

Author response to the referees

Manuscript: egusphere-2026-2076

Title: Slope-Constrained Orography (SCO v1.0) for explicit discrete-slope control in terrain preprocessing: a case study with WRF-ARW (v4.6.1)

We thank both referees for their careful review and constructive suggestions. We have corrected the geometric interpretation and solver description, clarified the time-step settings, and revised the interpretation of the stability diagnostics. The revised manuscript retains the experiment set, reported numerical results and central finding of stable integration with substantially less terrain modification.

Response to Referee #1

We thank the referee for the careful review and the positive assessment of the method. We have also revised unclear wording and clarified the explanations identified by the referee. We respond to each comment below and describe the corresponding changes in the revised manuscript. Line numbers in the referee comments refer to the reviewed manuscript. Page and line numbers given under “Changes” refer to the revised manuscript without tracked changes. Section, figure, table and equation numbers are the same in both versions.

R1.1 — General comment 1: spectral response

The authors briefly investigate the spectral response of their filter compared to the WPS filter, but they only depict two SCO examples in their plot (Figure 8). Global spectral filters remove a well-defined portion of the spectrum of variation, local filters such as diffusion operators or gaussian filters have a spectral response that concentrates at the scale of the kernel size but increases with the number of filter applications. Do the authors have a feel for whether the SCO method has a relatively stable spectral response, or is it likely to vary significantly from one terrain dataset to another? How do they see it sitting between the extremes of spectral filters (well-defined response function) and local filters (response function depends on how aggressively the filter is applied).

Response. We expect the spectral response of SCO to vary between terrain datasets and have made this dependence explicit in the revised manuscript. SCO is a nonlinear, local adjustment driven by exceedances of neighbour-pair working bounds derived from θ . An input terrain field that satisfies all initial working pairwise bounds requires no adjustment. The spectral effect of SCO therefore depends on the amplitude and spatial structure of the input terrain, the grid spacing and θ , rather than on a terrain-independent transfer function.

In relation to the two types of filter described by the referee, SCO resembles a local filter in that it acts through neighbouring grid points, but controls the amount of adjustment through θ rather than through a filter width or pass count. A lower θ modifies more terrain in the tested sequence: across the five targets, the modified fraction rises from 6.5 % at $\theta = 26^\circ$ to 26 % at $\theta = 18^\circ$, and MAE_h from 3.3 to 17.9 m (Table 1). SCO has no prescribed spectral cutoff, and its user-facing control is a slope target rather than a smoothing-kernel width. In the spectra reported in Fig. 8, SCO22 and SCO24 retain most of the RAW power at both 2 and 10 km, and SCO24 retains more than SCO22, whereas WPS120 removes most of the power at these scales.

Figure 8 shows SCO22 and SCO24 because they lie on either side of the sampled SCO stability boundary: the integration with $\theta = 22^\circ$ passed and that with $\theta = 24^\circ$ failed (Table 1). The figure serves this comparison and the one with W1+SCO24 in Sect. 4.2. It was not intended as a general characterisation of the SCO response, and it describes the present terrain only.

Changes. Sect. 4.2 (p. 9, lines 196–200) now explains why SCO22 and SCO24 are shown in Fig. 8 and states that the spectral response depends on terrain amplitude and structure, grid spacing and θ rather than a fixed transfer function. It also notes that lower θ increases terrain modification in the tested sequence. The Fig. 8 caption (p. 13) now states that row-mean terrain power is normalised by RAW at each wavenumber.

R1.2 — General comment 2: variability of the simulated fields

The authors, in evaluating the quality of the simulations, only test bias and RMS errors vs. observations in this study. But these are not necessarily the most informative metrics for testing performance of convection-permitting resolution simulations, and it is the reproduction of qualitatively realistic detail evidencing the better reproduction of hitherto unresolved processes (such as convection, small scale gravity waves, complex terrain boundary layer behaviour) that is often desirable on increasing resolution. Are they able to comment on the variability that occurs in meteorological fields within the simulations? For instance, Sheridan et al. (2023) were able to demonstrate increased variability (resolved turbulence, essentially) in wind fields with an improved smoothing method (one similar in effect to SCO).

Response. We agree that bias and RMSE alone do not characterise the spatial and temporal variability of high-resolution flow. The primary contribution of this study is stable integration with substantially less terrain modification. Stability outcomes and terrain-retention diagnostics directly support this contribution, while hourly station errors provide a complementary assessment of the successful configurations.

The terrain diagnostics are relevant to the referee’s concern because they quantify the ridge–valley structure retained in the model’s lower boundary. SCO22 retains 97–99 % of RAW ridge-to-valley relief along two representative transects, compared with 79–86 % for WPS120 (Fig. 4), and 91 % of RAW row-mean terrain power near 10 km wavelength, compared with 25 % (Fig. 8; Table 2). SCO22 thus meets the integration criterion while preserving substantially more terrain structure available to force local flows.

Sheridan et al. (2023) illustrate the value of examining local-flow variability in addition to bulk errors. Scale-dependent atmospheric variability was not evaluated here. The lower near-surface errors of the successful SCO configuration (Table 3) support its practical performance in this case, but do not establish improved reproduction of convection, gravity waves or boundary-layer flow structure. We have clarified this distinction in the Discussion.

Changes. The final paragraph of Sect. 5 (p. 14, lines 230–234) now explains the relevance of retained terrain structure to local-flow forcing and distinguishes the hourly station-error evaluation from an assessment of resolved-flow variability, with reference to Sheridan et al. (2023).

Specific comments

R1.3 — Lines 41–42.

41-42. The workflow aspects seem to have little to do with the constraint “theta”, and it’s confusing that these aspects are conflated in a single sentence.

Response. We have separated these two aspects in the overview (Sect. 2.1). The text first describes the constraint implied by θ , followed by a separate paragraph explaining where SCO sits in the WPS workflow.

Changes. Sect. 2.1, p. 2, lines 40–51.

R1.4 — Line 42.

42. Some introduction of what “geogrid” and “metgrid” represent (perhaps within an overview of the WRF LAM workflow) would be useful to non-WRF users.

Response. We have added a brief explanation of geogrid and metgrid and the position of SCO to Sect. 2.1. geogrid interpolates static geographical data, including terrain, to the model grids, while metgrid interpolates meteorological fields horizontally to the same grids. SCO modifies the terrain in the geo_em files after geogrid and before metgrid. For completeness, ungrib extracts the meteorological fields before metgrid, and real.exe subsequently produces the initial and lateral boundary conditions.

Changes. Sect. 2.1, p. 2, lines 47–49.

R1.5 — Lines 43–44.

43-44. *The demonstration would be best at the end of 2.2, not here.*

Response. We have moved the description of the idealised example (Fig. 1) to the end of Sect. 2.2, after the discrete constraint has been defined.

Changes. Sect. 2.2, p. 3, lines 71–72; Fig. 1 now appears on p. 4.

R1.6 — Line 46.

46. *“the local bound” should be “a local bound” - you haven’t explained the local bound yet.*

Response. We have changed “the local bound” to “a local bound”.

Changes. Sect. 2.1, p. 2, line 42.

R1.7 — Line 49.

49. *Instead of redirecting the reader to the Appendix here, say “explanation to follow”. The overview should be simple and gentle.*

Response. We have replaced the Appendix reference in the overview with a forward reference to Sect. 2.2. Appendix A1 is now cited with the geometric explanation in that section.

Changes. Sect. 2.1, p. 2, lines 43–44; Appendix A1 is cited in Sect. 2.2, p. 3, line 65.

R1.8 — Lines 49–50.

49-50. *Mention of curvature constraints is distracting here when the reader is trying to understand the existing algorithm - by all means mention it somewhere more appropriate such as the conclusions.*

Response. We have moved the mention of curvature constraints from the overview to the Conclusions, where we identify them as a possible extension.

Changes. Removed from Sect. 2.1; added to the Conclusions, p. 14, lines 246–247.

R1.9 — Line 60.

60. *Equation (1) appears to indicate the maximum of the slope over the eight neighbour pairs. What’s the significance of this maximum? It’s not discussed in the text.*

Response. The maximum detects an exceedance in any of the eight sampled directions: the local slope is defined by the steepest neighbour-pair height-difference ratio. An average over the eight directions could hide a single steep pair behind gentler ones. We have added this explanation after Eq. (1).

Changes. Sect. 2.2, p. 3, lines 57–58.

R1.10 — Line 63.

63. *$S_{\max, \text{int}}$ was already introduced at line 58, why is it mentioned again here? Why are you referring to section 2.3?*

Response. We have removed the repetition and now define $S_{\max, \text{int}}$ once, after Eq. (1). The reference to Sect. 2.3 pointed to the definition of the active interior. We now give a brief definition at this point instead.

Changes. Sect. 2.2, p. 3, lines 58–59.

R1.11 — Lines 65–66.

65-66. *If the cardinal elevation difference constraint is $\Delta x \cdot \tan(\theta)$ then the diagonal constraint is $\sqrt{2} \cdot \Delta x \cdot \tan(\theta)$. This constraint is exactly in keeping with the slope θ , so why do you say the diagonals are under-constrained? The octagonal correction is described as removing the directional anisotropy, but all it does is multiply the constraint by $\cos(\pi/8)$ - for both the cardinal and diagonal pairs. The authors cite this as the reason $S_{\max, \text{int}}$ settles near θ_{eff} instead of θ . So it merely tightens the constraint uniformly, across both cardinal and diagonal directions, making no alteration to any directional / anisotropy aspect.*

Response. The referee is correct. The cardinal and diagonal bounds each correspond to θ in their sampled direction. The common factor $\cos(\pi/8)$ tightens both bounds by the same proportion and does not remove the directional anisotropy of the stencil. We have removed the incorrect claims and revised the geometric explanation.

The factor provides a conservative bound for a locally planar surface. Its gradient direction lies within $\pi/8$ of one of the eight sampled directions, so the largest sampled height-difference ratio, $|h_q - h_p|/d_{pq}$, is at least $|\nabla h| \cos(\pi/8)$. If every sampled ratio satisfies the bound $\tan \theta \cos(\pi/8)$, then $|\nabla h| \leq \tan \theta$ for any orientation of that plane. This is a locally planar interpretation, not rotational invariance of the stencil. The angle corresponding to the initial working bounds is $\theta_{\text{eff}} = \arctan[\tan \theta \cos(\pi/8)]$. The achieved maximum also depends on local tightening and inner convergence, as explained in R1.14.

Changes. Sect. 2.2 (p. 3, lines 60–65) and Appendix A1 (pp. 14–15, lines 252–259) now describe the factor as a conservative geometric scaling with the justification above. The statements that the diagonals are under-constrained and that the correction removes directional anisotropy have been deleted.

R1.12 — Line 76.

76. *Presumably the authors mean columns as well as rows.*

Response. We have corrected both descriptions to read “rows and columns”, covering the frozen rim and the taper zone.

Changes. Sect. 2.3, p. 3, lines 75–76; Algorithm A1, p. 16.

R1.13 — Line 84.

84. *L_{pq} is not defined, or linked to the explanation of the bounds in section 2.2.*

Response. We have defined L_{pq} where it first appears in Sect. 2.3: $L_{pq} = (d_{pq}/\Delta x) \min(f_p, f_q) \Delta h_{\max}^{(0)}$, where $d_{pq} = \Delta x$ for cardinal and $\sqrt{2} \Delta x$ for diagonal pairs, $\Delta h_{\max}^{(0)}$ is the cardinal bound of Eq. (2), and f is the local tightening factor ($f = 1$ initially). The text distinguishes f , which changes the bound, from the mobility weights w , which only determine how the excess is shared between the two endpoints.

Changes. Sect. 2.3, p. 4, lines 84–87.

R1.14 — Lines 98–99.

98-99. *“Stopping rule”, “criterion”: mixing terminology like this is confusing, pick a term and stick with it. If the convergence criterion references θ , and not θ_{eff} , why do the results in Table 1 cluster at θ_{eff} ?*

Response. Our previous description did not clearly distinguish the inner stopping condition from the outer acceptance test. We now use “inner stopping condition” for the sweep loop and “outer acceptance test” for the check $S_{\max, \text{int}} \leq \theta + 0.01^\circ$.

The initial neighbour-pair bounds correspond to $\tan \theta_{\text{eff}} = \tan \theta \cos(\pi/8)$. These bounds control the pairwise adjustments, whereas the outer test determines whether the resulting terrain is accepted. The achieved maximum can therefore lie near θ_{eff} even though the outer test references θ , as seen for SCO24 and SCO26

in Table 1. Local tightening can reduce the working bounds further. Passing the outer test does not, by itself, establish convergence to those internal bounds.

We also omitted the stagnation condition from the previous algorithm description. The inner loop stops when the maximum pairwise exceedance monitored during sweep s , E_s , is below 0.1 m, when $E_s \geq 0.999E_{s-1}$ for three consecutive comparisons, or after 500 sweeps. The monitor includes movable pairs adjoining the frozen boundary. Persistent exceedances there can trigger this exit before the interior has fully converged. The final maxima reported for SCO18–SCO22 satisfy θ but remain above θ_{eff} (Table 1; Fig. 5c), so these fields meet the nominal target without fully satisfying the tighter internal working bounds. The terrain fields used in the reported WRF experiments are unchanged.

Changes. Sect. 2.1 (p. 2, lines 44–45) no longer states that the solver iterates until every constraint is satisfied. Sect. 2.3 (p. 3, lines 78–81; p. 5, lines 96–106) and Appendix A2 (pp. 15–16, lines 261–265; Algorithm A1 on p. 16) now define the monitored quantity, the tolerance, stagnation and iteration-limit stopping conditions, and the monitoring and outer acceptance regions. Appendix A2 reports nominal-target acceptance separately from inner stopping. Sect. 2.2 (p. 3, lines 68–70), Sect. 4.1 (p. 8, lines 173–175) and the Fig. 5 caption (p. 11) distinguish attainment of the nominal target from convergence to the internal bounds.

R1.15 — Section 2.4.

Section 2.4 - explain the purpose of this important subsection at its outset.

Response. We have added an opening sentence to Sect. 2.4 explaining its purpose: the diagnostics quantify achieved slopes, the magnitude and spatial extent of terrain modification, and the small-scale structure retained.

Changes. Sect. 2.4, p. 5, lines 108–109.

R1.16 — Line 106.

106. This is not “spatial clustering”, but clustering in the slope parameter space?

Response. The referee is correct that the count describes concentration in the slope distribution, not spatial clustering. We have renamed it the “near-maximum count”: the number of active interior points within 0.2° below the member’s own $S_{\text{max,int}}$. Because it is referenced to the diagnosed maximum, it can also be affected by incomplete inner convergence (R1.14). We retain it as a descriptive diagnostic.

Changes. Sect. 2.4 (p. 5, lines 109–110) and the Table 2 header and caption (p. 12) now use the corrected name and interpretation. The justification of ε by a residual convergence gap of about 0.01° has been removed. Its use in explaining stability is addressed in R1.20.

R1.17 — Lines 109–110.

109-110. The MAE and modified area fraction should be explained in separate sentences.

Response. We have separated the definitions of MAE_h and the modified-area fraction into two sentences.

Changes. Sect. 2.4, p. 5, lines 111–113.

R1.18 — Line 112.

112. “short-wave”. “small scale” would make more sense.

Response. We have replaced “short-wave” with “small-scale” in general descriptions of terrain structure throughout the manuscript. We retain wavelength-specific wording where the spectral analysis refers to wavelengths.

Changes. Sect. 2.4, p. 5, lines 113 and 117; Sect. 4.1, p. 8, line 170; Sect. 4.2, p. 9, lines 189 and 191; Sect. 5, p. 13, lines 216 and 219; Sect. 6, p. 14, line 241.

R1.19 — Lines 124–125.

*124-125. The suggestion here is that the timestep must be reduced when more smoothing is applied in D01, which is the opposite of what I'd expect (smoothing makes the model more stable and may permit use of a longer timestep). I think what you mean is that this tactic improves stability in *D02*. But doesn't this undermine the comparison of different treatments of D02 terrain, since the experiments' stability in D02 now depends on both the terrain in D02, and the differences in the driving model D01 for different experiments?*

Response. Our original wording conflated terrain treatment with time stepping. D01 terrain was fixed after one smth-desmth_special pass, with all terrain-treatment differences confined to D02. WPS120 was numerically unstable at the 18/6 s D01/D02 time steps used for SCO and passed at 12/4 s. This reduction responded to instability, not to additional D01 smoothing. With the D01/D02 time-step ratio fixed at 3, the reduction also changed the parent-domain integration.

The comparison therefore concerns the successful configurations. The offline terrain-modification diagnostics remain directly comparable, while forecast-error differences cannot be attributed to D02 terrain alone. The time-step settings are summarised in Fig. 3.

Changes. The Abstract (p. 1, lines 9–10) now refers to the successful SCO configuration. Sect. 3.1 (p. 6, lines 126–129) now distinguishes the fixed D01 terrain from changes in time stepping and states how D01 supplies the nested integrations. The time-step settings are summarised in the Fig. 3 caption (p. 9), and the Table 1 caption (p. 8) links the stability outcomes to these configurations. Sect. 3.2 (p. 6, lines 148–151) defines the comparison in terms of the successful configurations, and Sect. 4.3 (p. 10, lines 207–208) states that their error differences cannot be attributed to D02 terrain alone.

R1.20 — Lines 177–178.

177-178. Are the authors able to speculate why this is?

Response. The passing WPS120 and failing SCO24 integrations used different time steps (R1.19; Fig. 3). Their contrast therefore does not demonstrate a higher terrain-only stability threshold for WPS.

Terrain structure beyond the maximum slope may still influence numerical behaviour, since the two methods produce different slope distributions and second-order terrain structure. The present diagnostics describe those differences but do not identify their effect on stability. The original causal attribution to the density of near-maximum slopes is not established, and we have withdrawn it.

Changes. Sect. 3.1 (p. 6, lines 126–129) and the Fig. 3 caption (p. 9) clarify the integration settings. Sect. 4.2 (p. 8, lines 177–180; pp. 9–10, lines 200–202) reports the sampled passing and failing outcomes under those configurations and describes terrain structure as a possible influence that the experiments do not isolate. The density-based attribution in Sects. 4.2, 5 and 6 has been removed (see also p. 13, lines 214–217, and p. 14, lines 241–243).

R1.21 — Figure 5 caption.

Figure 5 caption - it is the solid line that denotes θ_{eff} in the plot.

Response. We have corrected the caption to state that the solid line marks θ_{eff} and the dashed line θ . We also label panel (c) as the angular difference $S_{\text{max,int}} - \theta_{\text{eff}}$. The caption and Sect. 4.1 explain how the positive values for SCO18–SCO22 distinguish nominal target attainment from convergence to the internal working bounds (R1.14).

Changes. Fig. 5 caption, p. 11; Sect. 4.1, p. 8, lines 173–175.

R1.22 — Line 179.

179. I struggle to make sense of this sentence - “surrounding these members”?

Response. We have replaced the ambiguous phrase with the names of the four members. The sentence now reads: “At the domain scale, the slope fields of SCO22, SCO24, WPS100 and WPS120 differ sharply.”

Changes. Sect. 4.2, p. 8, line 181.

R1.23 — Lines 242–245.

242-245. It strikes me that the retention of terrain complexity with SCO, in addition to slope - in other words small scale spectral content - is what differs from the smoothly varying terrain of many-pass WPS. Have the authors considered looking at the second derivative of the terrain instead of just the slope?

Response. We have clarified the second-order terrain diagnostic included in Table 2 of the reviewed manuscript. The quantity previously labelled “Laplacian RMS” is the RMS of the five-point second difference. On the uniform grid this difference equals Δx^2 times the discrete Laplacian. SCO22 and SCO24 retain values of 279.8 and 289.6 m, compared with 47.9 and 49.9 m for WPS120 and WPS100.

The existing W1+SCO24 experiment reduces the second-difference RMS through pre-smoothing, but the integration still fails. SCO24 and W1+SCO24 both used D01/D02 time steps of 18/6 s, so their comparison does not introduce a time-step change. It shows that the tested pre-smoothing was insufficient to restore stability under these settings. Because pre-smoothing changes several terrain properties together, it does not isolate the role of second-order structure or exclude its influence on stability.

Changes. Sect. 2.4 (p. 5, lines 113–117) and Table 2 (p. 12) now name the measure the “RMS of the unnormalised five-point second difference”. Sects. 4.1 (p. 8, lines 170–171) and 4.2 (p. 9, lines 191–195) report this diagnostic, and Sect. 4.2 describes the pre-smoothing comparison, with the shared time steps listed in Fig. 3. Curvature constraints are mentioned as a possible extension in the Conclusions (p. 14, lines 246–247).

Response to Referee #2

We thank the referee for the positive assessment and for the constructive suggestions on the description of the method and its practical use. Our responses follow. Page and line numbers given under “Changes” refer to the revised manuscript without tracked changes.

R2.1 — Suggestion on WPS integration.

The authors describe a simple and practical approach to a long-standing problem in high-resolution WRF simulations over complex terrain. Steep slopes can cause numerical instability, while repeated domain-wide smoothing removes much of the terrain structure that motivates the use of a fine grid in the first place. SCO addresses this problem by prescribing a discrete-slope constraint and adjusting only the parts of the terrain that exceed it. This method should be useful to a broad WRF community. The authors might eventually consider packaging SCO as an optional part of the standard WPS terrain-processing workflow and exploring its contribution to the official WRF/WPS code base in the future.

Response. We thank the referee for this suggestion. Because SCO reads and writes standard geo_em files, it can already be used between geogrid and metgrid without code changes. Packaging it as an optional WPS component is a direction we would like to pursue in future releases. We now mention this possibility in the Conclusions.

Changes. Conclusions, p. 14, lines 246–247.

R2.2 — General comment #1.

#1 The manuscript describes θ as the “sole control variable” or as a single parameter. From a user’s perspective, θ is clearly the main attraction of the method, but the solver also contains fixed settings for the boundary treatment, local tightening, and convergence tolerances. It would be more accurate to describe θ as the primary user-facing parameter, or as the only parameter varied in this study, with the remaining quantities identified as fixed solver settings.

Response. We agree that θ is the main user-facing parameter, while the solver also contains fixed settings. We now describe θ as the “primary user-facing parameter” in the Abstract, Introduction and Conclusions. Sect. 2.3 gives the fixed values used for the boundary treatment (r , t), local tightening (β , $f_{\min}$) and iteration tolerances, and Sect. 3.2 states that only θ varies among the SCO members.

Changes. Abstract, p. 1, lines 5–6; Introduction, p. 2, line 31; Sect. 2.3, p. 3, lines 75–77, and p. 5, lines 97–106; Sect. 3.2, p. 6, lines 144–145; Conclusions, p. 14, line 236.

R2.3 — General comment #2.

#2 SCO22 and WPS120 are the least-modified stable members among the discrete candidates tested, rather than exact stability limits or globally optimal solutions. The wording should reflect this. It would also help to state more directly that the comparison is workflow-based: how much terrain modification does each preprocessing pathway require to obtain a tested member that satisfies the same 48 h stability criterion?

Response. We agree that the selection should be qualified by the candidates tested. SCO22 and WPS120 are the least-modified passing members among the discrete candidates tested under their reported configurations, not exact stability limits or globally optimal terrain solutions. Sect. 3.2 now makes this workflow-based comparison explicit. The configuration differences are described in R1.19 and Fig. 3. Sect. 4.3 states that their forecast-error differences cannot be attributed to D02 terrain alone.

Changes. Sect. 3.2, p. 6, lines 148–151; Fig. 3 caption, p. 9; Sect. 4.3, p. 10, lines 207–208.

R2.4 — General comment #3.

#3 The paper already contains most of the elements of a useful practical workflow, but they are spread across several sections. A short summary in the Discussion could describe how a user would diagnose the raw terrain, generate a coarse sequence of θ values, bracket the passing–failing interval with stability integrations, refine the interval if needed, and select the least-modified stable tested member.

Response. We agree that a concise workflow would make the method easier to use. We have added a short paragraph to the Discussion describing the practical sequence. A user first diagnoses the raw terrain and generates a coarse sequence of θ values. The resulting fields are checked offline for $S_{\max, \text{int}}$, internal constraint residuals and terrain modification. Stability integrations with a fixed model configuration and time step then bracket the sampled passing–failing interval, which can be refined if needed. The user selects the least-modified passing member among those tested. This selection applies to that configuration and should be reassessed when the time step or other model settings change.

Changes. Discussion, p. 14, lines 225–229.

R2.5 — General comment #4.

#4 The conclusions should remain tied to the case-study evidence. The results show that θ is a useful construction and search coordinate within SCO, but not that one slope value provides a universal WRF stability threshold across preprocessing methods, domains, or configurations.

Response. We agree that the conclusions should remain tied to the evidence from this case. We have revised them to describe θ as a parameter for generating terrain candidates and organising stability tests. They retain the findings of stable integration with substantially less terrain modification, greater small-scale terrain retention and lower near-surface errors for the successful configurations. The sampled passing range is explicitly tied to the tested case and model configuration; it is not presented as a universal WRF slope threshold.

Changes. Conclusions, p. 14, lines 236–249.

R2.6 — General comment #5.

#5 Add a concise practical workflow for users. One of the strengths of SCO is that candidate terrain fields can be generated rapidly and organised using a directly interpretable target angle. The discussion would benefit from a short recommended workflow for practical use.

Response. We have added the requested practical workflow to the Discussion as one concise paragraph, as described in our response to comment #3. The Discussion and Conclusions also explain how rapid terrain generation and an interpretable target angle help users organise candidates and select a passing tested member.

Changes. Discussion, pp. 13–14, lines 221–229; Conclusions, p. 14, lines 245–246.

Specific comments

R2.7 — Lines 141 and 159.

Line 141 and 159. “The least-modified stable member” should be qualified as referring to the candidates tested.

Response. We have qualified both instances to specify the candidates tested. Sect. 3.2 defines the selection under the reported configurations, and Sect. 4.1 identifies the selected members.

Changes. Abstract, p. 1, line 7; Sect. 3.2, p. 6, lines 148–149; Sect. 4.1, p. 7, lines 167–168.

R2.8 — Line 170.

Line 170. “Every unstable member terminates” makes the terrain member sound like the integrating system. The WRF integrations using those terrain members are what terminate.

Response. We have made the WRF integration the subject: “The WRF integrations with the unstable members terminated prematurely”.

Changes. Sect. 4.2, p. 8, line 177.

R2.9 — Line 175.

Line 175. “The two marginal stable members” is difficult to interpret. These are the stable members nearest the sampled stability boundaries.

Response. We have replaced “the two marginal stable members” with “the selected stable members, SCO22 and WPS120”, following the selection defined in Sect. 3.2. Their respective settings are identified in the Fig. 3 caption. This avoids implying that both identify exact or common fixed-step stability boundaries.

Changes. Sect. 4.2, p. 8, line 180.

R2.10 — Line 205.

Line 205. “Statistically indistinguishable” seems stronger than warranted because the paired test does not account for spatial correlation among stations. A statement about no detectable difference under the approximate stationwise test would be more consistent with Sect. 3.3.

Response. The referee is correct that “statistically indistinguishable” was too strong. We have revised the sentence to read “SCO20 and SCO22 show no detectable difference in wind and temperature RMSE under the approximate stationwise paired test (Sect. 3.3)” and made the Table 3 caption consistent. The test does not account for spatial correlation among stations, as stated in Sect. 3.3, and we do not interpret the result as demonstrating equivalence. The original sample and test are retained.

Changes. Sect. 4.3, p. 11, lines 210–211; Table 3 caption, p. 13.

R2.11 — Figure 5 caption.

Figure 5: Correct the line-style description in the caption.

Response. We have corrected the caption to state that the solid line marks θ_{eff} and the dashed line θ .

Changes. Fig. 5 caption, p. 11.

R2.12 — Abbreviations.

A final check of abbreviations would be useful, some abbreviations are used before they are defined.

Response. We have checked the abbreviations and added definitions at first use where they were missing, including WRF-ARW, ERA5, ASTER, Noah-MP, CPU, HPC and LP DAAC.

Changes. p. 2, lines 33–34 (WRF-ARW); p. 6, lines 130–133 (ERA5, ASTER) and 137–138 (Noah-MP); p. 13, line 223 (CPU); p. 17, lines 277 (LP DAAC) and 291 (HPC).

References

Sheridan, P., Xu, A., Li, J., and Furtado, K.: Use of targeted orographic smoothing in very high resolution simulations of a downslope windstorm and rotor in a sub-tropical highland location, *Adv. Atmos. Sci.*, 40, 2043–2062, doi:10.1007/s00376-023-2298-0, 2023.