Response to Ferdinand

Andrea Zonato and co-authors

30 June 2026

In preparing this revision, we regenerated all model outputs after updating the wind-zone length coefficients in the preprocessor and after adding new global building-data sources to the Dortmund input workflow. The revised simulations, figures, and skill metrics therefore reflect the updated coefficient settings and building morphology inputs.

Response to Ferdinand

Dear Ferdinand,

Opening Comment

> Dear authors,
>
> I think this work is a meaningful and well-directed step towards making high-resolution thermal comfort modelling accessible on a large scale.
>
> The manuscript introduces GLIDE-SOL: a fully scripted, globally deployable Python workflow that utilizes SOLWEIG to enable the rapid and repeatable mapping of thermal comfort across diverse urban environments. By integrating harmonized global inputs with physics-based diagnostics, GPU acceleration and scalable tiling, GLIDE-SOL addresses a significant gap in current practice by providing fast, globally applicable prediction tools for urban heat stress. I strongly support this approach and recommend publication of the manuscript after minor revisions.

Response to Opening Comment

We thank the reviewer for the supportive assessment. We have addressed the requested clarifications as minor revisions to the Methods, Results, Discussion, and figure captions.

Main Comments

Comment 1

> UHI correction. Potential double-accounting (vegetation). The current formulation, which combines SVF and vegetation fraction (fveg), raises the question of whether the effect of vegetation on the urban heat island (UHI) signal is accounted for twice. It would be helpful to explicitly justify why these two terms appear together in the formula

> and to clarify the radius used to compute fveg. It should also be clarified whether fveg refers to trees only or includes low-lying vegetation too (or only), and whether SVF accounts for buildings and vegetation simultaneously.

Response 1

Manuscript changes.

Revised manuscript lines 270–285.

We agree. We revised the UHI section to clarify that the SVF used in the UHI modulation accounts for buildings, terrain, and tree-canopy obstructions simultaneously, while low-lying vegetation is not included in SVF because it does not obstruct the sky at pedestrian level. Vegetation is also represented by `fveg`. In the Dortmund implementation, `fveg` is a local 0/1 indicator on the 2 m SOLWEIG grid: it is one for tree-canopy cells and vegetated land-use cells, including low vegetation where mapped as vegetation, and zero for built, water, and bare cells.

We also added the rationale for keeping SVF and `fveg` together. SVF represents radiative enclosure and reduced nocturnal sky exposure, while `fveg` represents the additional cooling tendency of vegetated surfaces through evapotranspiration and reduced heat storage. Including tree canopy in `fveg` is physically consistent with this diagnostic role: the vegetation term damps the applied UHI perturbation, whereas the tree contribution to sky obstruction acts through the SVF/radiative geometry and can reduce longwave loss to the sky, thereby tending in the opposite direction. The two terms therefore represent competing effects of vegetation rather than a simple duplicate contribution.

Comment 2

> Spatial resolution of adjusted air temperature (Ta). While Figure 6 provides some indication, the spatial resolution of the adjusted Ta field is not clearly stated in the Methods section (2m?). Could a slightly coarser spatial resolution in fact be more appropriate for the UHI correction, given that micro-scale Ta variability can introduce noise rather than a meaningful signal at very fine resolutions?

Response 2

Manuscript changes.

Revised manuscript lines 288–291.

We clarified that the adjusted Ta field is evaluated and written on the same 2 m grid as the SOLWEIG radiative fields because the model requires co-registered meteorological inputs for each tile. We also added that this should not be interpreted as a claim that physically meaningful UHI variability is resolved exactly at a 2 m atmospheric scale. The UHI amplitude is a city-scale daily value, and the 2 m pattern distributes that amplitude according to local morphology. At the same time, the diagnostic is applied at night, when stable boundary-layer conditions and weak turbulent mixing can support stronger near-surface spatial contrasts than during daytime; recent observations of urban park cooling also show that boundary-layer stability strongly modulates local nocturnal cooling intensity (Haeffelin et al., 2024). Therefore, even if the relevant Ta variability is not strictly a 2 m process, it can still occur at scales of the order of metres to tens of metres in sheltered urban settings. We agree that coarser or spatially smoothed UHI fields may be more appropriate for some applications and now identify this as a future sensitivity test.

Comment 3

> Day vs. nighttime diagnostics. The decision to adjust Ta only for night-time conditions is well-founded, given that Ta is the dominant driver of heat stress during the night. However, a separate evaluation of day versus night Ta performance would be a valuable addition, as it would enable readers to better isolate the effect of the diagnostic correction.

Response 3

Manuscript changes.

Revised manuscript lines 515–541.

We agree. We added text below the aggregate skill table explaining that all-hour Ta statistics blend daytime and night-time behavior. During daytime, Ta is affected mainly by ERA5 forcing and the elevation-based correction, while at night it is additionally affected by the diagnostic UHI cycle. We did not add a full day–night Ta stratification in this revision because the improvement in Ta is modest compared with the improvement in UTCI and wind speed, suggesting that the nocturnal UHI signal is not particularly large for Dortmund over the evaluated period. A seasonal day–night separation could be more informative, especially in winter when stable conditions are more frequent, but it would substantially extend the analysis and move the manuscript away from its main objective, which is the evaluation of UTCI-oriented heat stress performance. We now identify day–night Ta stratification as the most direct next diagnostic for isolating the nocturnal correction.

Comment 4

> Nighttime cooling in open areas (Fig. 4). Figure 4 appears to show higher Ta values in open areas than in forested areas (e.g. the north-west). Since the UHI adjustment is only applied at night, open vegetated areas where [2 - SVF - fveg] approaches zero would be expected to cool more rapidly, showing a reduced or negligible UHI effect. It would be worth discussing whether this behavior is correctly captured by the current formulation. Where does the [2 - SVF - fveg] term come from?

Response 4

Manuscript changes.

Revised manuscript lines 510 and 270–285.

We revised the discussion of the JJA map to clarify that the Ta panel is an all-hour summer mean, not a direct map of the night-time UHI term alone. It therefore combines the elevation-based temperature correction, ERA5 forcing, the local radiative environment, and the nocturnal UHI perturbation. For the specific behaviour noted in the north-west, our interpretation is that cooler forested areas are not inconsistent with the formulation: they remain vegetated cells, so `fveg=1` reduces the applied UHI perturbation, while their lower SVF reflects tree-canopy obstruction and a different local radiative environment from more open areas. We also clarified in the methods that the `[2 - SVF - fveg]` term follows the diagnostic formulation of Theeuwes et al. (2016), which uses SVF and vegetation fraction to modulate the daily maximum UHI intensity. In GLIDE-SOL we apply that formulation on the common SOLWEIG grid as a diagnostic, theory-guided proxy rather than as a locally fitted equation.

Comment 5

> Stratified evaluation by urban context. Presenting the model–observation comparison stratified by land use or urban structure type would be particularly insightful. A breakdown along an urban–rural gradient or by functional zone (e.g. dense urban, suburban, park, peri-urban or canopy over and SVF) would reveal the areas in which the model performs well and where the diagnostic corrections have the greatest impact.

Response 5

Manuscript changes.

Revised manuscript lines 544–553 and Appendix lines 745–751.

We agree. The revised manuscript now makes clearer that the station-level UTCI figure sorted by SVF provides a first stratification along an open-to-enclosed urban gradient, while the exact station-by-station values are retained in an appendix table. We initially tested a more explicit stratified analysis by LCZ and by an urban–rural/open-to-enclosed gradient, but the station sample did not show a clear or robust trend that justified adding a separate stratified analysis in the main paper. We therefore kept the station-level figure in the main text, retained the detailed table in the appendix, added the LCZ class for context, and state that a fuller LCZ, land-use, canopy-cover, or functional-zone stratification should be included in future validation work with a network specifically designed to sample these classes.

Comment 6

> Wind speed comparison. It is noted that 3.3 m wind observations are being compared with 10 m model output. For the sake of completeness, please clarify what the expected outcome would be if observations and model output were compared at the same height (e.g. 3 m). Reporting the R^2 value for wind speed would also be useful.

Response 6

Manuscript changes.

Revised manuscript lines 457–465 and Table 1 lines 518–539.

We clarified this point in the case-study section. The GLIDE-SOL wind variable keeps the technical output name `Va10m` for compatibility with the forcing and SOLWEIG workflow, but in this application it is interpreted as a diagnostic canopy wind speed after urban roughness rescaling and multiplication by the directional morphology-based wind-reduction coefficients. In the observation-based evaluation, we now denote this compared canopy-wind diagnostic by V_{can}, while retaining `Va10m` only as the technical output name. The diagnostic does not resolve a vertical wind profile inside the urban canopy; instead, it assumes that the corrected canopy wind is vertically uniform over the observation height range considered here. Under this assumption, the 3.3 m station observations and the modeled canopy-wind diagnostic can be compared directly without applying an additional logarithmic height conversion.

We did not retain R^2 for wind speed in the revised manuscript. In this dataset, near-canopy wind observations at 3.3 m are strongly sheltered by local morphology and therefore exhibit a narrow dynamic range with many low-wind values. Under these conditions, R^2 does not add useful information: it is dominated by variance matching around a small mean state, is highly unstable to small timing or amplitude mismatches, and can become strongly negative even when the modeled magnitudes are much more realistic. For this diagnostic canopy-wind application, bias and RMSE are therefore the more meaningful metrics, because they directly quantify whether the urban wind correction suppresses the unrealistic high-wind tail inherited from the coarse ERA5 forcing and brings modeled wind magnitudes closer to the observed canopy-scale exposure.

Additional Comments

Comment 7

> The observed shift in UTCI heat stress classes > 26°C appears to be driven primarily by the elevation-based Ta correction and wind attenuation. Is it possible to quantify the relative contributions of each component?

Response 7

Manuscript changes.

Revised manuscript lines 644–644 and 658–660.

We agree that component attribution would be valuable. In the revised manuscript we keep this point concise: the present diagnostics indicate that wind is the dominant all-hour contributor to the UTCI improvement, while future work should further improve the reproduction of air temperature. We retain the broader limitation and future-ablation perspective in the outlook section.

Comment 8

> The 10-cell padding corresponds to approximately 20 m at the stated grid spacing. Why was this threshold chosen and are 20 m sufficient to avoid boundary artifacts in both SVF calculations and wind field estimations?

Response 8

Manuscript changes.

Revised manuscript lines 227–230.

We performed sensitivity tests on the tile halo. The 20 m halo was sufficient to compute SVF correctly at tile edges and avoid visible boundary effects in the simulations, while keeping memory overhead small. We revised the manuscript accordingly.

Comment 9

> In several figures, green and red color ramps are used in close proximity, which can make the plots difficult to interpret. Please revise the color scales to improve visual clarity and accessibility.

Response 9

Manuscript changes.

Revised manuscript figure captions and discussion, lines 544–583 and 599–637.

We agree. We revised the figure captions and text so that interpretation does not depend on red/green naming, and we will use a color-blind-safe palette for the final figure export.