Response to Reviewers

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 Fredrik

Dear Fredrik,

Opening Comment

> Thank you for an interesting read and use of the SOLWEIG-model. A full workflow using GPU-enabled technology for Tmrt calculation is proven to be very efficient for this type a 2.5D-model. Many of you know me as the creator and maintainer of the SOLWEIG-model and based on that, I would like to add some comments regarding the setup and results produced. I must confess that I have only done one read-through, so some of my comments might be explained in the text already. Sorry for that.

Response to Opening Comment

Thank you very much for reading the manuscript and for these important comments. We agree that the SOLWEIG setup and the observational reference for Tmrt needed to be described more explicitly. We have revised the manuscript accordingly.

Comment F1

> Please improve the description on the settings using the model as this have large implications on your results and the interpretation of your findings. You state that you are using v2022a and this includes anisotrophic schemes for the long- as well as the diffuse shortwave sky. Did you implement these settings? Also, did you use the Reindl et al method to partition diffuse and direct shortwave radiation from global or did you get that from ERA5? More details are needed! Especially since the settings could be an explanation to your bias in Tmrt shown in e.g. fig 7, especially in dense urban areas because of the anisotrophic skies (Wallenberg et al. 2020; 2023)

Response F1

Manuscript changes.

Revised manuscript lines 154–204.

We agree. We expanded Sect. 2.1 ("Overview of the SOLWEIG core") to state the exact radiative schemes used in the Dortmund experiments. The simulations use the SOLWEIG v2022a anisotropic diffuse shortwave scheme, where diffuse shortwave radiation is distributed over sky-vault patches using the Perez et al. (1993) sky-luminance formulation rather than treated as an isotropic hemispheric term. They also use the v2022a anisotropic longwave scheme, where sky longwave emission is resolved over sky patches rather than prescribed as a spatially uniform hemispheric contribution.

We also clarified the treatment of radiative forcing in the meteorological-forcing description. ERA5 downward surface shortwave radiation is used as global horizontal shortwave irradiance. We do not take separate direct and diffuse shortwave components from ERA5 for these runs. Instead, SOLWEIG partitions the global shortwave irradiance into direct and diffuse components using clearness-index-based atmospheric transmissivity and the Reindl et al. (1990) relations. We additionally clarified that ERA5 downward longwave radiation is retained for reproducibility and optional experiments, but was not prescribed in the Dortmund simulations; longwave exchange was computed by SOLWEIG and distributed with the anisotropic longwave scheme.

Comment F2

> Using this version of the model requires the correct referencing which should include Wallenberg et al. 2020 and 2023.

Response F2

Manuscript changes.

Revised manuscript lines 140–160 and 617–620; reference list lines 810–846.

We added the requested references to the SOLWEIG setup description and to the discussion of remaining Tmrt uncertainties: Perez et al. (1993), on the sky-luminance distribution used for diffuse shortwave anisotropy; Wallenberg et al. (2020), on anisotropic diffuse shortwave radiation; and Wallenberg et al. (2023), on anisotropic longwave irradiance. We also cite Reindl et al. (1990) where the diffuse/direct shortwave partitioning is described.

Comment F3

> You are using black globe temperature sensors to derive Tmrt which needs to be discussed further. Did you calibrate these against more accurate observations and if not, do you think the bias in e.g. fig 7 could be explained based on this rather that my points in bulletpoint 1?

Response F3

Manuscript changes.

Revised manuscript lines 470 and 691–697.

We agree that the observational derivation of Tmrt should be discussed more carefully. The Dortmund biometeorological network uses Decentlab DL-BLG black-globe sensors, from which Tmrt is derived. The resulting Tmrt observations were not independently calibrated against 3-D net radiometers within this study. We therefore added a statement in the case-study section clarifying that no calibration correction was applied to the derived Tmrt observations.

We also added a short paragraph in the Discussion noting that the Dortmund comparison should be read as an evaluation against the operational biometeorological network rather than as a radiometer-level calibration experiment. We further clarified that the Dortmund simulations already use the anisotropic shortwave and longwave sky schemes, so the remaining Tmrt bias cannot simply be attributed to an omitted isotropic-sky approximation.

References Provided By Reviewer

> Wallenberg, Nils, Lindberg F, Holmer B, and Thorsson S. (2020) "The Influence of Anisotropic Diffuse Shortwave Radiation on Mean Radiant Temperature in Outdoor Urban Environments." Urban Climate 31 (2020). `https://doi.org/10.1016/j.uclim.2020.100589.`
>
> Wallenberg, N., Lindberg, F., Holmer, B., and Rayner, D. (2023) An anisotropic parameterization scheme for longwave irradiance and its impact on radiant load in urban outdoor settings. International journal of biometeorology. `https://doi.org/10.1007/s00484-023-02441-3.`

Both references are now present in `references.bib`, cited in the revised manuscript, and visible in the manuscript diff.

Best wishes,

Andrea and co-authors

Response to Jeremy

Opening Comment

> Dear Authors,
>
> The article describes a GPU based model which is inspired from state-of-the art model. The article aim, literature review, method and results are clearly presented. It is simple to read and useful to the community. The model outputs is rather well examplified with a very interesting dataset.
>
> In my opinion, only three minor issues raise:

Response To Opening Comment

Dear Jeremy,

We thank the reviewer for the positive assessment and for the detailed technical comments. We have expanded the manuscript where the setup, assumptions, and limitations needed to be made more explicit.

General Comments

Comment 1

> Several informations are missing to justify choices made by the authors (shortly listed here but more detailed below)
>
> number of wind directions used
>
> using 10-cells padding
>
> modulation terms for UHI
>
> lapse rate calculated via ERA5 multi-level data
>
> new values or schemes for URock are used
>
> If these choices have been made arbitrarily, a specific paragraph with current limitations should be added. This paragraph could be seen as a set of investigations that could be performed in the future and thus adressed to the research community). If the choice can be clearly justify, please explicitly write it at the corresponding location in the manuscript.

Response 1

Manuscript changes.

Revised manuscript lines 189–193, 217–220, 245–320, 323–415, and 639–652.

We agree. We added explicit text explaining which choices are physically motivated, which are pragmatic engineering choices, and which require future sensitivity analysis. In particular, we clarified the 12 wind directions, the UHI modulation term, the parcel-based elevation correction, the use of ERA5 forcing, and the empirical nature of several wind-model settings. We also clarified that sensitivity tests on the tile halo showed the 20 m padding to be sufficient for SVF calculations at tile edges and for avoiding visible boundary effects. We expanded the limitations paragraph to state that the remaining settings should be tested across additional cities, climates, input datasets, wind-sector resolutions, smoothing lengths, wake-length parameters, and terrain-following wind formulations.

Comment 2

> More specifically concerning the wind model, I think that a short paragraph (or / and table) dedicated to summarizing (and justifying) the differences with other existing Röckle models would make sense (eg. only along-x wind modification, no mass-balance

algo, upwind and wake models for trees, no street-canyon scheme, no backward zones, etc.).

Response 2

Manuscript changes.

Revised manuscript lines 323–335 and 343–415.

We agree. We added a dedicated table in the wind-model section summarizing the main simplifications: 12 precomputed wind sectors, a one-step attenuation field, no iterative mass-conservation correction, building effects applied along the wind-aligned x direction, no ellipsoidal lateral/vertical wake shape, no explicit street-canyon or backward-zone parameterization, vegetation treated as porous obstacles with additional upwind and leeward influence zones, and no terrain-following wind-coordinate transform. We also revised the text to emphasize that the resulting field is a diagnostic wind-attenuation proxy, not a dynamically complete wind solution.

Comment 3

There is no need to have a forge such as a gitlab or a github to publish in gmd but I wonder the objective of the authors if they do not have a platform where they can interact with a user community if there are any bug to solve ?

Response 3

Manuscript changes.

Revised manuscript lines 673–677.

We agree that a citable archive alone is not ideal for user interaction. The revised code-availability section now states that the Zenodo archive provides the citable release, while ongoing user support, bug reports, and community interaction are intended to be handled through the associated public code repository if the paper is officially published.

Detailed Comments

Comment 4

« Tolan et al., 2024 » ("egusphere-2026-776", p. 3) There is no corresponding reference

Response 4

Manuscript changes.

Revised manuscript lines 184–187 and bibliography lines 837–840.

The reference issue has been resolved.

Comment 5

> « the three-dimensional distribution of urban and peri-urban vegetation » ("egusphere-2026-776", p. 3) The previous paragraph definitely deals with urban since buildings are treated. This paragraph deals with vegetation (which is mainly located outside urban areas). Vegetation is known to be more difficult to identify in cities due to shadow and from what I have read, Tolan et al. (2024) do not evaluate the potential difference of results obtained inside or outside cities. I suppose a warning sentence here could be welcome.

Response 5

Manuscript changes.

Revised manuscript lines 80–84.

We agree. We added a warning sentence noting that urban tree detection is affected by shadows, mixed pixels, and building occlusion, and that this is treated as an input-data limitation rather than a locally validated tree inventory.

Comment 6

> « Platforms such as Google Earth Engine enable planetary-scale access and processing of these datasets in a reproducible, scriptable manner (Gorelick et al., 2017) » ("egusphere-2026-776", p. 3) I suppose this sentence located one sentence before would make more sense ?

Response 6

Manuscript changes.

Revised manuscript lines 88–89 and 195–205.

We revised the surrounding text so that Google Earth Engine is presented as the platform enabling reproducible access to the geospatial and meteorological products.

Comment 7

> « igure 1. » ("egusphere-2026-776", p. 5) Processes : clip, merge AND reclassify

Response 7

Manuscript changes.

Revised manuscript line 138.

Typo solved, thanks

Comment 8

> « 12 directional wind-sheltering rasters » (“egusphere-2026-776”, p. 7) Any motivation for using 12 ? Have you tested (or read something) about using more or less splitting to evaluate the differences in terms of accuracy and computation time ?

Response 8

Manuscript changes.

Revised manuscript lines 189–193.

We performed an internal sensitivity test with alternative directional discretizations and found that 12 sectors offered a good compromise between representing the dominant directional anisotropy of the urban fabric and keeping pre-processing/storage moderate. To keep the manuscript focused, we did not include the auxiliary sensitivity-test results in the paper; the revised text only states the rationale for using the 30 degree discretization.

Comment 9

> « Sudharsan » (“egusphere-2026-776”, p. 7) I suppose the reference is only the github ?

Response 9

Manuscript changes.

Revised manuscript lines 208–211 and 676–678.

We retained the GitHub repository as the practical reference for SOLWEIG-GPU and clarified that it served as a reference point for GPU-oriented refactoring.

Comment 10

> « While GLIDE-SOL re-architects the surrounding 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 » (“egusphere-2026-776”, p. 7) I would have started with the second part of the sentence

Response 10

Manuscript changes.

Revised manuscript lines 208–211.

We reordered and clarified the SOLWEIG-GPU paragraph. The revised text first states that the open-source SOLWEIG-GPU implementation served as a practical reference point for GPU-oriented refactoring of the SOLWEIG radiative geometry and time-stepping, and then explains how GLIDE-SOL re-architects the surrounding workflow.

Comment 11

« 10-cell padding is added » ("egusphere-2026-776", p. 7) Same as wind directional, if you have any results to share about the choice of this value, it would be interesting to mention. Or mentionning that it has not be done since it could be something to investigate in the future

Response 11

Manuscript changes.

Revised manuscript lines 217–220.

We added that the 10-cell padding corresponds to 20 m at the 2 m grid spacing used here. We performed sensitivity tests on the tile halo and found that this value is sufficient to compute SVF correctly at tile edges and avoid visible boundary effects in the simulations, while keeping the memory overhead small.

Comment 12

« d,a » ("egusphere-2026-776", p. 8) space missing

Response 12

Manuscript changes.

Typographic correction in revised manuscript lines 245–265.

We corrected the typographic issue.

Comment 13

« ready file to LERC+ZSTD (variable-specific MAX_Z_ERROR budgets of 0.15 °C for UTCI, 0.10 °C for Tmrt, 0.005 m s-1 for Va10m, and a 2 W m-2 default for radiative terms) » ("egusphere-2026-776", p. 8) Not clear

Response 13

Manuscript changes.

Revised manuscript lines 233–236 and 633–636.

We clarified the sentence by spelling out that the listed values are variable-specific maximum absolute compression errors used by the LERC+ZSTD repacking step.

Comment 14

« Rather than using fixed clock times » ("egusphere-2026-776", p. 9) it is not only that you do not use fix clock times, it is also that you define the night according to a specific definition. It would be easier to read if you clearly state it.

Response 14

Manuscript changes.

Revised manuscript lines 255–258.

We revised the UHI section to state explicitly how night is defined: the window around the daily zero-shortwave interval, extended from four hours before SWDOWN reaches zero until two hours after SWDOWN becomes positive after sunrise.

Comment 15

> « SVF(x, y) is the sky-view factor, fveg(x, y) » ("egusphere-2026-776", p. 9) SVF can be a ponctual information, vegetation fraction cannot. Did I miss the information about the zone considered for the calculation of these geographical indicators ? I suppose you do not use the 2 m resolution used for the other calculation, it would lead to far too much heterogeneity for air temperature...

Response 15

Manuscript changes.

Revised manuscript lines 260–266 and 645–650.

We clarified that SVF and the vegetation term are evaluated on the common SOLWEIG grid after resampling. In the Dortmund implementation, the vegetation term is a local 0/1 indicator for tree-canopy cells and vegetated land-use cells, including low-lying vegetation where mapped as vegetation, not a radially averaged vegetation fraction. The SVF is computed from buildings, terrain, and tree-canopy obstructions simultaneously; low-lying vegetation is not included in the SVF because it does not obstruct the sky at pedestrian level. We now state this explicitly and identify neighborhood-scale smoothing as a future sensitivity test.

Comment 16

> « The modulation term (2 - SVF - fveg) attenuates the applied UHI in open or highly vegetated areas, where nocturnal cooling is more effective. » ("egusphere-2026-776", p. 9) Is this a proposed Equation based on observed UHI behavior or the result of any fitting with observation ? I suppose the first (if 2nd, please specify the reference), it might be highlight as a limitation.

Response 16

Manuscript changes.

Revised manuscript lines 267–274 and 639–652.

We now state that (2 - SVF - fveg) is a diagnostic, theory-guided proxy rather than a coefficient fitted to observations, taken from (Theeuwes et al., 2016). We also identify this as a limitation and future sensitivity target.

Comment 17

> Instead of using SVF and fveg which are proxy of radiative and convective processes, what about using directly the wind and radiation calculated by the model for each time step ?

Response 17

Manuscript changes.

Revised manuscript lines 247–274 and 650–652.

We agree that directly using modeled wind and radiation would be more physically complete. However, in the present implementation we kept the UHI modulation consistent with the diagnostic framework of Theeuwes et al. (2016), where the daily UHI amplitude is linked to daily-scale atmospheric conditions rather than to the instantaneous local state at each time step. For this reason, GLIDE-SOL currently uses SVF and `fveg` as lightweight morphology-based proxies to distribute the daily UHI signal in space. We now clarify in the manuscript that a tighter coupling of the UHI modulation to modeled local wind and radiation is a logical future extension.

Comment 18

> « moist-adiabatic lapse rate Γm(T, p) » ("egusphere-2026-776", p. 9) Formula used ?

Response 18

Manuscript changes.

Revised manuscript lines 289–304.

We revised the lapse-rate description to match the implementation more closely. The manuscript now describes the parcel-based elevation correction: dry-adiabatic adjustment below the lifting condensation level, LCL diagnosis from temperature, pressure, and relative humidity, and saturated pseudo-adiabatic integration above the LCL. We added the governing equations and defined the variables.

Comment 19

> « at every time step » ("egusphere-2026-776", p. 9) Same formula for day and for night-time ? What about using the ERA5 multi-level data to derive a lapse rate ?

Response 19

Manuscript changes.

Revised manuscript lines 300–304 and 645–650.

We now state that the same elevation-correction formulation is used during daytime and night-time. We also clarify that the current lightweight workflow does not derive lapse rates from ERA5 pressure-level or model-level data; doing so would be a useful future extension but would add a heavier download time and storage to the automated workflow.

Comment 20

> « Reference wind » ("egusphere-2026-776", p. 9) How do you deal with domains intersecting several ERA5 tiles ? You may end with two (or more) different reference wind over your domain, do you smooth the reference wind values over the domain area ?
>
> If I understand well, when your domain is entirely included within a single ERA5 cell, the reference wind is the same for the whole domain, you do not recalculate the reference wind for each of your subdomain (tile) ?
>
> Would it make sense to interpolate wind over the whole domain cells based on wind value in ERA5 cells located around the concerned ERA5 ?

Response 20

Manuscript changes.

Revised manuscript lines 309–320.

We checked the implementation and revised the manuscript to avoid overstating the meteorological interpolation. For the Dortmund ERA5 setup, GLIDE-SOL samples ERA5 at the centre of the model domain and uses one city-scale background wind time series for all computational tiles. The wind is therefore not recalculated per tile, and no tile-to-tile discontinuity is introduced by the meteorological forcing. We also state explicitly that this does not retain gradients between neighbouring ERA5 grid cells; a smoothly varying, spatially interpolated ERA5 background is a future extension for larger domains.

Comment 21

> « with a simplified version » ("egusphere-2026-776", p. 10) Specify that it is a one step approach, the wind flow resulting is not balanced.

Response 21

Manuscript changes.

Revised manuscript lines 323–335.

We now specify that the approach is one-step and not mass balanced. The resulting field is described as a diagnostic attenuation field rather than a full flow solution.

Comment 22

> « wake » ("egusphere-2026-776", p. 10) using wake here is confusing because in the next section you differentiate upwind and wake regions. Maybe removing wake here make it simple ?

Response 22

Manuscript changes.

Revised manuscript lines 323–329 and 344–365.

We revised the wording around the wind section and describe the upwind and leeward zones explicitly in their own subsections.

Comment 23

> « effective obstacle width W , depth D (along-wind), and height H » (“egusphere-2026-776”, p. 10) what is the definition of these effective length ? The shape of the building may play a role on the way you calculate these parameters. The methods differ between QUIC-URB and URock for example.

Response 23

Manuscript changes.

Revised manuscript lines 325–329.

We added definitions for W, D, and H. W and D are the dimensions of the connected obstacle’s axis-aligned bounding box in wind-aligned coordinates, and H is the mean positive rasterized building or canopy height in the connected component.

Comment 24

> « reduction coefficient is vertically uniform » (“egusphere-2026-776”, p. 10) It seems it is although uniform along y right ? This is definitely a difference with the Röckle approach using ellipsoid shapes (and reduction) both along y and z. This difference should be specified.

Response 24

Manuscript changes.

Revised manuscript lines 330–335.

We added that the reduction is vertically uniform up to obstacle height and laterally uniform across the affected obstacle column before smoothing. We explicitly note that this differs from Röckle/URock models using ellipsoidal or tapered lateral/vertical wake shapes.

Comment 25

> « 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, » (“egusphere-2026-776”, p. 10) My feeling: make it simple, the first sentence was clear, you can directly say that it works differently in the canopy-flow parameterization

Response 25

Manuscript changes.

Revised manuscript lines 330–341 and 372–394.

We simplified this part and moved the emphasis to the vertical-uniformity assumption and its difference from the canopy-flow treatment.

Comment 26

> « 2W » (“egusphere-2026-776”, p. 10) Any motivation for setting 2W instead of the 1.5W used in Kaplan and Dinar (1996) or Bernard and al (2023) ?

Response 26

Manuscript changes.

Revised manuscript lines 344–350.

We thank the reviewer for raising this point. We revised the manuscript to state the implemented height-dependent expression Lf = 1.5 W / (1 + 0.8 W/H). We describe it as a short diagnostic upwind attenuation ramp rather than a full Röckle displacement-zone parameterization, because the leading 1.5 factor is used within an aspect-ratio-dependent expression and not as a fixed 1.5W distance.

Comment 27

> « 4 » (“egusphere-2026-776”, p. 11) Same as previous section, why 4 instead of 3 used by Kaplan and Dinar (1996) or Bernard et al. (2023) ?

Response 27

Manuscript changes.

Revised manuscript lines 355–364.

We thank the reviewer for catching this. The values given in the previous manuscript version were incorrect. We revised the manuscript to use Lr = 3 * 1.8 W / [(D/H)^0.3 (1 + 0.24 D/H)]. The factor 1.8 defines the cavity-zone scale, while the factor 3 is used to represent the combined cavity-zone plus wake-zone influence length, consistent with the Röckle/Kaplan–Dinar convention adopted in Bernard et al. (2023), where the wake-zone extent is three times the cavity-zone length. We also corrected the implementation accordingly.

Comment 28

> « while ensuring that overlapping wakes combine conservatively » (“egusphere-2026-776”, p. 11) Not clear, wake coefficient multiply each over when they overlap ?

Response 28

Manuscript changes.

Revised manuscript lines 364–365 and 405–414.

Thank you for pointing this out. We agree that the original wording was ambiguous. In the revised implementation and manuscript, overlapping wake contributions are combined multiplicatively and then clipped to the prescribed coefficient range. This applies to overlapping contributions within the same obstacle class as well as to the final combination of building-induced and tree-induced wake fields. We have clarified this explicitly in the revised text.

Comment 29

> « The model depends on an effective leaf area index » (“egusphere-2026-776”, p. 11) This is not the case in Bernard et al. (2023). If I understand well, the bigger a vegetation patch, the bigger the attenuation ? I think it is rather relevant and could lead to improved results, it should be noted that it differ from Bernard et al. (2023).

Response 29

Manuscript changes.

Revised manuscript lines 367–394 and 639–652.

We added that the vegetation treatment differs from URock: trees are treated as porous obstacles with canopy attenuation plus additional upwind and leeward influence zones. We agree that this is a potentially useful extension, but also identify it as a choice requiring further validation.

Comment 30

> « Given Ct,base at the obstacle location, we apply the same wake geometry as for buildings, with an upwind deceleration over » (“egusphere-2026-776”, p. 12) There is no such zones in URock or Röckle, please specify it. How these choices were made ? Arbitrarily, did you have tests and validations on a limited subset, etc. ?

Response 30

Manuscript changes.

Revised manuscript lines 367–405 and 639–652.

We now explicitly clarify that the application of building-like upwind displacement and leeward wake zones to trees is not part of the original URock/Röckle formulation, but a simplified diagnostic extension introduced in GLIDE-SOL. We adopted this choice because it appeared physically reasonable that trees, as porous roughness elements, also induce upstream flow adjustment and downstream wake effects. This assumption was further motivated by recent literature showing coherent wake structures behind isolated trees, such as Grandoni et al. (2026), “The Dynamics of a Wake Behind an Isolated Model Tree Within an Atmospheric Boundary Layer”. At the same time, we acknowledge that this parameterization is a pragmatic first-order approximation for urban preprocessing, rather than a formally validated vegetation scheme within URock/Röckle. We therefore revised the manuscript to make clear both the rationale for this choice and its present limitation, and we note that its quantitative sensitivity and validation should be addressed in future work on dedicated test cases.

Grandoni, L., Michard, M., Grosjean, N., and Salizzoni, P.: The Dynamics of a Wake Behind an Isolated Model Tree Within an Atmospheric Boundary Layer, Boundary-Layer Meteorology, 192, 26, 2026. `https://doi.org/10.1007/s10546-026-00971-y`

Comment 31

> « interiors » (“egusphere-2026-776”, p. 12) what do you mean by building interiors ? Courtyards ? If so why ?
>
> If it really makes reference to the building interior, I do not think it makes sense to keep it since obviously you do not consider building interior...

Response 31

Manuscript changes.

The wording was removed from the revised wind-model section; see the revised section at lines 323–415.

We agree and removed this wording from the manuscript.

Comment 32

> « effective kernel of about 40 m » (“egusphere-2026-776”, p. 12) A reason for 40 m ?

Response 32

Manuscript changes.

Revised manuscript lines 405–414 and 645–650.

We clarified that the wind coefficients are first smoothed at a short scale and then with a 40 m neighborhood-scale Gaussian smoother. The 40 m scale is used to reduce raster/rotation artifacts, blend overlapping wakes, and partially compensate for the fact that GLIDE-SOL does not prescribe an ellipsoidal lateral/vertical wake envelope as in fuller Röckle/URock implementations. It is therefore a pragmatic plan-view regularization scale, not a calibrated physical wake length. It is included among future sensitivity tests.

Comment 33

> « to blend overlapping wakes and suppress blocky artifacts » (“egusphere-2026-776”, p. 12) I do not clearly see what do you mean

Response 33

Manuscript changes.

Revised manuscript lines 405–414.

We clarified that this refers to reducing discontinuities caused by rasterized obstacle edges and directional rotation, while merging adjacent wake footprints into smoother attenuation fields. This smoothing also regularizes the simplified non-ellipsoidal wake geometry used by GLIDE-SOL.

Comment 34

> « Decentlab DL-SHT35 sensor » (“egusphere-2026-776”, p. 13) Could you please specify the shield used since it might be more important than the sensor itself ?

Response 34

Manuscript changes.

current station-instrument description is at revised manuscript lines 423–426.

Thank you for pointing this out. The shield used is the manufacturer’s default which is a relatively small shield with a diameter of approximately 8–10 cm. Due to the lack of active ventilation and the small size of the shield, a small radiation error is likely. When deploying the network, four sites were equipped with both types of sensors, a DL-ATM41 as well as a DL-SHT35. Initial, hourly comparisons of both sensors over almost two years show a median positive deviation of 0.4 K between 9–12 UTC. However, it should be noted that the DL-ATM41 also performs an internal radiation error correction which may, under certain circumstances, over-correct the air temperature, hence, arbitrarily increasing the described error. An outdoor field comparison with an independent, non self-correcting, actively ventilated radiation shield was not performed.

Comment 35

> « Figure 2. » (“egusphere-2026-776”, p. 14) You specify that stations were selected based on LCZ. It might be interesting to show either the LCZ map to appreciate the surrounding of each station, either coloring each circle with the LCZ color scheme.

Response 35

Manuscript changes.

Revised manuscript Figure 2 caption line 472; LCZ context in lines 464, 547–553, and Appendix Table lines 750–755.

We agree. We revised the domain figure so that the station labels are colored according to the observed LCZ class reported in the station metadata retrieved from the public station API (`https://api.data2resilience.de/v1/stations/metadata`), using the standard LCZ color scheme. We now cite Demuzere et al. (2022) as the literature reference for the LCZ framework and color convention, while clarifying that the station classes themselves come from the field-based metadata rather than from the global raster product. We also revised the station-error table so that the LCZ column uses the same field-based station classifications rather than values sampled from the global LCZ raster, because the metadata assignment better represents the immediate surroundings of the sensors and avoids local mismatches at some station locations.

Comment 36

> « Figure 3 » (“egusphere-2026-776”, p. 15) I am quite surprised by the shape of the acceleration zones. It does not look like the first step of a Röckle based model but rather after the second step (mass-conservation).
>
> More specifically, how do you explain the quite low wind reduction downwind the large parking building (5 storey in OSM) at coordinate 51.5069409°N, 7.4598144°E ?
>
> Same near a large (but probably not that high) "industrial" building at coordinate 51.5208217°N ,7.4848979E,
>
> etc.

Response 36

Manuscript changes.

Revised manuscript Figure 3 caption line 492 and discussion lines 495–498.

We thank the reviewer for this observation. We agree that the patterns in Figure 3 deserve further clarification. First, we note that a wind reduction is indeed present in the cited locations. For the large parking building at 51.5069409°N, 7.4598144°E, the wind coefficient immediately downwind is approximately 0.40, while for the industrial building at 51.5208217°N, 7.4848979°E it is approximately 0.80. Therefore, both buildings generate a wake region characterized by reduced wind speed relative to the background flow.

The extent of the wake is strongly controlled by building height. In the input dataset, the parking building is represented with a height of approximately 7 m, while the industrial building is represented with a height of approximately 3 m. The former is likely underestimated by the GBA buildings dataset. Since wake length scales with obstacle height, the relatively low height assigned to the industrial building results in a very short wake region. Consequently, further downwind the local wind coefficient is primarily influenced by the surrounding urban morphology rather than by the building itself. In particular, the industrial building is located within a relatively wide street canyon, where the attenuation associated with the obstacle rapidly vanishes.

In addition, the wind direction shown in Figure 3 is not exactly aligned with the east–west orientation of the buildings. The figure corresponds to a wind direction of 240°, implying a substantial southerly component. As a result, some areas that appear geometrically downwind when viewed on the map are not located directly within the wake of the buildings for the simulated wind direction and therefore experience only limited influence from these obstacles.

Finally, we emphasize that local patterns around individual buildings should be interpreted in light of the rasterized building geometry, wind-sector discretization, rotation back to the native grid, and the subsequent smoothing procedures applied in the workflow. The resulting coefficient field is intended to represent neighbourhood-scale variations in wind attenuation rather than validated recirculation structures around specific buildings.

Comment 37

> « Its effect is therefore » (“egusphere-2026-776”, p. 15) Not sure that I understand the reason for not having the map. A map is calculated for each time step, right ?

Response 37

Manuscript changes.

Revised manuscript lines 498–498.

The reviewer is right: the previous wording was confusing. The model output includes air temperature with the diagnostic UHI signal already applied. If the UHI perturbation is needed as a separate map, it can be obtained directly by subtracting the corresponding background air-temperature field from the adjusted air-temperature output. We revised the manuscript text accordingly.

Comment 38

> « air temperature » (“egusphere-2026-776”, p. 16) I may have already had asked but how do you deal with the discontinuity of ERA5 variables between 2 cells if they intersect your domain ? Do you interpolate some of the ERA5 variables over the whole domain ?

Response 38

Manuscript changes.

Revised manuscript lines 322–336.

We corrected and clarified this response to match the implementation. For the Dortmund ERA5 setup, the current workflow uses one domain-centre ERA5 sample to derive a domain-consistent background forcing time series. This avoids discontinuities between computational tiles, because all tiles share the same background forcing, but it does not represent spatial gradients between neighbouring ERA5 cells. We identify spatial interpolation of ERA5 fields over very large domains as a future improvement. However, this would also introduce a substantial storage and data-management burden, because the interpolated meteorological fields would need to be carried on, or downscaled to, the GLIDE-SOL grid resolution. We are therefore considering how to implement this without compromising the lightweight storage strategy of the workflow.

Comment 39

> « wind speed » (“egusphere-2026-776”, p. 16) It is surprising that none of the station is located in a high wind speed area. None of them are located in a low vegetation, bare soil or impervious LCZ type ? This question is consistent with my previous concern about the missing LCZ map (at least station)

Response 39

Manuscript changes.

Revised manuscript lines 495–505 and 544–553.

We agree that station exposure helps interpret why stations do not necessarily sample the full range of mapped wind-speed conditions. We clarified that the biometeorological stations were selected primarily for urban thermal-comfort monitoring, with coverage of LCZ and public-interest settings, rather than to sample the full range of modeled ventilation corridors. Therefore, the highest modeled wind-speed areas are not necessarily represented by station locations. At the same time, because the network is intended for urban-comfort applications, the stations are generally placed in settings where buildings and urban morphology influence the local exposure for at least some wind directions, even if not for all directions.

Comment 40

> « OHS 0.23 0.64 » ("egusphere-2026-776", p. 18) same as next comment
>
> « HAP 0.38 0.86 » ("egusphere-2026-776", p. 18) what is the reason for the sometimes large differences between SVF obs and mod ?

Response 40

Manuscript changes.

Revised manuscript lines 553–553 and Appendix Table lines 748–755.

We added text explaining that differences between observed and modeled SVF mainly reflect representativeness: observed SVF characterizes the exact sensor surroundings, while modeled SVF samples a 2 m grid cell from globally derived building and vegetation rasters. We also clarified that open-source input datasets are not exact at this scale, especially for vegetation. Therefore, discrepancies can be larger in urban canyons with street trees or small vegetation elements, where the local obstruction seen by the sensor may differ from what is represented in the building and vegetation rasters.

Comment 41

> « Figure 5. » ("egusphere-2026-776", p. 19) On Figure 5 and following, it might be difficult to see the SOL_STD points due to the high density of points. Any chance to replace the scatter by curves showing the mean line + confidence interval (might be 0.25 and 0.75 percentiles) ?

Response 41

Manuscript changes.

Revised manuscript lines 556–583.

We agree that density can obscure SOL_STD points in high-density scatter plots. We nevertheless prefer to retain the scatter plots because they provide an immediate overview of the overall improvement and of the point-wise bias structure, even if individual points cannot always be inspected precisely. At the same time, the following PDF figure provides the complementary distributional view: it summarizes the observed and modeled distributions and therefore supports the physical interpretation of the scatter-plot improvement.

Comment 42

> « incorporating flow-dependent predictors is a promising avenue for future development. » ("egusphere-2026-776", p. 24) excellent !

Response 42

Manuscript changes.

Revised manuscript lines 248–265 and 644–649.

We thank the reviewer for this positive feedback. We retained and sharpened this point by clarifying that the UHI amplitude already uses daily ERA5 radiation, temperature range, and wind-speed normalization, while more explicitly spatial and time-varying local wind/radiation coupling remains a future development target.

Comment 43

> « the use of a simplified urban wind-reduction model that cannot fully capture local channeling or sheltering » ("egusphere-2026-776", p. 25) I suppose terrain elevation might also something missing and which could be improved for example using approaches such as proposed in Robinson et al (2022)
>
> Robinson, David, Sara Brambilla, Michael J. Brown, Patrick Conry, Bryan Quaife, et Rod R. Linn. « QUIC-URB and QUIC-Fire Extension to Complex Terrain: Development of a Terrain-Following Coordinate System ». Environmental Modelling & Software 159 (janvier 2023): 105579. `https://doi.org/10.1016/j.envsoft.2022.105579.`

Response 43

Manuscript changes.

Revised manuscript lines 642–660; bibliography entry added in references.bib.

We expanded the limitations to include the lack of terrain-following wind coordinates and listed terrain-aware diagnostic wind formulations as a future improvement. We now cite the Robinson et al. (2023) QUIC-URB/QUIC-Fire terrain-following coordinate-system paper in the manuscript limitations section and added it to the bibliography.

Comment 44

> « Code and data availability » ("egusphere-2026-776", p. 26) There is no need to have a forge such as a gitlab or a github to publish in gmd but I wonder the objective of the authors if they do not have a platform where they can interact with a user community ?

Response 44

Manuscript changes.

Revised manuscript lines 673–676.

We added that the Zenodo archive provides a citable release, while ongoing user support, bug reports, and community interaction are intended to be handled through the associated public code repository and issue tracker if the paper is officially published.

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.