https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

From spectrogram to range-weighting function: a unified 2-D convolution model for coherent Doppler wind lidar

Pierre E. Allain[1], Andrew Hastings Black[1], David Tomline Michel[2], Théo Martin[2], Laurent Mugnier[2], and Matthieu Valla[2]

[1]Vaisala, France
[2]ONERA/DOTA, Palaiseau, France

Correspondence: Pierre E. Allain (pierre.allain@vaisala.com)

Abstract. The range-weighting function (RWF) of a pulsed coherent Doppler wind lidar is commonly identified with the emitted pulse intensity, giving a range resolution of half the pulse length. We show that this identification is the last step of a hierarchy of approximations whose assumptions are rarely stated.

Starting from the discrete sliced-atmosphere model, the expected Doppler spectrogram is a two-dimensional convolution between the atmosphere and a range–frequency point-spread function. Its frequency marginal is a one-dimensional range kernel—the cross-correlation of the squared analysis window with the pulse intensity—that governs integrated power and any scalar gate-by-gate velocity estimator when the spectrum is symmetric. This window-completed kernel is wider than the pulse alone—by 31% for the Gaussian pulse and Hann window examined here at $N_w \simeq N_\tau$. When backscatter varies across the range gate (cloud edges, aerosol layers), the velocity is additionally weighted by the local carrier-to-noise ratio. When the spectrum is asymmetric (chirp) or multi-modal (severe turbulence, wakes, vortices), no one-dimensional weighting is sufficient and the full two-dimensional model is required, either for interpretation or for spectrogram inversion.

We tabulate each model with its assumptions and domain of validity, so that practitioners can pick the simplest model consistent with the observable of interest and identify when classical one-dimensional models silently fail.

1 Introduction

Pulsed coherent Doppler wind lidars measure line-of-sight wind speed from the Doppler shift of laser light backscattered by atmospheric aerosols. They are deployed in wind-energy resource assessment and turbine control (Mann et al., 2010; Sathe et al., 2011; Peña et al., 2017), atmospheric boundary-layer (ABL) meteorology, aviation-safety wake-vortex monitoring (Hannon and Thomson, 1994; Köpp et al., 2004), and turbine-wake characterisation (Bingöl et al., 2010; Trujillo et al., 2011). In all these applications, a range-gated Doppler estimate is not a point measurement: the received signal is a coherent sum over atmospheric slices weighted by the emitted pulse and further shaped by the receiver processing window. This range mixing is commonly described by a one-dimensional range-weighting function (RWF), identified with the pulse intensity profile (Frehlich, 1997;

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

Frehlich and Cornman, 1999). A common engineering simplification is then to set the range resolution to $c\tau/2$, where τ is the pulse duration.

However, a coherent Doppler lidar does not directly produce a velocity. Its primary output is a local Doppler spectrum, computed as a windowed periodogram at each range gate. A spectrogram is the collection of these spectra along the line of sight. A velocity, a carrier-to-noise ratio (CNR), or a spectral width are different scalar reductions of this spectrum. They need not share the same range response unless additional assumptions hold. For integrated power, frequency information is discarded and a one-dimensional range kernel suffices. For velocity, the same kernel applies only when the spectrum is symmetric in frequency. When the Doppler spectrum itself must be interpreted, as in turbulence characterisation (Frehlich, 1994; Frehlich and Cornman, 1999; Sathe et al., 2011; Peña et al., 2017), wake studies (Bingöl et al., 2010; Trujillo et al., 2011), or vortex detection (Hannon and Thomson, 1994; Köpp et al., 2004), a one-dimensional RWF is no longer the appropriate forward model.

The two-dimensional convolution structure of the expected spectrogram was derived by Martin et al. (2025) (building on the sliced-atmosphere model of Salamitou et al. (1995) and Valla (2005)) and used as the forward operator of a super-resolution inversion, but the relationship between this 2-D model and the classical 1-D RWF was not analysed.

The purpose of this paper is to fill that gap. We show that the frequency marginal of the 2-D PSF is a one-dimensional range kernel h_c, and that every 1-D RWF in common use is a further reduction of this kernel under explicit assumptions. We organise the models into a hierarchy (Table 1) and identify the regime in which each level is valid.

2 Discrete signal model

We work with a discrete sliced-atmosphere model (Salamitou et al., 1995; Valla, 2005; Martin et al., 2025). The line of sight is sampled at period $T_s = 1/F_s$ into range slices of width $\Delta z = cT_s/2$; the emitted pulse spans N_τ samples (so $\tau = N_\tau T_s$) and the local Doppler spectrum is computed with an analysis window of length N_w. Each atmospheric slice u carries a backscatter power $\sigma_a^2[u]$, a single local Doppler frequency $f_D[u]$ (related to the line-of-sight velocity by $f_D = v/(\lambda/2)$ for a heterodyne receiver, λ being the laser wavelength), and is independent from its neighbours ("mono-Doppler" slices). The complex baseband signal at pulse m is the coherent sum of the emitted pulse weighted by the slice amplitudes, Doppler-shifted at the observation time:

$$s_m[n] = \sum_{\ell=0}^{N_\tau-1} \sqrt{p_\ell}\, a_m[n-\ell]\, e^{j2\pi f_D[n-\ell]\, n\, T_s} + \eta_m[n], \tag{1}$$

with circular Gaussian noise η_m of variance σ_η^2 and decorrelated slices

$$\mathbb{E}\{a_m[u] a_m^*[v]\} = \sigma_a^2[u]\, \delta_{u,v}. \tag{2}$$

We define the slice carrier-to-noise ratio

$$\mathrm{CNR}_a[u] = \sigma_a^2[u]/\sigma_\eta^2. \tag{3}$$

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

This model is the discrete counterpart of the heterodyne covariance framework of Frehlich (1994); Frehlich and Cornman (1999) when the range mixing induced by the pulse is retained explicitly but the turbulent intra-gate covariance is not expanded.

3 Forward model: a 2-D convolution

The local Doppler spectrum at gate n is the windowed periodogram averaged over M pulses,

$$\hat{S}[n,q] = \frac{1}{M}\sum_{m=1}^{M}\left|\sum_{i=0}^{N_w-1} w_i\, e^{-j2\pi qi/N_w}\, s_m[n+i]\right|^2, \tag{4}$$

with normalised window weights $\sum_i w_i^2 = 1$ and bin frequencies $f_q = q/(N_w T_s)$. Substituting (1) into (4), the speckle decorrelation (2) collapses the resulting double sum over slice indices to its diagonal; the remaining slice-only phase is killed by the modulus squared, and grouping by the range offset $k = n - u$ (Appendix A) yields a single sum over atmospheric slices,

$$\mathbb{E}\{\hat{S}[n,q]\} = \sum_u \sigma_a^2[u]\,\mathrm{PSF}[n-u, f_D[u] - f_q] + \sigma_\eta^2, \tag{5}$$

where the range-frequency point-spread function (PSF) reads

$$\boxed{\mathrm{PSF}[k,\Delta f] = \left|\sum_{i=0}^{N_w-1} w_i\sqrt{p_{k+i}}\, e^{j2\pi\Delta f\, iT_s}\right|^2}, \tag{6}$$

with the convention $p_q = 0$ for $q < 0$ or $q > N_\tau - 1$. The support in range offset is $-(N_w - 1) \le k \le N_\tau - 1$. We will refer to the product $g_i = w_i\sqrt{p_i}$ as the *overlap function* between window and pulse; the summand in (6) is its shifted version $g_i(k) = w_i\sqrt{p_{k+i}}$.

Introducing the range-frequency object

$$o[u,f] = \sigma_a^2[u]\,\delta(f - f_D[u]), \tag{7}$$

(5) is a 2-D discrete convolution,

$$\boxed{\mathbb{E}\{\hat{S}[n,q]\} - \sigma_\eta^2 = (\mathrm{PSF} *_{z,f} o)[n, f_q].} \tag{8}$$

Equation (8) is the forward model used in the rest of the paper. On the frequency axis, each Doppler line is written as a single DFT bin; in general a line spreads over several bins, and the operator then samples the PSF at the true line frequency without any grid rounding (Appendix A). Two properties of the PSF are central. First, when the pulse envelope and the analysis window are real, the PSF is symmetric in frequency,

$$\mathrm{PSF}[k,\Delta f] = \mathrm{PSF}[k,-\Delta f]. \tag{9}$$

Second, the PSF is generally not separable in range and frequency:

$$\mathrm{PSF}[k,\Delta f] \neq A[k]B[\Delta f], \tag{10}$$

because the summand $w_i\sqrt{p_{k+i}}$ in (6) couples the range offset k to the Fourier index i. Approximate separability only holds when the window is much narrower than the pulse.

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

4 1-D reductions of the 2-D model

Three observables of practical interest—integrated power, CNR, and centroid velocity—are scalar reductions of the spectrogram. Each reduction summed over frequency makes the spectral dimension disappear and exposes the same 1-D range kernel,

$$\boxed{h_c[k] = \frac{1}{N_w} \sum_q \mathrm{PSF}[k, f_q] = \sum_{i=0}^{N_w-1} w_i^2 \, p_{k+i}.} \tag{11}$$

The kernel h_c is the discrete cross-correlation of the squared analysis window with the pulse intensity—the *window-completed* range-weighting function. It is the discrete counterpart of the pulse-and-gate weighting used to interpret volume-averaged line-of-sight velocities (Frehlich, 1997; Frehlich and Cornman, 1999; Mann et al., 2010). When the window is much shorter than the pulse, (11) collapses to the pulse-only approximation,

$$h_c[k] \approx p_k. \tag{12}$$

The derivations of (11)–(17) from (8) are given in Appendix B.

4.0.1 Measured CNR.

Defining the wideband CNR from the noise-subtracted spectrogram,

$$\mathrm{CNR}[n] = \frac{\sum_q \mathbb{E}\{\hat{S}[n,q]\} - N_w \sigma_\eta^2}{N_w \sigma_\eta^2}, \tag{13}$$

one obtains

$$\boxed{\mathrm{CNR}[n] = \sum_k \mathrm{CNR}_a[n-k] \, h_c[k].} \tag{14}$$

The measured CNR is the atmospheric CNR filtered in range by h_c. No further assumption is needed because no spectral information is used.

4.0.2 Centroid velocity with variable CNR.

The spectral centroid of the noise-subtracted spectrum,

$$\hat{f}_D[n] = \frac{\sum_q f_q \, (\mathbb{E}\{\hat{S}[n,q]\} - \sigma_\eta^2)}{\sum_q (\mathbb{E}\{\hat{S}[n,q]\} - \sigma_\eta^2)}, \tag{15}$$

is a weighted average of $f_D[u]$ *provided the PSF is symmetric in frequency* (9): the change of variable $\Delta f = f_D[u] - f_q$ in the numerator splits it into an $f_D[u]$-term and a Δf-term that vanishes by (9) (Appendix B). The local backscatter power then sets the weighting:

$$\boxed{\hat{f}_D[n] = \frac{\sum_k f_D[n-k] \, \mathrm{CNR}_a[n-k] \, h_c[k]}{\sum_k \mathrm{CNR}_a[n-k] \, h_c[k]}.} \tag{16}$$

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

When CNR_a is approximately constant over the support of h_c, it factorises out and the centroid velocity reduces to a pure 1-D convolution,

$$\boxed{\hat{f}_D[n] = \frac{\sum_k f_D[n-k]h_c[k]}{\sum_k h_c[k]}.} \tag{17}$$

Equation (17) is the discrete counterpart of the pulse-weighted velocity v_{wgt} introduced by Frehlich (1997, Eq. 23) and routinely used to interpret lidar velocities as pulse-and-gate-averaged line-of-sight velocities (Mann et al., 2010). Equation (16) makes the CNR weighting explicit; it becomes quantitatively significant as soon as the backscatter varies appreciably across the gate (Sec. 6).

4.0.3 Beyond the centroid.

Other Doppler estimators do not realise the same 1-D reduction in general. Peak detectors depend on the location of the spectral maximum; Gaussian or maximum-likelihood (MLE) fits depend on the assumed spectral shape (Rye and Hardesty, 1993a, b; Dabas, 1999; Smalikho, 2003; Held and Mann, 2018). However, when the local spectrum is unimodal and symmetric in frequency (the regime in which the 1-D reductions above hold), the asymptotic frequency returned by a centroid and by an MLE coincide, because the mean of a symmetric unimodal distribution equals its mode (finite-sample MLE differs by an estimator-variance term independent of h_c). The kernel h_c therefore applies to any scalar gate-by-gate Doppler estimator in that regime. The distinction becomes material when the spectrum is asymmetric (chirped pulse, multi-component returns) or when the estimator exploits spectral structure across gates, as in the maximum a posteriori (MAP) super-resolution of Martin et al. (2025).

5 Hierarchy of models and assumptions

The preceding derivation gives a hierarchy of models. Each level is obtained from the level above by discarding information or by adding an assumption. The hierarchy is therefore tied to the observable being interpreted: integrated power, centroid velocity, or the full Doppler spectrum. In Table 1, $*$ denotes discrete convolution and $\tilde{p}_k = p_{-k}$, so that the cross-correlation in (11) is written as a convolution.

The table can be read top-down as a decision procedure: start with the 2-D model and simplify only when the assumptions of the next row are satisfied.

The common approximation that associates the range resolution with $c\tau/2$ corresponds to the last row with an additional top-hat pulse assumption. The table should not be read as a list of interchangeable models: each row corresponds to a different observable or to an additional approximation.

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 1. Hierarchy of models, kernels, assumptions, and domains of validity.

Model	Formula	Kernel	Additional assumptions	Valid when	Typical use
2-D spectrogram PSF	$\mathbb{E}\{\hat{S}\} = \mathrm{PSF} * o$ (8)	$\mathrm{PSF}[k, \Delta f]$ (6)	Mono-Doppler slices; speckle decorrelation across slices.	Whenever the mono-Doppler slice assumption holds.	Spectral shape, turbulence, shear, wakes, vortices, deconvolution (Martin et al., 2025).
Frequency marginal	$\mathrm{CNR} = \mathrm{CNR}_a * h_c$ (14)	$w^2 * \tilde{p}$	Only total spectral power is interpreted.	Always (no frequency information used).	CNR, signal availability, backscatter proxy.
CNR-weighted centroid	$\hat{f}_D = (f_D\,\mathrm{CNR}_a) * h_c \,/\, \mathrm{CNR}_a * h_c$ (16)	$w^2 * \tilde{p}$	(9) holds; scalar gate-by-gate estimator (centroid, fit, MLE).	Symmetric PSF and unimodal spectrum.	Variable CNR: cloud edges, aerosol layers, backscatter gradients.
Constant-CNR	$\hat{f}_D = f_D * h_c \,/\, \sum h_c$ (17)	$w^2 * \tilde{p}$	Additionally: CNR_a constant over support of h_c.	Backscatter smooth over the gate $L_\sigma \gg L_h$.	Homogeneous ABL, wind-energy volume averaging.
Pulse-only	$\hat{f}_D = f_D * p \,/\, \sum p$ (12)	p	Additionally: processing window negligible.	$N_w \lesssim N_\tau/4$.	Back-of-the-envelope range resolution.

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

6 When can a one-dimensional RWF be used?

Table 1 lists five models. In practice one wants the simplest that is still consistent with the observable of interest. Three length scales govern the choice.

Let L_h be the RMS range width of the window-completed kernel h_c (Eq. 11); Let $v_{1/2}$ be the half width at half maximum (HWHM) of $\mathrm{PSF}[k{=}0,\cdot]$ in velocity units. For Gaussian p and w with RMS durations σ_p and σ_w (in samples) and a laser wavelength λ, the overlap function $g_i = w_i\sqrt{p_i}$ introduced after (6) is itself Gaussian and all widths have closed forms:

$$v_{1/2} = \frac{\lambda\sqrt{\ln 2}}{4\pi T_s}\sqrt{\frac{1}{\sigma_w^2}+\frac{1}{2\sigma_p^2}}, \qquad L_h = \sqrt{\sigma_p^2+\tfrac{1}{2}\sigma_w^2}\,\Delta z. \tag{18}$$

The spectral resolution is set by the shorter of σ_w and $\sqrt{2}\sigma_p$; the range width grows with both. For non-Gaussian shapes, $v_{1/2}$ is read from the PSF directly.

Let L_σ be the range scale over which the atmospheric backscatter power changes appreciably. Formally, $L_\sigma = \sigma_a^2/|\partial_z\sigma_a^2|$: it is large in a well-mixed boundary layer (hundreds of metres) and small at a cloud edge (a few metres).

6.1 Window broadening of the range kernel

The pulse-only RWF (12) ignores the analysis window entirely: it sets $h_c[k] = p_k$. The true kernel (11) is wider because it is the cross-correlation of p with w^2.

Let L_p denote the RMS width of the pulse alone (the $N_w = 1$ limit of L_h).

Figure 3 (left) shows the ratio L_h/L_p, which measures how much the analysis window widens the effective range kernel beyond the pulse-only value. At $N_w = N_\tau/4$ the widening stays below 3%; at $N_w = N_\tau$ it reaches 31% for the Gaussian pulse and Hann window used here—the only pulse/window pair tested numerically in this paper, so the exact value for other shapes remains to be checked.

Wind-energy and boundary-layer processing commonly set $N_w \simeq N_\tau$, so the pulse-only RWF overstates range resolution by almost a third here—a difference large enough to matter in practice.

6.2 Variable-backscatter velocity bias

The constant-CNR formula (17) treats all slices inside the gate as equally weighted. In reality, backscatter gradients weight the centroid toward the stronger-signal side. Equation (16) accounts for this;

The velocity difference between the two models can be estimated as follows: within the gate, the backscatter gradient shifts the effective centre of the range kernel by an amount proportional to L_h/L_σ, while the velocity change across the kernel is $|\partial_z f_D|\,L_h$. Their product gives

$$\max_z\left|\hat{f}_D^{(\mathrm{cst})} - \hat{f}_D^{(\mathrm{CNRw})}\right| \simeq |\partial_z f_D|\,L_h\,\frac{L_h}{L_\sigma}, \tag{19}$$

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

as confirmed numerically in Fig. 3 (right) over $L_h/L_\sigma \in [0.01, 0.5]$.

For $|\partial_z v| = 0.1$ m/s/m and $L_h \simeq 8.3$ m, the bias stays below 0.1 m/s for $L_\sigma \gtrsim 70$ m. In a well-mixed boundary layer L_σ exceeds several hundred metres and the bias is entirely negligible; this is presumably why the constant-CNR model has survived so long without challenge.

At a cloud edge or capping inversion L_σ can drop to a few metres, making $L_h/L_\sigma \sim 1$ and the constant-CNR model unreliable. The CNR-weighted formula (16) is then the correct 1-D reduction.

6.3 Intra-gate velocity dispersion

Each atmospheric slice in (7) is assigned a single Doppler frequency. If the true intra-gate velocity dispersion σ_v is large enough to broaden the spectrum beyond the instrument width, the mono-Doppler assumption breaks down and any scalar RWF becomes incomplete.

When σ_v is non-negligible, the observed spectral width is approximately $\sqrt{v_{1/2}^2 + \sigma_v^2}$, broader than the instrument-only width $v_{1/2}$.

Practically, the 1-D reductions of Sec. 4 remain adequate as long as $\sigma_v \lesssim v_{1/2}/3$, i.e. the intra-gate dispersion contributes less than $\sim$10% to the observed spectral width; beyond that, in turbulence, strong shear, wakes or vortices, the full 2-D PSF (8) is needed.

6.4 Chirped pulse

The centroid formula (16) relies on the frequency symmetry (9). A deterministic intra-pulse phase $\psi[\ell]$ breaks it. For a linear chirp, the first moment $\bar{\Delta f}(k)$ of $\mathrm{PSF}[k, \cdot]$ depends on the range offset k, breaking (9) and forcing the full 2-D model (Sec. 7.4).

The scaling laws (19) and $L_h/L_p(N_w/N_\tau)$ are verified numerically in Sec. 7.2; the resulting thresholds appear in the "Valid when" column of Table 1.

7 Numerical verification

All simulations use a Gaussian pulse of $N_\tau = 64$ samples, a Hann analysis window of length $N_w = N_\tau$, $F_s = 250$ MHz, and $\lambda = 1.55$ μm ($\Delta z \simeq 0.6$ m per slice). The atmosphere is a synthetic boundary layer with a logarithmic wind profile, a low-level jet, and a cloud layer whose leading edge is sharper than its trailing side. These choices are not intended to reproduce a specific instrument but to exercise every transition in the hierarchy of Table 1.

7.1 PSF and marginalisation identity

Figure 1 (top) shows the two-dimensional PSF (6). Its frequency marginal is computed bin by bin and compared with the closed-form kernel $h_c[k] = \sum_i w_i^2 p_{k+i}$ (bottom). The two agree to machine precision, confirming (11). The same identity has been verified experimentally on a WindCube lidar using hard-target returns at multiple spatial resolutions; we reproduce the

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

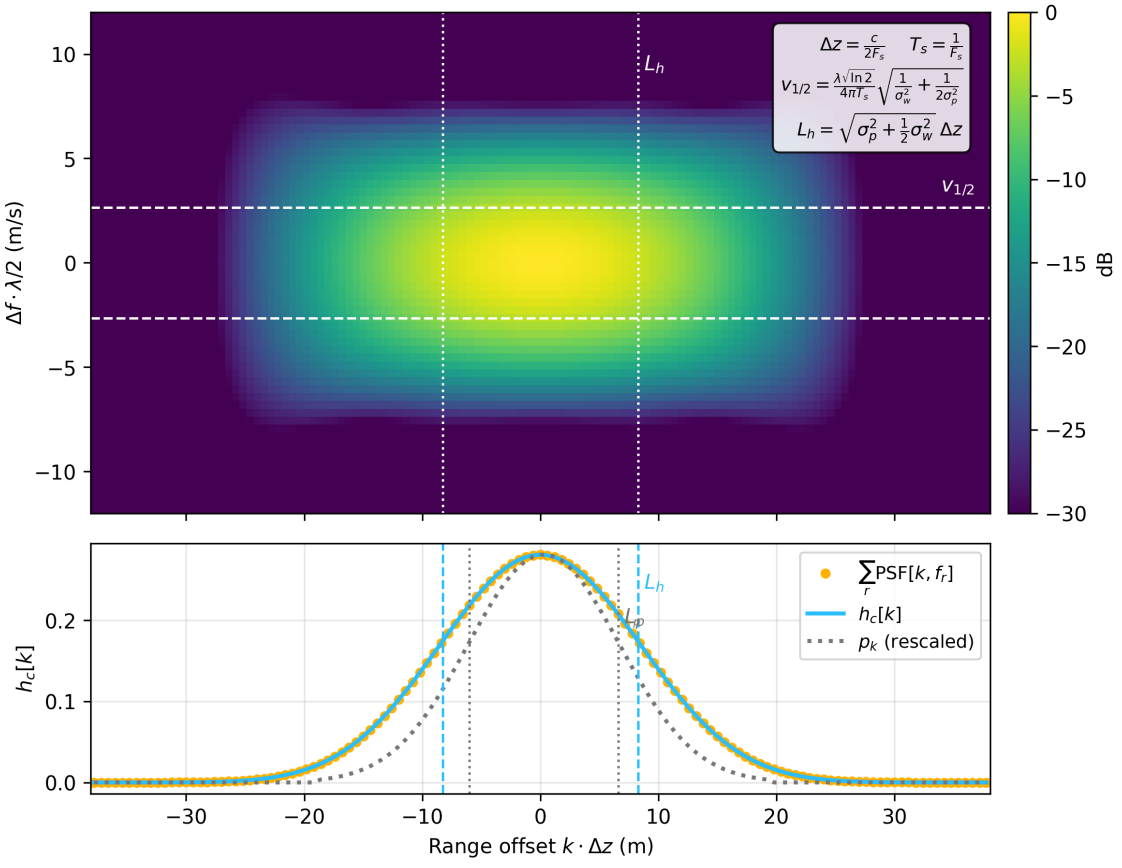


Figure 1. Top: two-dimensional PSF (6) for the reference parameters of Sec. 7. Dashed lines mark $\pm v_{1/2}$ (spectral HWHM) and $\pm L_h$ (range RMS of h_c). Bottom: the frequency marginal matches $\sum_i w_i^2 p_{k+i}$ to machine precision. Dotted vertical lines mark $\pm L_p$ (range RMS of p); h_c is visibly wider.

experimental RWF measurement of Allain et al. (2022) in Fig. 2 for comparison with the closed-form kernel h_c. The pulse-only RWF, rescaled and recentred, is narrower—the difference becoming the 31% range-resolution bias discussed in Sec. 6.

7.2 Scaling laws of the validity criteria

The two non-trivial scaling laws of Table 1 are tested independently (Fig. 3).

The left panel sweeps N_w/N_τ and plots L_h/L_p, the ratio of the range RMS of h_c to that of the pulse alone. The broadening is below 3% for $N_w \le N_\tau/4$ and reaches 1.31 at $N_w = N_\tau$.

The right panel places a Gaussian backscatter bump of width L_σ on a linear wind shear and measures the maximum velocity difference between the constant-CNR formula (17) and the CNR-weighted formula (16). The error follows $|\partial_z f_D|\, L_h\, (L_h/L_\sigma)$ over the tested range.

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

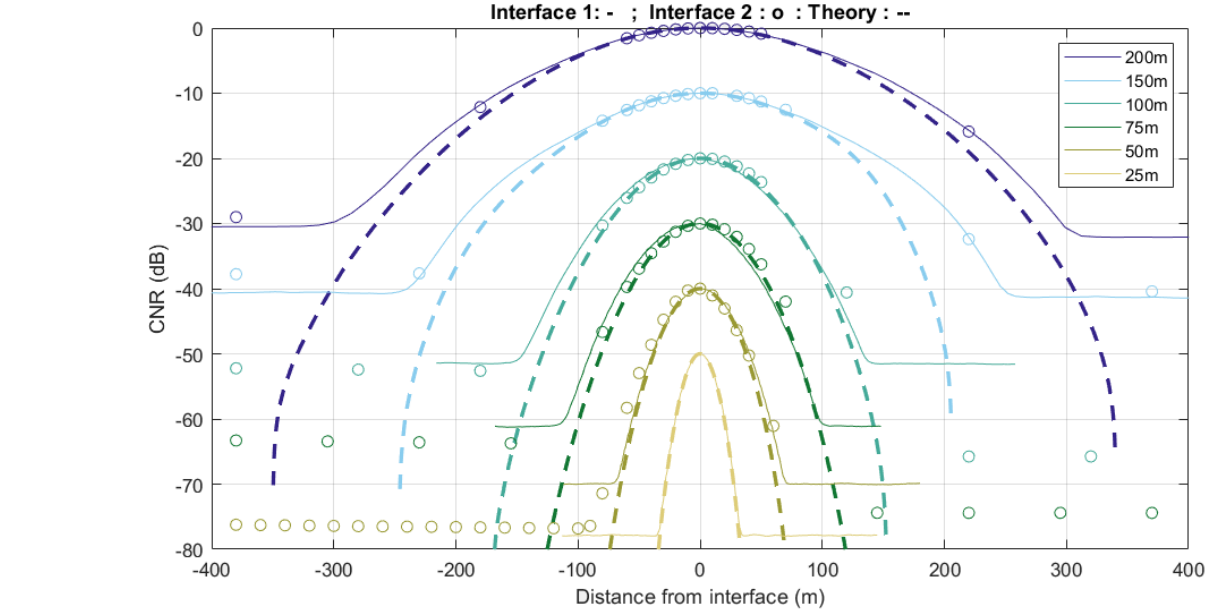


Figure 2. Experimental RWF measured on a WindCube lidar with hard-target returns at several range resolutions (data reproduced from Allain et al., 2022), overlaid with the closed-form kernel h_c of (11) computed with the corresponding instrument parameters.

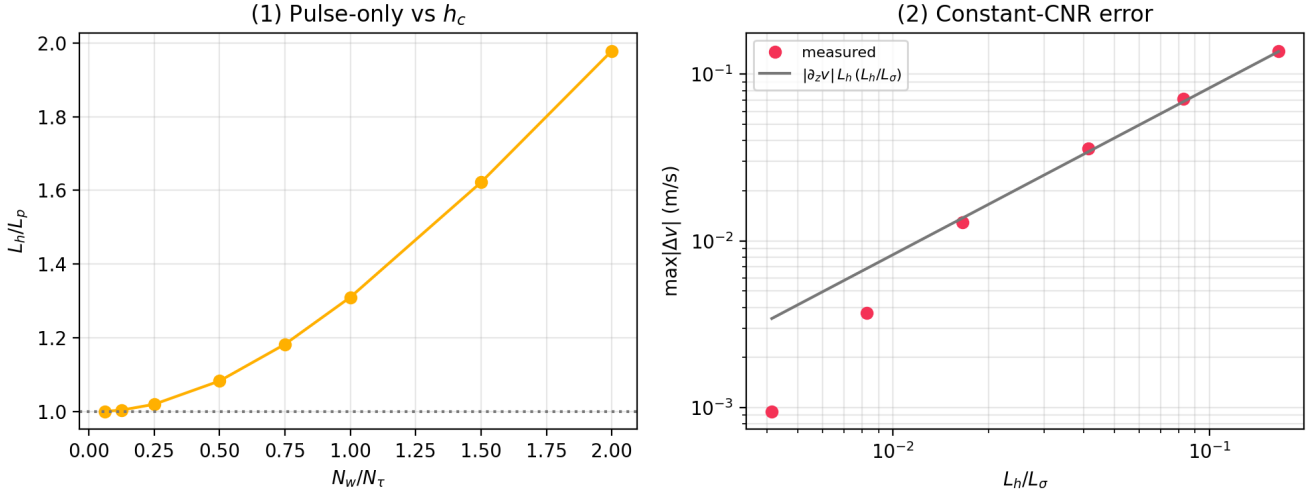


Figure 3. Validity boundaries of the two simplest 1-D models. *Left:* ratio of the window-completed range width L_h to the pulse-only width L_p as a function of the window-to-pulse length ratio N_w/N_τ; the pulse-only model ($L_h/L_p = 1$) is accurate only for $N_w \lesssim N_\tau/4$. *Right:* maximum velocity error introduced by assuming constant backscatter (Eq. 17 vs. Eq. 16), plotted against the ratio of kernel width to backscatter scale L_h/L_σ; the grey line is the scaling law (19).

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

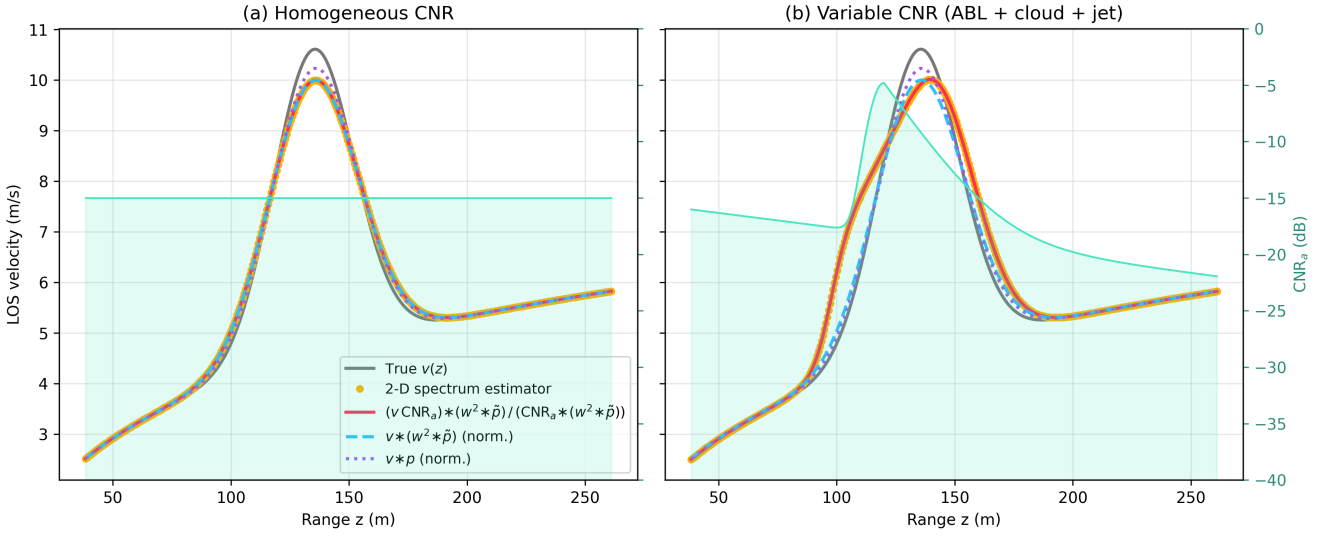


Figure 4. Centroid velocity from the 2-D spectrogram and from the 1-D reductions. (a) Homogeneous CNR_a: all models agree. (b) Variable CNR_a (ABL + cloud + jet): only the CNR-weighted RWF (16) matches the 2-D centroid. The shaded area shows $\mathrm{CNR}_a(z)$ in dB (right axis).

7.3 Velocity averaging on a realistic profile

Figure 4 applies the four 1-D reductions of Sec. 4 to two atmospheric scenarios sharing the same wind profile. Under homogeneous CNR_a (panel a), all four reductions collapse onto the 2-D centroid: the hierarchy is invisible. Under variable CNR_a (panel b), the constant-CNR RWF and the pulse-only RWF deviate near the cloud edge where $L_\sigma \sim L_h$, while the CNR-weighted formula (16) tracks the 2-D centroid exactly. This is a direct illustration of the L_h/L_σ criterion.

7.4 Chirped pulse: breakdown of the 1-D reductions

A linear chirp $\psi[\ell] = \pi\alpha(\ell T_s)^2$ is added to the reference pulse, keeping the envelope and the atmosphere identical to Fig. 4(b). The PSF becomes asymmetric in Δf, and the first moment $\overline{\Delta f}(k)$ of $\mathrm{PSF}[k,\cdot]$ is no longer zero. Crucially, it depends on the range offset k, because the analysis window sees a different portion of the chirp phase as k varies.

Subtracting a single constant, measured on a zero-velocity calibration target, removes only the k-average of this offset. A range-dependent residual remains: in the 2-D centroid, each gate averages $\overline{\Delta f}(k)$ weighted by the local $\mathrm{CNR}_a \cdot h_c$. Where CNR_a is smooth the weighting barely changes with range and the residual is small; at the cloud edge, the weights shift and the calibrated centroid departs from (16) (Fig. 5; the global offset is $\overline{\Delta f} = \alpha\tau/2 \approx 7.7$ MHz, or $\approx$6.0 m/s for the chirp rate used here). No 1-D reduction captures this effect: when (9) is broken, the full 2-D spectrogram model (8) is needed.

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

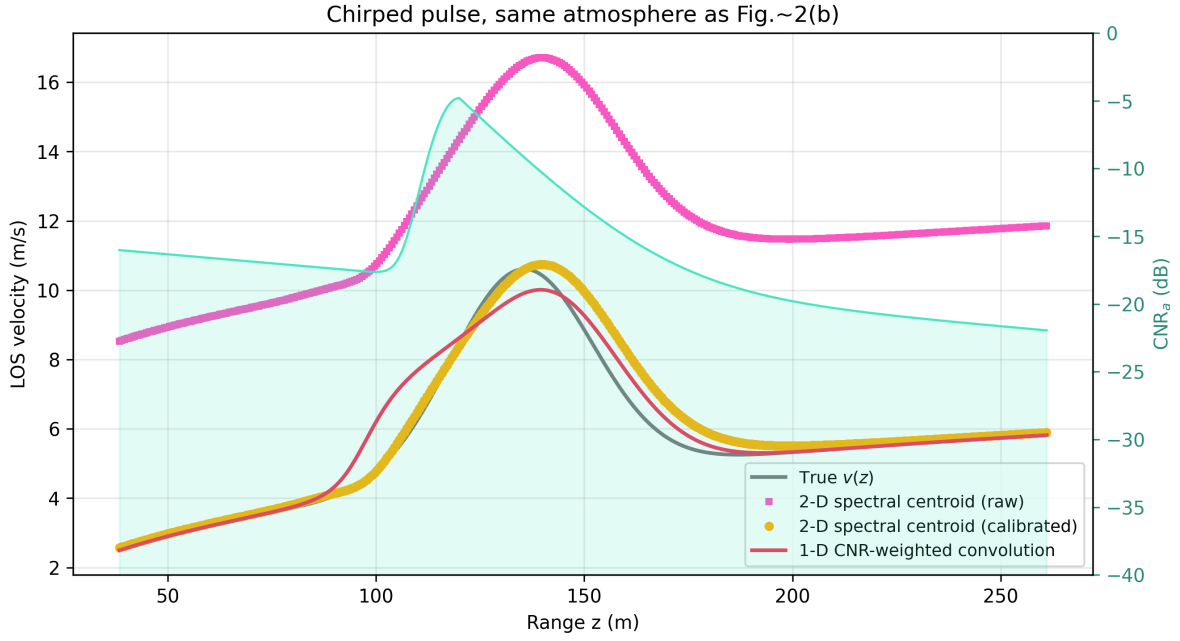


Figure 5. Same scenario as Fig. 4(b) with a chirped pulse. The raw 2-D spectral centroid (squares) carries a range-dependent bias from the asymmetric PSF. Subtracting a global offset calibrated on a zero-velocity target (circles) removes the mean bias but a residual remains wherever CNR_a varies within the gate. The 1-D CNR-weighted convolution (solid) cannot capture this effect: with a chirped pulse, no 1-D reduction is sufficient.

8 Discussion

The numerical examples of Sec. 7 use a single Gaussian-pulse / Hann-window parameter set. The closed-form widths (18) and the scaling law (19) are independent of pulse shape, so the hierarchy itself generalises; the quantitative thresholds (N_w/N_τ, L_h/L_σ) should be re-evaluated for non-Gaussian designs. Experimental validation of h_c on a WindCube lidar with hard-target returns is reported in (Allain et al., 2022); the spectrogram-level model is exercised by the super-resolution inversion of Martin et al. (2025).

8.1 Forward model versus Doppler estimator: a simulation pitfall

A common engineering practice is to validate a Doppler estimator (centroid, Gaussian fit, MLE) by generating a synthetic signal in which the wind field is convolved with a one-dimensional RWF—typically the pulse-only RWF or h_c. The estimator is then applied to this synthetic signal and its bias and variance are compared with the true wind.

This procedure is inconsistent when the estimator is not a pure centroid. The 1-D convolution $f_D * h_c$ produces a velocity profile, not a spectrogram: it has no spectral width, no multi-component structure, and no intra-gate broadening. A Gaussian-fit or MLE estimator, which operates on the local Doppler spectrum, is therefore tested against a forward model that does not

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

contain the very features it is designed to handle. The resulting performance figures are not representative of the estimator behaviour on real lidar returns.

The forward model that matches the data seen by a spectral estimator is the two-dimensional spectrogram of Sec. 3, generated either from (5) or from a time-domain sliced-atmosphere simulator (Salamitou et al., 1995; Valla, 2005). The same 2-D PSF serves as the forward operator in spectrogram-inversion approaches (Martin et al., 2025; Michel et al., 2022), where it acts as the convolution kernel to be inverted under a regularising prior. The 1-D reductions of Sec. 4 are valid simulation targets only when the lidar itself reports a centroid-type velocity and the regime assumptions of Table 1 hold.

9 Conclusions

The main point of this paper is simple: the expected spectrogram of a pulsed coherent Doppler wind lidar is a 2-D convolution, and every scalar range-weighting function in common use is a projection of the underlying 2-D PSF. Which projection is adequate depends on what is being estimated and on how the atmosphere behaves inside the range gate.

For routine velocity retrieval in a well-mixed boundary layer, the classical models work. They work not because they are general but because two conditions happen to be met: the spectrum is unimodal and the backscatter is smooth over the gate. When either condition fails, for example at a cloud edge, in a wake, or in strong shear, the one-dimensional picture breaks down and the hierarchy of Table 1 indicates which additional structure must be retained.

Perhaps the most practically relevant finding is that the window-completed kernel $h_c = \sum_i w_i^2 p_{k+i}$, not the pulse envelope alone, governs range averaging in all regimes where a 1-D reduction is valid.

For the processing parameters common in wind-energy applications ($N_w \simeq N_\tau$), ignoring the window overstates range resolution by 31% for the Gaussian pulse and Hann window tested here.

Code availability. The Python implementation of the sliced-atmosphere simulator, the closed-form kernel h_c (11), the two-dimensional PSF (6), and all numerical figures of this paper (`simu_papier_RWF_V5.py`) is available from the corresponding author upon reasonable request. The authors intend to deposit a citable, versioned release on Zenodo prior to journal submission and to update this statement with the corresponding DOI.

Data availability. No new experimental data sets were generated for this study. All range-weighting-function curves, point-spread functions, and velocity profiles shown are computed analytically or from the synthetic sliced-atmosphere simulator described in the "Code availability" section above. The experimental range-weighting-function measurements reproduced in Fig. 2 were originally published by Allain et al. (2022) and are available from that reference.

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

Appendix A: Detailed derivation of the 2-D PSF

This appendix expands the step from the signal model (1) to the expected spectrogram (5) and the PSF (6). The derivation follows the time-domain heterodyne model of (Salamitou et al., 1995; Valla, 2005); an equivalent continuous-time version was given in French by Martin et al. (2023).

A1 Expansion of the windowed periodogram

The signal model (1) is evaluated at the running time n; the windowed DFT in (4) reads $s_m[n+i]$ for $i = 0, \dots, N_w - 1$, so each occurrence of n in (1) is replaced by $n + i$. This gives, for a single pulse m,

$$X_m[n,q] = \sum_{i=0}^{N_w-1} \sum_{\ell=0}^{N_r-1} w_i \sqrt{p_\ell}\, a_m[n+i-\ell]\, e^{j2\pi f_D[n+i-\ell]\,(n+i)\,T_s}\, e^{-j2\pi q i/N_w} + N_m[n,q], \tag{A1}$$

where $N_m[n,q]$ collects the noise contribution. The periodogram is $\hat{S}[n,q] = M^{-1}\sum_m |X_m[n,q]|^2$. Writing the modulus as a product of the sum with its conjugate, $|z|^2 = z\,z^*$, produces a double sum over slices: the direct factor carries a slice index ℓ, the conjugate factor a second index ℓ'.

A2 Expectation and slice decorrelation

Only the speckle amplitudes a_m are random; the weights w, the pulse p and the exponentials are deterministic. The expectation of the double sum therefore reduces to that of the amplitude product. With $u = n + i - \ell$ and $v = n + i' - \ell'$, the decorrelation (2) reads

$$\mathbb{E}\{a_m[u]\, a_m^*[v]\} = \sigma_a^2[u]\, \delta_{u,v}. \tag{A2}$$

Two distinct slices scatter independently, so their cross-product averages to zero; the Kronecker $\delta_{u,v}$ forces $u = v$ and collapses the double slice sum into a single sum over u. This is the only step where the squared modulus is opened. The noise being independent of the signal and circular, its cross-terms vanish and $\mathbb{E}\{|N_m[n,q]|^2\} = \sigma_\eta^2$ for the normalised window.

A3 From one slice to the PSF

Introduce the range offset $k = n - u$. Keeping a single slice u (the sum over u and the noise term are reinstated at the end), its contribution is

$$\sigma_a^2[u] \left| \sum_i w_i \sqrt{p_{k+i}}\, e^{j2\pi (f_D[u]-f_q)\, i T_s}\, e^{j2\pi f_D[u]\, n T_s} \right|^2 .$$

The factor $e^{j2\pi f_D[u]\, n T_s}$ does not depend on the summation index i; it is constant in front of the sum and, using $|c\,z|^2 = |c|^2|z|^2$ with $|e^{j\theta}| = 1$, it drops out entirely. This removes the only dependence on n other than through $k = n - u$, leaving

$$\sigma_a^2[u] \left| \sum_i w_i \sqrt{p_{k+i}}\, e^{j2\pi (f_D[u]-f_q)\, i T_s} \right|^2 = \sigma_a^2[u]\, \mathrm{PSF}[n-u,\, f_D[u] - f_q],$$

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

with $\mathrm{PSF}[k, \Delta f]$ as in (6). Reinstating the slice sum and the noise variance,

$$\mathbb{E}\{\hat{S}[n,q]\} = \sum_u \sigma_a^2[u]\,\mathrm{PSF}[n-u,\, f_D[u]-f_q] + \sigma_\eta^2. \tag{A3}$$

Equation (A3) is exact and fully discrete; it is the core result of this appendix.

A4 Reading Eq. (A3) as a 2-D convolution

Equation (A3) is a weighted sum of shifted PSF copies, but the slice index u enters the two PSF arguments in two different ways, and this difference decides whether each axis is a convolution.

Range axis. Here u appears as the difference $n-u$. Writing out consecutive slices,

$$\cdots + \sigma_a^2[u{-}1]\,\mathrm{PSF}[n-u+1,\cdot] + \sigma_a^2[u]\,\mathrm{PSF}[n-u,\cdot] + \sigma_a^2[u{+}1]\,\mathrm{PSF}[n-u-1,\cdot] + \ldots,$$

each slice contributes the *same* kernel, rigidly shifted by one range sample from its neighbour and weighted by $\sigma_a^2[u]$. A weighted sum of integer-shifted copies of a fixed kernel is a discrete convolution. The range axis is a convolution exactly.

Frequency axis. Here u enters through $f_D[u]$ inside the argument $f_D[u]-f_q$. As $f_D[u]$ is an arbitrary sample of the velocity profile, it is not of the form (constant $-\,u$): the PSF is not shifted from slice to slice along frequency but *evaluated* at $f_D[u]-f_q$. Evaluation is not a convolution, so (A3) is not, as written, a 2-D convolution.

Operator form. A genuine 2-D operator is nonetheless what the deconvolution framework of Martin et al. (2025) inverts. It is obtained by defining the discrete range–frequency object on the DFT grid, the on-grid counterpart of (7),

$$o[u,q] = \sigma_a^2[u]\,\delta_{q,\mu_D[u]}, \qquad \mu_D[u] = \left\lfloor f_D[u]\,N_w T_s \right\rceil, \tag{A4}$$

a single nonzero bin per slice, and the discrete 2-D convolution

$$(\mathrm{PSF} *_{z,f} o)[n,f_q] = \sum_u \sum_q \mathrm{PSF}[n-u,\, f_q - f_{\mu_D[u]}]\,o[u,q] = \sigma_a^2[u]\,\mathrm{PSF}[n-u,\, f_q - f_{\mu_D[u]}].$$

The Kronecker $\delta_{q,\mu_D[u]}$ picks the single bin $\mu_D[u]$;

This matches (A3) once $f_{\mu_D[u]} \simeq f_D[u]$, the DFT bin nearest the Doppler line (the convolution sign convention on the frequency axis is absorbed into the definition of o). This condition is strong: an off-grid line is not a single bin but a Dirichlet kernel spread over all bins, and the single-bin replacement is accurate only for $|\varepsilon| \ll 1$, with $\varepsilon = f_D[u]N_wT_s - \mu_D[u] \in [-\frac{1}{2}, \frac{1}{2}]$, which an arbitrary velocity profile does not guarantee.

The single-bin object (A4) is therefore a notational convenience, not a requirement: keeping the PSF evaluated at the true $f_D[u]$ (as in (A3)) is exact for any ε and carries the multi-bin spread. In either reading (A3) is the operator form (8) used in the main text.

$$\mathbb{E}\{\hat{S}[n,q]\} - \sigma_\eta^2 = (\mathrm{PSF} *_{z,f} o)[n,f_q]. \tag{A5}$$

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

Where the line sits. The spectral line is the object (A4); it never appears inside a squared modulus, which was opened once and for all at the speckle step (A2). When (9) is broken (chirped pulse, Appendix A Sect. A5) the convolution rewriting fails and the exact evaluation form (A3) must be kept.

A5 Extension to a complex pulse envelope

A deterministic intra-pulse phase $\psi[\ell]$ is included by replacing $\sqrt{p_{k+i}}$ with $\sqrt{p_{k+i}}\,e^{j\psi[k+i]}$ in (6). The frequency marginal (11) is unchanged since $|e^{j\psi}| = 1$ inside the Parseval sum, but (9) is generally broken; this is the chirped-pulse regime of Sec. 7.4.

Appendix B: Detailed derivation of the 1-D reductions

This appendix derives (11), (14), (16) and (17) from the 2-D forward model (8).

B1 Frequency marginal: from PSF to h_c

Summing the PSF (6) over DFT bins and applying Parseval to the inner DFT—setting $x_i = w_i\sqrt{p_{k+i}}$:

$$\sum_{q=0}^{N_w-1}\left|\sum_i x_i\, e^{-j2\pi qi/N_w}\right|^2 = N_w\sum_i |x_i|^2 = N_w\sum_i w_i^2 p_{k+i}.$$

Dividing by N_w yields (11). The $1/N_w$ is absorbed by the N_w at the denominator of the CNR definition (13), so the downstream 1-D reductions (14)–(17) carry no N_w factor.

B2 CNR profile

Summing (5) over q, using the previous subsection, and substituting into (13) gives directly

$$\mathrm{CNR}[n] = \sum_u \frac{\sigma_a^2[u]}{\sigma_\eta^2}\, h_c[n-u] = (\mathrm{CNR}_a * h_c)[n],$$

which is (14).

B3 Centroid velocity, CNR-weighted form

Substitute (5) into (15). The numerator becomes

$$\sum_q f_q \sum_u \sigma_a^2[u]\,\mathrm{PSF}[n-u, f_D[u]-f_q].$$

The PSF carries the bin index q only through the offset $\Delta f = f_D[u] - f_q$. Switching summation variable from q to Δf (which leaves the sum range unchanged since $f_D[u]$ is a u-dependent shift) and writing $f_q = f_D[u] - \Delta f$ splits the numerator into two pieces:

$$\sum_u \sigma_a^2[u]\left[f_D[u]\sum_{\Delta f}\mathrm{PSF}[n-u,\Delta f] \;-\; \sum_{\Delta f}\Delta f\,\mathrm{PSF}[n-u,\Delta f]\right].$$

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

The second sum is the first moment of $\mathrm{PSF}[n-u,\cdot]$, which vanishes by the frequency symmetry (9); the first sum equals $N_w\, h_c[n-u]$ by the frequency-marginal identity above. The denominator follows the same algebra without the f_q factor, so N_w cancels in the ratio (16).

B4 Reduction to constant-CNR

When $\mathrm{CNR}_a[n-k]$ is approximately constant over the support of h_c, it factorises out of both sums in (16) and (17) follows.

B5 Pulse-only limit

When $N_w \ll N_\tau$, w_i^2 is sharply peaked on a single index; (11) then collapses to $h_c[k] \approx p_k$, i.e. (12), and (17) reduces to the classical pulse-only formula.

Author contributions. PEA conceived the unifying hierarchy of models, implemented the numerical verification, and prepared the manuscript. AHB derived the CNR-weighted centroid formula (Eq. 16) and the associated variable-backscatter validity analysis (Sect. 6), and contributed to the manuscript. DTM, TM, LM, and MV developed the two-dimensional sliced-atmosphere and point-spread-function framework on which this work builds and reviewed the manuscript. All authors discussed the results and commented on the manuscript.

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

Financial support. This research received no external grant funding; it was carried out as part of the authors' regular duties at Vaisala and ONERA/DOTA.

Acknowledgements. The authors thank their colleagues at Vaisala and ONERA/DOTA for helpful discussions during the preparation of this manuscript. As several authors are not native English speakers, AI-based language tools were used to help with the English wording. The derivations, the analysis, and the conclusions are entirely the authors' own.

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.

References

P. E. Allain, T. Martin, M. Valla, C. Benzo, V. Garbarini, D. T. Michel, L. Thobois, A. Dolfi-Bouteyre, and J.-P. Cariou, Comparison of analytical and numerical derivation of spatial resolutions of Coherent Pulsed Doppler Lidar and experimental direct measurement, in Proc. 21st Coherent Laser Radar Conf. (CLRC), 2022.

F. Bingöl, J. Mann, and G. C. Larsen, Light detection and ranging measurements of wake dynamics. Part I: one-dimensional scanning, Wind Energy, 13, 51–61, https://doi.org/10.1002/we.352, 2010.

A. Dabas, Semiempirical model for the reliability of a matched filter frequency estimator for Doppler lidar, J. Atmos. Oceanic Technol., 16, 19–28, https://doi.org/10.1175/1520-0426(1999)016<0019:SMFTRO>2.0.CO;2, 1999.

R. G. Frehlich, Coherent Doppler lidar signal covariance including wind shear and wind turbulence, Appl. Opt., 33, 6472–6481, https://doi.org/10.1364/AO.33.006472, 1994.

R. G. Frehlich, Effects of wind turbulence on coherent Doppler lidar performance, J. Atmos. Oceanic Technol., 14, 54–75, https://doi.org/10.1175/1520-0426(1997)014<0054:EOWTOC>2.0.CO;2, 1997.

R. G. Frehlich and L. Cornman, Coherent Doppler lidar signal spectrum with wind turbulence, Appl. Opt., 38, 7456–7466, https://doi.org/10.1364/AO.38.007456, 1999.

S. M. Hannon and J. A. Thomson, Aircraft wake vortex detection and measurement with pulsed solid-state coherent laser radar, J. Mod. Opt., 41, 2175–2196, 1994.

D. P. Held and J. Mann, Comparison of methods to derive radial wind speed from a continuous-wave coherent lidar Doppler spectrum, Atmos. Meas. Tech., 11, 6339–6350, https://doi.org/10.5194/amt-11-6339-2018, 2018.

F. Köpp, S. Rahm, and I. Smalikho, Characterization of aircraft wake vortices by 2-μm pulsed Doppler lidar, J. Atmos. Oceanic Technol., 21, 194–206, https://doi.org/10.1175/1520-0426(2004)021<0194:COAWVB>2.0.CO;2, 2004.

J. Mann, A. Peña, F. Bingöl, R. Wagner, and M. S. Courtney, Lidar scanning of momentum flux in and above the atmospheric surface layer, J. Atmos. Oceanic Technol., 27, 959–976, https://doi.org/10.1175/2010JTECHA1389.1, 2010.

T. Martin, M. Valla, L. Mugnier, P. E. Allain, A. Lessuise, and D. T. Michel, Super-résolution pour les lidars vent hétérodynes par approche inverse, in 29^{e} Colloque Francophone de Traitement du Signal et des Images (GRETSI), https://hal.science/hal-04297203, 2023.

T. Martin, L. M. Mugnier, M. Valla, P. E. Allain, and D. T. Michel, Bayesian approach for spatial super-resolution of heterodyne wind lidars, Opt. Lett., 50, 6445–6448, https://doi.org/10.1364/OL.566273, 2025.

D.-T. Michel, L. Mugnier, and M. Valla, Détection et télémétrie par impulsions de rayonnement électromagnétique, patent FR 3 110 002 B1, 2022.

A. Peña, J. Mann, and N. Dimitrov, Turbulence characterization from a forward-looking nacelle lidar, Wind Energ. Sci., 2, 133–152, https://doi.org/10.5194/wes-2-133-2017, 2017.

B. J. Rye and R. M. Hardesty, Discrete spectral peak estimation in incoherent backscatter heterodyne lidar. I. Spectral accumulation and the Cramer–Rao lower bound, IEEE Trans. Geosci. Remote Sens., 31, 16–27, https://doi.org/10.1109/36.210440, 1993a.

B. J. Rye and R. M. Hardesty, Discrete spectral peak estimation in incoherent backscatter heterodyne lidar. II. Correlogram accumulation, IEEE Trans. Geosci. Remote Sens., 31, 28–35, https://doi.org/10.1109/36.210441, 1993b.

P. Salamitou, A. Dabas, and P. H. Flamant, Simulation in the time domain for heterodyne coherent laser radar, Appl. Opt., 34, 499–506, https://doi.org/10.1364/AO.34.000499, 1995.

https://doi.org/10.5194/egusphere-2026-5090
Preprint. Discussion started: 19 November 2026
© Author(s) 2026. CC BY 4.0 License.



A. Sathe, J. Mann, J. Gottschall, and M. S. Courtney, Can wind lidars measure turbulence?, J. Atmos. Oceanic Technol., 28, 853–868, https://doi.org/10.1175/JTECH-D-10-05004.1, 2011.

I. N. Smalikho, Techniques of wind vector estimation from data measured with a scanning coherent Doppler lidar, J. Atmos. Oceanic Technol., 20, 276–291, https://doi.org/10.1175/1520-0426(2003)020<0276:TOWVEF>2.0.CO;2, 2003.

J.-J. Trujillo, F. Bingöl, G. C. Larsen, J. Mann, and M. Kühn, Light detection and ranging measurements of wake dynamics. Part II: two-dimensional scanning, Wind Energy, 14, 61–75, https://doi.org/10.1002/we.402, 2011.

M. Valla, Étude d'un lidar Doppler impulsionnel à laser Erbium fibré pour des mesures de champ de vent dans la couche limite de l'atmosphère, Ph.D. thesis, Télécom Paris (ENST), 2005.