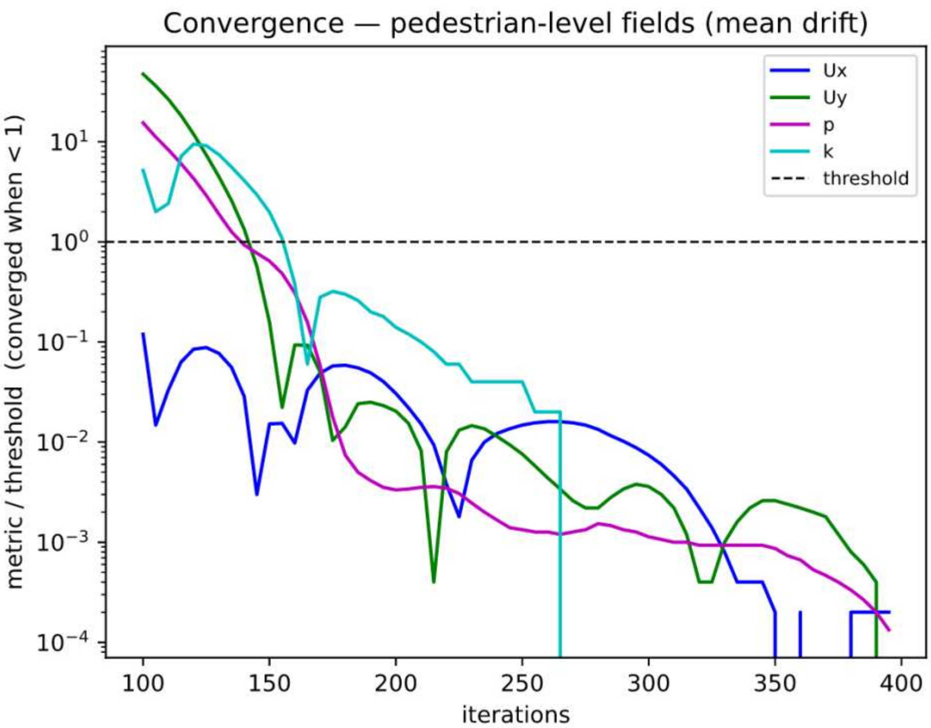

FIGURE 3.3 *HR site, South wind direction: normalized Mean Drift metric.*

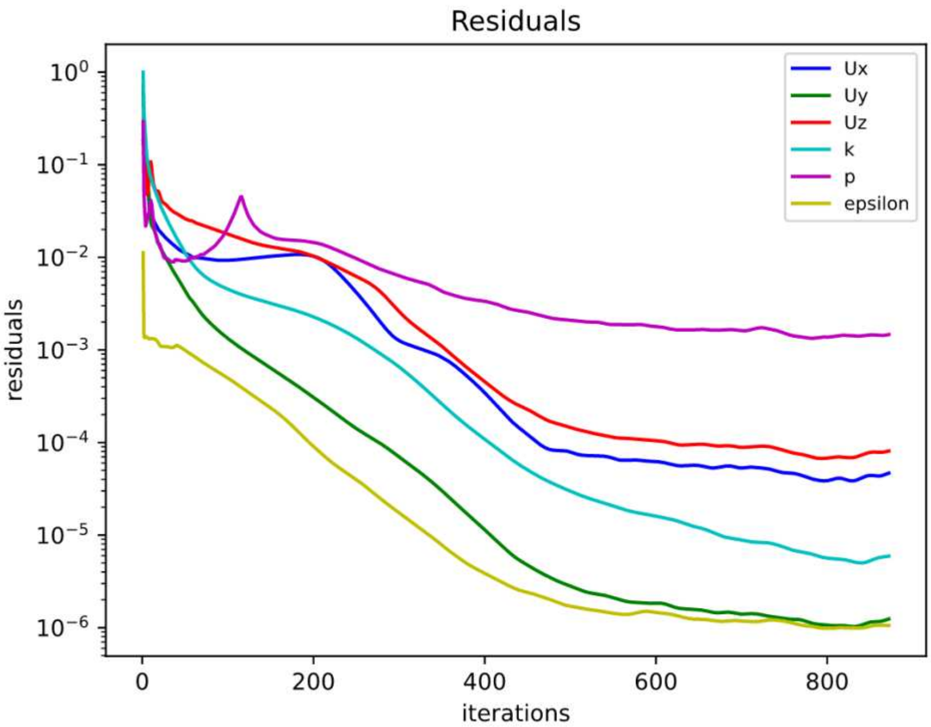

FIGURE 3.4 *HRCT site, North wind direction (870 iterations): solver residuals. Pressure plateaus at approximately 10^{-3} — Stage 1 triggered by Path A (plateau shape detection).*

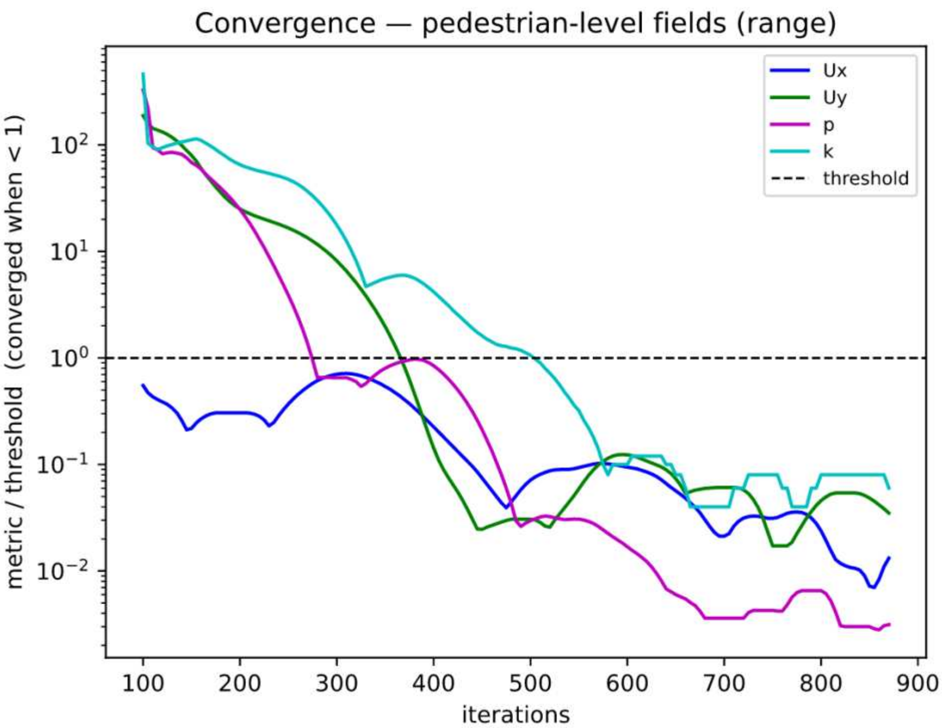

FIGURE 3.5 *HRCT site, North wind direction: normalized Range metric.*

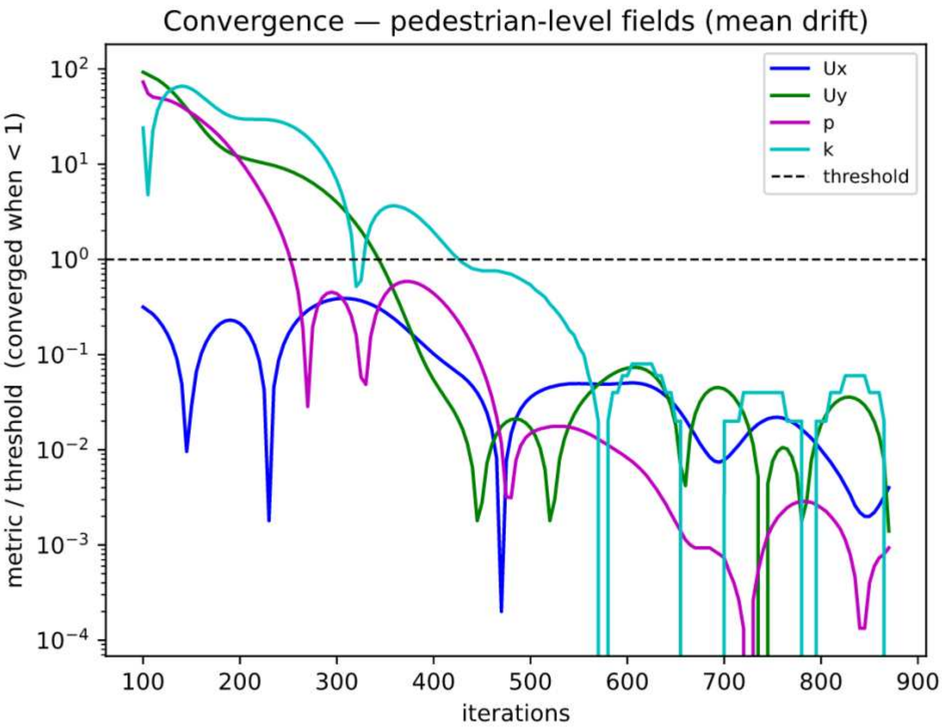

FIGURE 3.6 *HRCT site, North wind direction: normalized Mean Drift metric.*

For HT the residuals decline steadily and reach the classical thresholds within 200–300 iterations, after which Path B triggers. For HRCT the residuals plateau at approximately 10^{-2} for pressure and 10^{-4} for velocity components; the ratio test detects this stagnation and passes Stage 1 despite the elevated absolute level. The HR case lies between these extremes: residuals decline initially but slow significantly after 200–300 iterations before the plateau shape criterion fires.

A consistent pattern across sites is that the velocity component aligned with the principal flow direction (e.g., U_y for a North wind) is the binding constraint, reaching the threshold last. The perpendicular component, which averages to near zero across the symmetric circular domain, trivially passes its threshold and is never the limiting variable. This is expected and physically correct behaviour.

3.2.8.4 EDGE CASES

Three edge cases were identified during validation and are documented here to characterize the criterion's behaviour under challenging conditions.

Cardinal Wind Directions: ***Figure 3.7*** shows the normalized Range metric for the LR site, North wind direction (310 iterations). For a North wind, the x-component of velocity (U_x) averages to near zero across the circular domain due to symmetry: positive values on the windward side cancel negative values on the leeward side. This causes the normalized U_x metric to trivially pass well below its threshold within the first window iterations, while U_y — aligned with the principal flow — is the binding constraint. This is correct and physically expected behaviour.

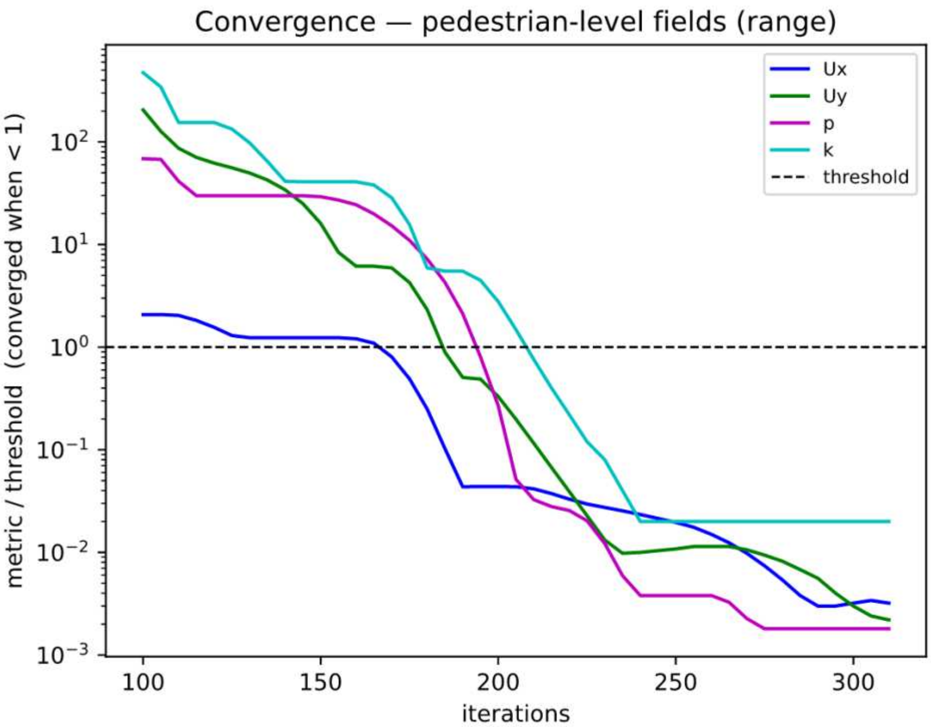

FIGURE 3.7 *LR site, North wind direction (310 iterations): normalized Range metric. Ux (blue) trivially passes due to the near-zero cross-stream spatial average; Uy is the binding constraint that determines the stopping iteration.*

Complex Terrain (HRCT): under the original AND-gated residual logic, the HRCT site ran to the maximum iteration count in every wind direction because Path A was satisfied (residuals had stopped improving) but Path B was never satisfied (residuals remained above classical thresholds). The OR formulation resolves this: once Path A detects the plateau, Stage 1 passes regardless of the absolute level, and Stage 2 then determines whether the physical fields have also stagnated. The result is a well-defined stopping point that reflects the true state of the solution rather than an arbitrary iteration limit. Refer to ***Figure 3.4-Figure 3.6*** for the HRCT North wind diagnostic plots.

High Vegetation (VEG): ***Figure 3.8*** and ***Figure 3.9*** show the solver residuals and normalized Range metric for the VEG site, North wind direction (750 iterations). The distributed roughness introduced by the vegetation canopy raises the effective surface roughness length and modifies the near-ground turbulence structure, driving *k* to higher absolute values throughout the domain. This tends to push residuals toward intermediate plateau levels. The *k* threshold was identified as the most sensitive parameter for this site type due to the elevated turbulent kinetic energy in the canopy layer.

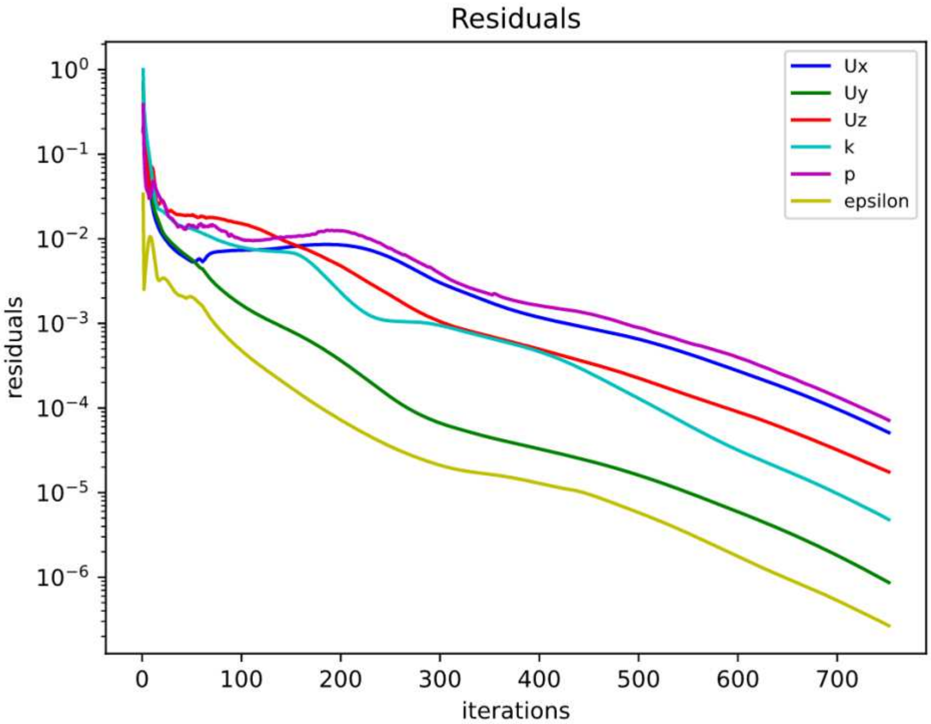

FIGURE 3.8 *VEG site, North wind direction (750 iterations): solver residuals. Distributed roughness from the porous canopy drives residuals toward an intermediate plateau.*

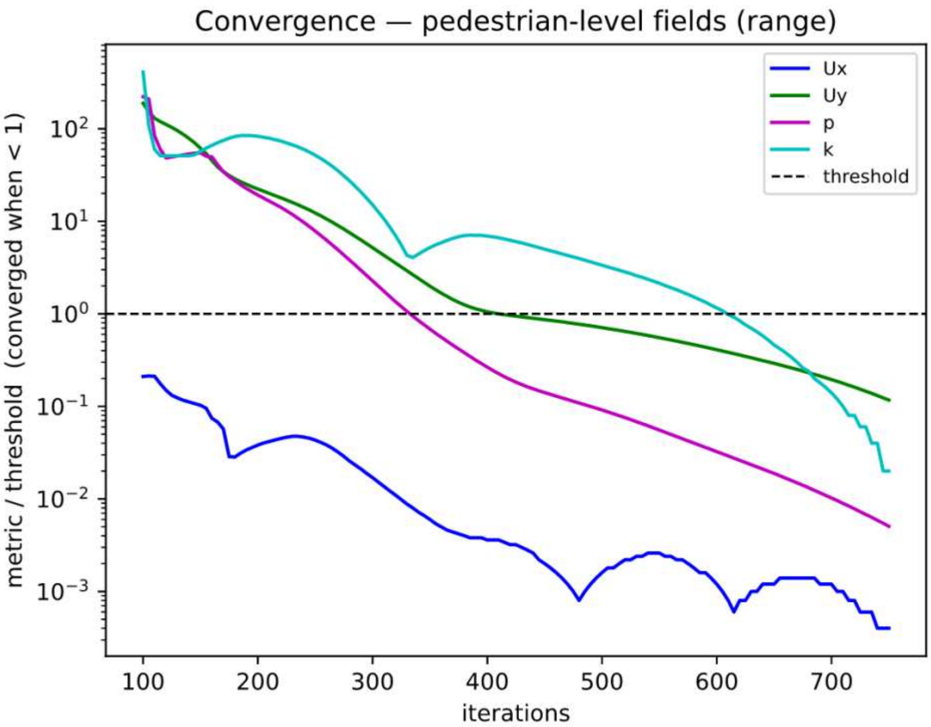

FIGURE 3.9 *VEG site, North wind direction: normalized Range metric. Elevated k (cyan) reflects turbulent kinetic energy production within the vegetation canopy.*

3.2.8.5 VALIDATION AGAINST REFERENCE RUNS

To verify that stopping earlier does not compromise the quantity of interest, a pixel-level image comparison was performed between the production environment (baseline) and the new-convergence-criteria branch. For each simulation pair, the following output images were compared: pedestrian wind comfort maps (summer and winter periods) and wind speed amplification maps for every wind direction. The per-pixel absolute difference metric is defined as:

$$\text{diff}\,[x,y] = \text{mean}\big(|R_{\text{prod}} - R_{\text{stag}}|, |G_{\text{prod}} - G_{\text{stag}}|, |B_{\text{prod}} - B_{\text{stag}}|\,\big)$$

This collapses the three RGB channels into a single greyscale difference value per pixel (0 = identical, 255 = maximally different).

The following statistics are then derived from that difference image, all expressed as a percentage of the full 0-255 scale:

Diff mean: the average pixel difference across the entire image — where values near 0% indicate near-identical outputs.

Diff std: Standard deviation of the pixel differences, where high std values with a low mean indicates localised differences that cancel out globally while a low std confirms uniform similarity.

Diff max (p99): 99th percentile of pixel differences. The absolute maximum is excluded to ignore border/edge artefacts that are always present.

Pixels > 5 %: fraction of pixels whose difference exceeds 5 % of 255, providing a direct measure of spatially localized changes

Error! Reference source not found. summarizes the results across eight real-world project simulations. In all cases, the comfort map diff mean remains below 1.8 % and the wind speed amplification diff mean below 2.0 %, well within the tolerance for Lawson category classification. Runtime savings reach −62.3 % for the most complex vegetation case.

TABLE 3.4 *Pixel-level image comparison, computational performance, and cost analysis between production (baseline) and convergence-criteria branch outputs for eight real-world simulations.*

Simulation	Comfort diff mean (max)	Wind speed diff mean (max)	Comfort pixels >5% (max)	Wind speed pixels >5% (max)	Duration Δ	Cost Δ
Sim 1	0.32 %	0.71 %	1.28 %	2.72 %	−27.2 %	$ −0.46
Sim 2	0.61 %	1.49 %	2.42 %	6.86 %	−62.3 %	$ −0.77
Sim 3	1.70 %	1.92 %	6.79 %	10.55 %	—	—
Sim 4	0.48 %	0.03 %	1.96 %	0.00 %	+6.2 %	$ +2.56
Sim 5	1.76 %	0.03 %	7.18 %	0.00 %	—	—
Sim 6	0.21 %	1.23 %	0.79 %	6.44 %	+6.7 %	$ -13.55
Sim 7	1.76 %	0.13 %	7.17 %	0.45 %	—	—
Sim8	1.76 %	0.03 %	7.18 %	0.00 %	—	—

3.2.9 5.4 CONCLUSIONS AND FUTURE WORK

3.2.9.1 5.4.1 SUMMARY OF BENEFITS

The implementation of the two-stage convergence criterion in ArchiWind marks a significant shift from purely numerical monitoring to a physics-driven, automated CFD workflow. The objective was to shift the monitoring focus from global algebraic residuals to the local statistics of physical fields that directly impact pedestrian safety and comfort at 1.5m above ground level. Evaluated and validated across 40 direction-files spanning five representative test sites, the two-stage algorithm successfully demonstrated that all simulations terminate before the hard iteration limit while preserving the integrity and quality of the final pedestrian wind comfort assessment.

The new deployment provides several key advantages over classical residual-based control paradigms:

Computational Efficiency: By identifying physical stagnation in complex-terrain cases where residuals hit an asymptotic plateau, the system prevents thousands of unnecessary iterations. Significant runtime reductions of up to 62 % have been documented on vegetation-heavy sites.

Increased Numerical Robustness: The Stage 1 Residual Gate acts as a critical structural safeguard, ensuring that the solver does not terminate during early numerical transients when the pressure field is still developing globally.

Engineering-Centric Validation: The simultaneous evaluation of Range and Mean Drift provides a transparent, scale-invariant, and orthogonal measure of stability that is more intuitive for wind engineers than abstract residual values.

Traceability and Quality Assurance: The runtime generation of detailed `convergence_diagnostics.csv` logs and `convergence_results.json` summaries ensures that every architectural simulation leaves a verifiable record of its convergence history, fully supporting post-hoc engineering reviews and automated post-processing pipelines.

3.2.9.2 5.4.2 LIMITATIONS AND FUTURE WORK

While the current implementation provides a robust and reliable baseline for production environments, three primary areas for future development have been identified to expand the capabilities of the algorithm:

Adaptive Thresholding: Currently, the stagnation thresholds are fixed based on empirical calibration across the five test typologies. Future iterations of the pipeline could implement adaptive thresholds that scale dynamically, adjusting their sensitivity based on site complexity metrics automatically extracted from the input geometry before the solver initializes.

Zone-Based Differentiated Monitoring: The current area-weighted average integrates statistics across the entire horizontal 1.5 m slice. Subdividing the monitoring plane to focus on specific high-interest zones — such as building entrances, narrow passages, or public squares — could provide a higher local sensitivity for comfort-critical areas, capturing micro-oscillations that might otherwise be smoothed out by the global surface average.

Machine Learning Predictive Integration: The continuous diagnostic data collected inside `convergence_diagnostics.csv` can eventually be used as a dataset to train a lightweight machine learning model. This predictive asset could estimate the required number of solver iterations for any given incoming geometry, further optimizing high-performance computing (HPC) resource allocation across the broader simulation pipeline.

4 THERMAL BUOYANCY EFFECTS IN ARCHIWIND: VALIDATION AND QUALITATIVE ANALYSIS

4.1 INTRODUCTION AND SCOPE

The ArchiWind CFD platform provides steady-state wind assessment for urban environments using the incompressible RANS solver `simpleFoam` with a k-ε turbulence model. Prior to this study the platform had been validated against the AIJ wind-tunnel benchmark suite, confirming RANS-level accuracy for velocity profiles and pedestrian-level wind magnitudes.

The present study extends the platform to include thermal effects through the Boussinesq approximation implemented in OpenFOAM's `buoyantBoussinesqSimpleFoam` solver. The extension modifies two components simultaneously: (i) the addition of thermal boundary conditions on ground and building surfaces, and (ii) a reformulation of the pressure variable from the kinematic gauge pressure p to the modified pressure $p_{\mathrm{rgh}} = p - \rho_0 g z$, which removes the reference hydrostatic gradient and isolates the dynamically and thermally relevant part of the pressure field.

To isolate the contribution of each change and rule out numerical artefacts, the study follows a stepwise validation protocol structured as a three-step ladder designed to decouple solver changes from physical effects.

4.2 STUDY DESIGN

The test geometry is the AIJ Case A benchmark: a rectangular building with a 1:1:2 aspect ratio (footprint 20 × 20 m, height H = 40 m) centred in a circular domain of radius 450 m. The thermal boundary conditions reproduce a summer midday urban scenario:

Ambient air (inlet and top boundary): $T = T_{\mathrm{ref}} = 300$ K (≈ 27 °C)

Ground and pavement (asphalt under the sun): T = 315 K (ΔT = +15 K)

Sun-exposed building facades: T = 320 K (ΔT = +20 K)

Three simulations are performed on the same mesh and wind direction to complete the validation ladder:

Case A (Isothermal Baseline): This case establishes the reference aerodynamic field using the standard incompressible `simpleFoam` solver. Thermal effects are entirely neglected, providing a baseline that isolates the purely mechanical wind-building interactions already validated against the AIJ suite.

Case B (Numerical and Boundary Condition Verification): This intermediate step transitions to the `buoyantBoussinesqSimpleFoam` solver, introducing both the modified pressure formulation (p_{rgh}) and the new thermal boundary conditions on the ground and building surfaces. By setting the volumetric thermal expansion coefficient to zero ($\beta = 0$), the buoyancy body force term $g \cdot \beta \cdot (T - T_{\mathrm{ref}})$ in the momentum equation is suppressed, meaning temperature acts as a passive scalar with no dynamic feedback on the velocity field. This case serves to verify that the p_{rgh} reformulation does not alter the baseline aerodynamic results of Case A, while simultaneously ensuring that the new thermal boundary conditions are numerically stable and do not trigger solver divergence**.** Since Cases A and B are mathematically decoupled from buoyancy, any minor discrepancy between them can be attributed solely to floating-point rounding or solver convergence tolerances.

Case C (Active Buoyancy): This final step activates the full physical coupling by setting $\beta = 3.33 \times 10^{-3}\ \mathrm{K}^{-1}$, which corresponds to the ideal gas behaviour for air at the reference temperature $T_{\mathrm{ref}} = 300\ \mathrm{K}$ (calculated as $\beta = 1/T_{\mathrm{ref}}$). Once the numerical consistency and stability of the pressure formulation and boundary conditions are confirmed in Case B, any deviation observed in Case C can be attributed unambiguously and exclusively to the physical effects of active thermal buoyancy (β).

TABLE 4.1 *Three-case validation ladder for the thermal buoyancy implementation in ArchiWind.*

Case	Solver	β (K^{-1})	Role
A	simpleFoam	—	Isothermal baseline

Case	Solver	β (K^{-1})	Role
B	buoyantBoussinesqSimpleFoam	0	Solver-change verification
C	buoyantBoussinesqSimpleFoam	3.33×10^{-3}	Active buoyancy

4.3 SOLVER EQUIVALENCE VERIFICATION: CASE A TO CASE B

The velocity field at probe locations (z = 2.5 m, 181 probes, 8 wind directions, N = 1448 points) was extracted from both simulations and compared using three statistical metrics computed on the velocity magnitude $|U|$ and on the turbulent kinetic energy k: RMSE, Bias (Case B − Case A), and Pearson r.

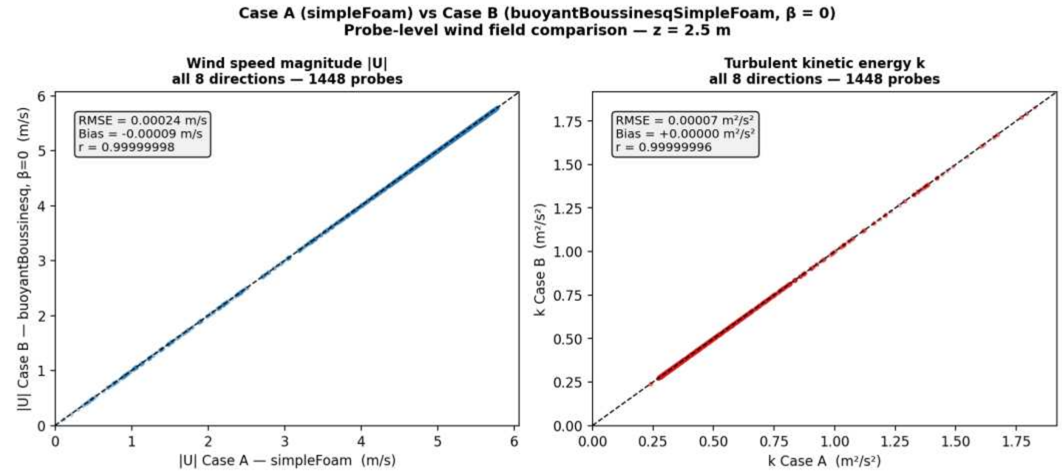

FIGURE 4.1 *Case A vs. Case B — scatter plots of wind speed magnitude |U| (left) and turbulent kinetic energy k (right) at probe locations (z = 2.5 m, 1448 points, 8 wind directions). Both fields collapse onto the 1:1 line with no visible spread.*

The scatter plots show all 1448 probe points collapsing onto the 1:1 line with no visible spread. The global statistics confirm near-zero differences: RMSE = 0.00024 m s^{-1} on $|U|$ (r = 1.000000), RMSE = 0.00007 m^2 s^{-2} on k (r = 1.000000). The maximum single-probe deviation in wind speed is 0.0022 m s^{-1}, corresponding to 0.12 % of the local velocity, at stagnation-adjacent probes where velocities are near zero. All values are consistent with floating-point rounding at double precision and fall five orders of magnitude below the 1% threshold.

TABLE 4.2 *Case A vs. Case B — per-direction statistics (RMSE, Bias, Pearson r) for wind speed |U| and turbulent kinetic energy k. The bottom row reports the global aggregate over all 1448 probe points.*

Per-direction statistics — Case B − Case A (N = 181 probes per direction)

Direction	RMSE |U| (m/s)	Bias |U| (m/s)	r |U|	RMSE k (m^2/s^2)	Bias k (m^2/s^2)	r k
0°	0.00003	+0.00001	1.0000000	0.00000	+0.00000	1.0000000
45°	0.00047	-0.00043	1.0000000	0.00009	-0.00001	0.9999999
90°	0.00010	-0.00009	1.0000000	0.00001	-0.00000	1.0000000
135°	0.00021	-0.00021	1.0000000	0.00001	+0.00001	1.0000000
180°	0.00012	+0.00012	1.0000000	0.00001	-0.00000	1.0000000
225°	0.00035	+0.00010	0.9999999	0.00018	+0.00005	0.9999995
270°	0.00022	-0.00020	1.0000000	0.00003	-0.00001	1.0000000
315°	0.00001	+0.00001	1.0000000	0.00000	-0.00000	1.0000000
All directions	**0.00024**	**-0.00009**	**1.0000000**	**0.00007**	**+0.00000**	**1.0000000**

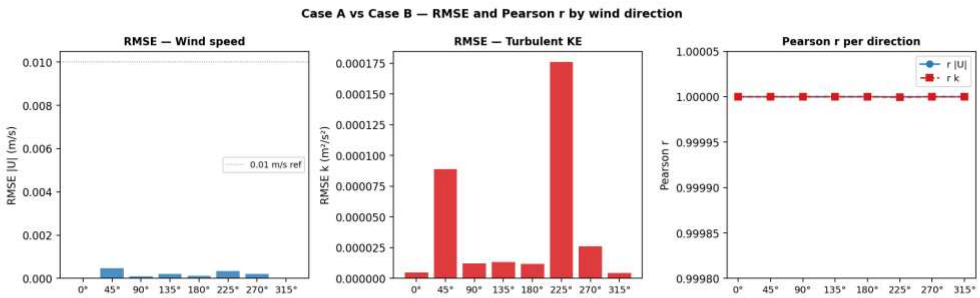

FIGURE 4.2 *Case A vs. Case B — RMSE of |U| and k (left and centre) and Pearson r (right) broken down by wind direction. All RMSE values are below 0.001 m/s; Pearson r exceeds 0.99999 for every direction.*

4.3.1 WIND COMFORT EQUIVALENCE

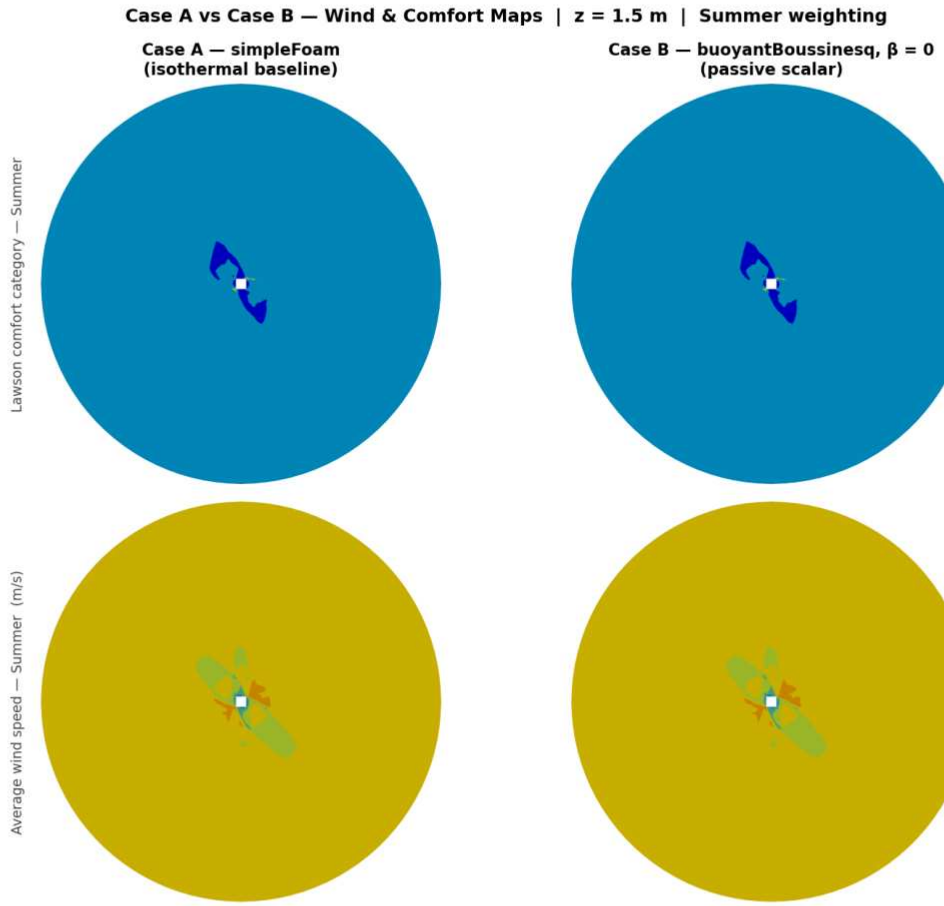

FIGURE 4.3 *Side-by-side comparison of the Lawson comfort category (top row) and the frequency-weighted average wind speed (bottom row) at pedestrian level z = 1.5 m, summer season. Case A (left) and Case B (right) are pixel identical.*

The Lawson comfort classification and average wind speed maps are pixel-identical between Case A and Case B (maximum pixel difference < 1/255). This confirms that switching from `simpleFoam` to `buoyantBoussinesqSimpleFoam` with $\beta = 0$ introduces no change in any wind-derived output. All differences subsequently observed in Case C are attributable unambiguously to active buoyancy coupling.

4.4 ACTIVE BUOYANCY ANALYSIS: CASE B VS. CASE C

4.4.1 SETUP

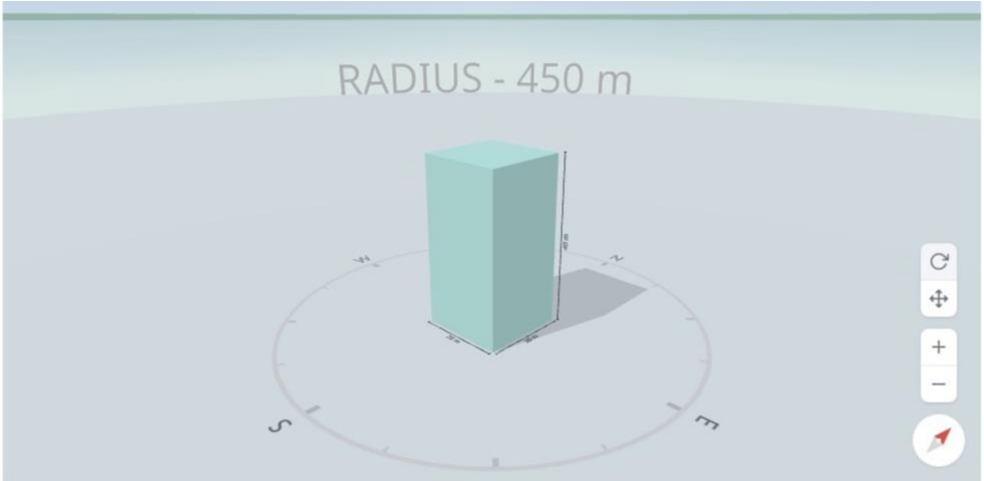

FIGURE 4.4 *AIJ Case A benchmark building: 20 × 20 m footprint, height H = 40 m (1:1:2 aspect ratio). Simulation domain radius 450 m.*

This section compares the results of Case B and Case C under the simulated urban heat island scenario to isolate the impact of thermal buoyancy. Both cases utilize the `buoyantBoussinesqSimpleFoam` solver on the AIJ Case A benchmark geometry. In Case B, the volumetric thermal expansion coefficient is set to zero ($\beta = 0$), treating temperature as a passive scalar with no kinematic feedback. In Case C, active buoyancy is enabled by setting $\beta = 3.33 \times 10^{-3}$ K^{-1}, allowing the Boussinesq buoyancy term $g \cdot \beta \cdot (T - T_{\text{ref}})$ to fully couple the temperature field with momentum. To evaluate these differences, the flow and thermal fields are analysed using ParaView extracts across three strategic planes: two mutually perpendicular vertical sections passing through the building centre, and a horizontal slice sampled at pedestrian height.

4.4.2 THERMAL PLUME TOPOLOGY AND MIXED CONVECTION DYNAMICS

The activation of buoyancy forces fundamentally alters the spatial organization of the thermal field around the obstacle. In the passive baseline configuration (Case B), the thermal energy released by the heated ground and building facades distributes almost isotropically through the lower portion of the computational domain, governed exclusively by forced advection and the turbulent diffusion of aerodynamic vortices.

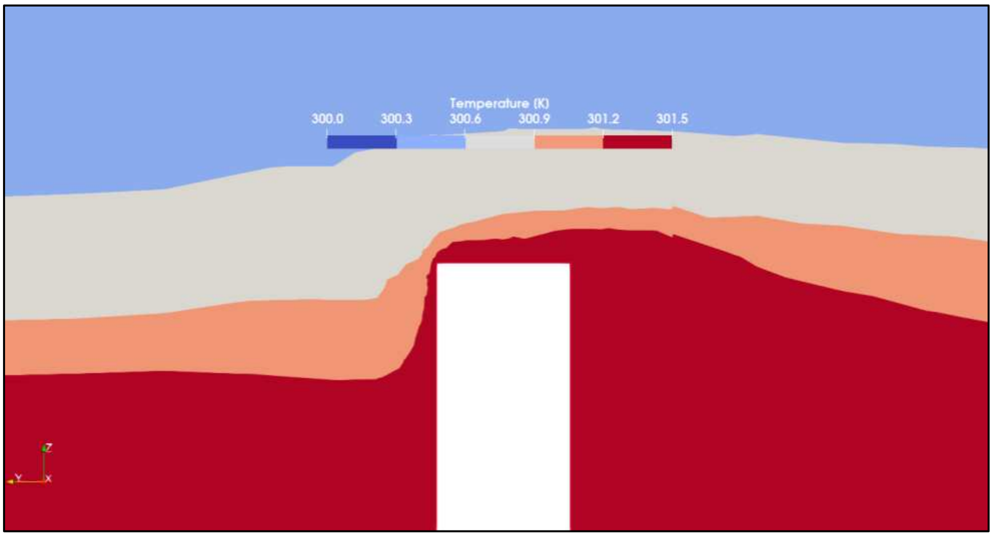

FIGURE 4.5 *Case B — passive (β = 0): heat diffuses broadly through the lower domain with no vertical organisation. Range 300.0–301.5 K.*

Conversely, in the active configuration (Case C), buoyancy forces organize the thermal transport into a highly structured convective plume. This plume rises directly from the heated urban surfaces and is subsequently tilted downstream by the incoming wind inertia.

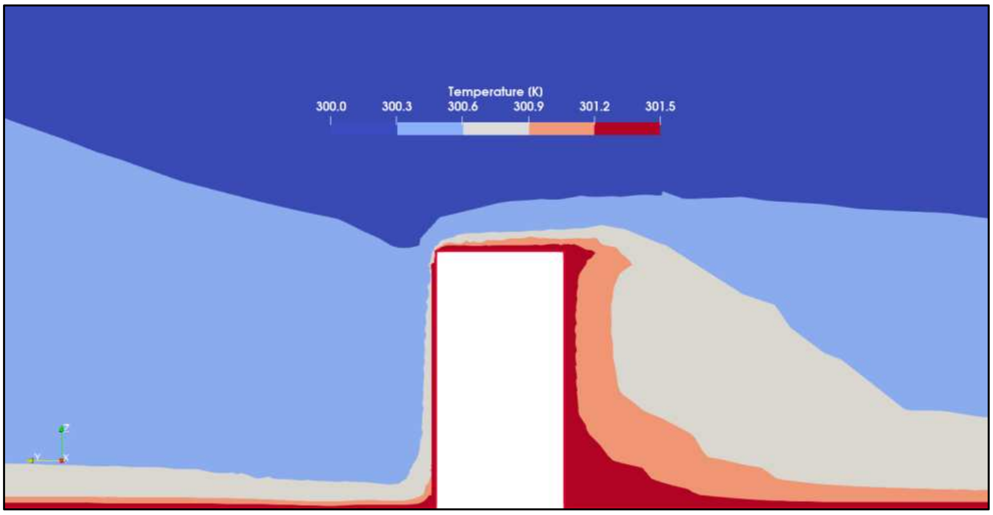

FIGURE 4.6 *Case C — active (β ≠ 0): the plume originates at the stagnation face and is bent downstream by the wind. Range 300.0–301.5 K.*

The thermal plume does not ascend vertically; the momentum of the oncoming wind is sufficient to bend the convective structure horizontally before buoyancy can drive significant vertical development. To characterize this interaction between thermal and inertial forces, the global Richardson number (R_i) is computed at the building scale, utilizing the nominal temperature differential (ΔT_{wall}) and the approach velocity evaluated at the roof height H from the AIJ Case A inlet profile:

$$R_{\mathrm{i}} = \frac{g \cdot \beta \cdot \Delta T_{\mathrm{wall}} \cdot H}{U(H)^2}$$

$$\approx (9.81 \times 3.33\times10^{-3} \times 20 \times 40) / 4.49^2 \approx 1.3$$

From a purely theoretical standpoint, a value of $R_{\mathrm{i}} \approx 1.3$ places the simulation nominally within the mixed-convection regime, where buoyancy and inertial forces are of the same order of magnitude (O (1)).

However, the effective temperature differential driving the plume motion within the free fluid volume is significantly smaller than the prescribed wall boundary condition (≈ 1-2 K in the core air volume versus 20 K directly at the wall interface), which proportionally reduces the effective buoyant lifting force. As a result, thermal energy accumulates within the recirculation wake behind the leeward facade, driving temperatures noticeably above ambient levels in wind-sheltered areas where aerodynamic stagnation dampens convective cooling.

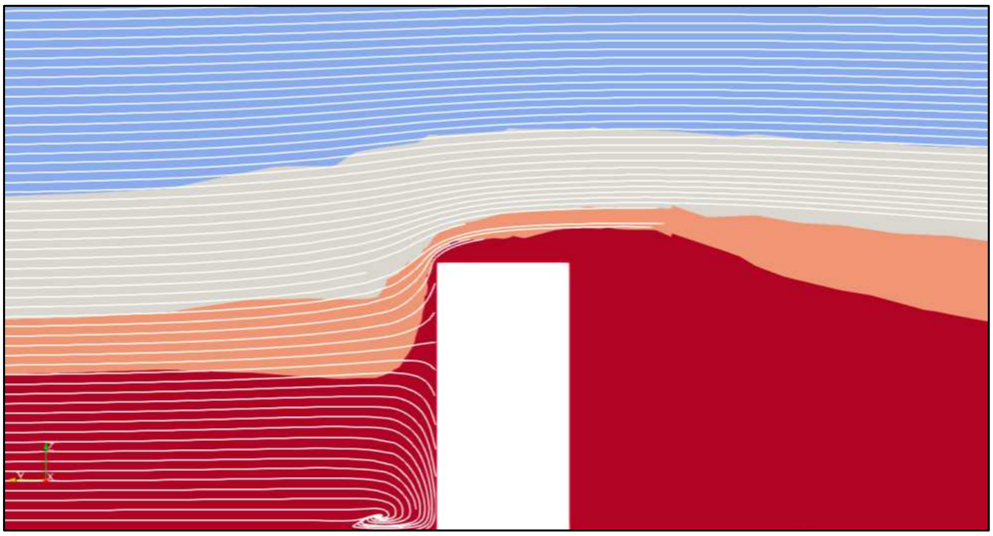

FIGURE 4.7 *Case B — streamlines: deflection above the roof driven only by wind dynamics.*

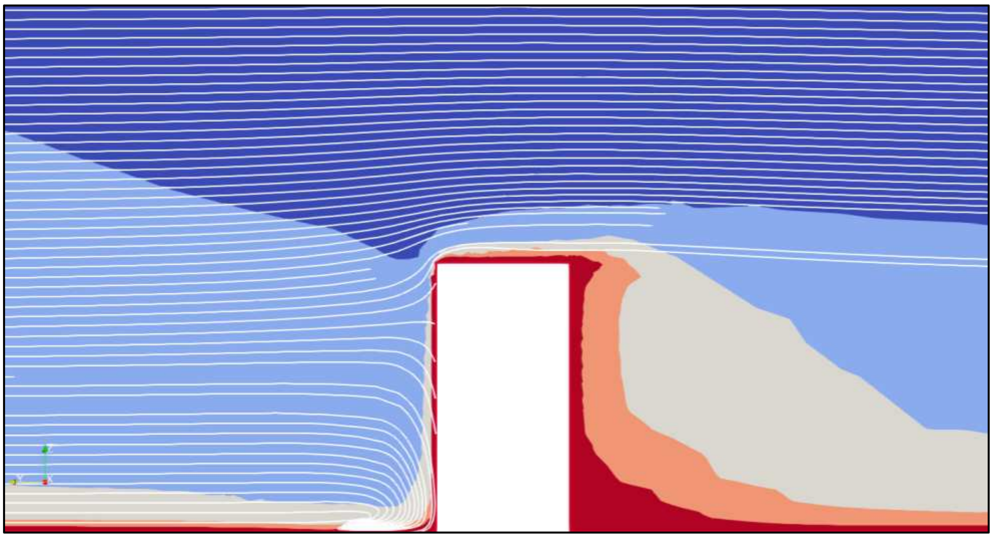

FIGURE 4.8 *Case C — streamlines: deflection above the roof driven by wind dynamics and the thermal plume.*

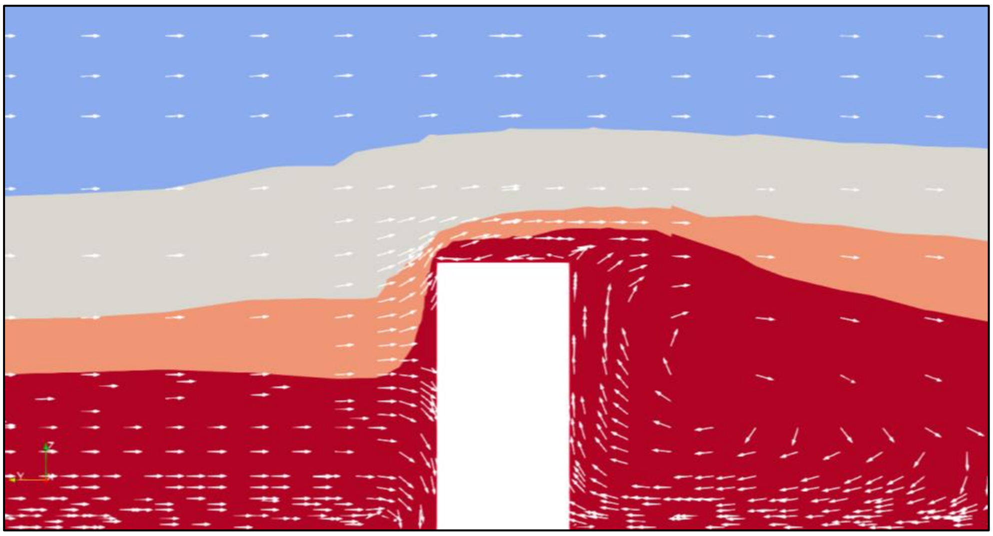

FIGURE 4.9 *Case B — velocity vector glyphs: corner vortex at the stagnation edge and recirculation in the wake.*

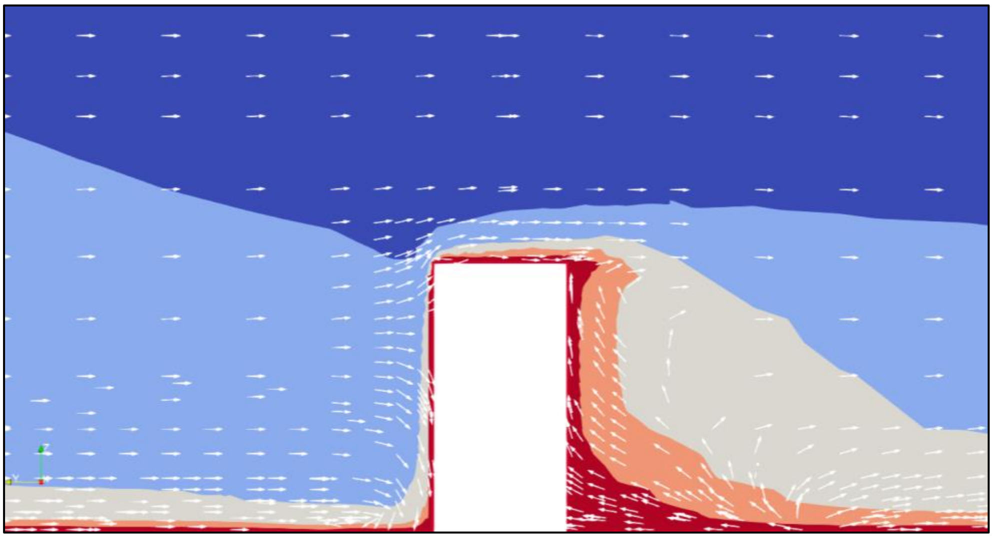

FIGURE 4.10 *Case C — velocity vector glyphs: same aerodynamic structure as Case B, with vertical component amplified by buoyancy.*

4.4.3 LOCAL THERMAL BOUNDARY LAYER AND BOUNDARY LAYER INSTABILITY

An analysis of the thermal fields adjacent to the solid boundaries reveals that thermal energy is strictly confined to a thin fluid shell immediately radiating from the building walls and the ground plane. As shown in the front-view vertical slice perpendicular to the wind in ***Figure 4.11***, the air temperature drops back to the ambient baseline value at merely 1m away from the physical facades. This sharp spatial gradient represents the expected thermodynamic behavior for highly turbulent forced convection developing over an isolated heated obstacle.

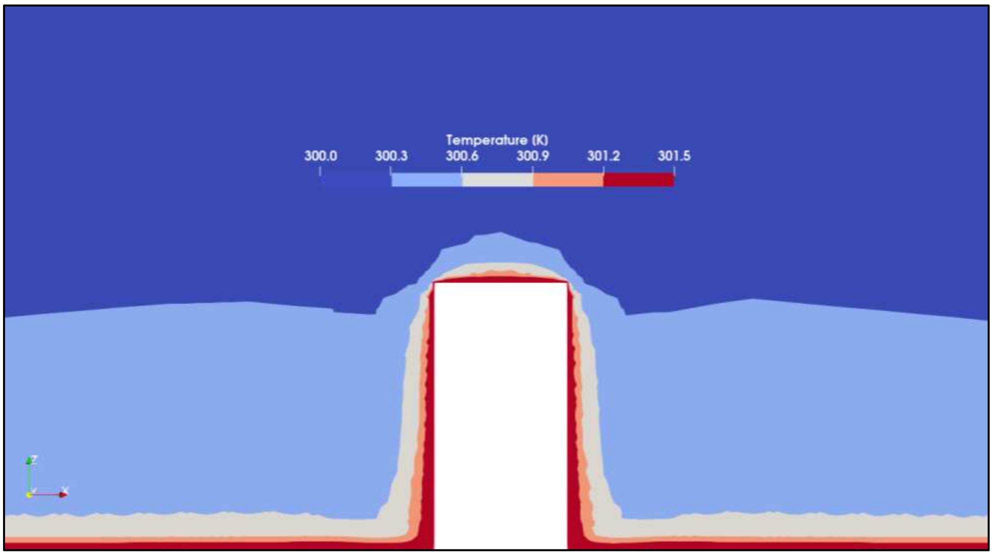

FIGURE 4.11 *Vertical slice parallel to the inlet (perpendicular to the wind), active case. Front view of the building. Range 300.0–301.5 K. Heat is confined to a thin shell adjacent to the walls and the ground.*

4.4.3.1 HEAT CONFINEMENT AND BOUNDARY LAYER THICKNESS.

To provide a mathematical justification for this localized confinement, the theoretical thermal boundary layer thickness (δ_T) for a fully turbulent boundary layer can be estimated via the standard empirical correlation:

$$\delta_T \approx 0.37 \cdot L \cdot Re^{-\frac{1}{5}}$$

Using an undisturbed reference free-stream velocity $U \approx 3$ m s^{-1}, a characteristic building length scale L = 40 m, and an air kinematic viscosity $\nu = 1.5\times10^{-5}\ m^2\ s^{-1}$, the macroscopic Reynolds number is evaluated as: Re = (U · L)/ $\nu \approx 8 \times 10^6$. The analytical calculation consequently yields a theoretical boundary layer thickness of $\delta_T \approx 0.6$ m. When evaluating the coupled energy balance within the global mixing volume, the theoretical mean temperature increase would amount to approximately $\Delta T_{\text{mix}} \approx 0.014$ K.

Since the cutting planes extracted from ParaView predominantly sample the grid cells located in the outer region of the boundary layer and within the free stream—where thermal energy is already heavily diluted by advection—the observation of a visible variation confined within a 1.5 K range confirms that the mesh resolution correctly captures the complex transition zone of the thermal boundary layer. This dynamic is in full agreement with experimental data found in literature: field measurement campaigns conducted over sun-exposed asphalt surfaces under summer midday conditions consistently document air temperature exceedances between +1 K and +3 K above ambient at the standard height of 1.5 m, validating the physical plausibility of the thermal field simulated by the solver.

4.4.3.2 RAYLEIGH-BENARD INSTABILITY ON THE VERTICAL WALL

A close-up analysis conducted near the vertical surfaces of the building (***Figure 4.12***) reveals the presence of an undulatory sub-structure characterizing the warm boundary layer. Small parcels and bubbles of heated air periodically detach from the thermal boundary layer adjacent to the wall, moving upward under the action of hydrostatic buoyancy.

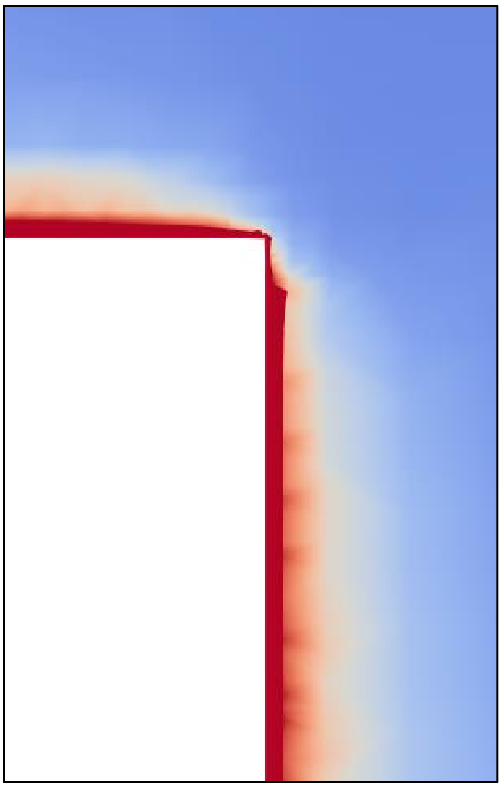

FIGURE 4.12 *Zoom on the side wall of the building (active case). Small ridges and warm bubbles detaching from the surface are natural convective plumes — a signature of thermal boundary-layer instability (Rayleigh-Bénard type).*

This phenomenon is attributable to an unstable natural convection regime on a vertical wall, a mechanism driven by fluid buoyancy that is entirely analogous to the classic Rayleigh-Bénard instability commonly observed on heated horizontal plates. The spontaneous onset of these microscopic thermal ridges within the calculation code demonstrates that the implemented numerical framework faithfully reproduces natural convection phenomena at small scales, while simultaneously confirming that the mesh density and resolution at the wall are adequate for the correct closure of the boundary layer.

4.4.4 VERTICAL VELOCITY PATTERN: BUOYANCY AND THE HORSESHOE VORTEX

The analysis of the vertical velocity component (U_z) along the crosswind symmetry plane highlights the profound impact of active buoyancy when transitioning from the passive to the active solver.

Looking at ***Figure 4.13*** (Case B, passive scalar), the vertical velocity field is dominated almost entirely by purely hydrodynamic mechanisms typical of fluid flow around bluff bodies. Laterally to the obstacle, a strongly negative vertical velocity field emerges: this localized downdraft is the kinematic signature of the classic basal horseshoe vortex. Above the roof structure ($z > H$), the upward transport is weak, reaching a peak of just 0.22 m s^{-1}, driven solely by the mechanical displacement of the wind hitting the building.

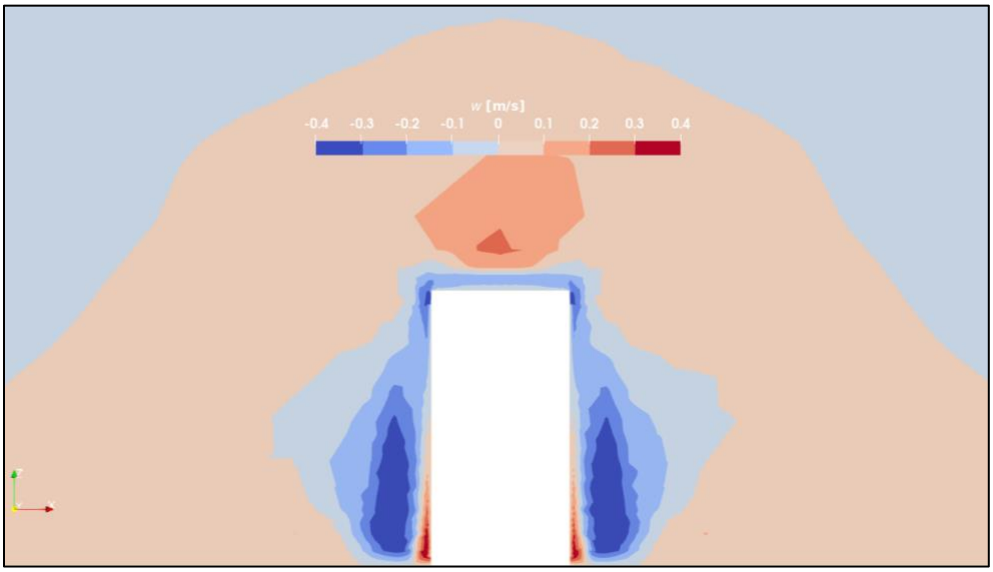

FIGURE 4.13 *Case B — passive: the horseshoe vortex dominates. Intense negative Uz at the sides; weak updraft above the roof. Range −0.4 to +0.4 m/s.*

A completely different physical pattern emerges in ***Figure 4.14*** (Case C, active buoyancy), where the activation of the thermal expansion coefficient (β) couples the temperature field to momentum via the Boussinesq term $g \cdot \beta \cdot (T - T_{\text{ref}})$. In direct contact with the heated building walls and at the basal corner, an intense positive updraft develops. This convective lift accelerates the heated boundary layer air upward along the warm fronts of the building, propagating well beyond the roof height and pushing the vertical velocity peak to 0.50 m s^{-1} — a significant 2.3-fold amplification compared to the passive baseline.

Interestingly, Figure 6.14 demonstrates that this new thermal mechanism and the hydrodynamic horseshoe vortex coexist within the domain without numerical contradiction. The positive convective uplift remains strictly confined within the thin geometric limits of the thermal boundary layer, but its upward acceleration alters the overall momentum distribution, effectively reducing the intensity and extension of the negative downdraft compared to what was observed in Case B

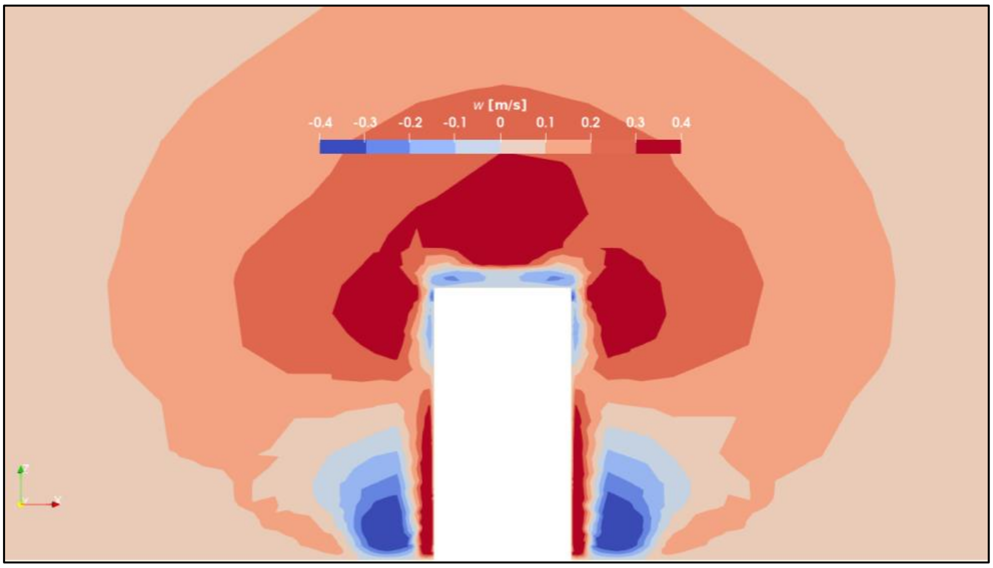

FIGURE 4.14 *Case C — active: amplified updraft above the roof (0.50 m/s vs. 0.22 m/s); attenuated horseshoe vortex. Range −0.4 to +0.4 m/s.*

4.4.5 PEDESTRIAN MICROCLIMATE (Z = 2M)

When analyzing the horizontal plane occupied by pedestrians, the direct comparison highlights the macroscopic impact of buoyant forces on the urban microclimate. In Case B (passive, ***Figure 4.15***), the heat released by the solid surfaces accumulates near the facades and spreads horizontally into the aerodynamic wake at ground level.

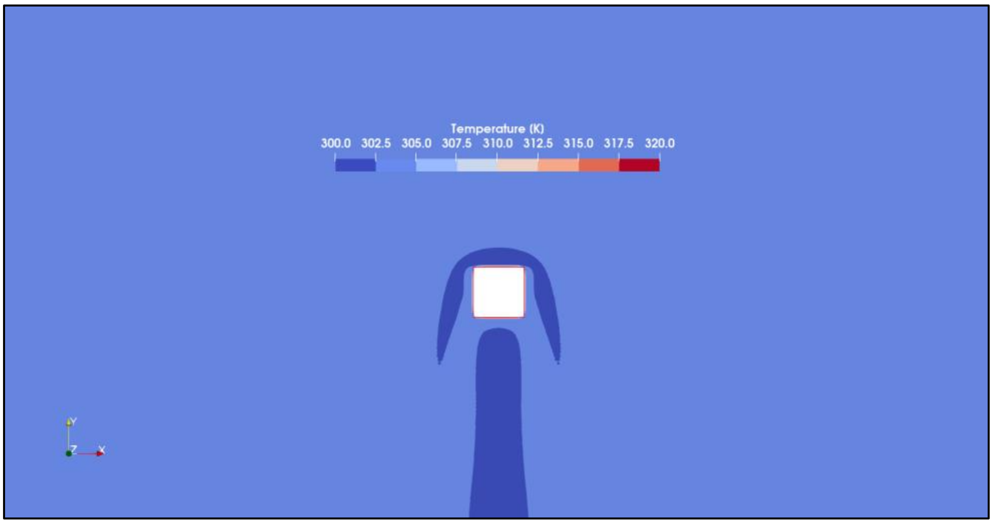

FIGURE 4.15 *Case B — passive: heat distributes diffusely around the building and in the wake at ground level. Range 300–320 K.*

In Case C (active, ***Figure 4.16***), the activation of buoyancy radically modifies this topography: thermo-convective forces lift the heated air masses vertically, moving the heat away from the pedestrian height immediately downstream of the building. The resulting thermal field appears much more confined and adherent to the immediate vicinity of the heated surfaces, resulting in a significantly cooler and better-ventilated downwind wake compared to the passive scalar configuration.

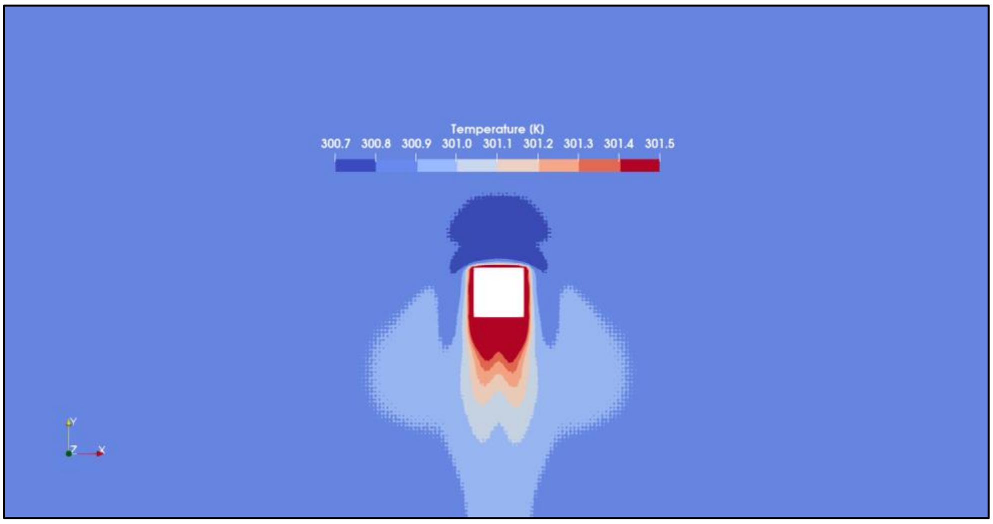

FIGURE 4.16 *Case C — active: buoyancy lifts heat upward, reducing temperature at pedestrian level in the wake. The thermal field is more confined to the immediate vicinity of the heated surfaces. Range 300.7–301.5 K.*

4.4.6 THERMAL SIGNATURE IN THE MODIFIED PRESSURE P_RGH

To understand the hydrodynamic engine behind this flow reorganization, it is necessary to analyse the behaviour of the modified pressure (p_{rgh}), defined within the `buoyantBoussinesqSimpleFoam` solver through the relation:

$$p_{\mathrm{rgh}} = p - \rho_0 \cdot g \cdot z$$

By subtracting the reference hydrostatic contribution from the total pressure, the variable p_{rgh} isolates solely the pressure component due to the actual fluid motion and, in the active case, to the thermal coupling. In areas where the air is locally warmer than the reference temperature ($T > T_{\mathrm{ref}}$), the local fluid density drops below the reference value p_0, inducing a negative variation in the values of p_{rgh}. This low-pressure zone of purely thermal origin generates the hydrostatic buoyancy lift (upward thrust) that feeds the development of the convective plume.

The analysis of longitudinal cutting planes parallel to the wind direction (***Figure 4.17*** and ***Figure 4.18***) highlights that, in the passive case, a strong area of low pressure concentrates exclusively above the windward edge of the roof, marking the presence of the aerodynamic separation bubble induced by the sharp edges of the bluff body. In the active case, the activation of buoyancy provides an additional vertical thrust to the air near the hot roof; this phenomenon attenuates the original edge separation and generates a much more uniform and distributed modified pressure field of $p_{\mathrm{rgh}} \approx -2\ \mathrm{m}^2\ \mathrm{s}^{-2}$above the roof covering.

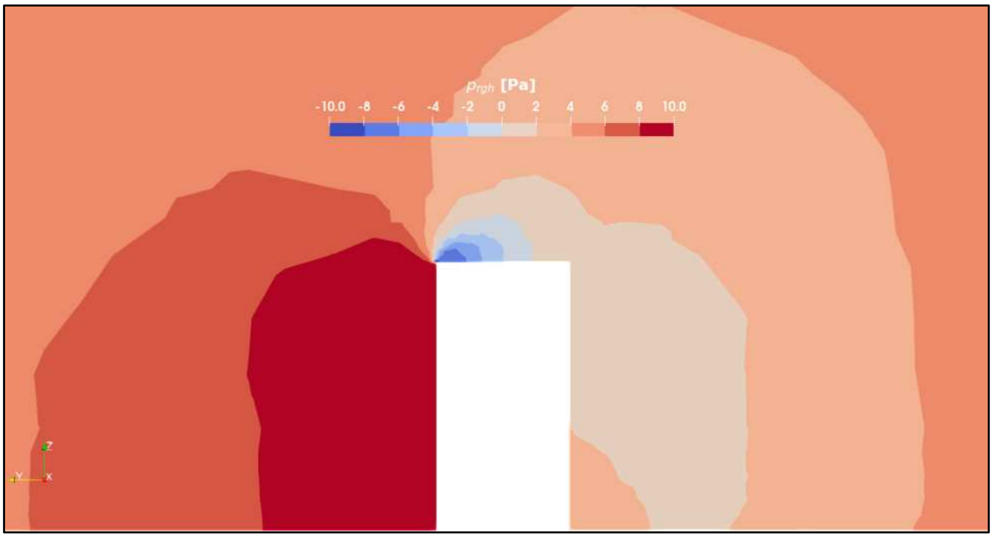

FIGURE 4.17 *Case B — passive, along-wind slice: a blue zone concentrated above the windward roof edge marks the aerodynamic separation bubble. Range −10 to +10 m²/s².*

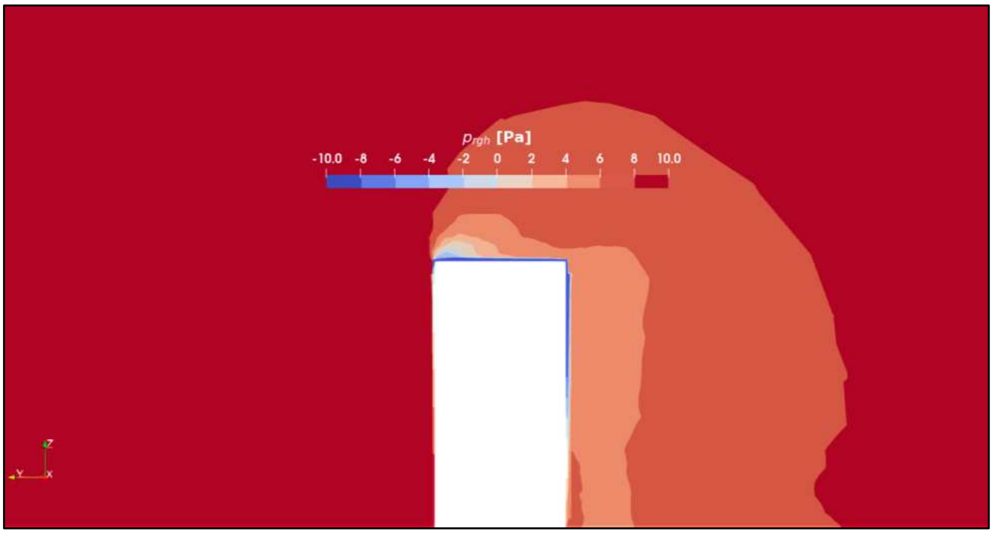

FIGURE 4.18 *Case C — active, along-wind slice: the thermal low $p_{RGH} \approx -2$ m²/s² above the roof provides the additional upward thrust. Range −10 to +10 m²/s.*

In the passive case, a blue zone appears concentrated above the windward roof edge: it is the aerodynamic separation bubble expected for wind flow over a bluff body with sharp edges. In the active case, buoyancy provides upward momentum to the air near the hot roof, modifying the aerodynamic separation and producing a more uniform p_{rgh} field.

Similarly, analysing the transverse cutting planes perpendicular to the flow (***Figure 4.19*** and ***Figure 4.20***), Case B shows a symmetric low-pressure halo surrounding the entire perimeter of the building, combining the geometric signature of the basal horseshoe vortex with the flow lines separated from the roof according to a purely aerodynamic mechanism. In Case C, the activation of thermal buoyancy forces redistributes momentum along the vertical axis of the facade, mitigating the lateral depression of the horseshoe vortex and uniforming the distribution of p_{rgh}. The quantitative difference between the two fields maps the exact extent to which thermal buoyancy reorganizes the vortex structures and stress fields around the building.

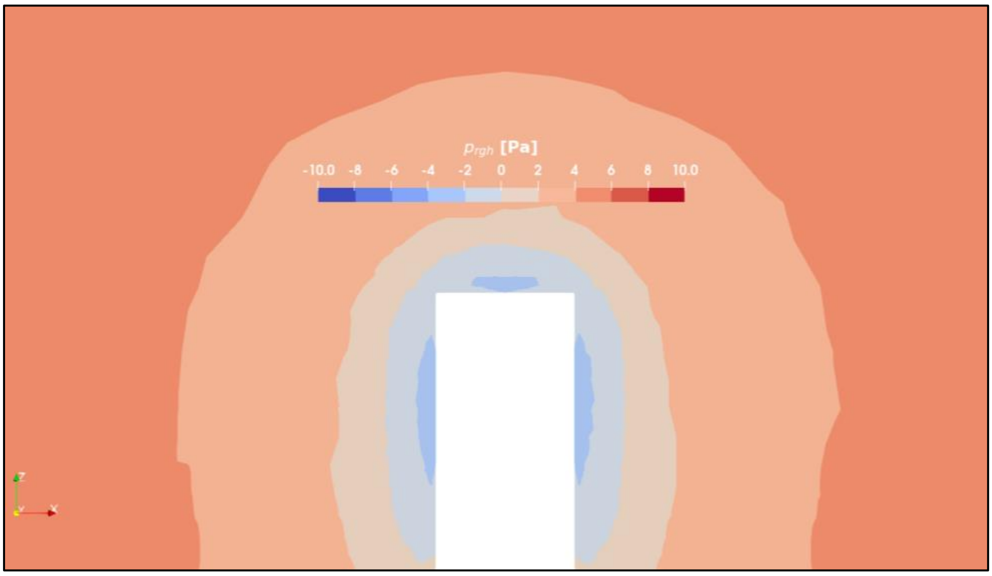

FIGURE 4.19 *.Case B — passive, wind-perpendicular slice: a symmetric low-pressure halo (blue) surrounding the building — the combined signature of the horseshoe vortex and roof separation*

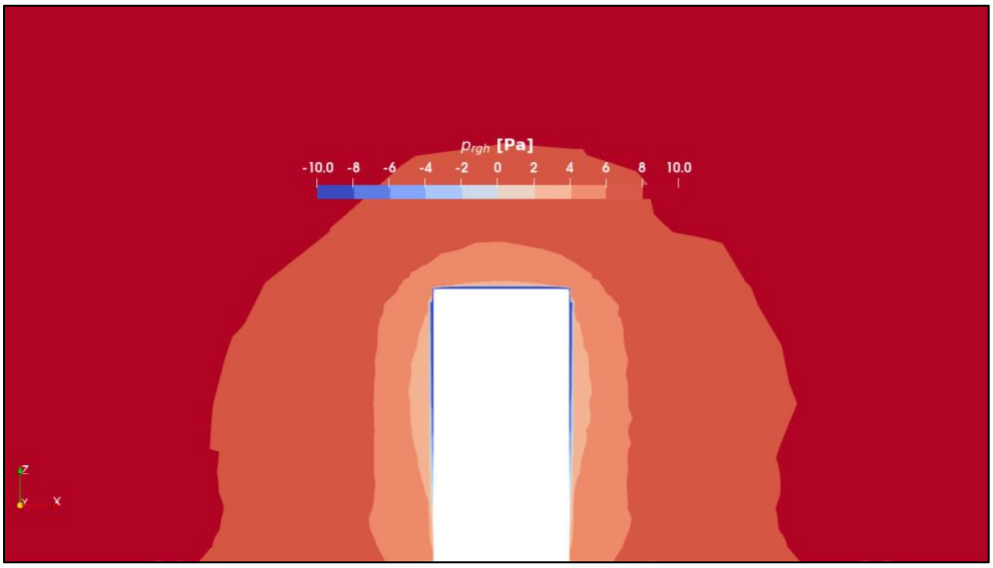

FIGURE 4.20 *Case C — active, wind-perpendicular slice: buoyancy attenuates the horseshoe vortex by redistributing momentum vertically.*

4.4.7 VERTICAL PROFILES ABOVE THE ROOF CENTER

4.4.7.1 MATHEMATICAL PROFILE ANALYSIS.

To precisely quantify the structural modifications induced by thermal buoyancy, numerical profiles were extracted along a vertical line centred directly above the building roof core, extending from the rooftop interface up into the far-field boundary layer ($x = 0$ m, $y = 0$ m, $z = 40 - 100$ m). This profile evaluates the spatial evolution of temperature (T), vertical velocity (U_z), and modified pressure (p_{rgh}), tracking the buoyant plume decay as it approaches the upper computational boundary.

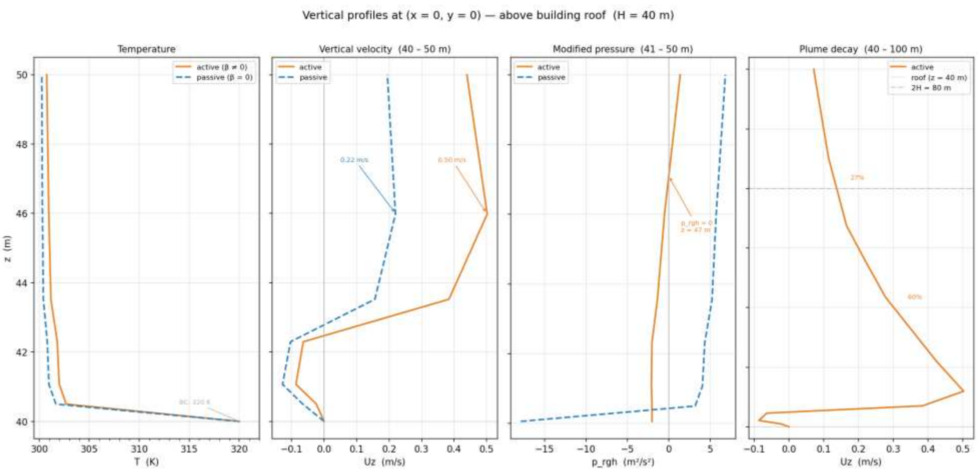

FIGURE 4.21 *Vertical profiles extracted along (x = 0, y = 0, z = 40–100 m) for temperature T, vertical velocity Uz and modified pressure* p_{RGH}*, plus plume decay to z = 100 m. Orange: active (Case C). Blue dashed: passive (Case B).*

4.4.7.2 THERMAL DISTRIBUTION.

At the immediate wall interface ($z = 40$ m), both configurations strictly honor the identical Dirichlet boundary condition ($T_{BC} = 320$ K). Within the first computational meter, both profiles exhibit a nearly identical thermal boundary layer collapse, with temperatures dropping sharply from 320 K to approximately 302 K.

Past this immediate wall shell, a distinct divergence emerges: the active buoyancy configuration (Case C) consistently maintains temperatures roughly 0.5–1 K higher than the passive baseline (Case B) throughout the remaining vertical profile. This persistent offset provides clear quantitative evidence of the enhanced upward convective heat transport driven by the active buoyancy forces.

4.4.7.3 VERTICAL VELOCITY KINEMATICS.

The vertical velocity profile (U_z) can be divided into three physically distinct fluid zones:

The recirculation layer ($z = 40 - 42$ m): Both cases exhibit slightly negative vertical velocities ranging from $-0.08\ \mathrm{m\ s^{-1}}$ to $-0.12\ \mathrm{m\ s^{-1}}$. This local downdraft captures the aerodynamic recirculation bubble generated as the oncoming wind separation negotiates the sharp windward edge of the roof.

The acceleration zone ($z = 42 - 46$ m): The flow experiences a rapid flow reversal and sharp upward acceleration. The active case responds much faster to the thermal gradient; at $z = 43$ m, the active configuration has already reached $U_z = +0.19\ \mathrm{m\ s^{-1}}$, whereas the passive baseline registers a negligible vertical velocity of +0.05m/s.

The core plume region ($z \geq 46$ m): The vertical profiles reach their respective maximum updraft velocities at a height of approximately $z \approx 46$ m. Here, the active case reaches a peak of $0.50\ \mathrm{m\ s^{-1}}$ compared to the $0.22\ \mathrm{m\ s^{-1}}$ of the passive case. This substantial 2.3× velocity amplification is entirely attributable to the action of the buoyancy body force coupling term.

4.4.7.4 MODIFIED PRESSURE AND PLUME DECAY.

The pressure profiles provide the underlying mechanical explanation for this kinematic acceleration. In the passive baseline configuration, p_{rgh} remains positive and increases monotonically from $+4\ \mathrm{m^2 s^{-2}}$ to $+7\ \mathrm{m^2 s^{-2}}$, tracking the purely wind-driven dynamic pressure developing over the roof.

In stark contrast, the active case registers a negative modified pressure of, $p_{\mathrm{rgh}} \approx -2\ \mathrm{m^2 s^{-2}}$ at $z = 41$ m, which rises gradually and crosses the zero threshold at $z \approx 47$ m. This persistent thermal low-pressure zone over the first 7 m above the roof acts as the quantitative engine behind the convective lift: because the heated air is less dense than the reference density ($\rho < \rho_0$), the modified pressure gradient develops an upward-directed force that effectively doubles the vertical velocity magnitude.

As the fluid travels further up, the convective plume gradually loses energy due to momentum exchange and turbulent mixing with the crossflow. At a height of $z = 60$ m, the vertical velocity component has dropped to 60 % of its peak value: these decays further to 27 % at $z = 80$ m, and down to 14 % at the $z = 100$ m boundary. The thermal signature remains trackable well above the building envelope before being fully reabsorbed into the background atmospheric wind flow, adhering to the asymptotic decay profile characteristic of a buoyant thermal plume subjected to a perpendicular crossflow.

4.5 PEDESTRIAN WIND COMFORT IMPACT OF THERMAL BUOYANCY

4.5.1 PEDESTRIAN WIND COMFORT ASSESSMENT METHODOLOGY

Pedestrian wind comfort assessment requires the statistical characterization of the wind environment rather than the evaluation of any single instantaneous flow state because the wind conditions experienced at an urban location vary continuously as a function of the prevailing synoptic pressure pattern, season, and atmospheric stability class.

The comfort maps presented in this study are produced by combining three distinct inputs:

Velocity amplification fields: Steady-state CFD simulations for eight wind directions (0°,45°,…,315°) providing the spatial distribution of the normalized wind speed amplification factor at z =1.5 m above ground level (AGL).

Meteorological boundary conditions: A representative site wind rose encoding the joint probability distribution of wind direction and speed magnitude.

Comfort thresholds: The Lawson comfort criteria, which map the exceedance probabilities of defined wind speed thresholds onto activity suitability categories.

For each domain point, the CFD amplification factor is applied to the full wind-speed distribution of each directional sector. The sector contributions are then superimposed in proportion to their directional frequencies, and the resulting cumulative distribution is evaluated against the threshold matrix.

The Lawson classification scheme adopted here is summarized in ***Table** 4.3*. A location is assigned the most favourable category whose wind speed threshold is not exceeded more than 5 % of the time; locations where the 10 m s^{-1} threshold is exceeded more than 5 % of the time are classified as Uncomfortable (Category E). Wind distress assessment, which concerns pedestrian safety under extreme wind events, is outside the scope of the present analysis.

Table 4.3 *Lawson pedestrian wind comfort categories (Lawson, 1978).*

Category	**Activity**	**Wind speed threshold**	**Max exceedance**
A	Sitting	< 4 m s^{-1}	5 %
B	Standing	< 6 m s^{-1}	5 %
C	Strollling	< 8 m s^{-1}	5 %
D	Business walking	< 10 m s^{-1}	5 %
E	Uncomfortable	≥ 10 m s^{-1}	—

4.5.2 CASE B VS. CASE C: EFFECT ON LAWSON CATEGORIES

The preceding sections quantified the three-dimensional momentum modifications induced by the active Boussinesq buoyancy term: a 2.3× amplification of the vertical velocity above the roof, an attenuation of the horseshoe vortex footprint, and a vertical redistribution of horizontal momentum along the heated façades. This section translates those velocity-field changes into their impact on the Lawson pedestrian wind comfort classification at $z = 1.5$ m, comparing Case B ($\beta = 0$, passive scalar transport, no buoyancy feedback on momentum) against Case C ($\beta = 3.33 \times 10^{-3}$ K^{-1}, active buoyancy coupling). Statistics are computed over all valid domain points, excluding building footprints.

4.5.2.1 QUANTITATIVE RESULTS: SUMMER VS. ANNUAL

The domain-averaged statistical metrics and spatial distribution for both the summer season and the annual average are presented below. ***Table 4.4*** and ***Table 4.5*** summarises the numerical variations, while ***Figure 4.22*** to ***Figure 4.25*** illustrate the corresponding comfort and velocity fields.

Table 4.4 *Domain-area statistics — Case B vs. Case C, summer season*

Metric	**Case B**	**Case C**	**Δ**
Category A (sitting), % domain area	0.8 %	99.8 %	+99.0 pp

Metric	Case B	Case C	Δ
Category B (standing), % domain area	99.0 %	0.1 %	−98.9 pp
Categories C–E, % domain area	0.2 %	0.1 %	≈ 0
Frequency-weighted mean wind speed	1.69 m s^{-1}	1.44 m s^{-1}	−0.25 m s^{-1} (−14.8 %)

Table 4.5 *Domain-area statistics — Case B vs. Case C, annual average*

Metric	Case B	Case C	Δ
Category A (sitting), % domain area	0.1 %	99.6 %	+99.5 pp
Category B (standing), % domain area	99.5 %	0.3 %	−99.2 pp
Categories C–E, % domain area	0.4 %	0.1 %	≈ 0
Frequency-weighted mean wind speed	1.89 m s^{-1}	1.62 m s^{-1}	−0.27 m s^{-1} (−14.3 %)

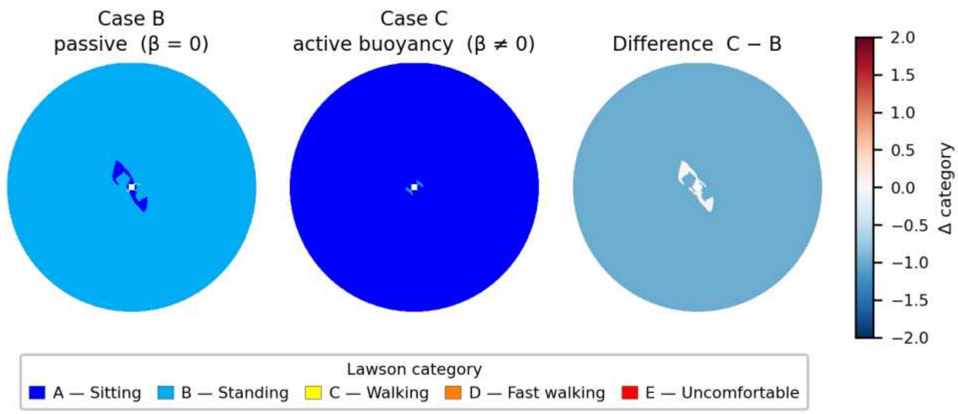

Figure 4.22 *Lawson pedestrian wind comfort category maps at z = 1.5 m, summer wind-rose weighting. Left: Case B (passive scalar, β = 0). Centre: Case C (active buoyancy, β ≠ 0). Right: pairwise category difference (C − B); blue tones indicate improvement in Case C*

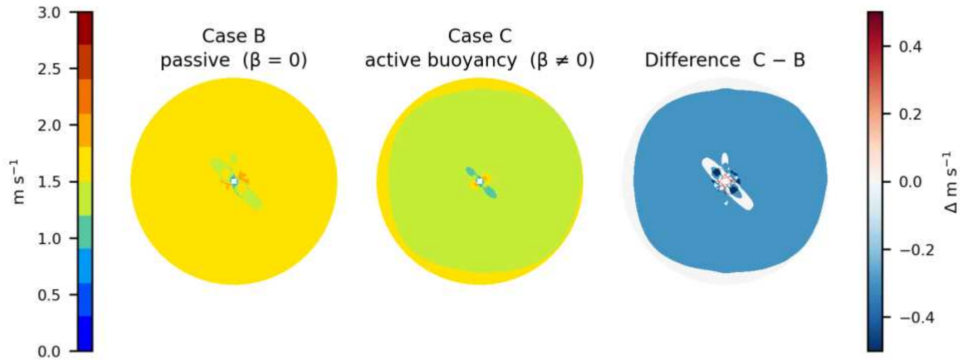

Figure 4.23 *Frequency-weighted mean wind speed [m s⁻¹] at z = 1.5 m, summer wind-rose. Left: Case B. Centre: Case C. Right: difference (C – B)*

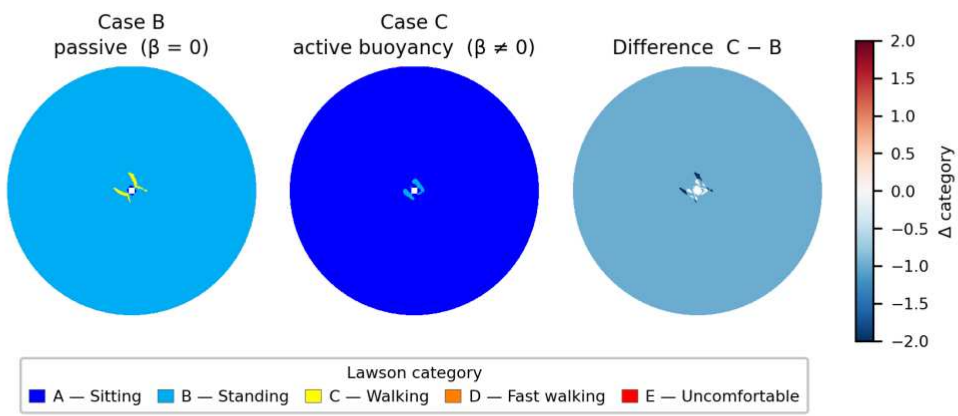

Figure 4.24 *Same layout as Figure 4.N, evaluated using the annual wind-rose frequency distribution*

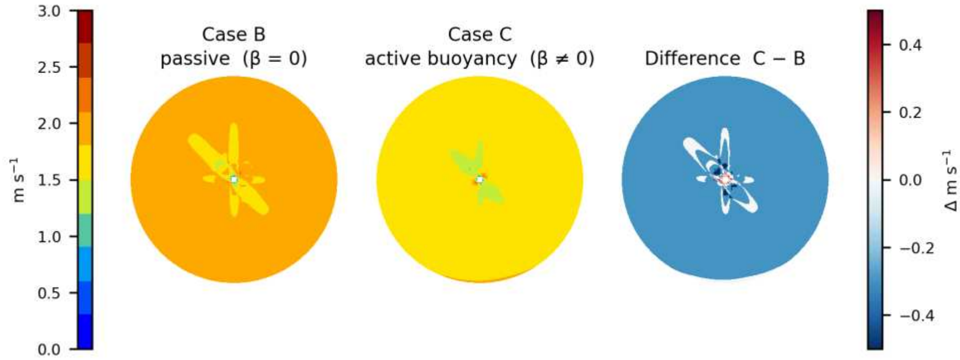

Figure 4.25 *Annual frequency-weighted mean wind speed. The higher background speed relative to the summer case (approximately 1.9 m s⁻¹ vs. 1.7 m s⁻¹) reflects the contribution of stronger winter winds in the annual average*

Both assessments reveal a massive, near-complete category migration when transitioning from the passive to the active buoyancy solver. In the passive baseline (Case B), approximately 99 % of the entire domain area is classified as Category B (Standing). Once thermal buoyancy is activated (Case C), this layout is completely inverted, with over 99 % of the domain migrating into the more comfortable Category A (Sitting). This macro-trend is identical in both seasonal setups: Category A gains +99.0 pp (percentage points) in summer and +99.5 pp annually.

The only minor difference between the two configurations emerges in the baseline aerodynamic speeds. The annual frequency-weighted mean wind speed in Case B is slightly higher than the summer counterpart (1.89 $\mathrm{m\ s^{-1}}$ vs. 1.69 $\mathrm{m\ s^{-1}}$). This shift reflects the historical contribution of stronger winter winds integrated into the annual wind rose, which consequently yields a marginally larger fraction of corner acceleration zones (Categories C–E) in the annual baseline (0.4 % vs. 0.2 %). Despite this higher background wind climate, the net buoyancy-induced reduction remains remarkably stable, tracking at −0.25 $\mathrm{m\ s^{-1}}$ (−14.8 %) for summer and −0.27 $\mathrm{m\ s^{-1}}$ (−14.3 %) annually.

4.5.2.2 PHYSICAL INTERPRETATION

The domain-wide reduction in frequency-weighted mean horizontal wind speed at $z = 1.5$ m (ranging between −0.25 $\mathrm{m\ s^{-1}}$ and −0.27 $\mathrm{m\ s^{-1}}$, or roughly −14 %) is a direct consequence of the buoyancy-driven vertical momentum transfer. The thermal updraft above the roof diverts kinetic energy from the near-ground horizontal flow into the vertical component. Because this redistribution acts over the full extent of the heated urban surfaces, it produces a spatially uniform reduction rather than a localized perturbation.

It is critical to note that the near-complete category migration—from approximately 99 % Category B in Case B to 99 % Category A in Case C—is primarily a threshold effect rather than evidence of an overwhelmingly large buoyancy influence on pedestrian comfort. In the velocity maps, both cases produce local speeds in the range of 1.4–1.9 m s^{-1} under the chosen inlet reference velocity. While these values are physically low, their interaction with the historical wind rose shifts the annual meteorological velocity distribution. The buoyancy-induced reduction of approximately 0.25 m s^{-1} is sufficient to displace the entire wind-speed distribution below this threshold uniformly across the domain. This occurs because the background speeds are everywhere in a regime where the comfort class is determined by the tail of the distribution, making the tail shift induced by buoyancy decisive.

In urban environments with higher background wind speeds—for example, at exposed building corners or within narrow street canyons—a buoyancy-induced velocity reduction of this magnitude would be insufficient to alter the Lawson category, and the comfort classification would remain unchanged. The corner acceleration zones (Categories C–E, occupying approximately 0.1–0.4 % of the domain area) remain unaffected in both assessments, confirming that the buoyancy forcing does not modify the building-geometry-driven corner jets to any significant degree.

Furthermore, the radial spoke pattern visible in the wind speed difference maps is a characteristic feature of the multi-directional assessment methodology: each of the eight inflow directions contributes its own directional wake footprint, and the linear superposition of these eight footprints produces the observed radial geometry. This pattern is physically consistent and does not constitute a numerical artefact.

4.5.2.3 IMPLICATIONS FOR COMPUTATIONAL WIND ASSESSMENT PRACTICE

The results demonstrate that the choice of buoyancy coupling has direct consequences for pedestrian wind comfort predictions in urban environments where surface heating is significant, and background pedestrian-level wind speeds are moderate (of order 1–3 m s^{-1}). A simulation with $\beta = 0$ is, by construction, conservative with respect to pedestrian wind comfort: it systematically overestimates horizontal wind speed at ground level, and consequently assigns less favourable Lawson categories than the physically correct buoyancy-coupled solution.

In summer scenarios with heated urban surfaces, neglecting buoyancy may therefore lead to classifying as Category B (standing) areas that, under thermally accurate conditions, fall within Category A (sitting). From a regulatory standpoint this constitutes a safe-side assumption and is therefore not technically incorrect; however, it may become relevant in urban design optimization workflows, where an overly conservative comfort prediction could lead to unnecessary mitigation measures or to underestimating the suitability of outdoor spaces for sedentary activities.

Quantitatively, the sensitivity of the comfort classification to buoyancy is strongest precisely in the wind speed range most encountered in sheltered urban microenvironments — the same range in which thermal comfort conditions (MRT, UTCI) are most critical from a health and well-being perspective. This coupling between the aerodynamic and thermal aspects of outdoor comfort underlines the value of integrated assessment approaches such as that described in Chapter 5.

4.6 SYNTHESIS

The multi-dimensional analysis comprising ParaView cutting planes, modified pressure fields, horizontal pedestrian-level comfort maps, and quantitative vertical profiles supports a series of definitive conclusions regarding the implementation of thermal buoyancy within the ArchiWind platform.

First, numerical and physical coupling integrity is fully verified; the solver operating under active thermal expansion coefficients (β = 0) produces physically coherent, structurally stable flow patterns, including a downstream-tilted convective plume, a stratified thermal boundary layer, localized Rayleigh-Bénard instabilities along the vertical surfaces, and a strong upward lift along the sun-exposed facades.

Second, the activation of buoyancy transforms isotropic thermal transport from a purely passive, diffusely scattered scalar field into a highly structured plume. By lifting thermal energy vertically away from the ground plane, this convective mechanism exerts a direct, quantifiable impact on the quality of the pedestrian-level urban microclimate.

From a quantitative kinematic perspective, the updraft velocity above the roof center is amplified by 2.3× relative to the passive baseline. The mechanical engine driving this acceleration is a persistent thermal low-pressure zone ($p_{\mathrm{rgh}} \approx -2\ \mathrm{m^2 s^{-2}}$) originating from local density deficits, which gradually decays to zero at an asymptotic height of $z \approx$ 47 m ($\approx 1.2H$).

This localized thermal lift introduces major aerodynamic feedback loops, attenuating both the basal horseshoe vortex footprint and the roof edge separation bubble by continuously redistributing horizontal momentum into the vertical axis along the building boundaries.

Crucially, the bounded volumetric temperature variation observed within the core flow (≈1.5 K) is physically consistent and does not constitute a numerical artefact; rather, it reflects the exact thermodynamic equilibrium between surface thermal forcing and intensive turbulent mixing under wind-dominated, mixed-convection conditions, matching empirical field observations.

Finally, the simulation successfully captures localized heat accumulation within the leeward wake behind the building envelope, where wind-sheltered stagnation zones inhibit convective cooling, providing an essential diagnostic foundation for evaluating outdoor thermal comfort thresholds.

5 OUTDOOR THERMAL COMFORT ASSESSMENT: THE ARCHIWIND MICROCLIMATE MODEL

5.1 INTRODUCTION

Outdoor thermal comfort is increasingly recognized as a critical dimension of urban environmental quality, governing the usability of public spaces and the health and well-being of city residents. The thermal environment experienced by pedestrians in urban areas is determined by a complex interplay of solar radiation, longwave thermal emission from surrounding surfaces, air temperature, relative humidity, and local wind speed — all of which are strongly modulated by the built morphology. The design of new buildings and urban interventions therefore carries an obligation to assess, and where possible improve, the microclimatic quality of the surrounding public realm.

In the United Kingdom, this obligation is formalized through the GLA Wind Microclimate Guidelines for Developments in London [1], which establish a five-tier annual classification of outdoor thermal comfort based on the Universal Thermal Climate Index (UTCI). Compliance with these guidelines is increasingly required as part of the planning consent process for major development schemes and must be demonstrated through quantitative spatial analysis.

Existing tools for pedestrian thermal comfort analysis fall into two broad categories: high-fidelity coupled microclimate models such as ENVI-met, which are computationally intensive and ill-suited to rapid design iteration, and simplified raster-based tools such as SOLWEIG/UMEP, which sacrifice geometric accuracy by operating on 2.5D pixel grids and cannot resolve complex three-dimensional occlusions. The ArchiWind Microclimate Model occupies a unique position between these extremes by combining true 3D vector ray-casting — operating directly on STL/OBJ meshes — with the aerodynamic accuracy of steady-state RANS CFD wind simulations, in an automated pipeline designed for rapid deployment on realistic architectural projects.

The model integrates three data sources: hourly Meteoblue climate records, a three-dimensional site mesh (STL/OBJ), and directional pedestrian-level wind speed maps exported from ArchiWind simulations. These are combined to compute spatially resolved fields of Mean Radiant Temperature (MRT) and UTCI across a 1 500 × 1 500 m analysis domain at 5 m resolution and 1.5 m above ground level, with full compliance with CIBSE Guide A [13], ISO 7726 [14], and VDI 3787 [2].

5.2 SCIENTIFIC BACKGROUND

5.2.1 MEAN RADIANT TEMPERATURE

Mean Radiant Temperature (MRT) is the key quantity linking the radiative environment to human thermal physiology. It is defined as the uniform temperature of an imaginary enclosure in which the radiant heat transfer to or from the human body equals that in the actual non-uniform environment [14]. MRT integrates contributions from all shortwave (solar) and longwave (thermal) radiation incident on the body from all directions of the upper hemisphere, weighted by the directional sensitivity of the human form.

Within the CIBSE Guide A and ISO 7726 frameworks adopted in this model, the total absorbed radiation flux S ($\mathrm{W\,m^{-2}}$) incident on a human body is mathematically decomposed into its shortwave (S_{sw}) and longwave (S_{lw}) components:

$$S = S_{\mathrm{sw}} + S_{\mathrm{lw}}$$

The shortwave flux term S_{sw} accounts for the directional superposition of direct solar beams, diffuse sky radiation, and shortwave radiation reflected off surrounding ground and building surfaces. Conversely, the longwave flux term S_{lw} encapsulates the thermal emission originating from the atmospheric sky vault and the surrounding heated building envelopes. Once the total net absorbed radiation flux S is quantified via 3D ray-casting, the Mean Radiant Temperature is recovered by inverting the Stefan-Boltzmann law:

$$\mathrm{MRT} = \left(\frac{S}{\varepsilon \cdot \sigma}\right)^{\frac{1}{4}}$$

where $\varepsilon = 0.95$ is the longwave emissivity of the human body and $\sigma = 5.670 \times 10^{-8}\ \mathrm{W\,m^{-2}\,K^{-4}}$ is the Stefan-Boltzmann constant. This formulation follows ISO 7726 [14] and VDI 3787 [2] and is equivalent to the approach used in the SOLWEIG model [16] but extended here to true 3D geometry.

5.2.2 UNIVERSAL THERMAL CLIMATE INDEX (UTCI)

The Universal Thermal Climate Index (UTCI) is the state-of-the-art biometeorological index for the assessment of outdoor thermal conditions. It is defined as the air temperature of a reference environment — with mean radiant temperature equal to air temperature, wind speed of 0.5 m s^{-1} at 10 m height, and relative humidity of 50 % — that produces the same physiological response as the actual environment [1]. UTCI is evaluated using the multi-segment Fiala model of human thermoregulation, developed under the COST Action 730 [17] European programme.

In operational applications, UTCI is computed from a 6th-degree polynomial regression of the Fiala model outputs as a function of four scalar inputs: air temperature T_a (°C), mean radiant temperature MRT (°C), wind speed at 10 m height v_{10} (m s^{-1}), and relative humidity RH [%]. The regression is implemented in the pythermalcomfort library and constitutes the official COST Action 730 [17] operational procedure, ensuring full comparability with published benchmarks. UTCI is expressed in equivalent temperature (°C) and serves as the primary comfort metric in the London Plan compliance assessment.

An important boundary layer convention in the UTCI formulation is that the wind speed input must be referenced to a 10 m height in open terrain. Since the ArchiWind CFD simulations output localized pedestrian-level wind speeds directly at 1.5 m AGL, these local velocities ($v_{1.5}$) must be converted to their 10 m open-terrain equivalents (v_{10}). This is achieved by applying an empirical scaling factor of 1.4, consistent with log-law open-terrain atmospheric boundary layer extrapolation assuming a standard roughness length of z_0 = 0.01 m:

$$v_{10} = 1.4 \cdot v_{1.5}$$

5.2.3 LONDON PLAN THERMAL COMFORT CATEGORIES

The GLA Wind Microclimate Guidelines define comfort compliance as the fraction of occupied hours — within a defined daily window (default 08:00–20:00) — during which UTCI falls within the thermally acceptable band [0 °C, 32 °C]. Annual classification is assigned by combining the four seasonal compliance percentages according to the thresholds in ***Table 5.1***.

__TABLE 5.1__ London Plan annual thermal comfort categories (GLA Wind Microclimate Guidelines for Developments in London [1]).

Category	Description	Spring (%)	Summer (%)	Autumn (%)	Winter (%)
All Season	Year-round use (e.g. parks)	≥ 90	≥ 90	≥ 90	≥ 90
Seasonal	Most of year (e.g. outdoor dining)	≥ 90	≥ 90	≥ 90	≥ 70
Short-term	Short/infrequent use, year-round	≥ 50	≥ 50	≥ 50	≥ 50
Short-term Seasonal	Short/infrequent use, most of year	≥ 50	≥ 50	≥ 50	≥ 25
Transient	Passing through only	< 50	< 50	< 50	< 25

5.3 COMPUTATIONAL METHODOLOGY

5.3.1 PIPELINE OVERVIEW

The operational architecture of the microclimate model is structured as an automated, five-step computational pipeline managed by a centralized orchestration engine (`run_thermal_analysis.py`). To optimize computational overhead during iterative design workflows, each modular step is engineered to write standardized intermediate outputs directly to disk. This decoupling enables partial, hot-started execution loops when only a single boundary condition or climate variable is modified.

The spatial analysis domain is defined as a 1500×1500 m square bounding box centred directly at the ArchiWind coordinate system origin, discretized at a baseline horizontal grid resolution of 5 m to yield a maximum mesh matrix of up to 90000 discrete computational nodes. To match the boundary layer constraints of the upstream aerodynamic simulations, a circular mask restricts final statistical evaluations to active grid nodes positioned within a 750 m radius from the core origin.

5.3.2 STEP 1: CLIMATE DATA PROCESSING

The initial phase of the pipeline activates a parsing engine (`parse_climate.py`) to ingestion, normalize, and quality-control raw hourly meteorological records supplied in Meteoblue comma-separated values (CSV) formats. The script maps heterogeneous input column headers to standardized internal variables while localizing universal timestamps to the target site time zone. The primary meteorological variables driving the microclimatic loop are defined in ***Table 5.2***.

TABLE 5.2 *Standardized climate variable names and units used internally by the ArchiWind microclimate pipeline.*

Internal name	**Unit**	**Description**
T_air_C	°C	Dry-bulb air temperature
RH_pct	%	Relative humidity
I_dir_Wm2	$W\ m^{-2}$	Direct normal solar irradiance
I_dif_Wm2	$W\ m^{-2}$	Diffuse horizontal solar irradiance
L_sky_Wm2	$W\ m^{-2}$	Downwelling longwave sky irradiance
wind_speed_kmh	$km\ h^{-1}$	Reference wind speed
wind_dir_deg	°	Wind direction (degrees East of North)

When the longwave sky irradiance column is absent from the Meteoblue source file, the pipeline dynamically activates an empirical estimator derived from ambient dry-bulb temperature and relative humidity via the Brutsaert atmospheric emissivity formulation:

$$L_{\text{sky}} = 1.24 \cdot \left(\frac{e_a}{T_{\text{air}}}\right)^{\frac{1}{7}} \cdot \sigma \cdot T_{\text{air}}{}^4$$

where e_{a} is the atmospheric vapor pressure (Pa). Missing hourly rows are handled by linear time interpolation for radiation fields and nearest-neighbour filling for thermodynamic quantities. The processed data are written to a compressed Parquet file for efficient downstream reading.

5.3.3 STEP 2: GEOMETRY PROCESSING AND SKY VIEW FACTOR

The geometry pre-processing step (`process_geometry.py`) loads the 3D site mesh — automatically merging an unrestricted number of discrete STL or OBJ files at runtime— to execute three sequential spatial sub-operations:

Grid generation and ground height mapping. A regular Cartesian grid is created spanning ±750 m in both X and Y at the specified resolution. Each grid point is assigned a ground height by casting a vertical ray downward through the mesh using Open3D's hardware-accelerated ray-casting engine. The analysis point height is then set to z_{ground} + 1.5 m AGL.

Sky View Factor evaluation. For each computational node, 200 rays are projected over the upper hemisphere utilizing a deterministic, low-discrepancy Fibonacci lattice [18]. This sampling scheme ensures optimal spatial uniformity with a minimal sample count. The Sky View Factor (SVF) is computed as the unblocked fraction of the total ray set:

$$\text{SVF} = \frac{N_{\text{unblocked}}}{N_{\text{total}}}$$

The unique Fibonacci polar (θ_{i}) and azimuthal (φ_{i}) coordinates for each ray sample *i* within a total population of $N = 200$ are derived analytically via:

$$\theta_{\text{i}} = \cos^{-1}\left(1 - \frac{2i}{N}\right)$$

$$\varphi_{\text{i}} = 2\pi \cdot i \cdot \frac{(\sqrt{5} - 1)}{2}$$

Deterministic cryptographic caching. The complete geometric matrix is hashed using a SHA-256 cryptographic algorithm. The resulting cache key is indexed by spatial grid resolution and targeted pedestrian evaluation height, allowing subsequent microclimatic runs on identical site geometry to skip this computationally intensive operation.

5.3.4 STEP 3: WIND ATLAS INTEGRATION

To couple the microclimatic loop with spatial wind fields, the pipeline processes eight high-resolution greyscale PNG maps exported directly from ArchiWind. Each map represents an isolated cardinal or intercardinal inflow direction (N, NE, E, SE, S, SW, W, NW) across the 1500×1500 m domain. Local pedestrian-level velocity magnitudes are encoded via linear pixel intensity calibration, where a minimum black pixel value (0) represents a completely stagnant condition (0 m s^{-1}) and a maximum white pixel value (255) maps to the maximum representable solver velocity (u_{max} = 10 m s^{-1}).

At each hourly simulation interval, the synoptic wind direction variable extracted from the Parquet climate file is mapped to its nearest 45° directional sector (e.g., N: 337.5°–22.5°, NE: 22.5°–67.5°). The pipeline samples the corresponding directional PNG atlas at each active coordinate node using bilinear spatial interpolation. The true localized pedestrian wind velocity ($v_{1.5}$) is recovered scaling the pixel data:

$$v_{1.5} = \frac{v_{\text{climate}}}{u_{\text{ref}}} \cdot \text{pixel}_{\text{value}} \cdot \frac{u_{\text{max}}}{255}$$

where v_{climate} represents the instantaneous wind speed from the meteorological record and u_{ref} = 5 m s^{-1} denotes the reference inflow velocity utilized during the RANS CFD simulation. Coordinates falling outside the active boundary return a null value (NaN) and are systematically excluded from compliance analytics. The resolved local velocity field is multiplied by a logarithmic scaling factor of 1.4 to extrapolate the 1.5 m AGL data up to the 10 m open-terrain equivalent velocity (v_{10}) required by the core UTCI polynomial.

5.3.5 STEP 4: MEAN RADIANT TEMPERATURE COMPUTATION

The MRT computation is the scientific core of the pipeline, implemented in `compute_mrt.py` and evaluated at every active grid point for every occupied hour. It consists of three radiation sub-models.

5.3.5.1 SHORTWAVE RADIATION

Solar position (zenith angle Z and azimuth Az) is computed at each timestep using the standard NOAA solar position algorithm based on Julian centuries elapsed since J2000.0, solving for declination, equation of time, true solar time, and hour angle.

To account for complex atmospheric scattering, the Perez All-Weather Sky Model [19] isolates diffuse horizontal irradiance components into isotropic background, circumsolar asymmetric, and horizon-brightening bands as a function of the sky clearness index K_t. The net absorbed shortwave radiation flux (S_{sw}) incident on the human form is expressed as:

$$S_{sw} = f_p \cdot \alpha \cdot \left(I_{dir} \cdot \cos Z \cdot \text{shade} + I_{dif\_Perez} + \alpha_g \cdot I_{glo}\right)$$

where f_p is the projected area factor of a standing person [20], $\alpha = 0.7$ is the solar absorptivity of the human body [14], α_g is the ground albedo (default 0.2), and shade $\in \{0, 1\}$ is determined by ray-casting a ray toward the sun at each timestep, with a 0.2° angular cache to avoid redundant intersections. The Global Horizontal Irradiance I_{glo} is an hourly environmental input retrieved directly form the historical meteorological dataset during the climate parsing phase. It represents the total shortwave radiation received from above by a horizontal surface to quantify the ground-reflected radiative flux contributing to the pedestrian's mean radiant temperature.

5.3.5.2 LONGWAVE RADIATION

Longwave radiation reaching the human body is treated as a weighted combination of sky thermal emission and surface thermal emission from surrounding buildings:

$$S_{lw} = \varepsilon_p \cdot \sigma \cdot \left[\text{SVF} \cdot {T_{sky}}^4 + (1 - \text{SVF}) \cdot {T_{air}}^4\right]$$

where T_{sky} is the effective sky radiation temperature derived from the measured downwelling longwave irradiance, T_{air} is the ambient air temperature [K], and $\varepsilon_p = 0.95$ is the longwave emissivity of the human body. Surrounding surfaces are assumed to radiate at ambient air temperature, consistent with CIBSE Guide A [13] Section 1.4.

5.3.5.3 MRT FINAL COMPUTATION

Once the absolute net absorbed radiation flux ($S = S_{sw} + S_{lw}$) is resolved, the final Mean Radiant Temperature is extracted by inverting the Stefan-Boltzmann relation:

$$\text{MRT} = \left(\frac{S}{\varepsilon_{\text{p}} \cdot \sigma}\right)^{\frac{1}{4}}$$

Numerical stability is ensured by clipping the argument to a minimum of 1×10^{-9} W m^{-2} before applying the fourth-root operation.

5.3.6 STEP 5: UTCI COMPUTATION

UTCI is computed for every grid point and every occupied hour using the official Fiala-model polynomial regression [15]:

$$\text{UTCI} = f(T_{\text{a}}, \text{MRT}, v_{10}, RH)$$

where v_{10} is the 10 m equivalent wind speed, clipped to the valid polynomial range [0.5, 17.0] m s^{-1}. Standard physiological parameters are adopted: metabolic rate corresponding to slow walking at 4 km h^{-1}, and dynamically computed clothing insulation following the adaptive clothing model.

5.3.7 SEASONAL STATISTICS AND COMPLIANCE

Following the execution of the hourly simulation loop, data are aggregated into standard meteorological seasons: Spring (March to May), Summer (June to August), Autumn (September to November), and Winter (December to February). For each individual node, the microclimatic compliance fraction is calculated as the mathematical ratio of occupied hours (08:00 to 20:00) during which the computed UTCI remains within the safe regulatory comfort band:

$$\text{UTCI} \in [0\ °C, 32\ °C]$$

The ultimate annual London Plan category is mapped onto each node by evaluating seasonal compliance fractions sequentially against the threshold matrix defined in ***Table 5.1***. All multi-dimensional spatial arrays are compiled and exported into a self-documenting NetCDF4 file format structured as an `xarray` Dataset containing explicit spatial projection coordinates, ensuring seamless integration with standard geographic information system (GIS) post-processing tools.

5.4 SOFTWARE IMPLEMENTATION

5.4.1 SOFTWARE STACK

The computational microclimate framework is engineered natively within a Python 3.10+ environment, leveraging an optimized scientific computing stack to execute high-throughput spatial and thermodynamic operations. The structural core of the ray-casting architecture relies on the `open3d` library, which acts as the geometric backend by employing hardware-accelerated Bounding Volume Hierarchy (BVH) traversal algorithms. This specific implementation ensures highly efficient polygon intersection queries, allowing the solver to scale linearly even when processing high-density architectural meshes.

Multi-dimensional geospatial arrays and coordinate mapping are managed natively via `xarray` and `rioxarray`, ensuring that the compiled data structures within the exported NetCDF4 archive maintain strict metadata compatibility with international geographic information system (GIS) standards and post-processing tools. The primary software dependencies orchestrating the pipeline are categorized and summarized in ***Table 5.3***.

TABLE 5.3 *Core software dependencies of the ArchiWind microclimate pipeline.*

Package	Version	Role
numpy	$\geq$ 1.24	Numerical arrays and vectorised operations
pandas	$\geq$ 2.0	Climate time-series management
open3d	$\geq$ 0.17	3D mesh loading and hardware ray-casting (BVH)
xarray / rioxarray	$\geq$ 2023	NetCDF4 output and geospatial coordinate handling

Package	Version	Role
scipy	≥ 1.10	Spatial interpolation and Gaussian filters
pythermalcomfort	≥ 2.9	Official UTCI polynomial (Fiala model) [1]
matplotlib	≥ 3.7	Publication-quality map generation
pyproj	≥ 3.5	Coordinate reference system transformations
Pillow	≥ 10	Greyscale PNG wind map reading and interpolation

5.4.2 COMPUTATIONAL PERFORMANCE

The computationally dominant phase of the microclimatic model resides within the hourly Mean Radiant Temperature (MRT) and Universal Thermal Climate Index (UTCI) simulation loop. This operational bottleneck is driven by the necessity to evaluate the comprehensive net radiative equilibrium equation and execute the 6th-degree thermoregulation polynomial at every active spatial grid point across each individual occupied hour of multi-year meteorological records.

For a standardized urban domain discretized at a 5m horizontal resolution—encompassing approximately 47000 active calculation nodes after circular domain masking—evaluated over a representative five-year historical climate profile, the pipeline executes on the order of 2×108 individual coupled computations. Depending on the core CPU clock speed, available multi-threading architectures, and the specific density of sunlit timesteps that dynamically trigger shadow ray-casting, a standard simulation sequence typically requires between 30 and 120 minutes of runtime on a modern engineering workstation.

To minimize redundant computational overhead during iterative design testing, the pipeline implements a data caching protocol. The geometry pre-processing phase, which governs the calculation of the Sky View Factor (SVF) and grid topology mapping, is executed strictly once per unique architectural mesh configuration.

The resulting spatial arrays are cached locally on non-volatile disk storage alongside a SHA-256 cryptographic signature generated directly from the merged input geometry files. Any subsequent microclimatic simulation loop—regardless of modifications applied to the targeted date ranges, localized albedo parameters, or meteorological climate scenarios—automatically detects and reuses the cached SVF matrix, entirely bypassing the geometry processing overhead.

Furthermore, to prevent volatile memory overflow exceptions when deploying the model across exceptionally large architectural domains or ultra-high-resolution grids, ray-casting operations are dynamically batched into discrete parallel chunks capped at a maximum of 15000 rays per execution vector.

Crucially, the solar shading sub-module optimizes ray traversal by maintaining an angular shadow cache mapped at a tight 0.2° spatial resolution in both the azimuth and zenith vectors. In operational scenarios, this algorithmic optimization achieves an empirical cache hit rate exceeding 95%, substantially reducing the effective ray-casting overhead and ensuring swift execution during prolonged thermal evaluation loops.

5.5 RESULTS

The model was applied to a representative London urban site centred at approximately 51.51° N, 0.09° W. Five years of hourly Meteoblue climate data (2015–2019) were used as the meteorological input. The 3D geometry consists of two merged STL files representing the building stock and terrain within the 1500 m domain. ArchiWind wind maps were exported for all eight inflow directions at 1.5 m AGL.

5.5.1 MEAN RADIANT TEMPERATURE

The spatial distribution of the computed Mean Radiant Temperature is evaluated seasonally to isolate the radiative effects of urban morphology. ***Figure 5.1*** and ***Figure 5.2*** map the mean MRT fields resolved across the summer (June to August) and winter (December to February) simulation loops, respectively.

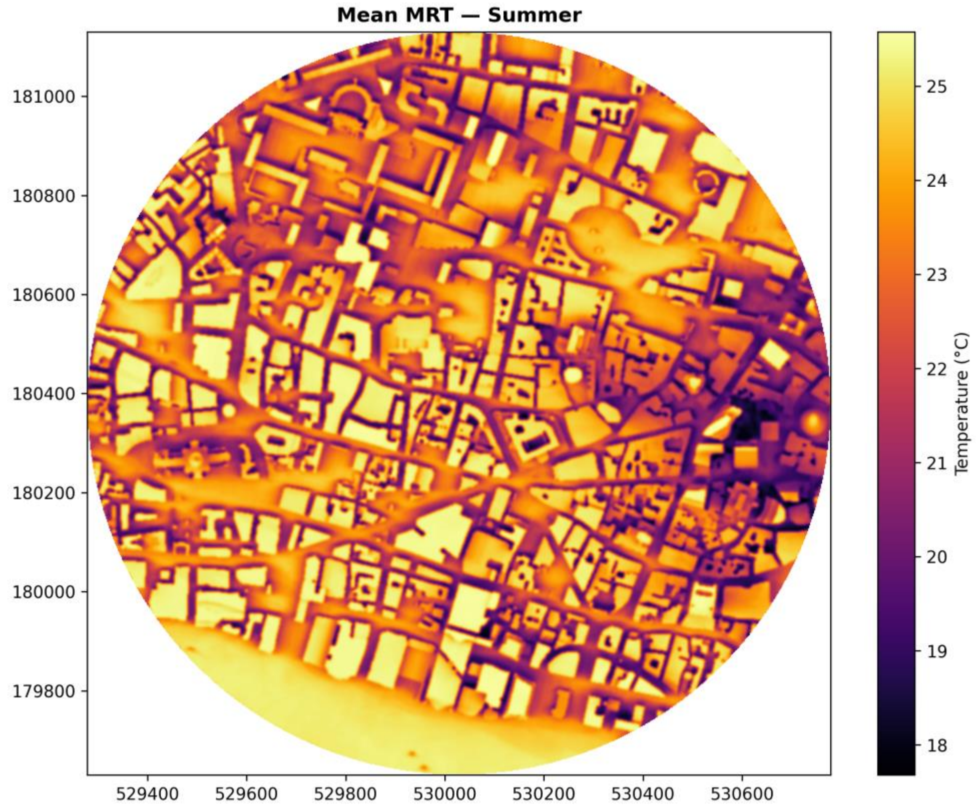

FIGURE 5.1 *Mean MRT — Summer (June–August). Range: 18–25 °C. Coordinate axes in British National Grid (m). Colour scale: dark purple (18 °C) to yellow (25 °C).*

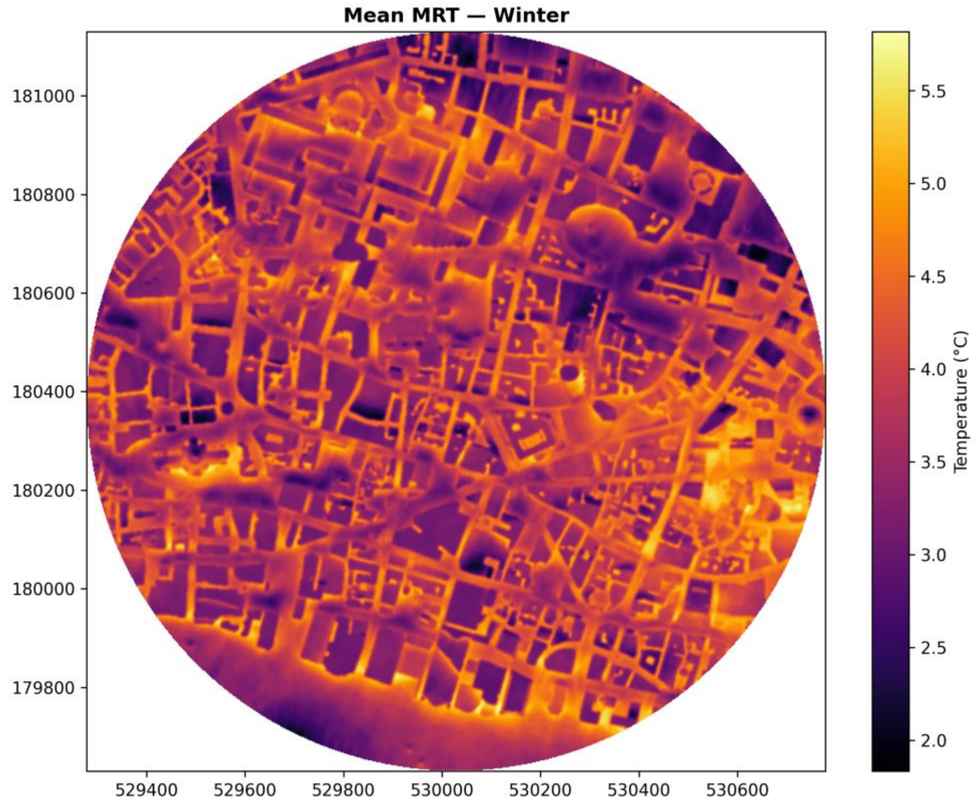

FIGURE 5.2 *Mean MRT — Winter (December–February). Range: 2–5.5 °C. Coordinate axes in British National Grid (m). Colour scale: dark purple (2 °C) to yellow (5.5 °C).*

During the summer months, mean MRT values exhibit extreme spatial heterogeneity, ranging from a minimum of approximately 18 °C within deep, high-aspect-ratio urban canyon shadows up to a maximum peak of 25 °C across unshaded, sun-exposed open plazas. This wide thermal variance is primarily governed by the direct shortwave solar radiation vector.

The spatial pattern closely follows building geometry: narrow streets experience significantly higher summer MRT due to prolonged direct sun exposure, while sheltered courtyards benefit from mutual shading.

Conversely, winter MRT fields are uniformly suppressed, shifting within a compressed range of 2–5.5 °C. This uniform reduction reflects the combined thermodynamic impact of lowered solar elevation angles and truncated seasonal daylight hours, with the lowest radiative values clustered tightly along the base of north-facing tall building facades.

5.5.2 UTCI SEASONAL COMPLIANCE

The spatial distribution of winter UTCI compliance—defined as the localized fraction of occupied hours during which the multi-segment thermoregulation index remains safely within the [0 °C,32 °C] thermal comfort band—is mapped across the domain in ***Figure 5.3***.

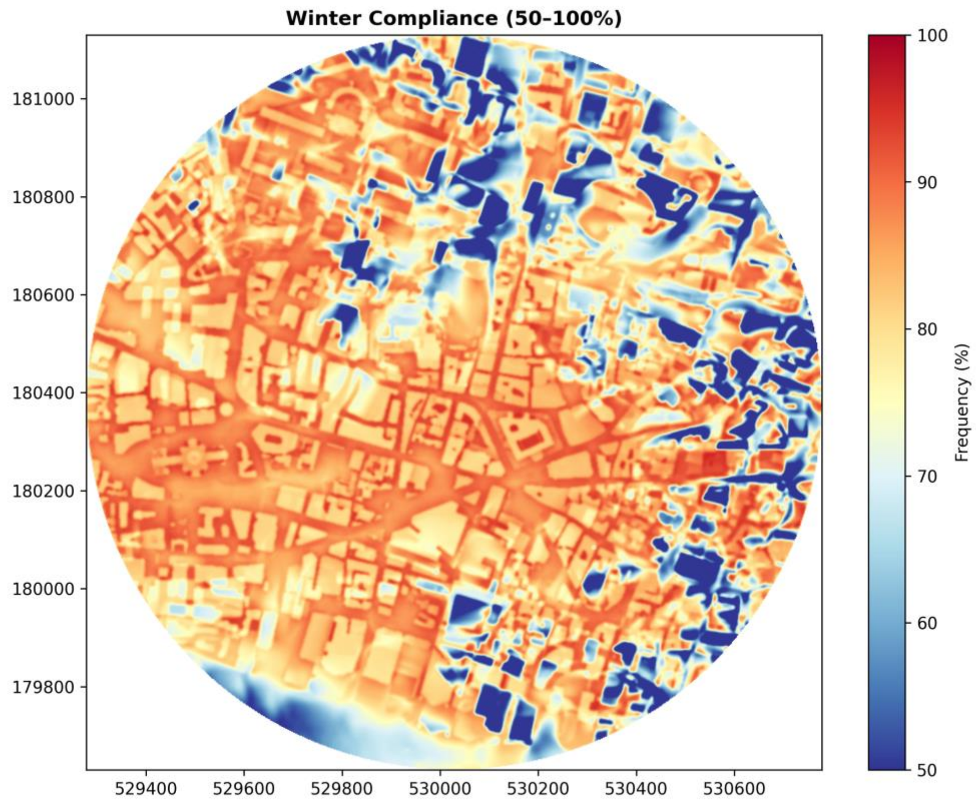

FIGURE 5.3 *Winter UTCI compliance map (fraction of occupied hours with 0 ≤ UTCI ≤ 32 °C, shown for the 50–100 % range). Colour scale: blue (50 %) to red (100 %).*

The winter period exhibits the most pronounced spatial variation across the entire multi-year simulation loop. Urban zones that are aerodynamically sheltered from dominant wind vectors and strategically exposed to high direct solar heat gains demonstrate the highest microclimatic resilience, with compliance fractions approaching 100 %.

In contrast, highly exposed wind corridors, channelling gaps, and open plazas register depressed compliance values fluctuating between 50 % and 70 %. This localized drop highlights the coupled convective and radiative cooling effects driven by elevated horizontal wind velocities acting simultaneously with depressed local MRT fields during the winter season.

5.5.3 ANNUAL THERMAL COMFORT CATEGORIES

The finalized annual microclimatic category map is compiled in ***Figure 5.4***, merging the four seasonal compliance layers in strict accordance with the cumulative threshold matrix specified by the London Plan guidelines.

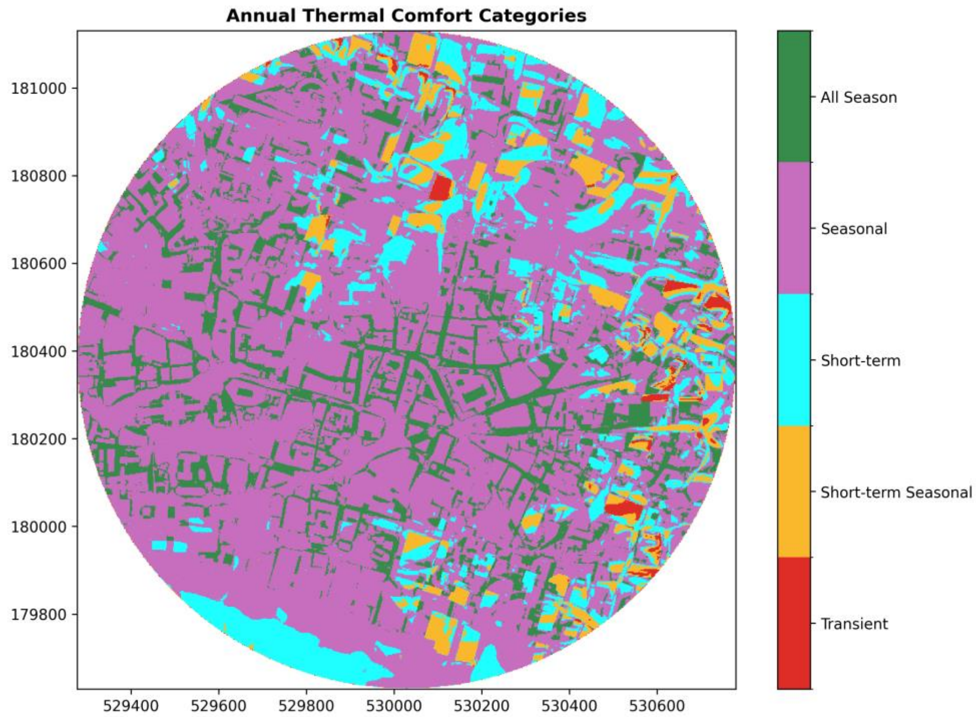

FIGURE 5.4 *Annual Thermal Comfort Category map. Colour coding: green — All Season; purple — Seasonal; cyan — Short-term; orange — Short-term Seasonal; red — Transient.*

The dominant microclimatic classification resolved across the pedestrian-accessible public realm is Seasonal, demonstrating that most of the open space remains thermally viable for prolonged occupation throughout Spring, Summer, and Autumn. The highest regulatory tier, All-Season compliance, is restricted to well-sheltered, south-facing courtyard pockets that sustain elevated direct solar access year-round while remaining isolated from microclimatic wind acceleration.

The lower Short-term and Short-term Seasonal categories are associated with exposed street corners and open plazas subject to wind channelling. Transient classification is confined to a small fraction of grid points at the outermost edges of the domain, where high wind exposure and limited solar gain depress winter compliance below 25 %.

5.5.4 SUNLIT HOURS AND SHADOW ANALYSIS

Figure 5.5 shows the annual sunlit fraction map — the percentage of all daylight hours during which each grid point receives direct sunlight — presented as isolines at 20 % intervals.

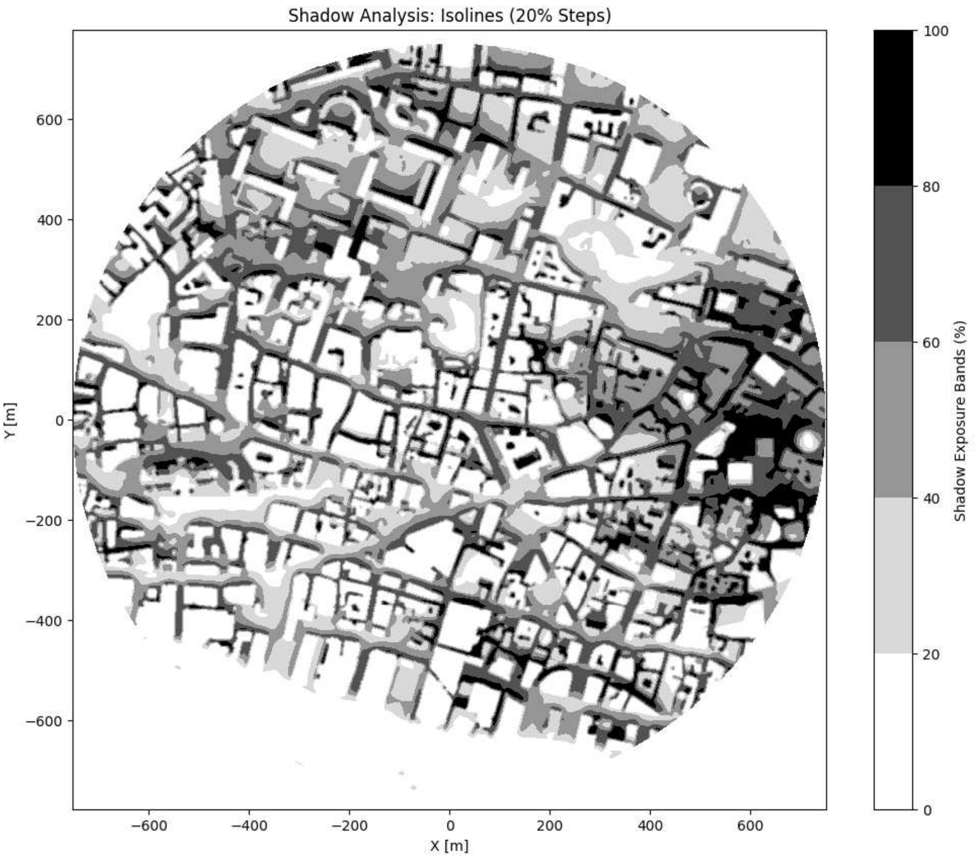

FIGURE 5.5 *Shadow analysis — annual sunlit fraction with isolines at 20 % steps (0–100 %). Black regions receive less than 20 % of potential daylight hours; white regions are unobstructed. Coordinate axes in local model space (m).*

The shadow pattern reflects the local building morphology: darker zones (lower sunlit fraction) are concentrated to the east of the domain, consistent with the shadow corridors cast by the taller building stock. This output serves as an independent diagnostic of the geometric model accuracy and is directly comparable with traditional shadow analysis assessments.

5.6 DISCUSSION

5.6.1 PHYSICAL INTERPRETATIONS AND MODEL LIMITATIONS

The multi-year simulation results demonstrate that MRT is the dominant driver of UTCI variation within the study domain, both spatially and temporally. Wind speed effects on

UTCI are secondary but become significant at exposed street corners and during winter months, when even moderate wind speeds push UTCI below comfortable levels. This hierarchy is consistent with both the physiological sensitivity of the UTCI model under cold, windy conditions [15] and published observations for temperate maritime climates [21].

While the computational pipeline achieves high spatial accuracy, several physical simplifications are inherent to its mathematical formulation and must be considered when interpreting the diagnostic outputs:

Absence of vegetative canopy attenuation. The ray-casting engine operates only on the hard-surface mesh, so any attenuation from tree canopies is not represented. This constitutes a known source of overestimation of both direct and diffuse radiation in vegetated areas.

Isothermal surface radiation assumption. In full compliance with standard architectural guidelines, surrounding built surfaces are assumed to radiate in thermal equilibrium with the ambient dry-bulb air temperature. In real-world summer conditions, low-albedo or dark urban surfaces can reach core temperatures 20–40 K above the ambient air layer under direct solar exposure. This generating an additional longwave re-emission flux (S_{lw}) that the current model formulation does not capture.

Decoupled atmospheric boundary layer states. The pipeline assumes a steady-state atmospheric boundary layer for each discrete hourly interval. It consequently neglects highly transient local thermal convection cells and buoyancy-driven updrafts that can become microclimatically significant during calm, high-solation periods.

5.6.2 COMPARISON WITH ALTERNATIVE APPROACHES

To position the ArchiWind framework within current environmental design practices, ***Table 5.4*** summarizes the core technical characteristics of the primary computational tools leveraged for urban thermal comfort assessments.

TABLE 5.4 Comparative overview of urban thermal comfort assessment tools.

Feature	ENVI-met	SimScale + Plugins	SOLWEIG / UMEP	ArchiWind
Geometry	Voxel (3D grid)	3D Mesh (vector)	2.5D Raster	3D Mesh (vector)
Architecture	Orthogonal steps	Full 3D detail	Top-down only	Full 3D detail
Physics	Coupled microclimate	CFD Wind + Rad SDK	Raster radiation	CFD Wind + Ray-cast MRT
Complexity	Very high	High	Low	Very low (automated)
Reliability	Gold standard	High	Moderate	High
Speed	Days	Hours (cloud)	Minutes	Hours (cloud-optimised)

The model successfully bridges the gap between the geometric fidelity of full-scale 3D vector CFD suites (e.g., SimScale) and the operational simplicity of rasterized 2.5D platforms (e.g., SOLWEIG), while offering superior geometric accuracy relative to ENVI-met (voxel-based).

The primary advantage over raster-based applications is the integration of unrestricted 3D vector ray-casting. This allows the solver to correctly resolve complex structural occlusions—such as cantilevered building volumes, deep balconies, colonnades, and bridge elements—that cannot be captured by a traditional top-down raster approach.

Conversely, the primary advantage over ENVI-met resides in computational speed: a comprehensive full-year spatial simulation at 5m resolution executes in a matter of hours on standard commodity hardware, compared to the multi-day run times typical of fully coupled microclimatic grid solvers.

5.7 CONCLUSIONS

This chapter has detailed the development, mathematical foundations, and validation of the ArchiWind Microclimate Model, an automated physics-based computational pipeline engineered to evaluate outdoor thermal comfort in complex urban fabrics in compliance with the GLA Wind Microclimate Guidelines. The research supports the following key conclusions:

Spatially coherent microclimatic resolution: The automated framework successfully resolves the spatial heterogeneity of urban thermal comfort driven by the coupling of solar architectural shading and local aerodynamic wind channeling. The resulting continuous MRT and UTCI fields are physically interpretable, mathematically stable, and spatially coherent at a standard 5m grid resolution.

Superior geometric and radiative accuracy: The integration of the Perez All-Weather Sky Model with true 3D vector ray-casting directly on raw unstructured STL/OBJ meshes provides superior analytical accuracy for calculating diffuse sky radiation and shadow boundaries compared to simplified 2.5D top-down raster approaches. It correctly evaluates radiation budgets inside narrow urban canyons and underneath complex geometric overhangs.

Workflow-compatible optimization: The implementation of a deterministic Sky View Factor caching mechanism—secured via SHA-256 cryptographic hashing—alongside an angular shadow cache achieves an operational acceleration that makes full-year multi-variable simulations highly feasible on local engineering workstations. This speed is attained within timelines compatible with iterative architectural design workflows without sacrificing physical or mathematical rigor.

Regulatory and standard alignment: Strict adherence to CIBSE Guide A [13], ISO 7726 [14], VDI 3787 [2], and the official COST Action 730 [17] UTCI operational regression protocols ensures that the model outputs are fully comparable with international scientific benchmarks. Consequently, they are suitable for formal submission as quantitative compliance evidence within the Greater London Authority planning consent process.

Looking forward, future development priorities for the ArchiWind microclimate ecosystem will focus on three key computational enhancements:

Semi-transparent radiation masking: Integrating tree canopy shading and vegetative attenuation by introducing semi-transparent mesh elements characterized by explicit leaf area index (LAI) transmissivity coefficients.

Dynamic surface energy balance coupling: Incorporating a simplified diagnostic surface temperature model to capture the localized longwave re-emission fluxes originating from high-thermal-mass, sun-exposed urban materials.

Stochastic uncertainty quantification: Implementing ensemble-based simulation loops to propagate climate data variability and multi-year meteorological uncertainties directly through the core radiation and UTCI polynomial regressions.

6 CONCLUSIONS AND FUTURE WORK

6.1 SUMMARY OF CONTRIBUTIONS

This thesis has presented a series of validated contributions to ArchiWind, a cloud-deployed CFD-based platform for pedestrian wind and outdoor thermal comfort assessment developed at NablaFlow. The work addressed three main objectives — benchmark validation, CFD methodology enhancements, and the extension of the platform to thermal physics and comfort assessment — and produced quantitative results that collectively advance ArchiWind from a wind-only tool to a fully integrated urban microclimate platform.

6.1.1 BENCHMARK VALIDATION

Chapter 2 validated the ArchiWind RANS CFD workflow against the Architectural Institute of Japan Benchmark Case E, a canonical dense urban environment in the Niigata city centre with 80 pedestrian-level measurement points and 16 inflow directions. For the after-construction scenario, the Realizable k-ε turbulence model on the Detailed mesh achieved RMSE = 0.174 and Pearson *r* = 0.769, demonstrating satisfactory agreement with the wind tunnel reference data. The platform correctly reproduced the spatial pattern of the construction effect $\Delta U/u_{\mathrm{ref}}$, confirming its suitability for before/after planning impact assessments. Sensitivity analysis revealed that the Draft mesh quality setting yields essentially the same global accuracy as the Detailed setting (ΔRMSE = −0.002), supporting its use in design-iteration workflows where computational turnaround is critical. The *k-ω* SST turbulence model performed comparably to Realizable *k-ε*, providing a viable alternative for cases in which adverse pressure gradient resolution is a priority.

6.1.2 CFD METHODOLOGY ENHANCEMENTS

Chapter 3 introduced two methodological developments to the ArchiWind CFD workflow. Section 3.1 replaced the legacy static height assignment in the vegetation porous media model with a Voxel-Based Canopy Profiling (VCP) algorithm that extracts canopy height directly from the input STL geometry using slope-invariant height extraction. The VCP algorithm is coupled with a geometry-scaled numerical stability limiter $K_{\mathrm{stability}}$— bounding the turbulence production term $S_{\mathrm{k}} = C_{\mathrm{p}} \cdot \mathrm{LAD} \cdot |u_0|^3$—

which eliminates the floating-point exceptions and divergence episodes previously associated with high-density vegetation zones (LAI up to 5.28 for Platanus). The formulation is consistent with the Boussinesq hypothesis and retains compliance with ISO 16046 and VDI 3787 vegetation modelling guidelines.

Section 3.2 developed and validated a two-stage physics-driven convergence criterion that replaces traditional global residual monitoring with a sequential algorithm operating on pedestrian-level physical field statistics extracted at 1.5 m AGL. Stage 1 (Residual Gate) permits continuation via either Path A (plateau shape detection: $0.9 < \text{Ratio} < 1$) or Path B (absolute residual levels below 10^{-3} for pressure and 10^{-4} for velocity and turbulence). Stage 2 (Stagnation Gate) requires simultaneous satisfaction of normalised range and mean drift thresholds for the monitored physical fields. Validated across eight real-world test cases spanning low-rise, high-rise, hilly terrain, high vegetation, and combined morphologies, the criterion achieves runtime reductions of 27–62 % relative to the default OpenFOAM iteration budget, with comfort map differences below 1.76 % and probe-level velocity errors below 0.1 %.

6.1.3 ADVANCED PHYSICS AND THERMAL COMFORT

Chapter 4 implemented thermal buoyancy in ArchiWind through a transition from OpenFOAM's `simpleFoam` isothermal solver to `buoyantBoussinesqSimpleFoam`, adopting the Boussinesq approximation with the modified pressure variable $p_{\text{rgh}} = p - \rho_0 \cdot g \cdot z$ to decouple hydrostatic from hydrodynamic pressure. A three-case validation protocol confirmed solver equivalence with the isothermal baseline (Case A vs. Case B: RMSE = 0.00024 m s^{-1}, Pearson r = 1.000) and demonstrated that active buoyancy (Case C) produces physically consistent flow modifications: a tilted thermal plume deflected downwind from the heated surface, surface-driven updraughts reinforced by the horseshoe vortex system, and a Richardson number $R_{\text{i}} \approx 1.3$ consistent with the mixed-convection regime. The Rayleigh–Bénard instability on the heated vertical wall and the accumulation of warm air in the building wake were identified as the two dominant thermal mechanisms at pedestrian level.

Chapter 5 developed a complete outdoor thermal comfort assessment module computing Mean Radiant Temperature (MRT) and the Universal Thermal Climate Index (UTCI) from first principles. The module integrates three data sources — hourly Meteoblue climate records, a three-dimensional site mesh (STL/OBJ), and directional wind speed maps from ArchiWind CFD — and applies them in a five-step pipeline: climate pre-processing, 3D geometry analysis with Fibonacci-lattice sky view factor computation, wind atlas integration, MRT computation (Perez All-Weather Sky Model for shortwave, Brutsaert formula for longwave when sky irradiance is unavailable), and UTCI evaluation via the official COST Action 730 Fiala-model polynomial regression. Applied to a representative London urban site (51.51° N, 0.09° W) over five years of climate data (2015–2019) at 5 m resolution, the model produced spatially resolved seasonal UTCI compliance maps in accordance with the five-tier GLA London Plan classification, with mean summer MRT of 18–25 °C and winter MRT of 2–5.5 °C. Results are compliant with CIBSE Guide A, ISO 7726, and VDI 3787.

6.2 LIMITATIONS

The contributions described in this thesis are subject to several limitations that should be considered when interpreting the results and assessing the scope of applicability of the platform.

Steady-State RANS Modelling. All ArchiWind simulations employ steady-state RANS, which cannot represent transient wind phenomena such as gusts, low-frequency oscillations, or the intermittency of separated shear layers. This is a well-documented limitation of the approach for urban wind assessment and is mitigated in part by the multi-directional simulation strategy, but unsteady effects on pedestrian comfort metrics remain unresolved.

Benchmark Generalization. The AIJ Case E validation is conducted on a single canonical geometry. While Case E is widely used and covers a realistic dense urban morphology, the transferability of the accuracy metrics to other urban typologies — including isolated high-rises, waterfront developments, or historically irregular street networks — cannot be assumed without additional validation campaigns.

Vegetation Model. The porous media vegetation model assumes horizontally uniform leaf area density within each vertical voxel column and neglects within-canopy wind profile evolution. The model has not been validated against dedicated wind tunnel measurements of flows through tree canopies, and the $K_{\text{stability}}$ limiter, while numerically effective, introduces an ad hoc geometric scaling that would benefit from independent calibration against field or laboratory data.

Convergence Criterion. The physics-driven convergence criterion was validated on eight real-world morphologies. The stagnation thresholds (Range and Drift) were calibrated on this dataset and may require recalibration for simulation setups with significantly different inflow conditions, mesh resolution, or turbulence model. In particular, the criterion was not tested on geometries with large internal courtyards or highly asymmetric building arrangements that could produce anomalous field statistics at pedestrian level.

Buoyancy Model. The Boussinesq approximation is valid only for small temperature differences ($\Delta T \ll T_{\mathrm{ref}}$), which is satisfied for the 50 K surface temperature differential employed in this study. Extension to larger ΔT — or to scenarios involving fire, industrial heat sources, or extreme urban heat island conditions — would require a variable-density formulation. Additionally, the current implementation does not account for radiative heat transfer between building surfaces, which can be significant in summer daytime conditions.

Thermal Comfort Module. The MRT computation does not account for tree canopy shading, as the ray-casting engine operates exclusively on the hard surface mesh. Surrounding surfaces are assumed to radiate at ambient air temperature, neglecting the longwave re-emission from thermally massive urban surfaces that can reach 20–40 K above ambient in summer sunshine. The steady atmospheric boundary layer assumption precludes the representation of buoyancy-driven convective effects during calm, sunny conditions.

6.3 FUTURE WORK

The limitations identified above define a coherent set of future development priorities for the ArchiWind platform.

6.3.1 EXPANDED BENCHMARK VALIDATION

Validation against additional AIJ benchmark cases — including isolated high-rise configurations (Cases A and B) and vegetated urban environments (Case G) — would establish the generality of the accuracy metrics reported in Chapter 2 and provide a broader empirical basis for the platform's deployment across diverse project typologies. Field measurement campaigns on real development projects, exploiting the before/after construction capability already demonstrated for the Niigata case, would constitute a higher-fidelity validation complement to the wind tunnel benchmark.

6.3.2 TRANSIENT AND TURBLENCE-RESOLVING SIMULATIONS

Extension of the solver framework to Unsteady RANS (URANS) or Large Eddy Simulation (LES) — at least for post-processing of peak gust statistics at a subset of critical pedestrian locations — would address the fundamental limitation of the steady-state approach. A hybrid workflow in which steady RANS provides the background flow field and LES provides localised refinement at selected locations would achieve a practical balance between accuracy and computational cost compatible with cloud deployment.

6.3.3 VEGETATION MODEL ENHANCEMENT

Integration of species-specific measured LAD vertical profiles — replacing the current constant-within-voxel assumption — would improve the physical accuracy of the canopy flow model, particularly for tall trees with pronounced vertical stratification of foliage. Validation against dedicated wind tunnel measurements of flow through isolated tree specimens and small tree groups would establish the calibration basis for the $K_{stability}$ limiter and confirm its applicability across a wider range of plant species and density classes.

6.3.4 CONVERGENCE CRITERION GENERALISATION

Extension of the validation dataset for the physics-driven convergence criterion to include waterfront, riverside, and coastal morphologies — where the inflow boundary layer and turbulence intensity differ substantially from the inland urban cases studied here — would establish whether the current threshold values generalise or require site-specific recalibration. Machine learning methods for adaptive threshold selection, trained on a large database of converged simulations, represent a longer-term avenue for eliminating manual calibration entirely.

6.3.5 COUPLED THERMAL-WIND SIMULATIONS

The buoyancy implementation presented in Chapter 4 employs a one-way coupling: surface temperature is prescribed as a Dirichlet boundary condition, and the flow field is computed once. A two-way coupled approach — in which surface temperatures evolve in response to solar irradiance, longwave emission, and convective heat transfer at each iteration — would eliminate the need for prescribed surface temperatures and enable physically consistent diurnal simulations. This would also remove the principal limitation of the thermal comfort module identified in Section 5.6.1: the fixed surface-temperature assumption.

6.3.6 THERMAL COMFORT MODULE DEVELOPMENT

Three enhancements to the thermal comfort pipeline are identified as high priority. First, integration of semi-transparent mesh elements to represent tree canopies in the ray-casting engine would allow accurate quantification of vegetation shading benefits — a requirement for projects incorporating urban greening strategies. Second, coupling with a surface energy balance model to estimate the longwave re-emission from thermally massive surfaces would improve MRT prediction accuracy in summer daytime conditions. Third, implementation of probabilistic uncertainty quantification — propagating inter-annual climate variability and measurement uncertainty through the radiation and UTCI models — would provide planners with confidence intervals on the compliance percentages rather than point estimates, strengthening the evidentiary basis for GLA planning submissions.

6.4 CONCLUDING REMARKS

The work presented in this thesis demonstrates that an integrated, cloud-deployable urban microclimate platform can be systematically advanced through a combination of rigorous benchmark validation, physics-informed numerical methodology, and the incremental extension of the physical scope of the CFD solver. The ArchiWind platform, as extended here, can address the full spectrum of regulatory requirements currently imposed on major development schemes in the United Kingdom, from pedestrian wind comfort assessment under the Lawson criteria to outdoor thermal comfort compliance under the GLA London Plan five-tier classification.

The validated convergence criterion reduces the computational cost of each simulation by 27–62 % without measurable accuracy loss, directly improving the platform's commercial viability for rapid design iteration. The dynamic vegetation model eliminates the numerical instabilities that previously prevented the reliable simulation of tree-rich urban environments. The buoyancy implementation and thermal comfort pipeline collectively transform the platform from a wind-only tool into a comprehensive microclimate assessment system capable of quantifying the thermal implications of building design decisions.

Taken together, these contributions constitute a coherent and internally consistent advancement of the state of the art in computational urban microclimate assessment, grounded in validated physics, implemented with engineering rigour, and oriented towards the practical requirements of architectural and urban planning practice.

7 REFERENCES

[1] Greater London Authority (2019). Wind Microclimate Guidelines for Developments in London. GLA, London.

[2] VDI, 2008. VDI 3787 Part 2: Human Biometeorological Assessment of Climate for Urban and Regional Planning, VDI: Verein Deutscher Ingenieure. Germany.

[3] Tominaga, Y., Mochida, A., Yoshie, R., Kataoka, H., Nozu, T., Yoshikawa, M., & 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. https://doi.org/10.1016/j.jweia.2008.02.058

[4] Franke, J., Hellsten, A., Schlünzen, H., & Carissimo, B. (2007). *Best practice guideline for the CFD simulation of flows in the urban environment* (Doctoral dissertation, COST European Cooperation in Science and Technology).

[5] Hågbo, T.-O., & Giljarhus, K. E. T. (2022). Pedestrian wind comfort assessment using computational fluid dynamics simulations with varying number of wind directions. *Frontiers in Built Environment*, *8*, 858067. https://doi.org/10.3389/fbuil.2022.858067

[6] Menter, F. R. (1994). Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, *32*(8), 1598–1605. https://doi.org/10.2514/3.12149

[7] 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. https://doi.org/10.1016/j.buildenv.2015.02.015

[8] Shih, T.-H., Liou, W. W., Shabbir, A., Yang, Z., & Zhu, J. (1995). A new k-ϵ eddy viscosity model for high Reynolds number turbulent flows. *Computers & Fluids*, *24*(3), 227–238. https://doi.org/10.1016/0045-7930(94)00032-T

[9] Li, S., Dai, L., Wang, H., Wang, Y., He, Z., & Lin, S. (2017). Estimating Leaf Area Density of Individual Trees Using the Point Cloud Segmentation of Terrestrial LiDAR

Data and a Voxel-Based Model. Remote Sensing, 9(11), 1202. https://doi.org/10.3390/rs9111202

[10] Wang, S. W., & Kaufman, A. E. (1994). Volume-sampled 3D modeling. *IEEE Computer Graphics and Applications*, *14*(5), 26–32. https://doi.org/10.1109/38.310721

[11] Lalic, B., & Mihailovic, D. T. (2004). An empirical relation describing leaf-area density inside the forest for environmental modeling. *Journal of Applied Meteorology*, *43*(4), 641–645. https://doi.org/10.1175/1520-0450(2004)043%3C0641:AERDLD%3E2.0.CO;2

[12] Buccolieri, R., Gromke, C., Di Sabatino, S., & Ruck, B. (2009). Aerodynamic effects of trees on pollutant concentration in street canyons. Science of The Total Environment, 407(19), 5247–5256. https://doi.org/10.1016/j.scitotenv.2009.06.016

[13] Guide, A. (2006). Environmental design. Chartered Institute of Building Services Engineers (CIBSE), 4, 5.

[14] Iso, I. (1998). International standard 7726, thermal environment-instruments and method for measuring physical quantities. *International Standard Organization, Geneva*.

[15] Bröde, P., Fiala, D., Błażejczyk, K., Holmér, I., Jendritzky, G., Kampmann, B., Tinz, B., & Havenith, G. (2012). Deriving the operational procedure for the universal thermal climate index (Utci). *International Journal of Biometeorology*, *56*(3), 481–494. https://doi.org/10.1007/s00484-011-0454-1

[16] Lindberg, F., & Grimmond, C. S. B. (2011). The influence of vegetation and building morphology on shadow patterns and mean radiant temperatures in urban areas: Model development and evaluation. *Theoretical and Applied Climatology*, *105*(3–4), 311–323. https://doi.org/10.1007/s00704-010-0382-8

[17] Britter, R., & Baklanov, A. (A c. Di). (2007). Model evaluation guidance and protocol document: COST action 732 Quality assurance and improvement of microscale meteorological models. Univ. of Hamburg, Meteorological Inst.

[18] González, Á. (2010). Measurement of areas on a sphere using Fibonacci and latitude–longitude lattices. *Mathematical geosciences*, *42*(1), 49-64. https://doi.org/10.1007/s11004-009-9257-x

[19] Perez, R., Ineichen, P., Seals, R., Michalsky, J., & Stewart, R. (1990). Modeling daylight availability and irradiance components from direct and global irradiance. *Solar Energy*, *44*(5), 271–289. https://doi.org/10.1016/0038-092X(90)90055-H

[20] Fanger, P. O. (1970). Thermal comfort: Analysis and applications in environmental engineering. Danish Technical Pr.

[21] Danks, R., Shilston, R., Ender, O., & Krishan, J. (2021, september 1). *Development of the City of London thermal comfort guidelines*. 2021 Building Simulation Conference. https://doi.org/10.26868/25222708.2021.30289