



Polarimetric QPE with an X-band radar on the northwest coast of Peru: a ray-coherence diagnostic and a reliability-weighted specific-attenuation estimator

Ivan Herrera-Casas¹, Rodolfo Rodriguez¹, Carlos Del Castillo², Nilson Diaz¹, and Yamina Silva²

¹Estación Científica Ramón Mugica Martínez, Universidad de Piura, Piura, Peru

²Subdirección de Ciencias de la Atmosfera e Hidrosfera, Instituto Geofísico del Perú, Lima, Peru

Correspondence: Ivan Herrera-Casas (ivanhc.pp@gmail.com)

Abstract. The northwest coast of Peru is an arid region punctuated by extreme rainfall events, making accurate quantitative precipitation estimation (QPE) critical for hydro-meteorological risk management. This study develops, validates, and selects locally calibrated polarimetric QPE algorithms for a new dual-polarization X-band radar in Piura, Peru, using raindrop size distribution (DSD) data from the 2024 rainy season applied to radar data from the February–April 2025 season and evaluated against a network of rain gauges.

The manufacturer’s default algorithm severely overestimated total rainfall (percentage bias, PBIAS, of +25.6%), an error removed by the local calibration. Among the locally calibrated candidates, four variants of a specific-attenuation (A_H) algorithm were formulated around different data-selection thresholds for the total differential-phase span, $\Delta\Phi_{DP}$, that constrains the ZPHI retrieval. At the gauges these variants proved nearly indistinguishable, yet noise in the span estimate perturbs each ray by a nearly multiplicative factor, imprinting radial streaks that daily point accumulations cannot register. To measure this artifact directly from the radar data, we introduce a gauge-independent, scan-level diagnostic based on the azimuthal anomaly of $\ln R$, whose primary metric is the ray-coherent fraction of anomaly variance (coh). Applied to more than 26,000 scans, it reveals contamination in all precipitation regimes and under every threshold configuration, whereas a span-independent $R(Z_H, K_{DP})$ control sits at the noise floor. This motivates a reliability-weighted estimator that blends $R(A_H)$ with the local composite through a per-ray inverse-variance weight driven by the estimated noise of the $\Delta\Phi_{DP}$ constraint itself. The weighted estimator matches the best threshold variant at the gauges while reducing the ray-coherent contamination by factors of 2.4–2.7 in the moderate and heavy regimes, approaching the control floor without blurring the field.

1 Introduction

The northwest coast of Peru is characterized by a complex climate system, marked by a sharp transition from its typically arid state to periods of extreme precipitation. These heavy rainfall events, capable of exceeding the mean annual total in only a few days, are primarily linked to anomalous warming of the sea surface temperature (SST). This warming occurs either locally in the eastern Equatorial Pacific, driving Coastal El Niño events, or more broadly across the central-western Pacific



during the warm phase of the El Niño–Southern Oscillation (ENSO) (SENAMHI, 2016; Horel and Cornejo-Garrido, 1986; Rodríguez-Morata et al., 2019).

25 Research into this pronounced climatic variability has relied on data from sparsely distributed rain gauges and satellite-based products (Goldberg et al., 1987; Bendix, 2000; Takahashi, 2004; Douglas et al., 2009). However, the limited spatio-temporal resolution of these datasets has hindered a detailed understanding of the fine-scale processes governing the region’s convective systems.

The deployment of a weather radar to the region marked a significant observational advancement. A single-polarization radar
30 installed in 2019 improved spatial coverage, but its operational limitations – including a single-elevation scanning strategy and non-polarimetric measurements – posed substantial challenges for reliable Quantitative Precipitation Estimation (QPE). Specifically, signal attenuation affecting reflectivity (Z) and the absence of a locally derived reflectivity-rainfall rate ($Z - R$) relationship limited its accuracy. Attempts to mitigate these issues, such as applying empirical calibrations or Machine Learning models, have proven complex and often result in a loss of the native spatio-temporal resolution of the radar data (Rollenbeck
35 et al., 2021). To overcome these fundamental challenges, a new X-band dual-polarization radar was installed in the region in early 2025.

Dual-polarization radars have been successfully used to mitigate part of the sources of errors of radar rainfall estimation (Ryzhkov et al., 2022). Particularly, at X-band where attenuation is a major obstacle, the specific differential phase (K_{DP}),
immune to miscalibration, attenuation and partial beam blocking, has been used to correct the reflectivity and for direct esti-
40 mation of light-to-moderate and heavy precipitations (Park et al., 2005; Matrosov et al., 2014; Yu et al., 2023). Furthermore, algorithms based on specific attenuation (A), have also proven successful in operational settings, offering accuracy across diverse rain regimes and high spatial resolution (Testud et al., 2000; Ryzhkov et al., 2014; Diederich et al., 2015).

Despite the physical robustness of these polarimetric techniques, the quantitative relationships that link observables like K_{DP} and A to the rainfall rate (R) are not universal. These relationships are explicit functions of the underlying raindrop size
45 distribution (DSD) parameters. The DSD is known to exhibit significant variability across different climatological regimes and even between distinct storm types within the same region (Tapiador et al., 2010; Narayana Rao et al., 2018; Ibañez et al., 2023). The extreme climatic transitions on the northwest coast of Peru strongly suggest that relying on pre-established coefficients would introduce large errors. Therefore, locally derived DSD information is essential for unlocking the full potential of polarimetric QPE in this unique environment.

50 Attenuation-based estimation carries an additional, specifically practical difficulty. In the ZPHI family of retrievals (Testud et al., 2000), the specific attenuation along each ray is constrained by the total span of differential phase over the precipitation path, $\Delta\Phi_{DP}$, so the quality of an entire ray of the retrieved rain field rests on a single scalar: the difference between the Φ_{DP} values at the two ends of the path. The reliability of that difference is not uniform across rays. It degrades as the span shrinks, since a given error on the boundary values weighs ever more heavily in relative terms, and it degrades when the boundary gates
55 are themselves contaminated – Diederich et al. (2015) trace ray-scale anomalies in $R(A_H)$ to segments sampled next to clutter. The established response is a minimum-span criterion, below which the retrieval reverts to an attenuation-corrected $R(Z_H)$;



the limit is placed empirically at 4° at X band (Ryzhkov et al., 2014; Diederich et al., 2015). Whether a fixed minimum span suffices to keep the resulting errors out of the retrieved rain field, and how this can be established, are examined here.

This study leverages DSD measurements from a field campaign during the 2024 rainy season (January–April) to develop, evaluate, and select polarimetric QPE algorithms for the new radar. Sect. 2 describes the data and validation setup. Sect. 3 derives the local rainfall relationships and quantifies their improvement over the manufacturer defaults. Sect. 4 develops the attenuation-based algorithms and the strategies investigated for a reliable estimation of $\Delta\Phi_{DP}$. Sect. 5 validates the resulting algorithm family against the gauge network and establishes the limits of what a sparse daily network can resolve. Sect. 6 addresses that limit with a gauge-independent, scan-level diagnostic of the spatial artifacts caused by $\Delta\Phi_{DP}$ noise, and Sect. 7 builds on the diagnostic to construct a reliability-weighted estimator and to compare the full method family over the season. Sect. 8 discusses operational implications and limitations.

2 Data and study area

2.1 Study area and observational network

This study focuses on the Piura region on the northwest coast of Peru (Fig. 1), an area characterized by a predominantly arid climate. The region's precipitation regime is heavily influenced by SST anomalies in the Niño 1+2 region and is subject to extreme convective rainfall during its rainy season, which typically occurs from December to April.

A critical component of this study's methodology is the use of DSD data from the 2024 rainy season to calibrate radar data from the 2025 rainy season. The validity of this approach is supported by the consistent climatic conditions across both periods. According to the Coastal El Niño Index (ICEN), based on the SST in the Niño 1+2 region, both the January–April 2024 and February–April 2025 periods were characterized by weak El Niño conditions (ENFEN, 2024), ensuring a comparable microphysical environment.

The observational network for this study comprises three main components: an X-band polarimetric radar, a laser disdrometer deployed during a 2024 field campaign, and a network of rain gauges used for validation.

2.2 Radar data

The primary radar dataset was collected by a Furuno WR-2120 dual-polarization X-band radar, located at the University of Piura (UDEP; 5.1707° S, 80.639° W, 52 m.a.s.l.). The radar operated continuously from February to April 2025, employing a scanning strategy of eight elevation angles ($2.5, 3.5, 4.5, 5.5, 7.5, 10.0, 12.5$, and 15.0°) completed every two minutes. Key technical specifications of the radar are listed in Table 1.

Before analysis, radar data, specifically reflectivity (Z_H), differential phase (Φ_{DP}), and specific differential phase (K_{DP}), underwent a quality control (QC) procedure, implemented in Python using the Python ARM Radar Toolkit (Helmus and Collis, 2016). The system applies its own noise filtering before the data are recorded, so gates of low signal-to-noise ratio are already absent from the archived fields. First, non-meteorological echoes were filtered using a threshold on the co-polar correlation




Figure 1. Topographic map of the study area in the Piura region on the northwest coast of Peru. The location of the WR-2120 X-band polarimetric radar at the University of Piura is marked by a magenta star, with its 70 km operational range indicated by a dashed magenta ring. The location of the Parsivel disdrometer is shown by the blue square. The locations of the rain gauges are also indicated with black triangles, with two co-located at the radar and disdrometer sites. Elevation data are from the SRTM global digital elevation model (NASA Shuttle Radar Topography Mission, 2013).

coefficient ($\rho_{hv} < 0.8$). Next, a custom algorithm was applied to remove remaining non-meteorological radial spikes. This filter identifies and removes any contiguous echo region with an azimuthal extent of less than 10° that begins within 10 km of the radar. The filter is tuned to the interference artifacts present in this dataset. As all measurements from the 2.5° elevation angle are below an altitude of 3.5 km, no additional considerations were made regarding the melting layer, which is expected to be above 4 km.

2.3 Disdrometer data and processing

Drop Size Distribution (DSD) data were collected with an optical disdrometer, Parsivel, during a field campaign led by the Geophysical Institute of Peru (IGP). The instrument was deployed in Chulucanas (5.1084° S, 80.169° W) from January to April 2024, recording DSDs at a one-minute sampling interval.



Table 1. Technical specifications of the Furuno WR-2120 radar.

Parameter	Value
Manufacturer	Furuno (Japan)
Model	WR-2120
Frequency	9.4 GHz
Polarization	Dual-polarimetric
Operational range	70 km
Radial resolution	75 m
Azimuthal resolution	0.5°
Amplifier	Solid-state module
Antenna gain	≥ 33 dB
Beamwidth	2.7°
Peak power	100 W
Max PRF	2000 Hz
Max pulse width	50 μ s

After excluding drops with fall speeds deviating more than 30 % from the terminal velocity relationship of Atlas and Ulbrich (1977), a final dataset of 1,917 DSDs with non-zero rain rates was used to simulate polarimetric radar variables using the T-matrix scattering code implemented in the PyTMatrix Python library (Mishchenko et al., 1996; Leinonen, 2014). The simulations were configured for the radar’s frequency (9.4 GHz) and assumed the following: (1) a temperature of 20 ° C; (2) a Gaussian canting angle distribution with a mean of 0° and standard deviation of 10°; and (3) a composite drop axis ratio model using Brandes et al. (2004) for drops $D < 1.5$ mm and Thurai et al. (2007) for drops $D > 1.5$ mm.

2.4 Rain gauge network and validation setup

Validation of the radar-based QPE products was performed against a network of rain gauges operated by the National Meteorological and Hydrological Service of Peru (SENAMHI) and by UDEP. Gauge sites were selected based on two criteria: (1) data availability of at least 85 % during the February–April 2025 study period; and (2) suitability for rainfall measurement, which involved excluding stations sited for hydrological purposes at river level.

This selection process yielded a final validation network of 12 gauges within the radar range (10 from SENAMHI and 2 from UDEP), whose locations are shown in Fig. 1. Because the majority of these stations provide data at a daily resolution, all radar-based precipitation estimates were accumulated to daily totals to facilitate a direct comparison. For each gauge location, radar estimates were extracted from the data of the lowest elevation scan (2.5°). A special consideration was made for the gauge co-located with the radar; since the QC procedure removed the gates directly overhead, radar data were instead extracted from a representative point approximately 600 m away.



115 Additionally, for spatial comparison of accumulated rainfall, 11 additional stations were selected within a square domain of ± 100 km around the radar, following the same criteria. The accumulated reference map was calculated using universal kriging interpolation with a linear variogram model following Rollenbeck et al. (2021).

The performance of the radar-based QPE algorithms was evaluated against rain gauge observations using four standard statistical metrics: the Pearson Correlation Coefficient (CC), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Percentage Bias (PBIAS). For a set of N radar-gauge data pairs, where R_i is the radar estimate and G_i is the corresponding gauge measurement, and $\bar{R}$ and $\bar{G}$ are the respective mean values, these metrics are defined as:

$$CC = \frac{\sum_{i=1}^N (R_i - \bar{R})(G_i - \bar{G})}{\sqrt{\sum_{i=1}^N (R_i - \bar{R})^2 \sum_{i=1}^N (G_i - \bar{G})^2}} \quad (1)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (R_i - G_i)^2} \quad (2)$$

$$MAE = \frac{1}{N} \sum_{i=1}^N |R_i - G_i| \quad (3)$$

$$PBIAS = \frac{\sum_{i=1}^N (R_i - G_i)}{\sum_{i=1}^N G_i} \times 100\% \quad (4)$$

125 Uncertainty in these metrics was quantified with a day-level block bootstrap: entire days ($N_{\text{days}} = 53$; $N_{\text{pairs}} = 202$) were resampled with replacement, preserving the spatial correlation of the radar-gauge pairs within each day, and the metrics were recomputed for every replicate ($B = 1,000$). The same resampled days were used for all algorithms in each replicate, so that differences between any two methods can be evaluated on paired replicates.

3 Local polarimetric relationships and default-vs-local calibration

130 3.1 Derivation of local polarimetric rainfall relationships

The formulation of the evaluated QPE algorithms relies on standard power-law relationships between rainfall rate (R) and polarimetric variables: $R(Z_H) = aZ_H^b$, $R(K_{DP}) = aK_{DP}^b$, and $R(A_H) = aA_H^b$, where a and b are empirically derived parameters for each relation. The coefficients for the locally derived rainfall relationships, obtained via linear regression on log-transformed data from the 2024 DSD dataset, are presented in Table 2 alongside the manufacturer's default values. The corresponding scatter plots and fitted lines are shown in Fig. 2. Each relationship was fitted over the range in which it is expected to be used operationally: $R(Z_H)$ for $Z_H < 40$ dBZ, $R(K_{DP})$ for $K_{DP} > 0.3^\circ \text{ km}^{-1}$, which is also the threshold below which the composite logic of Sect. 3.2 does not invoke it, and $R(A_H)$ for $A_H > 0.05 \text{ dB km}^{-1}$. Below this last value the simulated pairs suggest a somewhat different slope, so the fitted power law is best read as an approximation tuned to the moderate-to-heavy regime that dominates the seasonal accumulation. Since the fallback to $R(Z_H)$ described in Sect. 4.2 is applied per ray rather



140 than per gate, the relation is in practice extrapolated slightly below its fitting range at the weaker gates of the rays that are retained.

In addition to the rainfall rate estimators, the DSD dataset was used to derive the two auxiliary parameters required by the specific-attenuation retrieval of Sect. 4: the coefficient α for the linear relationship between specific attenuation and specific differential phase ($A_H = \alpha K_{DP}$), which was found to be 0.26 dB deg^{-1} , and the exponent b for the relationship between linear
145 reflectivity and specific attenuation ($A_H = cZ_H^b$), determined to be 0.74. All parameters consumed by the algorithms evaluated in this study are collected in Table 2.

Table 2. Coefficients for rainfall relationships ($R = aX^b$) for manufacturer (default) and locally derived parameters, together with the auxiliary parameters of the specific-attenuation retrieval. Local relationships were derived from DSDs collected during the 2024 field campaign.

Relationship	Form ($R = aX^b$)	Default		Local	
		a	b	a	b
$R(Z_H)$	$R = aZ_H^b$	0.0365	0.625	0.0531	0.594
$R(K_{DP})$	$R = aK_{DP}^b$	23.52	0.815	16.22	0.795
$R(A_H)$	$R = aA_H^b$	-	-	41.03	0.698
$A_H = \alpha K_{DP}$		-		$\alpha = 0.26 \text{ dB deg}^{-1}$	
$A_H = cZ_H^b$ (ZPHI exponent)		-		$b = 0.74$	




Figure 2. Scatter plots of rain rate versus polarimetric variables simulated from the 2024 disdrometer DSD data (blue dots). Panels show relationships for (a) horizontal reflectivity (Z_H), (b) specific differential phase (K_{DP}), and (c) specific attenuation (A_H). The locally derived power-law fits are shown as solid red lines, and the manufacturer's default relationships are shown as dashed gray lines.

It is observed from Fig. 2 that the default $R(K_{DP})$ relationship (Fig. 2b) results in a substantial overestimation of rainfall, highlighting a significant regional bias in the default parameters. Within the fitted ranges the scatter is visually smallest for $R(K_{DP})$ and largest for $R(Z_H)$, with $R(A_H)$ intermediate; the $R(A_H)$ pairs are in fact most tightly grouped at the lower end



150 of the range, below the interval used for the fit, where the retrieval nevertheless relies on $R(Z_H)$ because the differential-phase constraint, and not the relation itself, becomes unreliable there (Sect. 4.2). The scatter of the $R(A_H)$ relation is expected to be more apparent at X band than at longer wavelengths: the DSD sensitivity of $R(A)$ is near its minimum at S band, where A -based estimation is notably robust to DSD variability, and increases toward shorter wavelengths (Ryzhkov et al., 2014). Part of the remaining dispersion could in principle be absorbed by adding Z_{DR} as a second predictor, but this was not pursued here, 155 as it would reintroduce the sensitivity to absolute calibration and differential-attenuation correction that A -based estimation is meant to avoid. Given the single-season, single-site nature of the campaign, the coefficients of Table 2 are presented as a first local approximation, not as a definitive characterization of the regional DSD; the diagnostic and estimator results of Sects. 6 and 7 concern the structure of the retrieval and are not sensitive to moderate changes in these values.

3.2 The composite algorithms: default and local

160 The first two algorithms evaluated in this study isolate the effect of local calibration on the manufacturer's product logic:

- $R(Z_H, K_{DP})_{DEF}$: The manufacturer's default composite algorithm. This product uses a standard Marshall–Palmer $R(Z_H)$ relationship for light precipitation and switches to an $R(K_{DP})$ relationship when $Z_H > 30$ dBZ and $K_{DP} > 0.3^\circ \text{ km}^{-1}$. The K_{DP} values used are those provided by the manufacturer's software.
- $R(Z_H, K_{DP})_{LOC}$: The locally calibrated version of the $R(Z_H, K_{DP})_{DEF}$ composite described above. It uses the same 165 switching logic but employs the specific $R(Z_H)$ and $R(K_{DP})$ relationships derived from the 2024 disdrometer campaign.

The DSD-level overestimation of the default $R(K_{DP})$ relationship (Fig. 2b) carries through to the gauges (Sect. 5): the default composite overestimates seasonal rainfall by +25.6%, and replacing its coefficients with the local ones – within the identical product logic – reduces the bias to +9.94%. Local calibration thus removes the dominant, DSD-attributable error at its source. The composite logic itself, however, is not the only way to exploit the polarimetric measurements; Sect. 4 develops 170 the attenuation-based alternative.

4 Attenuation-based estimation

4.1 Specific-attenuation retrieval and the $\Delta\Phi_{DP}$ estimation problem

Attenuation-based QPE is attractive at X band for well-established reasons: derived from the path-integrated differential phase, $R(A_H)$ relationships are inherently immune to radar miscalibration and partial beam blockage, provide an explicit correction 175 for signal attenuation, and preserve the radar's native spatial resolution (Ryzhkov et al., 2014; Diederich et al., 2015). The specific attenuation A_H is calculated with the ZPHI rain-profiling algorithm (Testud et al., 2000), following the formulation of Ryzhkov et al. (2014):

$$A_H(r) = \frac{Z_H^b(r) C(b, \Delta\Phi_{DP})}{I(r_1, r_2) + C(b, \Delta\Phi_{DP}) I(r, r_2)}, \quad (5)$$



with

$$180 \quad I(r, r_2) = 0.46b \int_r^{r_2} Z_H^b(s) ds, \quad (6)$$

$$C(b, \Delta\Phi_{DP}) = 10^{0.1b\alpha\Delta\Phi_{DP}} - 1, \quad (7)$$

where Z_H is the measured (attenuated) linear reflectivity, $[r_1, r_2]$ is the precipitation path along the ray, $b = 0.74$ and $\alpha = 0.26 \text{ dB deg}^{-1}$ are the locally derived parameters of Table 2, and $\Delta\Phi_{DP} = \Phi_{DP}(r_2) - \Phi_{DP}(r_1)$ is the total span of differential phase over the path. The structure of Eqs. (5)–(7) makes the central dependency of the method explicit: the entire radial profile
185 of A_H , and hence of $R(A_H)$, is scaled by a single per-ray constraint, $\Delta\Phi_{DP}$.

For the span calculation, the Φ_{DP} field is first smoothed along range with a median filter with a window of 15 gates, which suppresses the gate-to-gate fluctuations of the raw phase. Even so, direct estimation of $\Delta\Phi_{DP}$ from the smoothed profile produces physically untenable results: strongly overestimated rain rates, most severely in light precipitation, caused by an unstable phase signal in weak echo regions or at the boundaries of storms – instabilities that are systematic over segments
190 much longer than the filter window and therefore survive the smoothing. For example, Fig. 3 shows the profile of a light rain event, where in the first segment the smoothed Φ_{DP} shows a pronounced increase, inconsistent with the low reflectivity values, below 25 dBZ; together with an erratic “jump” at the end of the path, this would lead to a $\Delta\Phi_{DP}$ value exceeding 15° if a direct calculation were applied, causing a high overestimation of the specific attenuation and rainfall rate.




Figure 3. Radial profiles of key polarimetric variables for a light rain event on 5 February 2025 at 10:52 UTC (Azimuth: 68.0°). Panels show: (a) Raw (gray) and smoothed (black; 15-gate median filter) differential phase (Φ_{DP}) ; (b) Horizontal reflectivity (Z_H); and (c) Co-polar correlation coefficient (ρ_{HV}). The vertical dashed lines indicate the path boundaries defined by a reflectivity threshold ($Z_H > 25 \text{ dBZ}$; gray) and a stricter threshold based on high correlation ($\rho_{HV} > 0.95$; orange). Additional case studies illustrating the behavior of the filtering strategies in moderate and heavy precipitation are provided in Appendix A.



4.2 Data-selection strategies for $\Delta\Phi_{DP}$

195 To address this challenge, three data-filtering strategies were investigated to identify the set of reliable gates used for the $\Delta\Phi_{DP}$ calculation, giving rise to four variants of the attenuation-based algorithm:

- $R(Z_H, A_H)_{BASE}$: The baseline configuration. The total span of differential phase, $\Delta\Phi_{DP}$, is calculated from the smoothed Φ_{DP} profile using the median of the first and last three gates with $Z_H > 25$ dBZ. This value, corresponding to a local rain rate of approximately 1.6 mm h^{-1} , discards gates of low signal-to-noise ratio while retaining most of the precipitation field. For rays with a total differential phase $\Delta\Phi_{DP} < 4^\circ$, it uses the local $R(Z_H)$ relationship, as suggested by Diederich et al. (2015).
- $R(Z_H, A_H)_{Z>30}$: Same as the baseline, but the reflectivity threshold for calculating $\Delta\Phi_{DP}$ is raised to $Z_H > 30$ dBZ (about 3.2 mm h^{-1}).
- $R(Z_H, A_H)_{\rho_{hv}>0.95}$: Same as the baseline, but a high co-polar correlation coefficient threshold of $\rho_{hv} > 0.95$ is also
205 required for gates used in the $\Delta\Phi_{DP}$ calculation, a physically based criterion that aims to identify regions of uniform rainfall where the differential phase signal should be most coherent.
- $R(Z_H, A_H)_{STRICT}$: Same as the baseline, but it uses the local $R(Z_H)$ relationship for rays with a total differential phase $\Delta\Phi_{DP} < 10^\circ$.

These thresholds do not define the rain path; they identify a reliable sub-path over which the total differential-phase span
210 $\Delta\Phi_{DP}$ can be estimated stably, and that span then constrains the attenuation over the full precipitation path as determined by the initial QC. Specific attenuation and rainfall rate are therefore calculated for every meteorological gate within the full path, whether or not it met the stricter thresholds.

The behavior of the three filtering strategies across precipitation regimes, examined through radial case studies (Fig. 3 and Appendix A), can be summarized as follows. In light rain, the moderate $Z_H > 25$ dBZ threshold is sometimes insufficient to
215 reject instability at the far boundary, while the stricter $Z_H > 30$ dBZ threshold is more robust but frequently yields a span too small ($< 4^\circ$) to be considered reliable, forcing the algorithm to revert to $R(Z_H)$. In moderate precipitation, both stricter strategies mitigate the boundary jumps, but the $Z_H > 30$ dBZ threshold tends to enclose a shorter segment than the $\rho_{hv} > 0.95$ filter, which can bias the span in either direction depending on the phase behavior of the excluded segments. In heavy precipitation, the co-polar correlation coefficient was observed to be unstable, often dropping quickly below the 0.95 threshold even within
220 valid rain cores and leading to a significant underestimation of the path-integrated attenuation, whereas the reflectivity-based thresholds consistently identify start and end gate windows that span the entire high-reflectivity region.

That no single filtering threshold is robust in every case study motivates the conservative $R(Z_H, A_H)_{STRICT}$ variant, which relies on the more stable $R(Z_H)$ relationship for a much wider range of light-to-moderate precipitation events where the phase signal is inherently less reliable. In the radial case studies the boundary instabilities appear largely confined to light
225 precipitation, evidence that favors STRICT as the most promising configuration. Sect. 5 evaluates the full algorithm family against the gauge network.



5 Gauge validation of the algorithm family

Figure 4 shows the daily radar-gauge scatter for the six algorithms formulated so far – the two composites and the four attenuation variants – and Table 3 collects the corresponding metrics with their bootstrap confidence intervals.




Figure 4. Scatter plots of daily accumulated rainfall from six QPE algorithms versus corresponding gauge measurements for the February–April 2025 study period, on logarithmic axes: (a) $R(Z_H, K_{DP})_{DEF}$, (b) $R(Z_H, K_{DP})_{LOC}$, (c) $R(Z_H, A_H)_{BASE}$, (d) $R(Z_H, A_H)_{STRICT}$, (e) $R(Z_H, A_H)_{Z>30}$, and (f) $R(Z_H, A_H)_{\rho_{hv}>0.95}$. Data points are color-coded to distinguish between gauges located within 60 km of the radar (blue) and those at farther distances (orange). Point metrics are provided in each panel; bootstrap confidence intervals are given in Table 3.

230 The validation results for the composite algorithms align with the DSD-based analysis of Sect. 3.1. As anticipated from the significant overestimation of the default $R(K_{DP})$ relationship (Fig. 2b), $R(Z_H, K_{DP})_{DEF}$ exhibits a large positive bias of +25.6%, which the stratified distributions of Fig. 5 show to be driven by the overestimation of moderate and heavy rainfall. Applying the locally derived DSD parameters within the identical product logic reduced the PBIAS to +9.94%, with essentially



Table 3. Gauge validation metrics for the February–April 2025 study period ($N_{\text{pairs}} = 202$ over $N_{\text{days}} = 53$), with 95 % day-block bootstrap confidence intervals on PBIAS.

Method	PBIAS (%) [95 % CI]	MAE (mm)	RMSE (mm)	CC
$R(Z_H, K_{DP})_{\text{DEF}}$	+25.63 [+11.39, +41.98]	5.55	9.63	0.836
$R(Z_H, K_{DP})_{\text{LOC}}$	+9.94 [−2.44, +25.70]	4.89	8.50	0.833
$R(Z_H, A_H)_{\text{BASE}}$	+2.59 [−9.22, +19.19]	4.88	8.69	0.838
$R(Z_H, A_H)_{\text{STRICT}}$	−1.10 [−12.66, +14.19]	4.69	8.70	0.838
$R(Z_H, A_H)_{Z>30}$	−6.88 [−17.45, +7.60]	4.64	8.90	0.838
$R(Z_H, A_H)_{\rho_{hv}>0.95}$	−4.33 [−16.48, +11.71]	4.93	9.53	0.804

unchanged correlation (CC = 0.833 versus 0.836), although a tendency to overestimate rainfall in the moderate intensity category (10–40 mm day^{−1}) persists.

The attenuation-based variants improve the calibration further. The baseline reduces the PBIAS to +2.59 %, and the STRICT variant $R(Z_H, A_H)_{\text{STRICT}}$ achieves the nominally best combination of metrics: PBIAS of −1.10 % with the lowest MAE and among the highest correlations. The $\rho_{hv} > 0.95$ variant produced the weakest correlation and among the highest errors of the family, consistent with the instability of ρ_{hv} inside heavy rain cores described in Sect. 4.2. Judged by Table 3 alone, STRICT is the most robust estimator, consistent with the radial case studies of Sect. 4.2. Beyond the clear rejection of the default, however, the differences among the locally calibrated candidates are small – a PBIAS range of about four percentage points, MAE within 0.2 mm, CC within 0.005 – and well within the uncertainty of what 12 daily gauges over a single season can resolve, as the confidence intervals in Table 3 indicate.

Two systematic residuals are shared by all the local methods (Fig. 5). In light precipitation, a consistent overestimation persists; because it survives even with a locally derived $R(Z_H)$ relationship, the cause is more likely physical or instrumental – in this arid environment, processes like virga and raindrop evaporation below the radar beam are significant possibilities, with subtle radar miscalibration or errors in manual gauge measurements as further candidates. In heavy precipitation, an underestimation concentrated at gauges beyond 60 km (orange points in Fig. 4) reflects the 2.5° beam overshooting the core of intense precipitation or sampling weaker rainfall aloft, as the beam samples at nearly 3 km altitude or higher at those distances.

The near-identity of the scatter plots in Fig. 4c–f measures the resolving power of the gauge network, not the equivalence of the variants. The four attenuation variants were formulated precisely because unreliable $\Delta\Phi_{DP}$ estimates corrupt the retrieval, and their thresholds intervene directly on that mechanism – yet nothing in the gauge statistics registers those interventions. Examination of individual scans shows that the instantaneous rain fields of the attenuation variants exhibit radial streaks: entire rays scaled up or down relative to their neighbors. Inspection of the candidate fields as rain-rate PPIs shows the streaks only in light rain, the same regime in which the unconstrained span calculation of Sect. 4.1 produced its grossest overestimates. A daily accumulation at a gauge point integrates over hundreds of scans and, at each of them, a different realization of the ray errors; much of the structure cancels, and what survives is diluted below the resolving power of a 12-gauge network. Point-daily




Figure 5. Box plots showing the distribution of daily rainfall values for the gauges (reference) and the QPE algorithms, stratified by gauge-measured daily rainfall intensity: light (bottom panel, > 1 mm), moderate (middle panel, > 10 mm), and heavy (top panel, > 40 mm) events. For each category, N represents the number of events, and S indicates the percentage of the total accumulated rainfall that these events contribute.

statistics are, in other words, structurally insensitive to the very error mode the variants were designed to suppress. Whether the streaks are in fact confined to light rain, as these displays suggest, and how much each variant actually suppresses them, are questions the gauges cannot answer – and answering them requires measuring the artifact directly from the radar fields. Sect. 6 develops such a measurement.

6 The ray-coherence diagnostic

6.1 Ray-coherent errors and their signature in $\ln R$

The error structure that span noise imposes on the retrieved rain field follows directly from Eqs. (5)–(7). A perturbation $\delta(\Delta\Phi_{DP})$ of the span estimate perturbs the constraint C for the entire ray at once, while the precipitation path $[r_1, r_2]$ and the measured reflectivity profile are unchanged. Writing $u \equiv 0.1 \ln 10 b \alpha \Delta\Phi_{DP}$, so that $C = e^u - 1$, direct differentiation of Eq. (5) gives the logarithmic sensitivity of the retrieval to the span,



$$S(r) \equiv \frac{\partial \ln A_H(r)}{\partial \ln \Delta \Phi_{DP}} = \frac{u e^u}{e^u - 1} \cdot \frac{I(r_1, r_2)}{I(r_1, r_2) + C I(r, r_2)}, \quad (8)$$

which holds exactly for any reflectivity profile. Because $I(r, r_2)$ decreases monotonically from $I(r_1, r_2)$ to zero along the ray, $S(r)$ increases monotonically from $u/(e^u - 1)$ at the path start to $u e^u/(e^u - 1)$ at its end. In the low-attenuation limit $u \ll 1$, the expansions of these endpoint values confine $S(r)$ to $1 \pm u/2 + \mathcal{O}(u^2)$ uniformly along the ray: the retrieved attenuation scales as $A_H \propto \Delta \Phi_{DP}$, and the fractional error of A_H equals the fractional error of the span at every gate. When attenuation along the path is substantial the sensitivity is genuinely range-dependent – below unity near the path start and above it near the end – but the total it redistributes is rigidly constrained: by construction of the ZPHI retrieval, $\int_{r_1}^{r_2} A_H dr$ is exactly linear in $\Delta \Phi_{DP}$ for any reflectivity profile, and equal to $\alpha \Delta \Phi_{DP}/2$ up to the rounding of the constant in Eq. (6) (Testud et al., 2000). A fractional span error therefore rescales the path-integrated attenuation in exact proportion, and $S(r)$ merely redistributes that fixed total about unity, damping the per-gate error near the path start and amplifying it near the end, where $S = u e^u/(e^u - 1)$ attains its maximum and so bounds the per-gate error – 1.24 at $\Delta \Phi_{DP} = 10^\circ$, and 1.51 at 20° . Equivalently, differentiating the constraint with respect to $\ln \Delta \Phi_{DP}$ gives $\int_{r_1}^{r_2} A_H S dr = \int_{r_1}^{r_2} A_H dr$: the A_H -weighted mean of $S(r)$ is exactly one for any profile and any attenuation, independently of that constant.

Propagating through the power law $R = a A_H^{b_R}$ with $b_R = 0.698$ (Table 2), and writing $c(r) \equiv R_{\text{est}}(r)/R_{\text{true}}(r)$ for the multiplicative rain-rate error, linearization in the perturbation gives the log error at each gate as

$$\ln c(r) = b_R S(r) \frac{\delta(\Delta \Phi_{DP})}{\Delta \Phi_{DP}} + \mathcal{O} \left[\left(\frac{\delta(\Delta \Phi_{DP})}{\Delta \Phi_{DP}} \right)^2 \right],$$

with $S(r)$ given exactly by Eq. (8). Two successive truncations now separate cleanly. First, since $S(r) = 1 + \mathcal{O}(u)$ uniformly along the ray, setting $S \approx 1$ yields, to leading order in the perturbation and in the span,

$$R_{\text{est}}(r) \approx c R_{\text{true}}(r), \quad \ln c \approx b_R \frac{\delta(\Delta \Phi_{DP})}{\Delta \Phi_{DP}}, \quad (9)$$

i.e., a nearly multiplicative, ray-coherent rain-rate error whose logarithm is a constant offset along the ray. Second, the accuracy of this ray-constant idealization is controlled by the span itself: $S(r)$ runs monotonically between the endpoint values $u/(e^u - 1)$ and $u e^u/(e^u - 1)$, whose difference is exactly u , so the along-ray spread of the linearized log error is, independently of the reflectivity profile,

$$\max_r \ln c(r) - \min_r \ln c(r) = b_R u \frac{\delta(\Delta \Phi_{DP})}{\Delta \Phi_{DP}}, \quad (10)$$

so the range-dependent remainder is a fraction u of the offset itself. For the spans that dominate light rain ($\Delta \Phi_{DP} \lesssim 5^\circ$, $u \lesssim 0.22$; Fig. 7b) the error is ray-constant to good approximation, the along-ray spread being about a fifth of the offset at most. At large spans the remainder becomes a substantial fraction of the offset, but it remains a smooth, monotone modulation



295 along the ray: distributed over hundreds of gates, it is far too gradual to register in gate-to-gate differences. Span errors of either component are therefore confined, for practical purposes, to the azimuthal structure of the field – a property the diagnostic below exploits from both sides, detecting the error azimuthally while verifying along range that nothing else has been altered.

Equations (9)–(10) have three consequences that shape the diagnostic. First, the natural coordinate for detecting the artifact is $L = \ln R$, in which the error is additive and, to leading order, ray-constant. Second, the offset is largest precisely where the rain is lightest: by Eq. (9) it scales with the inverse span, so for a given level of endpoint phase noise the fractional rain-rate error grows as $\Delta\Phi_{DP}$ shrinks. Rain-rate displays with discrete class palettes – the standard operational rendering, and the one used here to inspect the candidate fields – therefore show the artifact almost exclusively in light rain, where the offset is large enough to move a ray across a class boundary; in moderate and heavy rain the same contamination is present but perceptually masked (Appendix B1). In the azimuthal anomaly of $\ln R$, by contrast, the offsets of every regime appear as radial streaks. Third, the offset is confined to a single ray, i.e., to azimuthal scales below the 2.7° beamwidth – scales at which genuine precipitation features cannot appear, because the beam smears any real structure across several adjacent rays. A sub-beam-scale, ray-constant anomaly is therefore non-meteorological by construction, and this is the property that the coherence metric of Sect. 6.2 isolates. Along range, conversely, the ray-constant offset cancels exactly in gate-to-gate differences, so an along-range roughness metric is blind to the artifact itself and can serve as an independent guard against spatial smoothing.

310 Throughout this analysis, $\delta(\Delta\Phi_{DP})$ denotes any error in the span scalar consumed by Eq. (5), irrespective of its origin. The variants of Sect. 4.2 differ only in where the span is read, so endpoint phase noise, unstable-phase segments at the path boundaries, sub-path truncation, and backscatter differential phase all enter through the same scalar and propagate through Eqs. (8)–(10) identically. Only the azimuthal structure of the aggregate error matters downstream. Its ray-to-ray fluctuating part produces the flickering streaks that the diagnostic of Sect. 6.2 measures and that daily accumulations largely cancel; its azimuthally coherent part – a miscalibration of α , say, which enters only through the product $\alpha\Delta\Phi_{DP}$ – shifts neighboring rays together, is absorbed into the smooth background of the anomaly decomposition, and surfaces instead at the gauges as bias. The two validation instruments of this study are in this sense complementary: gauge PBIAS responds to the mean of the span error, the ray-coherence diagnostic to its ray-to-ray variance.

6.2 Scan-level metrics and scoring design

320 For every scan and every method, metrics are computed on the field $L = \ln R$, restricted to a common scoring mask: gates that pass the QC of Sect. 2.2 and have $R \geq R_{\min}$, with $R_{\min} = 0.1 \text{ mm h}^{-1}$ acting as a log floor; only scans with more than 1000 valid gates are included. All methods are evaluated on the identical mask within each scan, so every comparison below is paired at the scan level.

To isolate ray-coherent behavior, a smooth azimuthal background is first removed. At each range gate, a wrap-aware running median across the $2w+1 = 5$ neighboring rays defines L_{smooth} , and the azimuthal anomaly is $L' = L - L_{\text{smooth}}$. Real azimuthal features wider than the beam survive in the smooth component; sub-beam-scale excursions, including the ray-constant offsets of Eq. (9), are captured by L' . For each ray i , the per-ray offset is estimated robustly as $o_i = \text{median}_j L'(i, j)$ over its n_i valid gates. Two complementary quantities are then formed. The primary metric is the ray-coherent fraction of the anomaly variance,



$$\text{coh} = \frac{\sum_i n_i o_i^2}{\sum_{i,j} L'(i,j)^2}, \quad (11)$$

330 which approaches 1 when the anomaly field is dominated by ray-constant offsets (the artifact) and approaches 0 when the anomalies are azimuthally incoherent noise.

A spatially blurred field would also score low in coh. The diagnostic is therefore complemented by a structure-retention metric, the along-range roughness

$$\gamma_r = \text{median} \left[(L(i, j+1) - L(i, j))^2 \right], \quad (12)$$

335 computed over all range-adjacent valid gate pairs. Because the ray-coherent offset cancels exactly in along-range differences, and the range-dependent remainder of Eq. (10) varies too gradually to register in them, γ_r is by construction insensitive to span errors and measures genuine small-scale variability; this was verified directly by injecting synthetic span perturbations into a subset of scans, which raised coh substantially while leaving γ_r statistically unchanged. A method that suppressed the streaks by blurring would depress γ_r toward the level of a smooth field, whereas a method that removes only the ray-coherent error
340 should retain γ_r at the level of the underlying precipitation structure.

Scans are stratified into intensity regimes using the 99th percentile of measured reflectivity within the scoring mask, a peak-intensity classifier that responds to the presence of convective cores: light ($p_{99} < 40$ dBZ; $n = 20,147$ scans), moderate ($40 \leq p_{99} < 50$ dBZ; $n = 4,560$), and heavy ($p_{99} \geq 50$ dBZ; $n = 1,887$). Lower percentiles were found to saturate below the convective range because cores typically occupy only a small fraction of the echo area. This scan-level, reflectivity-based
345 stratification is distinct from the gauge-accumulation-based stratification of Sect. 5: the two share the light/moderate/heavy terminology but classify different objects (scans versus gauge-days) on different axes.

Finally, the locally calibrated composite $R(Z_H, K_{DP})_{\text{LOC}}$ of Sect. 3.2 is scored alongside the attenuation-based methods as a negative control. It shares their geometry, sampling, and microphysical inputs, but neither of its branches uses the span, so by the theory of Sect. 6.1 its coh should sit at the noise floor.

350 6.3 The diagnostic applied to case studies

Figure 6 shows the azimuthal anomaly field L' for a representative scan, recorded on 26 February 2025 at 20:10 local time and classified as moderate by the criterion of Sect. 6.2, processed by the four attenuation-based variants and by the span-independent control. The ray-coherent artifact is unmistakable in $R(Z_H, A_H)_{\text{BASE}}$ (panel a): coherent red and blue radial streaks, each an entire ray shifted by a constant $\ln R$ offset, exactly as Eq. (9) predicts. On this scan the baseline scores $\text{coh} = 0.52$ – more
355 than half of its sub-beam anomaly variance is carried by ray-constant offsets – against $\text{coh} = 0.004$ for the control (panel e), whose anomalies are azimuthally incoherent speckle with no preferred radial organization. The threshold variants (panels b–d) fall between the extremes, and much closer to the baseline than to the control: $\text{coh} = 0.44$ for $R(Z_H, A_H)_{\text{STRICT}}$, 0.34 for $R(Z_H, A_H)_{Z>30}$, and 0.50 for $R(Z_H, A_H)_{\rho_{\text{hv}}>0.95}$ – reductions of at most a third relative to the baseline, and, for the



ρ_{hv} criterion, of no practical consequence. All five panels correspond to rate fields that appear individually plausible in a
360 conventional rate display; the artifact becomes legible only in the anomaly coordinate (cf. Appendix B1).




Figure 6. Azimuthal anomaly of $\ln R$ (L' , Sect. 6.2) for the scan of 26 February 2025, 20:10 local time (moderate regime, Sect. 6.2), for (a) $R(Z_H, A_H)_{BASE}$, (b) $R(Z_H, A_H)_{STRICT}$, (c) $R(Z_H, A_H)_{Z > 30}$, (d) $R(Z_H, A_H)_{\rho_{hv} > 0.95}$, and (e) the span-independent control $R(Z_H, K_{DP})$. Ray-coherent errors appear as radial streaks of constant color. The corresponding rain-rate field is shown in Appendix B1. Per-panel values of the ray-coherence metric are quoted in the text. This is a single scan; the full-period statistics for the complete method family are given in Sect. 7.3.

The reliability of a span estimate is a continuous property of each ray: it improves smoothly with the magnitude of $\Delta\Phi_{DP}$ relative to the phase noise, and no binary gate criterion can track it. Any hard threshold either discards rays whose spans are perfectly reliable or admits rays whose spans are noise-dominated, and usually does both within the same scan.

7 The reliability-weighted estimator

365 7.1 Formulation

The diagnostic of Sect. 6 indicates that the attenuation retrieval should be trusted in proportion to the reliability of its per-ray constraint. We therefore define, for each ray i , a reliability-weighted blend of the specific-attenuation estimator and the locally calibrated composite:

$$R(Z_H, K_{DP}, A_H)_W(r) = w_i R(A_H)(r) + (1 - w_i) R(Z_H, K_{DP})_{LOC}(r), \quad (13)$$



370 where the weight follows from an inverse-variance (minimum mean-square-error) combination of two estimators with independent fractional errors:

$$w_i = \frac{\varepsilon^2}{\sigma_{A,i}^2 + \varepsilon^2}, \quad \sigma_{A,i} = b_R \frac{\sigma_{\Delta\Phi,i}}{\Delta\Phi_{DP,i}}. \quad (14)$$

Here $\sigma_{A,i}$ is the fractional error that span noise imposes on $R(A_H)$ along ray i , obtained by propagating the phase noise through Eq. (9), and ε is the assumed fractional uncertainty of the composite estimator, treated as a single tuning constant. The
375 phase noise $\sigma_{\Delta\Phi,i}$ is estimated from the scatter of the raw Φ_{DP} within the same two endpoint windows used to compute the span itself.

Equation (14) replaces the hard span thresholds of Sect. 4.2 with their continuous limit. As $\Delta\Phi_{DP,i} \rightarrow 0$, $\sigma_{A,i} \rightarrow \infty$ and $w_i \rightarrow 0$ smoothly, so no minimum-span rule is required; as the span grows relative to its noise, $w_i \rightarrow 1$ and the estimator approaches the pure attenuation retrieval.

380 One design detail of the composite branch matters for the integrity of the blend. If the $R(Z_H)$ branch of the composite is evaluated on reflectivity that has itself been attenuation-corrected by the ZPHI pass, the correction inherits the same noisy span, and the ray-coherent error leaks back into the product through the supposedly independent fallback. In the configuration evaluated here, the composite branch therefore uses reflectivity that is not corrected by the ZPHI pass.

The cost of this choice is a possible residual attenuation bias on light-rain gates shadowed by upstream heavy cells, where
385 uncorrected Z_H underestimates; the benefit is that the two branches of Eq. (13) have independent error structures, as both the inverse-variance logic and the diagnostic floor require. The exposed population is in any case small: behind convective cells, gates with $Z_H < 25$ dBZ rarely survive the noise filtering and QC of Sect. 2.2, and that censoring acts on the measured reflectivity, upstream of any correction, so it is common to both configurations of the branch. The gauge validation of Sect. 7.4 indicates the residual bias is small in practice.

390 A physical caveat concerns the backscatter differential phase δ . In heavy rain, δ contributes real, non-noise structure to the raw Φ_{DP} within the endpoint windows, inflating $\sigma_{\Delta\Phi,i}$ and biasing w_i downward exactly where the attenuation retrieval is at its best. The observed heavy-rain weights (median $w \approx 0.75$ in the largest span bin; Sect. 7.2) indicate that this effect is not crippling at the spans encountered in this dataset, but it is a recognized limitation of estimating phase noise from raw Φ_{DP} , and we return to it in Sect. 8.

395 7.2 Weight behavior and selection of ε

Figure 7 shows how the weight behaves as a function of the attenuation span and of the tuning constant ε , with rays pooled into four $\Delta\Phi_{DP}$ bins. The weight behavior and the selection of ε were characterized on a subset of 10 rainy days drawn from the study period, restricted to rays with a positive span and more than 10 valid gates; the full-period evaluation follows in Sect. 7.3. The weight increases monotonically with the span in every configuration, as it must. At the adopted value $\varepsilon = 0.1$, the
400 median weights are 0.02, 0.13, 0.36, and 0.75 for the $0 < \Delta\Phi_{DP} < 5^\circ$, $5-10^\circ$, $10-20^\circ$, and $\geq 20^\circ$ bins respectively: rays with



negligible spans are handed almost entirely to the composite ($w \approx 0$), rays with strong spans are dominated by the attenuation retrieval ($w \geq 0.75$), and the transition is smooth rather than gated.

Where the transition should sit can be argued a priori. Below roughly 30 dBZ neither polarimetric advantage is available: K_{DP} is below its usable threshold and the phase span is small relative to its noise, so the locally calibrated $R(Z_H)$ is the most reliable estimator – the same judgment embodied in the 30 dBZ switching threshold of the composite itself and in the small-span $R(Z_H)$ fallback recommended by Diederich et al. (2015). The reflectivity composition of the span bins (Fig. 7b; medians 24.8, 27.1, 29.8, and 33.2 dBZ) shows that the 5–10° bin is dominated by such gates – its upper quartile, 30.6 dBZ, sits barely above the switch – while the 10–20° bin straddles it, with a median of 29.8 dBZ; the weight should therefore remain low in the former and intermediate in the latter. Of the two candidate values that bracket the intended behavior, $\varepsilon = 0.2$ commits too much weight too early – medians of 0.38 in the 5–10° bin and 0.69 at 10–20° – whereas $\varepsilon = 0.1$ commits less (0.13 and 0.36). Between the two, the lower one is additionally favored by the asymmetry of the risks: over-trusting a noisy span reintroduces the ray-coherent artifact the estimator is designed to suppress, whereas over-trusting the composite forfeits comparatively little, since on the affected rays the composite reduces largely to the locally calibrated $R(Z_H)$. ε was selected on these a priori weight-behavior criteria alone, not by optimizing any validation score; no intermediate value was tuned.

Identifying which branch of the composite is actually invoked clarifies the limiting behavior of the blend. In the two lowest span bins, most gates fall below the composite’s 30 dBZ switch, so the $w_i \rightarrow 0$ limit reproduces the $R(Z_H)$ fallback of the threshold variants. In the 10–20° bin, the composite’s internal switching routes the gates above 30 dBZ to $R(K_{DP})$, so even at a moderate ray weight the heavier gates within these rays are handled by attenuation-immune estimators on both sides of the blend. This also confines the uncorrected-reflectivity caveat of Sect. 7.1 to gates below 30 dBZ, where the forgone attenuation correction is smallest.

7.3 Full-period diagnostic results

Figure 8 presents the coh distributions over the 26,594 scans of the study period, for the complete method family, each method scored on the same gates within each scan. Three results follow. The first concerns the reach of the artifact. The baseline carries median coh values of 0.81, 0.69, and 0.58 in the light, moderate, and heavy regimes respectively – in light rain, more than three quarters of its sub-beam anomaly variance is spurious ray-coherent contamination. The contamination is therefore not confined to light rain, contrary to what visual inspection of conventional displays suggests (Sect. 5).

The second concerns the thresholds. Every hard threshold reduces the contamination, but none approaches the control, which sits between 0.043 and 0.060 across all regimes. The stricter reflectivity threshold ($Z > 30$) and the STRICT fallback reach medians of 0.25 and 0.27 in light rain, yet rebound to 0.36–0.42 in moderate and heavy rain; the $\rho_{hv} > 0.95$ variant remains the worst performer throughout (0.68, 0.60, and 0.48 from light to heavy). The regime dependence of STRICT deserves isolating, because it exposes the mechanism behind its apparent success at the gauges: in light rain, most rays fail the $\Delta\Phi_{DP} \geq 10^\circ$ requirement and the method silently falls back to the span-independent $R(Z_H)$, so its low light-regime coh reflects abandoning the attenuation method. In moderate and heavy rain, where the method actually uses A_H – and where the operationally critical




Figure 7. (a) Distribution of the per-ray blend weight w (Eq. (14)) as a function of ϵ , for rays pooled into four attenuation-span ($\Delta\Phi_{DP}$) bins; whiskers span the 10th–90th percentiles, and the adopted configuration ($\epsilon = 0.1$) is the orange box in each panel. (b) Distribution of reflectivity per span bin, each box drawn below the panel of (a) for the same bin, showing that larger spans are associated with heavier precipitation. Both panels are computed on the 10-day subset described in Sect. 7.2, using rays with a positive span and more than 10 valid gates; the four bins contain 385,968, 183,732, 119,303, and 82,521 rays. Panel (a) is a per-ray distribution, panel (b) a per-gate distribution over the same rays.



rainfall occurs – roughly 40 % of its anomaly variance remains ray-coherent. Thresholds modulate the artifact; none of them
435 removes it.

The third concerns the weighted estimator, which changes the picture qualitatively. Its median coh is 0.25 in light rain and drops to 0.17 and 0.13 in the moderate and heavy regimes – factors of 2.4 and 2.7 below the best-performing threshold variant, in exactly the regimes where the attenuation branch carries the rain and where the thresholds rebound. In light rain the weighted estimator is instead level with the best threshold variant, as the design implies: there the weight of Eq. (14) leaves most rays to
440 the composite and the thresholds likewise revert to $R(Z_H)$, so the two converge on nearly the same span-independent behavior.




Figure 8. Distributions of the ray-coherence metric coh per scan, for the (a) light ($n = 20,147$), (b) moderate ($n = 4,560$), and (c) heavy ($n = 1,887$) regimes defined in Sect. 6.2. Boxes span the interquartile range, the line inside is the median and the whiskers reach the 10th and 90th percentiles; outliers are not drawn. Blue: baseline; orange: threshold variants and the weighted estimator; green: the span-independent control. All methods are scored on identical gates within each scan (paired design).

The structure-retention metric confirms that this reduction is not achieved by over smoothing (Fig. 9). In the light regime, γ_r is statistically indistinguishable across all methods. In the heavy regime, where the trade-off would live, the weighted estimator's along-range roughness (median $\approx 5.0 \times 10^{-3}$) sits modestly below the cluster of A_H -threshold methods ($\approx 7 \times 10^{-3}$). Since γ_r is insensitive to span errors (Sect. 6.2), this difference does not reflect artifact removal; it is the expected
445 signature of blending. The convex combination of Eq. (13) averages two fields whose gate-scale fluctuations are partially independent – the statistical fluctuations of the Z_H^b profile in the attenuation branch, and the heavily range-filtered K_{DP} of the composite – so some reduction of gate-to-gate variance, part noise and part genuine texture, is inherent to the estimator.

The visual counterpart of these statistics returns to where the problem was first seen. Figure 10 shows the case-study scan of Sect. 6.3, at the same color scale and range rings as Fig. 6, processed by the weighted estimator. The radial streaks that
450 dominate the baseline panel are mostly gone; what remains is azimuthally incoherent speckle superposed on the retained storm




Figure 9. As Fig. 8, but for the along-range roughness γ_r ; note the 10^{-3} factor on the axis.

structure. On this scan coh falls to 0.088: six times below the baseline, roughly four times below the best of the threshold variants on the same field, and closer to the control (0.004) than to any of them.




Figure 10. Azimuthal anomaly of $\ln R$ for the same scan, color scale, and range rings as Fig. 6, processed by the reliability-weighted estimator $R(Z_H, K_{DP}, A_H)_W$. The ray-coherent streaks of the baseline are absent while the storm structure is retained (coh = 0.088, against 0.52 for the baseline panel of Fig. 6 and 0.004 for its control panel).



7.4 Gauge validation of the weighted estimator

At the gauges, the weighted estimator performs on par with the best member of the threshold family: PBIAS = +0.20% (95 %
455 CI [−11.57, +16.26] %), MAE = 4.78 mm, RMSE = 8.66 mm, and CC = 0.835.

Evaluated on paired bootstrap replicates (Sect. 2.4), its calibration improvement over the local composite is clear (Δ PBIAS
= −9.74%, 95 % CI [−12.84, −6.28] %), while no difference from STRICT is detected (Δ PBIAS = +1.30%, 95 % CI
[−0.27, +3.53] %): any undetected calibration difference between the two is bounded to a few percent, far below the dif-
ferences that separated the default from the local family. The systematic residuals identified in Sect. 5 survive in the weighted
460 product, as they must, since it inherits the error characteristics of its branches by regime: the light-rain overestimation of the
composite, and the far-range underestimation of heavy rain, which reflects a sampling geometry that no rainfall-rate algorithm
operating on the same beam can be expected to remove.

In short, the gauge network sees the weighted estimator and STRICT as equivalent, while the diagnostic of Sect. 7.3 separates
them decisively in the moderate and heavy regimes that carry the operationally relevant rainfall. It is the diagnostic, not the
465 gauges, that supplies the selection criterion, and on that criterion the weighted estimator is adopted as the final product.

7.5 Spatial distribution of accumulated rainfall

Figure 11 shows the total accumulated rainfall from February to April 2025 as estimated by the gauge-interpolated reference
field, the manufacturer’s default algorithm, the weighted estimator, and the IMERG Final v7 satellite product (Huffman et al.,
2023). The gauge-interpolated map (Fig. 11a) indicates that the highest precipitation totals, exceeding 500 mm, were concen-
470 trated in the northeastern border of the radar domain. The map produced by the manufacturer’s default algorithm (Fig. 11b)
reveals a widespread overestimation with accumulations approaching 800 mm, confirming that the large positive PBIAS identi-
fied statistically is a systematic, domain-wide issue. The weighted estimator (Fig. 11c) reproduces the gauge-derived reference
closely across the accumulation range, tracking both the dry southwestern sector and the wet northeastern gradient, while
resolving spatial detail unavailable to the interpolated reference. For context, the IMERG product (Fig. 11d) captures the
475 broad-scale pattern and correctly identifies the northeast as the wettest area, but underestimates the peak accumulation there
by over 200 mm relative to the gauge reference, overestimates rainfall in the dry southwestern region, and cannot resolve the
fine-scale structure available to the radar.

8 Discussion

The use of manufacturer-default parameters resulted in a severe, systematic overestimation of rainfall. This error is primarily
480 attributed to the $R(K_{DP})$ relationship, whose coefficients coincide with those derived from DSD measurements in Japan
(Maesaka, 2011; Park et al., 2005). Such parameters are fundamentally ill-suited for the unique climate of northwest Peru,
which is shaped by complex interactions between Amazonian moisture, sea breezes, and warming events like El Niño (Caviedes




Figure 11. Cumulative rainfall maps for the February–April 2025 study period. Panels show: (a) Interpolated gauge data, (b) the $R(Z_H, K_{DP})_{DEF}$ algorithm, (c) the $R(Z_H, K_{DP}, A_H)_W$ weighted estimator, and (d) the IMERG Final Run satellite precipitation product. All panels display the 70 km radar range ring for a consistent spatial reference. Rainfall contours derived from the respective data are shown as dashed lines.

and Endlicher, 1989). The implementation of QPE relationships from locally derived DSDs, collected during the 2024 “weak” El Niño conditions, proved essential and is a prerequisite for everything that follows.

Beyond calibration, the central methodological lesson of this study concerns validation itself. The gauge network was decisive where the errors were large (default versus local) and uninformative where they were subtle within the attenuation-based candidates, as daily point accumulations integrate away exactly the spatially structured, scan-level error mode that distinguishes them. A selection based on gauge statistics alone would adopt STRICT – and carry substantial spurious radial structure into every downstream application. The ray-coherence diagnostic fills this gap with a reference-free measurement grounded in the error theory of the ZPHI retrieval. This style of internal, physics-based diagnostic – identify the dominant artifact mode of a retrieval, derive its signature, and measure that signature at scale – is transferable beyond this radar and this algorithm family, and offers a practical complement to gauge validation wherever reference networks are sparse.



For operational users of the weighted product, the practical consequence of the blend design is that error characteristics are inherited by regime. In light precipitation, where the weight hands most rays to the composite, the product inherits the composite's behavior, including the persistent light-rain overestimation whose candidate causes – virga and sub-beam evaporation in this arid environment, subtle radar miscalibration, or manual gauge errors – are physical and instrumental. In moderate and heavy precipitation, the product inherits the characteristics of the attenuation retrieval: immunity to miscalibration and partial beam blockage and an explicit treatment of attenuation, now with the span-noise artifact largely suppressed. Users should likewise retain the caveats that no estimator can escape: underestimation of heavy rainfall at far ranges, as the beam at the 2.5° elevation angle overshoots the precipitation, and possible total signal extinction behind intense convective cores, an attenuation limit inherent to X band radars.

Several limitations qualify these results and define future work. First, the analysis relies on the justifiable assumption that DSD characteristics were similar between the 2024 and 2025 seasons due to comparable weak El Niño conditions. Given the strong modulation of the regional climate by the SST, rainfall microphysics likely vary during strong El Niño or La Niña phases, and long-term disdrometer monitoring is essential to develop a dynamic QPE framework with relationships tailored to different climatic states. Second, the backscatter differential phase δ is not explicitly handled: it contaminates the raw-phase noise estimate in heavy rain and biases the weight downward where the attenuation retrieval is most valuable. Explicit δ screening of the endpoint windows, or a noise estimator applied to δ -corrected phase, is a natural refinement of Eq. (14). Third, the validation was constrained by the daily resolution of the gauge network, which prevents an assessment of algorithm performance at the radar's high temporal resolution; future validation efforts should incorporate data from high-resolution tipping-bucket rain gauges.

Finally, a word on transferability across radar bands. The DSD sensitivity of $R(A)$ decreases from X toward S band, where specific attenuation is close to DSD-invariant (Ryzhkov et al., 2014), while K_{DP} -based estimation is at its most favorable at X band. The reliability-weighting framework itself is band-agnostic: it requires only the ZPHI constraint and a raw-phase noise estimate. If the weighted blend is competitive at X band, where its composite partner is strongest and its attenuation branch most DSD-sensitive, there is reason to expect the balance to shift further in its favor at C and S bands; verifying this on other systems is left to future work.

9 Conclusions

This study developed, validated, and selected polarimetric QPE algorithms for a new X-band dual-polarization radar in the climatically complex region of northwest Peru. Rainfall relationships derived from local, in situ raindrop size distribution measurements proved to be a prerequisite for credible QPE: the manufacturer's default composite overestimated seasonal rainfall by +25.6%, an error removed at its source by the 2024 disdrometer-derived coefficients. Within the locally calibrated family, attenuation-based estimation delivered the best calibration, but its dependence on the total differential-phase span $\Delta\Phi_{DP}$ introduced ray-coherent errors that the daily gauge network – which rated all locally calibrated variants as essentially equivalent – was structurally unable to detect. A gauge-independent, scan-level diagnostic derived from the error structure



of the ZPHI retrieval and applied to more than 26,000 scans made the artifact measurable, and showed both that it affects all precipitation regimes and that it persists under every hard data-selection threshold, because the reliability of the span is a continuous per-ray property that no binary criterion can track. The reliability-weighted estimator constructed in response – a per-ray inverse-variance blend of the specific-attenuation retrieval and the local composite, driven by the span’s own noise estimate – matches the best threshold variant at the gauges while reducing the ray-coherent contamination by factors of 2.4–2.7 in the moderate and heavy regimes, approaching the noise floor of a span-independent control and retaining the along-range structure of the precipitation field. Following this evaluation, the weighted estimator has replaced the manufacturer’s default composite in the radar’s operational near-real-time product.

The final, validated QPE product represents a substantial advancement in regional rainfall monitoring, offering quantitatively reliable precipitation maps at a high spatio-temporal resolution and a more detailed and accurate depiction of precipitation events than both the default radar configuration and global satellite products like IMERG. For a region so vulnerable to hydro-meteorological hazards, particularly flooding during El Niño events, this enhanced capability is of direct and immediate importance for operational forecasting, water resource management, and civil defense.



Appendix A: Additional radial case studies

540 Figures A1–A3 present the radial case studies for moderate and heavy precipitation summarized in Sect. 4.2.




Figure A1. Same as Fig. 3, but for a moderate rain event on 26 February 2025 at 00:00 UTC (Azimuth: 26.0°). The unstable phase at the far boundary of the path illustrates the boundary-jump problem under the baseline threshold.




Figure A2. Same as Fig. 3, but for a second moderate rain profile from the same scan (Azimuth: 20.0°), illustrating the shorter path segments enclosed by the $Z_H > 30$ dBZ threshold relative to the $\rho_{HV} > 0.95$ filter.




Figure A3. Same as Fig. 3, but for a heavy rain event on 29 March 2025 at 02:12 UTC (Azimuth: 35.0°). The co-polar correlation coefficient drops below 0.95 within the rain core, causing the ρ_{HV} -based filter to truncate the path and underestimate $\Delta\Phi_{DP}$.

Appendix B: Additional diagnostic material

B1 Perceptual masking of the ray-coherent artifact

Figure B1 shows the $\ln R$ field of the scan analyzed in Figs. 6 and 10, for the baseline method. The rendering is deliberately favorable to the artifact: the palette is discrete and binned uniformly in $\ln R$, so a ray-constant offset of Eq. (9) displaces a ray by the same number of color classes at any rain rate, making the streaks detectable in every regime on close inspection. Even so, they remain subtle, since each offset spans only a small fraction of the field's dynamic range. Operational displays are far less favorable: their palettes are binned in rain rate (here 0.1–1, 1–5, 5–10, 10–20, 20–50, and 50–200 mm h^{-1}), classes that a multiplicative offset crosses mainly at low rates, where the fractional error of Eq. (9) is largest. Under such a palette the artifact is apparent only in light rain, as inspection of the candidate fields confirmed (Sect. 5). Under even a purpose-built static palette, then, the artifact that dominates the anomaly field of Fig. 6a is at best marginally perceptible: it has to be measured, not inspected – which is precisely why the threshold variants could not be separated by display alone.




Figure B1. The $\ln R$ field of the scan of Fig. 6, for $R(Z_H, A_H)_{BASE}$, rendered with a discrete palette binned uniformly in $\ln R$ – a display chosen to be as favorable as possible to visual detection of ray-coherent offsets. Even under this rendering, the artifact that dominates the corresponding anomaly field (Fig. 6a) is barely perceptible; in discrete rain-rate palettes it is apparent only in light rain.



Data availability. The radar, disdrometer, and UDEP rain gauge data used in this study are available from the corresponding authors upon reasonable request. Rain gauge data from SENAMHI are available via their institutional portal at <https://www.senamhi.gob.pe/?p=estaciones>. Near real-time visualizations of the radar data can be accessed at <https://ecrmm.udep.edu.pe/Radar/lluvia.html> (last access: 20 August 2026). The SRTM global digital elevation model shown in Fig. 1 was obtained from OpenTopography (NASA Shuttle Radar Topography Mission, 2013).

Author contributions. Ivan Herrera-Casas: conceptualization, methodology, software, formal analysis, data curation, writing – original draft preparation. Rodolfo Rodriguez: resources, supervision, project administration, writing – review and editing. Carlos Del Castillo: writing – review and editing. Nilson Diaz: software, writing – review and editing. Yamina Silva: resources, data curation, writing – review and editing.

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

Acknowledgements. The authors express their gratitude to the Estación Científica Ramón Mugica Martínez, Universidad de Piura (UDEP) for the continuous operation of the X-band radar. We gratefully acknowledge the Instituto Geofísico del Perú (IGP) for the deployment and maintenance of the Parsivel disdrometer during the 2024 field campaign, and the Servicio Nacional de Meteorología e Hidrología del Perú (SENAMHI) for providing the rainfall records from their gauge network.

Financial support. This research has received no external funding.



References

- Atlas, D. and Ulbrich, C. W.: Path- and Area-Integrated Rainfall Measurement by Microwave Attenuation in the 1–3 Cm Band, *Journal of Applied Meteorology*, 16, 1322–1331, [https://doi.org/10.1175/1520-0450\(1977\)016<1322:PAAIRM>2.0.CO;2](https://doi.org/10.1175/1520-0450(1977)016<1322:PAAIRM>2.0.CO;2), 1977.
- Bendix, J.: Precipitation dynamics in Ecuador and northern Peru during the 1991/92 El Nino: A remote sensing perspective, *International Journal of Remote Sensing*, 21, 533–548, <https://doi.org/10.1080/014311600210731>, 2000.
- Brandes, E. A., Zhang, G., and Vivekanandan, J.: Drop Size Distribution Retrieval with Polarimetric Radar: Model and Application, *Journal of Applied Meteorology*, 43, 461–475, [https://doi.org/10.1175/1520-0450\(2004\)043<0461:DSDRWP>2.0.CO;2](https://doi.org/10.1175/1520-0450(2004)043<0461:DSDRWP>2.0.CO;2), 2004.
- Caviedes, C. and Endlicher, W.: The precipitation in northern Peru during the El Nino oscillation of 1983 | Die Niederschlagsverhältnisse in Nordperu während des El-Nino Oscillation-Ereignisses von 1983, *Erde*, 120, 81–87, 1989.
- 575 Diederich, M., Ryzhkov, A., Simmer, C., Zhang, P., and Trömel, S.: Use of Specific Attenuation for Rainfall Measurement at X-Band Radar Wavelengths. Part II: Rainfall Estimates and Comparison with Rain Gauges, *Journal of Hydrometeorology*, 16, 503–516, <https://doi.org/10.1175/JHM-D-14-0067.1>, 2015.
- Douglas, M. W., Mejia, J., Ordinola, N., and Boustead, J.: Synoptic Variability of Rainfall and Cloudiness along the Coasts of Northern Peru and Ecuador during the 1997/98 El Niño Event, *Monthly Weather Review*, 137, 116–136, <https://doi.org/10.1175/2008MWR2191.1>,
580 2009.
- ENFEN: Definición operacional de los eventos El Niño y La Niña costera en el Perú, Tech. rep., ENFEN, <https://enfen.imarpe.gob.pe/download/nota-tecnica-enfen-01-2024-definicion-operacional-de-los-eventos-el-nino-costero-y-la-nina-costera-en-el-peru>, last access: 20 August 2026, 2024.
- Goldberg, R. A., Tisnado, G. M., and Scofield, R. A.: Characteristics of extreme rainfall events in northwestern Peru during the 1982–1983 El Niño period, *Journal of Geophysical Research: Oceans*, 92, 14 225–14 241, <https://doi.org/10.1029/JC092IC13P14225>, 1987.
- 585 Helmus, J. J. and Collis, S. M.: The Python ARM Radar Toolkit (Py-ART), a Library for Working with Weather Radar Data in the Python Programming Language, *Journal of Open Research Software*, 4, 25, <https://doi.org/10.5334/jors.119>, 2016.
- Horel, J. D. and Cornejo-Garrido, A. G.: Convection along the Coast of Northern Peru during 1983: Spatial and Temporal Variation of Clouds and Rainfall, *Monthly Weather Review*, 114, 2091–2105, [https://doi.org/10.1175/1520-0493\(1986\)114<2091:CATCON>2.0.CO;2](https://doi.org/10.1175/1520-0493(1986)114<2091:CATCON>2.0.CO;2), 1986.
- 590 Huffman, G. J., Stocker, E. F., Bolvin, D. T., Nelkin, E. J., and Tan, J.: GPM IMERG Final Precipitation L3 Half Hourly 0.1 degree x 0.1 degree V07, Goddard Earth Sciences Data and Information Services Center (GES DISC) [data set], <https://doi.org/10.5067/GPM/IMERG/3B-HH/07>, 2023.
- Ibañez, M. P. A., Martinez, S. C., Pura, A. G., Sajulga, R. A., Cayan, E. O., Jou, B. J.-D., and Chang, W.-Y.: Development of Quantitative Precipitation Estimation (QPE) Relations for Dual-Polarization Radars Based on Raindrop Size Distribution Measurements in Metro
595 Manila, Philippines, *Terrestrial, Atmospheric and Oceanic Sciences*, 34, 24, <https://doi.org/10.1007/s44195-023-00056-0>, 2023.
- Leinonen, J.: High-Level Interface to T-matrix Scattering Calculations: Architecture, Capabilities and Limitations, *Optics Express*, 22, 1655, <https://doi.org/10.1364/OE.22.001655>, 2014.
- Maesaka, T.: Operational Rainfall Estimation by X-band MP Radar Network in MLIT, Japan, in: 35th Conference on Radar Meteorology, AMS, 2011.
- 600 Matrosov, S. Y., Kennedy, P. C., and Cifelli, R.: Experimentally Based Estimates of Relations between X-Band Radar Signal Attenuation Characteristics and Differential Phase in Rain, *Journal of Atmospheric and Oceanic Technology*, 31, 2442–2450, <https://doi.org/10.1175/JTECH-D-13-00231.1>, 2014.



- Mishchenko, M. I., Travis, L. D., and Mackowski, D. W.: T-Matrix Computations of Light Scattering by Nonspherical Particles: A Review, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 55, 535–575, [https://doi.org/10.1016/0022-4073\(96\)00002-7](https://doi.org/10.1016/0022-4073(96)00002-7), 1996.
- 605 Narayana Rao, T., Amarjyothi, K., and Rao, S.: Attenuation Relations for Monsoonal Rain at the X Band from Disdrometric Measurements: Dependency on Temperature, Raindrop Size Distribution and Drop Shape Models, *Quarterly Journal of the Royal Meteorological Society*, 144, 64–76, <https://doi.org/10.1002/qj.3291>, 2018.
- NASA Shuttle Radar Topography Mission: Shuttle Radar Topography Mission (SRTM) Global, Distributed by OpenTopography [data set], <https://doi.org/10.5069/G9445JDF>, 2013.
- 610 Park, S.-G., Bringi, V. N., Chandrasekar, V., Maki, M., and Iwanami, K.: Correction of Radar Reflectivity and Differential Reflectivity for Rain Attenuation at X Band. Part I: Theoretical and Empirical Basis, *Journal of Atmospheric and Oceanic Technology*, 22, 1621–1632, <https://doi.org/10.1175/JTECH1803.1>, 2005.
- Rodríguez-Morata, C., Díaz, H. F., Ballesteros-Canovas, J. A., Rohrer, M., and Stoffel, M.: The anomalous 2017 coastal El Niño event in Peru, *Climate Dynamics*, 52, 5605–5622, <https://doi.org/10.1007/s00382-018-4466-y>, 2019.
- 615 Rollenbeck, R., Orellana-Alvear, J., Rodríguez, R., Macalupu, S., and Nolasco, P.: Calibration of X-Band Radar for Extreme Events in a Spatially Complex Precipitation Region in North Peru: Machine Learning vs. Empirical Approach, *Atmosphere*, 12, 1561, <https://doi.org/10.3390/atmos12121561>, 2021.
- Ryzhkov, A., Diederich, M., Zhang, P., and Simmer, C.: Potential Utilization of Specific Attenuation for Rainfall Estimation, Mitigation of Partial Beam Blockage, and Radar Networking, *Journal of Atmospheric and Oceanic Technology*, 31, 599–619, <https://doi.org/10.1175/JTECH-D-13-00038.1>, 2014.
- 620 Ryzhkov, A., Zhang, P., Bukovčić, P., Zhang, J., and Cocks, S.: Polarimetric Radar Quantitative Precipitation Estimation, *Remote Sensing*, 14, 1695, <https://doi.org/10.3390/rs14071695>, 2022.
- SENAMHI: Caracterización climática y escenarios climáticos de la región Piura, <http://siar.regionpiura.gob.pe/documentos/normativa/2585.pdf>, last access: 20 August 2026, 2016.
- 625 Takahashi, K.: The atmospheric circulation associated with extreme rainfall events in Piura, Peru, during the 1997–1998 and 2002 El Niño events, *Annales Geophysicae*, 22, 3917–3926, <https://doi.org/10.5194/ANGE0-22-3917-2004>, 2004.
- Tapiador, F. J., Checa, R., and de Castro, M.: An Experiment to Measure the Spatial Variability of Rain Drop Size Distribution Using Sixteen Laser Disdrometers, *Geophysical Research Letters*, 37, <https://doi.org/10.1029/2010GL044120>, 2010.
- Testud, J., Bouar, E. L., Oblis, E., and Ali-Mehenni, M.: The Rain Profiling Algorithm Applied to Polarimetric Weather Radar, *Journal of Atmospheric and Oceanic Technology*, 17, 332–356, [https://doi.org/10.1175/1520-0426\(2000\)017<0332:TRPAAT>2.0.CO;2](https://doi.org/10.1175/1520-0426(2000)017<0332:TRPAAT>2.0.CO;2), 2000.
- 630 Thurai, M., Huang, G. J., Bringi, V. N., Randeu, W. L., and Schönhuber, M.: Drop Shapes, Model Comparisons, and Calculations of Polarimetric Radar Parameters in Rain, *Journal of Atmospheric and Oceanic Technology*, 24, 1019–1032, <https://doi.org/10.1175/JTECH2051.1>, 2007.
- Yu, Z., Xian-tong, L. I. U., Bing-hong, C., Jia-bao, F., Lin, C., and Cong-cong, T.: Application of X-band Polarimetric Phased-array Radars in Quantitative Precipitation Estimation, *Journal of Tropical Meteorology*, 29, 142–152, <https://doi.org/10.46267/j.1006-8775.2023.011>, 2023.
- 635