

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. Therefore, I recommend publication after Minor Revisions.

Specific Comments and Suggestions

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.

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?

3. 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.

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.

5. Application to Velocity (V_r): 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.

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.

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.