

Radar Data Smoothing using the Discrete Cosine Transform: A Fast Spectral Domain Algorithm

Jairo M. Valdivia¹, Will Chapman¹, and Katja Friedrich¹

¹Department of Atmospheric and Oceanic Sciences, University of Colorado Boulder, Boulder, CO, USA

Correspondence: Jairo M. Valdivia (jairo.valdiviaprado@colorado.edu)

Abstract. This study introduces a computationally efficient and methodologically robust approach for smoothing radar data using the Discrete Cosine Transform (DCT). Traditional spatial convolution methods for noise reduction in polar coordinates suffer from geometric inconsistencies and prohibitive computational costs, particularly when implementing range-dependent dynamic kernels to maintain physical scale. We propose a spectral-domain alternative that utilizes the convolution theorem to perform equivalent smoothing operations. By deriving analytical transfer functions for various kernels—including Boxcar, Gaussian, and Savitzky-Golay—we demonstrate that the DCT method achieves identical performance to spatial convolution while effectively handling boundary conditions. Performance benchmarks on real C-band weather radar data reveal that the DCT-based approach offers speedup factors exceeding 800 times for large kernel sizes. Furthermore, for large-scale datasets (180 million pixels), equivalent processing time is reduced from over 1 hour to under 18 seconds. The proposed method ensures physically consistent smoothing across ranges, preserving small-scale meteorological features while enabling real-time data quality improvement.

Copyright statement. TEXT

1 Introduction

Weather radar observations often show large spatial variability, which is typically exacerbated by various sources of noise and observational artifacts. Some gradients are related to the nature of precipitation, others result from oversampling, and non-meteorological features such as partial or complete beam blockage (Wilson et al., 2012; Atlas, 1990), data gaps related to data filtering, non-uniform beam filling around cloud edges (Bringi and Chandrasekar, 2001), and range-dependent beam broadening (e.g. Ivić et al., 2013). While significant progress has been made in removing structured non-meteorological echoes (e.g., ground clutter filtering), high-frequency fluctuations in the signal remains a computational expensive process, particularly for improving data quality for visualization and automated analysis.

Traditional spatial domain smoothing techniques, such as the moving window average (boxcar filter; Smith, 1997), have been widely used to mitigate this "noise". However, applying these filters in the native polar coordinate system of radar data presents geometric challenges. A constant window size in the range-angle domain corresponds to a physically varying spatial




Figure 1. Conceptual comparison of coordinate distortion for (a) a window of constant size in grid space (range-angle indices) that corresponds to (b) physically different areas in a physical space depending on the range. At close range, the physical area is small, whereas at far range, it expands significantly. This distortion necessitates dynamic kernel sizing to maintain a constant physical smoothing scale.

scale; a small window at close range covers a much smaller physical area than the same window at long range. This leads to inconsistent smoothing: data at long ranges may be undersmoothed while data at close ranges may be oversmoothed, or vice-versa. To visualize this issue, Fig. 1 illustrates how a constant angular window corresponds to drastically different physical areas at different ranges, highlighting the need for range-dependent kernel adaptation which is computationally expensive in the spatial domain.

In this work, we propose a spectral-domain approach using the Discrete Cosine Transform (DCT; Ahmed et al., 1974) to perform radar data smoothing in the native polar coordinate space. This smoothing technique help us to filter out high-frequency oscillations in the signal that is treated as "noise". We demonstrate that DCT-based convolution offers a computationally efficient and mathematically equivalent alternative to spatial convolution, with the added benefit of easily handling range-dependent kernel scaling and boundary conditions that mitigate artifacts like the Gibbs phenomenon. Specifically, the DCT method allows us to perform convolution in $O(N \log N)$ time, which is significantly faster than the $O(N \cdot W^2)$ complexity of spatial convolution with large windows. This speed advantage allows for real-time processing of large radar volumes with physically correct range-variant smoothing, which was previously computationally prohibitive. DCT method requires a complete, continuous dataset. Unlike spatial convolution which can simply skip invalid pixels (NaNs) or adjust normalization, spectral transforms technically require defined values everywhere. 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.

This paper is organized as follows: Section 2 establishes the theoretical framework, detailing the convolution theorem and its application in polar coordinates. Section 3 demonstrates the application of this method to real radar observations. Section 4 presents a performance benchmark comparing the proposed DCT method against traditional spatial algorithms. Finally, Section 5 discusses the implications and Section 6 concludes the study.

2 Theoretical Framework and Geometric Considerations

2.1 Convolution Theorem and Transfer Function

The fundamental mathematical principle underlying this work is the Convolution Theorem, which states that convolution in the spatial domain is equivalent to point-wise multiplication in the spectral domain (Oppenheim et al., 1999). For two discrete signals $f[n]$ and $g[n]$, their linear convolution $(f * g)[n]$ corresponds to:

$$\mathcal{T}\{f * g\} = \mathcal{T}\{f\} \cdot \mathcal{T}\{g\} \quad (1)$$

where $\mathcal{T}$ denotes a linear transform. While this is commonly associated with the Discrete Fourier Transform (DFT) under periodic boundary conditions (circular convolution), it also applies to the Discrete Cosine Transform (DCT) under the assumption of symmetric boundary conditions (Makhoul, 1980). We utilize the DCT-II, defined for a signal $x[n]$ of length N as:

$$X[k] = 2 \sum_{n=0}^{N-1} x[n] \cos\left(\frac{\pi k(2n+1)}{2N}\right), \quad 0 \leq k < N \quad (2)$$

The DCT-II corresponds to a discretization of the Neumann boundary condition (zero derivative at the endpoints) in the continuum limit. In the discrete domain, this implies a *half-sample symmetric extension* of the data at the boundaries (Strang, 1999).

Ng et al. (1999) demonstrated that while the DFT diagonalizes circulant matrices (periodic convolution), the DCT-II diagonalizes (or approximately diagonalizes, depending on the kernel symmetry) matrices with Toeplitz-plus-Hankel structure. This structure mathematically represents the sum of a standard linear convolution (Toeplitz) and the convolution with the reflected data at the boundaries (Hankel). Therefore, point-wise multiplication in the DCT domain is equivalent to spatial convolution where data outside the domain is modeled as a symmetric reflection of the data inside, effectively mitigating the edge artifacts associated with periodic DFT assumptions.

For a symmetric filter $h[n]$ of width M (where M is odd), centered at $n = 0$, the coefficients are $h[n] = \frac{1}{M}$ for $|n| \leq \frac{M-1}{2}$ and 0 otherwise. The frequency response (or transfer function) $H[k]$ in the DCT domain corresponds to the DFT of the symmetric extension of the filter. Since $h[n]$ is real and even, its frequency response is given by the sum of cosine terms:

$$H[k] = \sum_{n=-(M-1)/2}^{(M-1)/2} h[n] \cos\left(\frac{\pi kn}{N}\right) \quad (3)$$

Substituting $h[n] = \frac{1}{M}$:

$$H[k] = \frac{1}{M} \left(\cos(0) + \sum_{n=1}^{(M-1)/2} \cos\left(\frac{\pi kn}{N}\right) + \sum_{n=-(M-1)/2}^{-1} \cos\left(\frac{\pi kn}{N}\right) \right) \quad (4)$$

70 Since $\cos(-x) = \cos(x)$, the negative indices sum is identical to the positive indices sum. Therefore, for a boxcar filter of width M in the DCT-II domain, the transfer function $H[k]$ is given by:

$$H[k] = \frac{1}{M} \left(1 + 2 \sum_{j=1}^{(M-1)/2} \cos\left(\frac{j\pi k}{N}\right) \right) \quad (5)$$

To generalize this for continuous widths W (essential for polar coordinate smoothing where the effective kernel width varies continuously with range), we utilize the Dirichlet kernel identity (Levi, 1974). This extension effectively defines a continuous

75 spectral filter that approximates a fractional-width moving average:

$$1 + 2 \sum_{j=1}^K \cos(j\theta) = \frac{\sin((K+1/2)\theta)}{\sin(\theta/2)} \quad (6)$$

This is a standard result in Fourier analysis (Oppenheim et al., 1999). Substituting $K = (W-1)/2$ and $\theta = \omega_k = \frac{\pi k}{N}$, we obtain the analytical transfer function:

$$H[k] = \frac{1}{W} \frac{\sin(W\omega_k/2)}{\sin(\omega_k/2)} \quad (7)$$

80 Note that $\lim_{\omega_k \rightarrow 0} H[k] = 1$. This closed-form expression allows us to compute the smoothing operation efficiently for any real-valued width W , mitigating discretization errors associated with forcing integer kernel sizes. It is important to emphasize that while we refer to this as an "analytical boxcar convolution" due to its spatial equivalent, in the DCT domain it is strictly a spectral filter. This distinction allows us to bypass the constraints of discrete grid indices entirely, operating on the continuous frequency response of the underlying function rather than a pixelated approximation.

85 2.2 2D Convolution and Separability

The 2D spatial convolution of a matrix $x[m, n]$ with a kernel $h[m, n]$ is defined as:

$$y[m, n] = (x * h)[m, n] = \sum_i \sum_j x[m-i, n-j] h[i, j] \quad (8)$$

In our experiments, we use a boxcar filter of size $M \times M$. A key property of this filter is its separability, meaning it can be decomposed into two 1D filters:

$$90 \quad h[m, n] = h_{ang}[m] \cdot h_{range}[n] \quad (9)$$

where h_{ang} and h_{range} are 1D boxcar filters. Similarly, the 2D DCT is separable:

$$X_{2D}[u, v] = \text{DCT}_{range}(\text{DCT}_{ang}(x[m, n])) \quad (10)$$

For a standard separable convolution with a constant kernel, the 2D transfer function is simply the outer product of the 1D transfer functions: $H_{2D}[u, v] = H_{ang}[u] \cdot H_{range}[v]$.




Figure 2. Comparison of synthetic ground truth, noisy observation, and smoothing performance. (a) Ground-truth synthetic Range-Height Indicator (RHI) radar image in polar coordinates. The field value mimics reflectivity (Z_e), but could be any dual-polarization variable. (b) Noisy observation ($\sigma = 2.00$). (c) The added noise pattern. (d) DCT-based smoothed result. (e) Spatial Dynamic Kernel (DK) smoothed result. (f) Difference between DCT and DK results. (g) DCT error map. (h) DK error map. (i) Difference of absolute errors ($|\text{Error}_{\text{DCT}}| - |\text{Error}_{\text{DK}}|$), where positive values indicate the DCT method performed worse and negative values indicate the DK method performed worse. Performance metrics include Root Mean Square Error (RMSE) and Mean Absolute Error (MAE).




Figure 3. Spectral response of different kernels for two window sizes. Window 5: (a)-(e) 2D angle transfer functions, and (f) Range transfer function. Window 15: (g)-(k) Angle transfer functions, and (l) Range transfer function.

95 2.3 2D Convolution in Polar Coordinates

Radar observations are collected on a polar grid (r, θ) through scanning strategies such as the Plane Position Indicator (PPI) for horizontal surveys or the Range Height Indicator (RHI) for vertical cross-sections. While these datasets are frequently projected onto Cartesian coordinates for integration with numerical models, performing smoothing directly in the native polar domain avoids interpolation artifacts but introduces geometric complexities. To ensure physically consistent smoothing, the effective kernel size must dynamically adapt to the range-dependent expansion of the radar sampling volume, accounting for

the linear increase in transverse resolution with distance from the sensor. 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.

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 angle direction at every range. The equivalent
105 width in pixels is $W = W_{\text{phys}}/(r\Delta\theta)$, 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 angle dimension the arc length $\Delta s = r\Delta\theta$ grows with range, so the discrete kernel applied at each angle index is $N_{\text{ang}}(r)$, the odd integer obtained by symmetrically rounding W to keep the kernel centered on the target cell:

$$N_{\text{ang}}(r) = 2 \left\lfloor \frac{W_{\text{phys}}}{2r\Delta\theta} \right\rfloor + 1 \quad (11)$$

110 The factor $1/r$ makes W decrease with range even though W_{phys} is constant, so $N_{\text{ang}}(r)$ is small near the radar and grows toward long range.

In the spectral domain this range dependence implies that the operation cannot be expressed as a single global 2D transfer function. Instead we treat it as a sequence of separable operations: for each angle α , angle-direction smoothing uses a range-specific 1D filter $h_{\text{ang}}[\alpha]$, which in the 1D DCT along the angle direction corresponds to multiplying by a range-dependent
115 transfer-function matrix $\mathbf{H}_{\text{ang}}[\alpha, v]$. Range smoothing uses a fixed filter, which in the 1D DCT along the range direction corresponds to multiplying by a constant transfer function $H_{\text{range}}[u]$.

The separable form is exact for a kernel that is itself separable: a 2D boxcar $h[m, n] = h_{\text{ang}}[m] \cdot h_{\text{range}}[n]$ 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, angle) to match the polar sampling volume (Brangi and Chandrasekar, 2001), so the anisotropy
120 of the smoothing is the anisotropy of the radar geometry, not an artifact of the method.

The full 2D smoothing operation Y can be decomposed into a sequence of separable steps. First, we perform the angle-direction smoothing, whose transfer function depends on range:

$$\tilde{X}_{\text{ang}} = \text{DCT}_{\text{ang}}(X) \quad (12)$$

$$\tilde{Y}_{\text{ang}} = \tilde{X}_{\text{ang}} \circ \mathbf{H}_{\text{ang}} \quad (13)$$

$$125 \quad Y_{\text{ang}} = \text{IDCT}_{\text{ang}}(\tilde{Y}_{\text{ang}}) \quad (14)$$

where $\circ$ denotes the Hadamard (element-wise) product and $\mathbf{H}_{\text{ang}}$ is the matrix of range-dependent transfer functions. Next, we apply the constant smoothing along the range direction:

$$\tilde{Y}_{\text{range}} = \text{DCT}_{\text{range}}(Y_{\text{ang}}) \quad (15)$$

$$\tilde{Y}_{\text{final}} = \tilde{Y}_{\text{range}} \cdot H_{\text{range}} \quad (16)$$

$$130 \quad Y = \text{IDCT}_{\text{range}}(\tilde{Y}_{\text{final}}) \quad (17)$$

This step-by-step formulation explicitly shows the intermediate return to the spatial domain (Y_{ang}) required to transition from the per-angle smoothing (where the transfer function varies along the range) to the per-range smoothing (where the transfer function is constant).

Figure 2 validates the equivalence between the spatial and spectral operations using a synthetic dataset with known geometry. Both, the DCT and dynamic kernel methods utilized a boxcar window of 15 pixels in the range direction (approx. 1.1 km for a 75 m range resolution). The detailed analysis in Fig. 2 shows the original ground truth (Fig. 2a) and the noisy signal (Fig. 2b) side-by-side with the denoised outputs (Fig. 2d, e). Crucially, the difference map (Fig. 2f) and error plots (Fig. 2g-i) confirm that the DCT method produces results that are nearly indistinguishable from the computationally expensive spatial convolution, solving the geometric smoothing problem without the performance penalty. The minor edge effects seen in the difference map are due to the difference in boundary conditions of the DCT (symmetric extension) versus the zero-padding or truncation often used in spatial implementations.

2.4 Alternative DCT Kernels: Gaussian, Savitzky-Golay, and Hanning

A key advantage of the spectral approach is the flexibility of the transfer function. We can instantly switch between different smoothing characteristics by changing the analytical transfer function $H[k]$:

For a Gaussian filter with standard deviation σ , the normalized transfer function is (Smith, 1997):

$$H_{Gaussian}[k] = e^{-\frac{1}{2}(\omega_k \sigma)^2} \quad (18)$$

where $\omega_k = \frac{\pi k}{N}$.

The Savitzky-Golay filter fits a polynomial of order p within a window W . It acts as a symmetric Finite Impulse Response (FIR) filter with coefficients $c[n]$ (Savitzky and Golay, 1964). Its transfer function is:

$$H_{SG}[k] = c[0] + 2 \sum_{j=1}^{(W-1)/2} c[j] \cos(j\omega_k) \quad (19)$$

Similarly, for a Hanning window $w[n]$ (Harris, 1978; Blackman and Tukey, 1959), calculation of the transfer function follows the general symmetric FIR case:

$$H_{Hann}[k] = w[0] + 2 \sum_{j=1}^{(W-1)/2} w[j] \cos(j\omega_k) \quad (20)$$

Figure 3 details the spectral characteristics of these kernels. Fig. 3a corresponds to the discrete boxcar transfer function calculated using Eq. (4), while (b) displays the analytical continuous formulation from Eq. (7). Comparison of these two confirms the accuracy of the analytical model. Fig. 3c shows the Gaussian kernel response (Eq. (18)), where the width parameter was set to $\sigma = W/\sqrt{12}$ to ensure variance consistency with the boxcar window. The Savitzky-Golay filter (Fig. 3d), implemented with polynomial order $p = 3$ via Eq. (19), exhibits a broader passband but introduces evident spectral ringing. Finally, the Hanning window (Fig. 3e), computed using Eq. (20), demonstrates superior sidelobe suppression. These behavioral properties remain consistent across different window sizes (Fig. 3g-l). Practical guidance for selecting among these kernels is given in Section 5.

3 Application to Real Radar Observations

Real radar data typically suffers from issues such as beam blockage, attenuation, and removal of data points during quality control, leading to gaps in the observational field. Figure 4 demonstrates the application of the DCT method to C-band radar reflectivity data collected on February 23, 2022. The raw data (Fig. 4a) exhibits characteristic speckle noise. Even with a compact window size of 5 pixels, the method achieves substantial noise reduction while retaining feature definition. A notable implementation difference is that the spatial kernel (Fig. 4b) explicitly ignores NaN values, whereas the DCT method (Fig. 4c) requires a continuous field, handled here via linear interpolation. Despite this, the DCT approach, using a range-adaptive continuous boxcar window, effectively preserves mesoscale (1-2 km) features like cloud-top generating cells shown at cloud top in Fig. 4e-h, where sharp gradients are maintained against the noise floor, solving a common challenge where aggressive smoothing often washes out such fine structures.

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 (The Astropy Collaboration et al., 2022) is a robust drop-in upgrade that addresses both the low-frequency bias and the kink. More generally, support-normalized filtering (Knutsson and Westin, 1993; Westin et al., 1994; Pham et al., 2006), 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 (Valdivia-Prado et al., 2026).

The uniqueness of the DCT approach compared to other methods lies in its ability to leverage the separable convolution property even when the physical kernel size varies dynamically. In spatial domain methods, a range-dependent kernel requires recalculating window indices for every range bin, a process that breaks vectorization and incurs heavy branch prediction penalties inside nested loops. In contrast, the DCT method handles this complexity through simple element-wise matrix multiplication in the spectral domain, utilizing the analytical transfer function to represent the continuous variation of the kernel width. This separation of geometry (kernel definition) from computation (filtering) is the core advantage that distinguishes DCT from traditional spatial iterative methods, allowing it to scale entirely differently with problem size.

4 Performance Benchmarking and Efficiency Analysis

To evaluate the computational efficiency, we compared the DCT method against a typical spatial dynamic kernel implementation.




Figure 4. Application to C-band radar observations on 23 Feb 2022 at 02:28 UTC. (a) Original observational field with noise and gaps. (b) Result of spatial dynamic kernel smoothing (window 5 pixels). (c) Result of DCT-based smoothing (analytical continuous boxcar, window 5 pixels). (d) Difference between spatial and spectral methods. The red dashed box indicates the zoom region shown in panels (e)-(h).

4.1 Experimental Setup

195 The benchmarks were conducted in two parts to isolate the effects of window size and total data volume:

- **Window Size Test:** Performed using RHI radar observations with dimensions 474×1180 pixels (approx. 0.56 million pixels). Here, the window size was varied across $W \in \{5, 9, 15, 31, 51, 101\}$ pixels to capture the realistic overhead of handling jagged array edges and memory access patterns in operational data.
- **Data Size Test:** Performed using a fixed window size of $W = 5$ pixels across random noise arrays with dimensions
200 representative of various operational scenarios:
 - Small (100×500 , 200×1000 pixels): Typical of single RHI scans or low-resolution grids.
 - Medium (360×1500 , 720×2500 pixels): Corresponding to standard PPI sweeps or single volume files.
 - Large (7200×2500 , 72000×2500 pixels): Representing aggregated datasets, such as a full volume scan or approximately one hour of continuous volume data (up to 180 million pixels).

205 This ensures a controlled environment to measure pure algorithmic scaling as a function of N .

4.2 Results

As shown in Fig. 5, the spatial convolution algorithm exhibits $O(N \cdot W^2)$ complexity. Our benchmarks reveal that as the window size increases from 5 to 101 pixels, the execution time for the spatial dynamic kernel escalates drastically from 8.5 seconds to over 44 seconds. In sharp contrast, the DCT method’s complexity is dominated by the $O(N \log N)$ transform step,
210 making the smoothing operation effectively independent of window size. The analytical kernels (Gaussian and Analytical Boxcar) maintain a near-constant execution time of approximately 60 ms regardless of the filter width due to the analytical representation of the kernel (see Eqs. 7 and 18). This results in massive performance gains: for a large 101×101 pixels window, the Gaussian kernel achieves a speedup factor exceeding 800 times, and the Analytical Boxcar over 700 times. Even discrete DCT kernels like Savitzky-Golay and Hanning, which require iterative summation, outperform the spatial approach by factors
215 of 60 – 100 times.

The scalability analysis on dataset size further underscores the operational viability of the spectral approach. For a massive dataset representing approximately one hour of volume data (180 million pixels), the spatial dynamic kernel required over 42 min (2518 s) to complete. The Gaussian DCT implementation processed the same volume in less than 12 s. This orders-of-magnitude improvement transforms computationally prohibitive geometric processing into a task feasible for real-time
220 operational pipelines.

5 Discussion

Noise removal remains a critical challenge in scientific signal processing. In radar meteorology, noise manifests in diverse forms, ranging from ground clutter in the Doppler spectrum to thermal noise in I/Q voltage data. This study specifically targets




Figure 5. Computational benchmark results. (a) Execution time versus window size for a fixed data volume. (b) DCT speedup factors relative to the spatial dynamic kernel as a function of window size. (c) Execution time versus total number of pixels for a fixed window size. (d) DCT speedup factors as a function of total number of pixels.

high-frequency fluctuations, or speckle, within moment data such as reflectivity. While we model this noise as randomly distributed and range-independent, the intrinsic polar geometry of radar data introduces a spatial dependency: data point density decreases inversely with range ($1/r$). Consequently, scaling the smoothing window with range implies using fewer data points at longer distances, potentially compromising noise reduction efficacy. Given that the signal-to-noise ratio (SNR) typically degrades with range and beam broadening obscures fine-scale features, there is a theoretical basis for adaptive window scaling.

Such an approach would need to balance noise suppression against the preservation of detectable spatial features. 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-angle 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-Prado et al., 2026), which extends the present method to local statistics in polar coordinates through support-normalized filtering.

A common approach to denoising is to simply truncate high frequencies in the spectral domain, effectively applying an Ideal Low-Pass Filter. However, a hard cutoff in the frequency domain takes the shape of a boxcar function; by the convolution theorem, the inverse transform of this boxcar is a Sinc function ($\sin(x)/x$) in the spatial domain. The Sinc function is characterized by prominent oscillating tails, which manifest as spurious ringing artifacts near sharp signal transitions—a distinct error known as the Gibbs phenomenon (Gottlieb and Shu, 1997). Unlike the hard cutoff, the proposed DCT method multiplies the spectrum by a smooth shaping function (e.g., Gaussian or Hanning, as shown in Eq. 6), which decays gradually. This avoids the abrupt spectral discontinuity that causes the spatial ringing, effectively preserving the monotonicity of step-like features. Furthermore, the DCT’s implicit even symmetry at boundaries significantly reduces edge artifacts compared to the DFT (FFT), which assumes periodic boundaries (Ahmed et al., 1974; Makhoul, 1980). This property is a primary reason why DCT is the standard for image compression algorithms like JPEG (ITU-T, 1992). To visualize the Gibbs phenomenon, Fig. 6 presents a comparison between an Ideal Low-Pass filter (which causes ringing) and the proposed DCT smoothing. The sharp cutoff in the frequency domain for the low-pass filter creates oscillations in the spatial domain near discontinuities. The DCT method, using a smooth Gaussian-like transfer function, avoids this issue entirely, preserving the monotonicity of the step function without artificial ringing.

The Gibbs phenomenon will behave differently on different radar variables. While this study focused on reflectivity (Z_e), the DCT smoothing method is applicable to all radar moments, provided their physical characteristics are respected. Doppler velocity presents a unique challenge due to aliasing, where the velocity value wraps around the Nyquist interval (e.g., jumping from $+V_{max}$ to $-V_{max}$), creating sharp, artificial discontinuities. Smoothing across these folded boundaries is problematic because it mathematically averages physically disparate velocities (e.g., averaging $+15$ and -15 m s^{-1} to 0 m s^{-1}), and the extreme artificial gradients at the fold trigger severe Gibbs ringing if not handled correctly. Consequently, velocity unfolding (de-aliasing) is a strict prerequisite for spectral smoothing to ensure the underlying field is continuous and physically meaningful before the transform is applied. For differential reflectivity (Z_{dr}) and differential phase (ϕ_{dp}), the smoothing process is identical to that of reflectivity. Furthermore, this technique could be advantageous for estimating the specific differential phase (K_{dp}), which typically requires substantial smoothing to mitigate noise. However, for correlation coefficient (ρ_{hv}), care must be taken to ensure smoothed values remain within the valid $[0, 1]$ interval, although a normalized convolution naturally preserves the convex hull of the data.

While the Fast Fourier Transform (FFT) is often the default choice for spectral convolution, the DCT offers distinct advantages for real-valued radar data. Since the DCT involves only real arithmetic, it avoids the storage and computation overhead




Figure 6. Illustration of the Gibbs phenomenon. (a) Reconstructed signal comparison: an ideal low-pass filter (red line) produces ringing artifacts (Gibbs effect) near sharp transitions, whereas the DCT-based smoothing with a continuous Gaussian taper (blue line) provides a clean, physically consistent result compared to the reference signal (dashed black). (b) Corresponding DCT spectral domain visualization showing the hard cut vs. the smooth Gaussian weighting.

of complex numbers inherent to the FFT. This results in faster execution and simpler implementation, as supported by the benchmarks (Section 4) and literature (Makhoul, 1980).

265 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 (Blackman and Tukey, 1959; Harris, 1978; Kaiser, 1966; Welch, 1967), are not the focus of the present study.

6 Conclusions

We have presented a comprehensive method for smoothing polarimetric radar data using the Discrete Cosine Transform. By deriving the analytical transfer function for continuous kernels and formulating the 2D range-dependent convolution in the spectral domain, we achieved a method that is both physically consistent and computationally superior to spatial domain alternatives.

The primary contribution of this work is the resolution of the conflict between geometric accuracy and computational efficiency. Traditional spatial methods require $O(N \cdot W^2)$ operations, meaning the computational cost explodes as the smoothing window W increases to match beam broadening at long ranges. In contrast, the DCT method operates in $O(N \log N)$ time, rendering the processing time effectively independent of the kernel size.

Our performance benchmarks demonstrate that this theoretical advantage translates into massive operational gains:

1. **Kernel Scalability:** For large smoothing windows (e.g., 101 pixels), the analytical DCT kernels (Gaussian and Boxcar) provided speedup factors exceeding 800 times compared to the spatial dynamic kernel.
2. **Data Volume Scalability:** For a dataset representing approximately one hour of continuous volume scans (180 million pixels), the spatial algorithm required over 42 min (2518 s) to complete. The DCT implementation processed the same volume in less than 12 s.

This orders-of-magnitude improvement reduces hours of processing time to seconds. Complex geometric corrections that were previously computationally prohibitive can now be integrated into standard operational pipelines with negligible latency.

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 (Valdivia-Prado et al., 2026). The DCT approach avoids the Gibbs phenomenon associated with the Ideal Low-Pass Filter discussed in Section 5 and preserves the convex hull of the data, making it safe for variables such as correlation coefficient (ρ_{hv}). 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. Overall, this spectral framework offers a robust, high-performance standard
300 for modern weather radar analysis.

Code availability. The code used for the benchmarks and smoothing algorithms is available at <https://github.com/JValdivia23/radar-dct-smoothing> and archived on Zenodo at <https://doi.org/10.5281/zenodo.18226677>.

Data availability. Data from the WINTRE-MIX campaign are available at https://data.eol.ucar.edu/master_lists/generated/wintre-mix/, and the COW dataset can be accessed at <https://data.eol.ucar.edu/dataset/612.037>. The specific weather radar file and benchmark results analyzed
305 in this study are available in the project repository at https://github.com/JValdivia23/radar-dct-smoothing/tree/main/_legacy_data.

Author contributions. JV, WC, and KF designed the study; JV performed the analysis and wrote the manuscript.

Competing interests. The authors declare that they have no conflict of interest.

Disclaimer. The views expressed in this article are those of the authors.

Acknowledgements. This research was sponsored by the National Science Foundation grant AGS 2114011. The authors acknowledge the
310 use of generative AI tools (Google Gemini and OpenAI GPT) for proofreading, improving writing clarity, and verifying mathematical formulations. All original ideas, analysis, and conclusions remain the sole responsibility of the authors.

References

- Ahmed, N., Natarajan, T., and Rao, K.: Discrete Cosine Transform, IEEE Transactions on Computers, C-23, 90–93, <https://doi.org/10.1109/T-C.1974.223784>, 1974.
- 315 Atlas, D.: Radar in Meteorology: Battan Memorial and 40th Anniversary Radar Meteorology Conference, American Meteorological Society, Boston, MA, <https://doi.org/10.1007/978-1-935704-15-7>, 1990.
- Blackman, R. B. and Tukey, J. W.: The Measurement of Power Spectra from the Point of View of Communications Engineering, Dover Publications, New York, 1959.
- Bringi, V. N. and Chandrasekar, V.: Polarimetric Doppler Weather Radar: Principles and Applications, Cambridge University Press, Cambridge, UK, <https://doi.org/10.1017/CBO9780511541094>, 2001.
- 320 Gottlieb, D. and Shu, C.-W.: On the Gibbs Phenomenon and Its Resolution, SIAM Review, 39, 644–668, <https://doi.org/10.1137/S0036144596301390>, 1997.
- Harris, F.: On the use of windows for harmonic analysis with the discrete Fourier transform, Proceedings of the IEEE, 66, 51–83, <https://doi.org/10.1109/PROC.1978.10837>, 1978.
- 325 ITU-T: Recommendation T.81: Information technology – Digital compression and coding of continuous-tone still images – Requirements and guidelines, Tech. rep., International Telecommunication Union, 1992.
- Ivić, I. R., Zrnić, D. S., and Torres, S. M.: Whitening in Range to Improve Weather Radar Signal Processing, Journal of Atmospheric and Oceanic Technology, 30, 2733–2748, <https://doi.org/10.1175/JTECH-D-13-00007.1>, 2013.
- Kaiser, J. F.: Digital Filters, in: System Analysis by Digital Computer, edited by Kuo, F. F. and Kaiser, J. F., chap. 7, pp. 218–285, John Wiley & Sons, New York, 1966.
- 330 Knutsson, H. and Westin, C.-F.: Normalized and Differential Convolution: Methods for Interpolation and Filtering of Incomplete and Uncertain Data, Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 515–523, <https://doi.org/10.1109/CVPR.1993.341081>, 1993.
- Levi, H.: A GEOMETRIC CONSTRUCTION OF THE DIRICHLET KERNEL*, Transactions of the New York Academy of Sciences, 36, 640–643, <https://doi.org/10.1111/j.2164-0947.1974.tb03023.x>, 1974.
- 335 Makhoul, J.: A Fast Cosine Transform in One and Two Dimensions, IEEE Transactions on Acoustics, Speech, and Signal Processing, 28, 27–34, <https://doi.org/10.1109/TASSP.1980.1163351>, 1980.
- Ng, M. K., Chan, R. H., and Tang, W.-C.: A Fast Algorithm for Deblurring Models with Neumann Boundary Conditions, SIAM Journal on Scientific Computing, 21, 851–866, <https://doi.org/10.1137/S1064827598341384>, 1999.
- 340 Oppenheim, A. V., Schafer, R. W., and Buck, J. R.: Discrete-Time Signal Processing, Prentice Hall, Upper Saddle River, NJ, 2nd edn., 1999.
- Pham, T. Q., van Vliet, L. J., and Schutte, K.: Robust Fusion of Irregularly Sampled Data Using Adaptive Normalized Convolution, EURASIP Journal on Advances in Signal Processing, 2006, <https://doi.org/10.1155/ASP/2006/83268>, 2006.
- Savitzky, A. and Golay, M. J. E.: Smoothing and Differentiation of Data by Simplified Least Squares Procedures, Analytical Chemistry, 36, 1627–1639, <https://doi.org/10.1021/ac60214a047>, 1964.
- 345 Smith, S. W.: The Scientist and Engineer’s Guide to Digital Signal Processing, California Technical Publishing, San Diego, CA, ISBN 0-9660176-3-3, 1997.
- Strang, G.: The Discrete Cosine Transform, SIAM Review, 41, 135–147, <https://doi.org/10.1137/S0036144598336745>, 1999.

- The Astropy Collaboration et al.: The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package, *The Astrophysical Journal*, 935, 167, <https://doi.org/10.3847/1538-4357/ac7c74>, 2022.
- 350 Valdivia-Prado, J. M., Chapman, W. E., and Friedrich, K.: Spectral-Domain Local Statistics with Missing-Data Support for Cartesian and Polar Grids, <https://doi.org/10.48550/ARXIV.2604.22721>, 2026.
- Welch, P. D.: The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms, *IEEE Transactions on Audio and Electroacoustics*, AU-15, 70–73, 1967.
- Westin, C.-F., Nordberg, K., and Knutsson, H.: On the Equivalence of Normalized Convolution and Normalized Differential Convolu-
 355 tion, in: *Proceedings of ICASSP '94. IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. V/457–V/460, <https://doi.org/10.1109/ICASSP.1994.389389>, 1994.
- Wilson, J. W., Feng, Y., Chen, M., and Roberts, R. D.: Weather Radar Ground Clutter. Part I: Identification, Modeling, and Simulation, *Journal of Atmospheric and Oceanic Technology*, 29, 1193–1207, <https://doi.org/10.1175/JTECH-D-11-00109.1>, 2012.