



GLIDE-SOL: A GPU-accelerated Global Lightweight Infrastructure for Diagnostic Environmental Modeling with SOLWEIG

Andrea Zonato¹, Harsh G. Kamath², Naveen Sudharsan², Luca Monaco¹, Jonas Kittner³, Luise Wolf³, Matthias Demuzere^{4,5}, Ariane Middel^{6,7}, Benjamin Bechtel³, and Massimo Milelli¹

¹CIMA Research Foundation, Savona, Italy

²University of Texas at Austin, Austin, USA

³Bochum Urban Climate Lab, Ruhr University Bochum, Germany

⁴B-Kode VOF, Ghent, Belgium

⁵Global Facility for Disaster Reduction and Recovery, World Bank Group, Washington, USA

⁶The GAME School, Arizona State University, USA

⁷School of Geographical Sciences and Urban Planning, Arizona State University, USA

Correspondence: Andrea Zonato (andrea.zonato@cimafoundation.org)

Abstract. GLIDE-SOL is a fully scripted and globally re-deployable Python workflow that operationalizes SOLWEIG for rapid and repeatable thermal-comfort mapping across diverse urban environments. GLIDE-SOL is built on the SOLWEIG radiative balance libraries, but rewrites the surrounding system—including automated input generation, the execution engine, and post-processing—so that the model can be driven by globally available datasets and executed efficiently on GPUs. All inputs (terrain, building morphology, canopy height, land cover, and meteorology) are automatically derived from global products, eliminating local preprocessing while enabling consistent applications from neighborhood-scale analyses to city-wide and multi-city experiments. In addition, GLIDE-SOL introduces lightweight physical diagnostics to improve realism when driven by coarse meteorological forcing, targeting key urban controls on wind and near-surface temperature.

The workflow incorporates two physical augmentations: (i) roughness- and obstacle-based directional wind attenuation to approximate near-surface ventilation; and (ii) diagnostic temperature adjustments that combine a simple urban heat island (UHI) cycle with an elevation-based correction using high-resolution DEM information, to better capture nocturnal warming and local lapse-rate effects.

To scale to large metropolitan areas, GLIDE-SOL uses explicit domain tiling with cross-tile synchronization to preserve radiative consistency across tile boundaries, enabling meter-scale simulations over tens to hundreds of square kilometers without sacrificing reproducibility. Daily outputs (24 radiative and meteorological fields) are stored as compressed GeoTIFFs to reduce disk usage and accelerate downstream processing.

GLIDE-SOL is implemented through three reproducible components: an automated global-input generator; a SOLWEIG execution engine with coordinated tiling; and a post-processing module for systematic sampling, time-series extraction, and visualization. An operational demonstration in Dortmund, using hourly measurements from 25 urban and peri-urban stations and simulations run at 2 m grid spacing between August 2024 and December 2025, shows that incorporating wind attenuation



and the diagnostic temperature corrections substantially improves UTCI performance (RMSE reduced from 9.9 °C to 2.7 °C), alongside improvements in T_{mrt} , air temperature, and wind speed simulations.

By integrating harmonized global inputs with physics-based diagnostics, GPU acceleration, and scalable tiling, GLIDE-SOL supports applications such as operational UTCI nowcasting, retrospective and climatological analyses of heat stress, sensitivity tests of urban morphology and greening strategies, and coordinated multi-city experiments requiring consistent modeling protocols.

1 Introduction

Fast and transferable urban thermal-comfort modeling is increasingly required for climate-adaptation planning, operational heat warnings, and systematic inter- and intra-city assessments. Exposure to urban heat stress is shaped by large-scale atmospheric conditions and fine-scale variations in radiation, wind, urban form, and surface cover, with pronounced differences between neighborhoods, street canyons, and open spaces (Emmanuel and Fernando, 2007; Mayer et al., 2008). Urban heat islands, particularly at night, and uneven daytime exposure to high radiant loads and locally reduced wind speeds contribute to elevated heat-stress risk (Oke et al., 2017). Radiant load, often expressed through mean radiant temperature (T_{mrt}), is a crucial meteorological driver of human energy balance and thermal comfort during warm, clear conditions (Mayer and Höppe, 1987; Thorsson et al., 2007). T_{mrt} is a key input to bio-climatic indices such as the Physiological Equivalent Temperature (PET) and Universal Thermal Climate Index (UTCI) (Höppe, 1999; Fiala et al., 2012; Jendritzky et al., 2012), which combine radiation, air temperature, humidity, and wind speed into a single thermal-stress metric (Cheng et al., 2012).

T_{mrt} arises from the net shortwave and long-wave radiation fluxes from all surrounding surfaces and the sky. While intra-urban air-temperature differences are often modest during the day, T_{mrt} can exhibit strong spatial variability over distances of only a few meters (Lindberg and Grimmond, 2011a, b; Thorsson et al., 2011; Middel and Krayenhoff, 2019). These microscale variations are primarily driven by shadow patterns created by buildings, trees, and topography, as well as by differences in surface radiative and thermal properties such as albedo, emissivity, and heat capacity (Thorsson et al., 2007; Konarska et al., 2014; Lindberg et al., 2016). The SOLWEIG (*SOLar and LongWave Environmental Irradiance Geometry*) (Lindberg et al., 2008) model is widely used to estimate T_{mrt} and related thermal-comfort indices in complex urban environments. It resolves shortwave and longwave radiation fluxes in three dimensions and has been applied to analyze the influence of urban geometry, ground surfaces, and vegetation on outdoor thermal comfort in a variety of climates, and validated against ground observations (Lindberg and Grimmond, 2011a, b; Briegel et al., 2024; Buo et al., 2023). SOLWEIG is also integrated in the Urban Multi-scale Environmental Predictor (UMEP) (Lindberg et al., 2018), which provides a QGIS-based environment for linking urban-climate models with geospatial data to support climate services, heat-risk assessments, and planning applications. Despite its strong theoretical foundation and broad adoption in case studies, SOLWEIG has not yet matured into a globally operational tool.

Several persistent barriers limit the wider deployment of SOLWEIG for city-wide or multi-city applications. First, realistic setups typically require detailed, locally assembled input data: high-resolution digital elevation/surface models, building land



cover, footprints and heights, vegetation structure, and representative meteorological forcing (Jänicke et al., 2016; Zonato et al.,
55 2020). These inputs are often derived from local LiDAR campaigns, city-specific GIS databases, or bespoke classifications,
which are not uniformly available across regions. Second, microscale radiative models can be computationally demanding
when applied at meter-scale resolution over large urban domains or long simulation periods (Toparlar et al., 2017; Briegel
et al., 2024). Third, near-surface wind is frequently prescribed from single meteorological stations or coarse reanalyses that do
not explicitly account for building-induced sheltering and channeling, even though wind strongly modulates thermal comfort
60 (Cionco, 1972; Chu and Wang, 2025). CFD models such as ENVI-met and PALM-4U can resolve pedestrian-level wind and
temperature fields (Simon, 2016; Maronga et al., 2020), but their computational cost generally restricts their use to micro-scale
domains or limited simulation periods.

In parallel, global geospatial and meteorological datasets have advanced to the point where they can serve as a common input
backbone for urban climate applications. ESA WorldCover provides 10 m global land-cover classes suitable for distinguishing
65 major urban and vegetated surfaces (Zanaga et al., 2022). On the terrain side, MERIT DEM (Yamazaki et al., 2017) offers
harmonized global elevation information at 3-arc-second (90 m) resolution, forming a consistent basis for topographic inputs.

Complementing these terrain-only products, a new generation of global building datasets has emerged, providing both
footprints and 3D attributes essential for urban climate modeling. OpenStreetMap (OpenStreetMap contributors, 2017) supplies
global building polygons, while dedicated building datasets now include UT-GLOBUS (Kamath et al., 2024), offering 3 m
70 building-height estimates worldwide; GlobalBuildingAtlas (Zhu et al., 2025), providing global LoD1 building geometries
at 3 m resolution; GeoClimate's height-completed OSM layers (Bernard et al., 2022), which add locally calibrated height
estimates; and the recently released OpenBuildingMap dataset (Oostwegel et al., 2025), which integrates multiple global
sources to produce harmonized building footprints enriched with semantic attributes such as height, floors, and function.
Together, these building datasets deliver near-global coverage of 3D urban morphology, enabling more realistic city- and
75 street-scale simulations worldwide.

At the same time, global canopy-height products have also matured rapidly: very-high-resolution (1 m) tree-height maps
derived from self-supervised deep learning trained on airborne LiDAR (Tolan et al., 2024), the global canopy height model by
Lang et al. (Lang et al., 2023) at 10 m resolution, and the GEDI-based Canopy Height Mapper (Potapov et al., 2021), which
delivers canopy heights at 30 m resolution through a Google Earth Engine workflow, collectively provide consistent estimates
80 of vegetation height and distribution. Together, these resources supply a coherent global basis for representing not only terrain
and buildings but also the three-dimensional distribution of urban and peri-urban vegetation.

On the meteorological side, the ERA5 reanalysis provides globally consistent, hourly atmospheric fields (Hersbach et al.,
2020). However, its coarse resolution inherently limits the representation of mesoscale and, especially, urban-scale processes:
even when dynamically down-scaled, ERA5 tends to smooth local thermal gradients and fails to capture city-specific features
85 and extreme conditions (Adinolfi et al., 2023). Platforms such as Google Earth Engine enable planetary-scale access and
processing of these datasets in a reproducible, scriptable manner (Gorelick et al., 2017).

Recent studies demonstrate that city-wide, building-resolving thermal-comfort mapping with SOLWEIG and UMEP is now
feasible using scripted workflows. Jänicke et al. (2016) performed a city-wide analysis of T_{mrt} at building resolution, while



more recent work has produced high-resolution thermal exposure and shade maps to support cool-corridor planning (Buo
90 et al., 2023). A key recent example is the study by Lindberg et al. (2025), who applied UMEP–SOLWEIG at 2 m resolution
across Stockholm, Gothenburg, and Malmö to quantify how building density, tree fraction, and ground cover shape PET, UTCI,
 T_{mrt} , and pedestrian wind speed during heatwaves.

At the same time, GPU-based approaches to urban heat mapping have shown that large-domain, high-resolution simulations
can be substantially accelerated. Li and Wang (2021) report over 99% speed-up in computing city-scale heat-exposure metrics
95 from high-resolution surface data. Yang et al. (2023) achieve near–real-time urban LES using a multi-GPU framework that
preserves building-resolving detail. Similarly, the multi-scale City-LES v2.0 model (Kusaka et al., 2024) exploits GPU clusters
to efficiently couple mesoscale and microscale LES, enabling long-duration, building-explicit simulations with full radiative
calculations. Likewise, Li et al. (2024) employs GPU-accelerated computation to assess heat sensitivity and vulnerability across
major U.S. cities, demonstrating that large-scale exposure analyses can now be carried out with operationally feasible runtimes.

100 In addition, recent open-source work has demonstrated how the core SOLWEIG radiative-geometry and time-stepping
calculations can be refactored for GPU execution. SOLWEIG-GPU (<https://github.com/nvnsudharsan/SOLWEIG-GPU>)
provides a Python implementation that runs SOLWEIG (v2022a) on GPUs via a simple API, CLI, or GUI, and targets meter-
scale, city-scale domains by combining domain tiling with GPU-parallelized computation of SVF, T_{mrt} , shadows, and UTCI
(while wall height and wall aspect remain CPU-side and are parallelized across cores).

105 Reported benchmarks show substantial acceleration for the SVF+UTCI core: for tile sizes of 1000×1000 , 1500×1500 , and
 2000×2000 cells, runtimes decrease from 1187, 3322, and 6487 s on a CPU workstation to 47, 105, and 158 s on an NVIDIA
A6000 GPU, corresponding to approximately $25\times$, $32\times$, and $41\times$ speed-ups. This work served as a practical reference point
for the GPU-oriented design choices adopted in GLIDE-SOL.

The goal of GLIDE-SOL (Global Lightweight Infrastructure for Diagnostic Environmental modeling with SOLWEIG) is
110 to transform SOLWEIG from a case-specific research tool into a reproducible, globally deployable workflow. GLIDE-SOL
rests on three pillars. First, all required input data are obtained from globally or regionally available datasets—including ESA
WorldCover (Zanaga et al., 2022), MERIT DEM (Yamazaki et al., 2017), GlobalBuildingAtlas (Zhu et al., 2025), and ERA5
accessed and processed via Google Earth Engine (Gorelick et al., 2017; Hersbach et al., 2020)—and automatically transformed
into SOLWEIG-ready rasters. Second, GLIDE-SOL incorporates diagnostic physical modules designed to compensate for
115 limitations of coarse meteorological forcing in urban settings: a Röckle-style directional wind-attenuation scheme (Röckle,
1990; Bernard et al., 2023) and a simple urban heat island adjustment informed by urban-climate theory (Theeuwes et al.,
2016). Third, GLIDE-SOL targets scalability and computational efficiency through GPU-accelerated radiative geometry and
explicit domain tiling. Entire cities, including large metropolitan areas spanning tens to hundreds of square kilometers, can thus
be simulated at meter-scale resolution, with daily outputs stored as compressed GeoTIFFs for efficient downstream analysis.

120 In this paper, we document the GLIDE-SOL workflow and demonstrate its capabilities in Dortmund, Germany, using a
dense observational network of 25 urban and peri-urban stations between August 2024 and December 2025. Simulations are
conducted at 2 m grid spacing over the full metropolitan area. We first describe the automated pre-processing pipeline and
its use of global datasets, then outline the SOLWEIG configuration, diagnostic wind and Urban Heat Island (UHI) modules,



and GPU tiling strategy, and finally present an evaluation of hourly UTCI, T_{mrt} , air temperature, and wind speed against
 125 observations. Beyond reporting skill metrics, we discuss operational aspects such as wall-clock time and memory usage, and
 consider how a global-input, SOLWEIG-based workflow can support operational heat-risk assessment, climatological analysis,
 urban-planning evaluation, and coordinated multi-city modeling experiments.

2 Materials and methods

This section presents the end-to-end methodology of the GLIDE-SOL framework. We first describe the SOLWEIG Preprocessor,
 130 which automatically builds a co-registered spatial input stack (rasters and vectors) and harmonizes ERA5 meteorological
 forcing into SOLWEIG-ready formats. We then outline the GPU-accelerated execution strategy, including tiling, memory
 management, and CPU/GPU responsibilities for scalable hourly simulations. Finally, we detail the additional physical parameterizations
 introduced in this work (diagnostic UHI cycle, elevation bias correction, and wind reduction coefficients) that improve realism
 while preserving a fully reproducible, script-driven workflow based on globally available data sources and standardized
 processing steps. Figure 1 provides an overview of the complete pipeline.

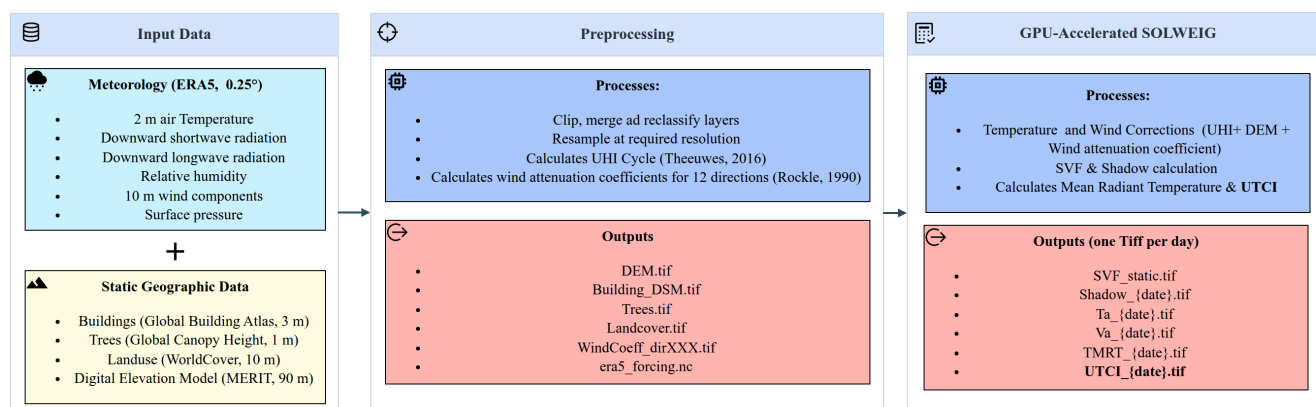


Figure 1. Workflow of the GLIDE-SOL framework.

135

2.1 Overview of the SOLWEIG core

SOLWEIG (*Solar and LongWave Environmental Irradiance Geometry*) is a three-dimensional radiative model developed to
 estimate mean radiant temperature (T_{mrt}) and related thermal-comfort quantities from urban geometry, vegetation, sun position,
 and near-surface meteorological forcing (Lindberg et al., 2008; Lindberg and Grimmond, 2011a, b; Lindberg et al., 2018). At
 140 each time step, the model combines direct, diffuse, and reflected shortwave radiation with longwave exchange from sky, ground,
 walls, and vegetation to resolve the radiative environment experienced by a pedestrian, from which T_{mrt} and derived indices
 such as UTCI can be computed. This representation has been widely used in urban-climate and outdoor-comfort studies,



including intercomparisons with other microclimate models and evaluations against observations (Aydin et al., 2019; Briegel et al., 2024).

145 In GLIDE-SOL, this logic is preserved but organized into an explicit, GPU-oriented execution sequence. The solver loads the static morphology layers (`Building_DSM`, `DEM`, `Trees`, `Landuse`, `Walls`, and `Aspect`), computes total and directional sky-view factors (SVF), evaluates building and vegetation shadow masks, and then advances the hourly SOLWEIG radiation balance on each tile before computing T_{mrt} , air temperature adjustments, wind speed, and UTCI. This implementation mirrors the standard SOLWEIG radiative structure while exposing the geometry and radiation terms needed for tiled GPU execution
150 in GLIDE-SOL.

2.2 Automated inputs with SOLWEIG Preprocessor

The input builder assembles a reproducible stack of rasters and vectors from a latitude/longitude seed, relying on globally accessible sources. The emphasis is to avoid bespoke digitization and to keep decisions transparent and scriptable:

- **Grid and Coordinate Reference System (CRS):** starting from a target latitude–longitude pair, the script automatically
155 selects the appropriate UTM zone, builds a geographic bounding box, and defines a regular metric grid (square pixels). All subsequent rasters (land cover, DEM, trees, buildings, wind coefficients) are reprojected and resampled onto this common grid, so that masking, sampling, and spatial statistics are perfectly co-registered.
- **Land cover, WorldCover reclassification, and elevation:** the script downloads ESA WorldCover 2020 version 2 (Zanaga et al., 2022) for the study area and reclassifies it into a small set of land-use classes tailored to SOLWEIG:
160
 - 1: urban (WorldCover code 50),
 - 2: vegetation (default class for all other codes),
 - 3: water (WorldCover codes 80, 70, 90),
 - 4: bare & sand (WorldCover code 60).

The reclassified raster is saved as `Landuse.tif`. In parallel, the script downloads and clips a digital elevation model (DEM), using a high-resolution global product such as MERIT DEM ((Yamazaki et al., 2017), 90 m), and resamples it to the same grid as `DEM.tif`. The DEM is later used to apply an elevation-based correction to ERA5 2 m air temperature (Section 2.4.2).
165

- **Urban form:** the script ingests LoD1 building heights from the Global Building Atlas ((Zhu et al., 2025), 3 m), together with building footprints and impervious surfaces from OpenStreetMap. These layers are rasterized on the common grid and stacked with land-use information, yielding `Buildings.tif` (building footprints) and `Building_DSM.tif`,
170 computed as `DEM + LoD1 building heights`.
- **Trees:** a high-resolution tree-canopy DSM is downloaded via Google Earth Engine from global canopy-height products, trying three sources in order of priority: (i) Meta Forest Monitoring / CanopyHeight ((Tolan et al., 2024), 1 m), (ii) ETH



175 Global Canopy Height 2020 ((Lang et al., 2023), 10 m), and (iii) a GEDI-based product ((Potapov et al., 2021), 30 m).
The selected canopy DSM is mosaicked and rasterized on the common grid to produce `Trees.tif`, representing tree
and woody vegetation height across the domain.

– **Wind coefficients:** using `Buildings.tif`, `Trees.tif`, the script computes 12 directional wind-sheltering rasters
`WindCoeff_dirXXX.tif` (one for each 30° sector) based on a Röckle-style wake model. For each wind direction,
the sheltering effect of buildings and trees is propagated downstream, then smoothed (e.g. via Gaussian filtering) and
180 clipped to the SOLWEIG domain to avoid artifacts (Section 2.4.3).

– **Meteorology:** meteorological forcing is harmonized from ERA5 reanalysis (Hersbach et al., 2020), retrieved via Google
Earth Engine (Gorelick et al., 2017). The key variables—near-surface air temperature at 2 m (T2M), relative humidity at
2 m (RH2M), downward shortwave and longwave radiation at the surface (SWDOWN and LWDOWN), surface pressure
(PSFC), the zonal and meridional wind components at 10 m (U10M and V10M), aerodynamic roughness length (Z0),
185 and surface altitude—are renamed and standardized onto a common time axis; from these, a diagnostic UHI cycle is
derived (Section 2.4.1). The final output is a NetCDF file ready to be used as forcing for SOLWEIG, such that the
downstream pipeline remains identical regardless of the original meteorological source.

2.3 GPU acceleration and parallelization

The GPU strategy in GLIDE-SOL is inspired by the open-source SOLWEIG-GPU implementation by Sudharsan et al. (https://github.com/nvnsudharsan/SOLWEIG-GPU/tree/main/solweig_gpu). While GLIDE-SOL re-architects the surrounding
190 workflow (global input generation, coordinated tiling, and post-processing), this prior work provided a valuable reference
point for GPU-oriented refactoring of the SOLWEIG radiative geometry and time-stepping.

GPU/CPU split and tiling

The most intensive geometric and radiative computations run entirely on the GPU: shadow casting, sky-view factor (SVF)
195 estimation, the hourly SOLWEIG radiation balance, and the UTCI calculation all take advantage of fast, auto-casted GPU
math. The CPU is responsible for coordinating the workflow. Before each run, GDAL splits the full spatial domain into
tiles—typically around one megapixel each—so that every piece comfortably fits within the available VRAM. Crucially, to
prevent edge effects during the simulation (particularly in neighborhood-dependent calculations such as the sky-view factor), a
10-cell padding is added to the perimeter of each tile. This ensures that the calculations for the cells near the tile boundaries use
200 accurate surrounding data. Each GPU processes its assigned tiles day by day. Once the geometric and radiative calculations are
complete, the added padding is removed during the final GeoTIFF clipping phase to ensure the output raster precisely matches
the original spatial domain, eliminating any boundary artifacts. One CPU process orchestrates the flow per device.



Minimizing host–device transfers

Data movement between CPU and GPU is deliberately kept small. For each tile, only the tile itself and its halo are transferred
205 from host to device. After the computation of an hour, the outputs—UTCI, mean radiant temperature, air temperature, wind
speed, and optionally shadows—are sent back to the CPU as float32 arrays. At the end of every simulated hour, both GPU and
CPU memory are cleaned (CUDA cache, Python garbage collection, malloc_trim) so that the memory footprint stays flat even
over long runs.

CPU I/O and deferred re-compression

210 All writing and compression duties remain on the CPU. Intermediate GeoTIFFs are written with light ZSTD compression and
marked as “ready” once all tiles for a given day and variable have arrived. A dedicated repacking thread then uses GDAL
CreateCopy to convert each ready file to LERC+ZSTD (variable-specific MAX_Z_ERROR budgets of 0.15 °C for UTCI,
0.10 °C for T_{mrt} , 0.005 m s⁻¹ for Va10m, and a 2 W m⁻² default for radiative terms), and writes a hidden “repacked” marker
when finished. GPU workers wait at day boundaries until these markers appear, ensuring I/O never lags behind the compute
215 tasks.

Memory stability

Overall, memory usage remains tightly controlled: the GPU never holds more than one (or a few) tiles at a time plus the 3D
shadow and SVF buffers, while CPU RAM is stabilized through one-time raster loads, periodic cleanup, and deferred heavy
compression. The result is a pipeline in which GPU computation and CPU I/O overlap smoothly, maximizing throughput
220 without ever exceeding VRAM or system memory limits.

2.4 Physical parameterizations

We summarize the three added parameterizations that improve on the baseline SOLWEIG formulation.

2.4.1 Diagnostic UHI cycle

To represent nocturnal warming that is not resolved in ERA5 near-surface meteorology, we impose a diagnostic UHI cycle that
225 perturbs the background air temperature before the SOLWEIG run. The daily maximum UHI amplitude is estimated using the
analytical formulation of Theeuwes et al. (2016), evaluated directly from ERA5 variables. For each day d , the maximum UHI
intensity is

$$\text{UHI}_{\text{max}}(d) = \left[\left(\frac{\overline{\text{SWDOWN}}_d}{\rho c_p} \right) \cdot \text{DTR}_d^3 \right]^{1/4}, \quad (1)$$

where $\overline{\text{SWDOWN}}_d$ is the daily-mean incoming shortwave radiation, DTR_d the daily temperature range, ρ the air density,
230 and c_p the heat capacity of air. The intra-night evolution is prescribed using an asymmetric sine-shaped function whose timing



follows observed UHI behavior. Rather than using fixed clock times, the night window is diagnosed directly from ERA5 shortwave radiation: it extends from the four hours before SWDOWN = 0 until the first two hours with SWDOWN > 0 after sunrise. Within this night window, the additive perturbation applied to air temperature is

$$UHI_{d,L}(t, x, y) = \begin{cases} [2 - SVF(x, y) - f_{veg}(x, y)] UHI_{max}(d) \sin\left(\pi \frac{s(t)}{L}\right), & t \in \text{night window}, \\ 0, & \text{daytime.} \end{cases} \quad (2)$$

235 where $SVF(x, y)$ is the sky-view factor, $f_{veg}(x, y)$ the vegetation fraction, $s(t)$ the elapsed time since the night-window start and L its duration. The modulation term $(2 - SVF - f_{veg})$ attenuates the applied UHI in open or highly vegetated areas, where nocturnal cooling is more effective.

This diagnostic scheme provides a computationally efficient way to introduce an approximate UHI effect—typically absent in ERA5 because of its coarse resolution and lack of an urban land-surface representation—while remaining compatible with
240 the global, automated nature of the GLIDE-SOL workflow.

2.4.2 Elevation bias correction

Before forcing SOLWEIG, 2 m air temperature is adjusted to the local DEM height through an effective, humidity-dependent lapse rate Γ_{eff} :

$$T_{2m,corr}(x, y) = T_{2m} + \Gamma_{eff}(T_{2m}, p, RH) \Delta H(x, y), \quad (3)$$

245 where p is air pressure, RH relative humidity from the driving meteorology, and ΔH the height difference between the ERA5 surface altitude and the high-resolution DEM, expressed as $\Delta H = H_{ERA5} - H_{DEM}$. In the implementation, Γ_{eff} linearly interpolates between the dry-adiabatic lapse rate $\Gamma_d = g/c_p \approx 9.8 \text{ K km}^{-1}$ and a moist-adiabatic lapse rate $\Gamma_m(T, p)$ computed from saturation specific humidity:

$$\Gamma_{eff} = (1 - RH)\Gamma_d + RH\Gamma_m(T, p), \quad (4)$$

250 Thus, the elevation correction smoothly transitions from dry- to moist-adiabatic behavior depending on humidity, and is applied at every time step before computing T_{mrt} and UTCI.

2.4.3 Wind reduction coefficients

We first harmonize the large-scale wind field from ERA5 and then derive directional, morphology-based wind reduction coefficients $C(\mathbf{x}, \theta)$.

255 Reference wind

ERA5 provides near-surface wind components at 10 m height (U_{10}, V_{10}) together with a roughness length $z_{0,ERA5}$. The script estimates a domain-average urban roughness length $z_{0,city}$ from buildings and trees, and rescales the 10 m wind from $z_{0,ERA5}$



to $z_{0,\text{city}}$ assuming a neutral logarithmic wind profile. For each wind component we apply

$$U_{10}^{\text{city}} = U_{10}^{\text{ERA5}} \frac{\ln(10/z_{0,\text{city}})}{\ln(10/z_{0,\text{ERA5}})}, \quad V_{10}^{\text{city}} = V_{10}^{\text{ERA5}} \frac{\ln(10/z_{0,\text{city}})}{\ln(10/z_{0,\text{ERA5}})}. \quad (5)$$

260 The resulting time series $(U_{10}^{\text{city}}(t), V_{10}^{\text{city}}(t))$ represent a spatially uniform, domain-consistent background canopy wind. In SOLWEIG this background wind is later multiplied by the spatially varying coefficients $C(\mathbf{x}, \theta)$ described below.

Directional wake model

Directional wind-reduction coefficients $C(\mathbf{x}, \theta)$ are computed with a simplified version of the Röckle wake model (Röckle, 1990; Bernard et al., 2023). For each wind direction θ (12 directions with 30° spacing), the building and tree fields are rotated into wind-aligned coordinates so that the flow effectively comes from west. In this rotated grid, each connected building (or tree cluster) is characterized by an effective obstacle width W , depth D (along-wind), and height H . These dimensions are then used to define upwind and leeward zones.

Vertical uniformity assumption

270 In the present implementation, the upwind deceleration and leeward wake formulations are applied assuming that the wind reduction coefficient is vertically uniform from the ground up to the obstacle height, i.e. up to the building height or the canopy top in the case of trees. In other words, within the upwind and wake regions, the wind speed is assumed to be reduced by the same factor at all heights below the obstacle top, while the height dependence of the wind profile is only accounted for locally at the obstacle location (through the canopy-flow parameterization for trees).

275 This assumption is consistent with the conceptual framework of simplified urban wake models, where the dominant effect of buildings and vegetation on the mean flow is represented through vertically integrated drag and displacement effects rather than through a fully resolved vertical structure, especially in dense urban canopies. Several studies have shown that, for practical urban-scale applications, the bulk momentum deficit induced by buildings and canopies can be reasonably approximated as vertically homogeneous within the roughness sub-layer, up to the obstacle height (Santiago and Martilli, 2010; Cheng and Porté-Agel, 2015; Zonato et al., 2023). Given the objectives of the present study and the spatial resolution considered, this simplification offers a robust compromise between physical realism and computational efficiency.

Upwind deceleration

Upwind of each obstacle, the model prescribes a monotonic ramp from full wind to a reduced value over a characteristic length

$$L_f = \frac{2W}{1 + 0.8(W/H)}. \quad (6)$$

285 For buildings, the along-wind profile is parameterized as

$$C_f(x) = C_{\min} + (1 - C_{\min}) \left(\frac{x}{L_f} \right)^{p_f}, \quad 0 \leq x \leq L_f, \quad (7)$$



where x is the upwind distance from the wall (against the mean wind), $p_f = 1.5$, and $C_{\min} = 0.1$ is a lower bound enforced during post-processing. Very close to the building, C_f approaches C_{\min} , while farther upwind it gradually relaxes to 1.

Leeward wake

290 Downwind of buildings, the wake length L_r follows Röckle (1990):

$$L_r = \frac{4 \times 1.8 W}{\left(\frac{D}{H}\right)^{0.3} \left[1 + 0.24 \left(\frac{D}{H}\right)\right]}. \quad (8)$$

Along the wake centerline, the leeward recovery is assumed linear with distance:

$$C_b(x) = \frac{x}{L_r}, \quad 0 \leq x \leq L_r, \quad (9)$$

with $C_b(x)$ truncated to the range $[C_{\min}, 1]$ and $C_b \rightarrow 1$ as $x \rightarrow L_r$. Immediately downstream of the wall, C_b tends towards
295 C_{\min} , representing the recirculation and vortex zone. The implementation applies this recovery profile along each obstacle column in the wind-aligned grid, while ensuring that overlapping wakes combine conservatively.

Trees

Trees are treated with the same geometric description (W, D, H), but as porous obstacles. For each grid cell containing canopy, a local, height-dependent wind reduction coefficient $C_{t,\text{base}}$ is first computed at the evaluation height ($z = 10$ m in this study)
300 using a canopy-flow parameterization following Bernard et al. (2023). The model depends on an effective leaf area index

$$\text{LAI}_{\text{eff}} = \text{LAI} \lambda_p, \quad (10)$$

where LAI is the leaf area index (expressed as $\text{m}^2 \text{m}^{-2}$ and set to $\text{LAI} = 5$) and λ_p is the canopy plan-area fraction. In each grid cell, the canopy height H is taken from the tree DSM, and we define a displacement height $d = 0.7 H$ and a canopy roughness length $z_{0,\text{tree}} = 0.1 H$, following standard canopy-flow parameterizations. The base attenuation coefficient k (denoted α in the
305 implementation) is then modeled as

$$k = a_0 + a_1 \text{LAI}_{\text{eff}}. \quad (11)$$

With $a_0 = 0.5$ and $a_1 = 0.2$. The base canopy reduction factor at height z is then defined piecewise. Let H be the local canopy height (from the tree DSM), $d = 0.7 H$ the displacement height, and $z_{0,\text{tree}} = 0.1 H$ the canopy roughness length (with a small lower bound to avoid singularities). Denoting the reference roughness by $z_{0,\text{ERA5}}$ (taken from the meteorological dataset), we
310 first define the reference neutral wind-profile denominator

$$D_{\text{ref}}(z) = \ln\left(\frac{z}{z_{0,\text{ERA5}}}\right), \quad (12)$$

and we bound the arguments of the logarithms in the numerator so that they remain larger than 1.



For points above or just at the canopy top ($z \geq H$), the base reduction factor is purely logarithmic:

$$C_{t,\text{base}}(z \geq H) = \frac{\ln((z-d)/z_{0,\text{tree}})}{D_{\text{ref}}(z)}. \quad (13)$$

315 For points inside the canopy ($z < H$), the model assumes that the wind speed at the canopy top follows the same logarithmic scaling, but decays exponentially towards the ground according to the attenuation coefficient k :

$$C_{t,\text{base}}(z < H) = \frac{\ln((H-d)/z_{0,\text{tree}})}{D_{\text{ref}}(z)} \exp\left[-k\left(1 - \frac{z}{H}\right)\right], \quad (14)$$

so that $C_{t,\text{base}} \rightarrow \ln((H-d)/z_{0,\text{tree}})/D_{\text{ref}}(z)$ at the canopy top ($z \rightarrow H$) and decreases towards $\exp(-k)$ near the ground.

Given $C_{t,\text{base}}$ at the obstacle location, we apply the same wake geometry as for buildings, with an upwind deceleration over

$$320 \quad L_f = \frac{2W}{1 + 0.8(W/H)}$$

and a leeward wake of length

$$L_r = \frac{4 \times 1.8W}{\left(\frac{D}{H}\right)^{0.3} [1 + 0.24\left(\frac{D}{H}\right)]}.$$

Upwind of the canopy edge (x measured against the mean wind), the coefficient is ramped from $C_{t,\text{base}}$ to 1 using the same functional shape as for buildings,

$$325 \quad C_{t,\text{up}}(x) = C_{t,\text{base}} + (1 - C_{t,\text{base}}) \left(\frac{x}{L_f}\right)^{p_f}, \quad 0 \leq x \leq L_f, \quad (15)$$

with $p_f = 1.5$. Downwind, the profile recovers linearly from $C_{t,\text{base}}$ towards 1 over the wake length,

$$C_{t,\text{down}}(x) = C_{t,\text{base}} \frac{x}{L_r}, \quad 0 \leq x \leq L_r, \quad (16)$$

Combination and smoothing

For each direction θ , building and tree wakes are computed on the wind-aligned grid and then rotated back to the original grid.

330 At each grid point, we combine the two contributions multiplicatively,

$$C(\mathbf{x}, \theta) = C_b(\mathbf{x}, \theta) C_t(\mathbf{x}, \theta), \quad (17)$$

with building interiors masked out. The resulting 2-D field is Gaussian-smoothed (with an effective kernel of about 40 m) to blend overlapping wakes and suppress blocky artifacts.

335 Finally, $C(\mathbf{x}, \theta)$ is clipped to the interval $[C_{\min}, 1]$ and stored as `WindCoeff_dirXXX.tif` (one file per direction). In SOLWEIG, the near-surface wind speed is obtained by multiplying the domain-consistent 10 m wind ($U_{10}^{\text{city}}, V_{10}^{\text{city}}$) by these directional coefficients $C(\mathbf{x}, \theta)$. For the computation of UTCI, a minimum wind speed of 0.5 ms^{-1} is imposed to avoid unrealistically low values.



2.5 Case study: Dortmund

The city of Dortmund ($51^{\circ}30'N, 7^{\circ}28'E$), Germany, was selected as a case study area due to the large city-wide urban biometeorological weather station network. Dortmund is located in the larger metropolitan Ruhr region and was selected as a case study area in the ICLEI Action Fund 2.0 with the Data2Resilience project, which supported the implementation of data-driven services to enhance the resilience to extreme heat. Following the approach in Freiburg, Germany (Plein et al., 2024), a large weather station network was deployed. It consists of 51 stations measuring air-temperature and relative humidity using a Decentlab DL-SHT35 sensor, as well as 25 additional stations equipped with biometeorological sensors, specifically a black globe temperature sensor (Decentlab DL-BLG with 125 mm diameter) and an all-in-one weather station sensor (Decentlab DL-ATM41), providing additional parameters besides air temperature and relative humidity such as black globe temperature, downward shortwave radiation, wind speed and wind direction, and precipitation. The stations are located throughout the city and cover a wide variety of urban structures based on Local Climate Zones. Biometeorological stations are deliberately placed at locations of public interest, such as main squares, urban parks and districts with combined vulnerabilities, to enable the local government to assess the heat exposure and guide data-driven adaptation measures. With few exceptions, stations are mounted on lamp posts 3.3 m above ground and 50 cm away from the pole. The black globe sensor always faces south, and this way avoids shading by the pole. In some cases (seven in total) where this is not possible due to other obstructions, it is precisely documented in the metadata (Hüser et al., 2026), where also each station setup is carefully described. Measured raw data is transmitted by individual sensors every five minutes via LoRaWAN (Long Range Wide Area Network) and sensor data are then combined during postprocessing, allowing for computation of derived biometeorological parameters such as mean radiant temperature (T_{mrt}) following DIN EN ISO 7726:2021-03 (2021) and the common thermal comfort indices UTCI (Jendritzky et al., 2012), PET (Höppe, 1999) and heat index (HI) (Steadman, 1979; Rothfusz, 1990). The maximum allowed temporal offset between sensor measurements (black globe and all-in-one sensor) is 5 minutes, labeling the combined station-measurement with the all-in-one sensor's original timestamp since the expected response time of the black globe sensor is much slower.

For this study, observations and metadata from 25 stations (more than 250,000 valid station-hour records) were derived via the public API (Application Programming Interface) which provides near-realtime data (<https://api.data2resilience.de>). Station coordinates are provided in WGS84 by the API, and the Dortmund domain is defined in UTM 32N (EPSG:32632) at 2 m grid spacing, covering the urban core and adjacent districts (tens of km^2). Simulations span the period from 6 August 2024 to 31 December 2025. Two model configurations are evaluated: a baseline run driven by unmodified ERA5 forcing (S_{01_STD}) and a configuration that activates both the morphology-based wind correction and the diagnostic UHI cycle ($S_{01_WC_UHI}$).

3 Results

We report computational performance and output volume for the Dortmund case study, illustrate the spatial imprint of the wind and UHI diagnostics, and evaluate station-based skill using aggregate statistics, distributions, and temporal and seasonal thermal stress metrics.




Figure 2. Dortmund model domain: land use, buildings, trees, and station points annotated with their respective ID on the common UTM32N grid. Generated during the input-build phase.

3.1 Computational performance

370 To quantify the practical impact of GPU acceleration, we benchmarked a single-day GLIDE-SOL run over Dortmund on a
 371 reduced 3129×3113 grid (2 m spacing). Using one NVIDIA A100 (40 GB) GPU, the 24 h simulation (saving UTCI, T_{mrt} ,
 372 T_a , and V_{a10m}) completed in 254 s (10.6 s per simulated hour). The same configuration executed in CPU-only mode on the
 373 same node required 1285 s (53.5 s per simulated hour), corresponding to a wall-clock speed-up of about 5.1. For the GPU run,
 PyTorch CPU threading was restricted to `num_threads=1` to avoid host-side oversubscription; in CPU-only mode all logical
 374 CPU threads were available.

In addition to compute acceleration, deferred repacking to LERC+ZSTD reduces I/O and storage requirements. For the
 Dortmund case-study grid, an uncompressed daily UTCI stack would occupy about 3.3 GB, whereas the repacked GeoTIFF is
 251 MB ($\sim 13\times$ smaller; 92 % reduction), using the compression error values described in 2.3.

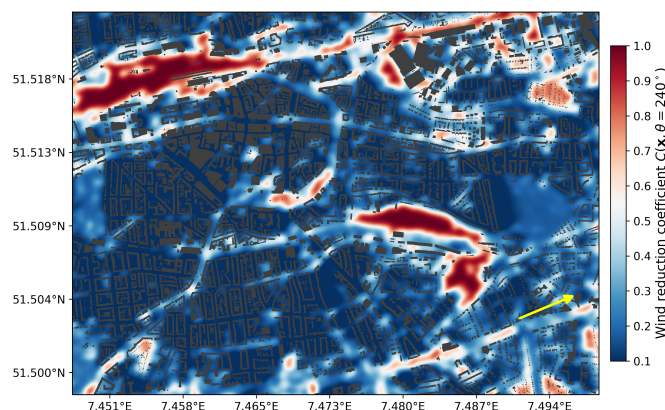


Figure 3. Example wind-reduction coefficient field $C(\mathbf{x}, \theta)$ for $\theta = 240^\circ$ (south-westerly flow) over Dortmund. Dark-blue regions indicate strong wind attenuation in dense urban fabric and in leeward wakes; red regions highlight ventilation corridors with limited reduction. The yellow arrow indicates the inflow direction used to compute upwind and downwind wakes.

3.2 Diagnostics of the wind and temperature corrections

380 The diagnostic wind module generates a directional attenuation coefficient field $C(\mathbf{x}, \theta)$ that converts the domain-consistent 10 m wind to a morphology-dependent ventilation proxy. For $\theta = 240^\circ$ (south-westerly flow; Fig. 3), C is lowest within dense blocks and in leeward wakes ($C \ll 1$), while major open corridors preserve higher effective wind speeds ($C \approx 1$). Although Reynolds-averaged and large-eddy CFD simulations provide higher-fidelity representations of shear layers and recirculation (Chu and Wang, 2025), the morphology-based $C(\mathbf{x}, \theta)$ provides a computationally inexpensive surrogate that retains the main anisotropy imposed by urban form. This is the physical mechanism by which the `Sol_WC_UHI` configuration compensates for the lack of explicit urban drag in coarse meteorological forcing.

The diagnostic UHI cycle acts on the background meteorological forcing rather than producing a standalone spatial field: it applies a night-time warming perturbation modulated by SVF and vegetation (Eq. 2). Its effect is therefore assessed through the station-based skill metrics and seasonal diagnostics presented below.

390 Figure 4 summarizes JJA mean fields from `Sol_WC_UHI` for central Dortmund. Station symbols report observed (left semicircle) and modeled (right semicircle) summer means at each location. Air temperature exhibits only modest spatial gradients and agrees well at station points. Wind speed shows strong morphology-driven heterogeneity, with low values in dense built-up areas and higher values in open corridors and peripheral zones; this pattern is consistent with the wake-based attenuation and with the improved wind-speed statistics (Sect. 3.3). T_{mrt} displays pronounced small-scale variability due to shading and land cover, and the remaining differences at station locations are consistent with the residual T_{mrt} errors in Table 1. UTCI integrates these drivers and shows close agreement of summer means at station points, indicating that the diagnostic corrections translate into realistic warm-season thermal-stress levels.



Sol_WC_UHI - JJA: model + station

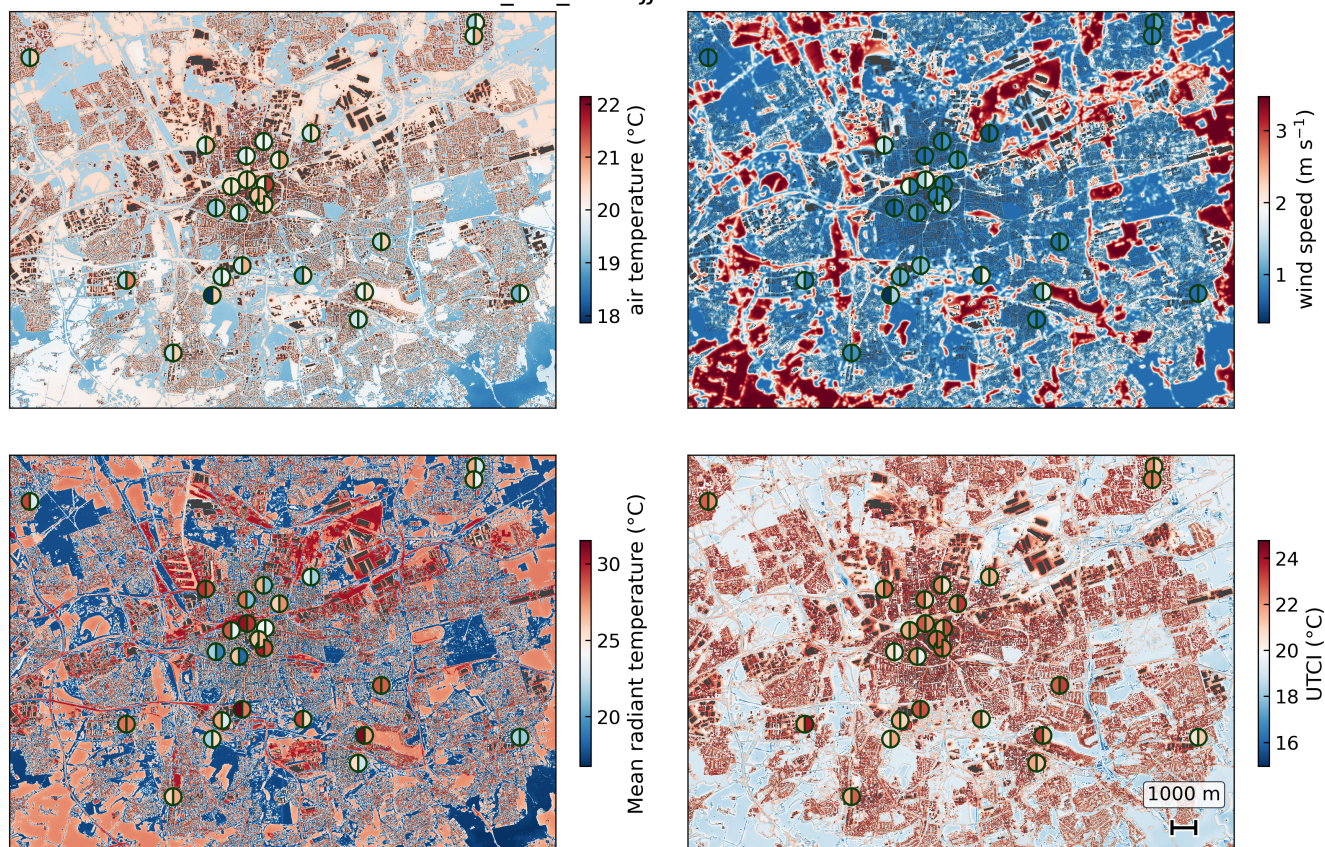


Figure 4. Summer (JJA) mean fields for the Sol_WC_UHI configuration. Panels show (top left) air temperature, (top right) wind speed, (bottom left) mean radiant temperature (T_{mrt}), and (bottom right) UTCI. Circles denote station locations; for each station, the left half of the symbol represents observed values, while the right half shows the corresponding modeled summer mean.

3.3 Aggregate skill across all stations

Model skill is evaluated against hourly measurements from 25 stations with heterogeneous but overlapping temporal coverage between August 2024 and December 2025 (more than 250,000 station-hour records). We compare two configurations: a baseline run driven by unmodified ERA5 meteorology (Sol_STD) and a run that activates the morphology-based wind correction and the diagnostic UHI cycle (Sol_WC_UHI). Table 1 summarizes the aggregate performance. The combined corrections reduce the UTCI cold bias from -8.85 to -0.82 °C and cut RMSE from 9.90 to 2.78 °C (72% reduction). Improvements are also evident for T_{mrt} (RMSE 7.34 °C \rightarrow 6.24 °C) and canopy air temperature (RMSE 1.78 °C \rightarrow 1.53 °C). The largest gain occurs for wind speed, where the strong positive bias in Sol_STD is nearly removed (bias 2.87 \rightarrow 0.02 m s $^{-1}$; RMSE 3.32 \rightarrow 0.74 m s $^{-1}$).



Table 1. Aggregate performance against observations. Bias, MAE, and RMSE are in °C for UTCI/ T_{mrt} / T_a and in m s^{-1} for wind speed.

Variable–configuration	Bias	MAE	RMSE
UTCI – Sol_STD	-8.85	8.94	9.90
UTCI – Sol_WC_UHI	-0.82	2.14	2.78
T_{mrt} – Sol_STD	-4.74	5.90	7.34
T_{mrt} – Sol_WC_UHI	-2.87	4.38	6.24
T_a – Sol_STD	-1.00	1.37	1.78
T_a – Sol_WC_UHI	0.16	1.10	1.53
V – Sol_STD	2.87	2.88	3.32
V – Sol_WC_UHI	0.02	0.53	0.74

Table 2 reports station-level UTCI errors (RMSE and mean bias, MB) for Sol_STD and Sol_WC_UHI, with sites sorted by sky-view factor (SVF) as a compact proxy for local enclosure and radiative exposure. In Sol_STD, the mean bias is uniformly negative across the network (between -6.8 and -10.0°C) and RMSE remains high (7.8 – 11.2°C) across the full SVF range, indicating that uncorrected forcing dominates the error (in particular, excessive wind cooling). With Sol_WC_UHI, RMSE decreases at every station (2.1 – 4.3°C) and biases collapse towards zero (range -3.1 to 1.9°C).

3.4 Scatter and distributional diagnostics

The scatter plots (Figs. 5 and 6) provide a compact view of systematic errors and conditional biases. For UTCI (Fig. 5), Sol_STD exhibits a pronounced cold bias, consistent with excessive convective cooling from overestimated winds and with missing nocturnal urban warming in the forcing. Activating wind attenuation and the UHI adjustment tightens the scatter and shifts the cloud towards the 1:1 line. In Fig. 6, the strongest improvement is in wind speed (center), where Sol_WC_UHI largely removes the positive bias. For canopy air temperature (right), the UHI cycle reduces the cold bias and yields lower MAE/RMSE, while T_{mrt} (left) improves more modestly.

The probability density functions (Fig. 7) provide a distribution-level evaluation of the main UTCI drivers and confirm that the corrections primarily act through wind and near-surface air temperature forcing.

For air temperature (top-left), both configurations reproduce the broad observed distribution, while Sol_WC_UHI shifts the PDF closer to the observations, consistent with the reduced T_a bias. The canopy wind speed PDF (top-right) reveals the major deficiency of Sol_STD: a strong shift towards higher wind speeds and an unrealistically long high-wind tail (up to ~ 10 – 12 m s^{-1}). Sol_WC_UHI largely removes this artifact, recovering the observed low-wind peak and damping the tail.

Consistently, the UTCI PDF (bottom-left) in Sol_STD is shifted towards lower values and exhibits an amplified cold tail, while Sol_WC_UHI moves the distribution back towards the observed mode and reduces the spurious cold tail. In contrast, T_{mrt}



Table 2. UTCI errors by station (sorted by SVF obs). Bold indicates lowest error for each station and metric.

Station	SVF obs	SVF mod	RMSE (°C)		MB (°C)	
			Sol_STD	Sol_WC_UHI	Sol_STD	Sol_WC_UHI
BHS	0.20	0.21	9.73	2.85	-9.08	-1.32
GAS	0.20	0.21	9.23	2.17	-8.11	-0.71
WEP	0.22	0.21	9.77	2.09	-8.81	-1.00
OHS	0.23	0.64	10.48	2.18	-9.42	-0.54
BRS	0.23	0.34	11.20	2.13	-10.05	-0.32
CLA	0.32	0.32	9.68	2.52	-8.62	-0.98
BOL	0.36	0.85	9.90	3.74	-7.91	-0.77
SBA	0.38	0.38	9.53	2.59	-8.50	-1.42
HAP	0.38	0.86	10.72	3.44	-9.28	-1.07
DKP	0.43	0.43	10.65	2.83	-9.88	-1.78
HKS	0.45	0.75	9.32	2.55	-8.21	0.31
DOU	0.46	0.43	8.71	3.10	-7.66	0.26
NOM	0.46	0.48	10.81	3.37	-9.91	-2.60
GWS	0.47	0.47	10.59	2.98	-9.73	-0.47
HAN	0.51	0.79	10.36	2.69	-8.99	-0.09
MSP	0.58	0.58	7.76	3.23	-6.78	1.93
PSW	0.62	0.61	10.69	3.30	-9.57	-1.99
WFP	0.64	0.90	9.76	4.31	-8.89	-3.21
DHS	0.64	0.87	10.15	2.51	-9.01	-0.90
KPW	0.65	0.65	10.27	3.54	-9.40	-0.49
WFH	0.65	0.65	9.17	2.58	-8.34	-0.28
FRP	0.66	0.92	9.00	3.07	-8.04	-1.15
HBF	0.71	0.93	9.93	3.88	-8.89	-2.54
SPS	0.71	0.78	8.93	3.47	-7.70	-1.55
PHX	0.71	0.95	10.26	4.17	-9.27	-3.12

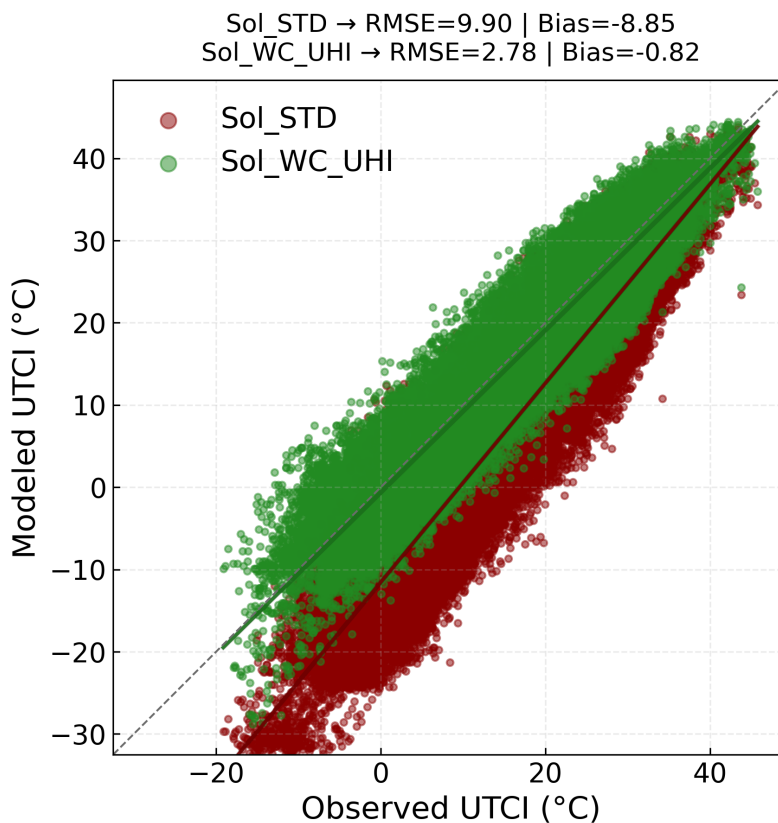


Figure 5. Observed vs. modeled UTCI for all stations and hours. The baseline `Sol_STD` (red) shows a systematic cold bias, whereas `Sol_WC_UHI` (green) concentrates closer to the 1:1 line, consistent with improved wind and temperature forcing.

PDFs (bottom-right) show only limited sensitivity to the applied corrections, indicating that remaining errors are dominated by radiative forcing, surface properties, and input uncertainties rather than by aerodynamic exposure.

3.5 Temporal behavior and thermal stress categories

430 To assess whether the improvements are persistent rather than limited to mean statistics, we examine the temporal evolution of spatially averaged conditions across all stations (Figs. 8 and 9) and the resulting UTCI category frequencies (Fig. 10). Because the series are aggregated over the entire station network, they primarily reflect synoptic-to-seasonal variability, so remaining offsets or amplitude errors indicate systematic bias.

The wind-speed time series (Fig. 8) highlights the dominant deficiency of the baseline configuration (`Sol_STD`, ERA5-
435 driven winds without local urban diagnostics): winds are systematically too strong for the urban canopy and exhibit frequent high-wind excursions that are not present in the observations. In `Sol_STD`, wind speeds repeatedly spike above $\sim 8\text{--}12\text{ m s}^{-1}$, whereas observed winds at the station network remain mostly below $\sim 2\text{ m s}^{-1}$. The combined roughness rescaling and wake-

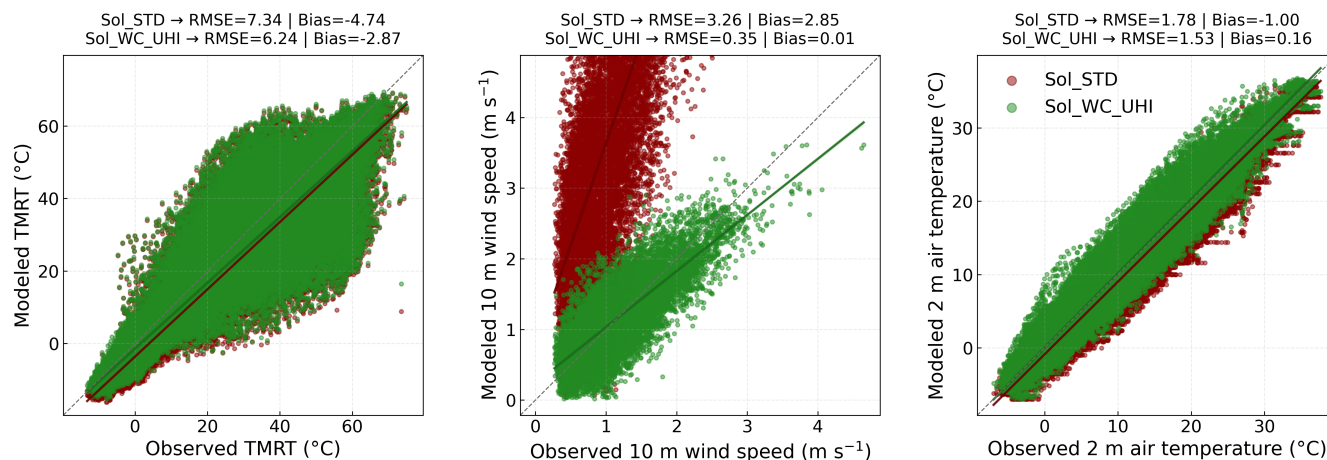


Figure 6. Observed vs. modeled (left) T_{mrt} , (center) wind speed, and (right) air temperature for all stations and hours. The baseline run (Sol_STD ; labeled Sol_STD in the panel titles) shows large wind-speed overestimation and a cold bias in T_a , while the diagnostic wind and UHI corrections in Sol_WC_UHI reduce bias and RMSE, particularly for wind speed.

based attenuation in Sol_WC_UHI suppress these excursions and constrain simulated winds to the observed envelope while preserving synoptic variability. This correction is critical because UTCI is highly sensitive to near-surface wind under cold and moderate conditions.

Consistent with the wind behavior, the UTCI time series (Fig. 9) shows a persistent cold bias in Sol_STD throughout the evaluation period, including exaggerated negative excursions during cold-season outbreaks and damped summer peaks. The corrected configuration Sol_WC_UHI substantially reduces this offset and reproduces the seasonal progression and much of the event-to-event variability. Periods of largest disagreement in Sol_STD coincide with intervals of excessive winds in Fig. 8, indicating that unrealistic aerodynamic exposure to ERA5 winds is a key driver of the baseline UTCI bias. The nocturnal UHI adjustment provides an additional contribution under stable night-time conditions.

The improved continuous UTCI signal translates into a more realistic partitioning into thermal stress categories (Fig. 10). In the baseline case, the systematic cold shift redistributes hours from “no thermal stress” into cold-stress classes: Sol_STD overestimates moderate and strong cold stress by +25.5% and +6.2%, and underestimates “no thermal stress” by -26.8%. With Sol_WC_UHI , category biases are within a few percentage points across the spectrum (e.g., slight cold stress +2.7%, moderate cold stress +1.7%, and no thermal stress -4.3%), and warm-season heat stress frequencies align closely with the observations. Beyond station-based aggregation, the impact of the diagnostics is reflected in the spatial mapping of category frequencies across the Dortmund domain (Figs. 11 and 12), which determines where cold and heat stress exposure accumulates.

In winter (DJF), the fraction of hours in the “slight cold stress” category (0–9°C UTCI) provides a sensitive diagnostic of the model’s cold shift: if UTCI is systematically underestimated, hours are redistributed from this intermediate class into moderate and strong cold-stress categories. Figure 11 shows that Sol_STD assigns uniformly low fractions across the domain,




Figure 7. Probability density functions for canopy air temperature, wind speed, UTCI, and T_{mrt} . Model curves are compared with observations (black). The combined wind and UHI corrections in `Sol_WC_UHI` suppress the spurious high-wind tail and remove most of the UTCI cold shift present in `Sol_STD`.

masking the expected center-to-periphery contrast. With `Sol_WC_UHI`, a clear urban-core signal emerges, with higher slight-cold frequencies over dense city fabric and lower values over parks, open areas, and the peri-urban fringe where ventilation is stronger. Station symbols reinforce this pattern, with inner-city sites moving toward observed frequencies while peripheral sites retain lower fractions. In summer (JJA), “moderate heat stress” (26–32 °C UTCI) shows a strong spatial contrast tied to urban form and ventilation (Fig. 12). In `Sol_STD`, moderate heat-stress frequencies are muted and relatively flat, consistent with the tendency to damp warm-season stress through excessive wind speeds even in dense neighborhoods. With `Sol_WC_UHI`, a center-to-outskirts gradient becomes visible: higher fractions emerge over compact, wind-sheltered urban fabric and heavily built corridors, while parks, green areas, and the outer suburban fringe retain low values. Station symbols indicate that the corrected configuration better captures both the elevated inner-city frequencies and the lower peripheral values without overstating the overall occurrence.




Figure 8. Time series of spatially averaged wind speed over all stations. The baseline `Sol_STD` shows frequent unrealistically high winds, whereas `Sol_WC_UHI` largely aligns with observed magnitudes after roughness rescaling and wake-based attenuation.

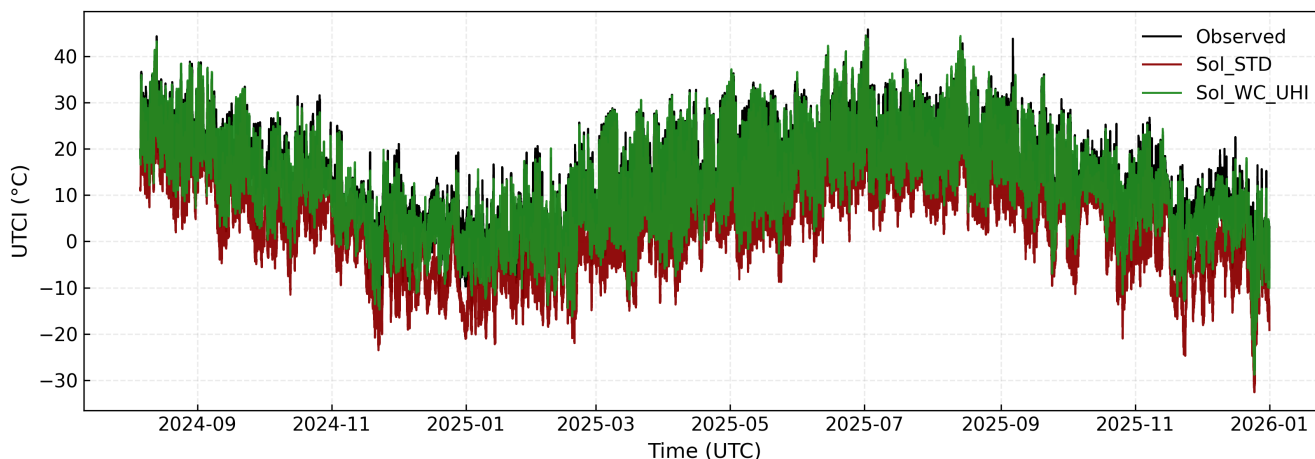


Figure 9. Time series of spatially averaged UTCI over all stations. `Sol_STD` remains systematically colder than observations, while `Sol_WC_UHI` reduces the offset and better reproduces the seasonal cycle and event-to-event variability.

4 Discussion

The Dortmund evaluation demonstrates that two lightweight diagnostics—directional wind attenuation and a nocturnal UHI adjustment—substantially reduce the cold bias that arises when SOLWEIG is forced directly with coarse reanalysis fields. In the baseline configuration (`Sol_STD`), near-surface winds are too strong for the urban canopy and night-time air temperature is underestimated, producing a persistent cold offset and an exaggerated cold tail in UTCI (Figs. 9 and 7). The temporal and category analyses indicate that the corrections improve not only mean skill but also event-scale variability and the partitioning

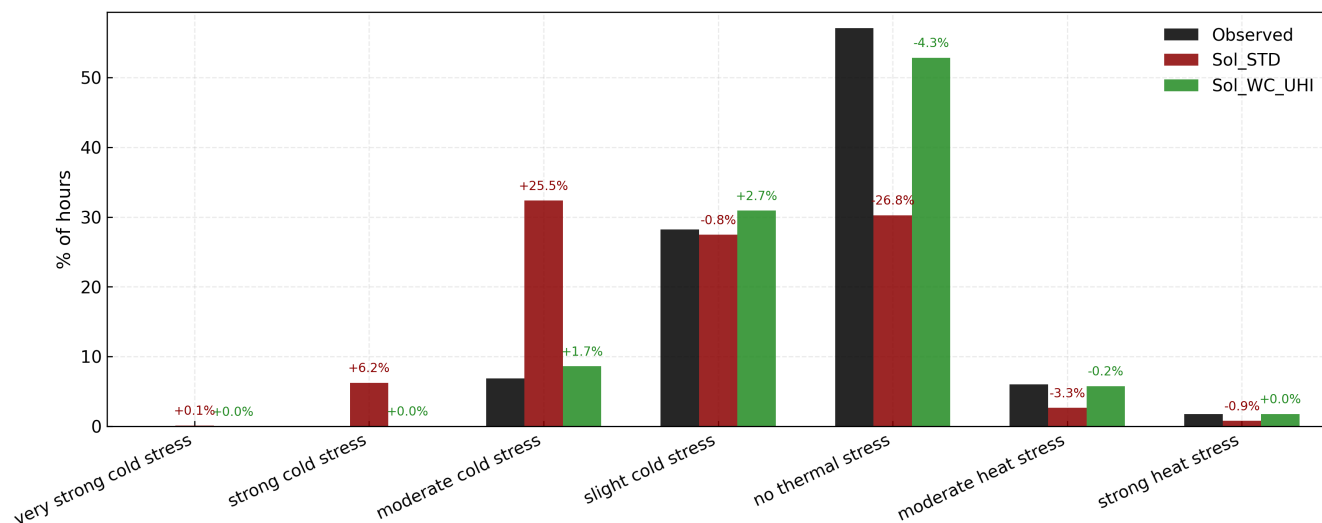


Figure 10. Hours per UTCI category across all stations. Bars show observed and simulated fractions; numbers above model bars indicate model–observation differences in percentage points. *Sol_WC_UHI* markedly improves the category distribution relative to *Sol_STD*, particularly by reducing spurious cold-stress frequencies and restoring the observed dominance of “no thermal stress”.

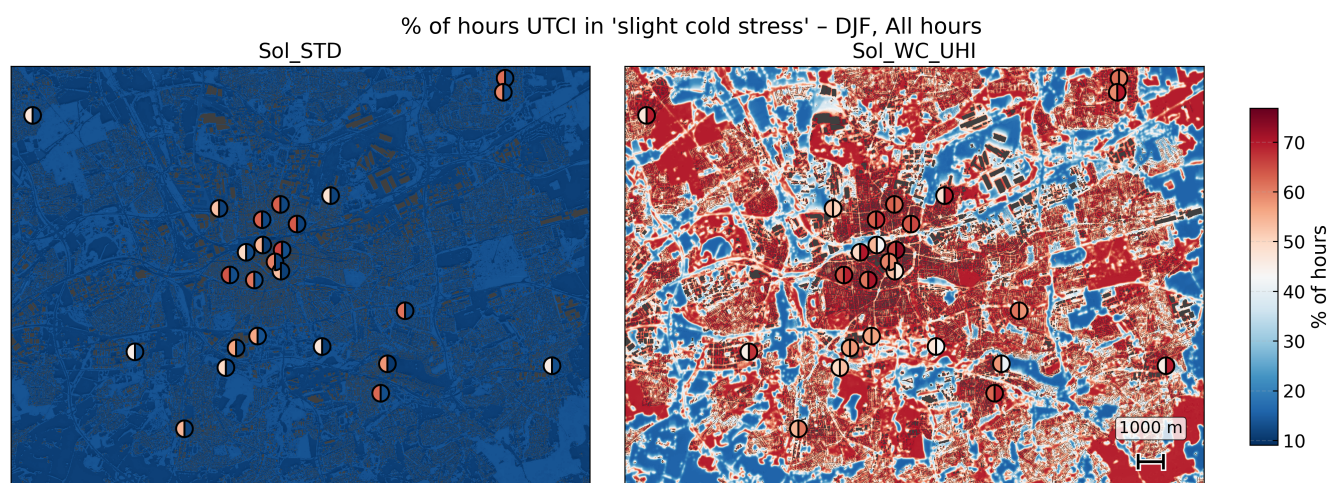


Figure 11. Fraction of hours classified as “slight cold stress” (UTCI) in DJF, considering all hours. Left: *Sol_STD*. Right: *Sol_WC_UHI*. Circles denote station locations; for each station, the left half of the symbol represents observed frequencies, while the right half shows the corresponding model value.

of hours into cold- and heat-stress classes, which is essential for exposure assessments. Because the workflow is driven
 475 by globally available inputs and standardized preprocessing, these gains are achieved without bespoke local data assembly,
 supporting consistent multi-city applications.

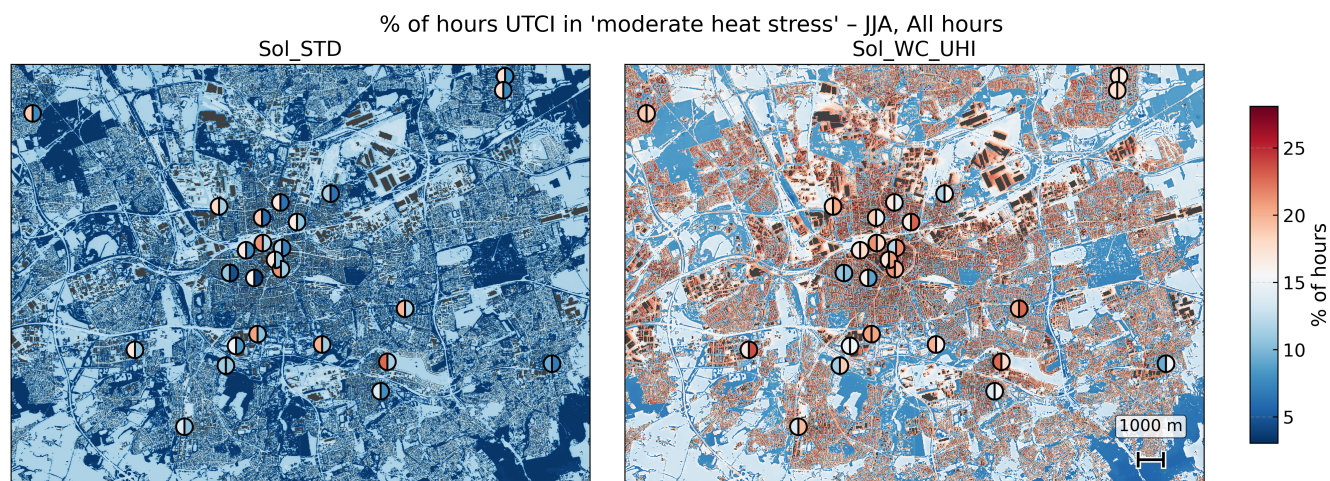


Figure 12. Fraction of hours classified as “moderate heat stress” (UTCI) in JJA, considering all hours. Left: `Sol_STD`. Right: `Sol_WC_UHI`. Circles denote station locations; for each station, the left half of the symbol represents observed frequencies, while the right half shows the corresponding model value.

The wind diagnostic is the dominant contributor to the UTCI improvement. By imprinting building and tree morphology onto the canopy wind through the coefficient fields $C(\mathbf{x}, \theta)$ (Fig. 3), `Sol_WC_UHI` removes the spurious high-wind tail (Fig. 7) and nearly eliminates the wind-speed bias (Table 1), consistent with Röckle-style conceptualizations of urban flow (Röckle, 1990).
 480 The resulting spatial gradients align with dense urban cores and more ventilated outskirts, reinforcing the physical consistency of the approach. The remaining station-to-station spread suggests that unresolved local exposure and representativeness (e.g., site-specific shielding) still matter, but to a much lesser extent than the forcing bias.

The diagnostic UHI cycle provides a complementary correction that is most relevant under stable nocturnal conditions, particularly in winter. It shifts the distributions of air temperature and UTCI towards the observations and reduces extreme
 485 cold-stress episodes (Figs. 7 and 9). Its simple structure keeps computational cost low and enables consistent application across regions, but the residual positive T_a bias in `Sol_WC_UHI` (Table 1) indicates that a single diagnostic parameter set cannot capture the full range of synoptic regimes and advective effects; **incorporating flow-dependent predictors is a promising avenue for future development.**

T_{mrt} improvements are comparatively modest (Table 1; Figs. 6 and 7), consistent with T_{mrt} being driven primarily by radiative
 490 forcing and urban geometry. Air temperature (T_a) influences T_{mrt} mostly indirectly: it modulates atmospheric longwave emission (through sky emissivity) and, in SOLWEIG, it is used to parameterize surrounding surface temperatures, including wall temperature (T_{wall}), which in turn affects the longwave radiation received by the pedestrian. Remaining discrepancies likely reflect uncertainties in morphology and surface properties and missing processes such as anthropogenic heat and its effect on surface temperatures and long-wave emission, especially during the heating season (Liu et al., 2023; Wallenberg



495 et al., 2025). In a global context, improving canopy and surface characterization layers is therefore a key path to further gains in radiative skill.

From a modeling and reproducibility standpoint, GLIDE-SOL addresses scalability constraints that often limit city-scale, meter-scale applications. GPU acceleration reduces wall-clock time by about a factor of five for the Dortmund benchmark. Deferred LERC+ZSTD repacking reduces output volumes by about 92 % while enforcing explicit error budgets per variable, 500 producing lightweight GeoTIFF archives that reduce I/O and storage pressure. Together with the global input pipeline, these features enable multi-year, city-scale UTCI mapping without prohibitive compute or storage requirements and make coordinated multi-city experiments operationally feasible.

Limitations of the present demonstration include uncertainty in open building and tree-height products, a simplified diagnostic representation of urban heat storage, and reliance on coarse ERA5 forcing that omits mesoscale circulations and city-wide 505 advection (Kittner et al., 2025). Additional limitations are the absence of an urban modification to relative humidity, the lack of anthropogenic heat contributions from building heating/cooling and traffic, and **the use of a simplified urban wind-reduction model that cannot fully capture local channeling or sheltering**. Future work should therefore evaluate the diagnostics across multiple cities and climates, quantify sensitivity to morphology and roughness assumptions, and explore coupling the UHI signal to mesoscale predictors or urban canopy schemes, with attention to regional variability in global input quality.

510 5 Conclusions

GLIDE-SOL extends SOLWEIG into a scripted, globally deployable workflow that (i) builds the full set of SOLWEIG inputs from global data sources with standardized preprocessing, (ii) adds computationally cheap diagnostics to mitigate known biases in coarse meteorological forcing over cities, and (iii) enables scalable simulations through GPU acceleration and bounded-error output compression that yields lightweight GeoTIFF archives for long-term analyses. The automated data-engine and 515 configuration-driven pipeline are key novelties that make city-scale UTCI mapping reproducible and transferable without bespoke local datasets.

For Dortmund (25 stations; more than 250,000 station-hour records), the enhanced configuration `SOL_WC_UHI` reduces UTCI RMSE from 9.90 to 2.78 °C and the mean bias from -8.03 to -0.55 °C. The improvement is driven primarily by correcting the wind speed (bias $2.87 \rightarrow 0.02$ m s⁻¹), with additional gains from the nocturnal UHI adjustment. Beyond mean 520 skill, the corrections restore the temporal evolution of UTCI, reduce spurious cold-stress hours, and recover realistic spatial contrasts in winter cold stress and summer heat-stress hotspots, indicating that globally driven diagnostics can reproduce both temporal and spatial exposure signatures.

On the technical side, a 24 h benchmark on a 3129×3113 grid completes in 254 s on a single NVIDIA A100 GPU (about $5 \times$ faster than CPU-only), and daily outputs can be reduced by over 90 % using LERC+ZSTD with explicit error budgets. The 525 combination of fast GPU execution and lightweight GeoTIFF storage supports routine, meter-resolution UTCI mapping at city scale over multi-year periods, enabling consistent multi-city benchmarking, scenario testing, and evaluation of heat-mitigation strategies at regional to global scales.



Code and data availability

The GLIDE-SOL source code is archived on Zenodo (Zonato et al., 2026) and is available at <https://doi.org/10.5281/zenodo.18671813>. The observational dataset used for the Dortmund evaluation is available on Zenodo at <https://zenodo.org/records/18221203>. The GPU implementation draws on the open-source SOLWEIG-GPU project, available at https://github.com/nvnsudharsan/SOLWEIG-GPU/tree/main/solweig_gpu.

Author contributions

A.Z. developed and integrated the GLIDE-SOL workflow, implemented the software, conducted the computational experiments, and wrote the original draft. M.D. contributed to the conceptualization and methodology and co-developed the software and experiments. H.G.K. and N.S. contributed to the GPU and software implementation and supported the computational experiments. L.M. contributed to the software and workflow implementation and supported the computational experiments. J.K. and L.W. contributed to the methodology and paper organization, curated and processed the observational datasets, contributed to the analysis, and contributed substantially to the writing of the text and figures. B.B. contributed to the curation and processing of the observational datasets. A.M. and M.M. contributed to the interpretation of the results and the revision of the manuscript. All authors reviewed and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Acknowledgements

Contributions from A.Z. and L.M. were supported by the RETURN project (Multi-risk Science for Resilient Communities under a Changing Climate), part of the Extended Partnership program funded by the European Union through NextGenerationEU under Italy's National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.3 (Decree No. 341 of March 15, 2022, Project Code PE00000005, CUP D53C22002510002). Contributions from M.D. were supported by the European Union's HORIZON Research and Innovation Actions under grant agreement No. 101137851 - CARMINE, and the European Union's HORIZON Innovation Actions under grant agreement No. 101147385 - Minority Report. Contributions from A.M. were supported by the U.S. National Science Foundation, grant number CMMI-1942805 (CAREER: Human Thermal Exposure in Cities - Novel Sensing and Modeling to Build Heat-Resilience). The design of the Dortmund station network was funded by ICLEI Europe through the ICLEI Action Fund 2.0, supported by Google.org, under the project "Data2Resilience" (D2R); the implementation was funded through Smart City Dortmund.



555 References

- Adinolfi, M., Raffa, M., Reder, A., and Mercogliano, P.: Investigation on potential and limitations of ERA5 Reanalysis downscaled on Italy by a convection-permitting model, *Climate Dynamics*, 61, 4319–4342, <https://doi.org/10.1007/s00382-023-06803-w>, 2023.
- Aydin, Y., Janke, J., and Middleton, R.: A comparative review of microclimate and thermal comfort models: RayMan, ENVI-met, SOLWEIG, and STEVE, *Sustainable Cities and Society*, 46, 101–126, <https://doi.org/10.1016/j.scs.2018.12.019>, 2019.
- 560 Bernard, J., Bocher, E., Le Saux Wiederhold, E., Leconte, F., and Masson, V.: Estimation of missing building height in OpenStreetMap data: a French case study using GeoClimate 0.0.1, *Geoscientific Model Development*, 15, 7505–7532, <https://doi.org/10.5194/gmd-15-7505-2022>, 2022.
- Bernard, J., Lindberg, F., and Oswald, S.: URock 2023a: an open-source GIS-based wind model for complex urban settings, *Geoscientific Model Development*, 16, 5703–5727, <https://doi.org/10.5194/gmd-16-5703-2023>, 2023.
- 565 Briegel, F., Wehrle, J., Schindler, D., and Christen, A.: High-resolution multi-scaling of outdoor human thermal comfort and its intra-urban variability based on machine learning, *Geoscientific Model Development*, 17, 1667–1688, <https://doi.org/10.5194/gmd-17-1667-2024>, 2024.
- Buo, I., Sagris, V., Jaagus, J., and Middel, A.: High-resolution thermal exposure and shade maps for cool corridor planning, *Sustainable Cities and Society*, 93, 104 499, <https://doi.org/10.1016/j.scs.2023.104499>, 2023.
- 570 Cheng, W.-C. and Porté-Agel, F.: Adjustment of Turbulent Boundary-Layer Flow to Idealized Urban Surfaces: A Large-Eddy Simulation Study, *Boundary-Layer Meteorology*, 155, 249–270, <https://doi.org/10.1007/s10546-015-0004-1>, 2015.
- Cheng, Y., Niu, J., and Gao, N.: Thermal comfort models: A review and numerical investigation, *Building and Environment*, 47, 13–22, <https://doi.org/https://doi.org/10.1016/j.buildenv.2011.05.011>, international Workshop on Ventilation, Comfort, and Health in Transport Vehicles, 2012.
- 575 Chu, R. and Wang, K.: CFD in Urban Wind Resource Assessments: A Review, *Energies*, 18, <https://doi.org/10.3390/en18102626>, 2025.
- Cionco, R. M.: A wind-profile index for canopy flow, *Boundary-Layer Meteorology*, 3, 255–263, <https://doi.org/10.1007/BF02033923>, 1972.
- DIN EN ISO 7726:2021-03: Ergonomics of the thermal environment - Instruments for measuring physical quantities (ISO 7726:1998); German version EN ISO 7726:2001, 2021.
- Emmanuel, R. and Fernando, H. J. S.: Urban heat islands in humid and arid climates: role of urban form and thermal properties in Colombo, Sri Lanka and Phoenix, USA, *Climate Research*, 34, 241–251, <https://doi.org/10.3354/cr00694>, 2007.
- 580 Fiala, D., Havenith, G., Bröde, P., Kampmann, B., and Jendritzky, G.: UTCI-Fiala multi-node model of human heat transfer and temperature regulation, *International Journal of Biometeorology*, 56, 429–441, <https://doi.org/10.1007/s00484-011-0424-7>, 2012.
- Gorelick, N., Hancher, M., Dixon, M., and et al.: Google Earth Engine: Planetary-scale geospatial analysis, *Remote Sensing of Environment*, 202, 18–27, <https://doi.org/10.1016/j.rse.2017.06.031>, 2017.
- 585 Hersbach, H., Bell, B., Berrisford, P., and et al.: The ERA5 global reanalysis, *Quarterly Journal of the Royal Meteorological Society*, 146, 1999–2049, <https://doi.org/10.1002/qj.3803>, 2020.
- Höppe, P.: The physiological equivalent temperature – a universal index for the biometeorological assessment of the thermal environment, *International Journal of Biometeorology*, 43, 71–75, <https://doi.org/10.1007/s004840050118>, 1999.
- Hüser, C., Wolf, L., Gottschalk, N., Kittner, J., Kraas, B., Mittelstädt, C., Reinhart, V., Sismanidis, P., Wawrzyniak, N., and Bechtel, B.: Data2Resilience - A Biometeorological Weather Station Network in Dortmund: Station Documentation, <https://doi.org/10.5281/ZENODO.18221202>, 2026.



- Jendritzky, G., de Dear, R., and Havenith, G.: UTCI — Why another thermal index?, *International Journal of Biometeorology*, 56, 421–428, <https://doi.org/10.1007/s00484-011-0513-7>, 2012.
- Jänicke, B., Meier, F., Lindberg, F., Schubert, S., and Scherer, D.: Towards city-wide, building-resolving analysis of mean radiant temperature, *Urban Climate*, 15, 83–98, <https://doi.org/https://doi.org/10.1016/j.uclim.2015.11.003>, 2016.
- 595 Kamath, H. G., Singh, M., Malviya, N., Martilli, A., He, L., Aliaga, D., He, C., Chen, F., Magruder, L. A., Yang, Z.-L., and Niyogi, D.: GLOBal Building heights for Urban Studies (UT-GLOBUS) for city- and street- scale urban simulations: Development and first applications, *Scientific Data*, 11, 886, <https://doi.org/10.1038/s41597-024-03719-w>, 2024.
- Kittner, J., Fenner, D., Demuzere, M., and Bechtel, B.: Analysis of nocturnal urban heat advection using crowd weather stations, *Quarterly Journal of the Royal Meteorological Society*, 151, e5065, <https://doi.org/https://doi.org/10.1002/qj.5065>, 2025.
- 600 Konarska, J., Lindberg, F., Larsson, A., Thorsson, S., and Holmer, B.: Transmissivity of solar radiation through crowns of single urban trees— application for outdoor thermal comfort modelling, *Theoretical and Applied Climatology*, 117, 363–376, <https://doi.org/10.1007/s00704-013-1000-3>, 2014.
- Kusaka, H., Ikeda, R., Sato, T., Iizuka, S., and Boku, T.: Development of a Multi-Scale Meteorological Large-Eddy Simulation Model for Urban Thermal Environmental Studies: The “City-LES” Model Version 2.0, *Journal of Advances in Modeling Earth Systems*, 16, e2024MS004367, <https://doi.org/https://doi.org/10.1029/2024MS004367>, e2024MS004367 2024MS004367, 2024.
- 605 Lang, N., Jetz, W., Schindler, K., and Wegner, J. D.: A high-resolution canopy height model of the Earth, *Nature Ecology & Evolution*, 7, 1778–1789, <https://doi.org/10.1038/s41559-023-02206-6>, 2023.
- Li, X. and Wang, G.: GPU parallel computing for mapping urban outdoor heat exposure, *Theoretical and Applied Climatology*, <https://doi.org/10.1007/s00704-021-03692-z>, 2021.
- 610 Li, X., Wang, G., Zaitchik, B., Hsu, A., and Chakraborty, T. C.: Sensitivity and vulnerability to summer heat extremes in major cities of the United States, *Environmental Research Letters*, 19, 094039, <https://doi.org/10.1088/1748-9326/ad6c64>, 2024.
- Lindberg, F. and Grimmond, C. S. B.: 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, 311–323, <https://doi.org/10.1007/s00704-010-0382-8>, 2011a.
- 615 Lindberg, F. and Grimmond, C. S. B.: Nature of vegetation and building morphology characteristics across a city: Influence on shadow patterns and mean radiant temperatures in London, *Urban Ecosystems*, 14, 617–634, <https://doi.org/10.1007/s11252-011-0184-5>, 2011b.
- Lindberg, F., Holmer, B., and Thorsson, S.: SOLWEIG 1.0 – Modelling spatial variations of 3D radiant fluxes and mean radiant temperature in complex urban settings, *International Journal of Biometeorology*, 52, 697–713, <https://doi.org/10.1007/s00484-008-0162-7>, 2008.
- 620 Lindberg, F., Onomura, S., and Grimmond, C. S. B.: Influence of ground surface characteristics on the mean radiant temperature in urban areas, *International Journal of Biometeorology*, 60, 1439–1452, <https://doi.org/10.1007/s00484-016-1135-x>, 2016.
- Lindberg, F., Grimmond, C. S. B., Gabey, A., Huang, B., Kent, C. W., Sun, T., Theeuwes, N. E., Järvi, L., Ward, H. C., Capel-Timms, I., Chang, Y., Jonsson, P., Krave, N., Liu, D., Meyer, D., Olofson, K. F. G., Tan, J., Wästberg, D., Xue, L., and Zhang, Z.: Urban Multi-scale Environmental Predictor (UMEP): An integrated tool for city-based climate services, *Environmental Modelling & Software*, 99, 70–87, <https://doi.org/10.1016/j.envsoft.2017.09.020>, 2018.
- 625 Lindberg, F., Wallenberg, N., Thorsson, S., Haeger-Eugensson, M., Lönn, J., Holmberg, B., Frid, M., and Fahlström, J.: Micro-scale, city-wide analysis of outdoor thermal comfort during heatwaves in high latitude cities: influence of building geometry and vegetation, *International Journal of Biometeorology*, <https://doi.org/10.1007/s00484-025-03030-2>, 2025.



- Liu, Y., Luo, Z., and Grimmond, S.: Impact of building envelope design parameters on diurnal building anthropogenic heat emission, *Building and Environment*, 234, 110–134, <https://doi.org/10.1016/j.buildenv.2023.110134>, 2023.
- Maronga, B., Banzhaf, S., Burmeister, C., and et al.: Overview of the PALM model system 6.0, *Geoscientific Model Development*, 13, 1335–1372, <https://doi.org/10.5194/gmd-13-1335-2020>, 2020.
- Mayer, H. and Höppe, P.: Thermal comfort of man in different urban environments, *Theoretical and Applied Climatology*, 38, 43–49, <https://doi.org/10.1007/BF00866252>, 1987.
- 635 Mayer, H., Holst, J., Dostal, P., Imbery, F., and Schindler, D.: Human thermal comfort in summer within an urban street canyon in Central Europe, *Meteorologische Zeitschrift*, 17, 241–250, <https://doi.org/10.1127/0941-2948/2008/0285>, 2008.
- Middel, A. and Krayenhoff, E. S.: Micrometeorological determinants of pedestrian thermal exposure during record-breaking heat in Tempe, Arizona: Introducing the MaRTy observational platform, *Science of the Total Environment*, 659, 129–143, <https://doi.org/10.1016/j.scitotenv.2018.12.166>, 2019.
- 640 Oke, T. R., Mills, G., Christen, A., and Voogt, J. A.: *Urban Climates*, Cambridge University Press, <https://doi.org/10.1017/9781139016476>, 2017.
- Oostwegel, L. J. N., Schorlemmer, D., and Guéguen, P.: From Footprints to Functions: A Comprehensive Global and Semantic Building Footprint Dataset, *Scientific Data*, 12, 1699, <https://doi.org/10.1038/s41597-025-06132-z>, 2025.
- OpenStreetMap contributors: Planet OSM, <https://planet.openstreetmap.org>, 2017.
- 645 Plein, M., Kersten, F., Zeeman, M., and Christen, A.: Street-Level Weather Station Network in Freiburg, Germany: Station Documentation, Tech. rep., Zenodo, <https://doi.org/10.5281/ZENODO.12732551>, 2024.
- Potapov, P., Li, X., Hernandez-Serna, A., Tyukavina, A., Hansen, M. C., Kommareddy, A., Pickens, A., Turubanova, S., Tang, H., Silva, C. E., Armston, J., Dubayah, R., Blair, J. B., and Hofton, M.: Mapping global forest canopy height through integration of GEDI and Landsat data, *Remote Sensing of Environment*, 253, 112–165, <https://doi.org/10.1016/j.rse.2020.112165>, 2021.
- 650 Röckle, R.: *Bestimmung der Strömungsverhältnisse im Bereich komplexer Bebauungsstrukturen*, Tech. rep., Fraunhofer-Institut für Bauphysik, 1990.
- Rothfusz, L. P.: *The heat index equation (or, more than you ever wanted to know about heat index)*, Fort Worth, Texas: National Oceanic and Atmospheric Administration, National Weather Service, Office of Meteorology, 9023, 640, 1990.
- Santiago, J. L. and Martilli, A.: A Dynamic Urban Canopy Parameterization for Mesoscale Models Based on Computational Fluid Dynamics Reynolds-Averaged Navier–Stokes Microscale Simulations, *Boundary-Layer Meteorology*, 137, 417–439, <https://doi.org/10.1007/s10546-010-9538-4>, 2010.
- 655 Simon, H.: *Modeling urban microclimate: development, implementation and evaluation of new and improved calculation methods for the urban microclimate model ENVI-met*, Ph.D. thesis, Johannes Gutenberg-Universität Mainz, Mainz, <https://doi.org/10.25358/openscience-4042>, 2016.
- 660 Steadman, R. G.: The assessment of sultriness. Part I: A temperature-humidity index based on human physiology and clothing science, *Journal of Applied Meteorology and Climatology*, 18, 861–873, 1979.
- Theeuwes, N. E., Steeneveld, G.-J., Ronda, R. J., and Holtslag, A. A. M.: A diagnostic equation for the daily maximum urban heat island effect for cities in northwestern Europe, *International Journal of Climatology*, 37, 443–454, <https://doi.org/10.1002/joc.4717>, 2016.
- 665 Thorsson, S., Lindqvist, M., and Lindberg, F.: Different methods for estimating the mean radiant temperature in an outdoor urban setting, *International Journal of Climatology*, 27, 1983–1993, <https://doi.org/10.1002/joc.1537>, 2007.



- Thorsson, S., Lindberg, F., Björklund, J., Holmer, B., and Rayner, D.: Potential changes in outdoor thermal comfort conditions in Gothenburg, Sweden due to climate change: the influence of urban geometry, *International Journal of Climatology*, 31, 324–335, <https://doi.org/https://doi.org/10.1002/joc.2231>, 2011.
- 670 Tolan, J., Yang, H.-I., Nosarzewski, B., Couairon, G., Vo, H. V., Brandt, J., Spore, J., Majumdar, S., Haziza, D., Vamaraju, J., Moutakanni, T., Bojanowski, P., Johns, T., White, B., Tiede, T., and Couprie, C.: Very high resolution canopy height maps from RGB imagery using self-supervised vision transformer and convolutional decoder trained on aerial lidar, *Remote Sensing of Environment*, 300, 113 888, <https://doi.org/https://doi.org/10.1016/j.rse.2023.113888>, 2024.
- Toparlar, Y., Blocken, B., Maiheu, B., and van Heijst, G.: A review on the CFD analysis of urban microclimate, *Renewable and Sustainable Energy Reviews*, 80, 1613–1640, <https://doi.org/https://doi.org/10.1016/j.rser.2017.05.248>, 2017.
- 675 Wallenberg, N., Holmer, B., Lindberg, F., Lönn, J., Maesel, E., and Rayner, D.: A simple step heating approach for wall surface temperature estimation in the SOLar and LongWave Environmental Irradiance Geometry (SOLWEIG) model, *EGUsphere*, 2025, 1–20, <https://doi.org/10.5194/egusphere-2025-2093>, 2025.
- Yamazaki, D., Ikeshima, D., Tawatari, R., Yamaguchi, T., O’Loughlin, F., Neal, J. C., Sampson, C. C., Kanae, S., and Bates, P. D.: A high-accuracy map of global terrain elevations, *Geophysical Research Letters*, 44, 5844–5853, <https://doi.org/https://doi.org/10.1002/2017GL072874>, 2017.
- 680 Yang, M., Oh, G., Xu, T., Kim, J., Kang, J.-H., and Choi, J.-I.: Multi-GPU-based real-time large-eddy simulations for urban microclimate, *Building and Environment*, 245, 110 856, <https://doi.org/https://doi.org/10.1016/j.buildenv.2023.110856>, 2023.
- Zanaga, D., Van De Kerchove, R., Daems, D., De Keersmaecker, W., Brockmann, C., Kirches, G., Wevers, J., Cartus, O., Santoro, M., Fritz, S., Lesiv, M., Herold, M., Tsendbazar, N.-E., Xu, P., Ramoino, F., and Arino, O.: ESA WorldCover 10 m 2021 v200, Zenodo, <https://doi.org/10.5281/zenodo.7254221>, version v200, 2022.
- 685 Zhu, X. X., Chen, S., Zhang, F., Shi, Y., and Wang, Y.: GlobalBuildingAtlas: an open global and complete dataset of building polygons, heights and LoD1 3D models, *Earth System Science Data*, 17, 6647–6668, <https://doi.org/10.5194/essd-17-6647-2025>, 2025.
- Zonato, A., Martilli, A., Di Sabatino, S., Zardi, D., and Giovannini, L.: Evaluating the performance of a novel WUDAPT averaging technique to define urban morphology with mesoscale models, *Urban Climate*, 31, 100 584, <https://doi.org/https://doi.org/10.1016/j.uclim.2020.100584>, 2020.
- 690 Zonato, A., Martilli, A., Santiago, J. L., Zardi, D., and Giovannini, L.: On a new one-dimensional $k - \epsilon$ turbulence closure for building-induced drag, *Quarterly Journal of the Royal Meteorological Society*, 149, 1674–1689, <https://doi.org/https://doi.org/10.1002/qj.4476>, 2023.
- 695 Zonato, A., Kamath, H. G., Sudharsan, N., Monaco, L., Kittner, J., Wolf, L., Demuzere, M., Middel, A., Bechtel, B., and Milelli, M.: GLIDE-SOL: A GPU-accelerated Global Lightweight Infrastructure for Diagnostic Environmental Modeling with SOLWEIG, <https://doi.org/10.5281/zenodo.18671813>, version v1, 2026.