Author Response — egusphere-2026-415 (Round 1)

Dear Editor,

We thank the editor and reviewers for their time and constructive feedback on our manuscript "*Radar Data Smoothing using the Discrete Cosine Transform: A Fast Spectral Domain Algorithm*". In this first revision, we have addressed all 12 reviewer comments (6 from each reviewer) through a combination of expanded discussion, additional references, and technical clarifications.

Reviewer comments are provided in black text, and our responses are marked in blue text. The editor and reviewers can map to PDF page numbers via the accompanying `latexdiff` PDF.

Reviewer 1

Recommendation:

Publication after Minor Revisions.

General Comments:

This manuscript presents a computationally efficient and theoretically sound approach for smoothing radar data in the native polar coordinate system using the Discrete Cosine Transform (DCT). The authors successfully address the geometric inconsistency of fixed-kernel spatial convolution and demonstrate massive performance gains. The topic is highly relevant to operational radar meteorology, and the manuscript is well-written and technically rigorous, mainly in the following aspects.

1. The paper effectively highlights the limitations of spatial-domain smoothing in polar coordinates, specifically the trade-off between physical consistency and computational cost when dealing with range-dependent beam broadening.
2. The derivation of analytical transfer functions for continuous kernel widths is a key strength. It elegantly bypasses the discretization errors of pixel-based indexing. Furthermore, leveraging the DCT-II's symmetric boundary conditions to mitigate the Gibbs phenomenon is a superior choice compared to standard FFT-based methods.
3. The performance benchmarks are thorough. The reported speedup factors and the reduction of processing time from hours to seconds for large datasets provide strong evidence for the practical utility of this method in real-time operational pipelines.

Answer: We thank the reviewer for the positive assessment. The points raised here frame the contributions the Specific Comments below ask us to expand on, and we have addressed each in turn.

Specific Comment 1:

Handling of Missing Data (NaNs): The requirement for a continuous data field (Lines 38-40, 161-166) is a significant limitation of spectral methods. The current solution relies on linear interpolation for gap-filling. Please expand the discussion on the impact of this choice. How does linear interpolation affect the spectral content in regions with extensive data loss (e.g., beam blockage or ground clutter removal)? Could this introduce artificial gradients or bias the DCT smoothing? A brief comparison with more advanced filling techniques (e.g., convolution-based interpolation mentioned in Line 166) would strengthen the manuscript.

Answer: We thank the reviewer for raising this point. The prefilling requirement of spectral transforms is a known limitation of DCT/FFT pipelines, and we have revised the manuscript at three locations to address it more explicitly.

1. Introduction, prefilling-requirement sentence.

Changed from: (at Introduction)

> *Therefore, data gaps due to filtering or beam blockage must be filled (e.g., via interpolation) prior to applying the transform. While this is a limitation, we discuss below how simple interpolation strategies are often sufficient for effective noise removal.*

To:

> Therefore, data gaps due to filtering or beam blockage must be filled (e.g., via interpolation) prior to applying the transform. This prefilling requirement is a known limitation of spectral methods; its consequences and the alternative strategies that address them are discussed in Section 3.

2. Sec. 3 (Application to Real Radar Observations), the linear-interpolation caveat paragraph (two sentences replaced).

Changed from: (at Sec. 3 (Application to Real Radar Observations))

> *One caveat of the DCT method is that it requires a continuous data field without any missing data. Radar observations, however, often contain NaN values where no signal was detected or where artifacts were removed. In this study, we employed linear interpolation to fill gaps prior to the DCT operation. While effective for small gaps, this method can introduce artifacts in regions with extensive data loss. Convolution-based interpolation is a powerful alternative, as implemented in the Astropy library.*

To:

> One caveat of the DCT method is that it requires a continuous data field without missing entries, and the present work fills gaps with linear interpolation prior to the DCT. Linear interpolation across a gap replaces the missing segment with a piecewise-linear bridge whose slope is unrelated to the underlying field, so two effects can bias the smoothed result. First, the bridge contributes a low-frequency trend across the gap; for gaps narrower than the local kernel width W this trend is small relative to the smoothing noise floor, but for wide gaps (large beam-blockage sectors, broad clutter removals) it can dominate the smoothed values within and near the gap. Second, the slope of the bridge generally differs from the local slope of the underlying field, producing a derivative discontinuity (a kink) at each gap boundary. The kink injects localized high-frequency content into the DCT spectrum; for W larger than the gap width this content is removed by the smoothing filter, but for narrow kernels residual kinks can survive as faint stripes parallel to the gap edge. The Astropy convolution-based interpolation cited above is a robust drop-in upgrade that addresses both the low-frequency bias and the kink. More generally, support-normalized filtering, which avoids prefilling by treating missing entries as a separate certainty mask, has been developed by the authors as a methodological extension of the present framework.

3. Conclusions, prefilling-cost sentence.

Changed from: (at Conclusions)

> *While the method requires a continuous data field, necessitating a preprocessing step to interpolate gaps (e.g., removed clutter or beam blockage), the benefits of physically correct, artifact-free smoothing far outweigh this limitation.*

To:

> While the method requires a continuous data field, necessitating a preprocessing step to interpolate gaps (e.g., removed clutter or beam blockage), the resulting low-frequency bias and gap-edge kinks are bounded for gap widths smaller than the local kernel and can be eliminated by the support-normalized extension developed by the authors.

We have also added three supporting references to the bibliography (Knutsson and Westin, 1993; Westin et al., 1994; Pham et al., 2006) and a forward citation to the companion preprint (Valdivia-Prado et al., 2026, arXiv:2604.22721) for the support-normalized extension.

Specific Comment 2:

The method treats azimuthal and range smoothing as independent processes (Eqs. 13-17). While computationally convenient, please justify why this decoupling is physically appropriate for meteorological fields. In cases of strong radial gradients (e.g., convective storm inflow), does the lack of 2D coupling in the kernel design lead to anisotropic smoothing?

Answer: We thank the reviewer for raising this point. The separable form of the polar operator is a property of the kernel rather than an approximation of the data, and we have added a brief clarification to Sec. 2.3 to make this explicit.

Inserted in Sec. 2.3 (2D Convolution in Polar Coordinates), paragraph inserted after the two-stage DCT pipeline:

> The separable form is exact for a kernel that is itself separable: a 2D boxcar h[m,n] = h_row[m] h_col[n] is mathematically identical to two sequential 1D boxcars, not an approximation of one. In polar geometry the resulting effective kernel is intentionally axis-aligned in (range, azimuth) to match the polar sampling volume (Bringi2001), so the anisotropy of the smoothing is the anisotropy of the radar geometry, not an artifact of the method.

For a field that varies primarily with range (e.g., a strong radial gradient in convective inflow), the azimuth filter is applied along lines of constant r where the field is approximately uniform, so the effective smoothing scale for radial gradients is set by the range kernel W_range alone; the larger azimuth width at long range has little effect on radial-gradient smoothing (Doviak1979). For features oriented diagonally in (range, azimuth) space, the axis-aligned kernel produces a slight directional asymmetry in smoothing that is shared with non-separable axis-aligned 2D boxcars and is in line with standard 2D smoothing practice. Within the scope of the present paper, the separable polar kernel matches the intended physical footprint at each range, and the anisotropy of the effective kernel is the anisotropy of the radar sampling volume rather than an artifact of the smoothing method.

Specific Comment 3:

Figure Resolution: Some figures (e.g., Fig. 3) appear somewhat compressed in the text. Higher resolution versions would aid in evaluating the subtle differences between the spatial and spectral methods.

Answer: We have generated all figures as vector PDFs so they display at the best quality the source data allows. Figures 2, 3, and 4 contain embedded raster for their data plot areas (axis frames, tick labels, and titles remain as vector paths); the resolution of the embedded raster should be sufficient at the column-width print size. We would kindly ask the reviewer to verify that the compressed appearance is not a rendering artifact of their PDF viewer rather than a property of the source file. If the issue persists we can increase the embedded raster resolution further, but we do not believe this is the most likely cause.

Specific Comment 4:

Variable Consistency: Ensure consistent notation throughout. For example, W is used for window size in pixels (Eq. 7) and W_phys for physical width (Eq. 11). Clarifying the distinction between "pixel width" and "physical width" early in Section 2.3 would improve readability.

Answer: We agree that the relationship between the pixel and physical widths was not made explicit at the point of first use, and we have added a short clarifying paragraph at the start of Sec. 2.3 to remove the ambiguity. As part of a broader Sec. 2.3 flow improvement, the new paragraph was further trimmed and integrated with the immediately following kernel-implementation paragraph.

Inserted in Sec. 2.3 (2D Convolution in Polar Coordinates), paragraph inserted at the opening of the section (after the angle/range equivalence statement added in response to R2-2, before Eq. 11):

> We adopt three related kernel-width quantities to describe this range dependence. Let W_phys denote the physical smoothing width in units of length (e.g., meters or kilometers), which is held constant in the azimuth direction at every range. The equivalent width in pixels is W = W_phys / (r Δθ), a continuous real-valued parameter identical to the W of Eq. 7. In the range dimension the resolution is constant, so a fixed kernel is used. In the azimuth dimension the arc length Δs = r Δθ grows with range, so the discrete kernel applied at each row is N_az(r), the odd integer obtained by symmetrically rounding W to keep the kernel centered on the target cell: [Eq. 11]. The factor 1/r makes W decrease with range even though W_phys is constant, so N_az(r) is small near the radar and grows toward long range.

The clarifying paragraph does not introduce a new symbol; it simply states the conversion W = W_phys/(rΔθ) at the point of first use, and binds the existing W of Eq. 7 to the W_phys of Eq. 11. All subsequent uses of W and W_phys in the manuscript follow this convention without further changes. No new citations are needed.

Specific Comment 5:

Application to Velocity (Vr): The discussion regarding the necessity of de-aliasing prior to smoothing (Lines 229-234) is excellent. Please ensure this warning is emphasized in the abstract or conclusion, as applying this method to folded velocity fields would yield catastrophic results.

Answer: We agree that the de-aliasing prerequisite deserves more visibility, and we have added an explicit one-sentence warning to the Conclusions section as a parallel to the existing per-variable caveats.

Inserted in Conclusions, one sentence inserted between the ρ_hv caveat and the "Overall" summary sentence:

> For Doppler velocity, de-aliasing is a strict prerequisite: the wrapped-velocity fold is an artificial discontinuity that the spectral transform will treat as a sharp physical edge and propagate as ringing.

The new sentence is placed in the Conclusions section because the per-variable caveats (NaN prefilling, ρ_hv convex-hull preservation, and now velocity de-aliasing) belong together as a single block describing the conditions of applicability. No new citations are needed.

Specific Comment 6:

Discussion and Future Work: The introduction mentions that SNR degrades with range (Line 208). Currently, the kernel scales geometrically (based on beam width) but not statistically (based on local SNR). A brief mention of whether the spectral framework could support adaptive smoothing (varying the transfer function based on local variance) would be a valuable addition to the future work section.

Answer: We agree that the present method is restricted to geometric range dependence, and that statistical adaptation to local SNR is a natural extension. Progress in this direction is already reported in the local-statistics extension of the present framework (Valdivia-Prado2026), and we have added a forward-looking note in the Discussion section to make this scope explicit.

Inserted in Sec. 4 (Discussion), two sentences appended to the existing adaptive-smoothing paragraph:

> Within the present spectral framework, statistical adaptation is a natural direction: the analytical transfer function H[k] is an explicit object that can be modulated on a per-row basis to widen or narrow the effective passband in regions of low or high local variance, respectively. Progress in this direction is already reported in the local-statistics extension of the present framework (Valdivia-Prado2026), which extends the present method to local statistics in polar coordinates through support-normalized filtering.

Closing:

Finally, this paper makes a significant contribution to the field of radar data quality control. The proposed DCT-based smoothing algorithm is elegant, fast, and addresses a real-world operational bottleneck. I look forward to seeing the revised version.

Answer: We thank the reviewer for the positive assessment and for the careful and constructive review throughout this process.

Reviewer 2

General Comments:

This work presents a technique for smoothing real-valued radar data in polar coordinates using a Discrete Cosine Transform. Not only does the DCT offer improved computational efficiency because convolution becomes multiplication in the frequency domain, but it also enables more efficient use of smoothing kernels that vary as a function of range. This is particularly useful for radar data where the beam widens with increased range and it is desirable to smooth over a wider window with increased range. The work demonstrates the improved efficiency of the DCT method over convolution in the spatial domain, and it demonstrates that the DCT technique produces an equivalent smoothed result on example C-band weather radar data. Overall this is a useful application of frequency domain filtering techniques which is relevant to multiple forms of radar data, and I recommend publication with minor revisions.

Answer: We thank the reviewer for the constructive comments.

Specific Comment 1:

Lines 103-106 (discrete vs continuous kernel width): Why do you use a discrete width here instead of the continuous-width transfer function shown in equation (7)? Is it just for comparison with the existing dynamic kernel technique for Figure 2?

Answer: Yes, the discrete width M=15 is used in the synthetic experiment to match the integer kernel width the spatial dynamic kernel method also requires, making Fig. 2 a direct one-to-one check of equivalence. The continuous-width analytical formulation (Eq. 7) is exercised separately in the real-data application (Fig. 4), where its fractional-width capability is the operationally relevant feature. This is the design intent: the synthetic experiment tests equivalence against an integer-width reference, while the real-data experiment exercises the new continuous-width capability that the paper contributes.

Specific Comment 2:

Line 109 (row/column vs azimuth/range terminology): Here and throughout this section, might I suggest not introducing the row/column terminology and instead just use azimuth/range? Or perhaps angle/range, to allow for the angle to be azimuth or elevation depending on the case. Particularly in the equations where the row operations are paired with the H_az transfer function, and the column operations paired with the H_range transfer function, I find the using both terminologies simultaneously to be a little confusing. Nevertheless, the meaning is clear in the end, so I leave this suggestion to the discretion of the author.

Answer: We thank the reviewer for the suggestion. We have standardized the notation on "angle" for the transverse polar coordinate and "range" for the radial polar coordinate throughout the paper, so that the analysis applies uniformly to PPI scans (where the angle is the azimuth) and to RHI scans (where the angle is the elevation). All math subscripts that previously used row/col/az have been renamed to ang/rng/ang, the row-index variable R has been renamed to α, and an equivalence statement has been added at the opening of Sec. 2.3 to make the new convention explicit.

Inserted in Sec. 2.3 (2D Convolution in Polar Coordinates), one-sentence equivalence statement inserted at the opening of the section, between the opening paragraph and the first use of "angle":

> Throughout this section, "angle" refers to the polar angular coordinate θ, which is the azimuth for PPI scans and the elevation for RHI scans; the analysis applies to either case.

The mathematical subscripts and prose in Sec. 2.3 have been updated in lockstep (DCT_row → DCT_ang, DCT_col → DCT_rng, H_az → H_ang, N_az → N_ang, h_row → h_ang, h_col → h_rng, and the intermediate variables X̃_row, Ỹ_col, Y_az → X̃_ang, Ỹ_rng, Y_ang; R → α for the row-index variable). The same convention is applied to H_az, N_az, and the row/column terminology wherever they appear elsewhere in the paper.

Specific Comment 3:

Line 262 (Gibbs phenomenon wording): "the Gibbs phenomenon associated with Fourier methods" can be read a bit too strongly to imply that the Gibbs phenomenon is unavoidable with Fourier techniques, where actually it is a hazard of jump discontinuities when summing finite Fourier series. Since the DCT uses a Fourier series, one can get Gibbs ringing with it as well. It is only avoided here because of the DCT's advantageous boundary conditions and by the use of kernels that avoid jump discontinuities. I would suggest rewording here to specify the particular Fourier method you want to draw a comparison to (i.e. the Ideal Low-Pass Filter mentioned in Section 5) or at least qualify it to "some Fourier methods".

Answer: We agree that the original phrasing was too broad. We have replaced "Fourier methods" in the Conclusions with the specific method being compared against, "the Ideal Low-Pass Filter discussed in Section 5", which is the hard-cutoff low-pass filter whose step discontinuity in the frequency domain produces Gibbs ringing. The cross-reference points the reader to the detailed treatment in Section 5, which already names the kernel-side mechanism (smooth shaping function, no jump discontinuity) that prevents the ringing.

Changed from: (at Conclusions)

> *the Gibbs phenomenon associated with Fourier methods*

To:

> the Gibbs phenomenon associated with the Ideal Low-Pass Filter discussed in Section 5

Specific Comment 4:

DCT kernel discussion: Different DCT kernels are presented and there is some discussion in section 2.4 about the strengths and weaknesses of each, but I think a more robust discussion of these is appropriate. It would be beneficial also to recommend a particular kernel for general use of this technique, or recommend different kernels for different situations. One particular aspect of the comparison between the given kernels is how some only offer discrete window sizes that can be varied with range while others can be continuously varied. Would one ever want to choose one of the discrete window kernels over the continuous ones? Are there other continuous kernels not mentioned that would be worth considering?

Answer: We agree that the original Sec. 2.4 listed the available kernels without a practical recommendation. We have added the recommendation at the end of the Discussion section, where it fits better with the other forward-looking content, with a one-sentence pointer in Sec. 2.4 cross-referencing it.

Inserted in Sec. 4 (Discussion), one paragraph added after the DCT vs FFT paragraph:

> For general-purpose range-dependent smoothing of polar radar data we recommend the analytical continuous Gaussian kernel (Eq. 18) as the default: it provides the smoothest taper (no spectral sidelobes, no ringing), it allows the fractional pixel widths required by the 1/r angular geometry, and its performance is superior. The analytical continuous boxcar (Eq. 7) is the appropriate choice when a hard cutoff kernel is desired, while the discrete boxcar (Eq. 4) should be reserved for cases that require an integer kernel width matching a spatial reference. The Savitzky-Golay kernel (Eq. 19) preserves low-order polynomial features such as the peak and area of reflectivity cores and is worth exploiting further in the present framework; its coefficients depend only on the window width W and the polynomial order p (not on the data values), are pre-computed once from the polynomial basis, and are applied as a fixed-weight FIR filter with the same O(W) per-pixel cost as any other kernel in this paper. Window families that are common in other application domains, such as the Hamming, Blackman, Tukey, Kaiser, and Welch windows (Blackman1959, Harris1978, Kaiser1966, Welch1967), are not the focus of the present study.

Specific Comment 5:

Kernel performance on C-band data: Along the same lines as the last comment, it would be illuminating to see how the different kernels perform on the C-band radar data shown in Figure 4, or at least whatever the "winning" kernel is as identified from prior discussion. The DCT boxcar shows an equivalent result to the existing Dynamic Kernel method at much less computational cost, but does another DCT kernel give better results?

Answer: We agree that comparing the four DCT-compatible kernels on the C-band data would be informative, but a quantitative comparison on real data is not feasible: the C-band observations have no ground truth, and the existing Fig. 4 uses the boxcar kernel as the reference for the equivalence check with the spatial dynamic kernel, which is not itself a benchmark. On the synthetic data where ground truth is available (Fig. 2), the continuous Gaussian kernel was able to reduce noise better than the boxcar under equivalent kernel size in our internal tests; those results are not shown in the manuscript, but the recommendation in the Discussion section to use the continuous Gaussian as the default (added in response to a related comment) is consistent with this finding. No manuscript change is required.

Technical Correction:

Line 266: The given Zenodo code link is nonexistent and should read:
https://doi.org/10.5281/zenodo.18226677

Answer: We thank the reviewer for catching this. The Zenodo DOI in the code-availability statement has been corrected to the correct identifier for the archived radar-dct-smoothing repository.

Changed from: (at Code availability section)

https://doi.org/10.5281/zenodo.10822667

To:

https://doi.org/10.5281/zenodo.18226677